PEX 8508

PCI Express Switch

Silicon Revisions and Errata List

CONFIDENTIAL PROPRIETARY INFORMATION

NDA REQUIRED

Version 1.7

October 2010

History:

Version 0.5, June 14, 2006
- Initial publication of the Errata list
- Added Errata 1 through 12

Version 0.6, August 22, 2006
- Added Erratum #13

Version 0.7, September 8, 2006
- Added Erratum #14

Version 0.8, October 27, 2006
- Added Erratum #15
- Added *Fixed in AB Silicon* column

Version 0.9, January 25, 2007
- Added Erratum #16
- Added Erratum #17
- Added Erratum #18
- Added Erratum #19
- Changed document title to PEX 8508

Version 1.0, February 26, 2007
- Added Erratum #20

Version 1.1, May 23, 2007
- Added Erratum #21
- Added Erratum #22

Version 1.2, September 26, 2007
- Added Erratum #23
- Added Erratum #24
- Added Erratum #25
- Added Erratum #26
- Added Erratum #27

Version 1.3, January 3, 2008
- Minor edit fixes

Version 1.4, May 28, 2008
- Added Erratum #28
- Added Erratum #29
- Added Erratum #30

Version 1.5, January 12, 2009

- Added Erratum #31

Version 1.6, June 3, 2010

- Added Erratum #32

Version 1.7, October 21, 2010

- Added Caution #1
- Added Caution #2

Errata Table

Errata in AA silicon	Description	Risk Category	Fixed in Rev AB Silicon	Fixed in Rev AC Silicon
1.	~~Cannot Write to Atmel 64Kbyte EEPROM~~ Retracted Erratum	~~Low~~	~~No~~	~~No~~
2.	Memory Mapped Access Not Functional	Low	No	No
3.	Incorrect Error Reported	Low	No	No
4.	UR Error for Requests Targeting Different Host Domain	Low	No	No
5.	CSR Access when BAR0 and BAR1 have same Values	Low	No	No
6.	~~Internal Debug Mode~~ Retracted Erratum	~~Low~~	~~No~~	~~No~~
7.	No ERR Message Reported when Malformed Packet Sent	Low	No	No
8.	FIFO Can Overflow when CFC Slower than SSC	Low	No	No
9.	Hot Reset Used with an EEPROM Enabled	Low	No	No
10.	PEX 8508 does not enter into L0s State as a Transmitter	Low	No	No
11.	PEX_LANE_GOOD Signals Report Incorrect Status	Medium	No	No
12.	Incorrect Value in Max Link Width Register	Low	No	No
13.	AC-JTAG Not Supported	Low	No	No
14.	12-byte Base register shift in NT Mode	Low	No	No
15.	Credit Overflow when Cut-Thru is Enabled	High	Yes	Yes
16.	EEPROM Load on Hot Reset Does Not Work Properly When Non-Transparency is Enabled	Low	No	No
17.	Role-Based Error Reporting Bit Cannot Be Set in Device Capability Register	Low	No	No
18.	Advisory Non-Fatal Error is not Supported	Low	No	No
19.	Coverage Deficiencies in PEX 8508 RAM Shadow Logic May Lead to Data Corruption	High	No	Yes
20.	Possible Link Down-Train During Link Initialization	Medium	No	Yes
21.	PEX 8508 RX Terminations Not Enabled in L2(Vaux) State	Low	No	No
22.	Incorrect Beacon Generation on	Low	No	No

	Downstream Port			
23.	Possibility for the LTSSM getting stuck in the FCInit state when in NT Mode	Low	No	No
24.	Dropped SKIP Ordered Set when Exiting L1 Power Management State	Low	No	No
25.	SDA and SCL I/O Pins can Load Down I^2C Bus when Chip is Not Powered	Low	No	No
26.	No Unsupported Request Response for Non-Posted TLP Targeted to Non-existent Downstream Device After an Upstream Port Secondary Bus Reset	Low	No	No
27.	PEX8508 Upstream Link Does Not Train Under Certain Configurations Involving Lane Reversal	Low	No	No
28.	I^2C Interface NAKs Last Expected Byte of Transaction	Low	No	No
29.	PEX8508 Internal Buffers Are Not Flushed When Link Goes Down and Cut-Thru is Enabled	Medium	No	No
30.	Link Up Failure of Identical Ports on Two Devices Connected by Crosslink and Lane Reversal	Low	No	No
31.	Vendor Defined Broadcast Messages Are Not Supported	Low	No	No
32.	Compliance Pattern Issues on Some Lanes	Low	No	No

Erratum #1 – Cannot Write to Atmel 64Kbyte EEPROM (AT25HP512)

Retracted Erratum – Documented in Data Book

Erratum #2 – Memory Mapped Access Not Functional

Risk Category: Low Silicon Revisions Affected: AA, AB, AC

Description

This erratum is only applicable to NT mode. The device implements a 'limit' register on NT virtual BARs. This option allows the user to use a segment of the full BAR range. When BAR0 (Base Address Register 0) on the upstream port is programmed with a value which is in the range of the full virtual BAR range (irrespective of the limit) then memory mapped accesses will not work for all ports.

