PEX 8749/48/47/33/32/25/24/23/17/16/13/12

Errata
Version 2.9

Table of Contents

Chapter 1: Errata List

NOTE: Throughout this document, unless specified otherwise, *Switch* is used to indicate the PEX 87xx switch (PEX 8749/48/47/33/32/25/24/23/17/16/13/12).

1.1 EEPROM Is Needed to Fine-Tune SerDes Settings

Risk Category: Low

Silicon Revisions Affected: BA

Description: Default setting in the device is not adjusted for best SerDes performance.

Solution/Workaround: Use EEPROM to fine tune the SerDes setting. PLX provides the EEPROM contents, which adjust the following SerDes parameters:

- DFE
- Lane VregH
- DFE_TERM

Impact: Higher receiver error count and link negotiation issues if the SerDes settings are not optimum.

1.2 PEX 8749/48/47/33 Ports 16 to 21 Fail to Return the Error Completion in DL_Down State (when the Slot Is Unpopulated)

Risk Category: Medium

Silicon Revisions Affected: BA

Description: If the PCI Express downstream port is unpopulated (DL_Down state), it should report unsupported request error for Memory Write and all Nonposted TLPs (including all Configuration transactions) that are destined to that port. The downstream port should also generate completion with completion status set to UR. Ports 16 to 21 do not report “unsupported request” error and do not generate error completion.

Solution/Workaround: Any one of the following solutions are valid:

- Disable Open Ports 16 to 21 (Port 0 offset 314h) using an EEPROM and enable when the slot is populated.
- Strap and use Port 16 as the upstream port, if port width is smaller than x16, ensure remaining ports 17–21 are populated. Avoid open slots on ports 16–21.
- Use BIOS version that aliases Completion timeouts to UR or apply the factory provided BIOS workaround.
- The following workaround keeps the unpopulated port in internal loopback mode to get the proper UR response back.
  - EEPROM puts station2 port (port 16) in internal loopback mode. For port16 being a x16 port:
    - Write port 16 Offset 204h Bit[31:16] to 0xFFFF
    - Port16 Offset 204h Bit[15:0] to 0xFFFF
    - Write Port16 Offset BD8h Bit[0] to 1b
    - Write Port16 Offset 228h Bit[8] to 1b
  - When BIOS detects PEX 87xx upstream port, it goes and removes all station2 ports out of Loopback.
    - Write Port16 Offset 204h Bit[31:16] to 0x0000
    - Write Port16 Offset 204h Bit[15:0] to 0x0000

  - Write Port16 Offset BD8h Bit[0] to 0
  - Write Port16 Offset 228h Bit[8] to 0
- Then it waits for 50 ms. Then it checks the DL_Active status on every station2 downstream port.
  - Read Port16 to Port21 offset 0x78 Bit[29]
- If status is not true for a port (port is down), software enables the corresponding PEX 87xx port into internal loopback mode.
  - Perform Strep1 programming sequence
- If status is true, device is populated to the slot. Do not enable the Internal loopback for this cause.

Impact: System may fail to boot up if BIOS does not alias Completion timeout to UR.

1.3 Recovery Equalization Phase 1 Timeout for Downstream Ports Is Not 24 ms

Risk Category: Low

Silicon Revisions Affected: BA

Description: In phase 1 of recovery equalization for downstream ports of the switch, if the remote device does not progress to phase 2 in 12ms, the switch downstream port times out and goes back to Recovery.speed.state. The specification defined timeout is 24ms. The specification says that the downstream device should lock to the incoming signal within 2 ms and then progress to the next state after looking at two TS1 ordered sets.

Solution/Workaround: None.

Impact: If the device connected to the downstream port of the switch is following the specification, it should not have an impact to the system.

1.4 Hot-Plug HP_PWREN*/HP_PERST# Pin State Is Unreliable During Fundamental Reset

Risk Category: Low

Silicon Revisions Affected: BA

Description: Hot-plug *power enable* pin and Hot-plug PERST# pin requires 300 µs from PEX_PERST# deassertion to stabilize to final value. They are unknown during and also 300 µs after fundamental reset.

Solution/Workaround: Delay PEX_PERST# deassertion by 300 µs using external circuit and do logical OR/AND operation with HP_PWREN/HP_PERST# pins respectively.

Impact: HP_PWREN/HP_PERST# pin is in unknown state for 300 µs after fundamental reset.

1.5 STRAP_I2C_SMBUS_EN_P Pin Does Not Enable/Disable the SMBUS ARP Mode as Expected

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: Strap pin STRAP_I2C_SMBUS_EN_P does not control the SMBUS ARP enable/disable functionality.

Solution/Workaround: Software/EEPROM/I2C can control the SMBUS ARP mode by programming Port0 Offset 2C8h bit[8].

Impact: STRAP_I2C_SMBUS_EN_P should not be used to control the SMBUS ARP mode.

1.6 Memory Read TLPs with Relax Ordering (RO) Attribute Gets Completion Timeout under Extreme Completion Credit Starvation Condition

Risk Category: Low

Silicon Revisions Affected: BA

Description: Completion without RO can block Completion with RO in a corner case scenario.

Errata is not applicable in the following conditions:

- If a switch port is ‘only’ transmitting RO completions or ‘only’ transmitting non-RO completions.
- Devices connected to switch advertise infinite credits for completions (usually the case for End Points and Root Complex).

The Memory read TLP with RO gets completion timeout if the switch destination port hits all the following conditions:

- Switch destination port gets both completion with RO and without RO TLPs.
- The device on the other end of the link advertises non-infinite completion credit. (It is not an endpoint.)
- Completion credit is throttled to the extreme (one completion TLP progress every few micro seconds).

Solution/Workaround: Disable *Completion special handling for RO* by programming Port0, Port8, and Port16 offset 760h Bit[29] to 1’b1.

Impact: Completion timeout for Memory Read TLP with RO.

1.7 Ports Do Not Advertise All the Supported Data Rates in Polling/Configuration State if Target Link Speed Is Programmed to Gen2 or Gen1 Rate, They Only Advertise the Programmed Target Link Value

Risk Category: Low

Silicon Revisions Affected: BA

Description: When the Target Link Speed in the Link Control 2 register is set to a value other than 3 (Gen3), the Switch ports do *not* advertise all their supported data rates during Polling and Configuration LTSSM substates.

Solution/Workaround: Program the Target Link speed to Gen3 to advertise all the supported data rates.

Impact: No real time issue working with other compliant devices expected.

1.8 When 2VCs Are Enabled in the System, VC1 Nonposted/Completion Traffic Could Block VC0 Posted if VC1 Nonposted/Completion Runs Out of Egress Credit

Risk Category: Low

Silicon Revisions Affected: BA

Description: VC1 nonposted/completion traffic blocks VC0 posted if VC1 nonposted/completion runs out of egress credit. VC0 posted queue makes progress when VC1 nonposted/Completion start progressing with egress credit.

Solution/Workaround: None.

Impact: VC0 queue blocked by VC1 nonposted/completion congestion. This errata does not affect posted traffic.

1.9 NT Virtual Endpoint Incorrectly Logs the Receiver Errors in Recovery State

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: The PEX 87xx ports should not log receiver error in recovery state according to the SPEC. But NT virtual endpoint logs the receiver error in recovery state. PEX 87xx transparent ports and link endpoints logs the receiver error properly.

Solution/Workaround: None.

Impact: Unexpected receiver error could be triggered due to rate change.

1.10 ASPM L0s Will Not Work at Gen1, Gen2, and Gen3 Rates

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: When the receive side of the device comes out of L0s, it goes into recovery. The PEX 87xx receivers need more time to lock than what the max NFTS allows. For this reason, the ASPM L0s will not work with the device.

Solution/Workaround: Use ASPM L1 for power saving in the system (with the workaround suggested in Section 2.10, PEX 87xx Downstream Port Does Not Change the Link Speed from Gen2 to Gen3 in Single Retrain for the Specified Rate Change Sequence).

