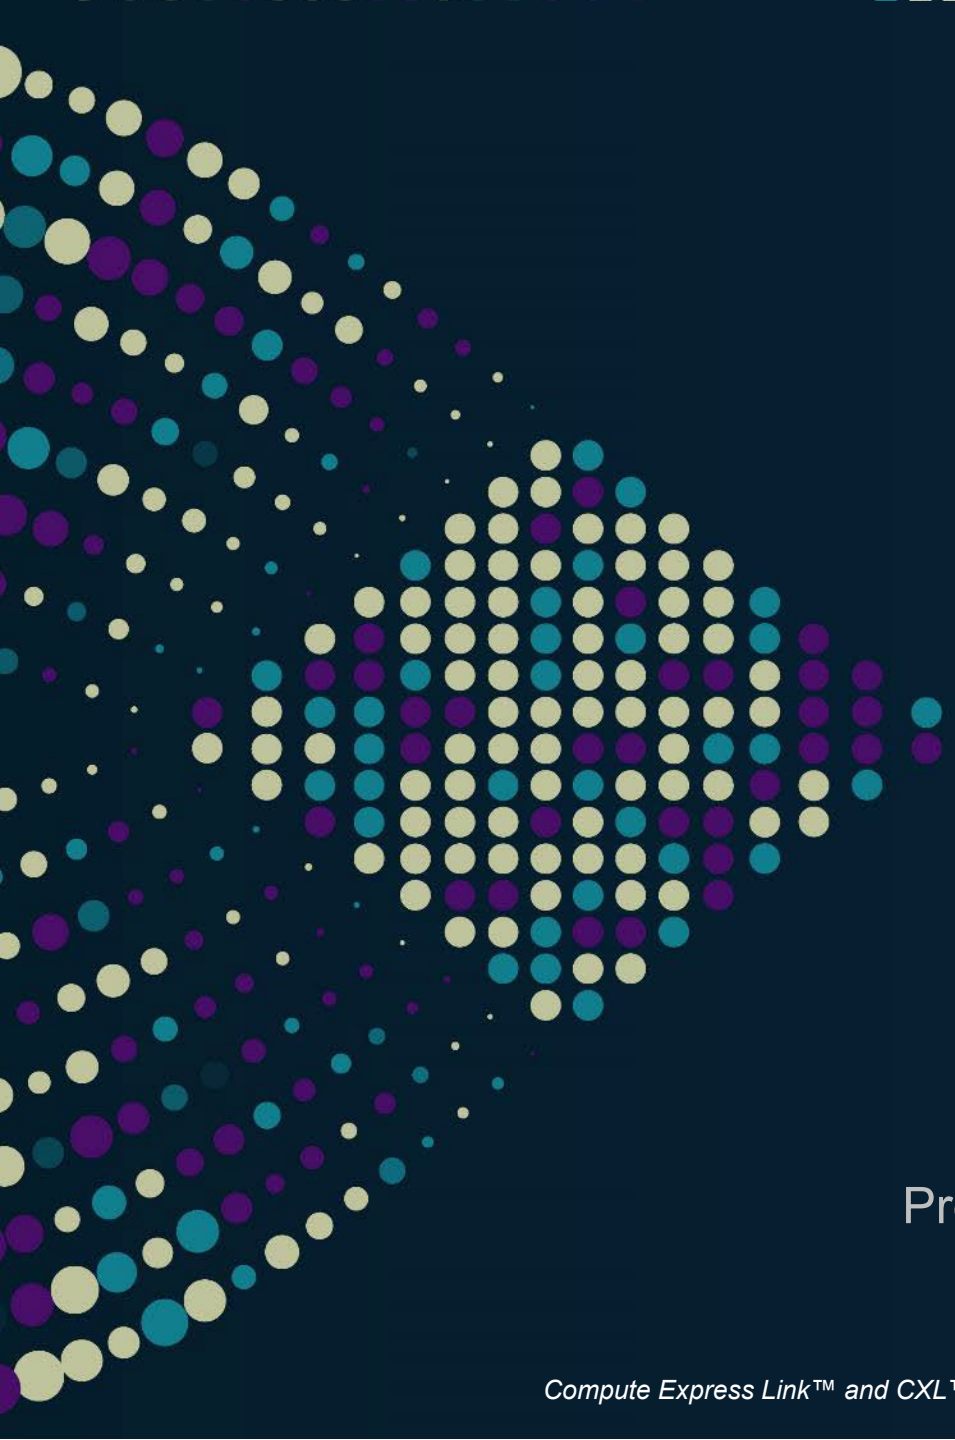

STORAGE DEVELOPER CONFERENCE

SDC 23

BY Developers FOR Developers

Advancing Coherent Connectivity With CXL

Moderator: Richelle Alvers (Intel)

Presented by: Sandeep Dattaprasad (Astera Labs), Steve Scargall (MemVerge), Gerry Fan (XConn Technologies)

About the Moderator

Richelle Ahlvers

Storage Technology
Enablement Architect, Intel

Richelle is a Storage Technology Enablement Architect at Intel, where she promotes and drives enablement of new technologies and standards strategies. Richelle has spent over 25 years in Enterprise R&D teams in a variety of technical roles, leading the architecture, design and development of storage array software, storage management software user experience projects including mobility, developing new storage industry categories including SAN management, storage grid and cloud, and storage technology portfolio solutions.

Richelle has been engaged with industry standards initiatives for many years and is actively engaged with many groups supporting manageability including SNIA, DMTF, NVMe, OFA and UCIe. She is Vice-Chair of the SNIA Board of Directors, Chair of the Storage Management Initiative, leads the SSM Technical Work Group developing the Swordfish Scalable Storage Management API, and has also served as the SNIA Technical Council Chair and been engaged across a breadth of technologies ranging from storage management, to solid state storage, to cloud, to green storage. She also serves on the DMTF Board of Directors as the VP of Finance and Treasurer.

About the Panelists

Sandeep Dattaprasad

Senior Product Manager,
Astera Labs

Sandeep Dattaprasad is a Senior Product Manager at Astera Labs with 15+ years of experience in semiconductor, software diagnostic tools, developing security strategies and firmware development for complex SoC product lines including Compute Express Link™ products, SAS RAID controllers, SAS expanders and PCIe® switches. He is also a contributing member of the CXL Consortium.

At Astera Labs, Sandeep focuses on driving product strategy for new market segments by translating data center bottlenecks into profitable and competitive hardware and software product roadmaps using CXL technology.

About the Panelists

Gerry Fan

Founder, Xconn Technologies

Gerry Fan is the founder of XConn Technologies, world leader of CXL/PCIe switch silicon manufacturer. He has spent 30 years of his career in ASIC product development. Prior to founding XConn, Gerry held both engineering R&D and management positions in Broadcom, Marvell, Cisco and a few successful startups. He went to Boston University for MS in EE.

About the Panelists

Steve Scargall

Senior Product Manager and Software Architect,
MemVerge

Steve Scargall, Senior Product Manager and Software Architect, MemVerge, delivering software-defined memory solutions using Compute Express Link (CXL) devices. Steve works with industry leaders in the CXL hardware vendor, Original Equipment Manufacturers (OEMs), Cloud Service Providers (CSPs), Enterprise, and System Integrator spaces to architect cutting-edge solutions.