Workaround

The space between the NT virtual limit and BAR range should not be used for BAR0 addresses. By default, the PCI enumeration software will never program it in such a way that this erratum is produced.

Customer Impact

When using EEPROM or I^2C to program the NT "Virtual" address ranges, ensure that there is no overlap in the address range. The 'limit' functionality can still be used to restrict the window across the non-transparent boundary.

Erratum #3 – Incorrect Error Reported

Risk Category: Low Silicon Revisions Affected: AA, AB, AC

Description

When a Transaction Layer Packet (TLP) using a Reserved Routing ID and an Undefined Message code is received, the PEX 8508 is supposed to flag the Reserved Routing ID under the "Local, Terminate at Receiver" implicit routing category and report an Unsupported Request (UR) error for the Undefined Message code. Instead, the PEX 8508 treats the Reserved Routing ID as an undefined type and reports a Malformed Error. Even if the message code is defined, the PEX 8508 will still report a Malformed Error for this TLP.

Workaround

None.

Customer Impact

Reserved Routing IDs and Undefined message codes cannot be used in the system.

Erratum #4 – UR Error for Requests Targeting Different Host Domain

Risk Category: Low Silicon Revisions Affected: AA, AB, AC

Description

This erratum is only applicable to NT 64-bit mode. When BAR0 is a 32 bit BAR and its address matches the upper 32 bits of the 64-bit BAR of the NT port (64-bit decode), the TLP is only supposed to hit the BAR4/BAR5 group. Instead, the PEX 8508 takes the upper 32 bits and compares them against the BAR0 value. It sees a match with BAR0 and the non-zero TC and it drops the TLP with an Unsupported Request error (UR). If the TC is 0, it will convert the transaction as a memory mapped access. This particular defect causes an UR error for requests targeting a different host domain.

Workaround

Examine memory windows after enumeration to ensure that there is no overlap. If there is an overlap, adjust the 64-bit BAR window.

Customer Impact

The lower 20 bits of the BAR0 space are not programmable, whereas the upper 32 bits of the 64-bit NT space are fully programmable. In practical systems there should be no overlap after normal enumeration given the extremely large window size that would need to occur to create the overlap condition.

Erratum #5 – CSR Access when BAR0 and BAR1 have same Values

Risk Category: Low Silicon Revisions Affected: AA, AB, AC

Description

This erratum is only applicable to NT mode. When both BAR0 (MEM space) and BAR1 (I/O space) have the same value when accessing the configuration using memory mapped configuration access, some Control & Status Registers (CSRs) will not be accessed properly.

Workaround

Disable the I/O space using EEPROM or I^2C prior to enumeration.

Customer Impact

Ensure memory (BAR0) & I/O (BAR1) BAR spaces do not overlap.

Erratum #6 – Internal Debug Mode Error

Retracted Erratum – Internal Debug Mode Error

Erratum #7 – No ERR Message Reported when Malformed Packet Sent

Risk Category: Low Silicon Revisions Affected: AA, AB, AC

Description

When a specific malformed packet is sent to the PEX8508 (a TLP with FMT & Type combination 5B), the packet is routed through the switch to the appropriate destination instead of being dropped with an ERR (error) message. The FMT of 10 & type field value of 11011 falls into the range of Advanced Switching messages according to the PCI Express base specification r1.0.

Workaround

None.

Customer Impact

Devices compliant to r1.0a or later should not have any problems.

Erratum #8 – FIFO Can Overflow when CFC Slower than SSC

Risk Category: Low Silicon Revisions Affected: AA, AB, AC

Description

When SSC isolation is enabled and the CFC clock is slower than the SSC clock, the internal FIFO could overflow and cause the links to get retrained.

Workaround

Restrict the clock frequencies to always ensure that the CFC clock is at least 1ppm faster than the SSC clock.

Customer Impact

Higher link error count if the CFC clock ppm cannot be controlled.