Impact: Power savings cannot be achieved when there are short periods of inactivity. For longer periods of inactivity, power savings can be achieved by ASPM L1.

1.11 If Software Initiates a Speed Change to Gen3 on the Downstream Port of the Switch after Links Are Up, DLLP and TLP Transmission Could Be Blocked for the Some Conditions

Risk Category: Low

Silicon Revisions Affected: BA

Description: The following are the conditions for this issue:

- Downstream port link partner is programmed to Gen1/2 rate and the link is up at Gen1/2 rate.
- Software programs the downstream port link partner “target link speed” to Gen3 and sets the switch downstream port Retrain Link bit.

Downstream Port goes to recovery state and does not change the data rate to Gen3 and incorrectly blocks DLLP/TLP transmission coming out of recovery.

Solution/Workaround: Always program the switch downstream port link partner “target link speed” to Gen3 and control the switch downstream port “target link speed” for speed change.

Impact: Switch port blocks DLLP and TLP transmission if switch downstream port initiates the speed change for the specified programming.

There is no impact if the upstream device initiates the speed change to the switch upstream port or on normal link ups to Gen3 out of reset.

1.12 When Virtual Switch Mode Is Enabled, VS_PERST# Does Not Initialize the Serial Hot-Plug Register

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: In virtual switch mode, the PEX 87xx scans IO expander and its values when PEX_PERST# deasserts. The PEX 87xx does not re-scan the IO expander again when VS_PERST# deasserts. In this case, the IO expander’s setting does not get loaded into the Serial Hot-plug register. In other words, the initial values can be wrong.

Solution/Workaround: Connect VS_PERST# to either HP_BUTTON or HP_MRLn inputs of one of the serial expanders with the corresponding hot-plug pin functionality disabled. The hot-plug pin functionality can be disabled either by EEPROM loading Slot Capability register or by software clearing the corresponding event enable bit in Slot Control register.

Impact: Serial Hot-plug register reset value could be incorrect after VS_PERST# deassertion.

1.13 Training Set Corruption during Link Recovery Could Cause PEX 87xx Port to Get Stuck in Recovery.speed State or Go to DL_Down State

Risk Category: Low

Silicon Revisions Affected: BA

Description: When the speed changes from Gen2/3 to a lower data rate, if the received training ordered set – ‘data rate field’ is corrupted due to noise on the link, the PEX87xx port could get stuck in recovery.speed state. This issue could also occur if a device connected to the PEX87xx is removed (hot-removal).

When the speed changes from Gen2/3 to a lower data rate, if the received training ordered set – ‘training control bits field’ (TCB) is corrupted due to noise on the link, the PEX87xx port could go to DL_Down state or get stuck.

Solution/Workaround: If the received training ordered set – ‘data rate field’ is corrupted, program the corresponding port’s ‘port disable’ CSR bit (program port0, port8, port16, offset 0x208 bits[5:0]) to transition the port to DL_Down state. After DL_Down state transition, clear ‘port disable’ CSR bit for PEX87xx port to link up.

For a workaround to the ‘training control bits field’ (TCB) corruption, before software initiates speed change to Gen1, set switch port0, port8, and port16 0x22C, bit[20]. After the speed change is done clear the bit.

Impact:

- Gen2/3 to a lower data rate may not work.
- Hot-removal does not consistently bring down the link.

1.14 PEX 8733/8717/8713 Reports More Lanes/Ports

Risk Category: Low

Silicon Revisions Affected: BA

Description: PEX8733/8717/8713 reports more lanes and ports than specified in the data sheet.

Solution/Workaround: PEX 87xx silicon revision CA corrects the number of ports and lanes reported.

Impact:

- Maximum Link Width (Link Capability) register is inaccurate; however links will come up at the expected widths as per datasheet. For example, x2 port will report a maximum link width of x4 but will correctly link up at x2.
- More Switch downstream ports appear in PCIe hierarchy

1.15 PEX 87xx Port “Link Status 2” Register’s Equalization Status Bits Have Incorrect Attribute

Risk Category: Low

Silicon Revisions Affected: BA

Description: Link Status 2 register's equalization complete, equalization phase1 successful, equalization phase2 successful, and equalization phase3 successful status bits are not implemented with ROS (read only sticky) attribute. They are implemented with RW1C attribute. Link Status 2 register’s Link Equalization Request status bit is not implemented with RW1CS attribute, instead it has RW1C attribute. Software should write port offset 0x98 Bit[5] to 1b to clear the Link equalization Request instead of writing Bit[21].

Solution/Workaround: None.

Impact: The status logged as expected as long software does not write to this offset.

1.16 PEX 87xx Disables Inactive Lanes upon Entry to Configuration Complete State

Risk Category: Low

Silicon Revisions Affected: BA

Description: In case of down configuration, the PEX 87xx disables inactive lanes upon exit from COFIGURATION.COMPLETE state.

Solution/Workaround: None.

Impact: Software initiated dynamic link width adjustment to save power may need to be disabled if the external device cannot handle this behavior.

1.17 PEX 87xx Link Goes Down If Lane0 Breaks when the Link Is Up in Lane Reversal Mode

Risk Category: Low

Silicon Revisions Affected: BA

Description: If the PEX 87xx port links up in the lane reversal mode and lane0 is broken, the link does not negotiate down. Instead, the Switch port times-out in configuration state and go back to detect state.

Solution/Workaround: None.

Impact: Surprise DL_Down.

1.18 PEX 87xx Port Does Not Properly Acknowledge the Link Partner's "Use Preset" Request

Risk Category: Low

Silicon Revisions Affected: BA

Description: If the link partner requests the switch port to use different preset (non-default preset), the switch port applies the new requested preset to its transmitter, but does not return the right value for use_preset/coefficients parameters for the requested preset while the protocol requires this to acknowledge the acceptance.

Solution/Workaround: Use coefficients that correspond to the desired preset for tuning.

Impact: The switch port does not properly acknowledge the "use preset" request.

1.19 PEX 87xx Port Does Not Reject Illegal Coefficients for the Specified Condition

Risk Category: Low

Silicon Revisions Affected: BA

Description: PEX 87xx should reject the illegal coefficients if the coefficients yield negative value for C0 -|C-1|-|C+1 |>= LF rule. PEX 87xx port does not reject these illegal coefficients.

Solution/Workaround: None.

Impact: Devices should not use illegal coefficients by the PCIe specification. PCI-SIG compliance software fails due to this issue.

1.20 PEX 87xx Software Initiated Up-Configuration Does Not Work Reliably with Electrical Idle Inference Mode

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: The PEX 87xx enables electrical idle inference mode by default (in Rev-CA). Hardware/software initiated link width down-configuration works fine with default mode. But link width up-configuration does not reliably bring-up the link to its full width due to analog CDR (Clock and Data Recovery) circuit limitation with electrical idle inference mode.

Solution/Workaround: PEX 87xx autonomous/software initiated link width up-configuration works without Electrical Idle inference mode. EEPROM or Software can be used to disable EI_inference mode (EI_inference is needed for long channels; see Section 2.1, A Separate EEPROM Image Is Recommended to Run the Links at Gen3 When the Channel Length Is Greater than 10 In.).

Software workaround for EI_inference mode:

Down-configuration:

- Change the target_link speed to Gen1 (write to port offset 0x98 Bits [3:0] to 0x1).
- Change the current operating link speed to Gen1 (write to port offset 0x78 Bit[5] to 1b).

– Disable the El_inference mode (PEX 87xx Port0/Port8/Port16 offset 0x204 Bits[15:0] to 0x0000[1]).
– Drop the Link Width.

Up-configuration:

– Bring-up the Link to full width.
– Enable the El_inference mode (PEX 87xx Port0/Port8/Port16 offset 0x204 Bits[15:0] to 0xFFFF[1]).
– Change the target_link speed to Gen3 (write to port offset 0x98 Bits[3:0] to 0x3).
– Change the current operating link speed to Gen3 (write port offset 0x78 Bit[5] to 1b).