Steve holds a bachelor's degree in BSc Applied Computer Science and Cybernetics from the University of Reading, UK. He has made significant contributions to the SNIA NVM Programming Technical Work Group, PMDK, NDCTL, IPMCTL, and other memory-centric open-source projects. Steve is the author of "Programming Persistent Memory: A Comprehensive Guide for Developers".

CXL Overview

Challenges

Industry trends driving demand for faster data processing and next-gen data center performance

Increasing demand for heterogeneous computing and server disaggregation

Need for increased memory capacity and bandwidth

Lack of open industry standard to address next-gen interconnect challenges

CXL

An open industry-supported cache-coherent interconnect for processors, memory expansion and accelerators

Coherent Interface

Leverages PCIe with 3 mix-and-match protocols

Low Latency

.Cache and .Memory targeted at near CPU cache coherent latency

Asymmetric Complexity

Eases burdens of cache coherent interface designs

CXL Protocol Overview

- Builds upon PCI Express® physical and electrical interface
  - Existing PCIe devices such as storage products can leverage CXL.io
- CXL transaction layer has three dynamically multiplexed sub-protocols

CXL.io

PCIe Discovery, configuration, interrupts, etc

- DMA
- Interrupts (MSI/MSIX)
- SR-IOV, ACS, ATS etc.
- NVMe

CXL.cache

- Device access to processor memory

CXL.mem

- Processor access to device attached memory

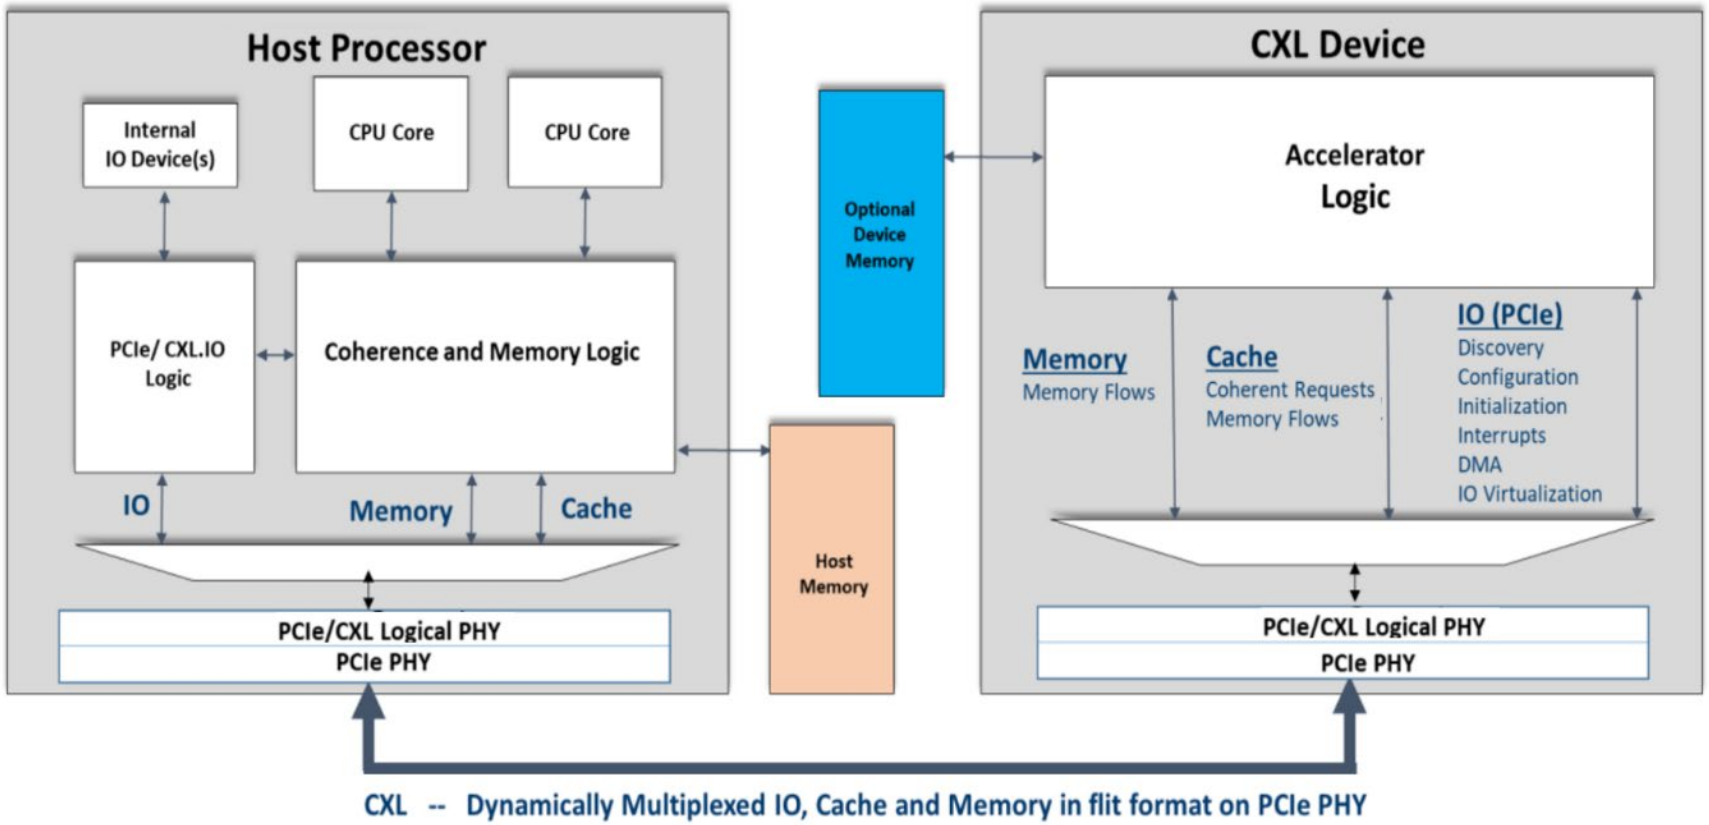

Representative CXL Usages

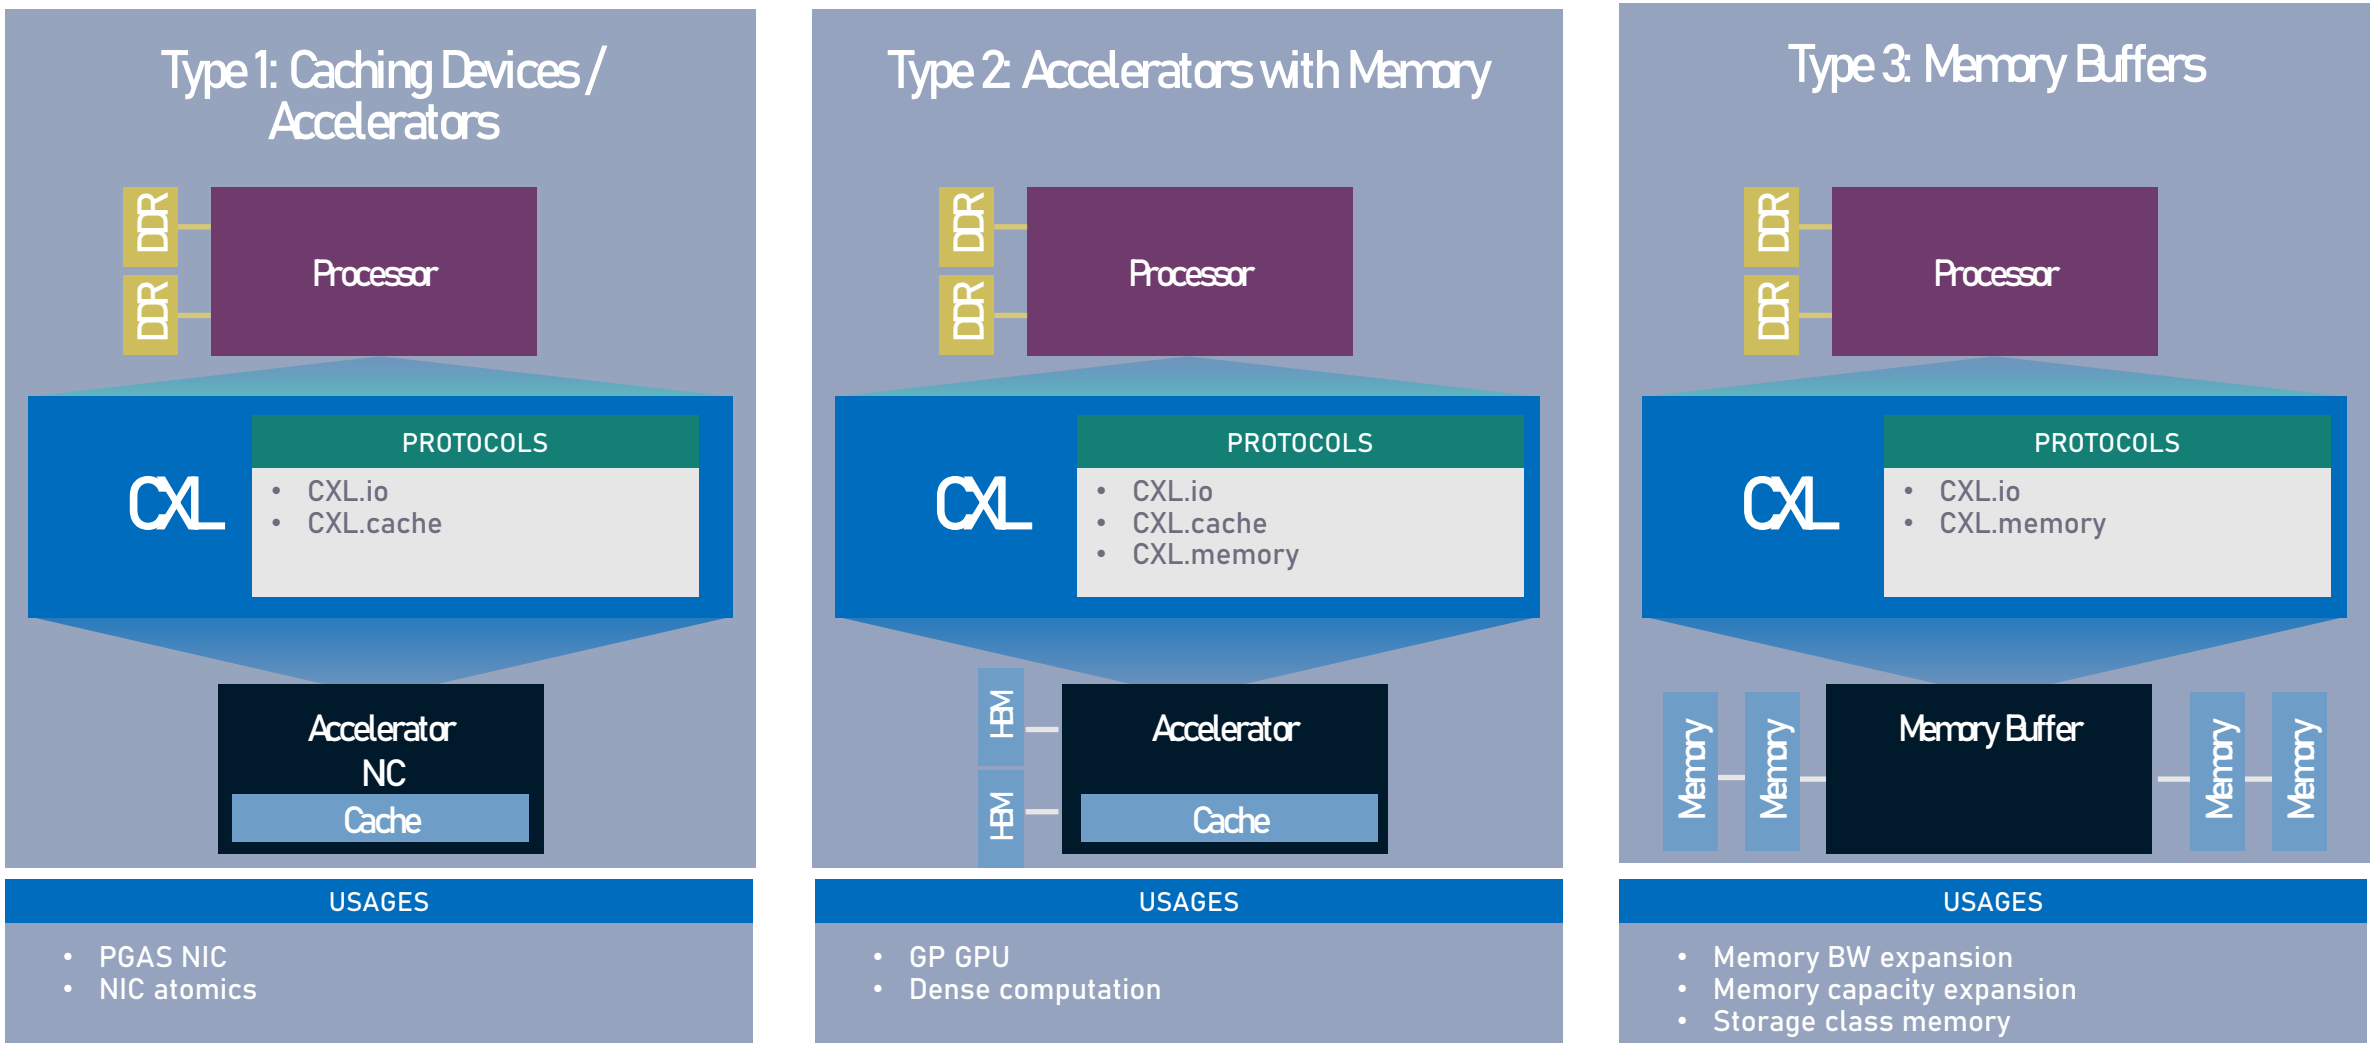

CXL Specification and Features

Features	CXL 1.0 / 1.1	CXL 2.0	CXL 3.0
Release date	2019	2020	1H 2022
Max link rate	32GTs	32GTs	64GTs
Flit 68 byte (up to 32 GTs)	✓	✓	✓
Flit 256 byte (up to 64 GTs)			✓
Type 1, Type 2 and Type 3 Devices	✓	✓	✓
Memory Pooling w/ MLDs		✓	✓
Global Persistent Flush		✓	✓
CXL IDE		✓	✓
Switching (Single-level)		✓	✓
Switching (Multi-level)			✓
Direct memory access for peer-to-peer			✓
Symmetric coherency (256 byte flit)			✓
Memory sharing (256 byte flit)			✓
Multiple Type 1/Type 2 devices per root port			✓
Fabrics (256 byte flit)			✓