Erratum #9 – Hot Reset Used with an EEPROM Load

Risk Category: Low **Silicon Revisions Affected: AA, AB, AC**

Description

When Hot Reset is used with an EEPROM enabled, the internal register 0x668 is reloaded. The value in this register reflects the number of active ports being used by the device. If the EEPROM does not have the correct value for the number of active ports, the device will not enumerate properly.

Workaround

Load the correct value for Port0 register 0x668.
Bit0 refers to port0 on the station
Bit1 refers to port1 on the station.
Bit2 refers to port2 on the station.
Bit3 refers to port3 on the station.
Bit4 refers to port4 on the station.

Customer Impact

Systems that implement a Hot Reset during enumeration may fail to boot without the correct EEPROM values.

Erratum #10 – PEX 8508 does not enter into L0s State as a Transmitter

Risk Category: Low **Silicon Revisions Affected: AA, AB, AC**

Description

No power saving will be observed when ASPM is enabled in the device, due to an errata the device will not actively place the transmitters in the power saving state when the transmitters are idle.

Workaround

None.

Customer Impact

The PEX 8508 will not initiate entry into L0s state even if the system enables the ASPM functionality. However, external devices can enter/exit L0s state. The system software can enable L0s in the platform but will not see the device actively enter into L0s state when the transmitting links are idle.

Erratum #11 – PEX_LANE_GOOD Signals Report Incorrect Status

Risk Category: Medium Silicon Revisions Affected: AA, AB, AC

Description

The PEX 8508 lane good signals give the correct status for lanes 0-4 in all port configurations. In the x4x1x1x1x1 configuration, the status of all lanes is correct. However, in all other port configurations, the last 3 "lane good" statuses are not correctly reflected by the PEX_LANE_GOOD signals.

Workaround

The correct link widths are reflected in the link status register which the software would need to read.

Customer Impact

PEX_LANE_GOOD[5-7] signals cannot be relied upon for configurations other than x4x1x1x1x1.

Erratum #12 – Incorrect Value in Max Link Width Register

Risk Category: Low Silicon Revisions Affected: AA, AB, AC

Description

PEX 8508 Max Link width register (Link Capability Register, 74h Bits[9:4]) value is incorrect.

Workaround

The correct link widths are reflected in the link status register which the software would need to read.

Customer Impact

This is a status register for reporting maximum capable link width. Since the PCI Express specification allows for devices to report link widths greater than the adapter/connector width, there should be no system issues as a result of this incorrect setting.

Erratum #13 – AC-JTAG Not Supported

Risk Category: Low Silicon Revisions Affected: AA, AB, AC

Description

The SerDes receiver circuitry contains a circuit defect that causes the scan flip-flop to be falsely cleared just before it is to be scanned. The receiver fails to operate reliably at conditions that would normally be used for system level AC-JTAG operations. The conclusion is that the SerDes AC-JTAG model is not reliable for typical system

environments and is no longer included in the BSDL file. DC-JTAG (1149.1) is unaffected.

Workaround

None. AC-JTAG has been removed from the BSDL file.

Customer Impact

Boundary-Scan instructions *EXTEST_PULSE* and *EXTEST_TRAIN* are not supported.

Erratum #14 – 12-byte Base register shift in NT Mode

Risk Category: Low Silicon Revisions Affected: AA, AB, AC

Description

This erratum applies only to NT enabled modes. When the NT Port Link Interface BAR4 or BAR2 is used in the 32-bit address space which is made prefetchable, the Base registers are erroneously shifted up by 12 bytes in the link Type-0 registers.

For example, assume a BAR with a size of 1MB is mapped to the 32 bit address space. During resource allocation, if the system assigns address range BBB0_0000h, this implies Base and Limit values of BBB0_0000h & BBBF_FFFFh. However, the PEX 8508 shifts the Base by 12-bytes to BBB0_000Ch. The Limit value remains unchanged at BBBF_FFFFh.

Workaround

The software needs to account for the 12-byte shift in Base when allocating resources or generating the transaction to the link NT port BAR4 or BAR2 32-bit BAR. If the BAR is not made prefetchable, it will not have any issues.

Customer Impact

When the 32-bit space is enabled, the resources allocated to the NT link BAR4 or BAR2 need to account for the 12-byte shift in the address space if the BAR is made prefetchable. Use BAR2 or BAR4 in the 32-bit space as non-prefetchable BARs.

Erratum #15 – Credit Overflow when Cut-Thru is Enabled

Risk Category: High Silicon Revisions Affected: AA

Description

When Cut-Thru mode is enabled and the PEX8508 has three or more ports configured at x2 link widths, a credit overflow violation might occur in the Cut-Thru concentrator port, typically the upstream port, of the PEX8508. This could result in a system failure. Both Virtual Channels in the PEX8508 are affected.