Impact: Software initiated link width up-configuration does not work reliably with Electrical Idle Inference mode. Note that this errata has no impact on hardware triggered link width changes (link up, bad links, and so on).

1.21 Gen3 x8 Link Could Cause LCRC Errors (Non-Fatal Recoverable) on Port1, Port9, Port17

Risk Category: Medium

Silicon Revisions Affected: BA

Description: At Gen3 speed on specific x8 ports, if there are at least eight back-to-back short TLP's followed by a long TLP going out of a link on a specific corner case sequence, the device could send the long TLP out with LCRC error. Short TLP is a 3 or 4 Double Word TLP.

Solution/Workaround: If it is possible to have less than eight short TLPs going back to back on the link, this issue could be avoided.

Impact: Gen3 x8 port1, port9, and port17 link could have LCRC error. The device connected to this x8 port could see a TLP with CRC error due to this issue. The connected device is required to "NAK" the TLP, which will cause a link re-transmission of the correct TLP and will be transparent to system software unless Bad TLPs are being monitored. Gen1 and Gen2 links are not affected by this errata.

1.22 PEX 87xx Broken Lanes Issue for the Specified Conditions

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description:

- If the PEX 87xx port links up in the lane reversal mode and lane0 is broken, the link does not negotiate down. Instead, the Switch port times-out in configuration state and goes back to detect state.
- When two PEX 87xx devices are connected to each other (in x2, x4, x8, and x16 configuration) and the middle lanes are broken, but both of the PEX 87xx transmitters still detect the receivers (possible if there is a retimer on the link), the link will never come up. The link should have down-trained and come up as a x1 link, but it does not.

1. Each bit controls the corresponding lane. If the lanes are not part of the specified port, leave the default value for the corresponding lanes

Solution/Workaround:

- None.
- Disable the termination seen by the PEX 87xx transmitter for the broken lanes.

Impact:

- Surprise DL_Down event.
- Link never comes up.

1.23 PEX 87xx Could Log Receiver Errors at Gen3 Data Rate in Asynchronous Systems with Protocol-Aware Retimers that Send Variable Length SKP OS

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: The PCIe device must transmit SKP OS block with 16 SKP symbols at Gen3 data rate. This rule does not apply to retimers. Protocol-aware retimers could be adding or deleting the SKP symbols when communicating to two different PCIe devices in an asynchronous clocking system. If the scrambled EDS (End of Data stream) pattern looks like SKP OS followed by the actual SKP OS, block alignment logic does not detect SKP at all. The block alignment works fine if the SKP OS has full 130-bit block for this case. But if the SKP OS block is not equal to full block size, the block aligner loses its alignment because it does not detect SKP OS properly.

Solution/Workaround:

- Synchronous clocking systems are not exposed to this errata.
- Ensure Retimer + RC/other PCIe device are in one clock domain and PEX 87xx is in a different clock domain.

Impact: When the port loses block alignment, it logs receiver error, eventually go to recovery state, and achieves the block alignment in recovery state. There is no speed/width change or no DL_Down due to this state transition.

1.24 Memory Read Requests to Expansion ROM Address xxxxxc90h Returns Incorrect Value Instead of Expansion ROM Value

Risk Category: Low

Silicon Revisions Affected: BA

Description: Memory Read Request to Expansion ROM space address xxxxxc90h should return the Expansion ROM content; however, the PEX87xx returns the corresponding port configuration space register offset C90h value.

Solution/Workaround: Do not write the execution code to the EEPROM for the above address. Shift writing the execution code to the EEPROM for the error address + 4.

Impact: Does not read the specified Expansion ROM content.

1.25 PEX 87xx Port Does Not Reliably Link-Up if Its Receiver Does Not Break EI within 28 ms in Gen3 Speed

Risk Category: Low

Silicon Revisions Affected: BA

Description: If a device detects receiver and does not break electrical idle (EI) within 28 ms during link bring-up, port supposed to transition to compliance state and stay until electrical idle breaks. PEX87xxBA port does transition to compliance state and links up properly for first three compliance state entries after fundamental reset. But the port does not link-up from fourth compliance state entry, because compliance data rate changes to Gen3 speed.

Solution/Workaround:

- Use fundamental reset instead of hot-reset or Link-disable.
- Set Enter Modified Compliance bit by programming offset 0x98h bit Bit[10].

Impact: If the port enters the compliance state four or more times, it does not link up.

1.26 STRAP_GEN1_GEN2_P Pin Does Not Enable Gen2 Only Mode with Port Bifurcation Feature Enabled

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: STRAP_GEN1_GEN2_P pin is supposed to enable Gen1 only, Gen1/Gen2 only, or Gen1/Gen2/Gen3 data rate mode.

STRAP_GEN1_GEN2_P pin does not enable Gen1/Gen2 only mode.

Solution/Workaround: EEPROM can be used to program Port0 offset 0x46C, Bit[12] to 0, and also to program each port's offset 0x74 Bits[3:0] to 2.

Impact: STRAP_GEN1_GEN2_P pin should not be used to configure Gen1/Gen2 only mode.

1.27 PEX 87xx NT Link/NT Virtual Endpoint PCI Compatible Vendor-Specific Compatibility Header Is Incorrect

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: PCI-compatible Vendor-specific capability header must present the capability length field a byte immediately following the Next Capability pointer field. But PEX 87xx NT Link/NT virtual endpoints PCI compatible vendor-specific capability header in offset 0xC8 present the length field in next CSR offset (0xCC).

Solution/Workaround: No issue is found so far across multiple operating systems. If an issue is found, the capability can be bypassed using EEPROM for configuration access and use memory mapped access to configure the capability.

Impact: None.

1.28 PEX 87xx Incorrectly Logs Receive Error During Link Width Up-Configuration

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: PEX 87xx transitions the lanes from electrical idle state to active state during link width up-configuration. The lanes' error reporting logic should be disabled until the lanes achieve bit lock and symbol lock/block alignment. But the port logs code error as soon as the lanes exit electrical idle state.

Solution/Workaround:

- Clear the receive error as soon as the link width up-configuration completes.
- Do not do link width up-configuration.

Impact: The PEX87xx ports log unexpected receiver error after link width up-configuration.

1.29 PEX 87xx Does Not Log DMA Payload Link List RAM ECC Error Status

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: PEX 87xx DMA generates TLPs targeting external devices. These TLPs are stored in DMA RAMs. Payload link list RAM is used to order the TLP payload within a TLP. Payload link list RAM content is ECC-protected with 1-bit error correction and 2-bit error detection. ECC 1-bit error correction works normally and 2-bit fencing works normally as defined in the "fencing document"/Databook. The ECC error status is not logged into the configuration space register for debug/software visibility.

Solution/Workaround: None.

Impact: No debug/software visibility for DMA payload link list RAM ECC error.

1.30 PEX 87xx DMA Channel Requires Hot-Reset If It Sends TLP to a PEX 87xx Port with Memory/Master Enable Cleared or in DL_Down State

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: If PEX87xx DMA TLPs target the PEX87xx port with memory/master enable cleared or in DL_Down state, the PEX87xx port should drop the DMA TLP with Unsupported Request (UR) error. The PEX87xx port should also inform the DMA channel to free up the memory resource with a DMA abort error status.

While the PEX87xx does report UR error status, the PEX87xx DMA channel does not free up the memory resource with a DMA abort error status.

Solution/Workaround: PEX87xx requires a hot-reset or fundamental reset to recover PEX87xx DMA operation after the failure.

Impact: It is not normal for the DMA channel to communicate to the PEX87xx port with memory/master enable cleared or in DL_Down state. If such a communication occurs, the PEX87xx needs a hot-reset/fundamental reset to recover.

1.31 PEX 87xx DMA Descriptor Size Less than 32 Bytes Could Cause Data Corruption

Risk Category: Low

Silicon Revisions Affected: BA

Description: PEX87xx is supposed to handle descriptor size up to 0 bytes (zero-length packet). But PEX87xx has a memory pointer de-allocation issue if DMA schedules back to back short TLPs less than 32 bytes payload. This incorrect memory pointer de-allocation corrupts TLP payload of any TLPs inside the device.

