
PCI
SIG™

Implementing PCI Express on a System-On-Chip (SOC)

Paul J Mattos

PCI Express IP Development Lead

IBM Systems & Technology Group

Copyright © 2004, PCI-SIG, All Rights Reserved

Agenda

- Introduction
- Example Product Description
- IP Selection Criteria
- Feature Selection Considerations
- Chip Integration Considerations
- Summary

Copyright © 2004, PCI-SIG, All Rights Reserved

Introduction

- Discuss PCI Express implementation considerations for Example SOC application
- Highlight critical trade-offs
  - ✓ Features supported
  - ✓ Die area
  - ✓ Performance
- Help sort through
  - ✓ Optional features
  - ✓ Implied requirements
  - ✓ Implementation specific choices

Copyright © 2004, PCI-SIG, All Rights Reserved

Agenda

- Introduction
- Example Product Description
- IP Selection Criteria
- Feature Selection Considerations
- Chip Integration Considerations
- Summary

Copyright © 2004, PCI-SIG, All Rights Reserved

SOC – Dual Endpoint/Root Complex

Product requires PCIE connectivity as an Endpoint in first application...

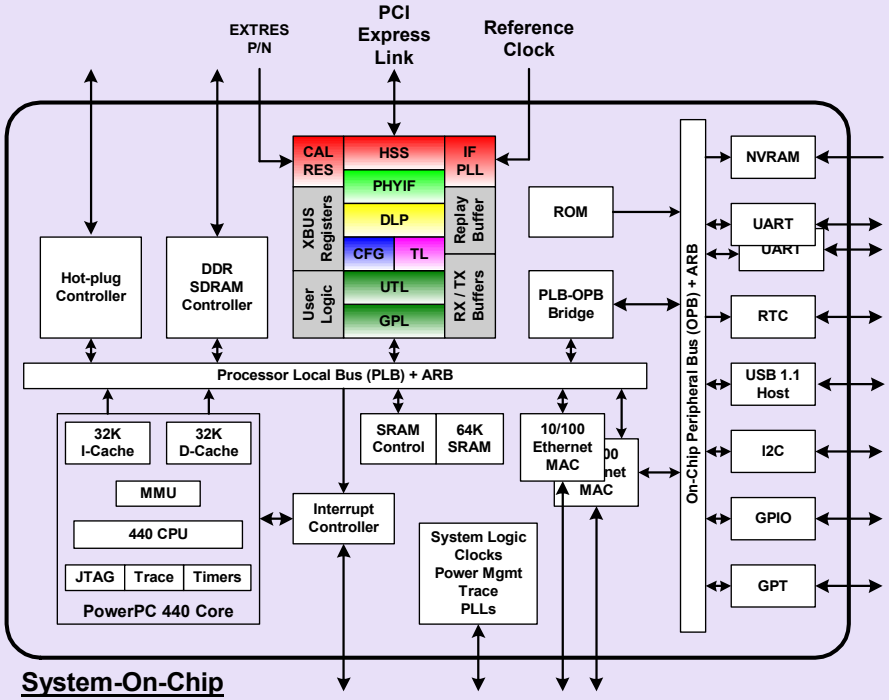

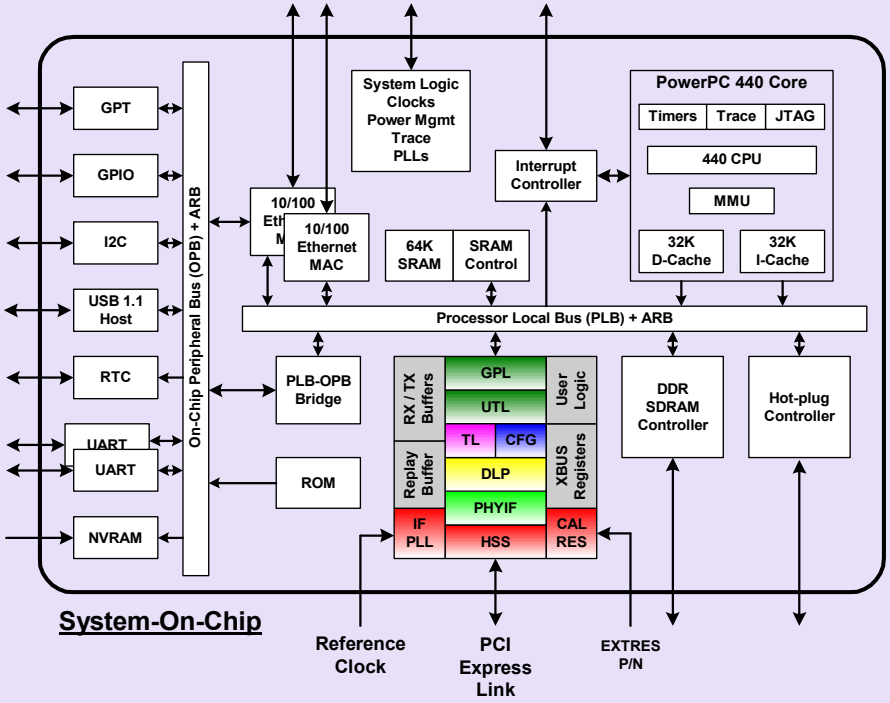

...and Root Complex in second application

Copyright © 2004, PCI-SIG, All Rights Reserved

Agenda

- Introduction
- Example Product Description
- IP Selection Criteria
- Feature Selection Considerations
- Chip Integration Considerations
- Summary

Copyright © 2004, PCI-SIG, All Rights Reserved

Typical PCIe SOC Building Blocks

- PCI Express Port
- Buffers (Replay, RX, TX)
- System Bus Bridge
- Programming Registers
- Controllers
  - ✓Hot-plug
  - ✓Interrupt
  - ✓Power Management
- Diagnostic Support

Copyright © 2004, PCI-SIG, All Rights Reserved

Logical PCI Express Hierarchy

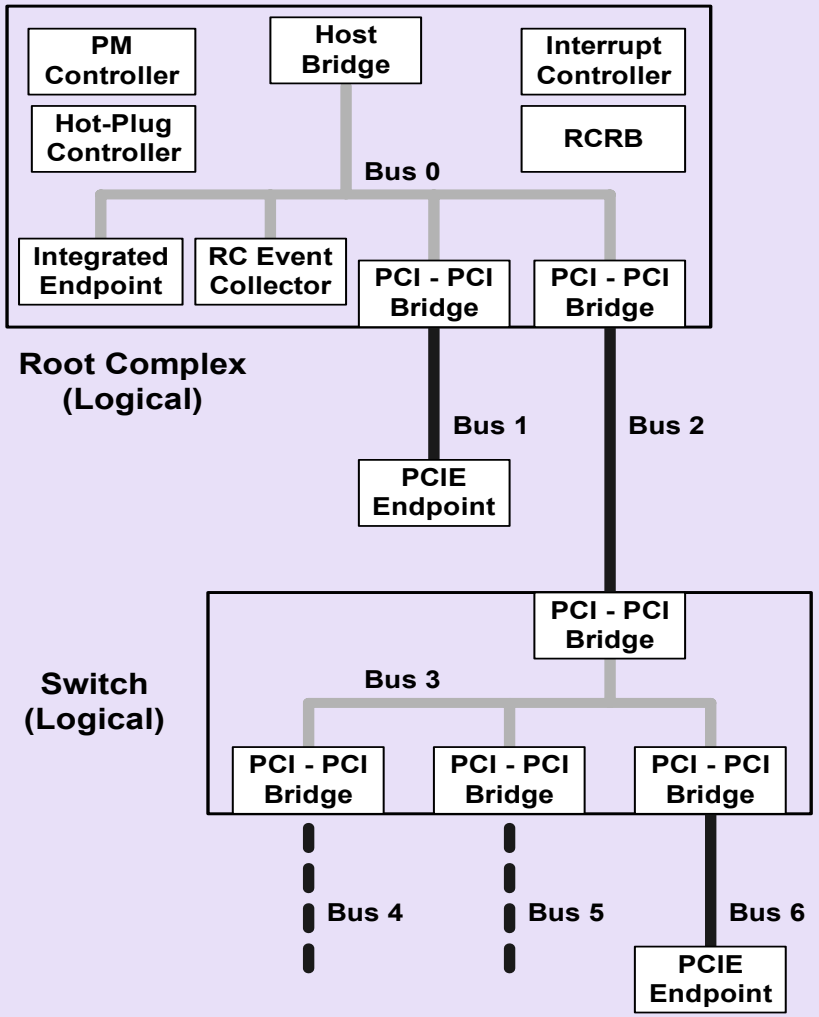

- The Base Specification refers to device types such as Root Complex, Switch, and Endpoint, and defines the architecture and protocols to interoperate in a PCIe system
- From a hardware implementation perspective, the majority of the Base Specification is represented by the port or port(s) on the device that incorporate the functionality of the protocol layers as defined in the specification
- Devices often have functionality in addition to and outside the scope of the specification

