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

PCI Express Architecture for Software Engineers

Training

Let MindShare Bring "PCI Express" 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 this protocol.

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

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 PCIe course. We can cover optional features listed below on request.

Who Should Attend?

This course is focused towards software engineers. It covers the standard interactions between software and PCIe hardware and provides the necessary background hardware behavior to understand what implications those interactions will have.

Course Length: 3 Days (but customizable to shorter duration)

Course Outline:

Topic	3-day Class
	PCIe Software
PCI Architecture Background Foundation	
PCI concepts important for understanding PCI Express	X
Physical Address Spaces	X
Traffic Types (System Memory, PIO and DMA)	X
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

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

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	
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	
TLP Format Details	
Normal TLP fields	
TLP Prefixes	
10-bit Tags	
PCI-SIG Vendor-Defined Messages	
Quality of Service and Arbitration	
TC/VC Mapping	
VC Arbitration	
Port Arbitration	
Multi-function Arbitration	
Flow Control	
Flow Control Protocol	
Scaled Flow Control	
Link Feature Exchange	
Flow Control Initialization	
Runtime Flow Control Update Mechanism	
Transaction Ordering	
Simplified Ordering Table	
Relaxed and ID-Based Ordering	
DLLP Format Details	
DLLPs	
NOP & Data Link Feature DLLPs	
ACK / NAK Protocol	
TLP Error Recovery Mechanism	
Simplified Replay Timer	
Examples of Numerous Error Scenarios	
Nullified Packets and Cut-Through Mode Switches	
Physical Layer Logic (2.5GT/s and 5.0GT/s)	
Block Diagram	
Ordered Sets	
Byte Striping/Unstriping	
Scrambling/Unscambling	
8b/10b Encoding/Decoding	
Serializing/Deserializing	
Physical Layer Logic (8.0GT/s, 16.0GT/s and 32.0GT/s)	
128b/130b Encoding/Decoding	
Control SKPs	
Ordered-Set Blocks and Data Blocks	
Data Streams and Packet Framing	

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

Data Parity Checking	
16.0 & 32.0 GT/s Data Parity Checking	
Precoding	

Physical Layer Electrical (all speeds)

Differential Tx / Rx	
2.5GT/s and 5.0GT/s De-emphasis	
8.0GT/s, 16.0GT/s and 32.0GT/s Equalization Concept	
Rx Equalization	
Electrical Conditions for different Link States	
Spread Spectrum Clocking (SSC)	
Separate Refclk Independent SSC (SRIS)	

Link Initialization and Training (LTSSM)

Detect, Polling, Configuration, L0 States	overview
Recovery: Link Speed Change	overview
Recovery: Equalization Process	overview
16.0 GT/s Equalization and Config Structures	overview
Negotiation for skipping parts or all of Tx Equalization	
32.0 GT/s Equalization and Config Structures	overview
Recovery: Link Width Change	
L0s, L1, L2, Hot Reset, Link Disable and Loopback States	
Modified TS1 / TS2s and Alternate Protocol Negotiation	

Interrupt Support

Legacy Interrupt Handling	X
MSI Interrupts	X
32-bit MSI Data	X
MSI-X Interrupts	X
HANDS-ON ARBOR LAB: Investigate source of MSI(-X) interrupt and delivery	

Error Detection and Handling

Correctable, Non-Fatal and Fatal Errors	X
Advisory Non-Fatal Errors	X
Error Subclass field for Correctable Error Messages	X
Advanced Error Reporting (AER)	X
HANDS-ON ARBOR LAB: Determine source and error reporting mechanism	

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

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	

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

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)	
Lane Margining (introduced in 4.0)	
Flattening Portal Bridge (FPB) (introduced in 4.0)	
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)	
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 Comprehensive PCI Express eLearning course (discounted pricing applies)
4) Add-On MindShare Arbor software tool, used for student labs in the class (discounted pricing applies)