Solution/Workaround: Perform the following steps to work around this issue:

1. DMA Descriptor size should be 32 bytes or more.
2. DMA data Read Request starting address to RCB/4KBoundary gap should be 32 bytes or more.
3. All descriptors use extended descriptor format (even for the 32-bit addressing).
4. DMA completion “completion status” is successful. If PEX87xx receives too many DMA completion with unsuccessful status, this error condition could trigger.

Impact: DMA descriptors of less than 32 bytes cannot be used.

1.32 PEX 87xx Downstream Port Incorrectly Sets Link Status Register’s “Link Autonomous Bandwidth Status” Bit for Any Successful Speed Change Event

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: PEX87xx is supposed to set downstream port Link Status register’s Link Autonomous Bandwidth Status field if successful_speed_negotiation is set to 1b and the Autonomous Change bit (bit 6 of Symbol 4) is set to 1b for eight consecutive TS2 Ordered Sets received in Recovery.RcvrCfg state. But PEX87xx downstream port always sets Link Status register’s Link Autonomous Bandwidth Status field for any successful speed change event.

Solution/Workaround: None.

Impact: Downstream Port Link Status register’s Link Autonomous Bandwidth Status field set for any successful speed change event.

1.33 Serial Hot Plug Could Fail if Implemented on More Than One Port

Risk Category: Medium

Silicon Revisions Affected: BA, CA

Description: If Serial Hot Plug is implemented on more than one Port, the I²C bus could send the I/O configuration data (from each port's Slot Control register, offset 7Ch) to the wrong I/O Expander.

Solution/Workaround: To implement Serial Hot Plug on more than 1 port, the following restrictions apply. These restrictions are not needed when only one port uses Serial Hot Plug.

- Any workaround requires custom software that writes to one I/O Expander at a time, avoiding I^2C arbitration.
- All slots must be powered up individually by software, rather than using the automatic power-up sequencing (enabled by MRL# being low). The EEPROM must Set the Power Controller Control bit (80h[10]) in all Hot Plug Ports, to disable automatic power-up sequencing. Then software can clear the bit in each Port. Command Complete interrupts or register polling can be used to determine that the I^2C write has completed (and so the next I/O Expander can be written).
- The LED Blink command must not be used for the Attention and Power LEDs, on more than one port at a time; if LEDs for more than one port need to blink concurrently, software must not write the value 10b to register 80h bits [9:8, 7:6]. Instead, software should turn the LEDs on and off to cause blinking (using bit values 01b and 11b, respectively).
- Alternatively, Hot Plug can be implemented without I/O Expanders, by using Presence Detect and Data Link Layer State Change interrupts.

Impact: Serial Hot Plug can be implemented for one slot only. This issue does not affect Parallel Hot Plug ports.

1.34 The DMA Payload RAM Assignment Among DMA Functions Need to Be Programmed Before Enabling the DMA Operation

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: PEX87xx DMA payload RAM is shared among the enabled DMA functions. The DMA payload RAM has 118 payload addresses. Each payload address can store up to 64 bytes of TLP payload. Software needs to divide the DMA payload RAM among the enabled DMA functions before enabling the DMA operation by programming DMA function's offset 0x258 Bits[23:16]. The sum of all DMA function's offset 0x258[23:16] should not exceed 118 (decimal) value.

Solution/Workaround:

- Example programming for 4 channel mode to get equal performance among the DMA functions. Program DMA function's offset 0x258 Bits [7:0] to 0x18 and Bits[23:16] to 0x1C.
- Example programming for 2 channel mode to get equal performance among the DMA functions. Program DMA function's offset 0x258 Bits [7:0] to 0x30 and Bits[23:16] to 0x3A.
- Example programming for 1 channel mode to get equal performance among the DMA functions. Program DMA function's offset 0x258 Bits [7:0] to 0x60 and Bits[23:16] to 0x75.

Impact: These registers must be programmed prior to using the DMA channel. Otherwise, the DMA behavior is undefined. One of the undefined behaviors is data corruption.

1.35 PEX 87xx Downstream Port Does Not Change the Link Speed from Gen2 to Gen3 in Single Retrain for the Specified Rate Change Sequence

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: If software changes PEX87xx downstream link data rate from Gen1 to Gen2 to Gen3 by changing both upstream and downstream component's target_link_speed, Gen1 to Gen2 speed change works fine. But PEX87xx downstream port does not initiate speed change from Gen2 to Gen3 with single retrain. It expects software to write Retrain Link bit twice to initiate Gen2 to Gen3 speed change.

Solution/Workaround:

- Keep the downstream component's target_link_speed to its maximum supported data rate. Change only PEX87xx downstream port's target_link_speed to change the link data rate.
- If Software needs to change the downstream component's target_link_speed, it needs to write PEX87xx downstream port Retrain Link bit two times.

Impact: PEX87xx downstream port does not initiate Gen2 to Gen3 speed change with single Retrain Link.

1.36 PEX 87xx Port16 to Port21 Do Not Log the Specified Errors

Risk Category: Low

Silicon Revisions Affected: BA

Description: PEX87xxBA Port16, Port17, Port18, Port19, Port20, and Port21 do not log the following errors:

- Replay Timeout.
- Replay Number Rollover.
- Downstream port does not log Poisoned TLP/ECRC Error status for downstream TLPs.
- Upstream port does not log Poisoned TLP/ECRC Error status for upstream TLPs.
- NT Port does not log Poisoned TLP/ECRC Error status for virtual to Link direction TLPs.
- The TLP targeting DL_Down port does not report Unsupported Request while dropping posted/nonposted TLPs. There is no error completion generated for nonposted TLP drop.

The preceding errors do not generate any error message. The Replay Number Rollover error status does initiate the Link recovery protocol.

Solution/Workaround: None.

Impact: Port16 to Port21 do not log the specified errors.

1.37 PEX 87xx Implements MC_Base_Address as Defined in the Description

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: PCIe specification says MC_Base_Address should be implemented as 2 ^ MC_Index_Position. But PEX8xxx implements MC_Base_Address as follows:

- If MC_Num_Group is 0x0, MC_Base_Address is implemented as 2 ^ (MC_Index_Position + 0)
- If MC_Num_Group is 0x1, MC_Base_Address is implemented as 2 ^ (MC_Index_Position + 1)
- If MC_Num_Group is 0x2 to 0x3, MC_Base_Address is implemented as 2 ^ (MC_Index_Position + 2)
- If MC_Num_Group is 0x4 to 0x7, MC_Base_Address is implemented as 2 ^ (MC_Index_Position + 3)
- If MC_Num_Group is 0x8 to 0xF, MC_Base_Address is implemented as 2 ^ (MC_Index_Position + 4)
- If MC_Num_Group is 0x10 to 0x1F, MC_Base_Address is implemented as 2 ^ (MC_Index_Position + 5)
- If MC_Num_Group is 0x20 to 0x3F, MC_Base_Address is implemented as 2 ^ (MC_Index_Position + 6)

This causes no multicast hit if the MC_Base_Address programming is not aligned to the power of MC_Index_Position+n (where n is 0 to 6).

Solution/Workaround: MC_Base_Address programming needs to be aligned to the power of MC_Index_Position+n.

Impact: No multicast hit if the MC_Base_Address programming is not aligned to the power of MC_Index_Position+n.

1.38 In the Default Mode, PEX 87xx Incorrectly Calculates the Gen3 TS1 Parity in Response to Use_Preset Equalization Request from the Link Partner

Risk Category: Low

Silicon Revisions Affected: CA

Description: If PEX87xx port receives the use_preset request from the link partner during Equalization Phase2 and Phase3, it responds by transmitting the TS1 with same preset from the use_preset request and the calculated coefficient with respect to PEX87xx port FS value. The parity equation uses the received coefficient fields instead of calculated coefficient fields causes the parity issue. If the link partner tunes the PEX87xx transmitter using coefficients request instead of preset requests, the parity is calculated right.