Copyright © 2004, PCI-SIG, All Rights Reserved

PCIe Stack/Device Relationship

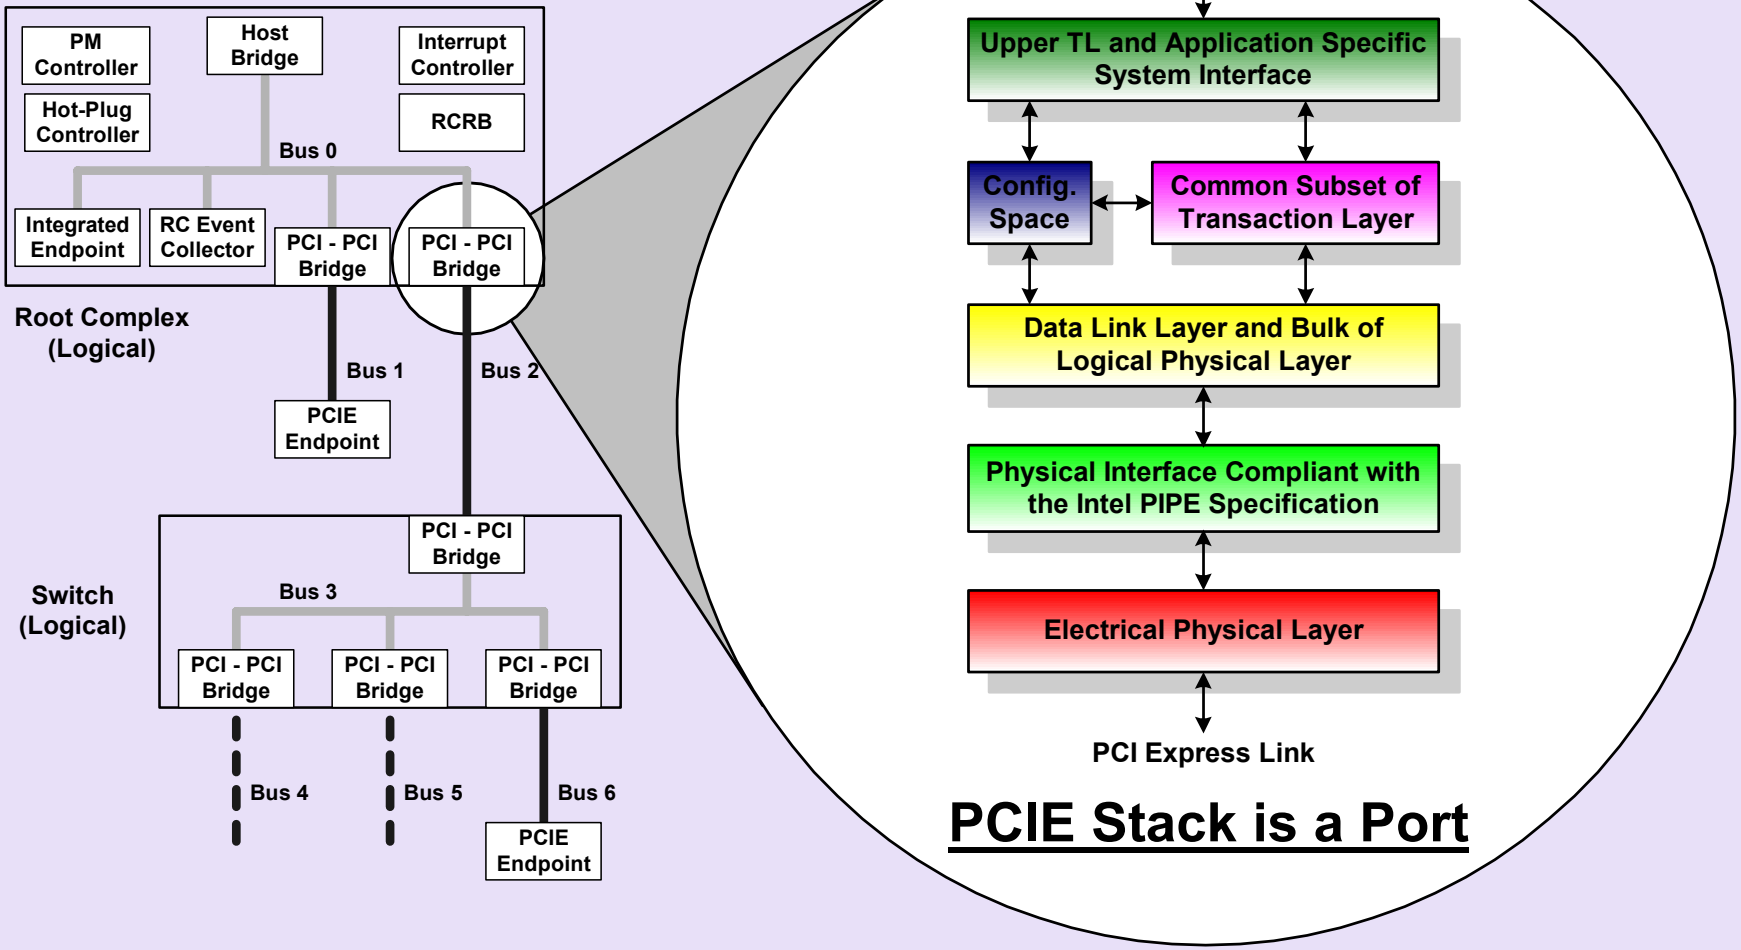

Copyright © 2004, PCI-SIG, All Rights Reserved

Profile of General Use PCIe IP

- Make or Buy?
  - ✓ What to look for in general use PCIe IP
- Provides multiple levels of flexibility in terms of features, area, and performance
  - ✓ Provides functional partitioning along natural protocol layer boundaries
    - Well-defined interfaces (e.g., PIPE)
    - Supports full and partial hardware stacks
      - For customers to incorporate proprietary functions/features
  - ✓ Supports extensive pre-synthesis configurability and pin programmability options
    - Minimizes or eliminates unused logic
    - Provides flexibility for multiple application support
  - ✓ Provides 'hooks' for expansion and support of optional features
    - Hot-plug, interrupts, error handling, power management

Copyright © 2004, PCI-SIG, All Rights Reserved

Major Configurable/ Programmable Items

- Configurable (pre-synthesis)
  - ✓ Maximum link width (performance & area)
  - ✓ Buffer sizes (performance & area)
  - ✓ Number of virtual channels (performance & area)
  - ✓ ECRC (area)
  - ✓ Configuration space capability structures (area)
- Programmable (flexibility)
  - ✓ Maximum link width
  - ✓ Buffer sizes
  - ✓ Number of virtual channels
  - ✓ ECRC
  - ✓ Configuration space capability registers
  - ✓ Device/port type support

Copyright © 2004, PCI-SIG, All Rights Reserved

Example of General Use PCIe IP

- Upper TL
  - ✓ Device specific TL functionality
- Configuration Space
- Lower Transaction Layer
  - ✓ Common TL functions for all device types
- Data Link & Logical Phys.
  - ✓ PIPE, LTSSM, DLCMSM
- Physical Interface
  - ✓ PIPE, 8b/10b, e-FIFO
- Electrical Physical Layer
  - ✓ SerDes, PLL

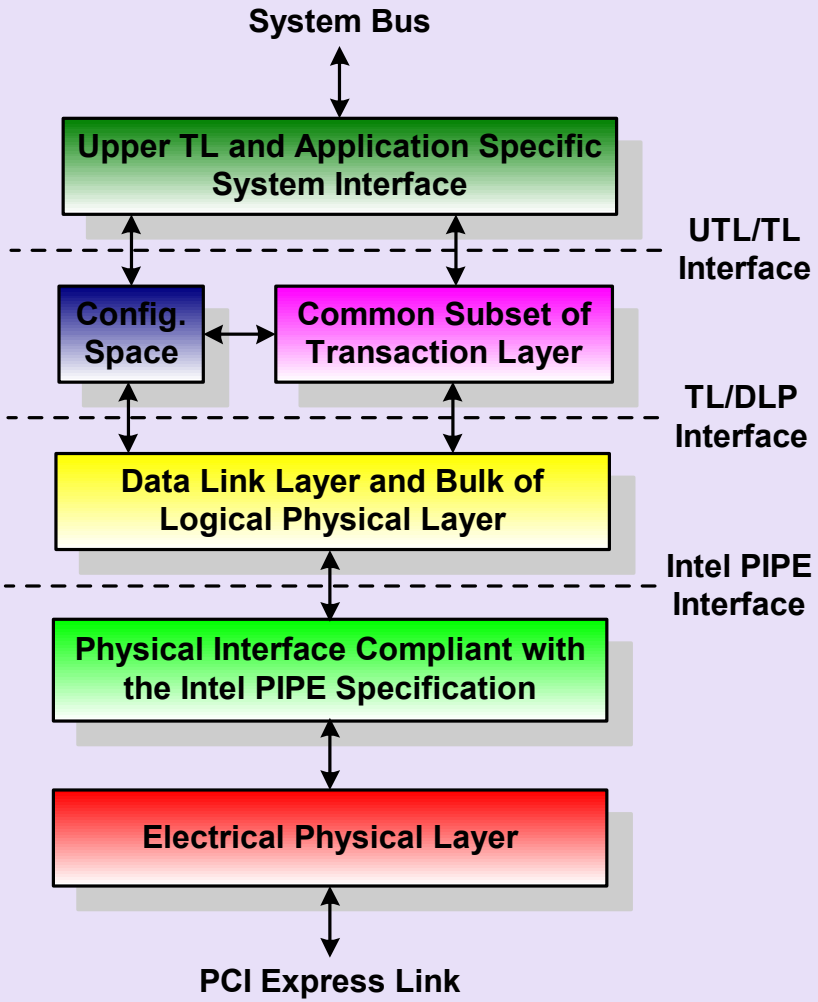

Copyright © 2004, PCI-SIG, All Rights Reserved

Agenda

- Introduction
- Example Product Description
- IP Selection Criteria
- Feature Selection Considerations
- Chip Integration Considerations
- Summary

Copyright © 2004, PCI-SIG, All Rights Reserved

Features with Significant Impact on Die Area

- Maximum link width
  - ✓ Anything which needs to be handled on a per lane basis is multiplied by number of lanes
    - – SerDes, Elastic FIFOs, 8b/10b encode/decode
  - ✓ LCRC generation/checking logic generally larger for wider links
  - ✓ LTSSM
- Buffer sizes (dependencies on Maximum Payload Size and number of requests and completions supported)
  - ✓ Replay, RX and TX
- Number of Virtual Channels
  - ✓ Buffers, buffer management, arbitration logic, flow control logic
- ECRC
  - ✓ Generation and checking logic generally larger for wider links
- AER
  - ✓ Many additional registers - for Root Ports there are even more registers and considerations for system notification

Copyright © 2004, PCI-SIG, All Rights Reserved

Area Optimization Options

- Consider port bifurcation when multiple ports are required
  - ✓ Enables reuse of SerDes lanes for more than one link
  - ✓ Must consider optional support for intermediate link widths
  - ✓ Also implies use of multiple protocol stacks
    - Good trade-off since protocol logic area < (Protocol + PHY) area
- Optimize replay buffer size if TX payloads < RX payloads
  - ✓ Since replay buffer is associated with TX, the local device determines maximum payload of outgoing packets
  - ✓ Need to keep in mind how the Maximum Payload Size parameter (Device Control Register) is used.
- ECRC
  - ✓ Typically would use both generation and checking, but if the application permits, use only one or the other.

