1-512-256-0197
www.mindshare.com
training@mindshare.com

PCI Express 5.0 Architecture

Training

Let MindShare Bring "PCI Express 5.0 Architecture" To Life For You

MindShare's PCI Express System Architecture course starts with a high-level view of the technology to provide the big-picture context and then drills down into the details for each topic, providing a thorough understanding of the hardware and software protocols.

The course describes additional features added to the architecture when moving through the PCIe specification revisions from 1.1 all the way to the latest 5.0. There are a large number of features and optional behaviors described in the PCIe spec. MindShare can customize the course to cover the topics that are most important for your group. Use the course outline below as a guide to request topics you want covered or removed.

You Will Learn:

- PCI Express features and capabilities
- The definition and responsibilities of each of the layers in the interface
- Details of the PCIe packet-based transaction protocol
- The error detection, reporting and possible correction mechanisms
- The address space and packet-routing methods used
- How the various power management techniques work
- The details of the configuration registers that provide control and status visibility to software
- Software enumeration process
- The process of transmitter equalization necessary to operate at 8.0GT/s, 16.0GT/s and 32.0GT/s
- The details of the LTSSM including the concepts important for each state
- Essential features added to Gen4 and Gen5 specifications

For SRIOV and ATS related topics, which are incorporated into the PCIe 5.0 spec, please request our 3-day IO Virtualization for Intel Platforms or IO Virtualization for ARM Platforms as these topics are **NOT** covered in this standard 5-day PCIe course. Additionally, PIPE 6.0 spec related topics are not covered but can be added upon request. We can cover optional features listed below on request.

Who Should Attend?

This course is hardware-oriented but is suitable for both hardware and software engineers because the configuration registers used to control the hardware are covered in detail. The course is ideal for RTL-, chip-, system- or system board-level design engineers who need a broad understanding of PCI Express. Given the in-depth architecture and design details covered, the course is also suitable for chip-level and board-level validation engineers.

Course Length: 5 Days (but customizable to shorter duration)

Course Outline:

Topic	5-day Class
	PCIe 5.0
PCI Architecture Background Foundation	
PCI concepts important for understanding PCI Express	X
Physical Address Spaces	X
Traffic Types (System Memory, PIO and DMA)	X

1-512-256-0197
www.mindshare.com
training@mindshare.com

Typical System Transactions (NVMe Example)	X
PCI Express Features and Architecture Overview	
Layered Architecture	X
ARM example topology	X
TLP, DLLP and Ordered Set Packet Format Overview	X
Protocol Overview	X
Configuration Overview	
Legacy and Enhanced Configuration Access Mechanism (ECAM)	X
Type 0 and Type 1 Headers, Capability and Extended Capability Structures	X
Bus Enumeration	X
HANDS-ON ARBOR LAB: Scan your system and determine topology	X
Address Space and Transaction Routing	
Clarification of Memory space	X
System memory vs MMIO	X
Prefetchable vs Non-prefetchable	X
IO space	X
Setting up the BARs as well as the Base and Limit registers	X
Switch Routing Mechanism	X
HANDS-ON ARBOR LAB: Debug problem with address mapping	X
TLP Format Details	
Normal TLP fields	X
TLP Prefixes	X
10-bit Tags	X
PCI-SIG Vendor-Defined Messages	X
Quality of Service and Arbitration	
TC/VC Mapping	overview
VC Arbitration	overview
Port Arbitration	overview
Multi-function Arbitration	overview
Flow Control	
Flow Control Protocol	X
Scaled Flow Control	X
Link Feature Exchange	X
Flow Control Initialization	X
Runtime Flow Control Update Mechanism	X
Transaction Ordering	
Simplified Ordering Table	X
Relaxed and ID-Based Ordering	X
DLLP Format Details	
DLLPs	X
NOP & Data Link Feature DLLPs	X
ACK / NAK Protocol	
TLP Error Recovery Mechanism	X
Simplified Replay Timer	X
Examples of Numerous Error Scenarios	X
Nullified Packets and Cut-Through Mode Switches	X
Physical Layer Logic (2.5GT/s and 5.0GT/s)	
Block Diagram	X
Ordered Sets	X
Byte Striping/Unstriping	X
Scrambling/Unscambling	X
8b/10b Encoding/Decoding	X

