Comprehensive update about CXL ecosystem and performance optimizations

Presenter:

Arun Bosco

Associate Technical Director,

Samsung Semiconductor India Research

Elavarasan Kamaraj

Senior Staff Engineer,

Samsung Semiconductor India Research

Mounika Kantamani

Senior Manager,

Samsung Semiconductor India Research

Karthikesha

Senior Staff Engineer,

Samsung Semiconductor India Research

©2024 Conference Concepts, Inc.
All Rights Reserved

Disclaimer

- This presentation and/or accompanying oral statements by Samsung representatives collectively, the “Presentation” is intended to provide information concerning the SSD and memory industry and Samsung Electronics Co., Ltd. and certain affiliates (collectively, “Samsung”).While Samsung strives to provide information that is accurate and up-to-date, this Presentation may nonetheless contain inaccuracies or omissions. As a consequence, Samsung does not in any way guarantee the accuracy or completeness of the information provided in this Presentation.

- This Presentation may include forward-looking statements, including, but not limited to, statements about any matter that is not a historical fact; statements regarding Samsung’s intentions, beliefs or current expectations concerning, among other things, market prospects, technological developments, growth, strategies, and the industry in which Samsung operates; and statements regarding products or features that are still in development. By their nature, forward-looking statements involve risks and uncertainties, because they relate to events and depend on circumstances that may or may not occur in the future. Samsung cautions you that forward looking statements are not guarantees of future performance and that the actual developments of Samsung, the market, or industry in which Samsung operates may differ materially from those made or suggested by the forward-looking statements in this Presentation. In addition, even if such forward-looking statements are shown to be accurate, those developments may not be indicative of developments in future periods.

©2024 Conference Concepts, Inc.
All Rights Reserved

Agenda

- Performance optimizations by CXL ecosystem
- CXL.MEM and CXL.IO Flow
- Command processing via In-Band(IB) path
- Fabric Manager(FM) API in Pooled Memory Management
- Command processing via Out-of-Band(OOB) path
- Summary

©2024 Conference Concepts, Inc.
All Rights Reserved

Performance Optimizations by CXL

- CXL (Compute Express Link) is a high speed cache coherent interconnect
- As a protocol CXL features helps in overall performance optimizations

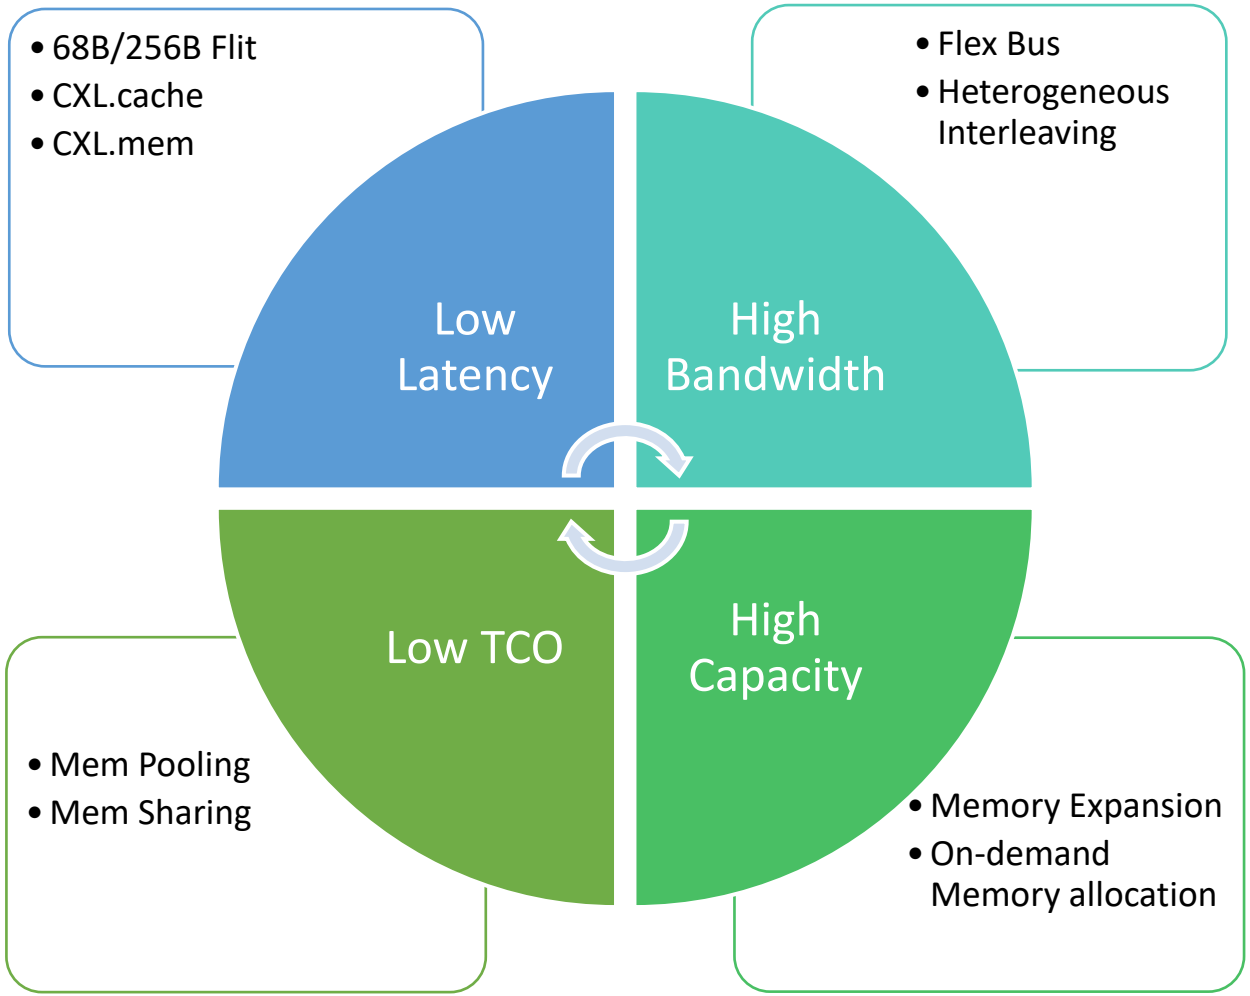

©2024 Conference Concepts, Inc.
All Rights Reserved

Flit (Flow Control Unit)

CXL Device Access by Host

- Host CPU uses the cxl.io, cxl.mem protocols to talk to the device via in-band path
- Host/Device uses cxl.cache to coherently access memories
- Host uses MCTP protocol to talk to the device via side-band path