Copyright © 2004, PCI-SIG, All Rights Reserved

Port Bifurcation Concept

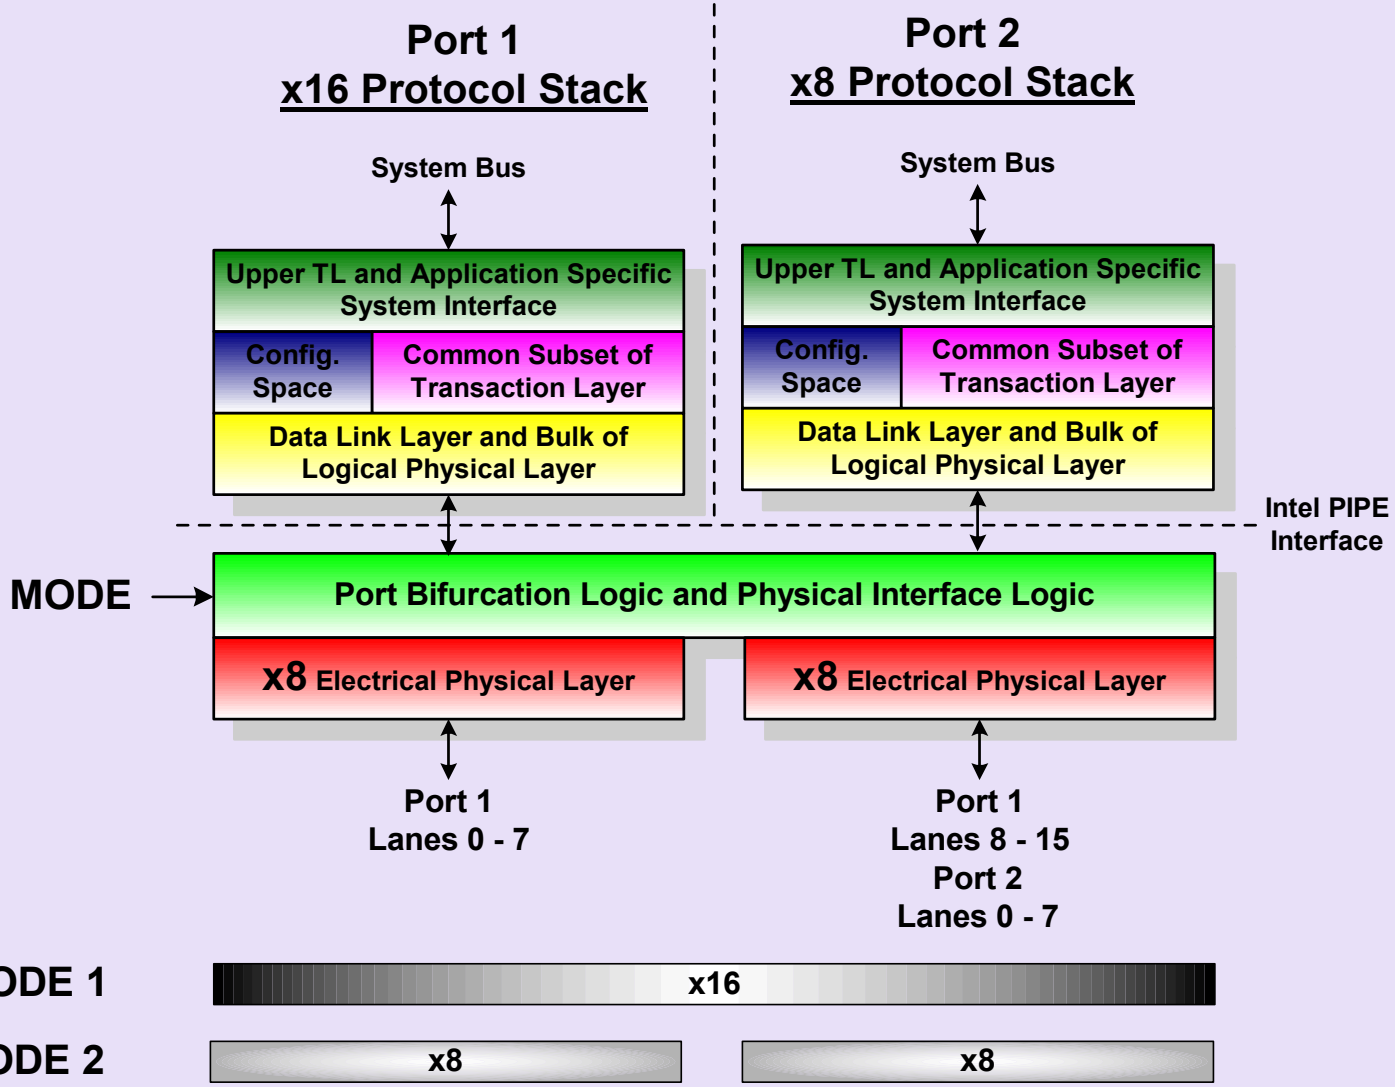

Copyright © 2004, PCI-SIG, All Rights Reserved

MPS in Device Control Register

- Maximum Payload Size parameter is used for:
  - ✓ Receive error checking
    - – Malformed TLP
    - – FCPE
  - ✓ Replay timer timeout
  - ✓ As Endpoint, for L1 and ASPM L1 minimum Flow Control credit accumulation
- In an Endpoint, if TX payloads are < RX payloads, then special system consideration may have to be given to PM Flow Control credit accumulation
  - ✓ Root Complex (supporting peer-to-peer), initial CPLD credit advertisement could be equivalent to largest Read Request Size issued by Root Complex and this could be smaller than MPS.

Copyright © 2004, PCI-SIG, All Rights Reserved

Agenda

- Introduction
- Example Product Description
- IP Selection Criteria
- Feature Selection Considerations
- Chip Integration Considerations
- Summary

Copyright © 2004, PCI-SIG, All Rights Reserved

Dual EP/RC Considerations

- Hot Reset
  - ✓ Initiated from Root Complex/received at Endpoint
- Configuration Requests
  - ✓ Only from Root Complex
  - ✓ Support for Configuration Request Retry as Endpoint
- Messages
  - ✓ Message flow reversed
- Interrupts
  - ✓ Generation (EP & RC) versus Termination (RC only)
- Power Management
  - ✓ Role reversals for L1 and ASPM L1 entry
- Hot Plug
  - ✓ Controller versus Adapter

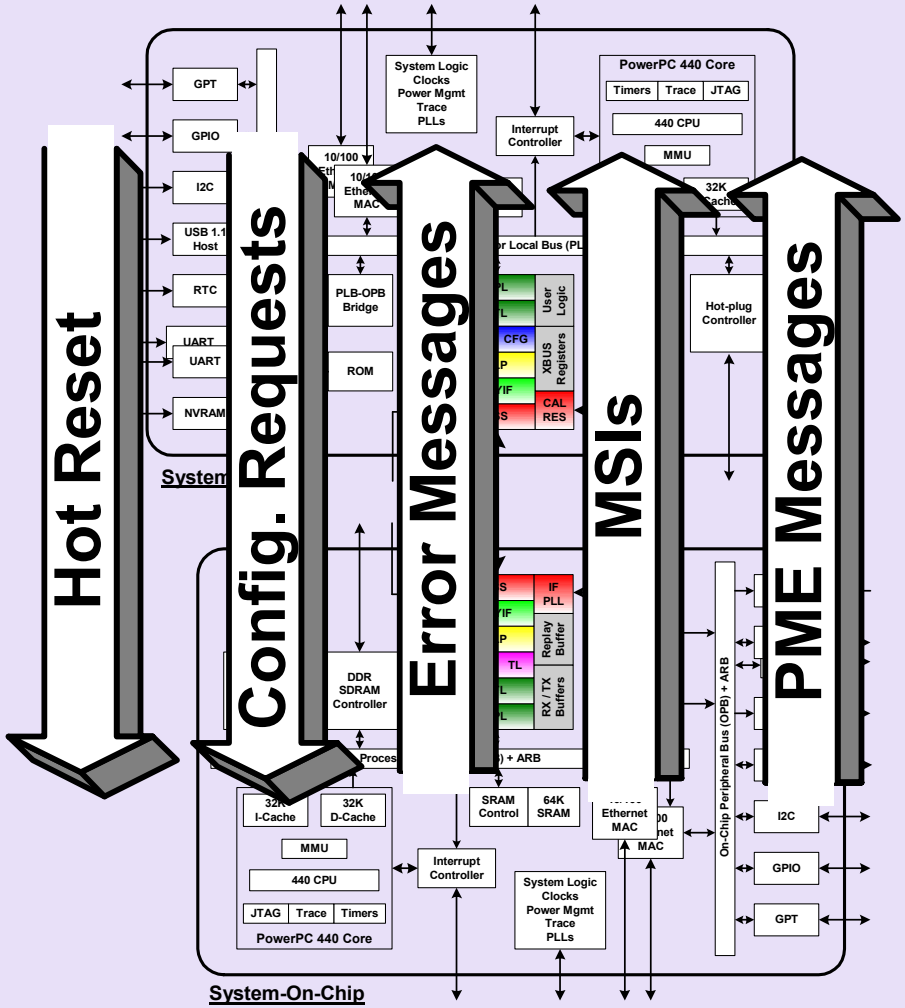

Copyright © 2004, PCI-SIG, All Rights Reserved

Reset Considerations

- Component must enter LTSSM Detect within 20ms of the end of Fundamental Reset
  - ✓ SOC initialization including internal PLL locks and PCIe stack initialization must complete in time
- Booting embedded processor from PCI Express
  - ✓ Need to ensure PCIe stack initialization does not require firmware support
- Minimum Hot Reset time is < 5ms
  - ✓ Endpoint must manage how it will respond to Hot Reset and complete internal ‘housekeeping’ in time
  - ✓ Physical layer must remain active to detect Hot Reset de-assertion

