Error Handling for Maximum Interoperability Between PCIe® Devices

Presented by Richard Solomon

On behalf of Terence Rodrigues, IBM

Copyright © 2007, PCI-SIG, All Rights Reserved

System Diagram

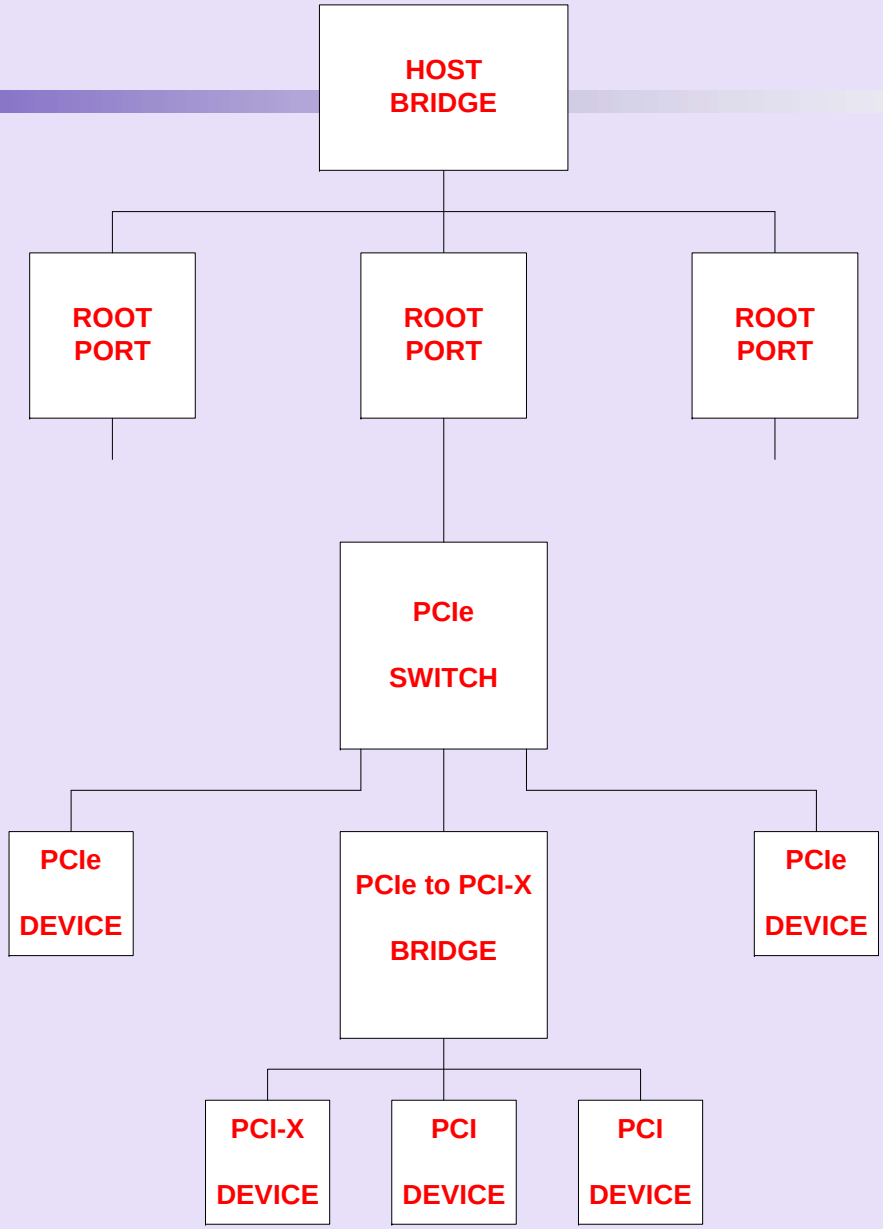

Copyright © 2007, PCI-SIG, All Rights Reserved

Error Classes

- Fatal Errors
  - ✓ Recovery through hardware reset
- Non-Fatal Errors
  - ✓ Recovery through device-specific software
- Correctable Errors
  - ✓ Automatic recovery implemented in hardware
- Advisory Non-Fatal Errors
  - ✓ Recovery through device-specific software
  - ✓ Requestor may promote the error as fatal/non-fatal
  - ✓ Regarded as Correctable Error by the Completor

Copyright © 2007, PCI-SIG, All Rights Reserved

PCI/PCI-X™ Error Handling Controls

- PCI Header Registers
  - ✓ Device Control
  - ✓ Bridge Control
- PCI-X Capability List Registers
  - ✓ PCI-X Command
  - ✓ PCI-X ECC Control & Status
  - ✓ PCI-X Bridge ECC Control & Status