	Not supported
✓	Supported

SDC 23

CXL 1.1/2.0/3.0 Memory Expansion

Direct-Attached Use Cases

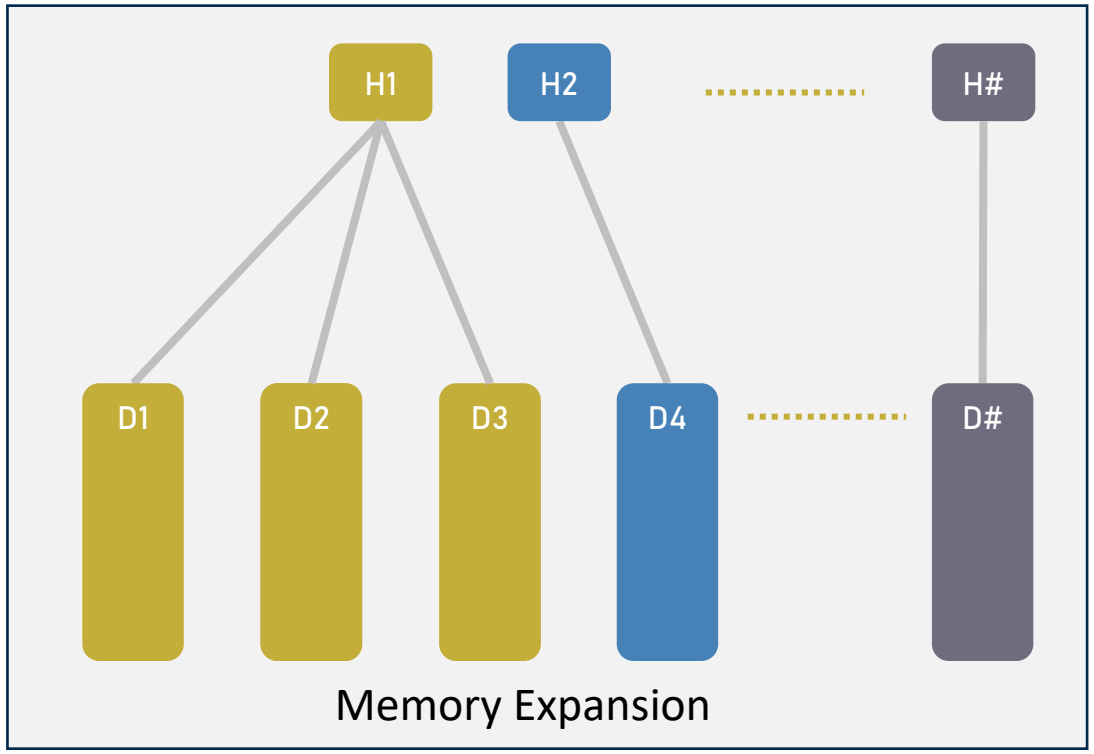

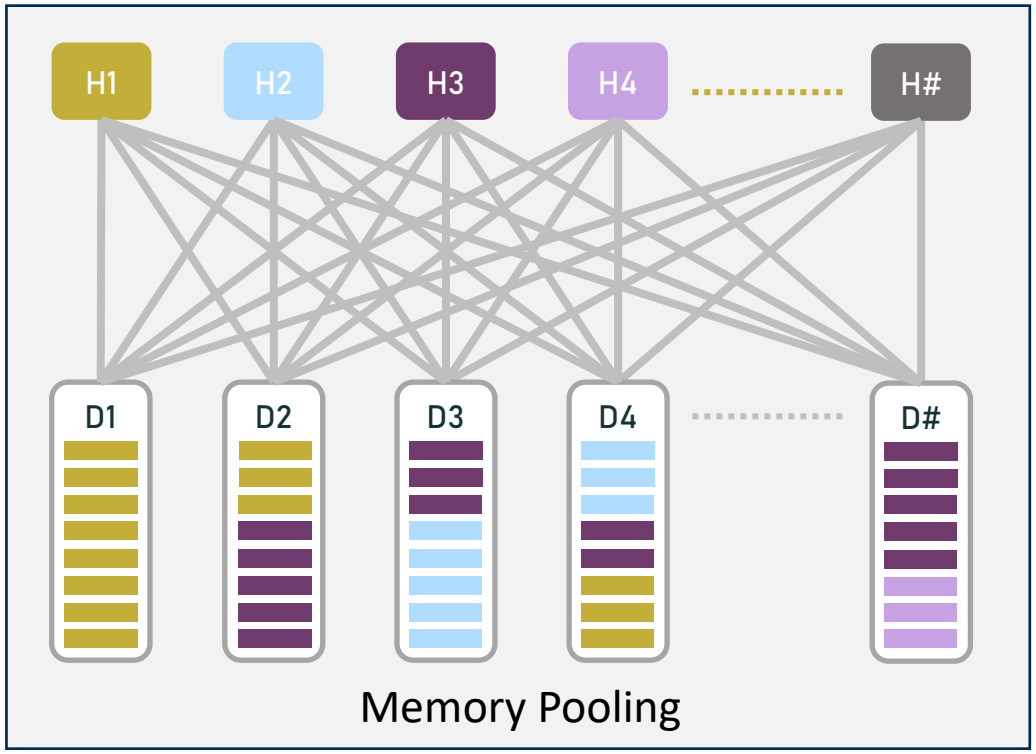

BENEFITS

- Provides lower latency
- Lower power
- Lower cost

CXL 2.0 for Heterogeneous Systems

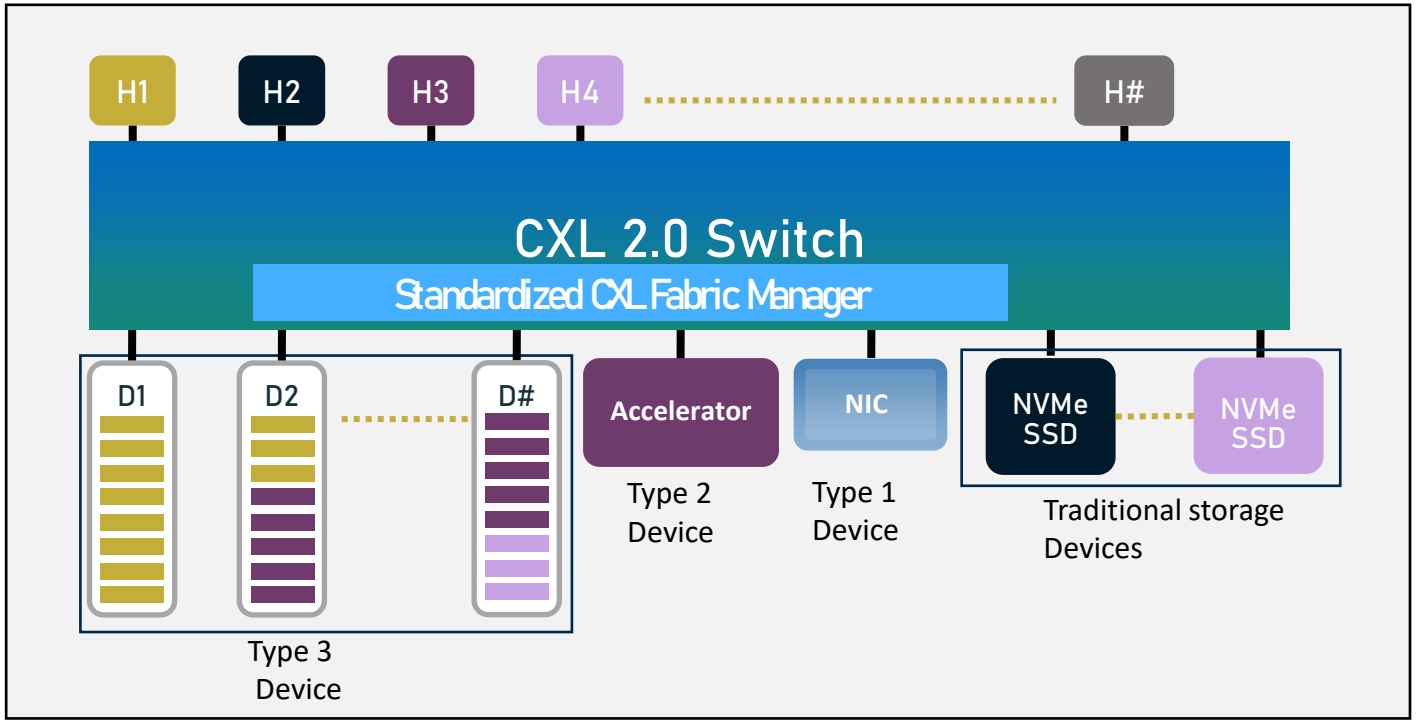

BENEFITS

- CXL 2.0 capabilities supports heterogenous systems
  - Supports systems with Type1, Type2 & Type3 CXL devices
  - CXL.io semantics supports NVMe storage devices facilitating Tiered memory
  - Fabric manager used for configuration & management of composable memory and storage applications
  - Cache coherency improves compute efficiency at higher bandwidth and lower latency
  - Supports SR-IOV MR-IOV use cases on traditional PCIe devices

