



PCIe[®] Protocol Updates Parts I & II

Joe Cowan / Mahesh Wagh
Hewlett-Packard Company / Intel Corporation

Objective

- PCIe[®] protocol working group has resolved PCIe 2.0 errata
 - ✓ Provide update on PCIe 2.0 errata, focus on key critical impact errata items
- PCIe protocol working group is actively developing protocol extension ECRs to PCIe 2.0 specification
 - ✓ ECRs available for PCI-SIG[®] membership review in first half 2008
 - ✓ Plan is to roll completed ECNs in to PCIe 3.0 specification
- Objective of this training material is to
 - ✓ Provide an overview of protocol extensions currently under development
 - ✓ Obtain early and intermediate feedback



PCIe 2.0 Errata

PCIe 2.0 Errata

- Spec errors, clarifications, typos, terminology and formatting grouped into four bins
 - ✓ Bin A: “non-straight-forward changes” – 23 approved as of 4/23
 - ✓ Bin B: “straight-forward changes” – 58 approved as of 4/23
 - ✓ Bin C: “trivial editorial changes” – 23 approved as of 4/23
 - ✓ Bin D: “unclassified” – 25 approved as of 4/23
- As of 4/23, 129 are approved & ~2 errata still being worked
- Formal errata release to PCI-SIG membership planned for Q2'08

PCIe 2.0 Errata Highlights

- Ordering Rules Summary Table and Descriptions overhauled to create a more suitable base for AtomicOps & ID-Based Ordering ECRs
 - ✓ Errata **do not** change ordering semantics, but subsequent ECRs **do**
 - ✓ More details in following slides...
- Several misc clarifications on 5.0 GT/s vs 2.5 GT/s operation
- Several misc clarifications on de-emphasis operation
- Several misc clarifications on cross-link operation
- Various fields: Reserved vs “required to be 0”
 - ✓ Generally prefer to keep cases Reserved to avoid squandering bits
 - ✓ A few cases left “required to be 0” in case some Receivers implemented checks, even though such checks were never intended
- Link Capability bits having same value in all Functions of MFD *associated with an Upstream Port*
 - ✓ Link Cap bits in Upstream Port Functions of MFD all refer to same Link
 - ✓ Link Cap bits in Downstream Port Functions of MFD refer to different Links
- ACS Violation bits in AER are optional normative
 - ✓ Not required to be implemented in AER unless ACS is implemented
- Earlier VC Enable bit clarifications for VC Cap struct carried over to same bit in MFVC Cap struct

Ordering Rules Summary

Table – Key Problems

- Not all components can distinguish Read Completions from I/O & Configuration Write Completions
- Some row/column labels enumerate individual transactions, causing maintenance problems as new transactions are defined, e.g., AtomicOps
- Associated descriptions are excessively wordy and inconsistent in style

Row Pass Column?		Posted Request	Non-Posted Request		Completion	
		Memory Write or Message Posted Request (Col 2)	Read Request (Col 3)	I/O or Configuration Write Request NPR with Data (Col 4)	Read Completion (Col 5)	I/O or Configuration Write Completion (Col 6)
Posted Request	Memory Write or Message Posted Request (Row A)	a) No b) Y/N	Yes	Yes	a) Y/N b) Yes	a) Y/N b) Yes
	Read Request (Row B)	No	Y/N	Y/N	Y/N	Y/N
Non-Posted Request	I/O or Configuration Write Request NPR with Data (Row C)	No	Y/N	Y/N	Y/N	Y/N
	Read Completion (Row D)	a) No b) Y/N	Yes	Yes	a) Y/N b) No	Y/N
Completion	I/O or Configuration Write Completion (Row E)	Y/N	Yes	Yes	Y/N	Y/N

Ordering Rules Summary Table – End Result

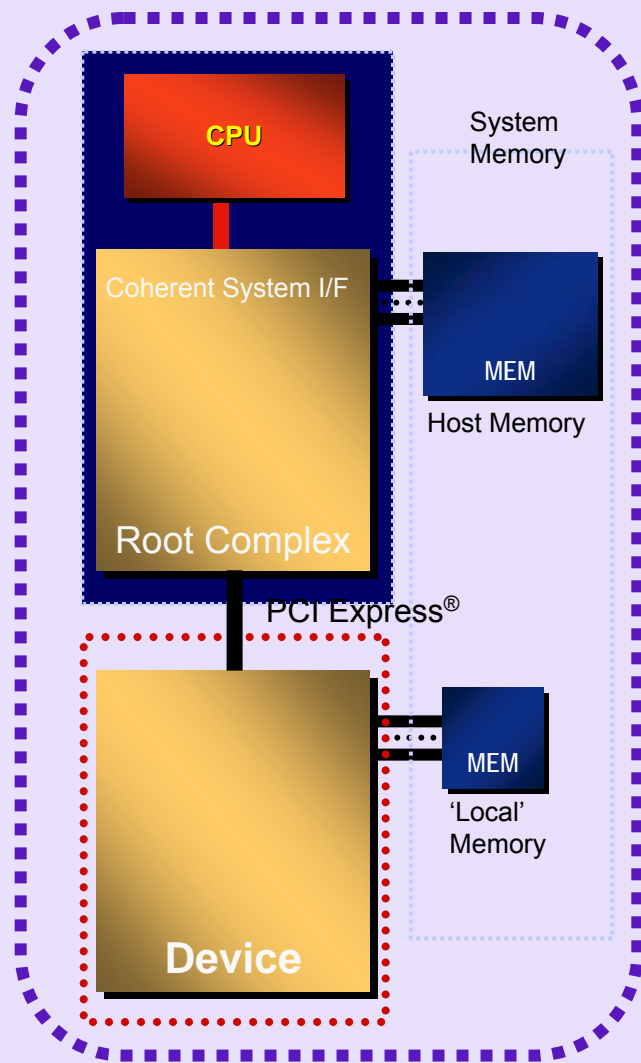
Row Pass Column?		Posted Request (Col 2)	Non-Posted Request		Completion (Col 5)
			Read Request (Col 3)	NPR with Data (Col 4)	
Posted Request (Row A)		a) No b) Y/N	Yes	Yes	a) Y/N b) Yes
Non-Posted Request	Read Request (Row B)	No	Y/N	Y/N	Y/N
	NPR with Data (Row C)	No	Y/N	Y/N	Y/N
Completion (Row D)		a) No b) Y/N	Yes	Yes	a) Y/N b) No

