PCIe® 2.0 Base Spec Protocol Updates

Joe Cowan
Computer Systems Architect

Dave Harriman
Sr Staff Architect

Today's Topics:

- Introduction
- Overview of Changes
- Completion Timeout ECN
- Function Level Reset ECR
- 2.0 Base Spec Link Speed Controls
- Link Bandwidth Notification ECR
- Access Control Services ECR

Introduction

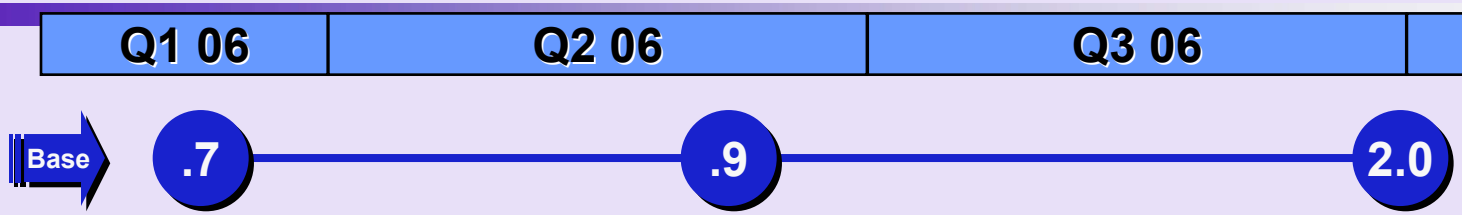

- PCI-SIG® updating PCI Express® (PCIe) specifications this year
- Single largest change – 5 gigabit/second (Gb/s) signaling speed
  - ✓ Optional new capability
- Several other improvements

Today's Focus

Overview of Changes

- Engineering Change Requests (ECRs), Engineering Change Notices (ECNs), and Errata
  - ✓ ECRs become ECNs after review/approval
  - ✓ Completed ECNs & Errata incorporated into 2.0 Base spec
- Errata – highlights:
  - ✓ Root Complex Event Collector Base Class Code conflict
  - ✓ Clarifications of uncommon error cases
  - ✓ Register bit clarifications on defaults, implementation requirements
  - ✓ Interrupt Disable bit consistency with conventional PCI
- 5 Gb/s signaling speed not an ECR/ECN
  - ✓ Included only in 2.0 Base spec

Completion Timeout ECN

- Required: architected disable bit
  - ✓ "Turns off" timeout
  - ✓ Not intended for use during normal operation on most systems
- Optional: programmability of Completion Timeout values
  - ✓ Devices indicate supported ranges from the four bins defined
  - ✓ Two selectable ranges for each bin

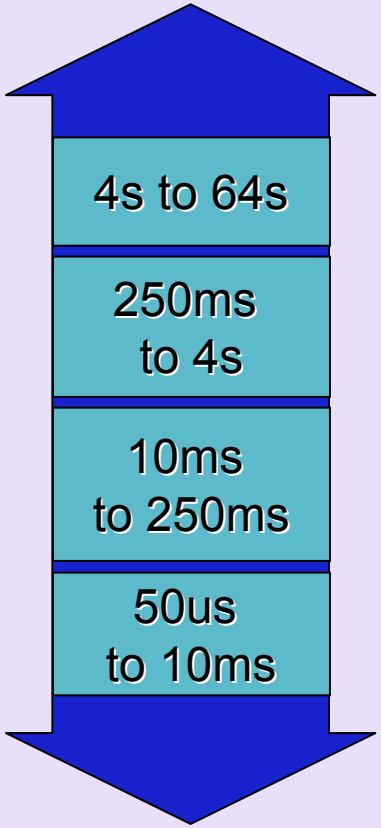

Function Level Reset (FLR) ECR: Background

- Background: new type of reset
  - ✓ Existing resets may (but not required to) reset function internals
  - ✓ FLR definition requires function internal reset
- General concept: SW initiated function-specific reset