Copyright © 2007, PCI-SIG, All Rights Reserved

PCIe Error Handling Controls

- PCIe Capability List Registers
  - ✓ Device Control
  - ✓ Device Control2
  - ✓ Root Control
- PCIe AER Capability List Registers
  - ✓ Uncorrectable Error Mask
  - ✓ Uncorrectable Error Severity
  - ✓ Correctable Error Mask
  - ✓ Advanced Error Capabilities and Control
  - ✓ Root Error Command
  - ✓ Secondary Uncorrectable Error Mask
  - ✓ Secondary Uncorrectable Error Severity

Copyright © 2007, PCI-SIG, All Rights Reserved

PCI Header
Error Control Registers

Copyright © 2007, PCI-SIG, All Rights Reserved

Device Control (PCI Header)

- Control Bits
  - ✓ SERR# Enable
  - ✓ Parity Error Response Enable
- Issues
  - ✓ PCIe devices should report an error
    - if SERR# is enabled **OR**
    - if reporting is enabled in the PCIe Device Control Register (and not masked off in the AER)
  - ✓ Some PCIe devices will report an error if SERR# is enabled, even though the reporting is masked in the AER
  - ✓ The effect of the masking is lost.

Copyright © 2007, PCI-SIG, All Rights Reserved

Device Control (PCI Header)

- Recommendations

  - ✓ If PCIe device supports AER, don't enable SERR#
    - Error status bits in AER map directly to Error Status Bits in the PCI Header, so no error information is lost.
  - ✓ If PCIe device doesn't support AER and RER, don't enable SERR#
    - Parity Error information can still be obtained from the PCI Header register.
    - SERR maps to Fatal/Non-Fatal errors in PCIe Status registers.
  - ✓ Enable Parity Error Response.

Copyright © 2007, PCI-SIG, All Rights Reserved

Bridge Control (PCI Header)

- Control Bits
  - ✓ SERR# Enable (forward SERR# from secondary to primary side of bridge)
  - ✓ Parity Error Response Enable
  - ✓ Master-Abort Mode
  - ✓ Discard-Timer SERR# Enable

- Issues
  - ✓ Master-Abort Mode:
    - Master abort is an expected error in PC Compatible systems during device scanning.
    - Setting this bit, causes the Master-Abort to also trigger an SERR#.
  - ✓ Discard-Timer SERR# Enable
    - Some PCI devices may not repeat a delayed transaction request in exactly the same form as the original request.
    - The bridge will discard the transaction when the Discard-Timer expires - and assert SERR# if the bit is enabled.
  - ✓ SERR# is treated as a fatal error by many chipsets and will cause a system crash (undesired behavior)

Copyright © 2007, PCI-SIG, All Rights Reserved

Bridge Control (PCI Header)

- Recommendations
  - ✓ Don't report a Master-Abort on one side of the bridge as an SERR# or Target Abort on the other side.
    - PCIe devices receiving an unsupported request
      - Report non-fatal error for write
      - Report advisory non-fatal error for reads
  - ✓ Don't enable SERR# upon discard timer timeout
    - PCI-X & PCIe
      - Completion protocol eliminates the need for the requester to retry the request.
    - PCI
      - Discarded DMA transactions (memory read) have no side-effects
      - Discarded MMIO transactions may have a side effect (system problem)
  - ✓ Enable SERR# and Parity Error Response.

Copyright © 2007, PCI-SIG, All Rights Reserved

PCIe Capability List Registers

Copyright © 2007, PCI-SIG, All Rights Reserved

Device Control (PCIe Capability List)

- Control Bits
  - ✓ Correctable Error Reporting Enable
  - ✓ Non-Fatal Error Reporting Enable
  - ✓ Fatal Error Reporting Enable
  - ✓ Unsupported Request Reporting Enable
- Issues
  - ✓ Non-Posted Unsupported Requests are expected in PC Compatible systems during device scanning.
  - ✓ Setting this bit, causes a non-fatal error message to be sent to the root complex.