- No longer separate rows/columns for Read Completions vs Config/IO Write Completions
- Row/column labels now consistently based on classes of transactions instead of enumerated individual transactions
- Associated descriptions now more consistently and concisely worded



PCIe Protocol Extensions

PCIe Protocol Extensions



Performance Improvements

- ✓ TLP Processing Hints – Hints to optimize system resources and performance
- ✓ ID-Based Ordering – Transaction-level attribute/hint to optimize ordering within RC and memory subsystem

Software Model Improvements

- ✓ Atomic Operations – Mechanisms to reduce synchronization overhead
- ✓ Page Request Interface – Mechanism incorporated into ATS 1.1 for a device to request faulted pages to be made available

Communication Model Enhancements

- ✓ Multicast – Mechanism to transfer common data or command set from one source to multiple recipients

Power Management

- ✓ Dynamic Power Allocation - Support for dynamic power operational modes through standard configuration mechanism
- ✓ Latency Tolerance & Bandwidth Requirement Reporting (LTBWR)
- ✓ Optimized Buffer Flush/Fill (OBFF)

Configuration Enhancements

- ✓ Resizable BAR– Mechanism for device to negotiate allocated BAR sizes
- ✓ Internal Error Reporting– Extend AER to report component internal errors and support multiple error logs



Protocol Extensions Summary



Extension	Description	Status (as of 4/23)
TLP Processing Hints (TPH)	Hints for optimized TLP processing within host memory/cache hierarchy	30-day Member Review closed 3/6; requires another 30-day Review
ID-Based Ordering (IDO)	New type of relaxed ordering semantics to improve performance	30-day Member Review closes 5/1
Atomic Operations (AtomicOps)	Atomic Read-Modify-Write mechanism	PWG approved final release 4/17
Page Request Interface (PRI; in ATS 1.1 spec)	Mechanism for a device to request for faulted pages to be made available	2-week BoD/WG Review closed 3/28
Resizable BAR	Mechanism to negotiate BAR size	30-day Member Review closed 3/14
Dynamic Power Allocation (DPA)	Mechanism to allow dynamic power mgmt for substates of D0 (active state)	30-day Member Review closed 3/21
Multicast	Address-Based Multicast of Posted Request TLPs	30-day Member Review closed 3/6
Internal Error Reporting	Extend AER to report component internal errors and support multiple error logs	30-day Member Review closed 3/6
LTBWR & OBFF	Mechanisms for platform to tune PM and to align device activities	2-week BoD/WG Review closed 4/3
TLP Prefix	Mechanism to extend TLP headers	Work-in-Progress



ID-Based Ordering (IDO)

Review: PCIe Ordering Rules

"No" entries caused by Producer/Consumer restrictions

Row Pass Column?	Posted Request (Col 2)	Non-Posted Request		Completion (Col 5)
		Read Request (Col 3)	NPR with Data (Col 4)	
Posted Request (Row A)	a) No b) Y/N	Yes	Yes	a) Y/N b) Yes
Non-Posted Request	Read Request (Row B)	No	Y/N	Y/N
	NPR with Data (Row C)	No	Y/N	Y/N
Completion (Row D)	a) No b) Y/N	Yes	Yes	a) Y/N b) No


Table is based on new 2.0 errata!

"Yes" entries are required for deadlock avoidance

- Maximum theoretical flexibility: All entries are "Y/N"
- Traditional Relaxed Ordering (RO) enables A2 & D2 "Y/N" cases
 - ✓ AtomicOps ECR defines an RO-enabled C2 "Y/N" case
- ID-Based Ordering (IDO) enables A2, B2, C2, & D2 "Y/N" cases