Solution/Workaround:

- PEX87xx port supports a nondefault mode by responding with receive preset and coefficients from the use_preset equalization request from the link partner. Program Port0/Port8/Port16 offset 0xBE8 bit[11] to 1’b1.
- Link partner to use coefficient request instead of use_preset request for tuning PEX87xx transmitter.

Impact: The TS1 transmitted by PEX87xx will get rejected by link partner.

1.39 False Detection of Receiver for Long Cable Lengths

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: PEX87xx devices can falsely detect receivers when there is no endpoint connected to them even though there is a long cable attached to the TX pins (high trace capacitance). This has been seen when cables attached to the TX pins are typically more than 5 feet in length. With this modification, cables in excess of 3m to 5m can still result in false detection.

Solution/Workaround: Write the following BFC register values to adjust the detection timing/pulse within each quad.

- For Quad 0:
  - 0x8A4E6F
  - 0x8A7BA0
  - 0x8A7EFC
  - 0x8A7FD3
- For Quad 1:
  - 0xAA4E6F
  - 0xAA7BA0
  - 0xAA7EFC
  - 0xAA7FD3
- For Quad 2:
  - 0xCA4E6F
  - 0xCA7BA0
  - 0xCA7EFC
  - 0xCA7FD3
- For Quad 3:
  - 0xEA4E6F
  - 0xEA7BA0
  - 0xEA7EFC
  - 0xEA7FD3

All lanes of a configured port must be set to the same value, even if a subset of lanes are connected to that port (that is, a x8 port with only a x4 cable connected will still need two quads programmed for alternate detection setting).

Impact: Customers using long cables to PCIe links can falsely detect receivers when there is nothing connected to the end of the cable.

1.40 PEX 87xx Upstream Port Lane Reversal with Link Width Down Negotiation Does Not Work for the Specified Conditions

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: The PEX 87xx upstream port does not link up if its lanes are reversed, the remote device tries to link-up in lower link width, and the remote device sends training sequence 1 (TS1) with Link number set to non-PAD on all the lower lanes that it intends to negotiate-out in configuration.Linkwidth.Start state. Typically, the remote device sends Link number set to PAD on all the lanes that it intends to negotiate-out in configuration.linkwidth.start. If the PEX87xx upstream port receives TS1 with Link number, set to PAD in configuration.linkwidth.start, it links-up fine.

Solution/Workaround:

- Disable the lane reversal.
- Configure the PEX 87xx upstream port programmed link width to the remote device's link width.

Impact: The PEX87xx upstream port does not link up for the preceding specified condition.

1.41 PEX87xx Does Not Enter ASPM-L1/PCI-PM-L1 State If the Nonposted Payload Credit Available from Remote Device Is Less than 1 MPS

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: PEX87xx is supposed to enter ASPM-L1 or PCI-PM-L1 state with other conditions satisfied if the nonposted payload credit available from remote device is sufficient enough to transmit a maximum payload architected nonposted TLP. Nonposted TLP's maximum payload size could consume up to two payload credits. But PEX87xx waits for 1-MPS size before it enters ASPM-L1/PCI-PM-L1 state.

Solution/Workaround: Have remote device to advertise at least 1 MPS payload credit for nonposted payload type.

Impact: No L1 state power saving if the nonposted payload credit available from remote device is less than 1 MPS. This issue was observed with the use of the Intel Haswell platform. With Haswell platform, L1 entry is not possible.

1.42 PEX87xx DMA Could Corrupt Memory Write TLP Payload of an Unaligned DMA Descriptor while More than One DMA Channel Is Actively Transferring TLPs with TLP Addresses Greater than 32-Bit Addressing

Risk Category: Medium

Description: PEX87xx DMA could corrupt one double word of memory write TLP payload when all the following conditions occur at the same time:

- PEX87xx DMA enabled for multi-channel operation and more than one DMA channel is actively transferring TLPs and the DMA descriptor destination address field is greater than 32-bit addressing.

- DMA descriptor's source_address, destination_address, or transfer size is not aligned to RCB boundary
- Root complex does not support ECRC generation for the completion TLPs
- Received Completion TLP occupies the full internal data bus. (CplD TLP size is 3DW Header + 2DW Payload + *n* × 4DW Payload, where *n* is 0, 1, 2, 3, and so on.)

Solution/Workaround: Use any one of the following workarounds.

- Use DMA in one channel mode only.
- Use only 32-bit destination addressing.
- DMA descriptor's source address, destination address, and transfer_size are RCB (Read Completion Boundary) aligned.
- Enable End-to-End CRC (ECRC) in the completer (Root Complex normally).
- For PEX8717/8713 enable Station 0/1, only use ports in Station 0. This is done by a combination of strap signals and Serial EEPROM.
  a. On the chip interface, set STRAP_STN0_PORTCFG[1:0] = 2'b00.
  b. On the chip interface, set STRAP_STN1_PORTCFG[1:0] = 2'b00.
  c. In the Serial EEPROM, set register 300h bits [2:0] (Port Configuration) Station 0 as desired; refer to Section 15.16.7 in the Databook for valid port configurations. Port 12 and Port 13 in Station 1 will be visible, but they are not functional and should not be used.
  d. In register 6Ch[2:0] (Device Capability), set the "maximum payload size" to match the number of enabled ports; refer to Section 19.8 in the Databook.
- Use the Broadcom DMA Descriptor validation function (software) to process the DMA descriptors. The Broadcom DMA Descriptor validation function examines the DMA descriptors and breaks the DMA descriptors on certain boundaries if the DMA descriptors have a possibility of encountering this hardware bug.

NOTE: Other possible methods for workarounds for PEX8733/25/17/13 that do not require board changes, software functions, or configuration limitations are being investigated. Contact your FAE for possible alternatives that might be in development.

Impact: Data corruption is possible for the specified condition.

1.43 PEX9700 and PEX8700 products Load the Port Arbitration Table without the Host Setting the Port Arbitration Table Load Bit

Summary: PEX9700 and PEX8700 products load the Port Arbitration Table without the host Setting the Port Arbitration Table Load bit.

Description: PEX9700 and PEX8700 products load the Port Arbitration Table without the host setting the Port Arbitration Table Load bit (register 0x15C bit [16]) under one of the following conditions:

- When link power state changes from L1 to L0.
- When link speed change occurs.
- When link width change occurs.

Impact: If invalid or erroneous values are programmed into the Port Arbitration Table entries, these values will automatically get loaded under the above conditions and unintentionally impact port arbitration and packet transport through the switch, based on the values of these bits. In worst case, this could block traffic on all ports.

Workaround: When software programs the Port Arbitration Table, it must program legal values based upon intended port arbitration, *or* bypass the entire VC Capability Registers, using EEPROM or I^2C to program the Next Capability Offset field of the previous capability in the list. (Because the previous capability is Secondary PCI Express Extended Capability, program 10Ch[31:20] to E00h, which addresses the Multicast Extended Capability Registers instead, bypassing the Virtual Channel Capability Registers.)

Risk: Low

Chapter 2: Cautions

NOTE: Cautions are not Errata. In Cautions, the Switch is not violating the PCIe Specification. However, our Switch may be behaving unexpectedly due to another non-compliant device connected to the Switch or due to ambiguity in the Specification itself.

2.1 A Separate EEPROM Image Is Recommended to Run the Links at Gen3 When the Channel Length Is Greater than 10 In.

Risk Category: Low

Silicon Revisions Affected: BA

Description: Use EEPROM Image: 1.X for short channels.

Total channel length < 10 in (8 in. with 1 PCIE connector, Assumption <9 dB loss)

NOTE: Some shorter channels could have more loss based on material and discontinuities.

Use Long Channel EEPROM Image: 2.X (Electrical Idle Inferred mode) for longer channels.

Total channel length greater than 10” or variable channel length

Note that in this mode:

- Software initiated Link Disable/Enable or hot-reset does not work properly if the device connected to PEX87xx does not send EIOS (Electrical Idle Ordered Set) before transitioning to electrical idle on its transmitter.