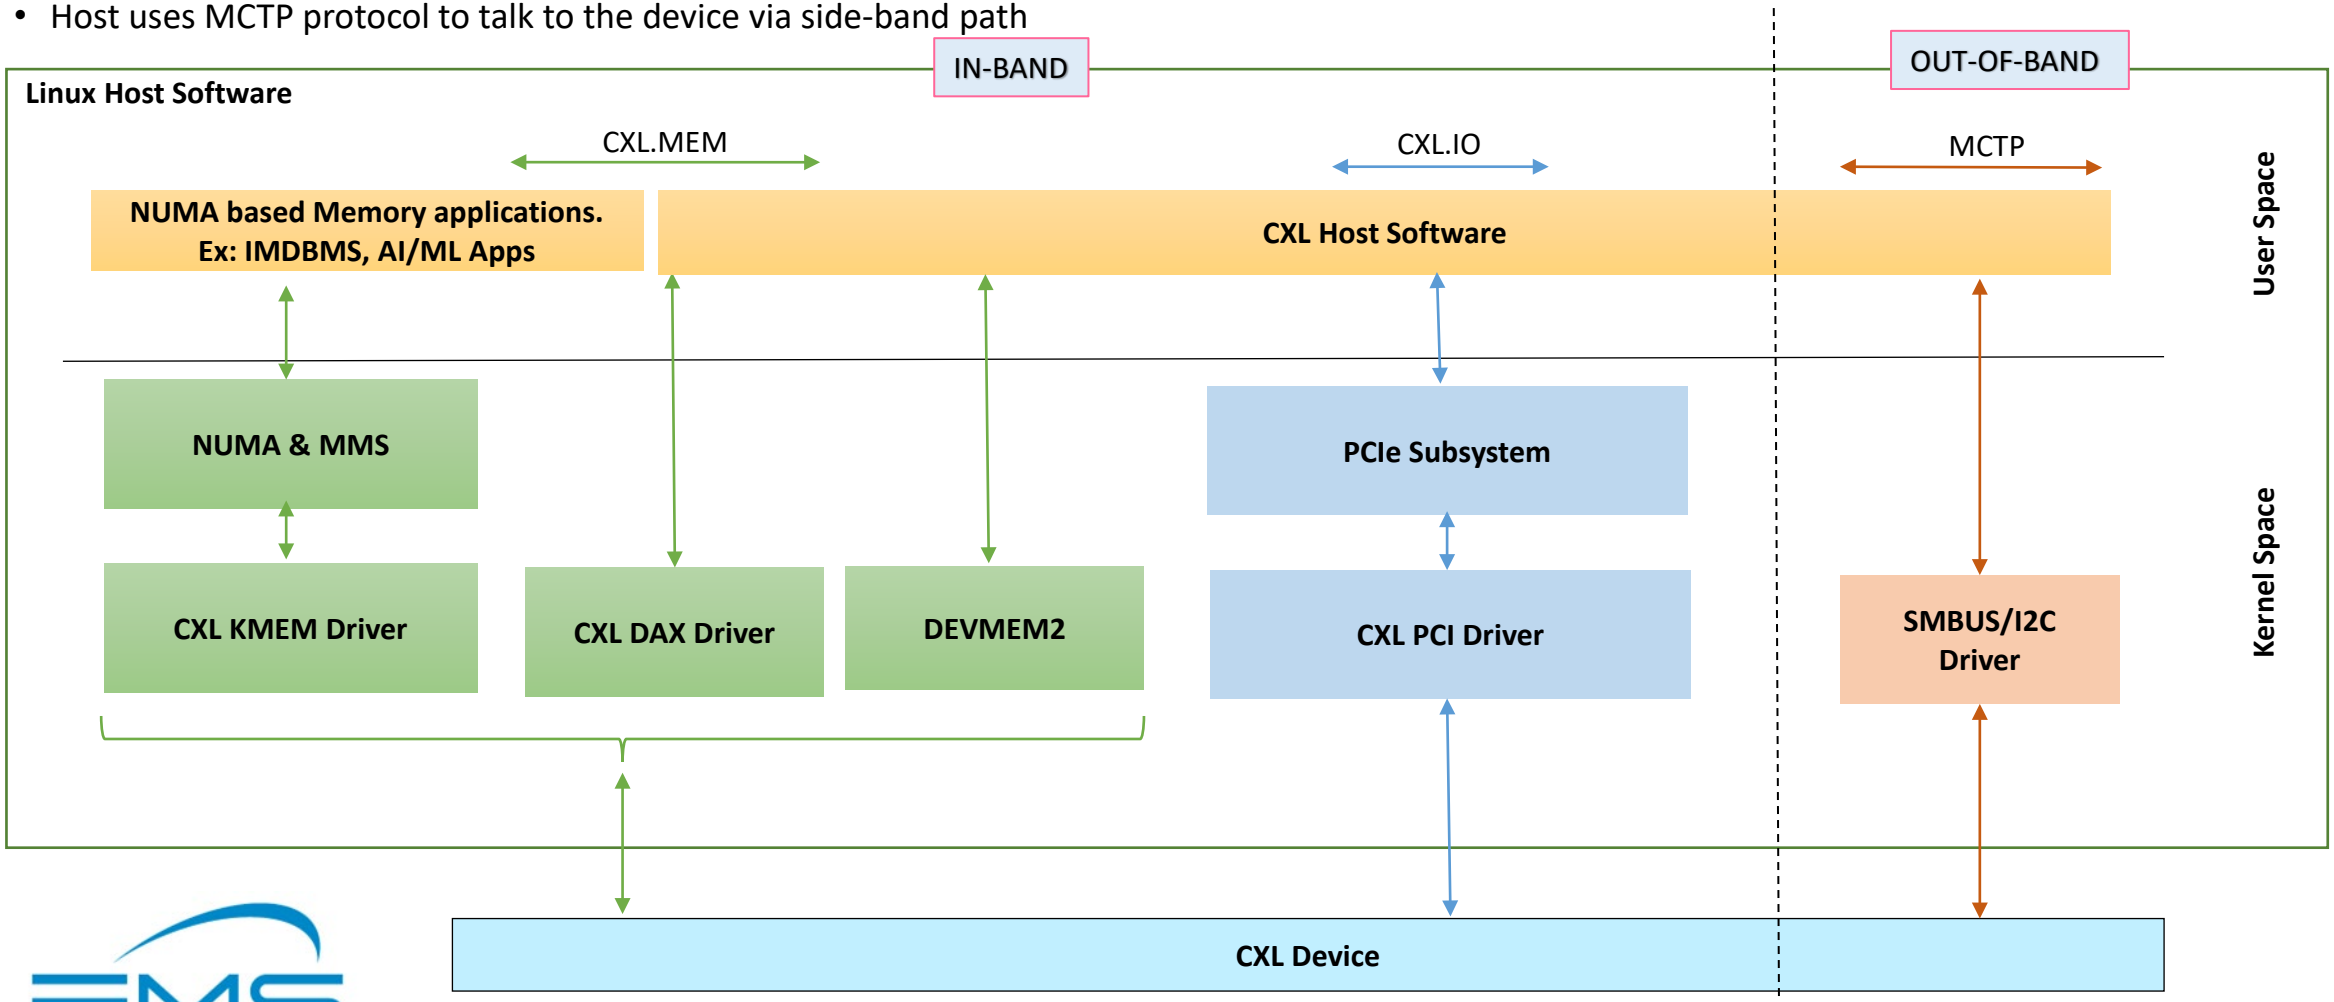

©2024 Conference Concepts, Inc.
All Rights Reserved

NUMA(Non-uniform memory architecture)
DAX(Direct Access Node)