The following conditions must be met in order to experience the failure.

1. PEX8508 must be strapped in one of the two following port configurations:
   - Port = x2; Port = x2; Port = x2; Port = x2
   - Port = x2; Port = x2; Port = x2; Port = x1; Port = x1

2. Only two header credits of any type remain available in the Cut-Thru concentrator port

3. There is no packet congestion in the switch such that the Cut-Thru path is activated

4. Packets resulting in failure must originate from two or more downstream ports.

5. Three packets from two or more of the downstream ports are received nearly simultaneously by the scheduling engine in the upstream port. As an example, Port X sends any packet while Port Y sends a sequence of packets consisting of a short packet (12-20 Bytes including header and payload) closely followed by another packet of the same type to the upstream port.

Workarounds

Disable Cut-Thru. This should be performed by clearing register 1DCh bit 21 in port 0. Since the Cut-Thru Enable bit is enabled by default, it is recommended that the change be performed during device initialization and before traffic is sent through the switch. Register 1DCh is accessible from various interfaces including EEPROM, I^2C and Memory-mapped.

Customer Impact

Cut-through will have to be disabled when three or more ports in PEX8508 are hardware configured at x2 link widths.

Erratum #16 – In Non-Transparency Mode, EEPROM Load on Hot Reset does not work properly

Risk Category: Low Silicon Revisions Affected: AA, AB, AC

Description

The PEX 8508 switch can be reset out-of-band via hardware reset signal (cold or warm reset) or in-band via a PCI Express reset message (hot reset) on the upstream port. The PEX 8508 can be configured to re-load the EEPROM configuration upon reset (cold, warm and hot).

When the 8508 is configured to load the serial EEPROM after a hot reset is received on its upstream port and the PEX 8508 is configured in Non-Transparent (NT) mode, the serial EEPROM will not be loaded properly.

Workaround

When in NT mode, do not load the serial EEPROM after a hot reset. Disable EEPROM load by writing a '1' to register 0x1DCh bit 17.

Issue a cold or warm reset instead.

Customer Impact

For an NT configures system, customer cannot issue a hot reset to the PEX8508.

Erratum #17 – Role-Based Error Reporting Bit Cannot Be Set in Device Capability Register

Risk Category: Low Silicon Revisions Affected: AA, AB, AC

Description

The PCI Express Base Specification r1.1 requires all compliant devices to Set the Role-Based Error Reporting bit in the Device Capability Register. This bit is cleared in the PEX 8508 switch.

Workaround

There is no work around. The PEX 8508 was designed before this ECN was made available by the PCI SIG.

Customer Impact

The Role-Based Error Reporting capability is not available to the user.

Erratum #18 – Advisory Non-Fatal Error not Supported

Risk Category: Low Silicon Revisions Affected: AA, AB, AC

Description

The Advanced Error Capability defines support for Advisory Non-Fatal Error. When present, the device which detects an uncorrectable error of non-fatal severity may signal the error using ERR_COR instead of ERR_NONFATAL. In this case the PEX 8508 will not set the Advisory Non-Fatal Error Status bit in the Correctable Error Status register.

Workaround

There is no work around. The PEX 8508 was designed before this particular ECN was made available by the PCI SIG. All other Advanced Error Reporting capabilities are supported by the PEX 8508.

Customer Impact

Non-fatal errors will be signaled with ERR_NONFATAL.

Erratum #19 – Coverage Deficiencies in PEX 8508 RAM Shadow Logic May Lead to Data Corruption

Risk Category: High Silicon Revisions Affected: AA, AB

Description

In reviewing the test programs and DFT coverage for the PEX 8508, it was discovered that the feature which allows the memories to be put in bypass mode during Automated Test Pattern Generation (ATPG) is disabled. Consequently, the logic surrounding the memories is not covered by the tests. This results in fault coverage loss of approximately 1% in the data path area of the chip.

Workaround

There is no workaround for silicon Revisions AA and AB. Silicon revision AC fixes this problem.

Customer Impact

Undetected manufacturing defects in the logic surrounding the RAMs could result in data corruption.

Erratum #20 – Possible Link Down-Train During Link Initialization

Risk Category: Medium Silicon Revisions Affected: AA, AB

Description

When the PEX 8508 is connected to devices that start transmitting training sets to the PEX 8508 after the PEX 8508 has been transmitting training sets (PEX 8508 in Polling state), there is a period of exposure where the receivers in the PEX 8508 are in 'electrical idle' state while the transmitters in the PEX 8508 are sending valid data . During this period the Elastic Buffer could be set in a state that prevents the lane from participating in the final link configuration.