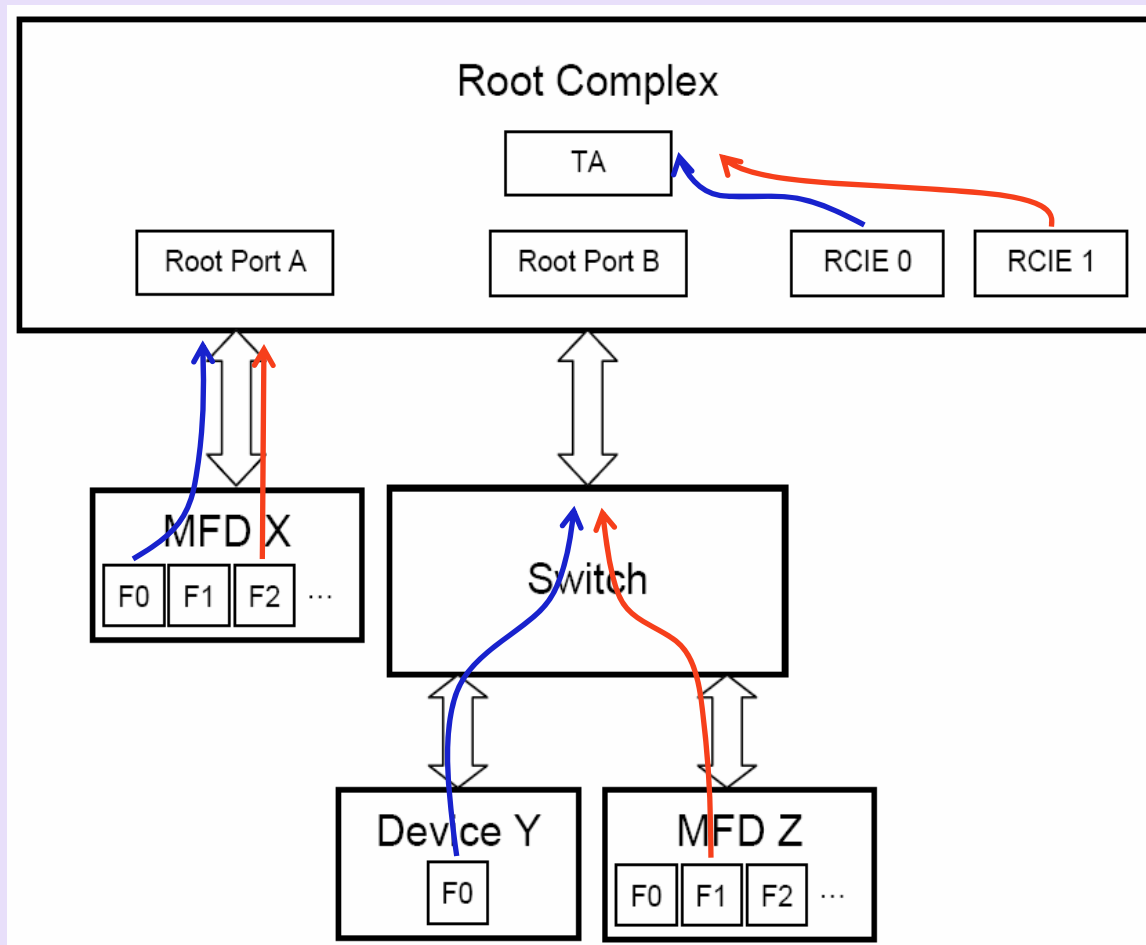
Motivation

- RO works well for single-stream models where a data buffer is written once, consumed, and then recycled
 - ✓ Not OK for buffers that will be written more than once because writes are not guaranteed to complete in order issued
 - ✓ Does not take advantage of the fact that ordering doesn't need to be enforced between unrelated streams
- Conventional Ordering (CO) can cause significant stalls
 - ✓ Observed stalls in the 10's to 100's of ns are seen
 - ✓ Worst case behavior may see such stalls repeatedly for a Request stream
- Consider case of NIC or disk controller with multiple streams of writes:



IDO: Perf Optimizations for Unrelated TLP Streams

- TLP Stream: a set of TLPs that all have the same originator
- Optimizations possible for unrelated TLP Streams, notably with:
 - ✓ Multi-Function device (MFD) / Root Port Direct Connect
 - ✓ Switched Environments
 - ✓ Multiple RC Integrated Endpoints (RCIEs)
- IDO permits passing between TLPs in different streams
- Particularly beneficial when a Translation Agent (TA) stalls TLP streams temporarily



IDO Usage

- IDO Use in Conjunction with RO
 - ✓ IDO and RO are orthogonal, though overlapping in some cases
 - ✓ Highly recommended to use both whenever safely possible
 - ✓ RO permits certain TLP passing within same TLP stream, which is never permitted by IDO
 - ✓ For traffic in different TLP streams, IDO permits control traffic to pass any other traffic, and generally it is not safe to Set RO with control traffic
- IDO safe & simple for Endpoints (EPs) to use with 2-party communication
 - ✓ E.g., an EP communicating directly only with the host
 - ✓ “Simple policy” for EP is to set IDO for all applicable TLPs it originates
- IDO is not simple to use with >2-party communication
 - ✓ E.g., 2 EPs doing peer-to-peer & also communicating via host memory
 - ✓ Generally it is not safe for EP to set IDO for all applicable TLPs
 - ✓ Appendix E documents example communication failure case
- No simple high-benefit use cases identified for RPs setting IDO in TLPs they originate
 - ✓ This is not prohibited, but is outside the scope of the spec