CXL.MEM Host to Device Memory Access

- Transactional interface between CPU and Memory
- Grouping of processors to memories maps memory to user space

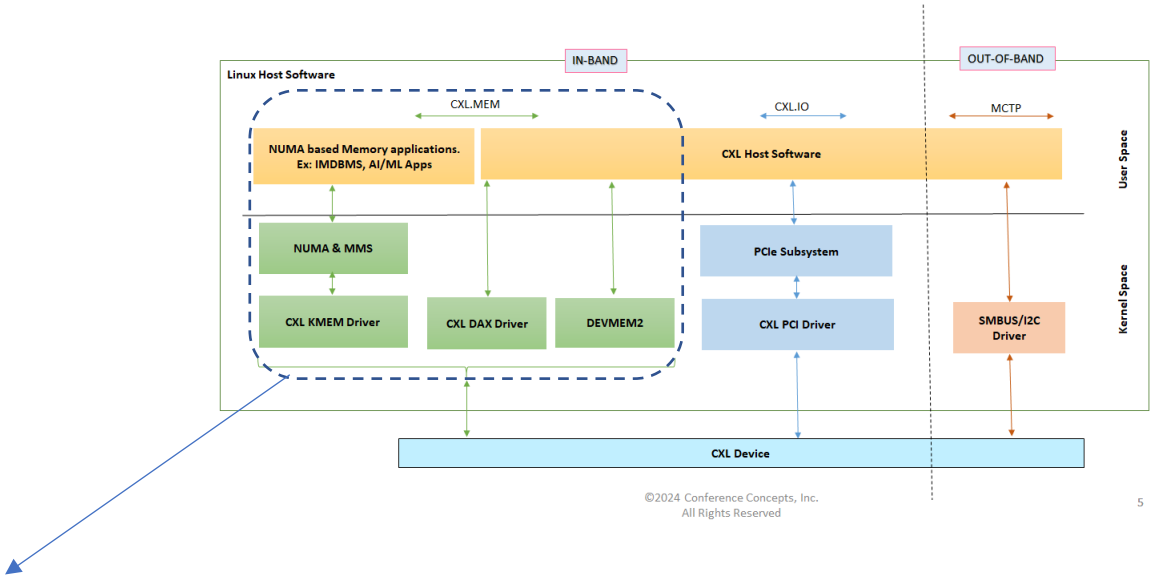

NUMA node

- Non-Uniform Memory Access node
- Kernel configuration as KMEM (Kernel Memory)
- Parallel access to System DRAM and CXL memory
  - Heterogeneous interleaving
  - Performance validation

DAX node

- Direct Access node
- Kernel configuration as PMEM (Persistent Memory)
- Parallel access to System DRAM and CXL memory
  - Applications directly access the memory with addresses

System Extended

- Extends CXL memory to the system memory as a single node
- Sequential access to System DRAM and CXL Memory
  - On-demand memory allocation

©2024 Conference Concepts, Inc.
All Rights Reserved

CXL.IO Command submission via In-Band(IB) path - 1

- Register Locator DVSEC maps to mailbox registers
- In-band commands submitted through mailbox registers