Copyright © 2004, PCI-SIG, All Rights Reserved

Root Port Error Consolidation

- Reporting PCIe System Errors due to either error messages (ERR_COR, ERR_NONFATAL, and ERR_FATAL) received by the Root Port or Root Port detected errors is controlled by the Root Control Register (PCIe System Error)
- Root Port generates PCIe AER related interrupts if enabled through the Root Error Command Register (AER related MSI)
- Root Error Status Register is the consolidation of errors from received error messages from the link AND errors detected within the Root Port itself

Copyright © 2004, PCI-SIG, All Rights Reserved

Root Port Receives Uncorrectable Error Messages

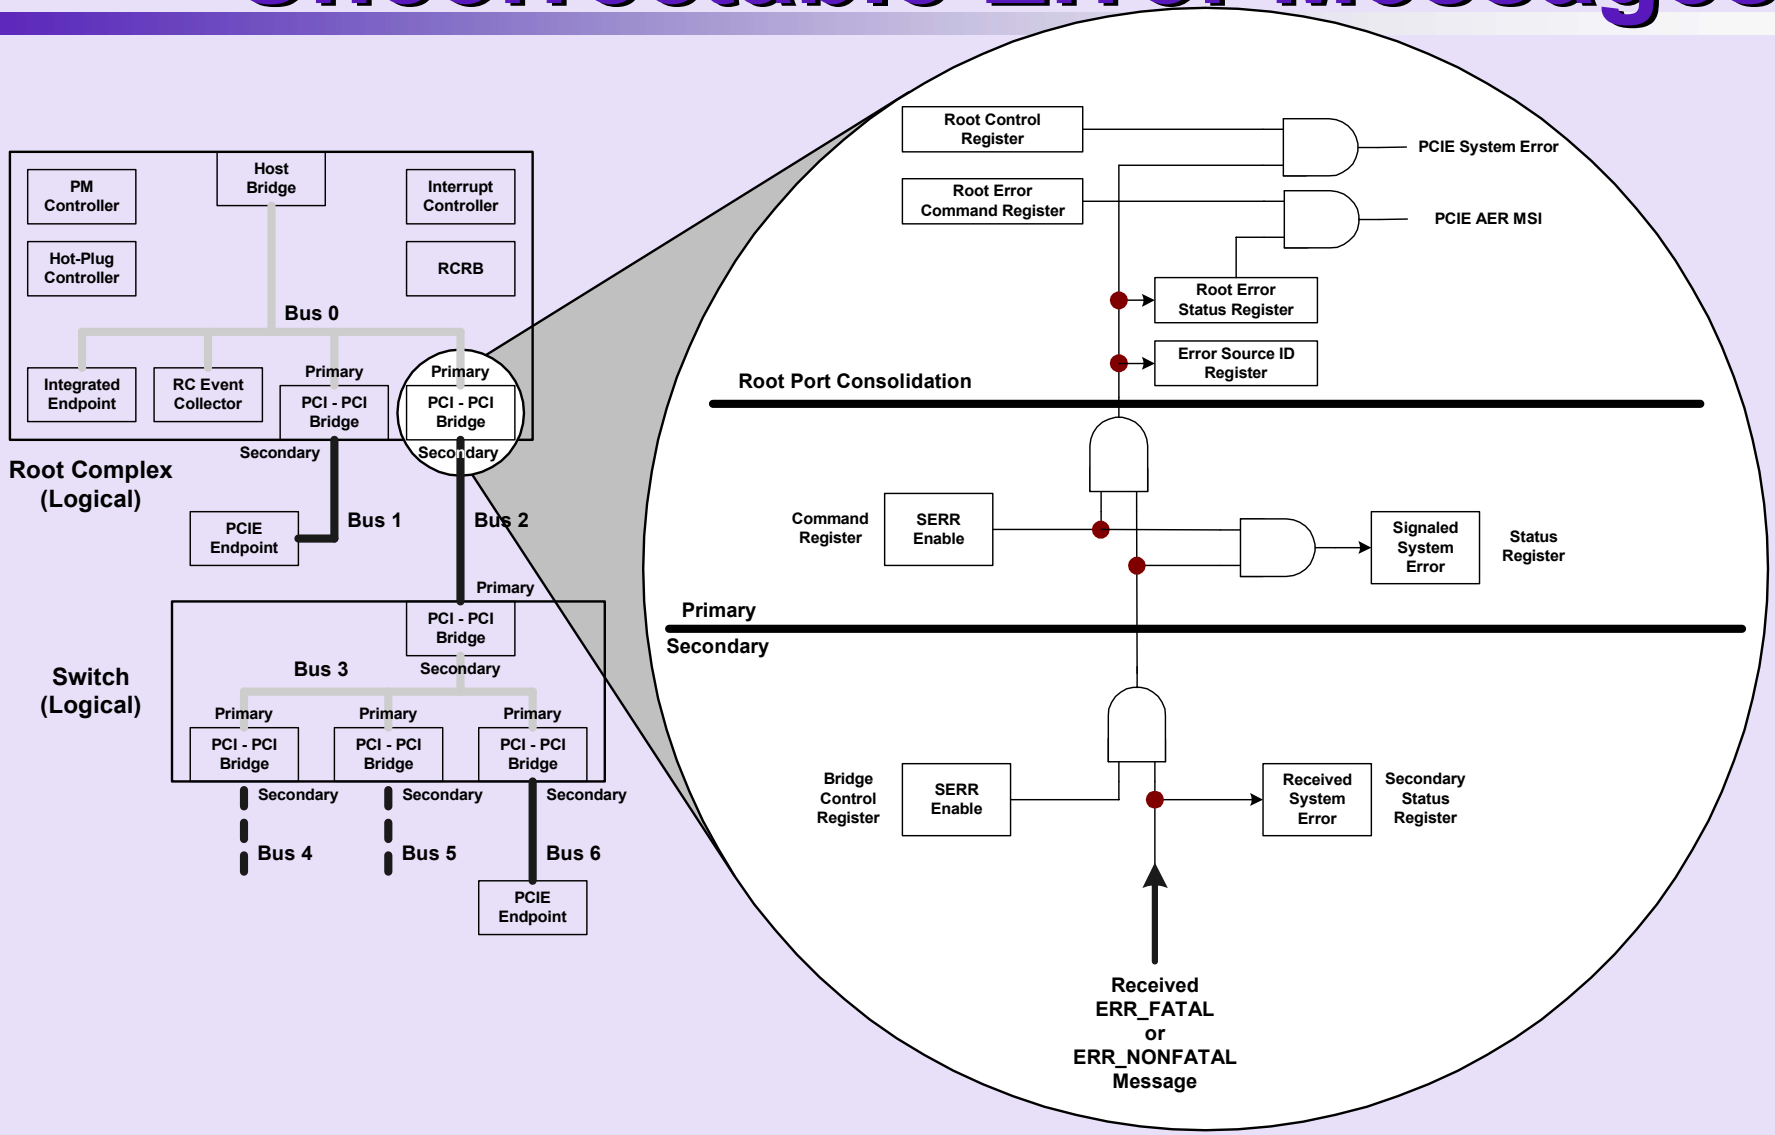

Copyright © 2004, PCI-SIG, All Rights Reserved

Root Port Receives Correctable Error Messages

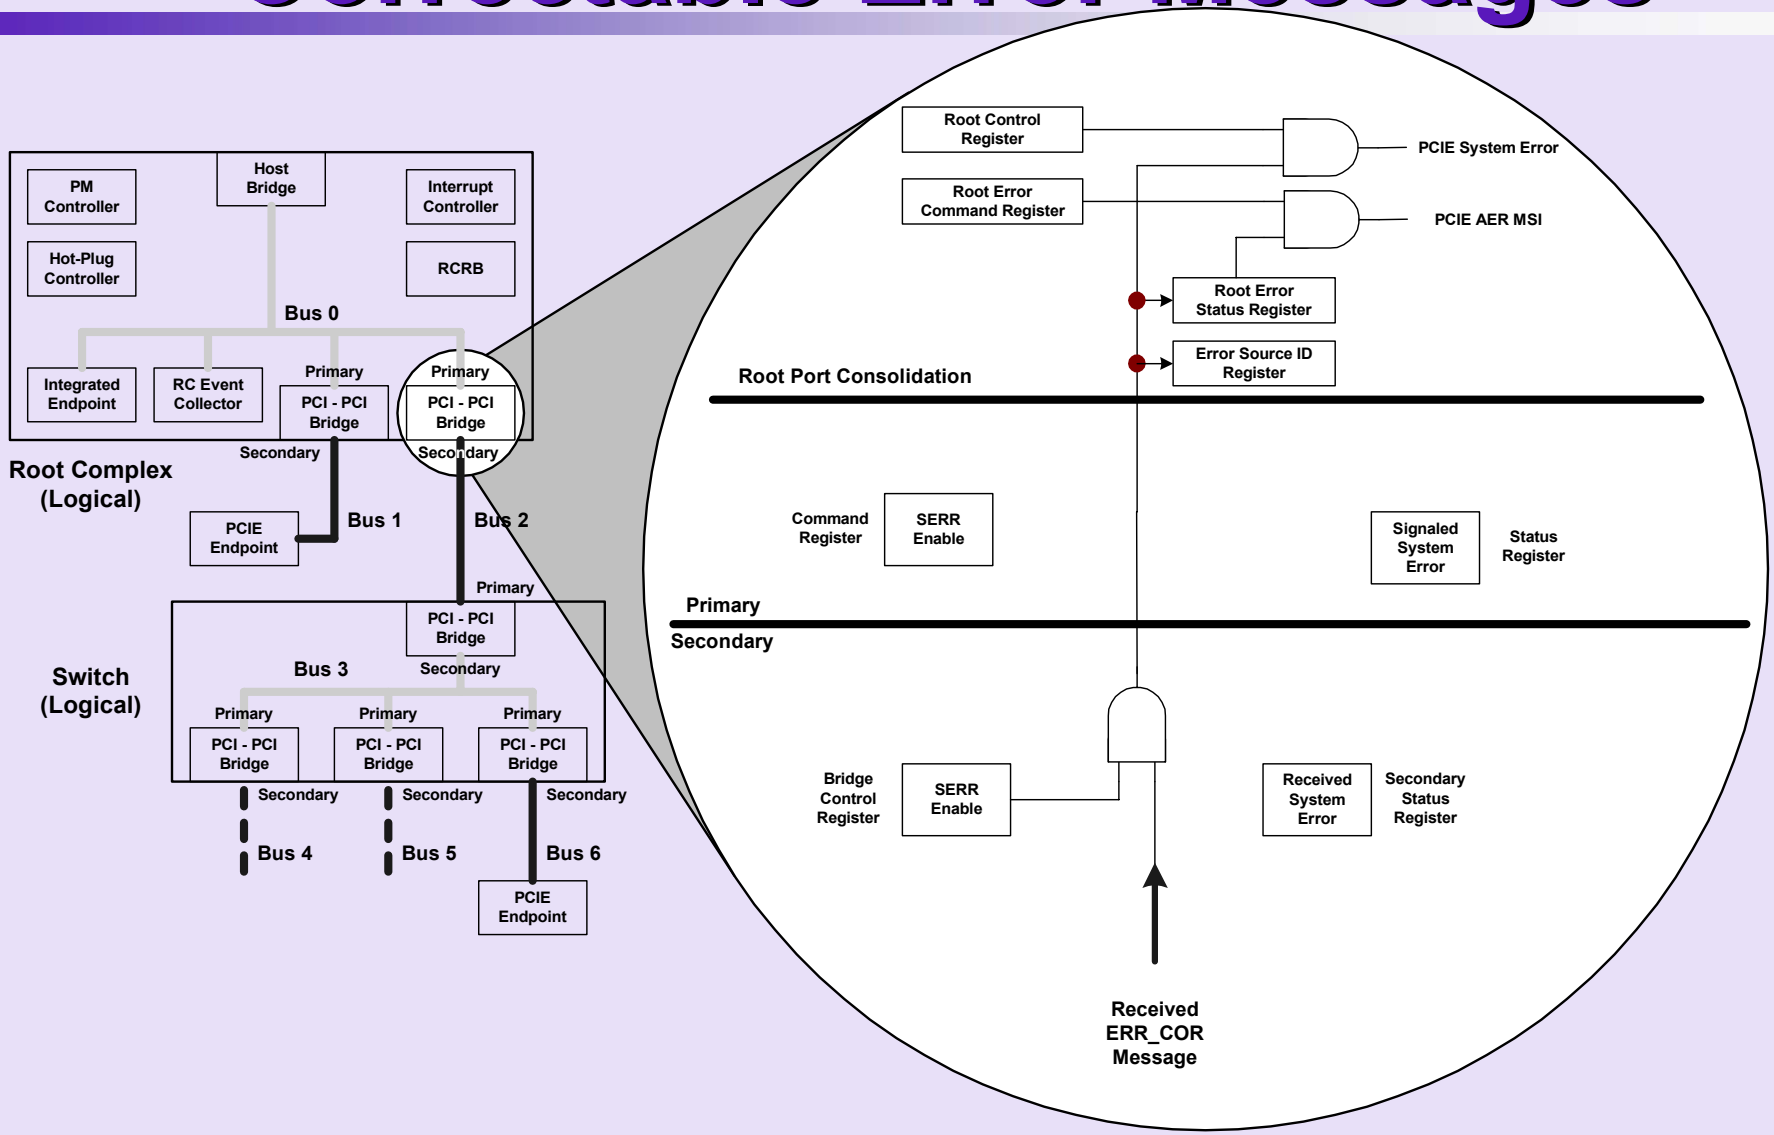

Copyright © 2004, PCI-SIG, All Rights Reserved

Root Port Detects Uncorrectable Error

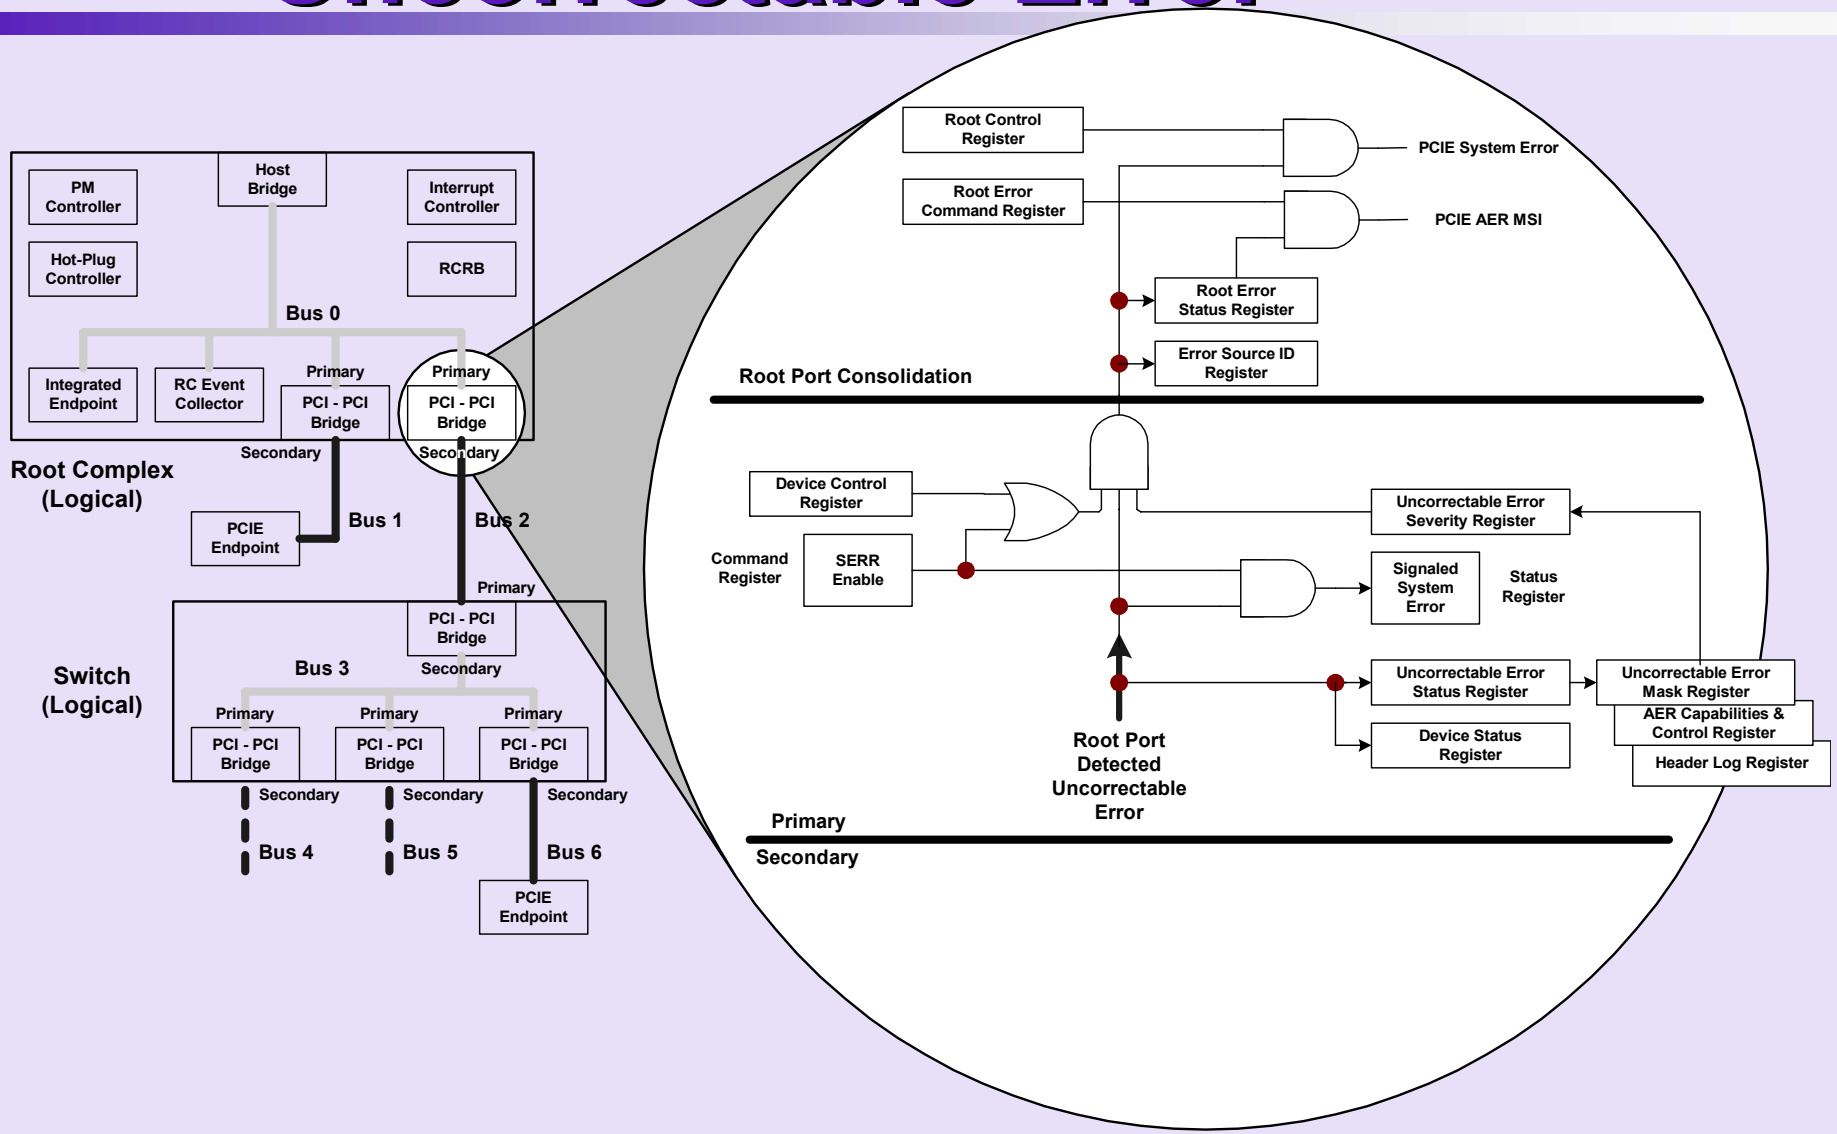

Copyright © 2004, PCI-SIG, All Rights Reserved

Root Port Detects Correctable Error

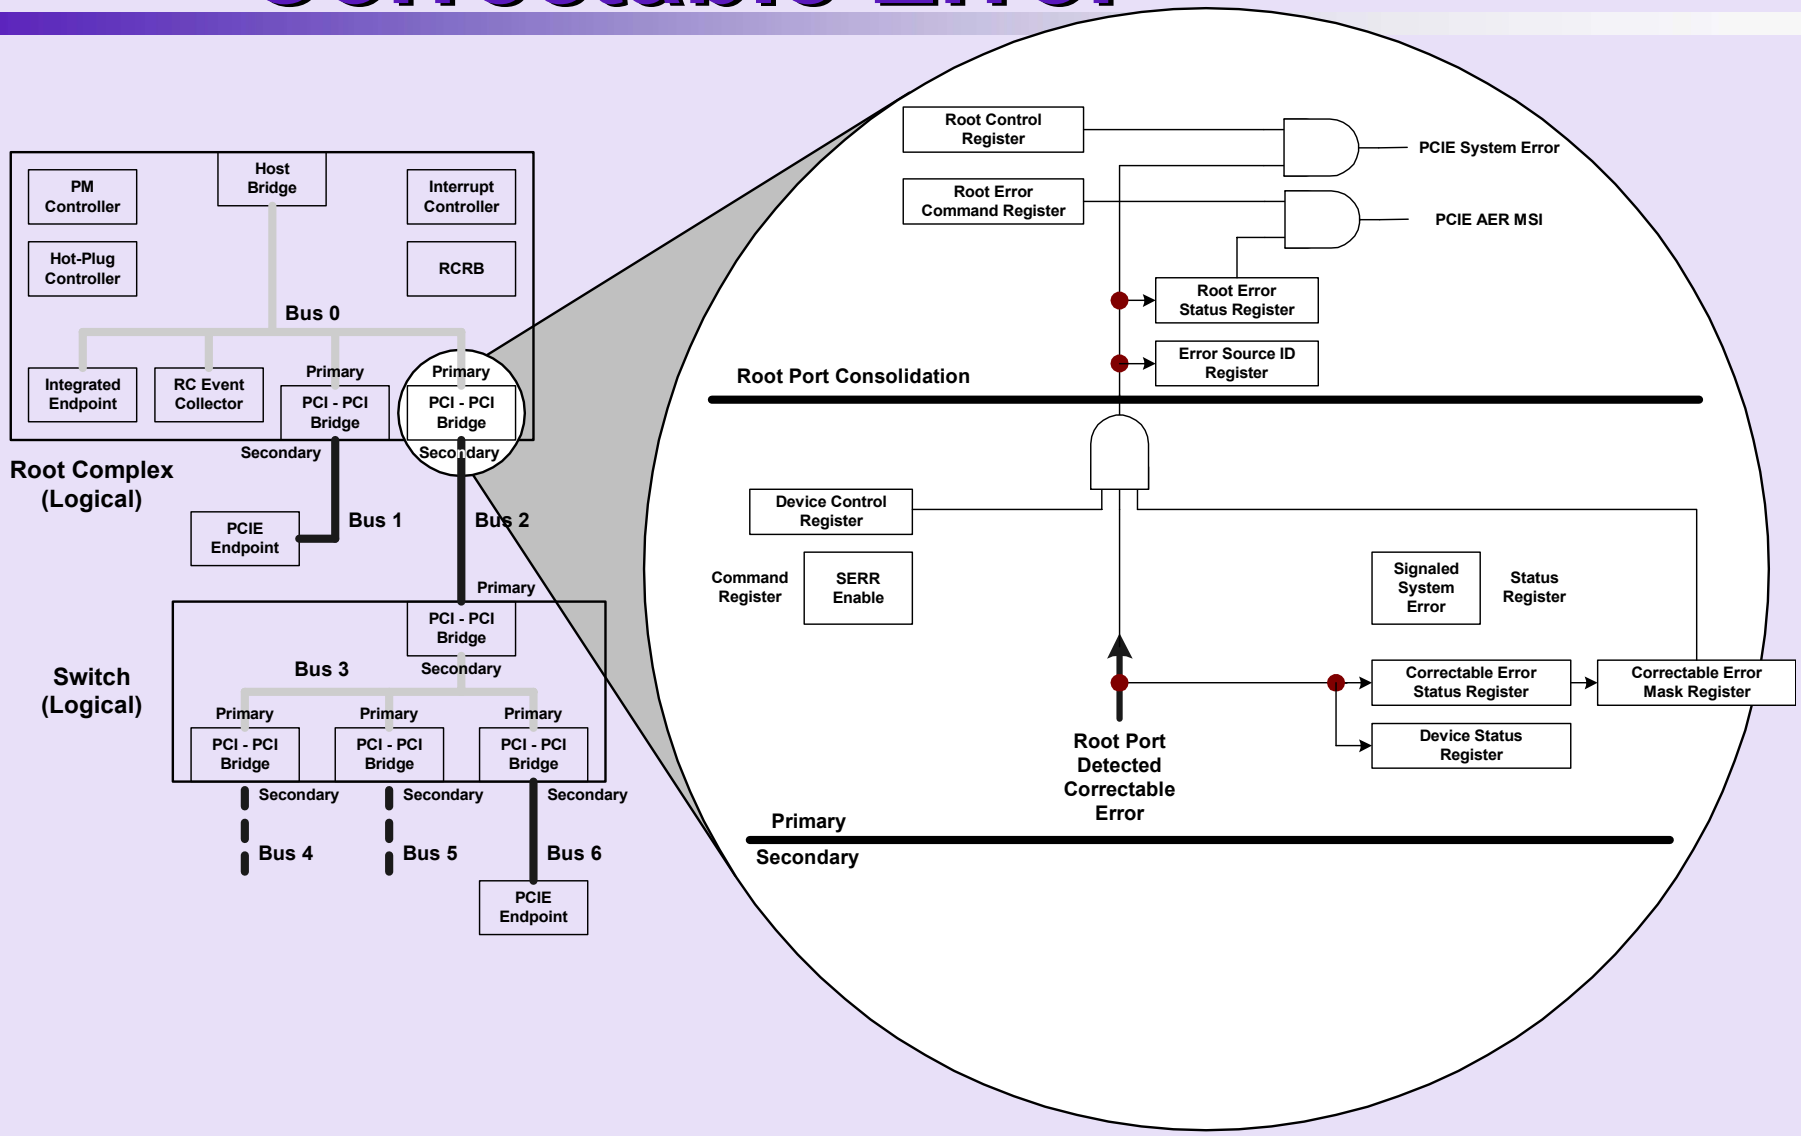

Copyright © 2004, PCI-SIG, All Rights Reserved

RP Error Signaling Example

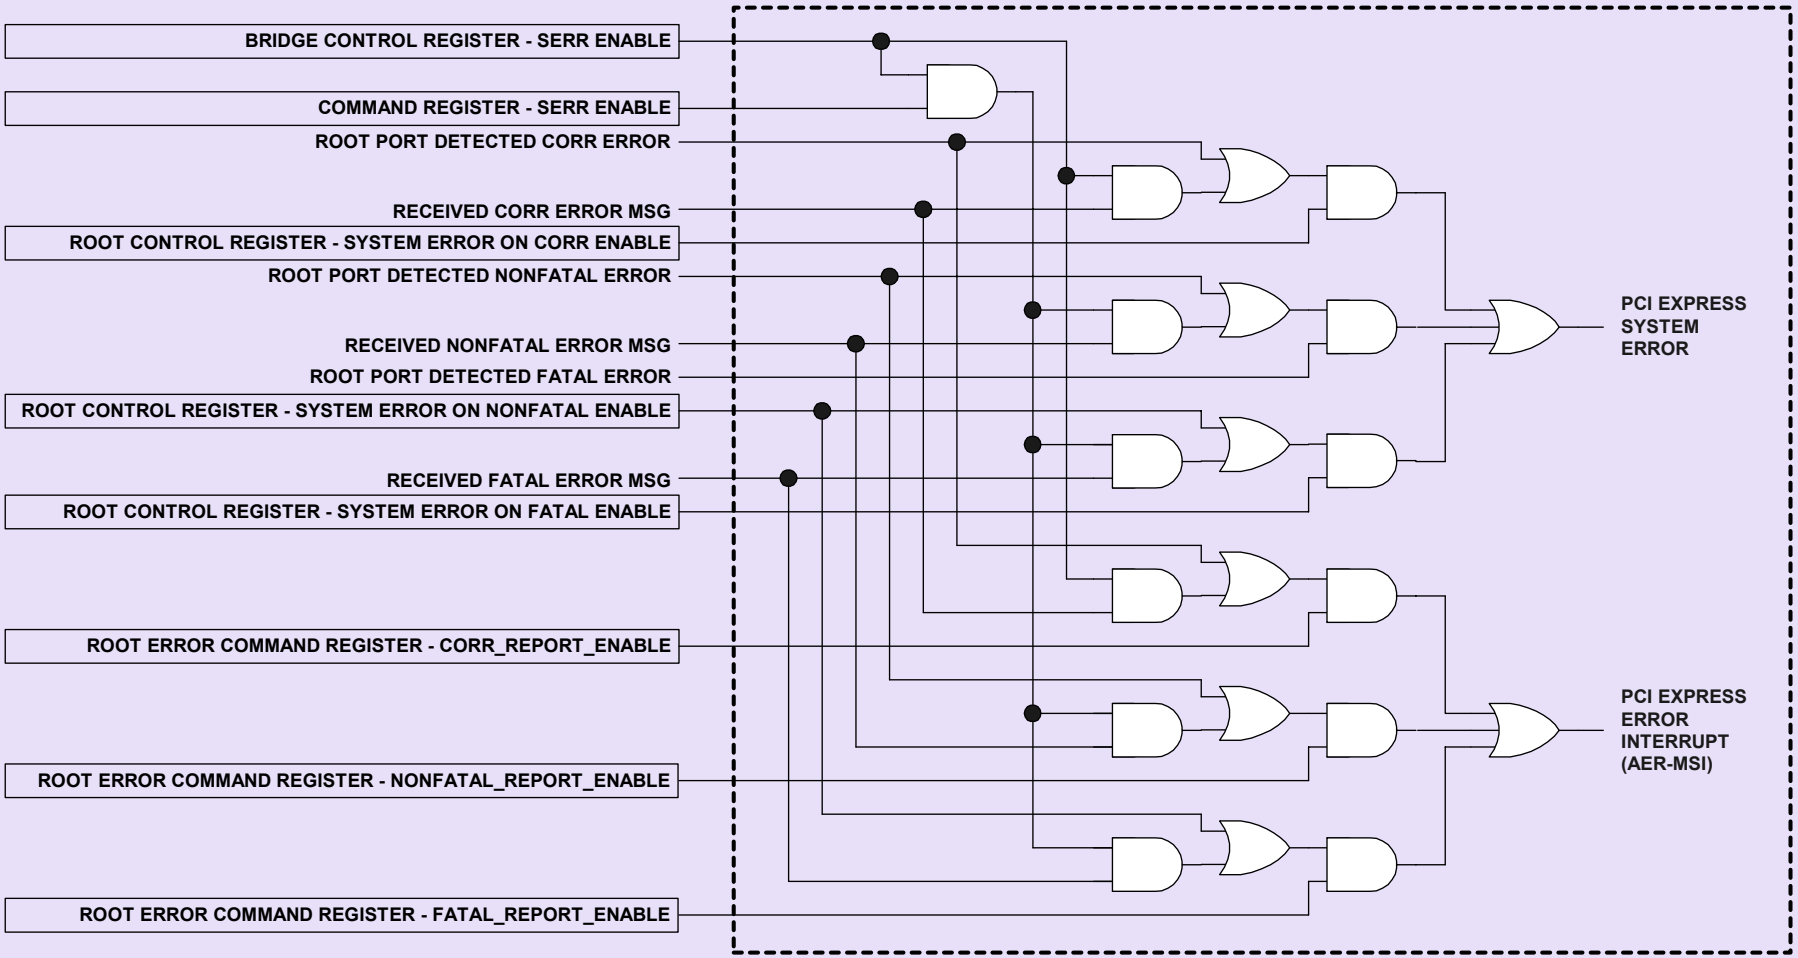

Copyright © 2004, PCI-SIG, All Rights Reserved

Interrupt Considerations

- Several potential interrupt sources
