2022 OFA Virtual Workshop

BENEFITS OF COMPUTE EXPRESS LINK™ (CXL™) FOR HIGH-PERFORMANCE COMPUTING

Sandeep Dattaprasad, Microchip

CXL Consortium

INDUSTRY LANDSCAPE

Proliferation of
Cloud Computing

Growth of
AI & Analytics

Cloudification of
the Network & Edge

CXL DELIVERS THE RIGHT FEATURES AND ARCHITECTURE

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

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 CONSORTIUM & GEN-Z CONSORTIUM

- CXL Consortium and Gen-Z Consortium signed the agreement to transfer Gen-Z specification and assets to CXL Consortium
- CXL Consortium will archive Gen-Z specification for five years. Gen-Z specification can be found here: https://www.computeexpresslink.org/projects-3
- Upon completion of the asset transfer, Gen-Z Consortium will finalize operations
- Gen-Z Specifications and all Gen-Z assets have been transferred to the CXL Consortium which will be used to inform the CXL family of specifications towards the development of a memory coherent interface.

CXL & OFA

- OFA had been actively collaborating with Gen-Z consortium
- In light of the recent agreement between CXL and Gen-Z consortiums, these OFA efforts are being redirected to focus on CXL
- OFMF WG have begun work on modelling CXL 2.0 memory sharing and is expected to eventually expand the scope to cover CXL 3.0 fabric
- We expect that a future version of Libfabrics will add support for CXL fabric
- CXL consortium is open to additional collaboration opportunities

HIGH PERFORMANCE COMPUTING (HPC)