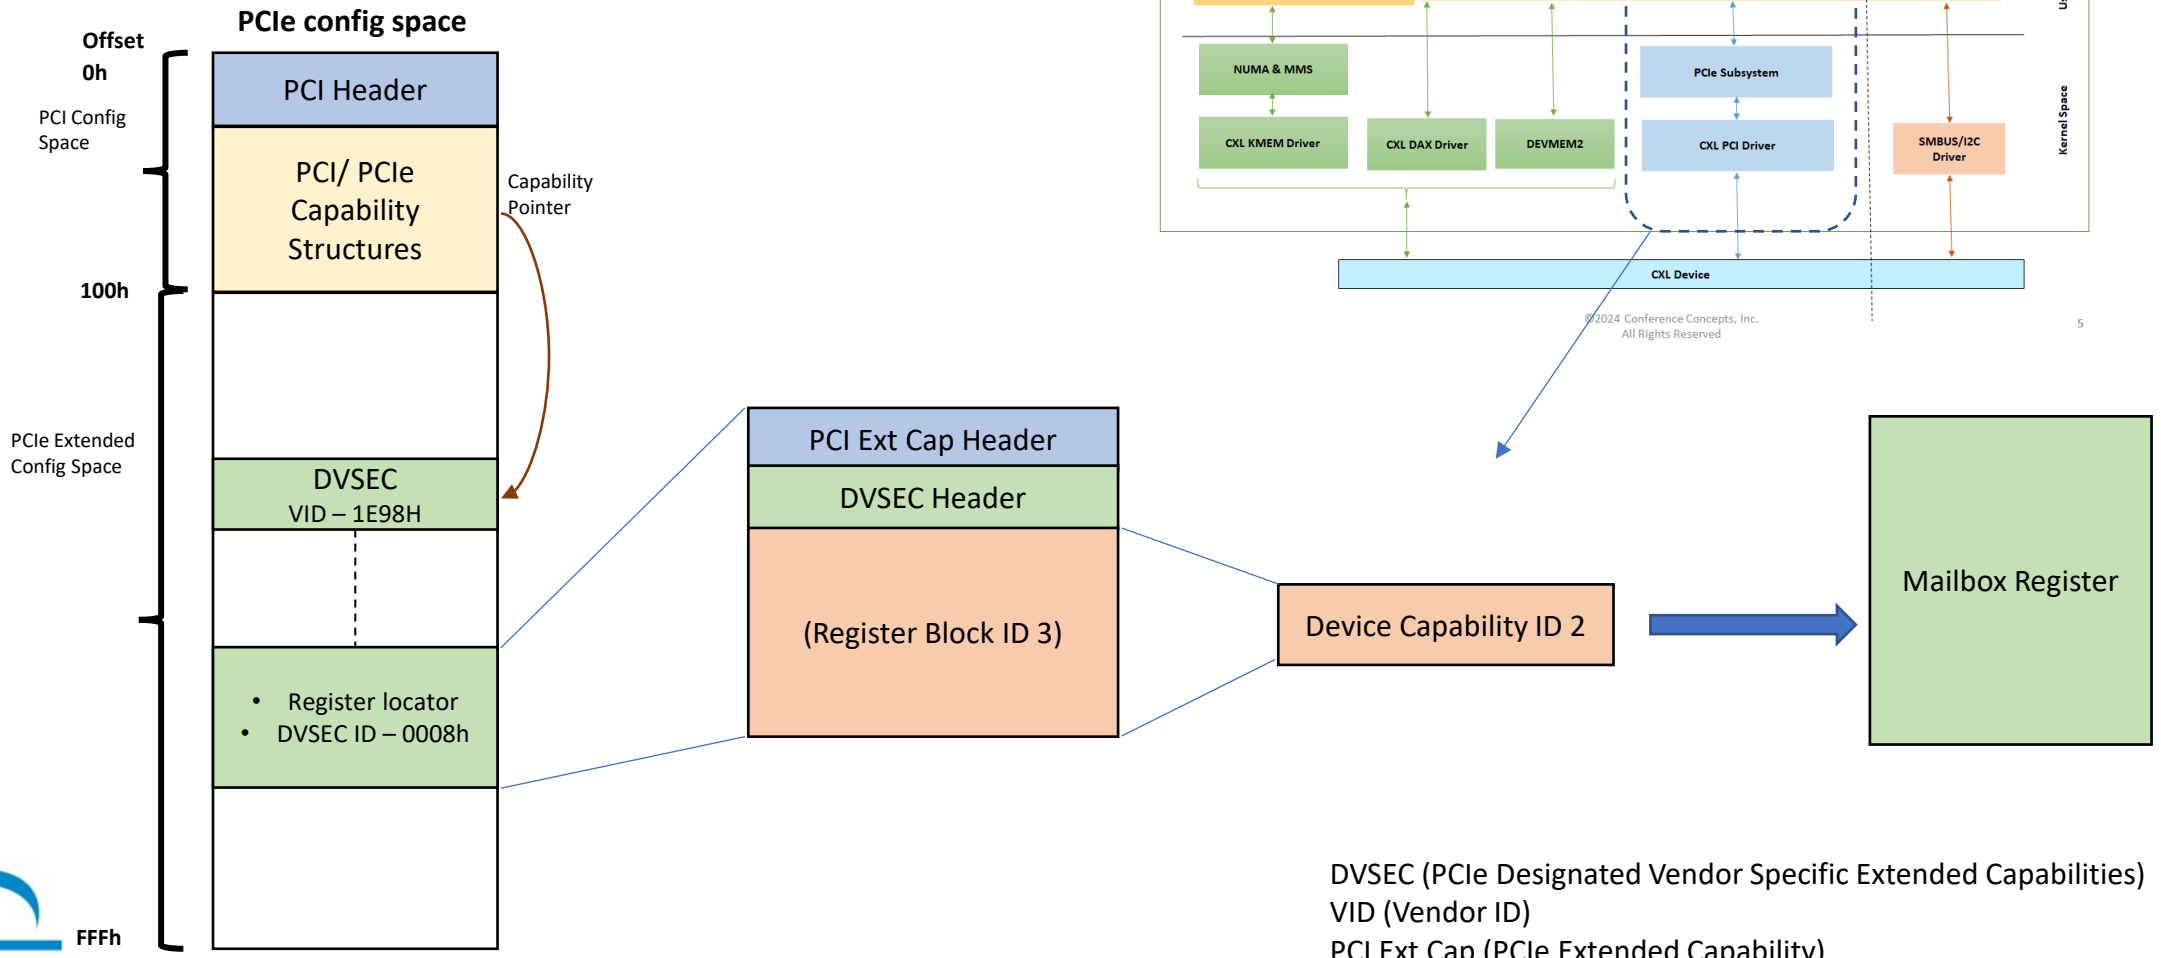

DVSEC (PCIe Designated Vendor Specific Extended Capabilities)
VID (Vendor ID)
PCI Ext Cap (PCIe Extended Capability)

©2024 Conference Concepts, Inc.
All Rights Reserved

CXL.IO Command processing via In-Band(IB) path - 2

- Command flow sequence from inception to completion and consumption
- Control Registers
- Command Registers
- Status Registers