  To workaround this issue, PEX87xx needs the following sequence:

  Step 1 – If the link is at Gen3 speed, then initiate speed change to Gen1 and clear the inferred mode bit (Port offset 98h Bit[3:0] to change the target link speed, Port0/Port8/Port16 offset 204h Bit[15:0] = 16’h0 for clearing inferred mode).

  Step 2 – Execute Link Disable/Enable/Reset sequences. After completion, set Inferred mode bit in Gen1 speed (L0 state) before initiating speed change back to Gen3.

NOTE: PEX 87xx does not end the hot-reset sequence with EIOS before entering electrical idle state (so back-to-back PEX87xx links in EI inferred mode need the workaround).

- Hot-plug insertion needs the following sequence:

  Step 1 – If the target_link_speed is at Gen3 speed, then change the target_link_speed to Gen1. After changing to Gen1 speed, clear the inferred mode bit. (Port offset 98h Bit[3:0] to change the target link speed, Port0/Port8/Port16 offset 204h Bit[15:0] = 16’h0 for clearing inferred mode). You can then safely remove the device.

  Step 2 – Insert the new device and wait for the link to come up in Gen1 speed. Set Inferred mode bit before initiating speed change back to Gen3.

2.2 Link to Intel Based Systems Is Limited to 19-dB Insertion Loss Channel

Risk Category: High

Silicon Revisions Affected: BA, CA

Description: Due to margin requirements when interfacing to Intel Xeon processor E5-2600 product family (formerly codenamed Romley) and future Intel Xeon processor E5 family products based PCIe Gen3 systems, it is recommended to limit the channel insertion loss at 4 GHz to 19 dB (note that this is 4 dB less than the maximum channel loss communicated in the *Intel Product Design Guide*). The transmitter may not provide enough boost and the proper proportions of pre-shoot and de-emphasis for the signal to be correctly received by the Intel receiver for channels with greater insertion loss than 19 dB. This 19 dB applies to the entire channel between PLX TX and Intel RX, including package loss. PLX switch package loss is 0.5 dB.

Note that the channel loss restriction does not apply in the other direction (Intel TX to PLX RX) because the PLX receivers do not have the same restriction as the Intel receivers. Loss between two PLX devices can be greater than 25 dB for well-designed channels.

2.3 PEX 87xx Port Replay Timer Equation Does Not Include “Extended Sync” Bit into the Calculation

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: The PCI Express device should adjust the replay timer value if software sets “extended sync” bit to 1’b1. This is to avoid unexpected replay timeouts and recoveries. The “extended sync” CSR bit intended for debug purpose only. PLX Switch supports the device specific register to multiply the “default” replay timer value by up to 256 times.

Solution/Workaround: Multiply the replay timer value by programming Port’s offset FA8h Bits[23:16] with the proper multiplication factor (e.g., Bits[23:16] == 8’h1, multiplies the replay timer by 2, Bits[23:16] == 8’h2, multiplies the replay timer by 4, and so on).

2.4 PEX 87xx Port Times Out in Configuration State in Lanes Breaks after Polling

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: The PEX 87xx port should negotiate the broken receive lanes in configuration state even if they are good in polling state. But the Port times out in configuration state and go to detect state. If the lanes are still broken in polling state, the PEX 87xx port negotiate down the port to the next legal width. The PEX 87xx port negotiate out the lane correctly if they are broken from polling state even if it detects receiver.

Solution/Workaround: None.

Impact: Link bring-up takes extra time for the dynamic broken lane condition.

2.5 PEX 87xx Port Sends TS1 Ordered Set with Incorrect Link Number for the Specified Broken Lane Condition with Dynamic Lane Reversal

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: The external device connected to the PEX 87xx port exchanged the right link number and lane numbers in configuration.lanenum.wait substate. If the external device receiver's lower lanes are broken in this configuration substate and is trying to lane reverse the link, the PEX 87xx port sends incorrect link number and finally timeout in configuration.lanenum.wait substate. The PEX 87xx port goes back to Detect state and come up with right link width and lane reversed configuration.

Solution/Workaround: None.

Impact: Port takes extra time to link up with dynamic broken lane and dynamic lane reversal condition.

2.6 PEX 87xx to PEX 87xx Link Does Not Come Up for Lane Reversal Configuration with Inner Broken Lanes

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: "Inner broken" is the lane number from 1 to n–2 (for example, for x8 link, lane 1 to lane 6 are the inner lanes). If PCIe link is in lane reversal mode with inner broken lane(s), it should negotiate out the broken lanes and link-up in lower width.

But PEX 87xx to PEX 87xx link does not link-up if the lanes are reversed and if one or more inner lanes are broken. The PEX 87xx port does link-up properly with non-PLX devices if lanes are reversed and one or more inner lanes are broken.

Solution/Workaround: None.

Impact: None.

2.7 Gen2 Loopback Entered from Recovery State Always Uses –6-db De-Emphasis

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: When Loopback state is entered from Recovery state and the port is a Loopback slave, the Selectable De-emphasis bit is not saved from the training sets received during Loopback.Entry state.

Solution/Workaround: Enter Loopback state from Configuration state.

Impact: When the link speed is Gen2, and Loopback state is entered from Recovery state, the PEX 87xx port will only transmit at –6 dB.

2.8 PEX 87xx PRBS Function Does Not Work

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: The PEX87xx PRBS checker expects the preamble detection followed by PRBS pattern check. The PEX87xx PRBS checker never detects the preamble pattern due to a design bug.

Solution/Workaround: None.

Impact: The PEX87xx PRBS function does not work.

2.9 PEX 87xx ASPM L1 Exit Logs Receive Error, Bad TLP, and Bad DLLP Errors at Gen3 Data Rate

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: PEX87xx ASPM/PCI-PM L1 exit logs Receiver Error, Bad TLP, and Bad DLLP errors at Gen3 data rate.

Solution/Workaround: Program the following registers using EEPROM, I2C, or software to fix the bit errors. These settings reduce the overall ASPM/PCI-PM L1 power savings by 5%:

- Port0/Port8/Port16 offset 0xBFC with 0x008E2120 – L0/L1/L2/L3
- Port0/Port8/Port16 offset 0xBFC with 0x00AE2120 – L4/L5/L6/L7
- Port0/Port8/Port16 offset 0xBFC with 0x00CE2120 – L8/L9/L10/L11
- Port0/Port8/Port16 offset 0xBFC with 0x00EE2120 – L12/L13/L14/L15
- Port0/Port8/Port16 offset 0xBFC with 0x008E7008 – L0/L1/L2/L3
- Port0/Port8/Port16 offset 0xBFC with 0x00AE7008 – L4/L5/L6/L7
- Port0/Port8/Port16 offset 0xBFC with 0x00CE7008 – L8/L9/L10/L11
- Port0/Port8/Port16 offset 0xBFC with 0x00EE7008 – L12/L13/L14/L15.

2.10 PEX 87xx Downstream Port Does Not Change the Link Speed from Gen2 to Gen3 in Single Retrain for the Specified Rate Change Sequence

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: If software changes the PEX87xx downstream link data rate from Gen1 to Gen2 to Gen3 by changing either upstream and downstream component's target_link_speed, Gen1 to Gen2 or Gen3 speed change is initiated immediately without waiting for a Retrain Link (register 78h[5]). Usually, a speed change should require a single Retrain Link. But the PEX87xx downstream port does not initiate speed change from Gen2 to Gen3 with single retrain. It expects software to write Retrain Link bit twice to initiate Gen2 to Gen3 speed change.

Solution/Workaround:

- Keep the downstream component's target_link_speed to its maximum supported data rate. Change only the PEX87xx downstream port's target_link_speed to change the link data rate.
- If software needs to change the downstream component's target_link_speed, it needs to write PEX87xx downstream port Retrain Link bit two times.

Impact: PEX87xx downstream port does not initiate Gen2 to Gen3 speed change with single Retrain Link.