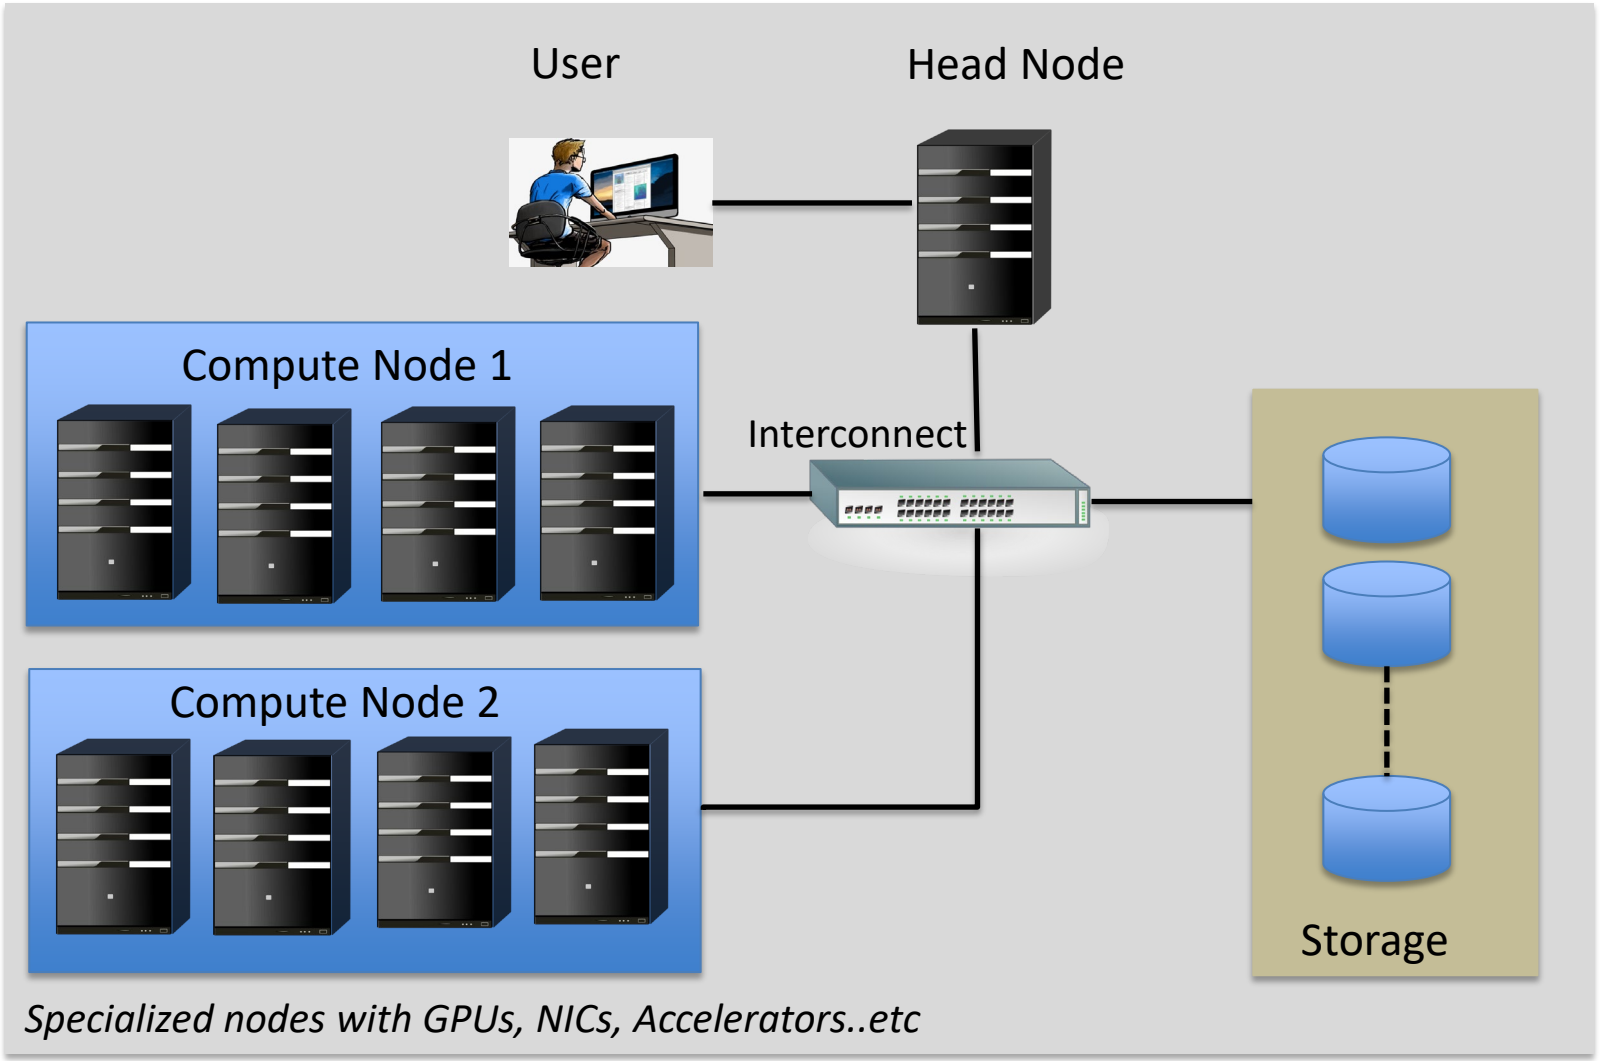

- High Performance Computing (HPC) has transformed industries
  - Aerospace
  - Finance
  - Healthcare
  - Engineering
  - Genomics
  - Autonomous driving
  - Seismic imaging
- Faster data processing and demands for heterogeneous computing drives the need for greater memory capacity and bandwidth

CXL 2.0 USAGE MODELS - RECAP

Caching Devices / Accelerators

Usages:
- PGAS NIC
- NIC atomics

Protocols:
- CXL.io
- CXL.cache

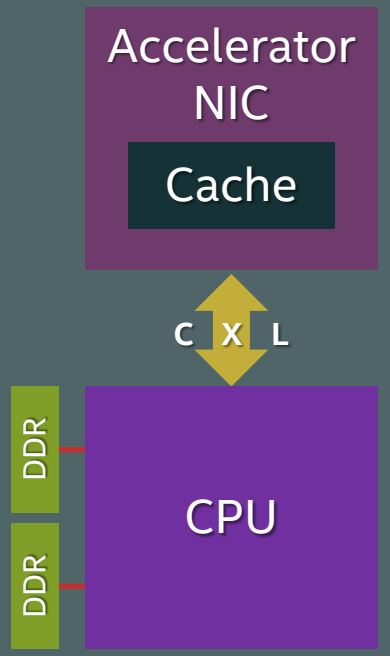

Accelerators with Memory

Usages:
- GPU
- FPGA
- Dense Computation

Protocols:
- CXL.io
- CXL.cache
- CXL.memory

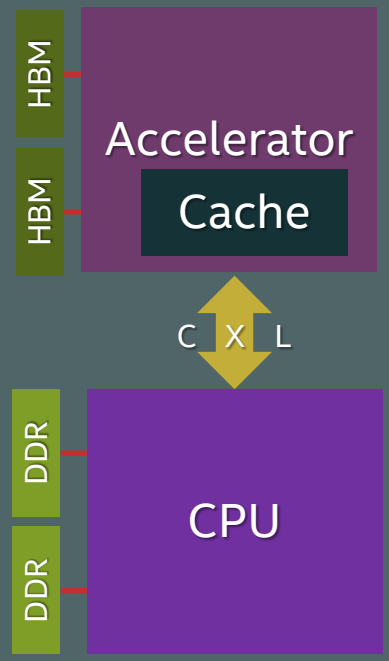

Memory Buffers

Usages:
- Memory BW expansion
- Memory capacity expansion
- Persistent Memory

Protocols:
- CXL.io
- CXL.mem

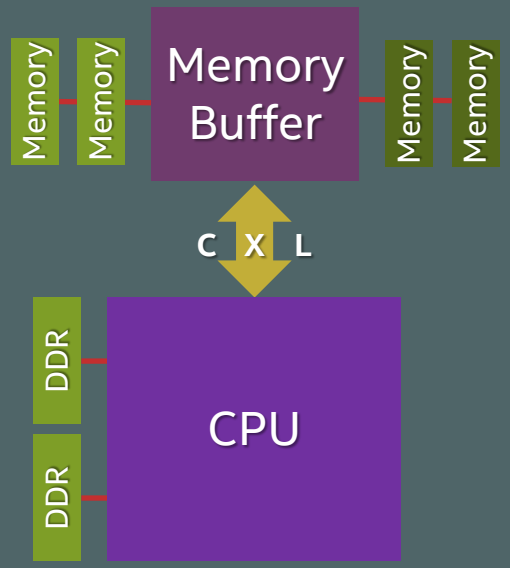

Benefits of CXL 2.0
and Persistent Memory

Moves Persistent Memory from Controller to CXL	Enables Standardized Management of the Memory and Interface	Supports a Wide Variety of Industry Form Factors