1-512-256-0197
www.mindshare.com
training@mindshare.com

Serializing/Deserializing	X
Physical Layer Logic (8.0GT/s, 16.0GT/s and 32.0GT/s)	
128b/130b Encoding/Decoding	X
Control SKPs	X
Ordered-Set Blocks and Data Blocks	X
Data Streams and Packet Framing	X
Data Parity Checking	X
16.0 & 32.0 GT/s Data Parity Checking	X
Precoding	X
Physical Layer Electrical (all speeds)	
Differential Tx / Rx	X
2.5GT/s and 5.0GT/s De-emphasis	X
8.0GT/s, 16.0GT/s and 32.0GT/s Equalization Concept	X
Rx Equalization	X
Electrical Conditions for different Link States	X
Spread Spectrum Clocking (SSC)	X
Separate Refclk Independent SSC (SRIS)	X
Link Initialization and Training (LTSSM)	
Detect, Polling, Configuration, L0 States	X
Recovery: Link Speed Change	X
Recovery: Equalization Process	X
16.0 GT/s Equalization and Config Structures	X
Negotiation for skipping parts or all of Tx Equalization	X
32.0 GT/s Equalization and Config Structures	X
Recovery: Link Width Change	
L0s, L1, L2, Hot Reset, Link Disable and Loopback States	X
Modified TS1 / TS2s and Alternate Protocol Negotiation	X
Interrupt Support	
Legacy Interrupt Handling	
MSI Interrupts	X
32-bit MSI Data	X
MSI-X Interrupts	X
HANDS-ON ARBOR LAB: Investigate source of MSI(-X) interrupt and delivery	X
Error Detection and Handling	
Correctable, Non-Fatal and Fatal Errors	X
Advisory Non-Fatal Errors	X
Error Subclass field for Correctable Error Messages	
Advanced Error Reporting (AER)	X
HANDS-ON ARBOR LAB: Determine source and error reporting mechanism	X
Power Management	
Device Power States	X
Link Power States	X
L1 Substates	X
Active State Power Management (ASPM) - hardware controlled	X
Software Controlled Power Management	X
Power Management Events (PME, Beacon and #WAKE)	X
Dynamic Power Allocation (DPA)	
Optimized Buffer Flush Fill (OBFF)	
Latency Tolerance Reporting (LTR)	X
System Resets	
Conventional Reset Mechanisms: Cold, Warm and Hot Reset	X
Function Level Reset (FLR)	X

1-512-256-0197
www.mindshare.com
training@mindshare.com

1b/1b Physical Layer (PCIe 6.0)		
	PAM4	
	Precoding, Gray Coding and Forward Error Correction (FEC)	
	LTSSM Updates (TS0s, Tx EQ changes, entering Flit Mode, etc.)	
Flit Mode		
	Flit Format	
	Error Correction and Flit Retry (Standard vs Selective)	
	DLP Bytes in Flit and Flit Types	
	Flow Control (Dedicated vs Shared)	
	TLP Structure (Header Base, OHC, TLP Trailer, etc.)	
	Segments	
	L0p	
Miscellaneous PCIe Features		
	Retimers (introduced in 4.0)	overview
	Lane Margining (introduced in 4.0)	overview
	Flattening Portal Bridge (FPB) (introduced in 4.0)	overview
	Hierarchy ID Reporting (introduced in 4.0)	
	Designated Vendor-Specific Extended Capability (DVSEC) (introduced in 4.0)	
	Enhanced Allocation (introduced in 4.0)	
	Emergency Power Reduction State (introduced in 4.0)	
	System Firmware Intermediary Support (introduced in 5.0)	overview
	Hot Plug	
	Power Budgeting	
	Multi-Casting	
	Protocol Multiplexing (PMUX)	
	Resizable BARs	
	Downstream Port Containment (DPC) and Enhanced DPC (eDPC)	
	Lightweight Notification (can be used for lightweight cache coherency)	
	Process Address Space ID (PASID)	
	Precision Time Measurement (PTM)	
	Device Readiness Status (DRS) and Function Readiness Status (FRS)	

Recommended Prerequisites:

A basic understanding of digital bus architectures such as PCI is recommended.

Course Material:

1) PCI Express Technology eBook (or hardcopy on request) by Mike Jackson and Ravi Budruk
2) Downloadable PDF version of the presentation slides
3) Add-On MindShare Arbor software tool, used for student labs in the class (discounted pricing applies)
4) Add-On Comprehensive PCI Express eLearning course (discounted pricing applies)