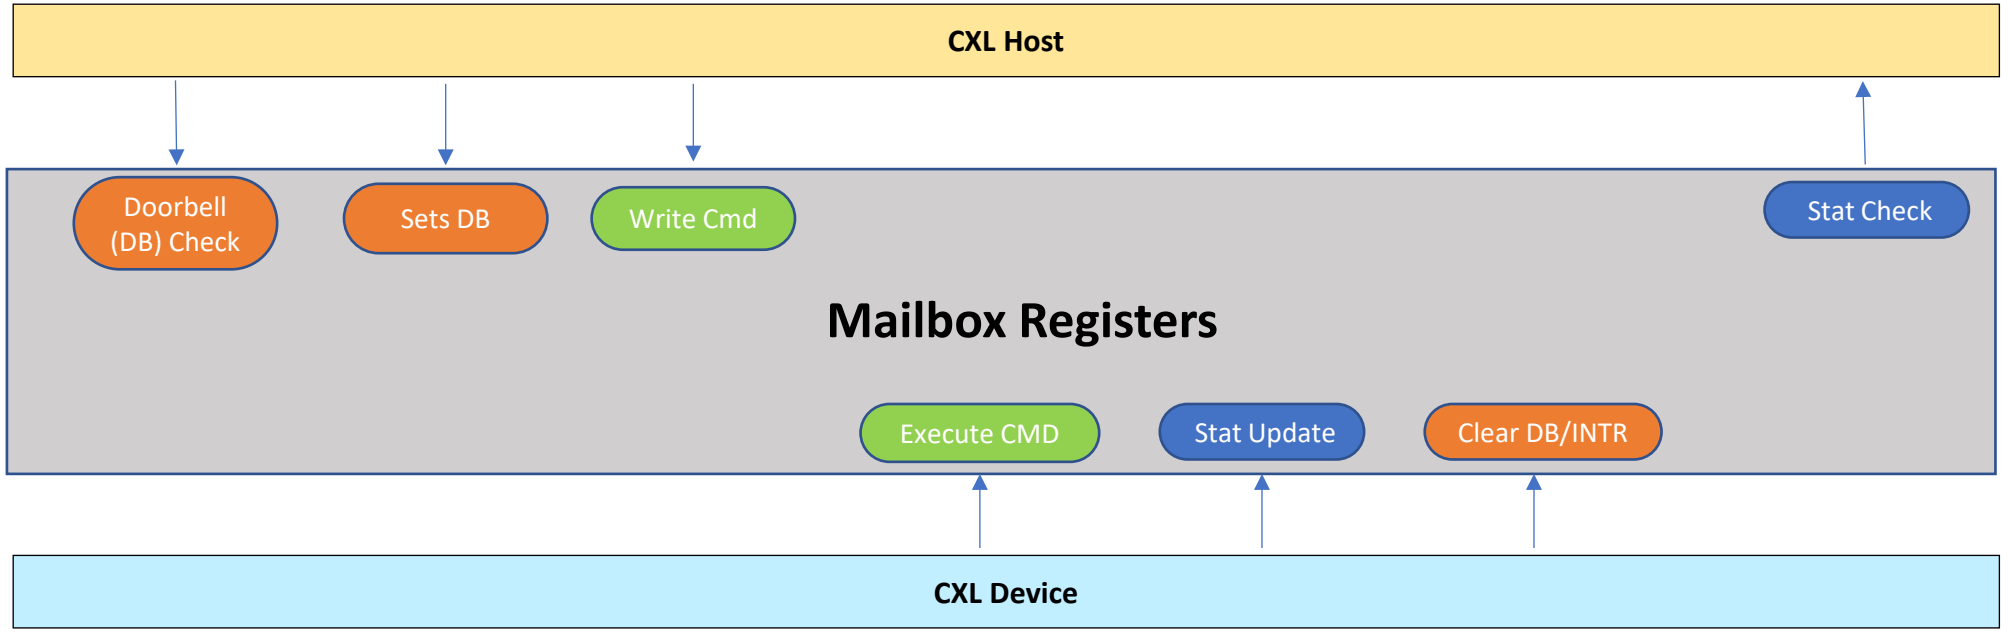

©2024 Conference Concepts, Inc.
All Rights Reserved

CXL Out-Of-Band Management Framework

- Open Community Defined Standards (DMTF/PCIe-SIG)
- Baseboard Management Controller (BMC) – Platform management independent of OS
- CXL Component Management
  - Type 07h - FM API commands
  - Type 08h - General and Memory Device CCI

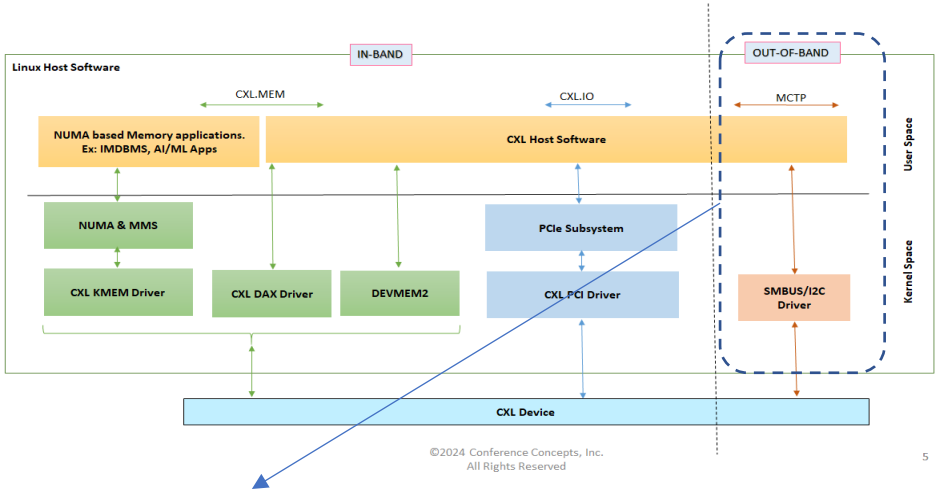

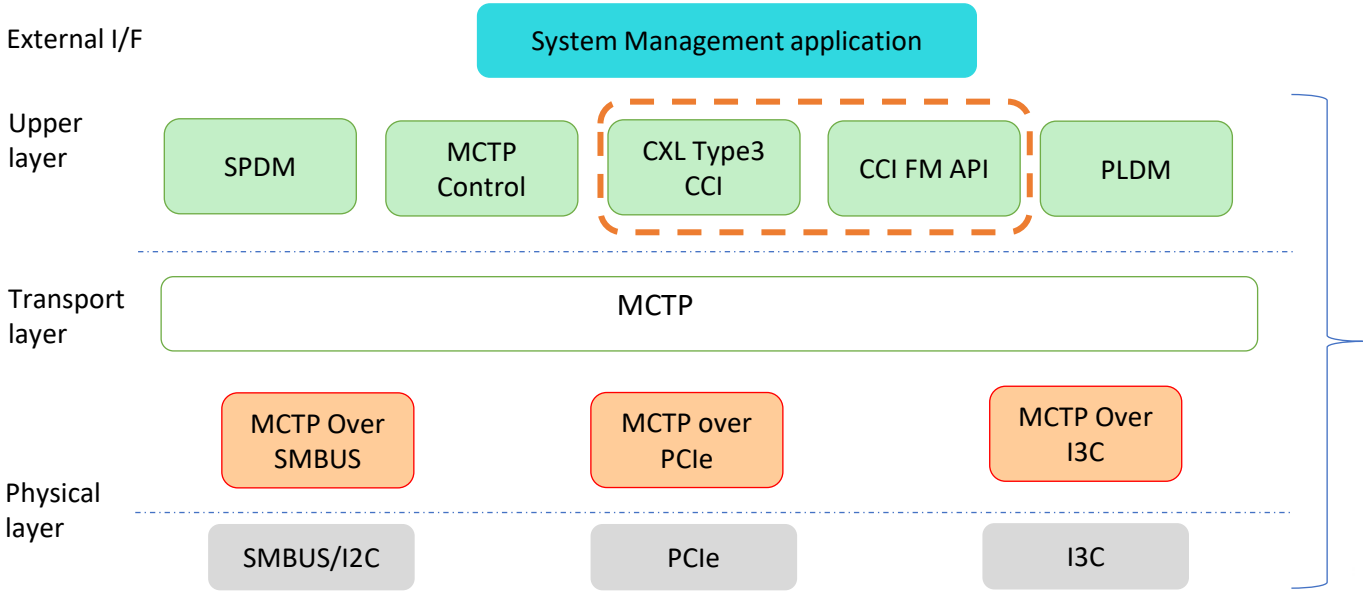

MCTP (Management Component Transport Protocol)
DMTF(Distributed Management Task Force)