CXL Switch Attached SSD

Optimizing SSD Performance

Memory/Storage Hierarchy

Fast
Expensive
Low latency
Not Persistent

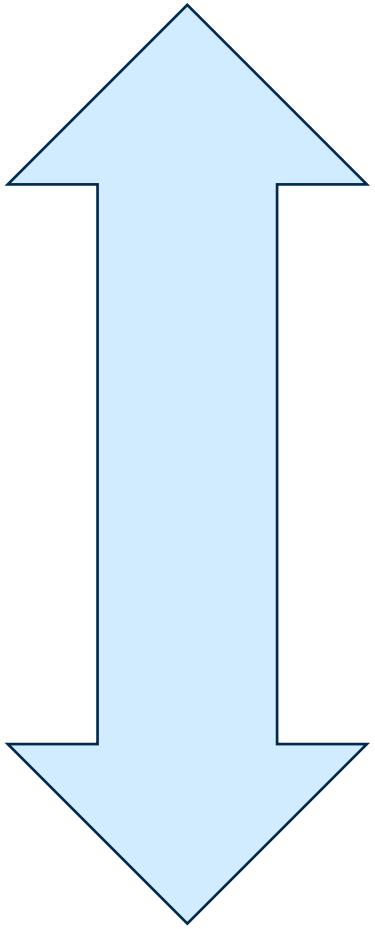

Slow
Cheap
High latency
Persistent

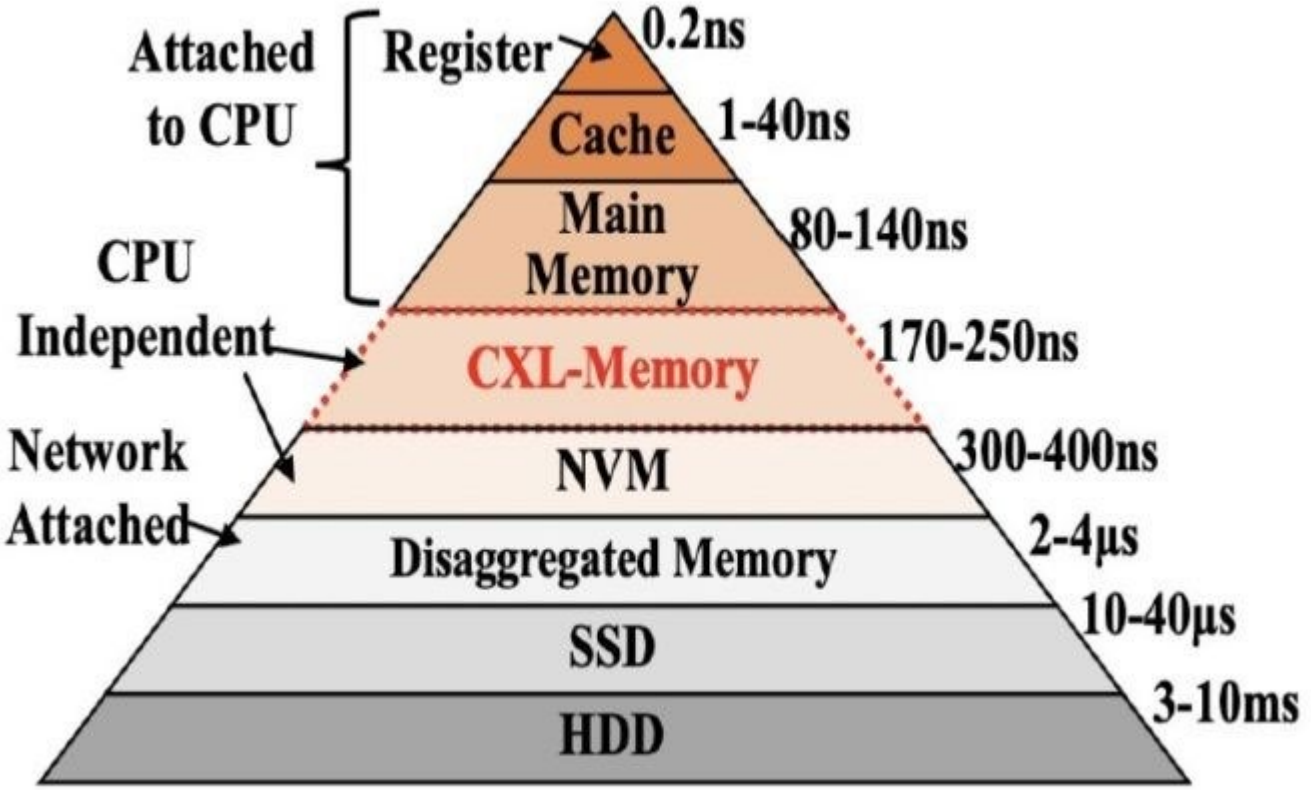

* Hot Chips 34

CXL Switch Enables Tiered SSDs with low latency and high capacity

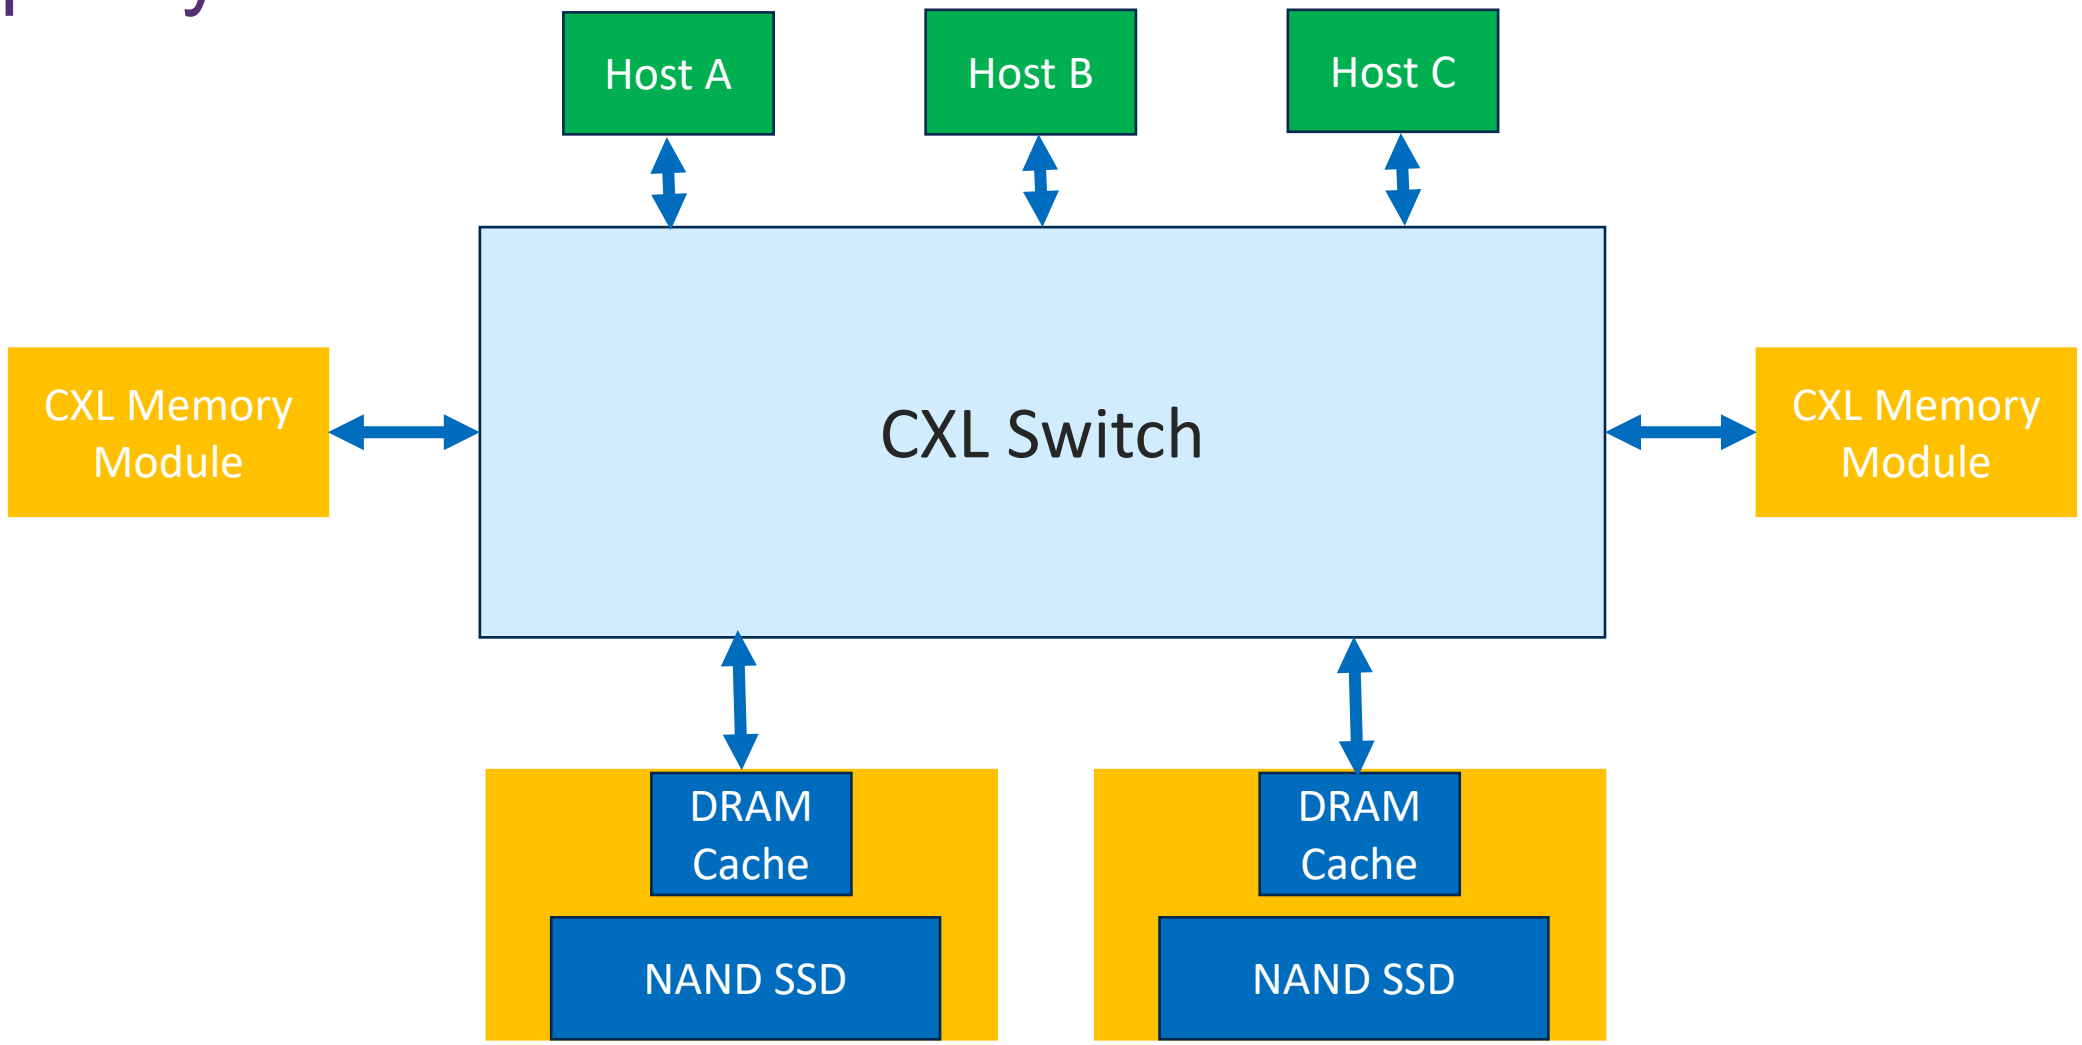

Software and Fabric Manager

What is a Fabric Manager?

- Fabric Manager (FM) is a conceptual term
- FM refers to the application–specific logic for:
  - Composing systems
  - Allocating pooled resources
  - Managing platforms, etc
- Can live anywhere & everywhere
  - Host BMC/IPU/DPU
  - Management software running on a host OS
  - Switch Firmware
  - Endpoint Device Firmware
  - Appliance or JBOM
- Communication can be in-band or out-of-band using a dedicated management Ethernet network.

What is a Fabric Manager?

- The FM framework is flexible by design
- FM building blocks enable a wide variety of deployment use cases
  - Enterprise Data Center, Embedded, Automotive, Hyperscaler, etc..
- Most management capabilities and features are optional to allow flexibility for the environment
- Advanced operations require using the Fabric Management
  - Provisioning logical devices from MLDs & DCDs
  - Configuring switch ports to host, endpoints (JBOM/Appliances), or another switch
  - Security Operations
  - Firmware Updates
  - etc.