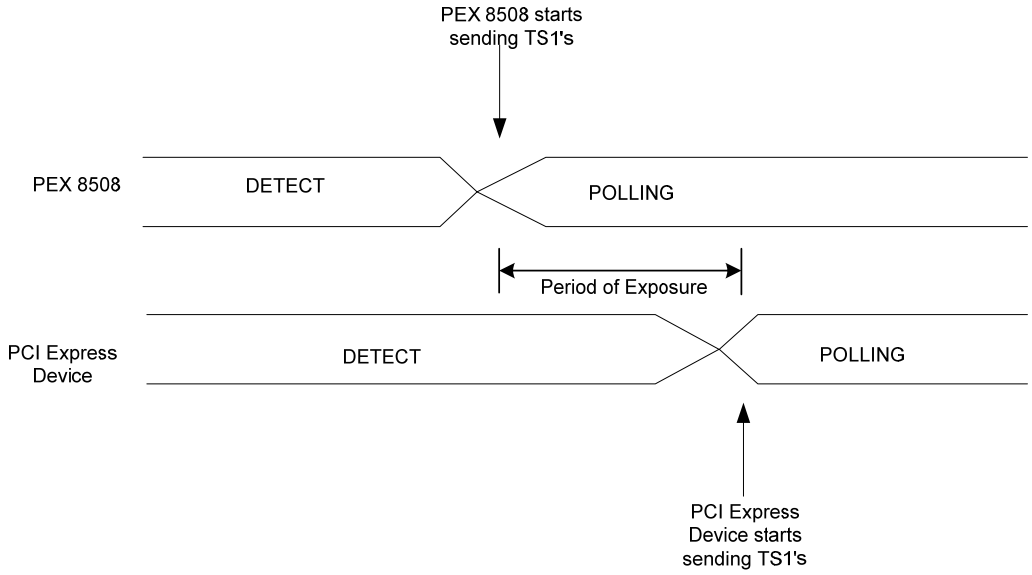

Workaround

1. Use Different Reference clocks. The Elastic Buffer is able to recover from this condition in the Polling state when different reference clocks are used. The link will be reliable once it transitions to the Active state.

2. Re-start training sequence on link. When the adjacent device attempts to train at a link width less than its programmed width, the PEX 8508 port can be instructed to re-start its link training sequence by issuing a link disable closely followed by a link enable. The PEX 8508 link can be disabled/enabled via register 0x78 bit 4 of the appropriate port configuration space.

3. Skew the reset inputs such that the adjacent device starts transmitting ordered sets first. The Elastic Buffer will not enter the illegal state when the PEX 8508 detects valid data on its receiver lanes during its transition from Detect state to Polling state.

Customer Impact

The PEX 8508 link may down-train to a smaller width when it enters the Polling state.

Erratum #21 – PEX 8508 RX Terminations Not Enabled in L2(Vaux) State

Risk Category: Low **Silicon Revisions Affected: AA, AB, AC**

Description

The SerDes for the PEX 8508 link receive signals are not terminated internally during L2 state. Sub-threshold levels on these signals could be incorrectly interpreted by the PEX 8508 as a Beacon sequence on that link. As a result, the system could inadvertently wake up.

Workaround

Bypass switch when implementing WAKE# functionality

1. Do not keep PERST# asserted to the PEX 8508
2. Maintain the PEX8508 in a powered-up state and route the WAKE# signal around the PEX8508 from downstream devices to the upstream devices

Customer Impact

Beacon functionality is not reliable

Erratum #22 – Incorrect Beacon Generation on Downstream Port

Risk Category: Low **Silicon Revisions Affected: AA, AB, AC**

Description

The PEX 8508 latches the upstream port information when the upstream port links up. If the link has not come up, a beacon signal could be detected on some lanes while the PERST# remains asserted.

Workaround

None

Customer Impact

A beacon signal could be detected on any lane while the PEX 8508 PERST# is asserted.

Erratum #23 – Possibility for the LTSSM getting stuck in the FCInit state when in NT Mode

Risk Category: Low **Silicon Revisions Affected: AA, AB, AC**

Description

During link initialization when the PEX 8508 is configured in NT mode, there is a possibility that the data-link layer state machine might get stuck in the FCinit state as a result of the following sequence of events:

1. PEX 8508 PHY link is up
2. Other PCIe device PHY link is up
3. Other PCIe device PHY link attempts to re-train prior to the completion of the current VC0 Flow-Control initialization between the PEX 8508 and the other device.