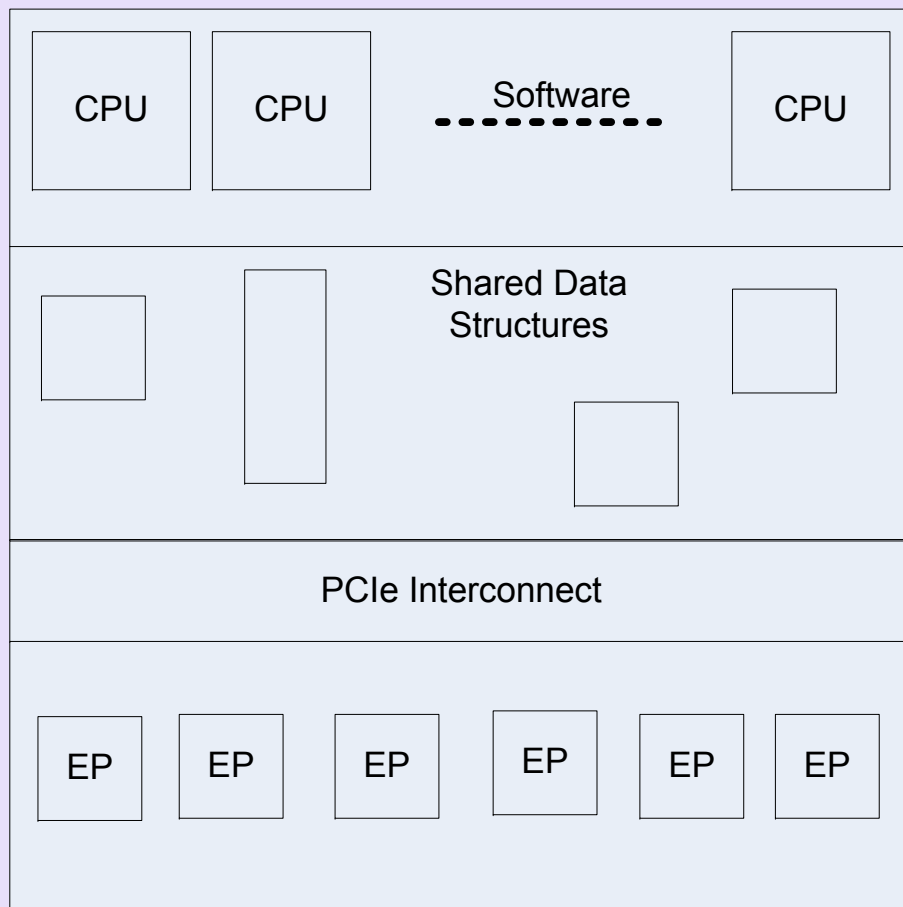
Software Control of IDO Use

- Separate IDO Enable bits for Requests and Completions
- By default, EPs and RPs are not enabled to Set IDO
- Envisioned that software drivers for EPs will enable IDO use
 - ✓ Driver should determine if it is safe to enable IDO use; e.g., if EP does >2-party communication
 - ✓ Driver should know the EP's policy for setting IDO when enabled; e.g., if the EP has advanced policies for >2-party communication
- Envisioned that RPs will not be enabled to Set IDO on most platforms
 - ✓ Enabling RPs to Set IDO is not prohibited, but is outside the scope of the spec



Atomic Operations (AtomicOps)

Motivation



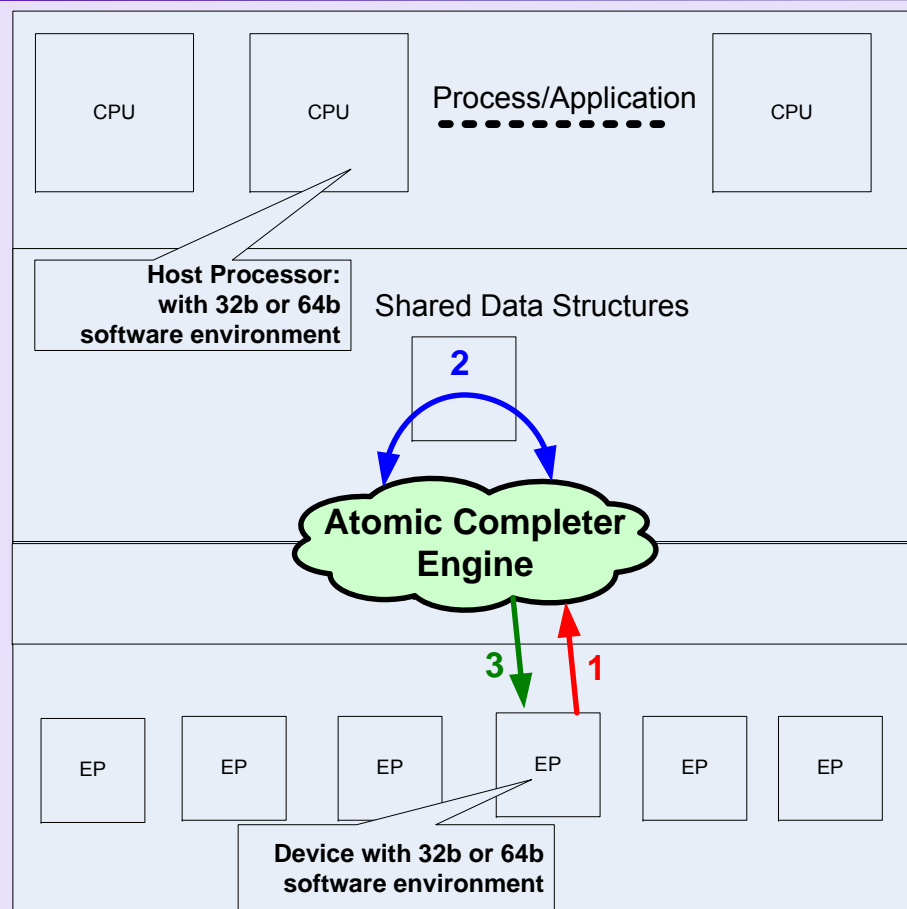
- Atomic transaction support for Host update of main memory exists today
 - ✓ Useful for synchronization without interrupts
 - ✓ Rich library of proven algorithms in this area
- Benefit in extending existing inter-processor primitives for data sharing/synchronization to PCIe interconnect domain
 - ✓ Low overhead critical sections
 - ✓ Non-Blocking algorithms for managing data structures e.g. Task lists
 - ✓ Lock-Free Statistics e.g. counter updates
- Existing application performance can generally be improved by providing synchronization enhancements
 - ✓ Faster packet arrival rates create demand for faster synchronization
- Emerging applications require PCIe synchronization enhancements to support multiple producer – multiple consumer co-ordination
 - ✓ Math, Visualization, Content Processing etc