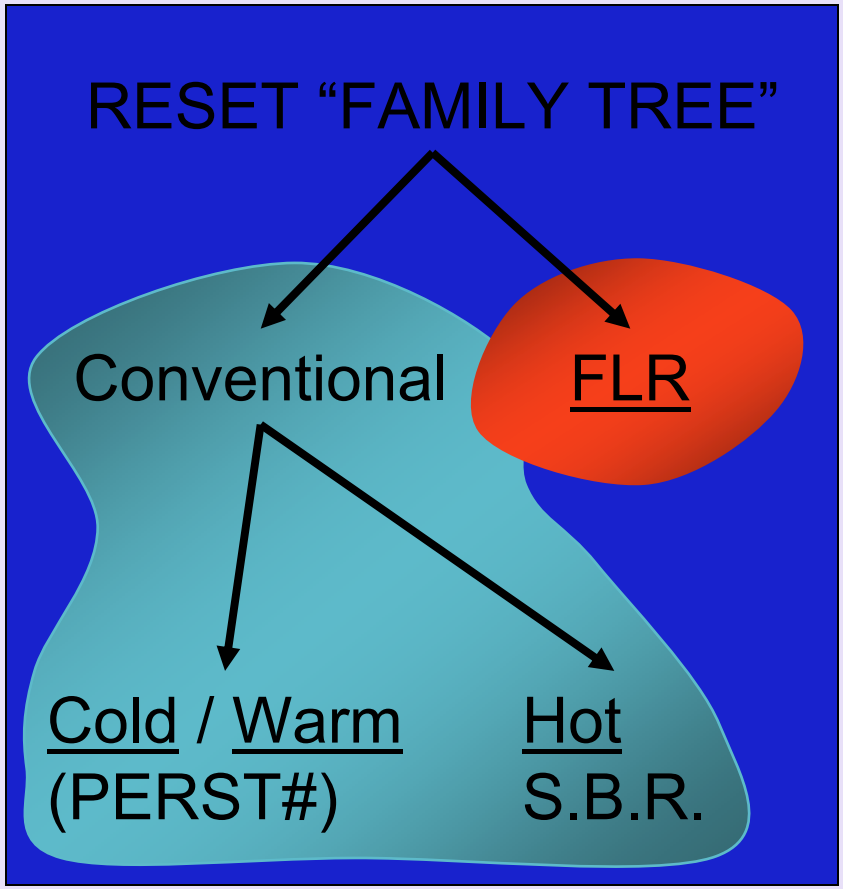

Function Level Reset (FLR) ECR: Details

- Endpoints only
  - ✓ All types: Legacy, Native, Integrated
- Register interface simple
- Implementation & effects potentially complex
  - ✓ Resets internal function-specific state
- Not all architected registers are reset
  - ✓ Hardware Initialized (HwInit), BIOS set, etc.

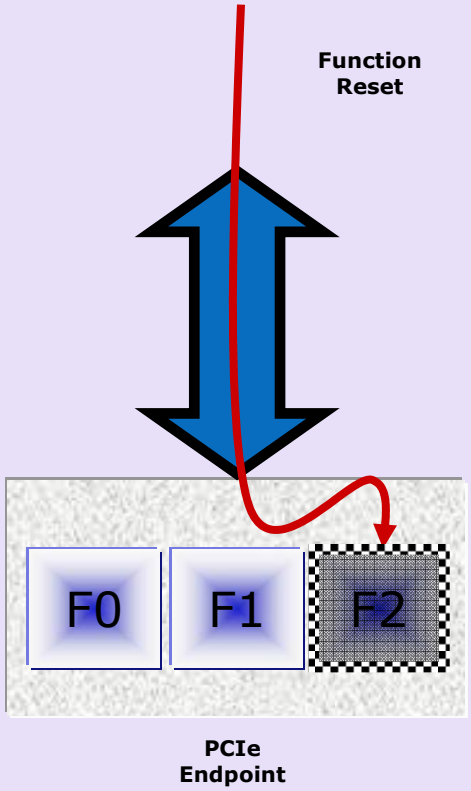

2.0 Base Link Speed Controls: External Link speed management

- By default, hardware automatically trains to the greatest common speed
  - ✓ Software can set an upper bound on the speed
  - ✓ Hardware may limit speed for Link reliability, e.g., if an out-of-spec Link encounters excessive errors
- By default, hardware is permitted to change the speed autonomously for other purposes, such as power management
  - ✓ Software can disable this
- There is a new mechanism supporting software control for entering/exiting Compliance Mode

2.0 Base Link Speed Controls: New/modified regs for external Links

- Link Capability register
  - ✓ Maximum Link Speed field renamed to Supported Link Speeds
- Link Status register
  - ✓ Link Speed field renamed to Current Link Speed
- (new) Link Control 2 register
  - ✓ Target Link Speed field
  - ✓ Hardware Autonomous Speed Disable bit
  - ✓ Enter Compliance bit