Copyright © 2007, PCI-SIG, All Rights Reserved

Device Control (PCIe Capability List)

- Recommendations
  - ✓ If PCIe device supports RER,
    - Non-posted unsupported requests treated as advisory non-fatal errors.
    - Posted unsupported requests treated as non-fatal errors.
  - ✓ If PCIe Device doesn't support RER, don't enable Unsupported Request Reporting, to get PC compatible behavior.
  - ✓ Enable the reporting of other errors (correctable, non-fatal and fatal)

Copyright © 2007, PCI-SIG, All Rights Reserved

PCIe Device Control2 (PCIe Capability List)

- Control Bits
  - ✓ Completion Time-out Disable
- Issues
  - ✓ Default Completion Time-out Range (50us – 50ms) may be too small for very large systems under stress.
  - ✓ Additional Completion Time-out Ranges (A, B, C, D) are optional for PCIe devices
- Recommendations
  - ✓ Disable Completion Timeout in PCIe devices
  - ✓ Let the system detect the Completion Timeout at a level above PCIe (Host Bridge)

Copyright © 2007, PCI-SIG, All Rights Reserved

Root Control (PCIe Capability List)

- Control Bits
  - ✓ System Error on Correctable Error Enable
  - ✓ System Error on Non-Fatal Error Enable
  - ✓ System Error on Fatal Error Enable
- Issues
  - ✓ Root Port behavior upon receipt of the above error messages is chipset – dependent (usually SERR# is asserted)
  - ✓ SERR# is treated as a fatal error by many chipsets and will cause a system crash (undesired behavior)
- Recommendations
  - ✓ Route SERR# to SMI or some other Interrupt.
  - ✓ If not possible, don't enable System Error on Correctable Error & Non-Fatal Error.

Copyright © 2007, PCI-SIG, All Rights Reserved

Uncorrectable Error Mask (PCIe AER Capability List)

- Control Bits
  - ✓ Training Error Mask
  - ✓ Unsupported Request Error Mask
  - ✓ Other Non-Fatal & Fatal Error Masks
- Issues
  - ✓ Training Error was not well-defined, and could result in a fatal error. It was subsequently undefined by an PCI-SIG ECN.
  - ✓ Non-Posted Unsupported Requests are expected in PC Compatible systems during device scanning.
  - ✓ Some PCIe devices always report errors to the Root Port unless the error is also masked in the AER.
  - ✓ So, even if Unsupported Request Reporting is disabled in PCIe Device Control reg, it still needs to be masked here.

Copyright © 2007, PCI-SIG, All Rights Reserved

Uncorrectable Error Mask (PCIe AER Capability List)

- Recommendations
  - ✓ Always set the Training Error Mask to avoid false triggering in the devices which may already have implemented this feature.
  - ✓ If PCIe device supports RER,
    - – Non-posted unsupported requests treated as advisory non-fatal errors.
    - – Posted unsupported requests treated as non-fatal errors.
  - ✓ If PCIe Device doesn't support RER,
    - – don't enable Unsupported Request Reporting, to get PC compatible behaviour.
  - ✓ Unmask the other non-fatal and fatal errors

Copyright © 2007, PCI-SIG, All Rights Reserved

Uncorrectable Error Severity (PCIe AER Capability List)

- Control Bits
  - ✓ Non-Fatal vs Fatal Severity
- Issues
  - ✓ Unsupported Requests ought to be reported as Non-Fatal or Advisory Non-Fatal Errors (Correctable)
  - ✓ However, many PCIe Root Ports don't provide the ability to route Non-Fatal Error separately from Fatal Error.
  - ✓ Unsupported Request conditions can thus inadvertently trigger a Fatal Error.
- Recommendations
  - ✓ If PCIe Device doesn't support RER, promote all non-fatal errors to fatal errors.
  - ✓ If RER is supported,
    - leave advisory non-fatal errors at non-fatal severity (will be treated as correctable errors)
    - promote remaining non-fatal errors to fatal.

Copyright © 2007, PCI-SIG, All Rights Reserved

Correctable Error Mask (PCIe AER Capability List)

- Control Bits
  - ✓ Advisory Non-Fatal Error Mask
  - ✓ Other Correctable Error Masks
- Issues
  - ✓ Advisory Non-Fatal Errors point to functional (s/w) problems rather than to problems with the Link.
  - ✓ Correctable Errors are used for link PFA.
- Recommendations
  - ✓ Mask Advisory Non-Fatal Errors
  - ✓ Unmask Other Correctable Errors

Copyright © 2007, PCI-SIG, All Rights Reserved

Advanced Error Capabilities & Control (PCIe AER Capability List)

- Control Bits
  - ✓ ECRC Generation Enable
  - ✓ ECRC Check Enable
- Issues
  - ✓ Many PCIe Adapters are unable to tolerate ECRC
    - – ECRC is optional, but ignoring it is not.
- Recommendations
  - ✓ Don’t enable ECRC Generation, unless all PCIe Devices below a Root Port support ECRC Generation & Checking.
  - ✓ The setting does not need to be the same across root ports
  - ✓ ECRC Check Enable has no effect, unless ECRC is being generated.