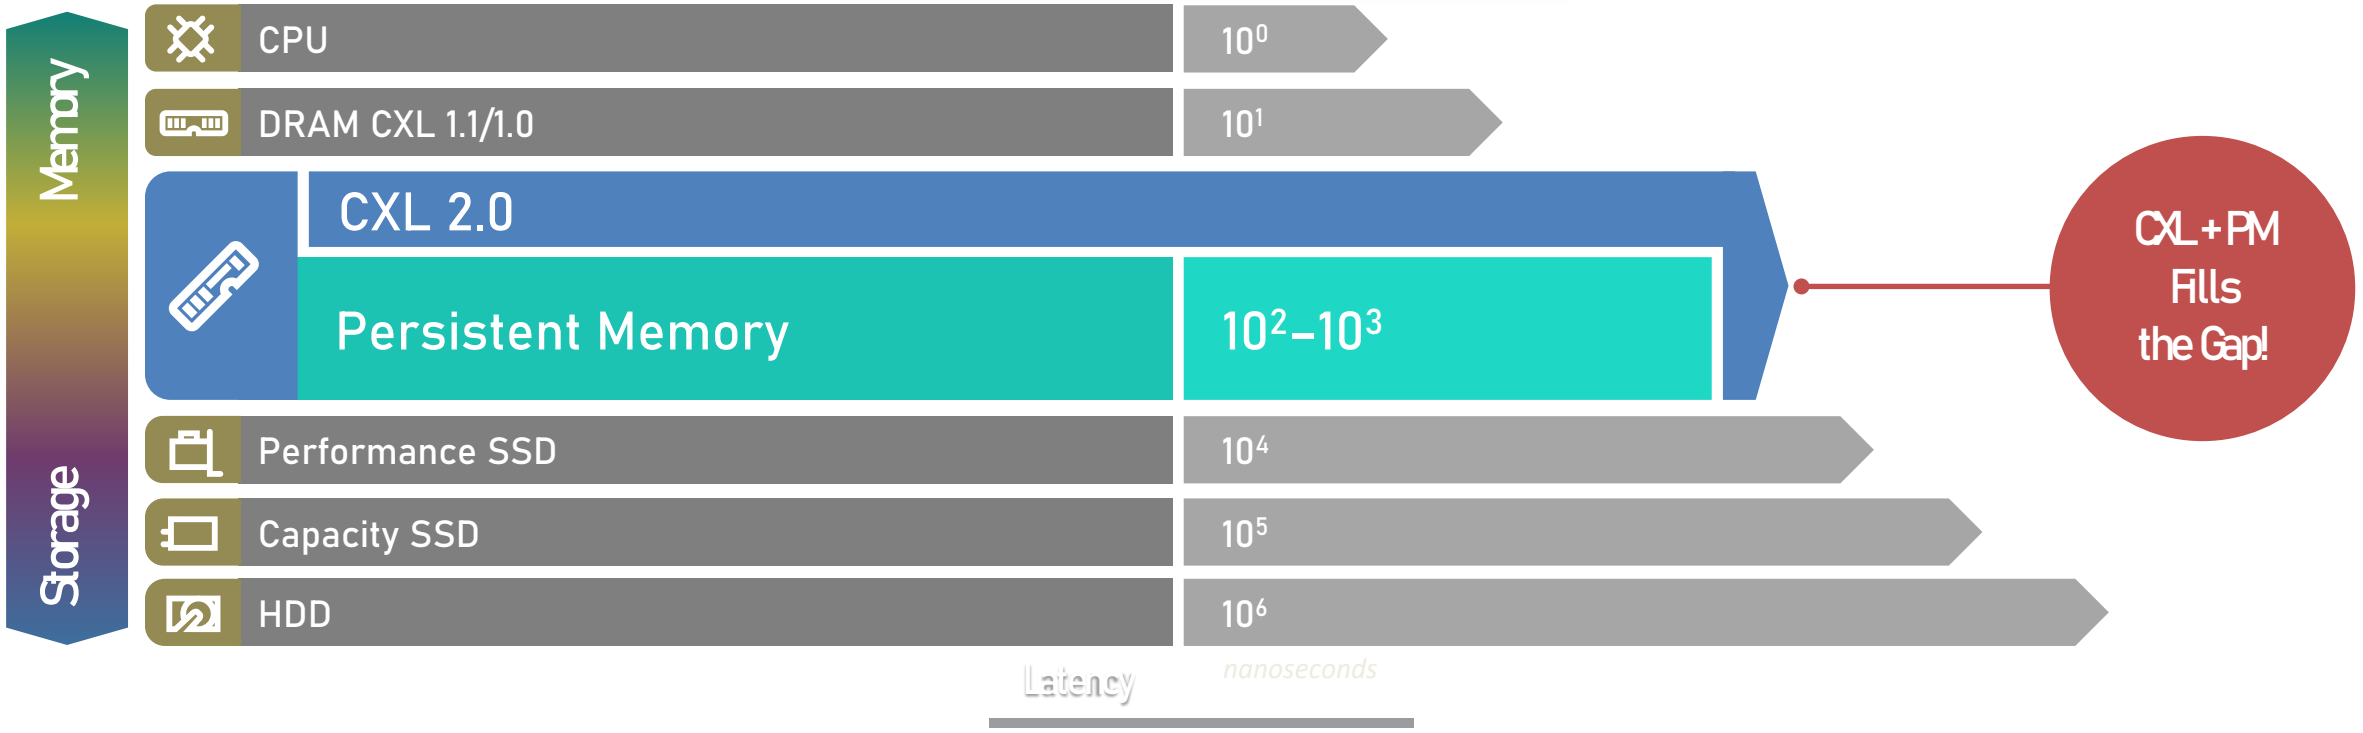

DATA CENTER: LOOKING OUTSIDE IN: SCOPE OF CXL 2.0 OVER CXL 1.1

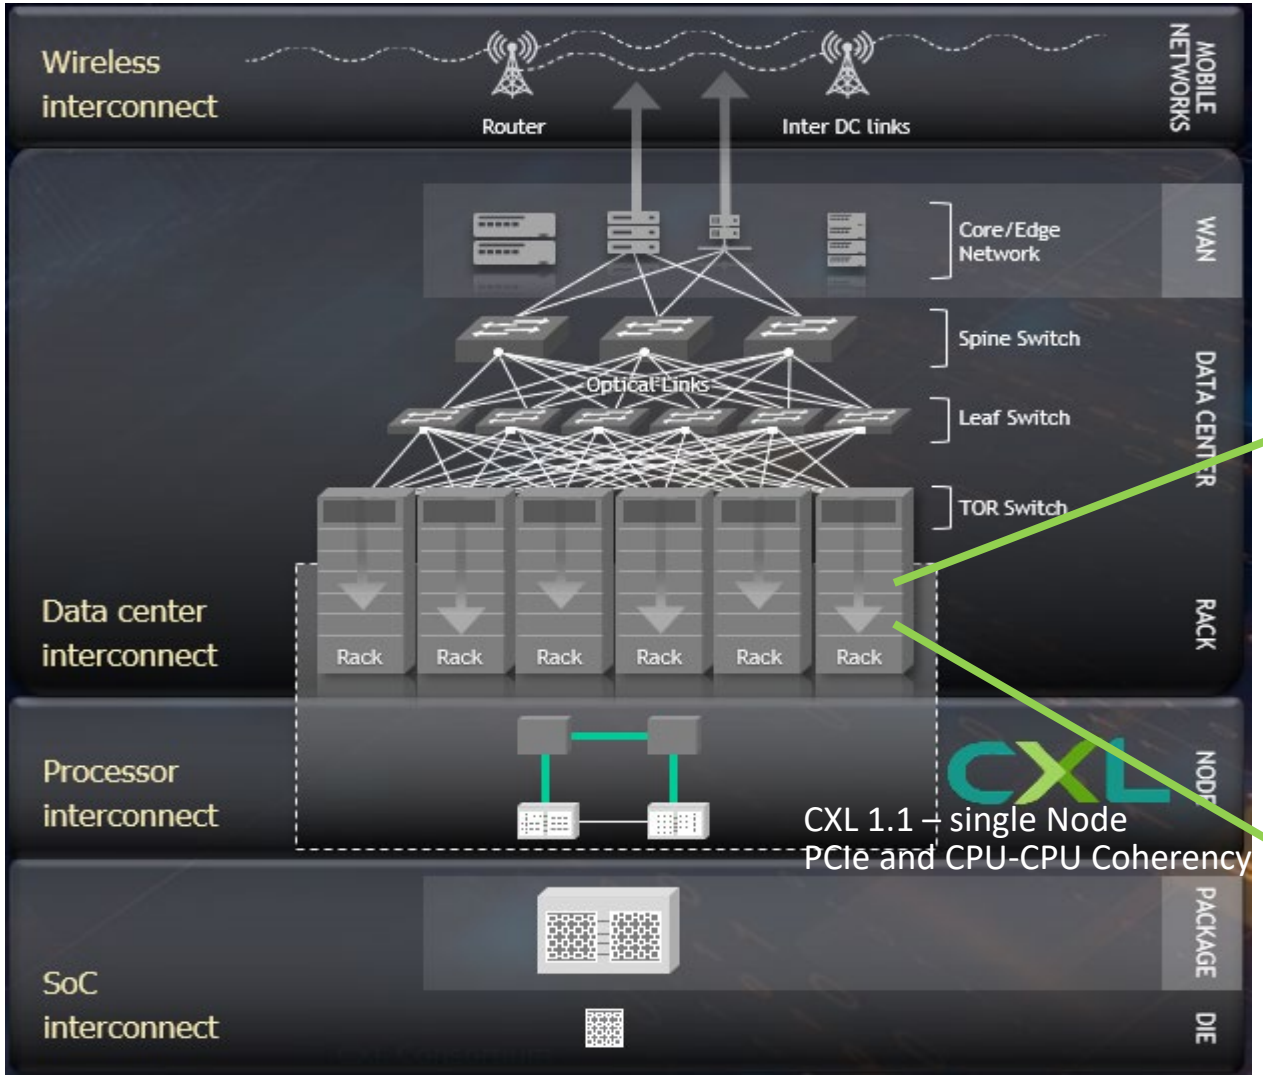

Memory/Accelerator Pooling with Single Logical Devices

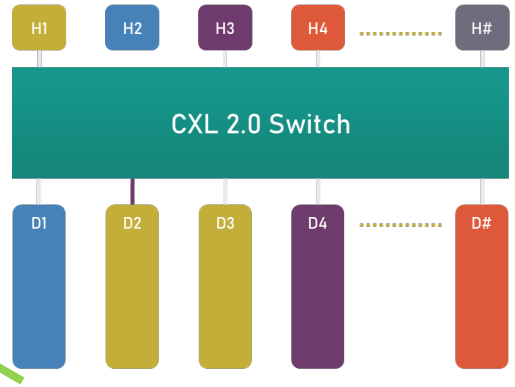

Memory Pooling with Multiple Logical Devices

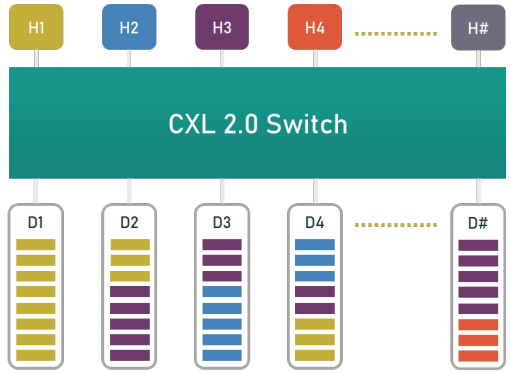

CXL 2.0: MEMORY POOLING

Memory/Accelerator Pooling with Single Logical Devices

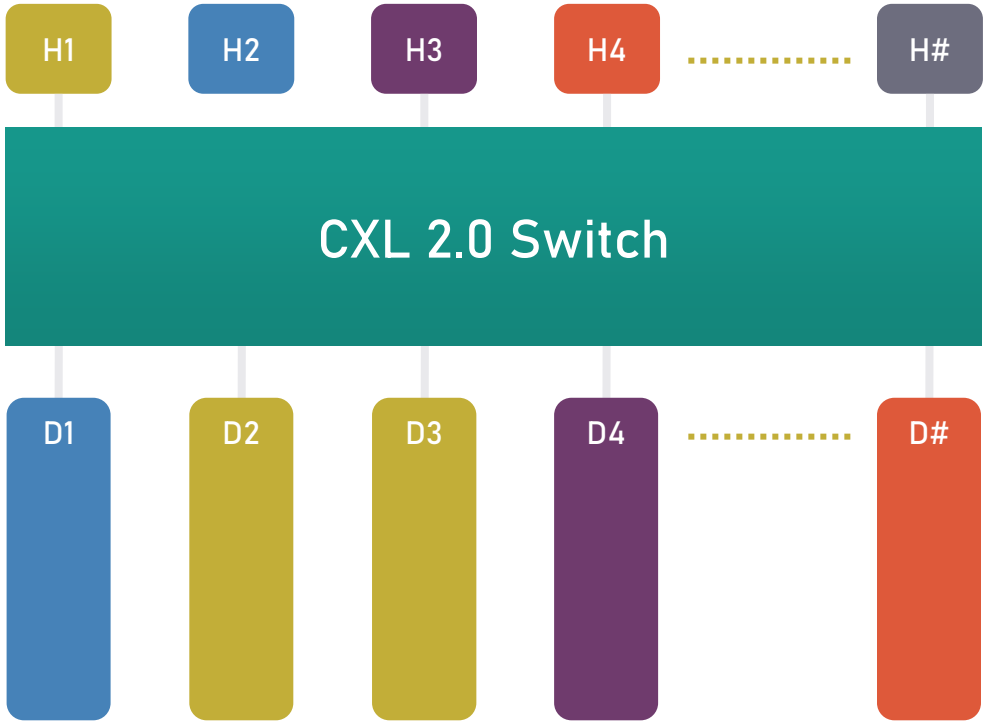

Memory Pooling with Multiple Logical Devices

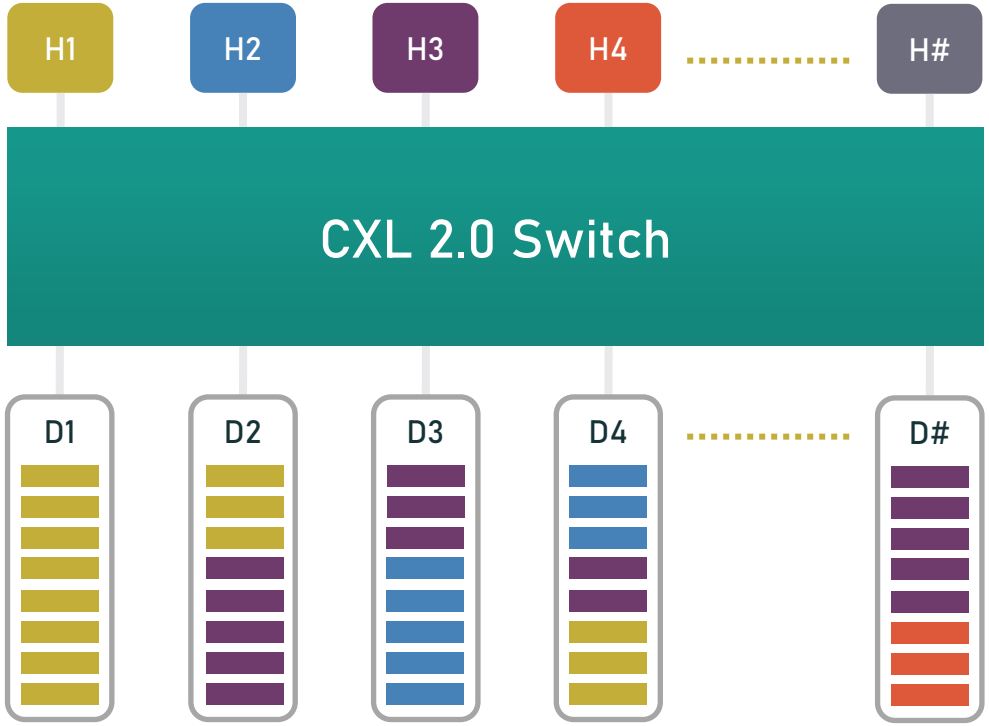

BENEFIT OF CXL 2.0 SWITCHING