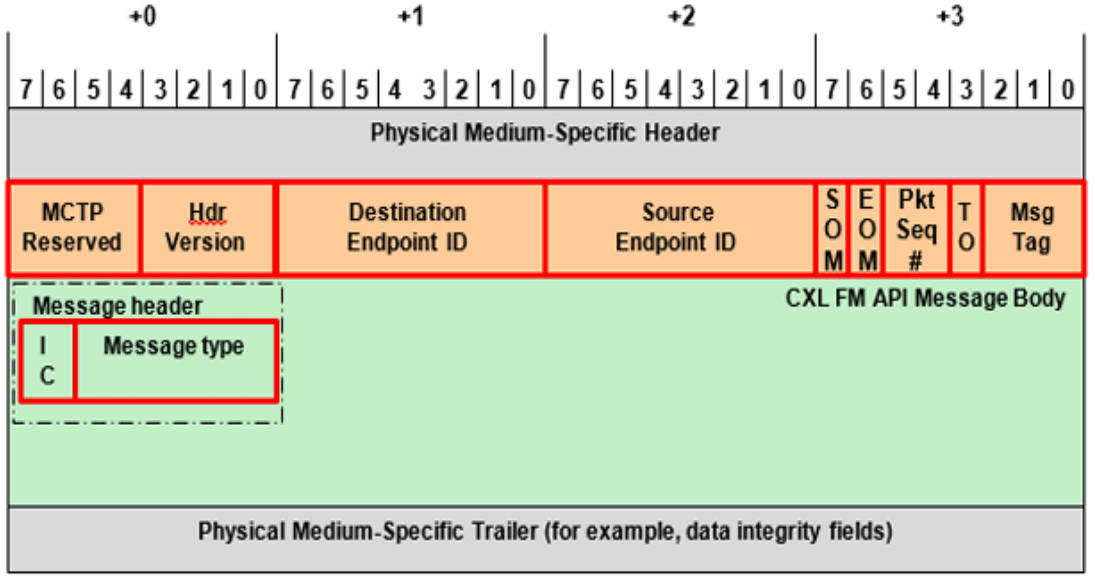

FM API (Fabric Manager Application Programming Interface)
CCI(Command component Interface)

©2024 Conference Concepts, Inc.
All Rights Reserved

CXL FM APIs use in Pooled Memory Management

- Fabric Manager (FM) a conceptual term
- Creation of MLD and assigning to Hosts

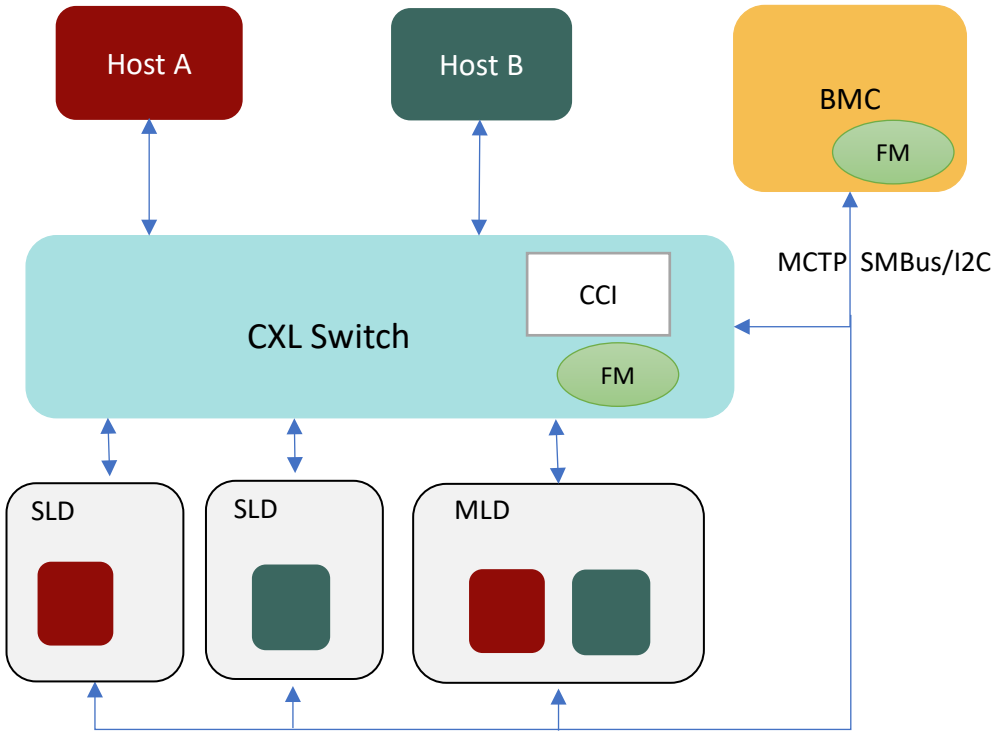

FM API Command Sets

- Switch Event Notification [Event Notification]
- Physical Switch [Identify, Get Port]
- Virtual Switch [Bind vPPB, Unbind vPPB]
- MLD Port [Bind vPPB, Unbind vPPB]
- MLD Components [Get/Set LDA, Get/Set QOS Status]