CXL Orchestration and Management Made Easy

- REST APIs can be built using the Fabric Manager
  - DMTF RedFish
  - MemVerge Memory Machine
- OS Tools
  - Linux: cxl & daxctl
- Data Center Infrastructure Management (DCIM) Software with CXL Module/Plugin integration

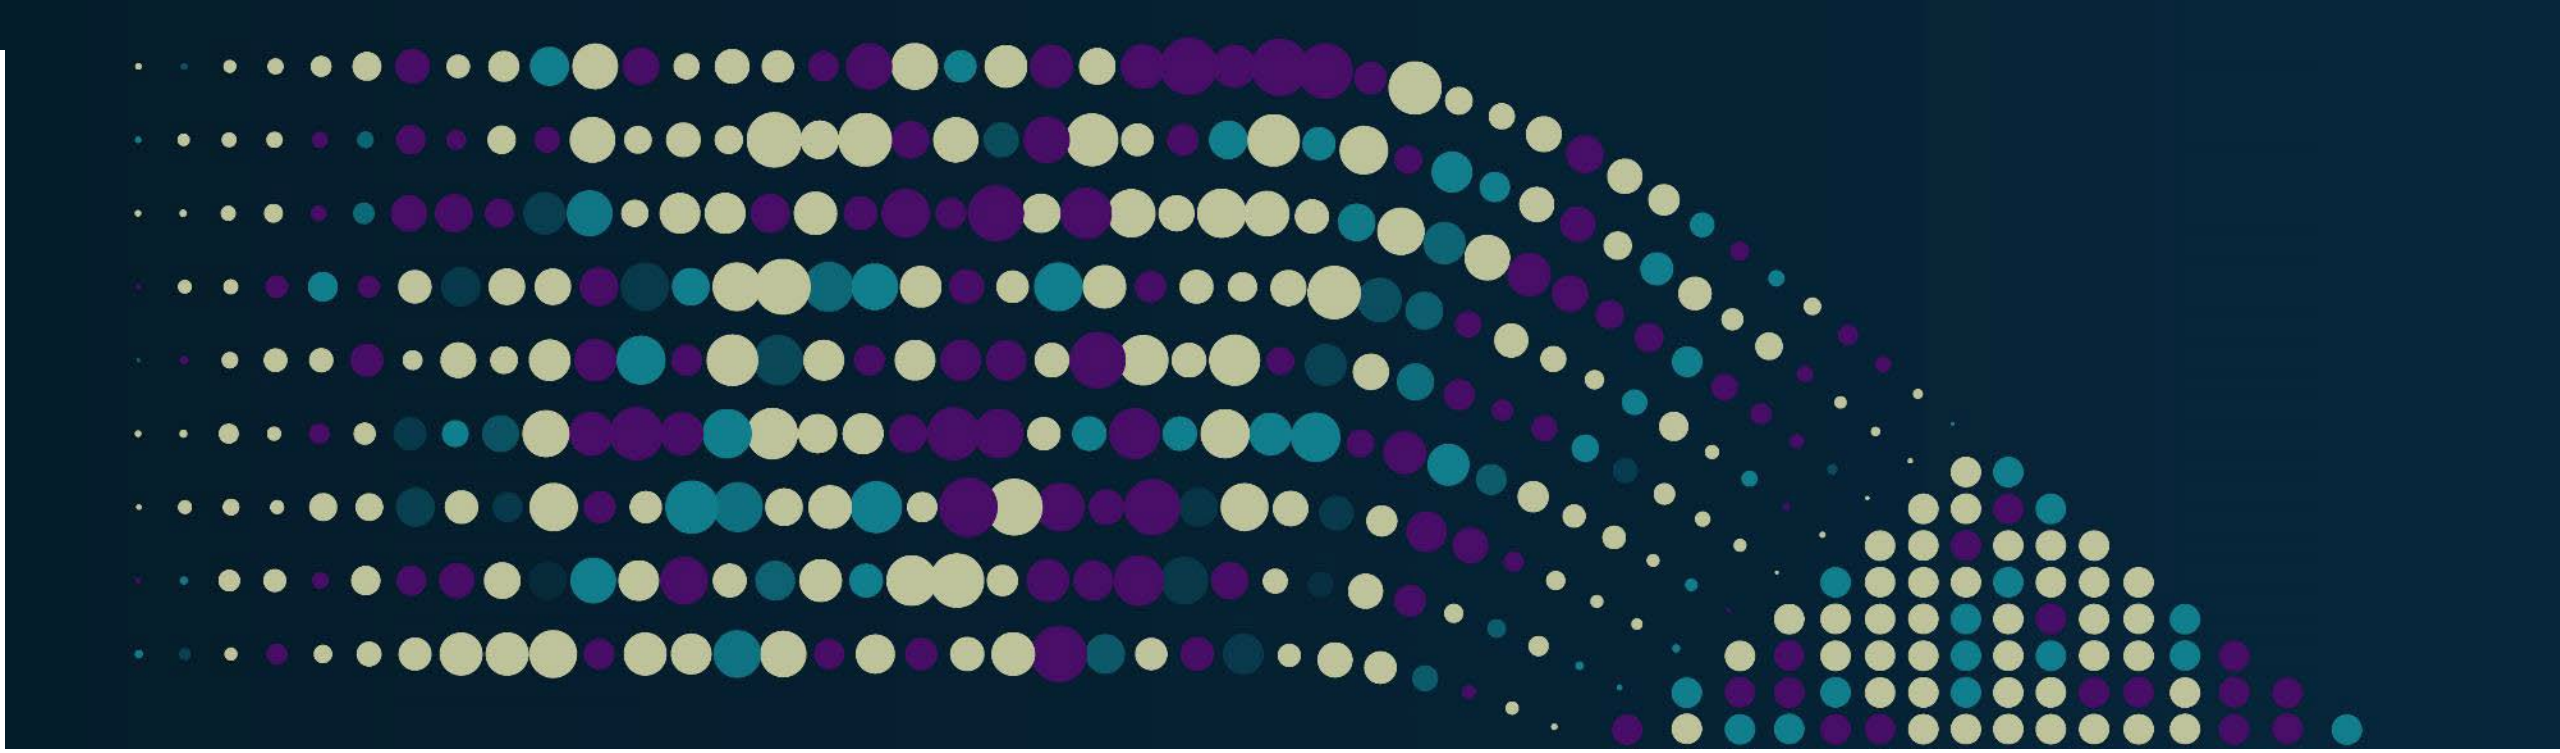

Please take a moment to rate this session.

Your feedback is important to us.