Atomic Operations

AtomicOp	Description
FetchAdd	$\text{Data}(\text{Addr}) = \text{Data}(\text{Addr}) + \text{AddData}$
Swap	$\text{Data}(\text{Addr}) = \text{SwapData}$
CAS (Compare & Swap)	If $(\text{CompareData} == \text{Data}(\text{Addr}))$ then $\text{Data}(\text{Addr}) = \text{SwapData}$

- Each AtomicOp returns initial Data(Addr)
- Operation sizes supported
 - ✓ 32b and 64b operation sizes for FetchAdd and Swap
 - ✓ 32b, 64b, and 128b operation sizes for CAS
- AtomicOp Request address must be naturally aligned to operation size
 - ✓ AtomicOp data access guaranteed to not cross cacheline boundaries

Atomic Operation Mechanism



- AtomicOp Request
 - ✓ Non-Posted Request with Data
 - ✓ FetchAdd, Swap, or CAS
 - ✓ Multiple outstanding AtomicOps supported
 - ✓ Requestor has to allocate space for Completion before issuing each AtomicOp Request
- Atomic Operation Execution
 - ✓ Atomic Completer Engine performs atomic read-modify-write operation
 - Read initial value
 - Compute and write new value
- AtomicOp Completion
 - ✓ Completion returns initial value
 - ✓ "Standard" Completion for Non-Posted Request; e.g. routed by ID, TID links the Completion with its associated Request

*Device to System memory shown in figure for illustration purposes

Atomic Operations Summary

- Atomic Operations provide capability symmetric to and consistent with capabilities supported by most host CPU architectures
 - ✓ Support core synchronization mechanisms (e.g. semaphores) for CPU/device (as for CPU/CPU), without use of bus-lock type mechanism and minimal use of critical sections
- Enable reuse of existing/debugged algorithms for critical sections, data sharing and queuing implementations
 - ✓ Device side re-use of algorithms already proven on hosts
- Atomic Operations supported: FetchAdd, Swap, and CAS
 - ✓ FetchAdd: Lock-Free Statistics
 - ✓ Swap: Sample & Reset
 - ✓ CAS: Essential for implementation of critical sections and non-blocking queuing algorithms
- Relaxed Ordering (RO) Attribute is permitted with AtomicOp Requests
 - ✓ AtomicOp Request with RO Set is permitted to pass Posted Requests
 - ✓ AtomicOp Completion with RO Set is permitted to pass Posted Requests
- Software discovery & control mechanisms (Device Capability 2 & Control 2 regs)
 - ✓ AtomicOp Routing Supported: necessary since current routing elements do not support
 - ✓ Various AtomicOp Completer bits: 3 bits indicate various type/size combos supported
 - ✓ AtomicOp Requester Enable: prevents AtomicOp use without enablement by software
 - ✓ AtomicOp Egress Blocking: prevents forwarding AtomicOp Requests to components that shouldn't receive them; e.g., those that would handle them as Malformed TLPs



Multicast


Motivation & Mechanism Basics

- Several key applications benefit from Multicast
 - ✓ Communications backplane (e.g. route table updates, support of IP Multicast)
 - ✓ Storage (e.g., mirroring, RAID)
 - ✓ Multi-headed graphics
- PCIe architecture extended to support address-based Multicast
 - ✓ New Multicast BAR to define Multicast address space
 - ✓ New Multicast Capability structure to configure routing elements and Endpoints for Multicast address decode and routing
 - ✓ New Multicast Overlay mechanism in Egress Ports allow Endpoints to receive Multicast TLPs without requiring Endpoint Multicast Capability structure
- Supports only Posted, address-routed transactions (e.g., Memory Writes)
 - ✓ Supports both RCs and EPs as both targets and initiators
 - ✓ Compatible with systems employing Address Translation Services (ATS) and Access Control Services (ACS)
 - ✓ Multicast capability permitted at any point in a PCIe hierarchy