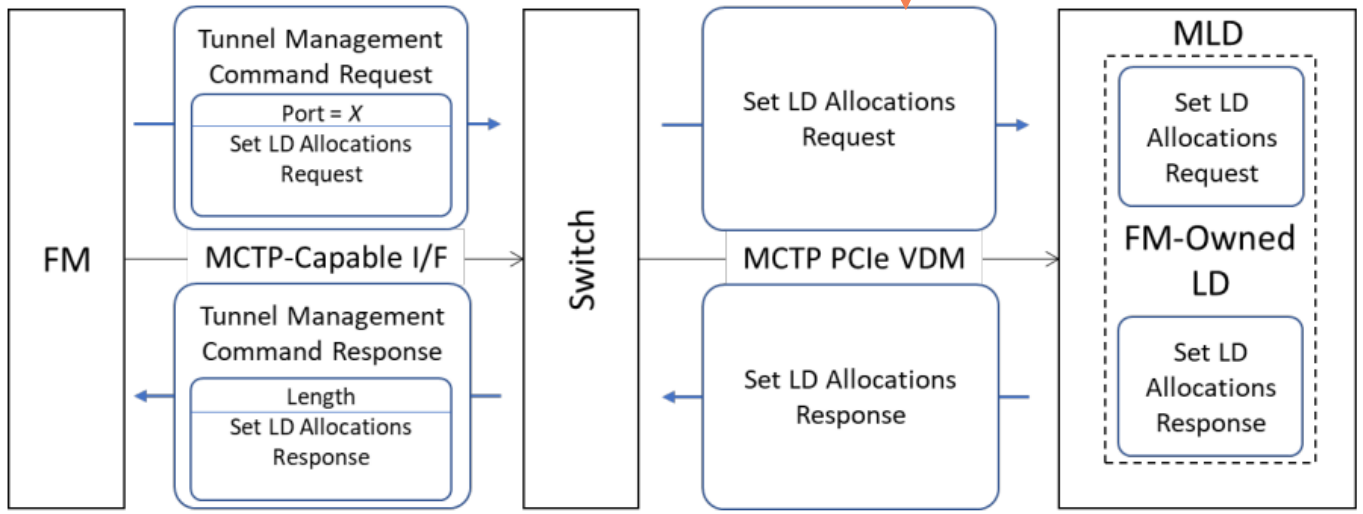

Credits: CXL Consortium

SLD – Single Logical Device
MLD – Multiple Logical Device
vPPB – Virtual Physical Port Binding