2.11 PEX 87xx NT0/NT1 Link and NT0/NT1 Virtual Endpoints Do Not Respond to the Specified Configuration Space Register Access while the NT0/NT1-Link Port Is in Reset State

Risk Category: Low

Silicon Revisions Affected: BA, CA

Description: If a Host (connected to the PEX87xx upstream ports) or I^2C/SMBus accesses the following NT0/NT1 endpoint configuration space registers during a reset of NT0/NT1-Link ports (by hot-reset condition, or by PEX_NT0/NT1_PERST# assertion and/or if either of the corresponding Control bits in the VSx_PERST# and NTx_PERST# Status Register (Port 0 register 3A8h[25:24]) are asserted (0)), PEX87xx does not respond to the corresponding configuration space access.

NT0/NT1 Link Endpoint Registers: All configuration space registers

NT0/NT1 Virtual Endpoint Registers: Doorbell and Scratch Pad registers (Offsets from C4Ch to C88h)

Solution/Workaround:

- Avoid I^2C or Host accessing the specified NT0/NT1 Endpoints configuration space registers when NT0/NT1-Link endpoints are in reset state. (Reset state is due to hot-reset condition, PEX_NT0/NT1_PERST# pin assertion, or VSx_PERST# and NTx_PERST# Status Register NT_PERST# (Port 0 register 3A8h[25:24]) is programmed to 0.)

  Or
- Disable PEX87xx hot-reset by programming Port0 offset 0xA30 Bit[4] and Avoid I^2C/Host accessing the specified NT0/NT1 Endpoints configuration space registers when NT0/NT1-Link endpoints are in NT0/NT1_PERST# reset state.

Impact: I^2C access timeout or completion timeout for the specified case.

Appendix A: Reference Information

A.1 Affected Products and Silicon Revisions

This document details Errata for the following products and silicon revisions.

Product	Description	Revision	Status
PEX 8712	12-Lane, 3-Port PCIe Gen3 Switch	CA	Production
PEX 8713	12-Lane, 10-Port PCIe Gen3 Switch	CA	Production
PEX 8716	16-Lane, 4-Port PCIe Gen3 Switch	CA	Production
PEX 8717	16-Lane, 10-Port PCIe Gen3 Switch	CA	Production
PEX 8723	24-Lane, 6-Port PCIe Gen3 Switch	CA	Production
PEX 8724	24-Lane, 6-Port PCIe Gen3 Switch	CA	Production
PEX 8725	24-Lane, 10-Port PCIe Gen3 Switch	CA	Production
PEX 8732	32-Lane, 8-Port PCIe Gen3 Switch	CA	Production
PEX 8733	32-Lane, 18-Port PCIe Gen3 Switch	CA	Production
PEX 8747	48-Lane, 5-Port PCIe Gen3 Switch	CA	Production
PEX 8748	48-Lane, 12-Port PCIe Gen3 Switch	CA	Production
PEX 8749	48-Lane, 18-Port PCIe Gen3 Switch	CA	Production

A.2 Device Documentation Version

The following documentation is the baseline functional description of the silicon.

Document	Version	Description	Publication Date
PEX 8712 Data Book	1.2	Data Book	September 2013
PEX 8713 Data Book	1.2	Data Book	September 2013
PEX 8716 Data Book	1.2	Data Book	September 2013
PEX 8717 Data Book	1.2	Data Book	September 2013
PEX 8723 Data Book	1.2	Data Book	September 2013
PEX 8724 Data Book	1.2	Data Book	September 2013
PEX 8725 Data Book	1.2	Data Book	September 2013
PEX 8732 Data Book	1.2	Data Book	September 2013
PEX 8733 Data Book	1.2	Data Book	September 2013
PEX 8747 Data Book	1.2	Data Book	September 2013
PEX 8748 Data Book	1.2	Data Book	September 2013
PEX 8749 Data Book	1.2	Data Book	September 2013

Revision History

Version 2.8, January 15, 2021

- Updated the workaround for erratum 43, Section 1.43, PEX9700 and PEX8700 products Load the Port Arbitration Table without the Host Setting the Port Arbitration Table Load Bit.
- Fixed typographical issues.

Version 2.7, January 20, 2017

- Updated template.
- Added erratum 43, Section 1.43, PEX9700 and PEX8700 products Load the Port Arbitration Table without the Host Setting the Port Arbitration Table Load Bit.

Version 2.6, April 26, 2016

Updated erratum 42, Section 1.42, PEX87xx DMA Could Corrupt Memory Write TLP Payload of an Unaligned DMA Descriptor while More than One DMA Channel Is Actively Transferring TLPs with TLP Addresses Greater than 32-Bit Addressing.

Version 2.5, September 25, 2015

- Added erratum 42, Section 1.42, PEX87xx DMA Could Corrupt Memory Write TLP Payload of an Unaligned DMA Descriptor while More than One DMA Channel Is Actively Transferring TLPs with TLP Addresses Greater than 32-Bit Addressing.
- Added erratum 41, Section 1.41, PEX87xx Does Not Enter ASPM-L1/PCI-PM-L1 State If the Nonposted Payload Credit Available from Remote Device Is Less than 1 MPS.
- Updated template.

Version 2.4, February 2015

- Corrected title for erratum 31, Section 1.31, PEX 87xx DMA Descriptor Size Less than 32 Bytes Could Cause Data Corruption.

Version 2.3, January 2015

- Updated workaround for erratum 34, Section 1.34, The DMA Payload RAM Assignment Among DMA Functions Need to Be Programmed Before Enabling the DMA Operation.

Version 2.2, August 2014

- Updated description erratum 40, Section 1.40, PEX 87xx Upstream Port Lane Reversal with Link Width Down Negotiation Does Not Work for the Specified Conditions.

Version 2.1, August 2014

- Added erratum 40, Section 1.40, PEX 87xx Upstream Port Lane Reversal with Link Width Down Negotiation Does Not Work for the Specified Conditions.

Version 2.0, June 2014

- Added erratum 39, Section 1.39, False Detection of Receiver for Long Cable Lengths.

Version 1.9, February 2014

- Updated workaround in erratum 33, Section 1.33, Serial Hot Plug Could Fail if Implemented on More Than One Port.

Version 1.8, February 2014

- Updated erratum 33, Section 1.33, Serial Hot Plug Could Fail if Implemented on More Than One Port.
- Added caution 10, Section 2.10, PEX 87xx Downstream Port Does Not Change the Link Speed from Gen2 to Gen3 in Single Retrain for the Specified Rate Change Sequence.

Version 1.7, August 2013

- Corrected workaround for erratum 10, Section 1.10, ASPM L0s Will Not Work at Gen1, Gen2, and Gen3 Rates.
- Updated the impact for erratum 34, Section 1.34, The DMA Payload RAM Assignment Among DMA Functions Need to Be Programmed Before Enabling the DMA Operation.
- Added errata 35 to 38:
  - Section 1.35, PEX 87xx Downstream Port Does Not Change the Link Speed from Gen2 to Gen3 in Single Retrain for the Specified Rate Change Sequence
  - Section 1.36, PEX 87xx Port16 to Port21 Do Not Log the Specified Errors
  - Section 1.37, PEX 87xx Implements MC_Base_Address as Defined in the Description
  - Section 1.38, In the Default Mode, PEX 87xx Incorrectly Calculates the Gen3 TS1 Parity in Response to Use_Preset Equalization Request from the Link Partner.

Version 1.6, April 2013

- Updated workaround for erratum 10, Section 1.10, ASPM L0s Will Not Work at Gen1, Gen2, and Gen3 Rates.
- Updated erratum 32, Section 1.32, PEX 87xx Downstream Port Incorrectly Sets Link Status Register's "Link Autonomous Bandwidth Status" Bit for Any Successful Speed Change Event.
- Added erratum 34, Section 1.34, The DMA Payload RAM Assignment Among DMA Functions Need to Be Programmed Before Enabling the DMA Operation.
- Added cautions 8 and 9, Section 2.8, PEX 87xx PRBS Function Does Not Work and Section 2.9, PEX 87xx ASPM L1 Exit Logs Receive Error, Bad TLP, and Bad DLLP Errors at Gen3 Data Rate.