WorkArounds

One of the following can help resolve the problem.

1. Allow hot-reset to propagate onto the NT link by setting device specific offset 0x1DCh[16] to 0.
2. Ensure Power-up and Reset sequencing of device to be compliant to the PCI Express Base specification; i.e. avoid link from randomly/unexpectedly re-training following the de-assertion of fundamental reset (PERST#) and before the initial link training sequence has completed.

Customer Impact

A PCIe device which allows its link PHY to re-train before the current FC initialization is complete will be exposed to the stuck FCinit state described above. A PCIe compliant device should not attempt to retrain its link before FC initialization is complete.

Erratum #24 – Dropped SKIP Ordered Set when Exiting L1 Power Management State

Risk Category: Low **Silicon Revisions Affected: AA, AB, AC**

Description

When exiting from the L1 Power Management (PM) state, there is a narrow window (recovery.idle state) during which SKIP ordered sets are not transmitted by the PEX8508. The exit from L1 PM state steps through 3 sub-states:

1) rcvry.rcvrlock
2) rcvry.rcvrconfig
3) rcvry.idle.

For the duration in which the link is in Recovery.idle state (waiting for receipt of 8 Idle symbols) there are no SKIP ordered sets transmitted by the PEX8508. Note that SKIP

symbols will be transmitted properly when entries into recovery sub-states are from another (L0 or L0s) state.

The Figure below illustrates the behavior.

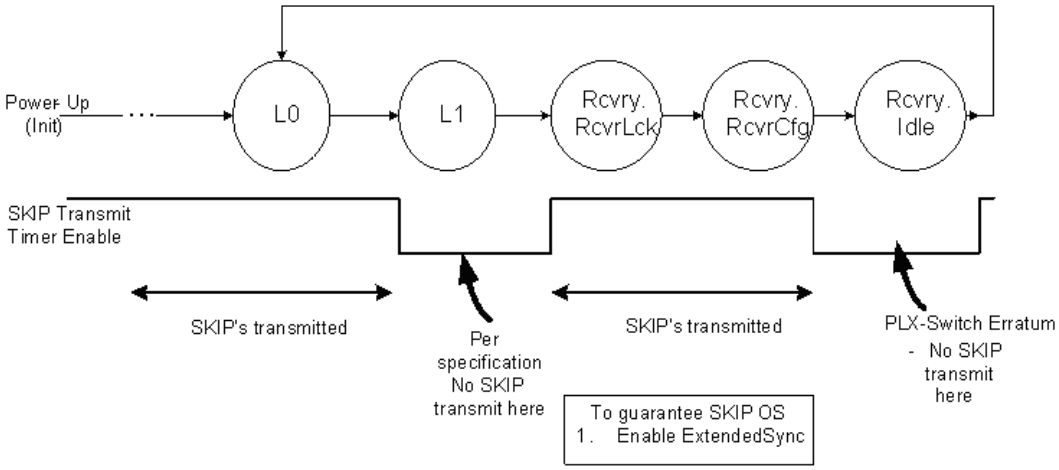

Workaround

One of the following workarounds can be used with upstream devices whose PHYs require SKIP ordered sets to exit its L1 state.

1. Disable the PEX 8508 upstream port's entry to L1 state for ASPM L1 [78h[1:0]] and PCI PM (device specific offset 1E0[7]]) mechanisms.
2. Enable the ExtendedSync mode (Offset 78h[7]), to guarantee SKIP ordered sets are transmitted in the RcvrLck and RcvrCfg states (shown above)

Customer Impact

PCI Express devices whose PHY require SKIP ordered sets to deskew and also to exit the recovery path shown above could fail L1 entry/exit and reach the L0 state >100ms later; consequently, taking the link down along the way. The workarounds listed above can be used to prevent the failing case. Note that when using Workaround 1, the PEX 8508 will not realize any L1 power savings.

Erratum #25 – SDA and SCL I/O Pins can Load Down I^2C Bus when Chip is Not Powered

Risk Category: Low **Silicon Revisions Affected: AA, AB, AC**

Description

The I^2C specification requires that the open-drain SDA and SCL I/O pins in slave mode be floating and present no load to the wire-AND bus when power is switched off to the PEX 8508.

The PEX I/O implementation contains a clipping diode to the VDD33. When the power is off, this internal diode can turn on when the SDA or SCL line is asserted high by a master or any other slave device hanging on the bus. The resulting forward-bias current can lead to abnormal operation within the PEX 8508, resulting in system hang-up.