Expansion

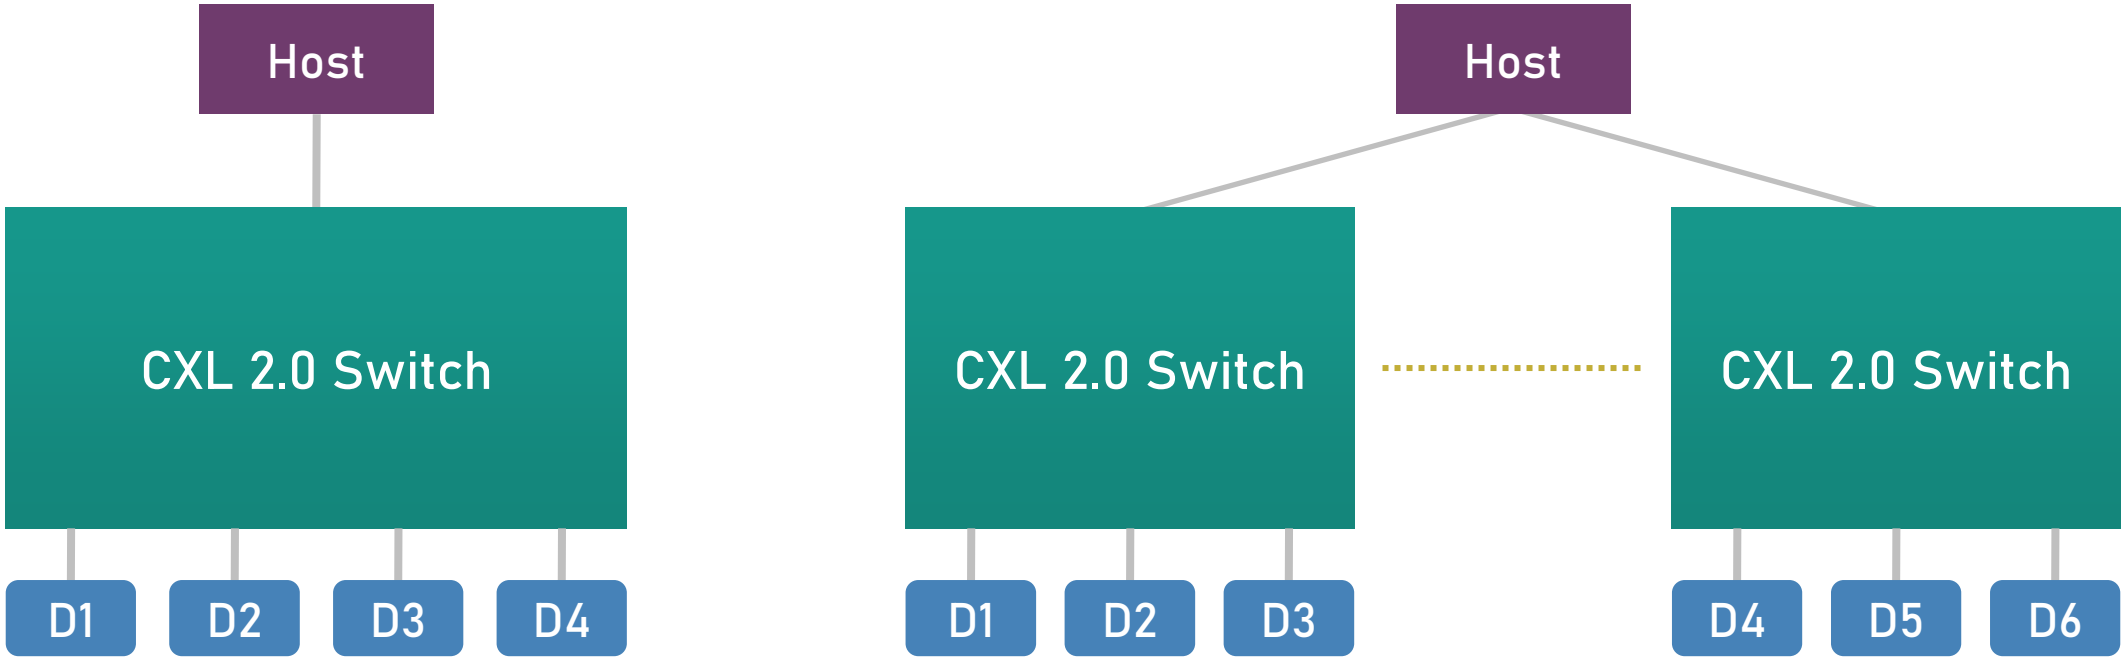

CXL 2.0: RAS CAPABILITIES

- In HPC applications, Reliability, Availability, and Serviceability (RAS) of attached memory is imperative
  - Reliability is the ability to provide correct service. One way to design a system to be reliable is to be able to detect and correct faults that may otherwise cause an error
  - Availability is the ability to be ready to provide correct service, possibly under degraded capability or performance.
  - Serviceability is the ability to diagnose and repair faults that may cause (or may have caused) an error.

RAS Capability	CXL 1.1	CXL 2.0
PCIe-based RAS mechanisms for link and protocol errors	★	★
Data poisoning	★	★
Viral	★	★
Error injection	★	★
Functional Level Reset (FLR) for CXL.io	★	★
Poison injection		★
Support for viral propagation through switches		★
Detailed memory error logging		★
Hot-plug (managed and surprise)		★
Global Persistent Flush and Dirty Shutdown tracking		★
Scan Media/Internal Poison List Retrieval		★

CXL 2.0 IDE

- **CXL 2.0 specification defined the protocol level changes needed to enable CXL.cachemem IDE.**
  - IDE encrypts and integrity protects the CXL.cache, .mem traffic on the link
- **Defines CXL IDE Key Management**
  - Requests by host software (or another agent like BMC) to IDE capable component
  - These requests provision IDE keys and instruct the components to initiate or tear down CXL IDE
  - Requests and responses are confidentiality and integrity protected
- **Enables architecture and code reuse by leveraging existing industry standards such**
  - Suite of Security Protocol and Data Model (SPDM) specifications from DMTF

