

Compute Express Link (CXL) Architecture

Let MindShare Bring “Compute Express Link (CXL) 3.0 with Switches” to Life for You

Compute Express Link (CXL) is a high-bandwidth, low-latency serial bus interconnect between host processors and devices such as accelerators, memory controllers/buffers, and I/O devices. CXL 3.0 is based on PCI Express® (PCIe®) 6.0 physical layer running at 64 GT/s with x16, x8 and x4 link widths. Degraded modes run at 32 GT/s, 16 GT/s and 8 GT/s with x2 and x1 link widths.

CXL interconnect adds coherency and memory semantics, thus allowing for its application in heterogeneous processing systems with a variety of host processors, memory subsystems and peripheral devices interconnected. CXL has applications in standard computer systems, Artificial Intelligence (AI), Machine Learning, communication systems, and High-Performance Computing (HPC). Emerging applications require a diverse mix of CPUs, GPUs, FPGAs, peripherals such as smart NICs, and other accelerators interconnected via an open industry standard protocol with the necessary features which CXL provides. CXL provides a rich set of three protocols that include 1) CXL.io based on PCIe TLP based transactions, 2) CXL.cache and 3) CXL.mem semantics. CXL uses the PCIe stack offering full interoperability with PCIe. CXL.cachemem protocols allow for coherent transactions in memory space. These transactions employ Flit-based packet routing on the Link.

MindShare’s comprehensive CXL 3.0 Architecture course provides a solid foundation of platform architectures and use cases of the three CXL protocols with Type 1, Type 2 and Type 3 devices in platform example. The course details the role of the Transaction Layer protocol, Link Layer including Flit packing and unpacking, ARB/MUX and Flex Bus Logical and Electrical Physical Layer of a CXL port design. We explain enumeration and configuration process during system bring-up with details of configuration and component registers. Other topics include switch architecture, reset, manageability, RAS features, power management, performance considerations and compliance testing. The course describes new features added to the CXL 3.0 spec such as 256B and Port Based Routing (PBR) Flits, CXL.mem Back-Invalidation, Dynamic Capacity Devices (DCD), LD-FAM and G-FAM Devices, Global Integrated Memory (GIM), Multi-Headed (MH) Devices, Hierarchy Based Routing (HBR) and Port Based Routing (PBR), Direct Peer-to-Peer Routing and lots more.

You Will Learn:

- CXL system architectures with Type 1, Type 2 and Type 3 devices
- CXL transaction protocol for CXL.io and CXL.cachemem including support for back-invalidation to CXL.cache protocol and device scaling
- CXL port design constituting Transaction, Link, ARB/MUX and Flex Bus Physical Layers
- CXL switch architecture including new features to CXL 3.0 HBR/PBR Switches (optional topic)
- Enumeration and initialization issues with Configuration and Component register definitions
- Reliability, Availability, Serviceability (RAS) and error handling features
- CXL 3.0 Spec added features such as 256B Flit support, G-FAM/LD-FAM Devices, Dynamic Capacity Devices (DCD), CXL Fabrics

Course Length: 5 Days (with many topics to be self-studied via self-paced eLearning course post class completion)

Who Should Attend: This course is hardware-oriented but is suitable for both hardware design and software engineers given the course covers CXL initialization topics, registers and command APIs. The course is ideal for RTL-, chip-, system- or system board-level design engineers who need a broad and in-depth understanding of CXL system architecture. The course is also suitable for chip-level and board-level validation engineers.

Course Outline:

Topics	5-Day Class	5-Day Class
	CXL 3.0 w/ Switches	CXL 3.0 w/o Switches
CXL Features and Architecture Overview		
CXL 1.1, CXL 2.0 and CXL 3.0 Features	X	X
CXL.mem, CXL.cache and CXL.io	X	X
Type 1 (with cache), Type 2 (with cache and memory, Type 3 (with memory) Devices	X	X
Switch and MLD device architecture overview	X	X
CXL 3.0 features overview	X	X
Layered architecture overview	X	X
Transaction flows in platforms with PCIe, CXL Type 1, 2, 3 Devices	X	X
Cache coherency and MESI protocol tutorial		X
Transaction Layer		
CXL.mem protocol	X	X
CXL.cache protocol	X	X
CXL.mem back-invalidation and CXL.cache device scaling	X	X
CXL.io protocol	X	X
PTH and UIO	X	X
Ordering Rules		
Link Layer		
IO Flit Packing	X	X
Cache/Mem Flit Packing	X	X
Control Flits	X	X
68B Flits	X	X
256B and PBR Flits	X	X
Cachemem Link Layer Retry (LLR) mechanism		X
Flow Control		X
Vital Feature		
ARB/MuX Layer		
Arbitration and Flit Multiplexing/DeMuxing	X	X
vLSM and Link Power Management		X
ALMP Flits		X
Physical Layer		
Protocol ID	X	X
Flit Header	X	X
NOP Flits, IDLE Flits, Latency Optimized Flits	X	X
CRC, FEC	X	X
Synch Header Bypass	X	X
1b/1b encoding	X	X
Link Training	X	X
Reset		
Cold reset		X
Warm reset		X
Hot reset		X
CXL reset		X
Function Level reset		X
Power Management		
Runtime control PM		
Physical Layer PM		
CXL.io PM		
CXL.cachemem PM		
RAS and Error Handling		
EFN Message		X
Vital Handling		X
Error Logging		X

Enumeration and Manageability		
Enumeration in 1.1 and 2.0 Platforms	X	X
Memory Interleaving	X	X
Memory Interleaving in CXL 3.0	X	X
Control and Status Registers		
Configuration registers and DVSEC	X	X
Component Registers	X	X
Backwards compatibility to CXL 1.1	X	X
CXL 3.0 registers	X	X
CXL 2.0 Switches		
Architectural overview with configuration options	X	
Various switch types	X	
Transaction forwarding	X	
Virtual Hierachy	X	
Fabric Manager and commands	X	
CXL 3.0 Switches		
Hierarchy Based Routing (HBR)	X	
Port Based Rounting (PBR)	X	
CXL Fabrics	X	
G-FAM Access Endpoint (GAE)	X	
FAST, LDST and ID-based Rerouting	X	
CXL 3.0 Device Types		
LD-FAM, G-FAM (GFD)	X	X
Multi-Headed SLD, MLD and GFD Devices	X	X
Dynamic Capacity Device (DCD)	X	X
Global Integrated Memory (GIM)	X	X
Shared Memory verses Pooled Memory	X	X

Recommended Prerequisites:

Good working knowledge of PCI Express a must. Computer architecture fundamentals. Knowledge of Intel, AMD or Arm processor architectures. Knowledge of cache coherency and MESI protocol

Course Material:

- 1) Downloadable PDF version of the presentation slides
- 2) CXL 2.0 eLearning course included with course fees unless reduced cost downgraded option