Example

— Multicast Address Route

- Address route Upstream—Upstream Port must be part of the forwarding Ports for Multicast



Multicast Memory Space



Multicast TLP Blocking

- Permits a routing element Ingress Port to block incoming MC TLPs, or an Endpoint Function to block outgoing MC TLPs
- For MC TLPs, extract the Multicast Group (MCG) number from the target address and perform MC Blocked TLP checking based on the MCG number:
 - ✓ If the Multicast TLP came in an Ingress Port of a routing element containing a Multicast Capability structure, block the Multicast TLP if:
 - The associated MC_Block_All bit is Set, or
 - The associated MC_Block_Untranslated bit is Set and the address is Untranslated
 - ✓ If the Multicast TLP is being sent by an Endpoint Function containing a Multicast Capability structure, block the Multicast TLP if:
 - The associated MC_Block_All bit is Set, or
 - The associated MC_Block_Untranslated bit is Set and the address is Untranslated

Multicast Overlay Mechanism

- Mechanism allows an Egress Port to remap the address of an outgoing MC TLP, effectively overlaying Multicast Memory Space on top of a target unicast address range
 - ✓ In a Downstream Port, this enables the MC TLP to target (unicast) Memory Space accepted by an Endpoint that lacks a Multicast Capability structure
 - ✓ At an Upstream Port, Multicast space can be remapped onto a Memory Space range associated with Host Memory
- Due to the address being modified, ECRC (if present) must either be stripped or regenerated
 - ✓ ECRC rules for Multicast Overlay are as shown in table below

MC_Overlay Enabled	TLP has ECRC	ECRC Regeneration Supported	Action if ECRC Check Passes	Action if ECRC Check Fails
No	x	x	Forward TLP unmodified	
Yes	No	x	Forward modified TLP	
Yes	Yes	No	Forward modified TLP with ECRC stripped	
Yes	Yes	Yes	Forward modified TLP with regenerated ECRC	Forward modified TLP with inverted regenerated ECRC



TLP Processing Hints (TPH) (formerly Opaque Processing Hints)

TPH Motivation

- Today's PCIe Requests are coherent with respect to system memory/caches
 - ✓ Root Complex maintains coherency via "snoop"
- Current communication between device and host does not take full advantage of system capabilities
 - ✓ Use of system cache hierarchy can help reduce access latency
 - ✓ Effective use of system resources (system interconnect and memory)
- TLP processing hints provide an opportunity to optimize use of system fabric and improve system efficiency
 - ✓ Facilitate data residency/allocation within system cache hierarchy
 - ✓ Minimize memory access latencies
 - ✓ Reduce memory & system interconnect bandwidth & associated power consumption

TPH Usage Models

Data Structures of interest

- Control Data; e.g., Descriptors
- Headers for protocol processing
- Payload for data copies

Usage Models

- Efficient processing of network I/O in systems with variable access latencies
 - ✓ I/O is increasingly becoming faster (E.g. 10G → 40G → 100G ...)
- Communications adapters in HPC clusters, Database System Clusters & Computational Accelerators
 - ✓ Multi-node barriers and collective operations are often a key performance limit.
 - ✓ Particularly exchange of lock information; often key barrier to scaling
 - ✓ Delay / overhead per operation limits efficiency in some operations (not all), reducing range of applicability.

Request Patterns

Detail

- Device Writes Host Reads (DWHR)
 - ✓ Push data to host agent
 - Reduce average overhead for access by host agent
 - ✓ Utilize opportunistic allocation policy
 - i.e. Update on line allocated and dirty
- DWHR (prioritized)
 - ✓ Push data to host agent
 - Reduce overhead for access by host agent
 - Minimize latency to single traversal through system fabric
 - ✓ Preferential allocation policy
 - i.e. Allocate/update on request
 - ✓ Prioritize retention of line
 - i.e. bias LRU to take advantage of temporal locality
 - ✓ Use for critical communication structures, flags, etc.

Request Patterns

Detail (cont)

- Host Writes Device Reads (HWDR)
 - ✓ Pull data from host agent
 - Minimize latency to single traversal through system fabric
- HWDR (prioritized)
 - ✓ Pull data from host agent
 - Minimize latency to single traversal through system fabric
 - ✓ Prioritize retention of line
 - Take advantage of temporal locality
 - ✓ Use for critical communication structures, flags, etc.

Request Patterns

Detail (cont)

- Device Reads/Writes, Device Reads/Writes (D*D*)
 - ✓ Mem-free state save
 - Reduced latency, improve number of supported contexts
 - Reduce required device resident memory
 - ✓ Utilized system caches for device “read mostly” data structures.
 - ✓ Steering tag is undefined
- Bi-directional shared structure access
 - ✓ Tag key data structures for symmetric frequent access from both Device and Host
 - ✓ System places line in optimal cache in path between Device and Host Agent
 - ✓ Example: Link list structures with insertions and deletions from both Device and Host

