1-800-633-1440
www.mindshare.com
training@mindshare.com

Hands-On PCI Express 5.0 Architecture

Training

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

The PCI Express (PCIe) architecture is a high-performance I/O bus used to interconnect peripheral devices in computing and communication platforms. PCI Express has been designed into consumer and high-end PCs, embedded computing, and communication markets and has established itself as the bus of choice for on-board I/O connections. This architecture is governed and defined by the PCISIG (Peripheral Component Interconnect Special Interest Group).

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 is a very large number of features and optional behavior for PCIe. MindShare can customize the course to cover the topics that are most important for your group.

You Will Learn:

- PCI Express features and capabilities
- The definition and responsibilities of each of the layers in the interface
- The error detection, reporting and possible correction mechanisms
- The details of the packet-based protocol used by PCIe
- 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
- 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

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.

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. Because the course contains practical examples of transactions on the various bus interfaces, the course is also suitable for chip-level and board-level validation engineers.

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

Course Outline:

- PCI Architecture Background Foundation
  - PCI concepts important for understanding PCI Express
- PCI Express Features and Architecture Overview
  - Layered Architecture
  - TLP, DLLP and Ordered Set Format Overview
  - Protocol Overview
- Configuration Overview
  - Legacy and Enhanced Configuration Access Mechanism (ECAM)
  - Type 0 and Type 1 Headers, Capability and Extended Capability Register Overview

1-800-633-1440
www.mindshare.com
training@mindshare.com

  - Bus Enumeration
  - **HANDS-ON ARBOR LAB**: Scan your system and determine topology
- Address Space and Transaction Routing
  - Clarification of Memory space
    - System memory vs MMIO
    - Prefetchable vs Non-prefetchable
  - IO space
  - Setting up the BARs (Base Address Registers) as well as the Base and Limit registers
  - Switch Routing Mechanism
  - **HANDS-ON ARBOR LAB**: Debug problem with plug-and-play address mapping
- TLP Format Details
- Quality of Service and Arbitration
  - TC/VC Mapping
  - VC Arbitration
  - Port Arbitration
  - Multi-function Arbitration
- Flow Control
  - Flow Control Protocol
  - Flow Control Initialization
  - Runtime Flow Control Update Mechanism
- Transaction Ordering
  - Simplified Ordering Table
  - Relaxed and ID-Based Ordering
- DLLP Format Details
- ACK / NAK Protocol
  - TLP Error Recovery Mechanism
  - Examples of Numerous Error Scenarios
  - Nullified Packets and Store-and Forward Switches vs. Cut-Through Mode Switches
- Physical Layer Logic (2.5GT/s and 5.0GT/s)
  - Ordered Sets
  - Byte Striping/Unstriping
  - Scrambling/Descambling
  - 8b/10b Encoding/Decoding
  - Serializing/Deserializing
  - Spread Spectrum Clocking (SSC)
    - SRIS (Separate Refclk Independent SSC)
- Physical Layer Logic (8.0GT/s, 16.0GT/s and 32.0GT/s)
  - 128b/130b Encoding/Decoding
  - Data Precoding for 32.0GT/s
  - Ordered-Set Blocks and Data Blocks
  - Data Streams and Packet Framing
- Physical Layer Electrical (all speeds)
  - Differences Between Speeds
  - 2.5GT/s and 5.0GT/s De-emphasis
  - 8.0GT/s, 16.0GT/s and 32.0GT/s Equalization
  - 16.0GT/s and 32.0GT/s Lane Margining at Receivers
- Link Initialization and Training (LTSSM)
  - Detect, Polling, Configuration, L0 States
  - Recovery (Retraining) State
    - Link Speed Change
    - Tx Equalization Process for 8.0GT/s, 16.0GT/s and 32.0GT/s
    - Link Width Change
  - L0s, L1, L2, Hot Reset, Link Disable and Loopback States
- Interrupt Support
  - Legacy Interrupt Handling
  - MSI Interrupts

1-800-633-1440
www.mindshare.com
training@mindshare.com

    - MSI-X Interrupts
    - **HANDS-ON ARBOR LAB**: Investigate source of MSI(-X) interrupt and delivery
- Error Detection and Handling
    - Correctable, Non-Fatal and Fatal Errors
        - Advisory Non-Fatal Errors
        - Error subclass field for Correctable Error Messages
    - Advanced Error Reporting (AER)
    - **HANDS-ON ARBOR LAB**: Determine source and error reporting mechanism
- Power Management
    - Device Power States
    - Link Power States
    - Active State Power Management (ASPM) – hardware controlled
    - Software controlled Power Management
    - Power Management Events (PME, Beacon and #WAKE)
    - Dynamic Power Allocation (DPA), Optimized Buffer Flush Fill (OBFF), Latency Tolerance Reporting (LTR), L1 Sub-States (L1.0, L1.1 and L1.2), Emergency Power Reduction State
- System Resets
    - Conventional Reset Mechanisms: Cold, Warm and Hot Reset
    - Function Level Reset (FLR)
- Hot Plug and Power Budgeting
    - Hot Plug Controller
- Other Features Introduced with PCIe 2.x and 3.x:
    - Internal Error Reporting
    - Multi-Casting
    - Atomic Operations
    - Resizable BARs
    - Alternative Routing-ID Interpretation (ARI)
    - TLP Processing Hints (TPH) and Steering Tags
    - 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)
    - Protocol Multiplexing (PMUX)
    - Address Translation Services (ATS)
    - Access Control Services (ACS)
    - Device Readiness Status (DRS) and Function Readiness Status (FRS)
- Other Features Introduced with PCIe 4.0
    - Support for Retimers
    - Flattening Portal Bridge (FPB)
    - Enhanced Allocation
    - Hierarchy ID Reporting
    - Designated Vendor-Specific Extended Capability (DVSEC)
- Other Features Introduced with PCIe 5.0
    - Negotiation for skipping parts or all of Tx Equalization
    - Alternate Protocol Negotiation
    - System Firmware Intermediary Support
    - Link Activation

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) MindShare Arbor software tool, used for student labs in the class
4) Optional: Comprehensive PCI Express 3.1 eLearning course