Workaround

At the board level, connect the SDA and SCL pins to the bus through a FET device controlled by the local PEX VDD33 supply rail. When the power is off, the two FET devices will isolate the PEX 8508 from the local I^2C bus. The SuperTex TN2501 is an example of a very good, low Vgs transistor.

Customer Impact

If the PEX 8508 is not powered up, the I^2C local bus may not function properly (i.e. it could result in system hang-up) without the two additional FET devices to isolate the PEX 8508 from the local I^2C bus.

Erratum #26 – No Unsupported Request Response for Non-Posted TLP Targeted to Non-existent Downstream Device After an Upstream Port Secondary Bus Reset

Risk Category: Low **Silicon Revisions Affected: AA, AB, AC**

Description

If the PEX 8508 is strapped in a configuration where the downstream ports are vacant (no device connected) and a Secondary-Bus-Reset is applied to the upstream port of the switch (Reset on the virtual bus, upstream port register 3Ch[22]), subsequent Non-Posted Transaction Layer Packets (TLPs) targeted to a non-existent downstream device are incorrectly silently discarded. The PEX 8508 should respond to such requests with an Unsupported Request (UR). If a device is present on the downstream port of the switch then the transactions on the corresponding down stream port are handled correctly. A Hot Reset received on the upstream port of the switch, or an upstream port link down condition (DL_Down), or a downstream port Hot Reset (Bridge Control 3C[22]) generated from the downstream port of the switch do not lead to this behavior.

Workaround

Rather than apply a Secondary Bus Reset to the upstream port, instead use a Hot Reset from the upstream device, or an upstream port DL_Down condition, or apply Secondary Bus Resets (3Ch[22]) from the downstream ports of the switch.

Customer Impact

The Device generating a Non-Posted TLP to a Non-Existent downstream device can experience a Completion Timeout, as the switch will not provide a Completion to the Non-Posted TLP.

Erratum #27 – PEX8508 Upstream Link Does Not Train Under Certain Configurations Involving Lane Reversal

Risk Category: Low **Silicon Revisions Affected: AA, AB, AC**

Description

On the link connected to the upstream port of the PEX 8508, if lane '0' of the upstream device is not connected to either the lowermost lane (lane 0) or the uppermost lane (lane n-1) of the PEX 8508 upstream port then the link does not train, even if the upstream device supports lane reversal.

Consider the figure below where a x2 port on the upstream device is connected to the x4 upstream port on the PEX 8508 and the lanes are routed according to the figure. Note that the upstream device supports Lane Reversal.

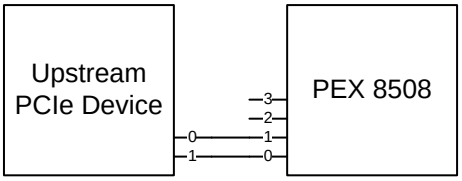

Since lane '0' of the PEX8508 upstream port is connected to lane '3' of the upstream device and the upstream device supports lane reversal, the link should train as x2. However, the erratum in the PEX 8508 prevents the link from training at all.

Workaround

1. Strap the PEX 8508 upstream port to match the port width of the external device. In the example above the upstream port of the PEX 8508 needs to be strapped with x2 width in order for the link to train.
2. Route the signals such that lane '0' from the upstream device connects to the uppermost lane on the PEX 8508. In the example above, lane '0' on the upstream device needs to connect to lane '3' on the PEX8508, etc.

Customer Impact

Link does not train on the upstream port of the PEX 8508 when the lanes are connected as described above. Note that the downstream ports of the PEX 8508 are not affected by this erratum.

Erratum #28 – I²C Interface NAKs Last Expected Byte of Transaction

Risk Category: Low Silicon Revisions Affected: AA, AB, AC

Description

This erratum applies to the I^2C interface of the device. When writing to the device, the I^2C interface always NAKs the last expected byte of the transaction. Instead, the device should ACK all bytes being written. If an external master tries to write more than the expected number of bytes, then the device should NAK those unexpected bytes.

Workaround

If the I^2C external master ignores the NAK on the last byte and does not retry the write of the last byte, the I^2C interface works as expected.

Customer Impact

After the I^2C interface NAKs the last expected byte of the transaction, if the I^2C master keeps retrying the I^2C write access, an infinite loop may be created.

Erratum #29 – PEX8508 Internal Buffers Are Not Flushed When Link Goes Down and Cut-Thru is Enabled

Risk Category: Medium Silicon Revisions Affected: AA, AB, AC

Description

When Cut-Thru mode is enabled in the PEX8508 and traffic is using the cut-thru path to a particular port, there is a possibility during a surprise link down event that the queued TLPs will not be flushed from the port that experienced the surprise down.