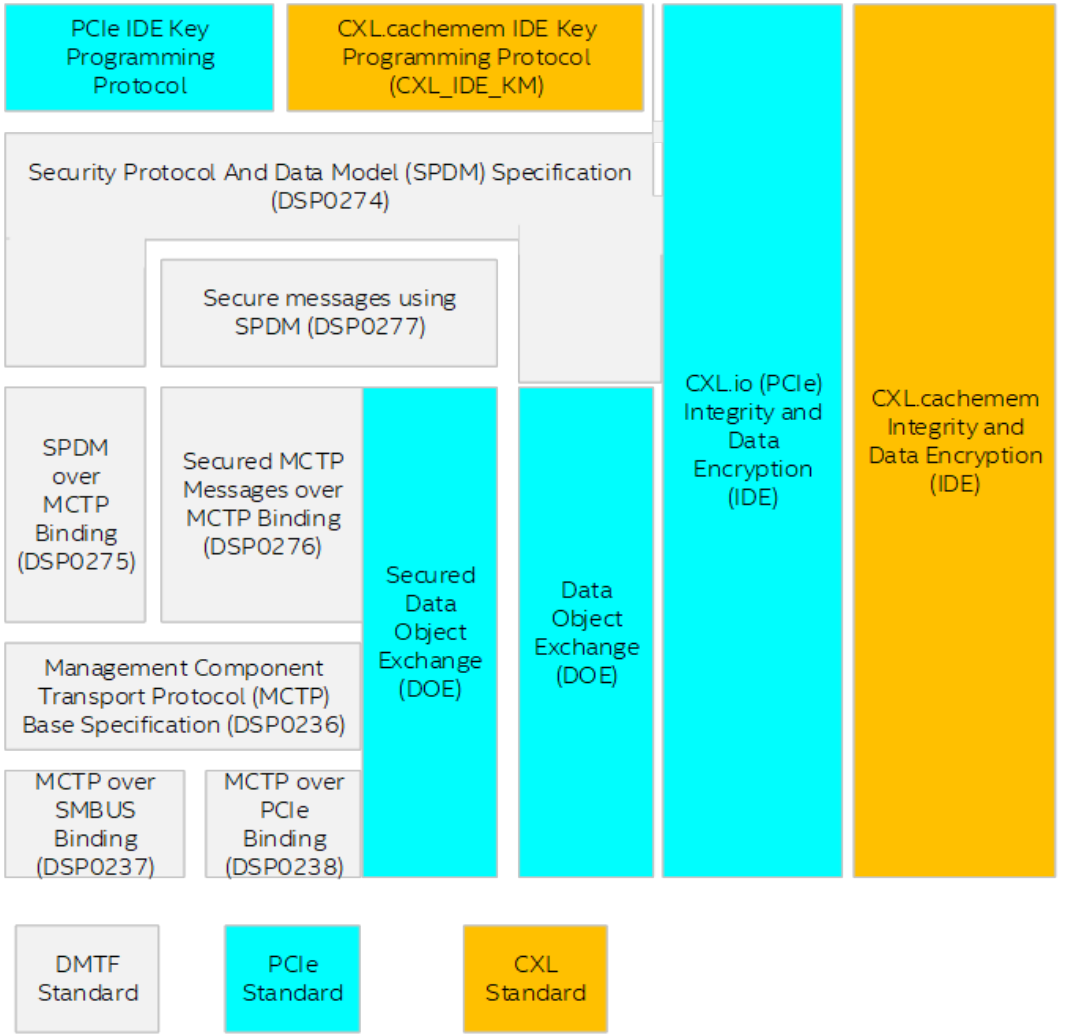

CXL 2.0 Security Benefits

CXL 2.0 provides Integrity and Data Encryption of traffic across all entities (Root Complex, Switch, Device)