Copyright © 2007, PCI-SIG, All Rights Reserved

Root Error Command (PCIe AER Capability List)

- Control Bits
  - ✓ Correctable Error Reporting Enable
  - ✓ Non-Fatal Error Reporting Enable
  - ✓ Fatal Error Reporting Enable
- Issues
  - ✓ Root Port behaviour upon receipt of the above error messages is chipset – dependent (usually SERR# is asserted)
  - ✓ SERR# is treated as a fatal error by many chipsets and will cause a system crash (undesired behaviour)
- Recommendations
  - ✓ If SERR# can be routed as desired, use the PCIe Root Command Register.
  - ✓ If not,
    - Use MSI to route the errors to SMI or to some other Interrupt.
    - Don't grant control of the PCIe Capability List to the OS.

Copyright © 2007, PCI-SIG, All Rights Reserved

Secondary Uncorrectable Error Mask (PCIe AER Capability List)

- Control Bits
  - ✓ Received Master-Abort Mask
  - ✓ Delayed Transaction Discard Timer Expired Mask
  - ✓ Other Non-Fatal & Fatal Error Masks
- Issues
  - ✓ Master-Abort Mode:
    - Master abort is an expected error in PC Compatible systems during device scanning.
    - Setting this bit, causes the Master-Abort to also trigger an SERR#.
  - ✓ Discard-Timer SERR# Enable
    - Some PCI devices may not repeat a delayed transaction request in exactly the same form as the original request.
    - The bridge will discard the transaction when the Discard-Timer expires - and assert SERR# if the bit is enabled.
  - ✓ SERR# is treated as a fatal error by many chipsets and will cause a system crash (undesired behavior)

Copyright © 2007, PCI-SIG, All Rights Reserved

Secondary Uncorrectable Error Mask (PCIe AER Capability List)

- Recommendations
  - ✓ Don't report a Master-Abort on one side of the bridge as an SERR# or Target Abort on the other side.
    - PCIe devices receiving an unsupported request
      - Report non-fatal error for write
      - Report advisory non-fatal error for reads
  - ✓ Don't enable SERR# upon discard timer timeout
    - PCI-X & PCIe
      - Completion protocol eliminates the need for the requester to retry the request.
    - PCI
      - Discarded DMA transactions (memory read) have no side-effects
      - Discarded MMIO transactions may have a side effect (system problem)
  - ✓ Enable Other Non-Fatal & Fatal Errors.

Copyright © 2007, PCI-SIG, All Rights Reserved

Secondary Uncorrectable Error Severity (PCIe AER Capability List)

- Control Bits
  - ✓Non-Fatal vs Fatal Severity
- Issues
  - ✓However, many PCIe Root Ports don't provide the ability to route Non-Fatal Error separately from Fatal Error.
- Recommendations
  - ✓If PCIe Device doesn't support RER, promote all non-fatal errors to fatal errors.
  - ✓If RER is supported,
    - leave advisory non-fatal errors at non-fatal severity (will be treated as correctable errors)
    - promote remaining non-fatal errors to fatal.

Copyright © 2007, PCI-SIG, All Rights Reserved

PCI-X Capability List Registers

Copyright © 2007, PCI-SIG, All Rights Reserved

PCI-X Command (PCI-X Capability List)

- Control Bits
  - ✓ Uncorrectable Data Error Recovery
- Issues
  - ✓ Device specific knowledge is required for recovery
- Recommendations
  - ✓ Enable recovery from uncorrectable data errors only if the system software knows how to recover and clean up.

Copyright © 2007, PCI-SIG, All Rights Reserved

PCI-X ECC Control & Status (PCI-X Capability List)

- Control Bits
  - ✓ Disable Singe-Bit-Error Correction
  - ✓ ECC Mode
- Issues
  - ✓ Not all devices on the bus may support ECC
- Recommendations
  - ✓ Enable ECC if all devices on the Bus support ECC
  - ✓ ECC Mode is configured by PCIXCAP initialization pattern
  - ✓ Single-Bit-Error Correction is only applicable in ECC mode.

Copyright © 2007, PCI-SIG, All Rights Reserved

PCI-X Bridge ECC Control & Status (PCI-X Capability List)

- Control Bits
  - ✓ Disable Singe-Bit-Error Correction
  - ✓ ECC Mode
- Issues
  - ✓ Not all devices on the bus may support ECC
- Recommendations
  - ✓ Enable ECC if all devices on the Bus (primary or secondary) support ECC
  - ✓ ECC Mode is configured by PCIXCAP initialization pattern
  - ✓ Single-Bit-Error Correction is only applicable in ECC mode

Copyright © 2007, PCI-SIG, All Rights Reserved

Thank you for attending the PCI-SIG Developers Conference Europe 2007.

For more information please go to www.pcisig.com

Copyright © 2007, PCI-SIG, All Rights Reserved