©2024 Conference Concepts, Inc.
All Rights Reserved

CXL Command processing via Out-of-Band(OOB) path

- Host to Mem Binding and MLD Configuration

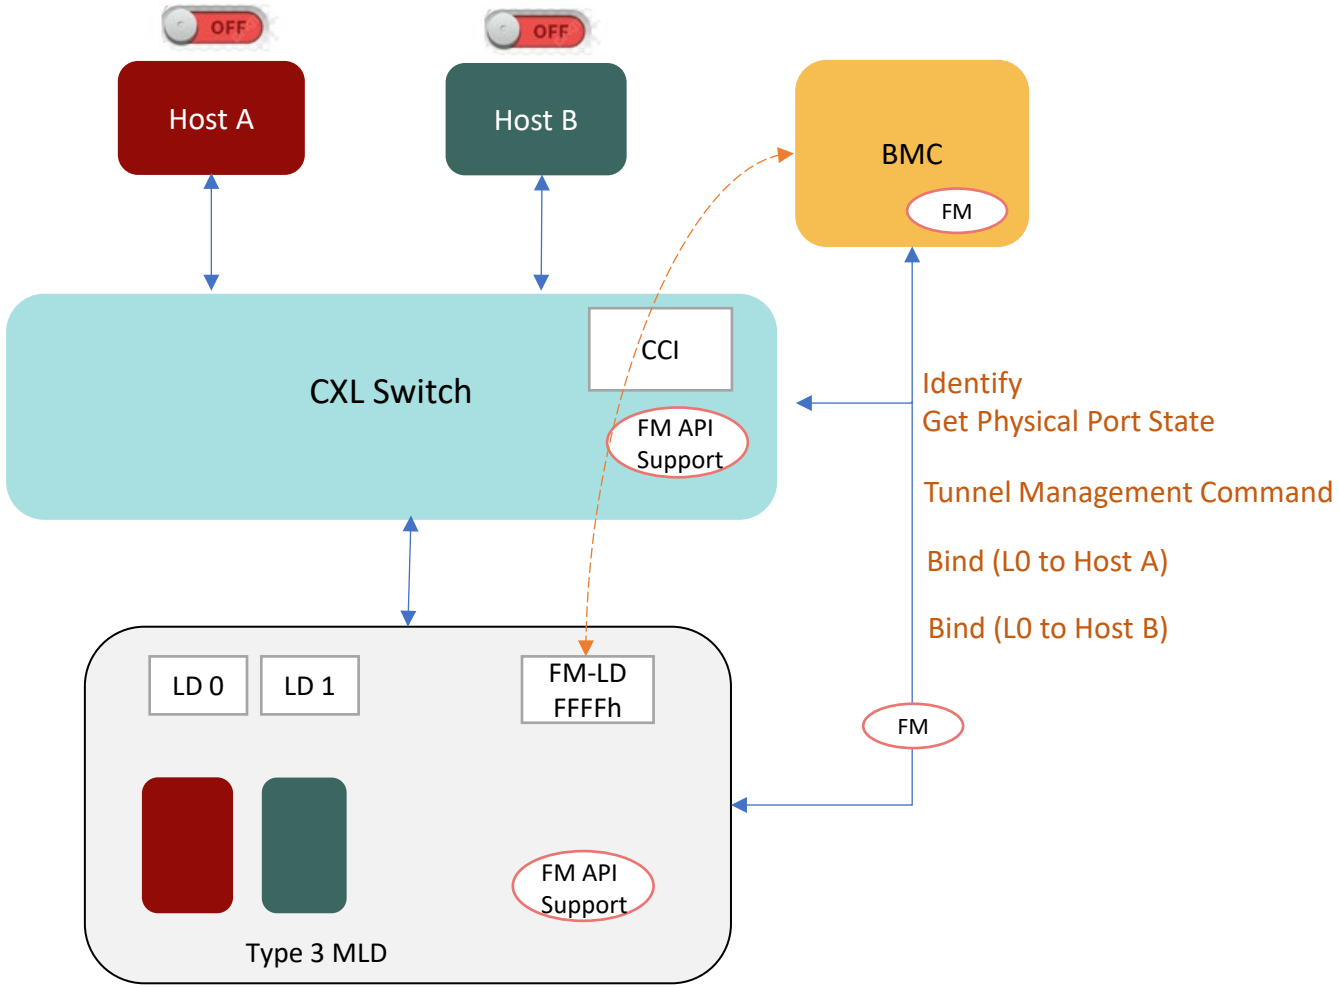

- Concurrent In-Band and Out-of-Band Command Flow

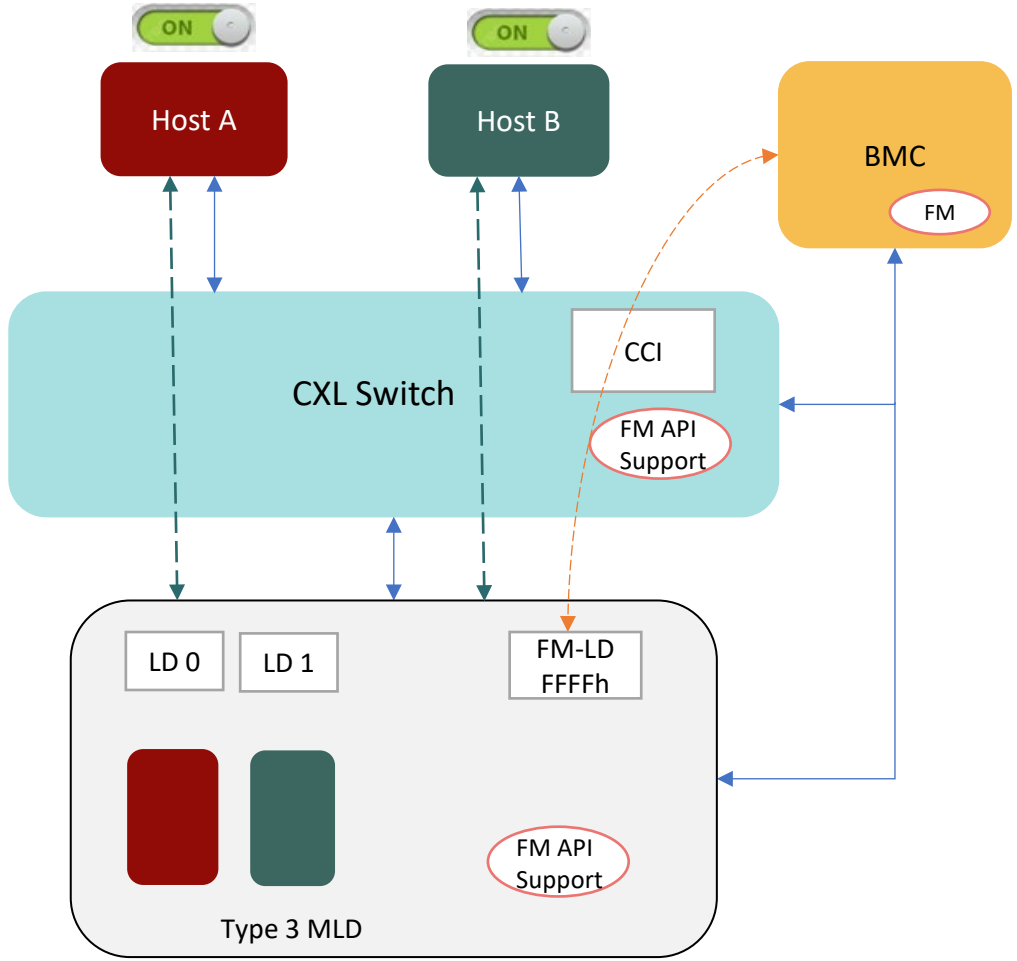

LD – Logical Device

©2024 Conference Concepts, Inc.
All Rights Reserved

CXL Summary

- CXL Performance Optimizations (Low Latency, High BW, High Capacity, Low TCO)
- Device access to Host (CXL.MEM, CXL.IO, In-Band and Out-of-Band Path)
- In-Band Processing via Mailbox Registers
- Pooled Memory Management via Out-of-Band (MCTP, Fabric Manager API)
- Out-Of-Band Processing of CMD (Binding, MLD Configuration)

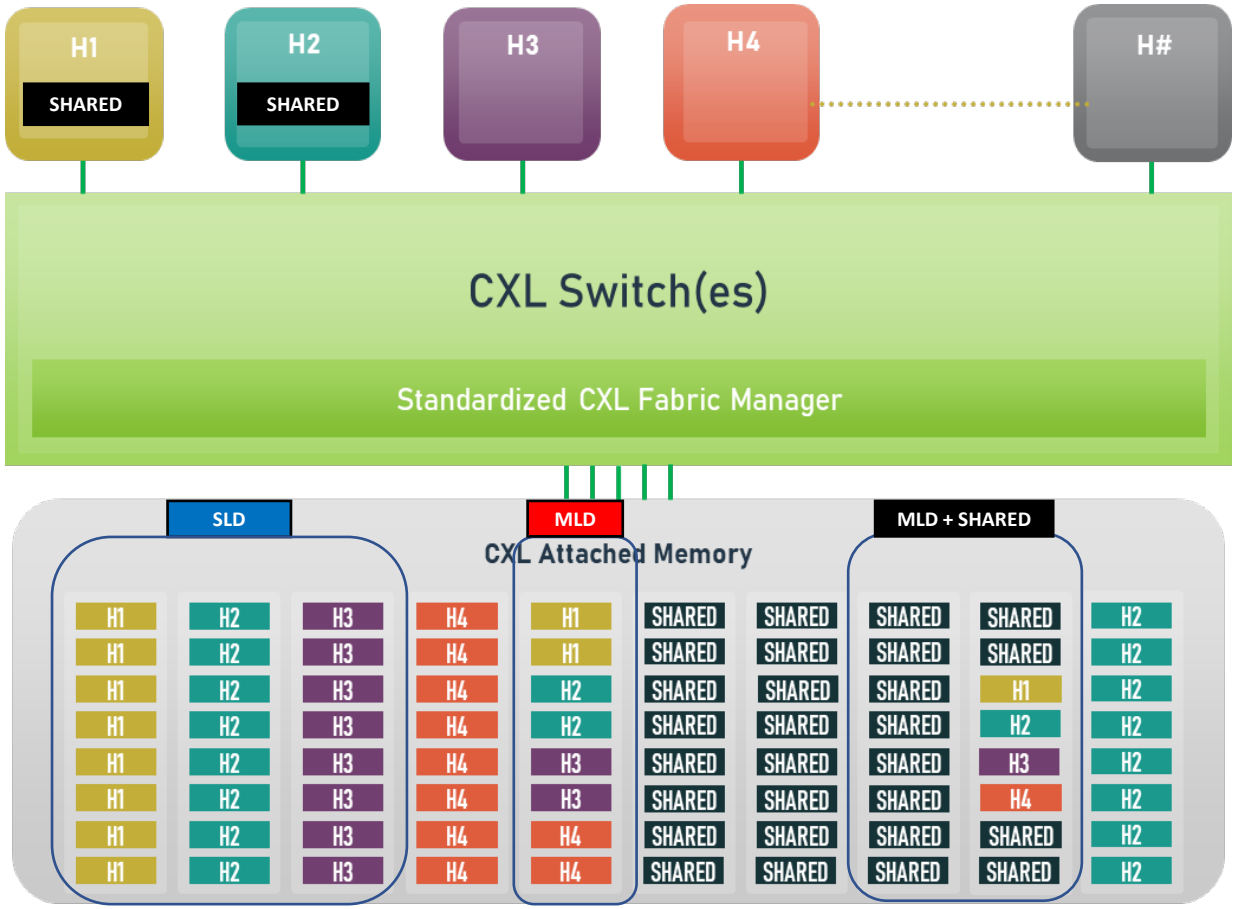

Credits: CXL Consortium

CXL 3.1 – Scalable, Composable & Shareable Memory Architecture

©2024 Conference Concepts, Inc.
All Rights Reserved

Questions & Answers

©2024 Conference Concepts, Inc.
All Rights Reserved

Thank You

©2024 Conference Concepts, Inc.
All Rights Reserved