2.0 Base Link Speed Controls: Root Complex internal Links

- Can report their supported and current speeds via similar changes to their Link Capability and Status registers
  - ✓ These registers are in the optional Root Complex Internal Link Control Capability, which was introduced by the Root Complex Topology Discovery ECN.
- Speed is not programmable via architected mechanisms

Link Bandwidth Notification ECR: General

- Motivation
  - ✓ Need mechanism for PCIe-aware software to be notified when Link bandwidth (speed or width) changes, due to hardware-autonomous Link retraining
  - ✓ Can help reduce vendor support costs by having software notify users if marginal Links retrain to a lower bandwidth, impacting system performance
  - ✓ Want notification mech available ASAP for all new PCIe components, not just those supporting 5 Gb/s signaling speed
- Though specified in this 1.1 ECR, still logically coupled with Link Speed Controls in 2.0 Base spec
- ECR timing is somewhat tied to Link Speed Controls in 2.0 Base spec stabilizing
- Though an optional ECN for 1.1, plan to make Link Bandwidth Notification a mandatory feature for 2.0

Link Bandwidth Notification ECR: Mechanism details

- Link Capability register
  - ✓ Link Bandwidth Notification Capability bit
- Link Control register
  - ✓ Hardware Autonomous Width Disable bit
  - ✓ Link Bandwidth Management Interrupt Enable bit
  - ✓ Link Autonomous Bandwidth Interrupt Enable bit
- Link Status register
  - ✓ Link Bandwidth Management Status bit
  - ✓ Link Autonomous Bandwidth Status bit

Access Control Services ECR: General

- Set of access control services for Downstream Ports and Functions in Multi-function devices (MFDs)
- New Extended Capability & new Status/Mask/Severity bits in Advanced Error Reporting registers
- Source Validation – Downstream Ports range check Requester ID Bus Number in upstream Requests
- Peer-to-peer (P2P) controls determine whether to forward directly, block, or redirect P2P Requests to the Root Complex for access validation
- Controls being considered for functionality defined by the Address Translation Services (ATS) specification, currently under development

Access Control Services ECR: Details

- Applicable to:
  - ✓ Root Complexes
  - ✓ Switches
  - ✓ Multi-function devices
- Planned services:
  - ✓ Source Validation
  - ✓ P2P Redirect & Upstream Forwarding
  - ✓ P2P Egress Controls
  - ✓ Translation Blocking (ATS)
  - ✓ Direct Translated P2P (ATS)

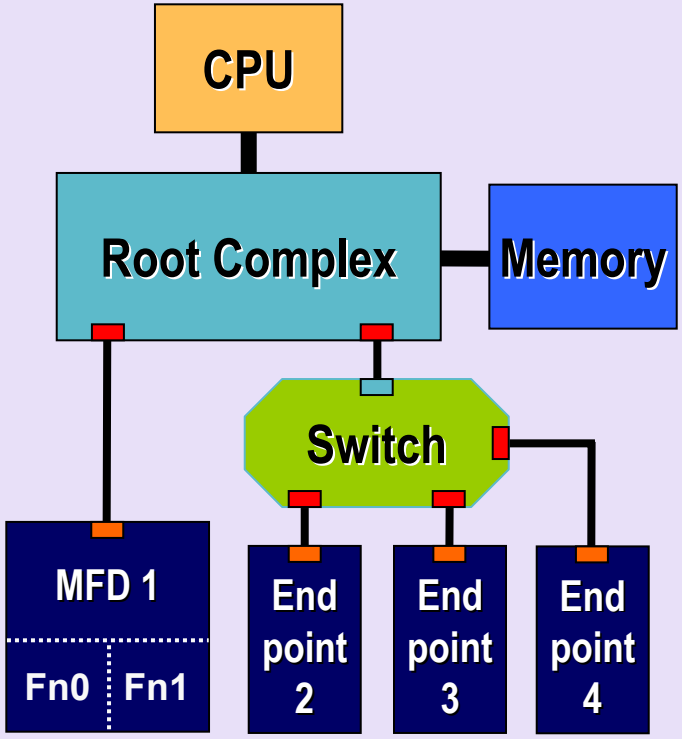

Call To Action

- Comprehend upcoming PCIe 2.0 Base spec improvements
  - ✓ Many enhancements in 2.0 besides 5 Gb/s signaling speed
  - ✓ 2.0 components are not required to implement 5 Gb/s signaling speed
- Start planning your 2.0 product now
- Keep in sync with PCI-SIG for further updates

Thank you for attending the
PCI-SIG Developers Conference 2006.

For more information please go to
www.pcisig.com