TPH Mechanism

- Mechanism to provide processing hints on per TLP basis for Requests that target Memory Space
 - ✓ Enable system hardware (ex: Root-Complex) to optimize on a per TLP basis
 - ✓ Applicable to Memory Read/Write and Atomic Operations
- TPH mechanism utilizes
 - ✓ 10 bits of existing header fields
 - 10-bit hint undergoing standardization in response to feedback
 - ✓ Potentially an additional 24 bits in a future ECR
 - “Extended Hints” via use of TLP Prefix deferred from initial TPH ECR
- Initial 10-bit hints being standardized to enable:
 - ✓ Temporal Re-Use
 - ✓ Association of host processing elements with Requests from specific devices
 - ✓ Location/Level of a targeted cache (host processor core identifier)
 - ✓ Allocation hints
 - ✓ <Not limited to examples provided above>

TPH Standardization

- TH: 1 bit in header to indicate presence of TPH in Memory Read, Memory Write and Atomic Operations TLP
- PH: 2 bits defined in header to indicate Processing Hints

PH[1:0]	Processing Hint	Usage Model
00	Bi-directional data structure	Bi-Directional data structure
01	Requestor	D*D*
10	Target	DWHR HWDR
11	Target with Priority	DWHR (Prioritized) HWDR (Prioritized)

- ST: 8 bits defined in header to carry System specific Steering Tag values
 - ✓ Use of Steering Tags is optional – ‘No preference’ value used to indicate no steering tag preference
 - ✓ Architected Steering Table for software to program system specific steering tag values

TPH Summary

- Mechanism to make effective use of system fabric and improve system efficiency
 - ✓ Reduce variability in access to system memory
 - ✓ Reduce memory & system interconnect BW & power consumption
- Ecosystem Impact
 - ✓ Software impact is under investigation - minimally may require software support to retrieve hints from system hardware
 - ✓ Endpoints take advantage only as needed → No cost if not used
 - ✓ Root Complex can make implementation tradeoffs
 - ✓ Minimal impact to Switches
- Architected software discovery, identification, and control of capabilities
 - ✓ RC support for processing hints
 - ✓ Endpoint enabling to issue hints



Dynamic Power Allocation (DPA)

Motivation

- Current PCIe provides standard Device & Link-level Power Management
- PCIe 2.0 adds mechanisms for dynamic scaling of Link width/speed
- Devices are increasingly higher consumers of system power & thermal budget
- No architected mechanism for dynamic control of device thermal/power budgets
- New PCIe power management mechanisms to complement support on-going industry wide efforts to optimize component and platform power management to meet new customer and regulatory operating requirements

DPA

- Extend existing PCIe device PM to provide active (D0) device power management sub-states
 - ✓ Support for up to 32 sub-states per Function
 - ✓ Maximum power allocation advertised per sub-state in Watts
 - ✓ Contiguously numbered sub-states 0 to N substate w/o gaps
 - ✓ Each successive sub-state reports maximum power allocation less than or equal to prior sub-state
- Software permitted to dynamically configure a Function to operate at any sub-state in any sequence
 - ✓ Function must operate at or below the maximum power allocation corresponding to the selected sub-state
- New DPA capability only for Endpoints
 - ✓ DPA capability to enable software to discover and actively manage Endpoint Function power usage



LTBWR and OBFF Overview (Developed as separate ECRs)

Reducing Platform Power through Latency Tolerance & BW Reporting

- Problem: Current platform PM policies guesstimate when devices are idle (e.g. w/inactivity timers)
 - ✓ Guessing wrong can cause performance issues, or even HW failures
 - ✓ Worst case: PM disabled to allow functionality at cost to power
 - ✓ Even best case not good – reluctance to power down leaves some PM opportunities on the table
 - Tough balancing act between performance / functionality and power
- Scope of problem is increasing
 - ✓ Power efficiency doesn't just apply to mobile any more
 - Such as Energy Star compliance

Wanted: Mechanism for platform to tune PM based on actual device service requirements


Latency Tolerance & Bandwidth Requirement Reporting (LTBWR)

- Concept: Msg for device service latency & BW requirement reporting
 - ✓ PCIe Message sent by Endpoint with required latency & bandwidth
 - Capability to report both snooped & non-snooped values
 - ✓ “Terminate at Receiver” routing
 - multi-Function devices & Switches coalesce Messages from below & send single Message up
- Provides device benefit
 - ✓ Devices can dynamically limit platform PM state as a function of device activity level – avoids performance pitfalls
- Provides platform benefit
 - ✓ LTBWR removes guess work from platform PM, enables greater power savings without impact to performance/functionality

LTBWR enables dynamic power vs. performance tradeoffs at minimal cost impact

Reducing Platform Power through Optimized Buffer Flush/Fill

- Problem statement: devices do not know power state of central resources
 - ✓ “Asynchronous” device activity prevents optimal power management of memory, CPU, RC internals by idle window fragmentation
 - ✓ Premise: If devices knew when to talk, most could easily optimize their Request patterns
 - Result: System would stay in lower power states for longer periods of time with no impact on overall performance
- Optimized Buffer Flush/Fill (OBFF) - a mechanism for broadcasting PM hint to device



How to do OBFF?