Version 1.5, December 2012

- Updated erratum 32, Section 1.32, PEX 87xx Downstream Port Incorrectly Sets Link Status Register's "Link Autonomous Bandwidth Status" Bit for Any Successful Speed Change Event.
- Added erratum 33, Section 1.33, Serial Hot Plug Could Fail if Implemented on More Than One Port.

Version 1.4. December 2012

- Updated erratum 25,Section 1.25, PEX 87xx Port Does Not Reliably Link-Up if Its Receiver Does Not Break EI within 28 ms in Gen3 Speed.
- Added erratum 32, Section 1.32, PEX 87xx Downstream Port Incorrectly Sets Link Status Register's "Link Autonomous Bandwidth Status" Bit for Any Successful Speed Change Event.

Version 1.3, November 2012

- Updated erratum 25, Section 1.25, PEX 87xx Port Does Not Reliably Link-Up if Its Receiver Does Not Break EI within 28 ms in Gen3 Speed.
- Added erratum 31, Section 1.31, PEX 87xx DMA Descriptor Size Less than 32 Bytes Could Cause Data Corruption.

Version, 1.2, October 2012

- Updated caution 1, Section 2.1, A Separate EEPROM Image Is Recommended to Run the Links at Gen3 When the Channel Length Is Greater than 10 In., and errata 13, 20, and 21, (Section 1.13, Training Set Corruption during Link Recovery Could Cause PEX 87xx Port to Get Stuck in Recovery.speed State or Go to DL_Down State, Section 1.20, PEX 87xx Software Initiated Up-Configuration Does Not Work Reliably with Electrical Idle Inference Mode, and Section 1.21, Gen3 x8 Link Could Cause LCRC Errors (Non-Fatal Recoverable) on Port1, Port9, Port17).
- Added errata 22 to 30:
  - Section 1.22, PEX 87xx Broken Lanes Issue for the Specified Conditions
  - Section 1.23, PEX 87xx Could Log Receiver Errors at Gen3 Data Rate in Asynchronous Systems with Protocol-Aware Retimers that Send Variable Length SKP OS
  - Section 1.24, Memory Read Requests to Expansion ROM Address xxxxxc90h Returns Incorrect Value Instead of Expansion ROM Value
  - Section 1.25, PEX 87xx Port Does Not Reliably Link-Up if Its Receiver Does Not Break EI within 28 ms in Gen3 Speed
  - Section 1.26, STRAP_GEN1_GEN2_P Pin Does Not Enable Gen2 Only Mode with Port Bifurcation Feature Enabled
  - Section 1.27, PEX 87xx NT Link/NT Virtual Endpoint PCI Compatible Vendor-Specific Compatibility Header Is Incorrect
  - Section 1.28, PEX 87xx Incorrectly Logs Receive Error During Link Width Up-Configuration
  - Section 1.29, PEX 87xx Does Not Log DMA Payload Link List RAM ECC Error Status
  - Section 1.30, PEX 87xx DMA Channel Requires Hot-Reset If It Sends TLP to a PEX 87xx Port with Memory/Master Enable Cleared or in DL_Down State
- Added cautions 6 and 7, Section 2.6, PEX 87xx to PEX 87xx Link Does Not Come Up for Lane Reversal Configuration with Inner Broken Lanes and Section 2.7, Gen2 Loopback Entered from Recovery State Always Uses –6-db De-Emphasis.

Version 1.1, July 2012

- Corrected the solution/workaround section in erratum 14, Section 1.14, PEX 8733/8717/8713 Reports More Lanes/Ports.
- Added errata 19 to 21:
  - Section 1.19, PEX 87xx Port Does Not Reject Illegal Coefficients for the Specified Condition
  - Section 1.20, PEX 87xx Software Initiated Up-Configuration Does Not Work Reliably with Electrical Idle Inference Mode
  - Section 1.21, Gen3 x8 Link Could Cause LCRC Errors (Non-Fatal Recoverable) on Port1, Port9, Port17
- Added cautions 4 and 5 (Section 2.4, PEX 87xx Port Times Out in Configuration State in Lanes Breaks after Polling and Section 2.5, PEX 87xx Port Sends TS1 Ordered Set with Incorrect Link Number for the Specified Broken Lane Condition with Dynamic Lane Reversal).

Version 1.0, May 2012

- Added errata 15 to 18:
  - Section 1.15, PEX 87xx Port “Link Status 2” Register’s Equalization Status Bits Have Incorrect Attribute
  - Section 1.16, PEX 87xx Disables Inactive Lanes upon Entry to Configuration Complete State
  - Section 1.17, PEX 87xx Link Goes Down If Lane0 Breaks when the Link Is Up in Lane Reversal Mode
  - Section 1.18, PEX 87xx Port Does Not Properly Acknowledge the Link Partner’s “Use Preset” Request
- Added caution 3, Section 2.3, PEX 87xx Port Replay Timer Equation Does Not Include “Extended Sync” Bit into the Calculation.
- Updated erratum 4, Section 1.4, Hot-Plug HP_PWREN*/HP_PERST# Pin State Is Unreliable During Fundamental Reset.

Version 0.96, March 2012

- Added erratum 14, Section 1.14, PEX 8733/8717/8713 Reports More Lanes/Ports.
- Updated workaround for erratum 12, Section 1.12, When Virtual Switch Mode Is Enabled, VS_PERST# Does Not Initialize the Serial Hot-Plug Register.

Version 0.95, February 2012

- Added erratum 13, Section 1.13, Training Set Corruption during Link Recovery Could Cause PEX 87xx Port to Get Stuck in Recovery.speed State or Go to DL_Down State.
- Added caution 2, Section 2.2, Link to Intel Based Systems Is Limited to 19-dB Insertion Loss Channel.

Version 0.94, January 2012

- Updated caution 1, Section 2.1, A Separate EEPROM Image Is Recommended to Run the Links at Gen3 When the Channel Length Is Greater than 10 In..

Version 0.93, January 2012

- Added caution 1, Section 2.1, A Separate EEPROM Image Is Recommended to Run the Links at Gen3 When the Channel Length Is Greater than 10 In..

Version 0.92, January 2012

- Updated erratum 2 description, Section 1.2, PEX 8749/48/47/33 Ports 16 to 21 Fail to Return the Error Completion in DL_Down State (when the Slot Is Unpopulated).
- Added errata 3 to 12:
  - Section 1.3, Recovery Equalization Phase 1 Timeout for Downstream Ports Is Not 24 ms.
  - Section 1.4, Hot-Plug HP_PWREN*/HP_PERST# Pin State Is Unreliable During Fundamental Reset.
  - Section 1.5, STRAP_I2C_SMBUS_EN_P Pin Does Not Enable/Disable the SMBUS ARP Mode as Expected.
  - Section 1.6, Memory Read TLPs with Relax Ordering (RO) Attribute Gets Completion Timeout under Extreme Completion Credit Starvation Condition.
  - Section 1.7, Ports Do Not Advertise All the Supported Data Rates in Polling/Configuration State if Target Link Speed Is Programmed to Gen2 or Gen1 Rate, They Only Advertise the Programmed Target Link Value.
  - Section 1.8, When 2VCs Are Enabled in the System, VC1 Nonposted/Completion Traffic Could Block VC0 Posted if VC1 Nonposted/Completion Runs Out of Egress Credit.
  - Section 1.9, NT Virtual Endpoint Incorrectly Logs the Receiver Errors in Recovery State.
  - Section 1.10, ASPM L0s Will Not Work at Gen1, Gen2, and Gen3 Rates.
  - Section 1.11, If Software Initiates a Speed Change to Gen3 on the Downstream Port of the Switch after Links Are Up, DLLP and TLP Transmission Could Be Blocked for the Some Conditions.
  - Section 1.12, When Virtual Switch Mode Is Enabled, VS_PERST# Does Not Initialize the Serial Hot-Plug Register.

Version 0.91, August 2011

- Initial publication of errata list.


BROADCOM®