Customer Impact

If the scenario above occurs, the internal queues for the affected port are not flushed resulting in the configuration space for the PEX8508 becoming inaccessible to the system. Consequently, the entire PCI Express hierarchy must be reset.

Workaround

1. Prior to disconnecting a device from the PEX8508 downstream transparent port, disable the link by setting the port's Link Disable bit (78h[4] = 1). This is applicable for cases in which a surprise link down is not expected (Hot Plug Surprise bit (7Ch[5]) is not set). This workaround is not available to the NT port since the NT port does not have a Link Disable bit.

2. Disable Cut-thru for the remaining cases where a surprise link down is expected.

Erratum #30 – Link Up Failure of Identical Ports on Two Devices Connected by Crosslink and Lane Reversal

Risk Category: Medium **Silicon Revisions Affected: AA, AB, AC**

Description

When two PLX switches are connected by ports that have the same Port Number, intermittent failure to linkup can occur. Affected ports must be connected using Crosslink and must be Lane Reversed. Ports not using both Crosslink and Lane Reversal, or that have different Port Numbers, are not affected.

The root cause is that the identical ports have the same crosslink time-out. During link training, both ports eventually start transmitting valid link numbers at the same time, which causes both ports to detect lane reversal at the same time. This in turn causes both ports to reverse lanes, which keeps the lanes reversed with respect to each other. Consequently the LTSSM continually cycles through the following states: Detect → Polling → Configuration

Customer Impact

If the scenario above occurs in a system with two back-to-back PLX switches, the PCIe link can intermittently fail to train.

Workaround

1. Disable Crosslink support in one of the switch devices. In a crosslink application, enable Crosslink support on only one of the two switches.
2. Design without Lane Reversal on this link.
3. Design such that the two switches are connected by non-identical port numbers. If for example the identical ports are NT ports (back-to-back), use a different port as the NT port on one of the switch devices.

Erratum #31 – Vendor Defined Broadcast Messages Are Not Supported

Risk Category: Low **Silicon Revisions Affected: AA, AB, AC**

Description

The PEX 8508 does not route vendor defined broadcast messages received on its upstream port to the downstream ports. Instead, the PEX 8508 switch silently drops the vendor defined Type 1 message TLPs and responds to Type 0 message TLPs with unsupported request.

Workaround

If possible, use memory write TLPs in place of broadcast TLPs.

Customer Impact

Endpoints connected to the downstream ports of the PEX 8508 will not receive vendor defined broadcast message(s) from the root-complex.

Erratum #32 – Compliance Pattern Generation Issue on Some Lanes

Risk Category: Low Silicon Revisions Affected: AA, AB, AC

Description

For certain port configurations, compliance pattern is not generated on certain lanes.

Port Configuration	Affected Lanes
0000	6,8,10
0010	8,10
0011	6,8,10
0100	6,8,10
0101	10
0110	8,10
1001	6,8

Workaround

Use port configuration 0x8h for compliance pattern testing.

Customer Impact

Port configurations other than 0x8h will have issues when doing compliance pattern testing.

List of Cautions

Caution #1

Description

Device/Vendor ID value in the NT Port Link Interface registers (offset 11000h) cannot be changed by EEPROM. The PEX 8508 Data Book indicates that the Device/Vendor ID (offset 11000h) of the NT Port Link Interface can be programmed by EEPROM. However, this register cannot be programmed by EEPROM (although the Device/Vendor ID (offset 10000h) of the NT Port Virtual Interface can be programmed by EEPROM).

Caution #2

Description

SerDes 6, 8 and 10 Drive Setting mismatch in register 248h, 24Ch and 254h. The correct register mapping is:

0248h SerDes Nominal Drive Current Select
[1:0]=SerDes0
[3:2]=SerDes1
[5:4]=SerDes2
[7:6]=SerDes3
[9:8]=SerDes4
[11:10]=SerDes6 <=
[13:12]=SerDes8 <=
[15:14]=SerDes10 <=

024Ch SerDes Drive Current Level Select1
[3:0]=SerDes0
[7:4]=SerDes1
[11:8]=SerDes2
[15:12]=SerDes3
[19:16]=SerDes4
[23:20]=SerDes6 <=
[27:24]=SerDes8 <=
[31:28]=SerDes10 <=

0254h SerDes Drive Equalization Level Select 1
[3:0]=SerDes0
[7:4]=SerDes1
[11:8]=SerDes2
[15:12]=SerDes3
[19:16]=SerDes4
[23:20]=SerDes6 <=
[27:24]=SerDes8 <=
[31:28]=SerDes10 <=