- Requirements:
 - ✓ Notify all Endpoints of optimal windows with minimal power impact
 - ✓ Keep it Simple – Maximize cost/benefit
- Solution 1: When possible, use WAKE# with expanded meanings
 - ✓ WAKE# currently defined for use in L2/L3 - Devices drive, input to platform
 - ✓ Define new semantic for OBFF - Platform central resource drives WAKE# for hint to devices
- Solution 2: WAKE# not available – Use PCIe Message
 - Con: Requires waking links – provide protocol mechanism(s) to avoid needless link wakeups
- Platform Firm/Software to Select 1 or 2 per Link
 - ✓ Switches will translate when needed

Greatest Potential Improvement When
Implemented by All Platform Devices



Resizable BAR Capability

Motivation

- **BAR Space Allocation**
 - ✓ Local memory is becoming quite large on some add-in cards
 - ✓ More frequently, all of the requested Memory Space resources are not being allocated for all of the add-in cards
 - Resource is not allocated or Function is forced to report smaller aperture size in order to be allocated
 - ✓ 256 MB seems to be a common limit in 32-bit systems
 - ✓ Current solutions are proprietary and not exposed to all software
- New mechanism allows requested Memory Space resources to be sized appropriately for the system it is in, so that all of the resources can be allocated
- Mechanisms are invoked during enumeration before BAR space allocation

Resizable BAR Capability

- New Resizable BAR Capability structure
 - ✓ Endpoint Functions advertise multiple supported sizes for device Memory Space resources
 - ✓ Controls enable software to select the device's requested BAR sizes
- Hardware utilizes this new capability to report multiple supported sizes for device BARs to resource allocation software
- Resource allocation software will attempt to program the largest acceptable size for each device BAR
 - ✓ System can be configured with optimal resource settings
- Backwards Compatible
 - ✓ Current resource allocation must work for software that does not comprehend this new capability



TLP Prefix Overview

Motivation

- Emerging usage models require increase in header size to carry new information
 - ✓ Example: Multi-Root IOV, Extended TPH
- TLP Prefix mechanism extends the header sizes by adding DWORDs to the front of headers
- Supports mixed topologies
 - ✓ i.e. fabrics with components that support different sets of Prefixes (including none)
- Components can selectively originate Prefixes
 - ✓ e.g. TLPs originating at component X and destined for component A might have a Prefix while TLPs originating at component X and destined for component B might not
 - ✓ Application/Software must ensure that the path from Requestor to Completer supports TLP Prefix (of the appropriate type)
 - e.g. TPH sender using Peer-to-Peer must know destination and adjust
- Baseline Error rules continue to apply
 - ✓ e.g. If a Base TLP Header is Malformed, Prefixes won't change that


Prefix Encoding

	+0				+1				+2				+3			
	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
Byte 0	1	0	0	0	Type	Local Prefix Contents										
Byte 0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
Byte 0	1	0	0	1	Type	End-End Prefix Contents										

	+0				+1				+2				+3			
	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
Byte 0	1	0	0	L	E	Type	Prefix Contents									
Byte 4	1	0	0	L	E	1111	Prefix Expanded Contents									
Byte 8	1	0	0	L	E	1111	Prefix Expanded Contents									

- Base TLP Prefix Size – 1 DW
 - ✓ Appended to TLP headers
- TLP Prefixes can be expanded or stacked
 - ✓ Type encoding of 1111 is used for expansion of previous TLP Prefix
 - ✓ End-End TLP Prefix has an upper limit of 4 DWORDs (16 Bytes)
- Link Local – Where routing elements may process the TLP for routing or other purposes.
 - ✓ Only usable when both ends understand and are enabled to handle link local TLP Prefix
 - ✓ ECRC not applicable
- End-End TLP Prefix
 - ✓ Requires support between the Requester, Completer and routing elements
 - ✓ End-End TLP Prefix not required to but is permitted to be protected by ECRC
 - If underlying Base TLP is protected by ECRC then End-End TLP Prefix is also protected by ECRC
 - ✓ Upper bound of 4DWORDs (16 Bytes) for End-End TLP Prefix
- Note: Bit 4 (LE) indicates Link Local vs. End-End TLP Prefix

Stacked Prefix Example:



- Link Local is first
 - ✓ Starts at 0
 - ✓ Type_{L1}
- End-End #1 follows Link Local
 - ✓ Starts at 4
 - ✓ Type_{E1}
- End-End #2 follows End-End #1
 - ✓ Starts at 8
 - ✓ Type_{E2}
- PCIe Header follows End-End #2
 - ✓ Starts at 12
- Switch routes using Link Local and PCIe Header
 - ✓ ... and possibly additional Link Local DWORDs
 - if more extension bits needed
 - ✓ Malformed TLP if don't understand
- Switch forwards End-End Prefixes unaltered
 - ✓ End-End Prefixes do not affect routing
 - ✓ Up to 4 DWORDs (16 Bytes) of End-End Prefix
- End-End Prefixes are optional
 - ✓ Different End-End Prefixes sequence can be unordered
 - affects ECRC but does not affect meaning
 - ✓ Repeated End-End Prefix sequence must be ordered
 - e.g. 1st Extended TPH vs. 2nd Extended TPH attribute
 - meaning of this is defined by each End-End Prefix



Call to Action

Call to Action

- Innovate and differentiate your products with PCI Express 2.0 industry standard
- Review and provide feedback on PCIe protocol extensions ECRs
- Contribute to the evolution of PCI Express architecture
- Visit www.pcisig.com for PCI Express specification updates

Thank you for attending the
PCI-SIG Developers Conference 2008

For more information please go to
www.pcisig.com