  - ✓ MSI/MSI-X and INTx
  - ✓ PME
  - ✓ Errors and System Errors
  - ✓ Hot-plug events
- Individual interrupts or consolidation?
  - ✓ Number of ports
  - ✓ Size of supported hierarchy
  - ✓ Number of interrupt controllers
- How are interrupts to be terminated in Root Complex?
  - ✓ MSI addresses mapped to interrupt controller?
  - ✓ INTx as physical wires from Root Port to interrupt controller?

Copyright © 2004, PCI-SIG, All Rights Reserved

Power Management Considerations

- L1 and ASPM L0s (as receiver) are required
- ASPM L0s (TX is PCIe optional)
  - ✓ Look for TX support since driver power can be a significant contributor
- ASPM L1 (PCIe optional)
  - ✓ Many early customers have this on their list
- L2 (PCIe optional)
  - ✓ Useful low power state if also enabled for shutting down unused ports
    - Active port (with no device connected) will perform Receiver Detect every 12ms
    - Must have method of shutting down unused ports with no device attached (i.e., not able to perform normal L-state transition)

Copyright © 2004, PCI-SIG, All Rights Reserved

Hot-Plug Considerations

- Controller on-chip or off-chip?
  - ✓ Number of chip-level I/O about the same in either case
- PCIe stack should provide necessary hooks to/from configuration space
  - ✓ Control outputs
  - ✓ Status inputs
  - ✓ Event notification