Backup

CXL Linux Architecture

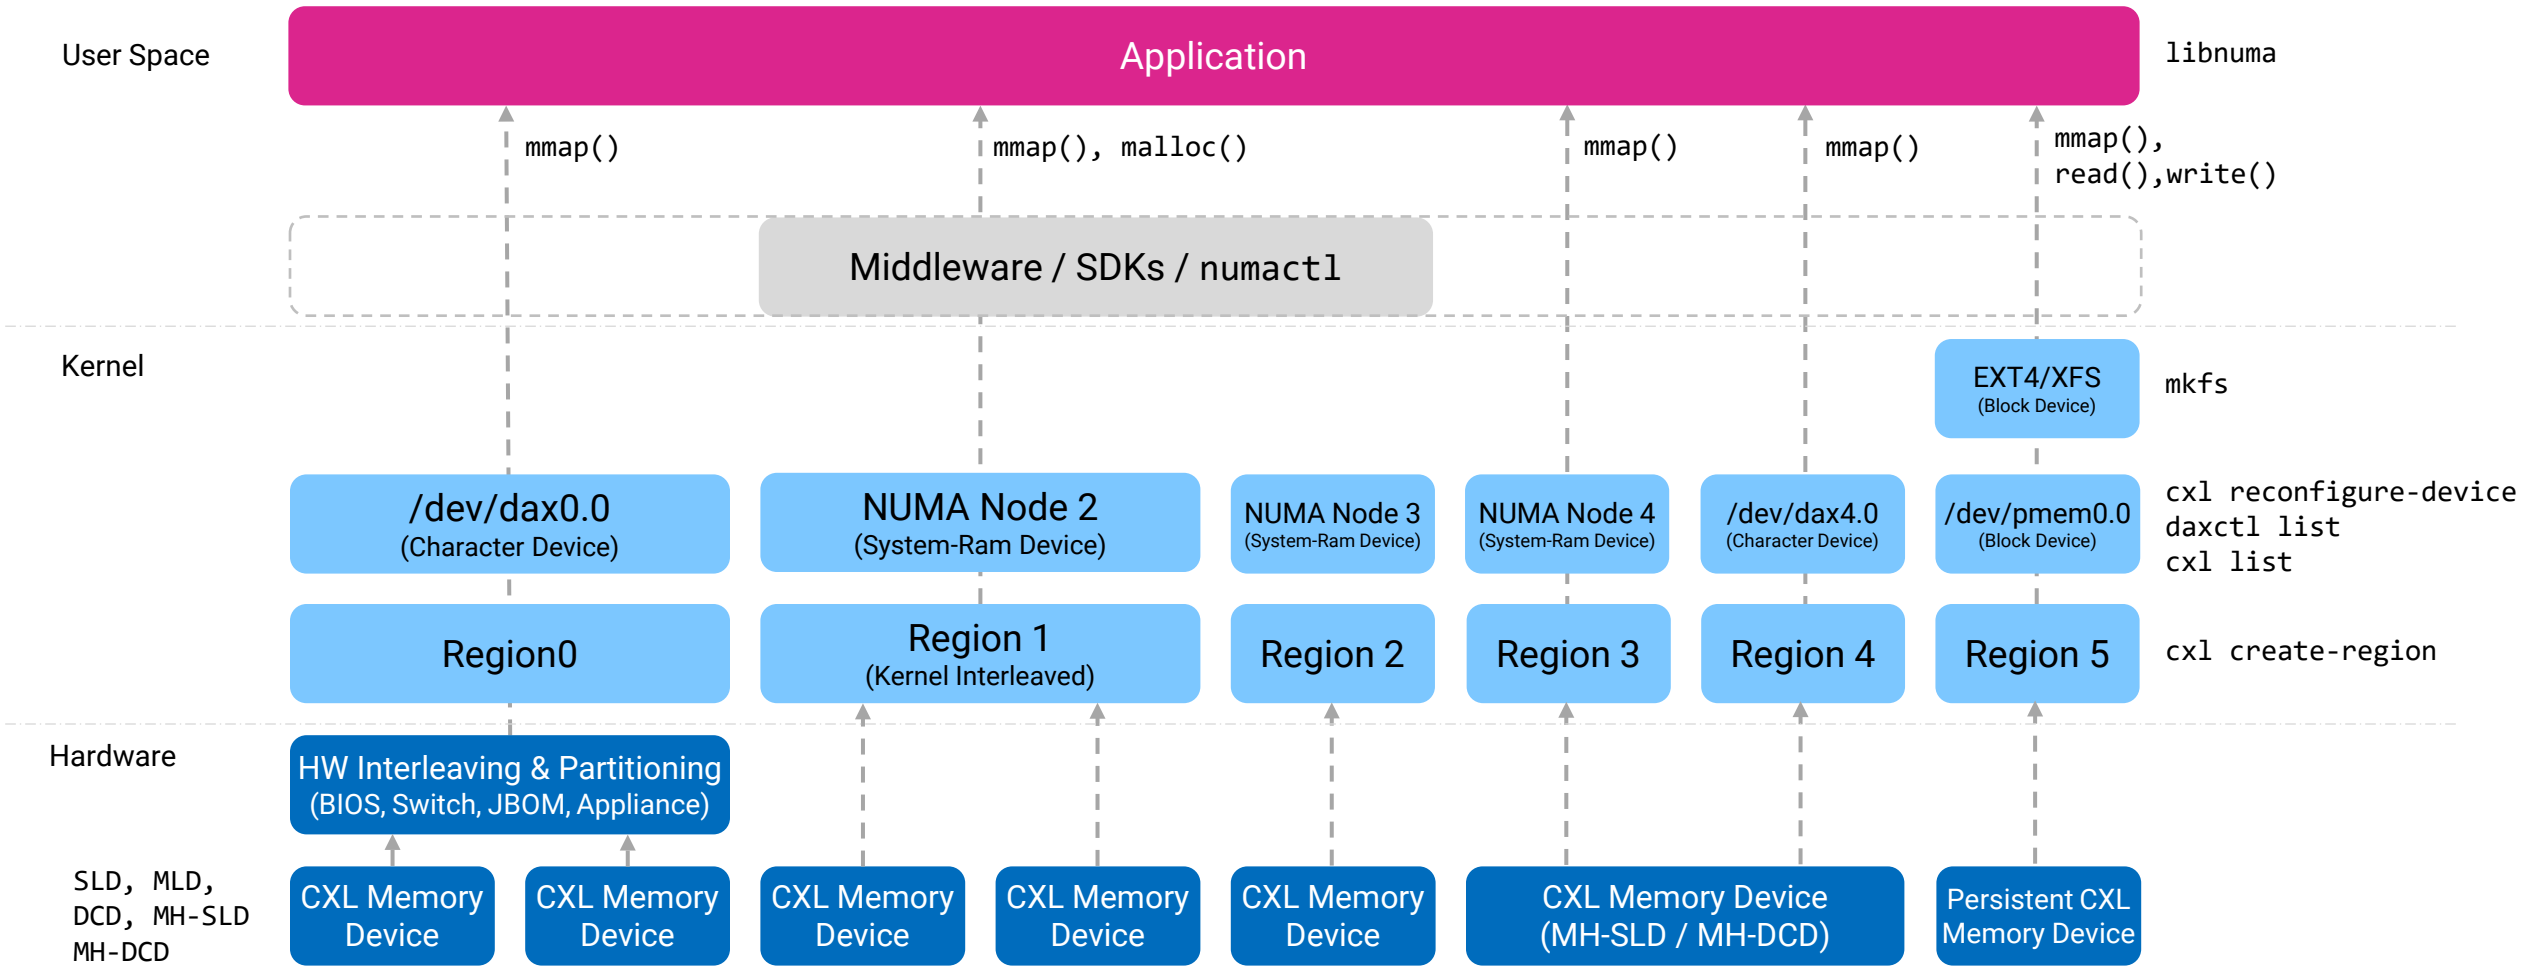

  Compute Express Link™ and CXL™ are trademarks of the Compute Express Link Consortium.