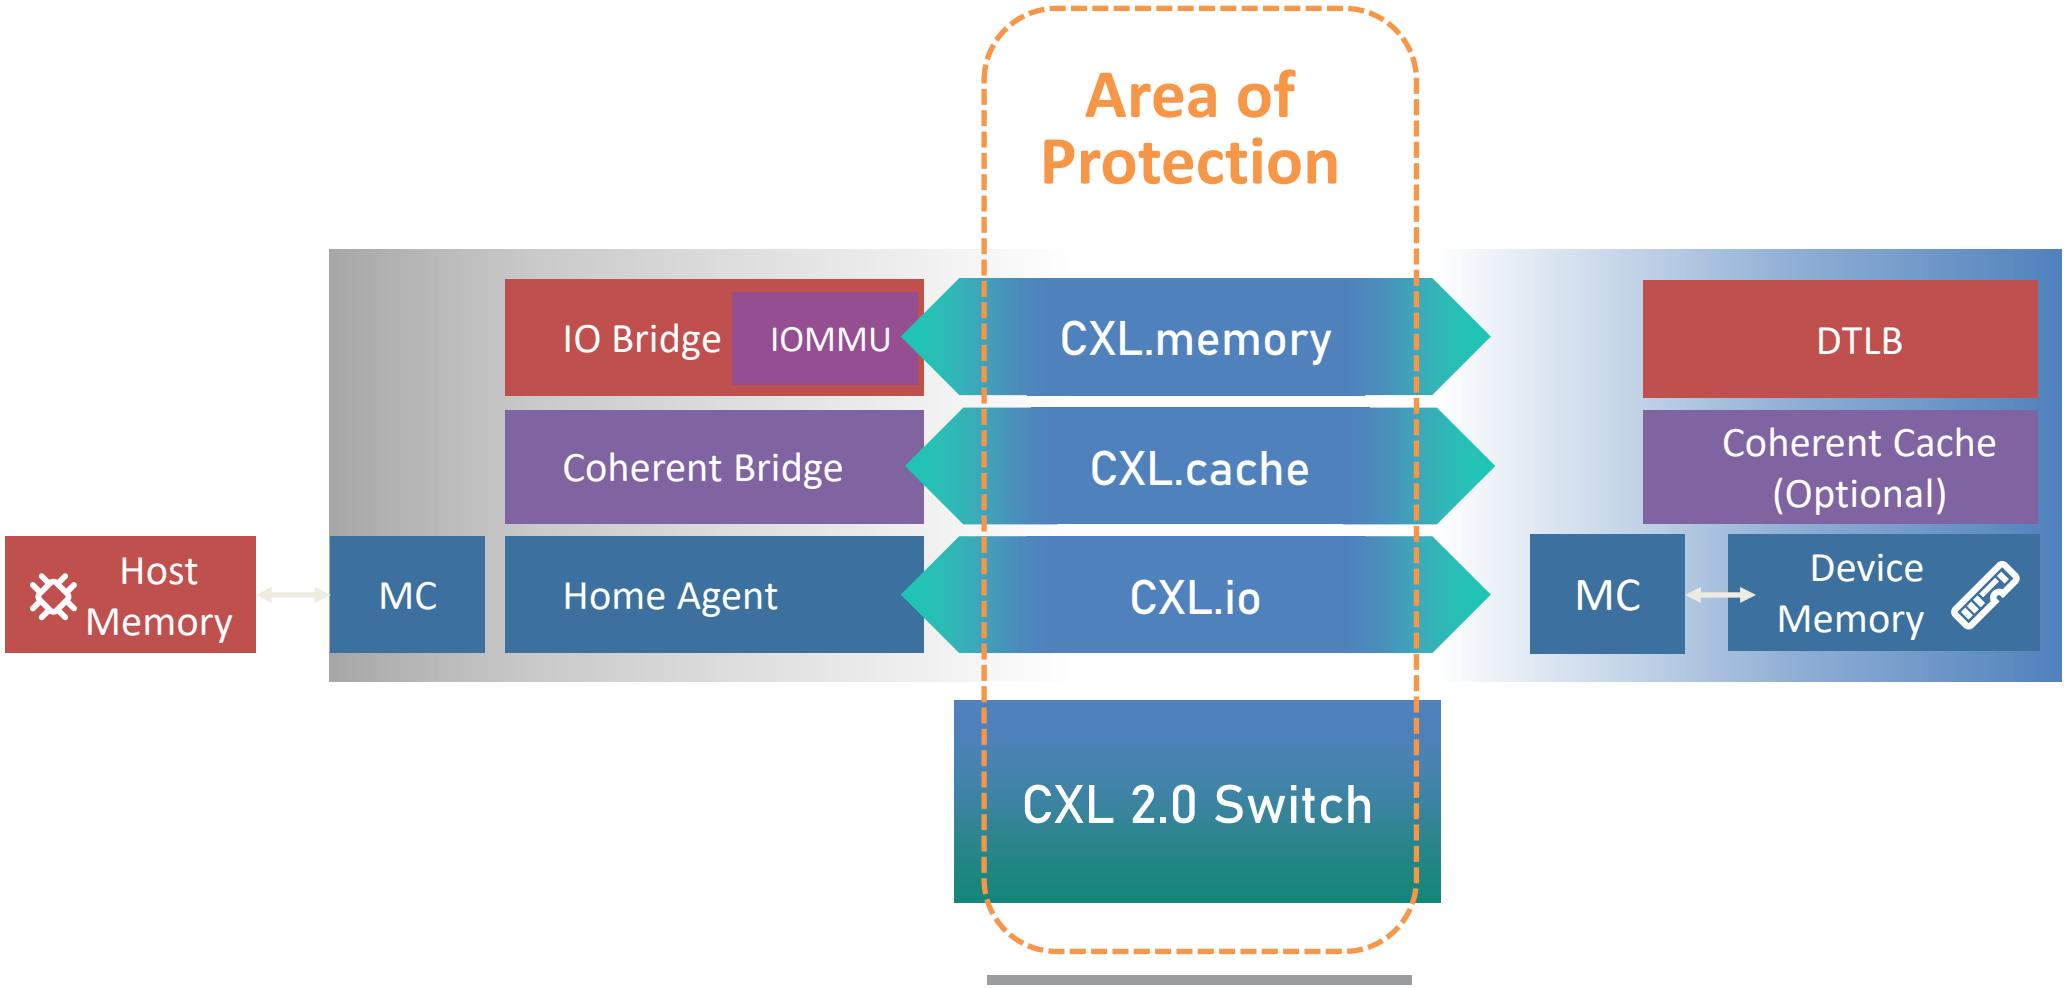

CXL 3.0 SPECIFICATION – WHAT IS COMING NEXT?

- **CXL 3.0 features will have…**
  - Double bandwidth
  - Improved capability for better scalability and improved resource utilization
    - Enhanced memory pooling and enables new memory usage models
    - Multi-level switching with multiple host and fabric capabilities and enhanced fabric management
    - New symmetric coherency capabilities
    - Improved software capabilities
  - Fully backward compatible with CXL 1.0, CXL 1.1, CXL 2.0
  - Specification available 2022

IN SUMMARY

CXL Consortium momentum continues to grow

- 180+ members and growing
- Responding to industry needs and challenges
- Developing on establishing an MOU with OFA to extend fabrics management for datacenter fabrics

CXL 2.0 introduces new features & usage models

- Switching, pooling, persistent memory support, security
- Backward compatible with CXL 1.1 and 1.0
- Built in Compliance & Interop program

Call to action

- Join CXL Consortium and get involved in our Working Groups
- Visit www.computeexpresslink.org for more information
- Follow us on Twitter and LinkedIn for more updates!

2022 OFA Virtual Workshop

THANK YOU

Microchip

CXL Consortium