Copyright © 2004, PCI-SIG, All Rights Reserved

Hot-Plug Controller Concept

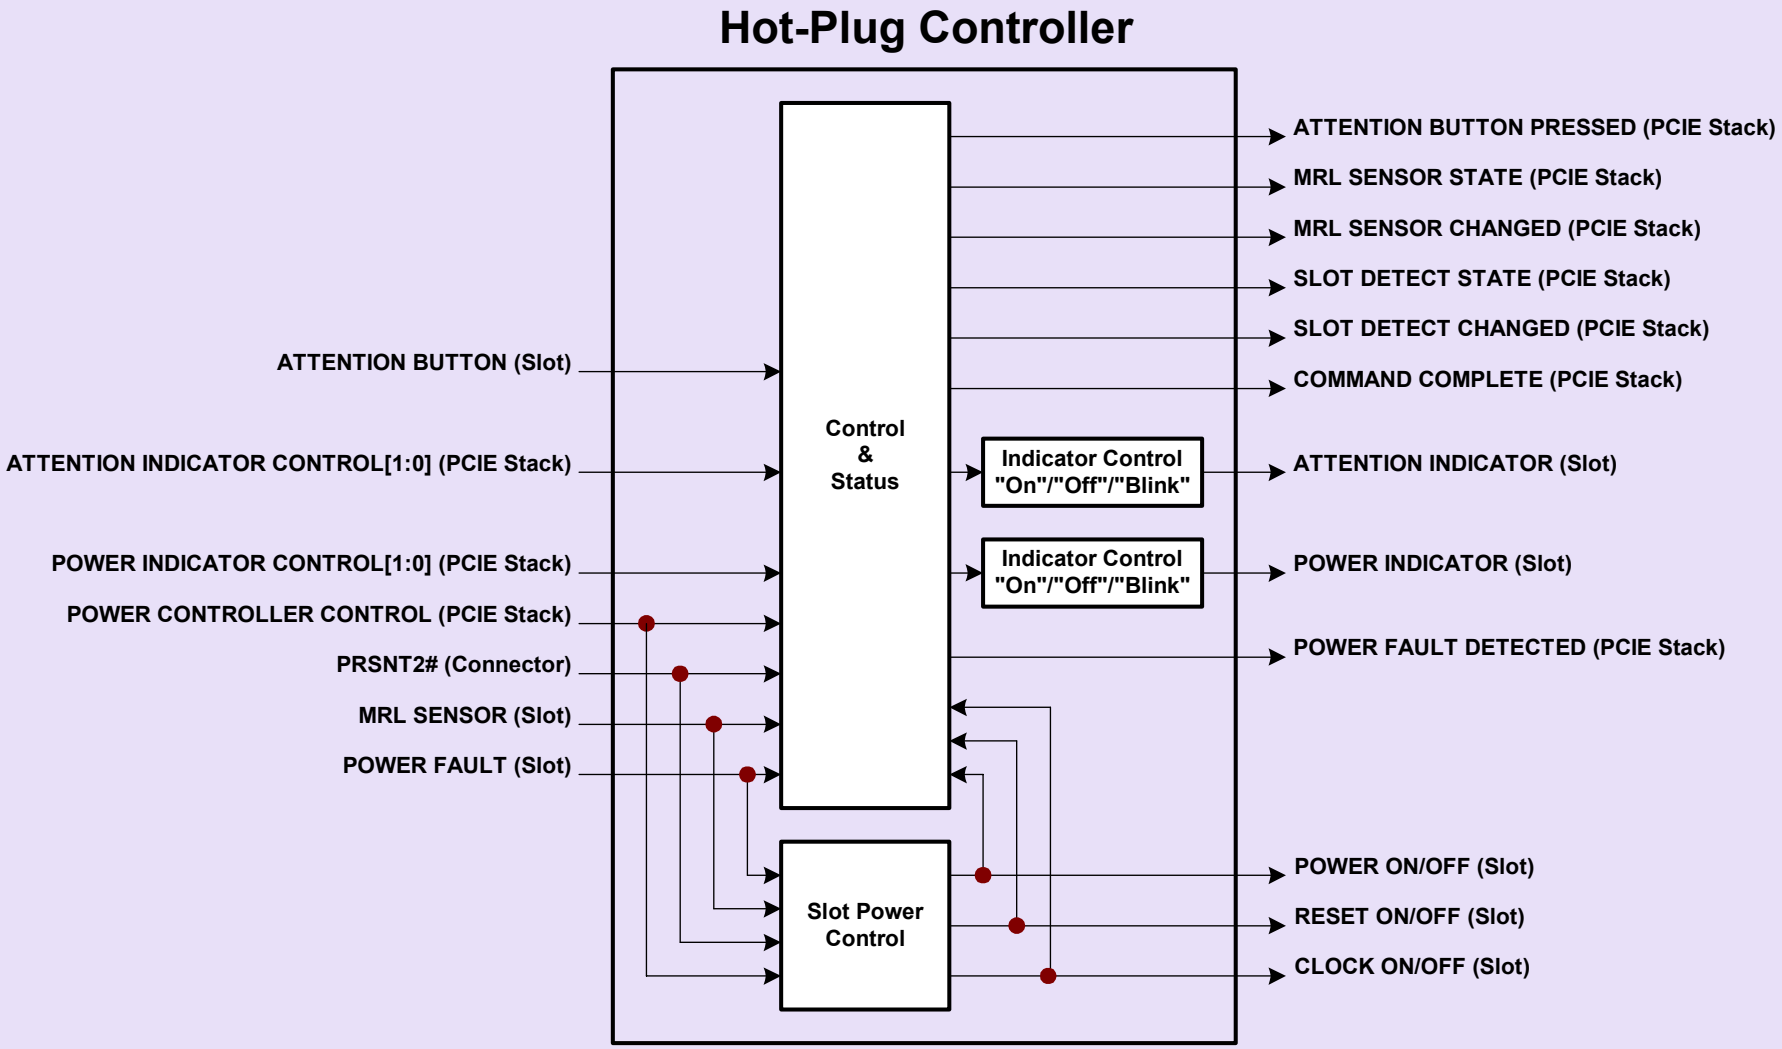

Copyright © 2004, PCI-SIG, All Rights Reserved

Debug Considerations

- Near-end loop back (SerDes specific)
  - ✓ Helpful diagnostic support for high-availability systems
- PCIe (far-end) loop back
  - ✓ Master mode (PCIe optional) support is important
- On-chip logic analyzer access points
  - ✓ Functional partitioning within PCIe stack should provide considerable access
- Performance monitors
  - ✓ Primarily consisting of counters and timers
  - ✓ For multi-port devices, could be consolidated (die area)
- Error Injection
  - ✓ Good for simulation and perhaps for initial system bring-up
  - ✓ Keep it simple – sophisticated mechanisms can be expensive

Copyright © 2004, PCI-SIG, All Rights Reserved

Miscellaneous Considerations

- PIPE (Intel PHY Interface for PCI Express)
  - ✓ If developing upper protocol layers, look for PIPE support from third party PHY vendors
  - ✓ Ensures level of portability between foundries
- PHY Integration Considerations
  - ✓ Please see your PHY provider
- System Memory Map
  - ✓ Address Decode
    - BARs and Base/Limit Registers
  - ✓ Address Translation
    - System specific

Copyright © 2004, PCI-SIG, All Rights Reserved

Agenda

- Introduction
- Example Product Description
- IP Selection Criteria
- Feature Selection Considerations
- Chip Integration Considerations
- Summary

Copyright © 2004, PCI-SIG, All Rights Reserved

Access to PCI Express Interconnect

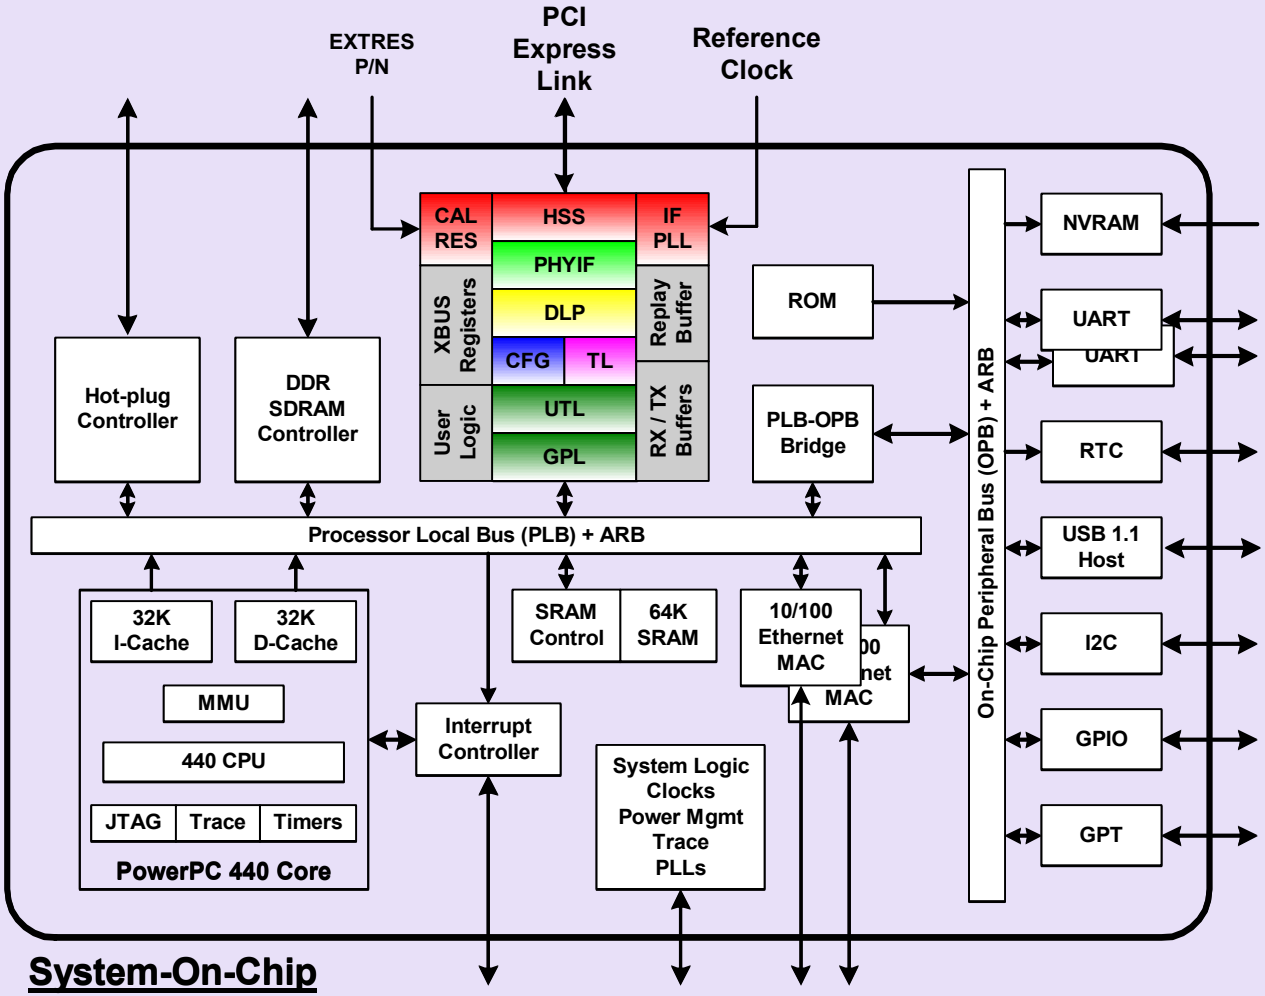

Copyright © 2004, PCI-SIG, All Rights Reserved

Concluding Remarks

- PCI Express SOC implementations should be straightforward
- Just need to keep in mind critical trade-offs…
  - ✓ Features supported
  - ✓ Die area
  - ✓ Performance
- …and consider what to select and integrate
  - ✓ Optional features
  - ✓ Implied requirements
  - ✓ Implementation specific choices

Copyright © 2004, PCI-SIG, All Rights Reserved

Thank you for attending the
PCI-SIG Developers Conference 2004.

For more information please go to
www.pcisig.com

Copyright © 2004, PCI-SIG, All Rights Reserved

Implementing PCI Express on a System-On-Chip (SOC)

Paul J Mattos

PCI Express IP Development Lead

IBM Systems & Technology Group

Copyright © 2004, PCI-SIG, All Rights Reserved


PCI
SIG™