

CXL™ Consortium Update

Siamak Tavallaei, CXL President



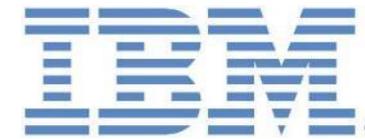
EMPOWERING OPEN.



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

Introducing the CXL Consortium



CXL Board of Directors

Industry Open Standard for
High Speed Communications

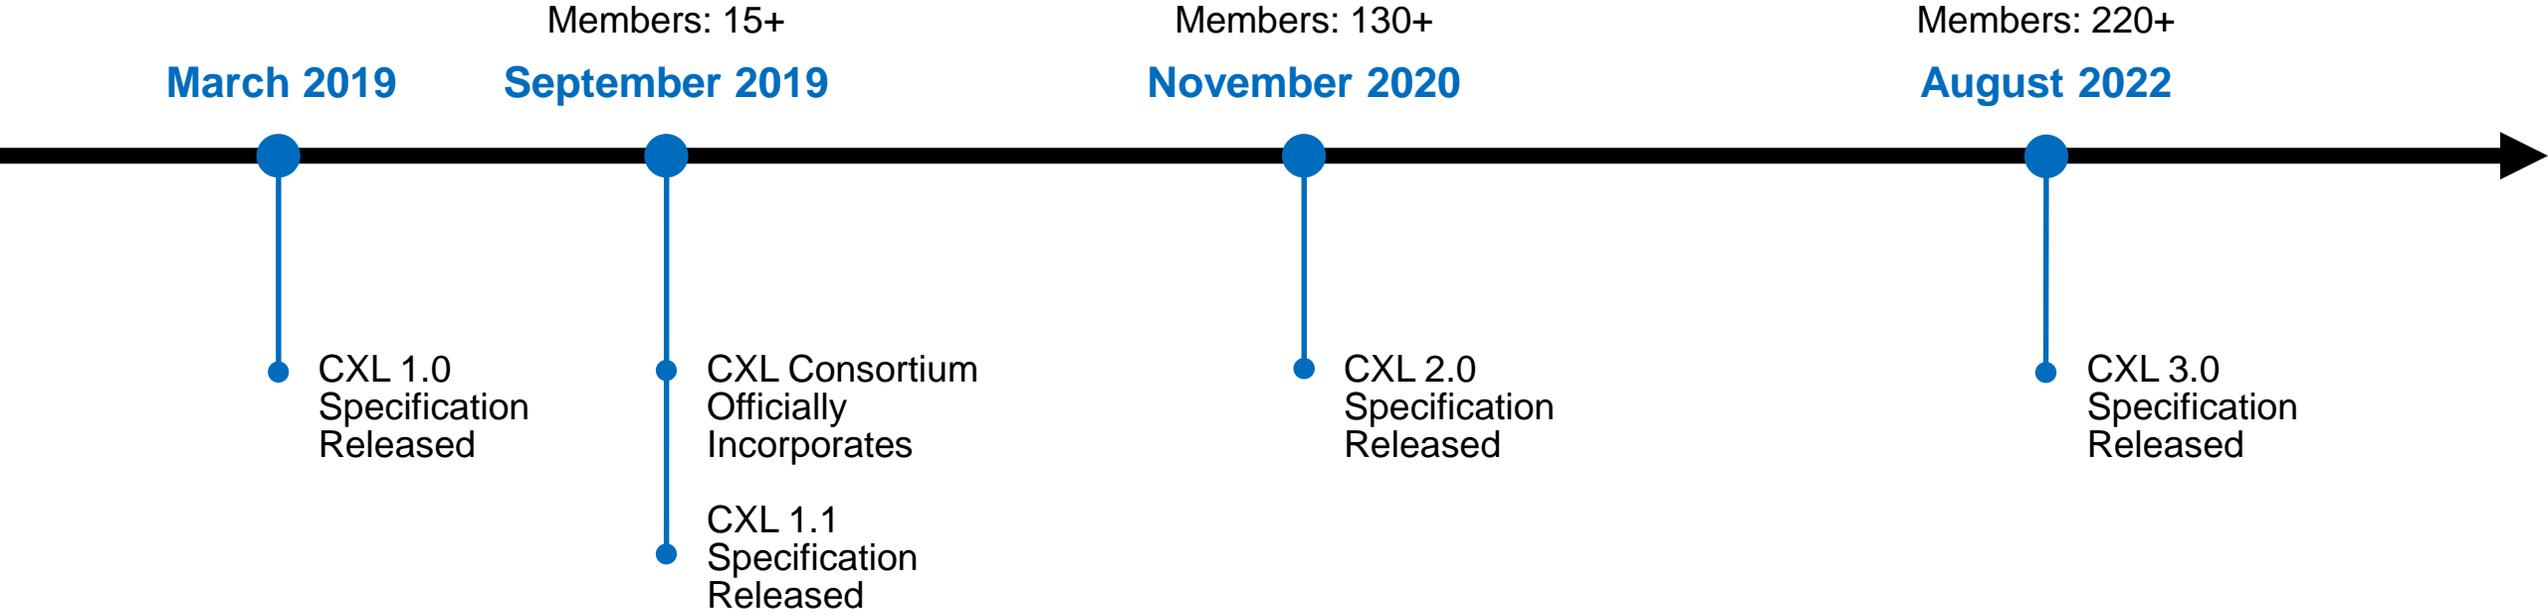
220+ Member Companies

CXL Specification Release Timeline



Press Release

August 2, 2022, Flash Memory Summit
CXL Consortium releases Compute Express Link 3.0 specification to expand fabric capabilities and management



Compute Express Link™ (CXL™) Overview

Industry Landscape

Proliferation of Cloud Computing



Growth of AI & Analytics



Cloudification of The Network & Edge

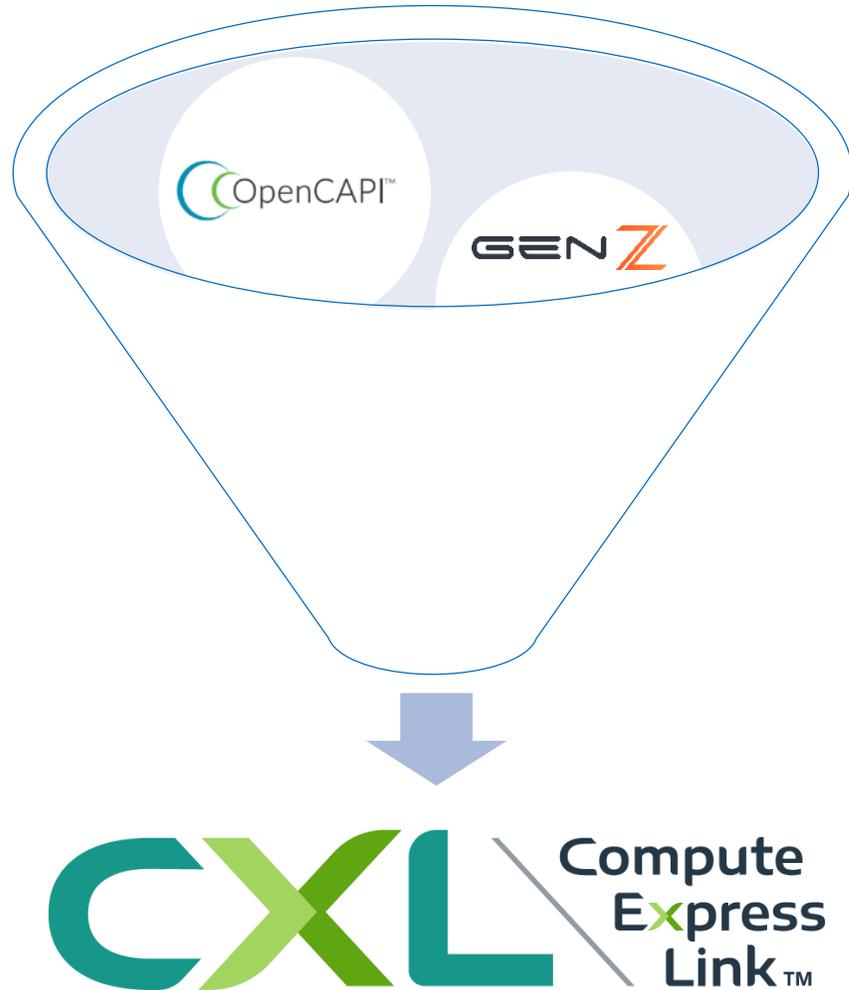


Growing Industry Momentum

- CXL Consortium showcased first public demonstrations of CXL technology at SC'21
- View virtual and live demos from CXL Consortium members here: <https://www.computeexpresslink.org/videos>
 - Demos showcase CXL usages, including memory development, memory expansion and memory disaggregation



Industry Focal Point



CXL is emerging as the industry focal point for coherent IO

- CXL Consortium and OpenCAPI sign letter of intent to transfer OpenCAPI specification and assets to the CXL Consortium



August 1, 2022, Flash Memory Summit
**CXL Consortium and OpenCAPI Consortium
Sign Letter of Intent to Transfer OpenCAPI
Assets to CXL**

- In February 2022, CXL Consortium and Gen-Z Consortium signed agreement to transfer Gen-Z specification and assets to CXL Consortium



Unveiling the CXL 3.0 specification



Press Release

August 2, 2022, Flash Memory Summit
CXL Consortium releases Compute Express Link 3.0 specification to expand fabric capabilities and management

CXL 3.0 Specification

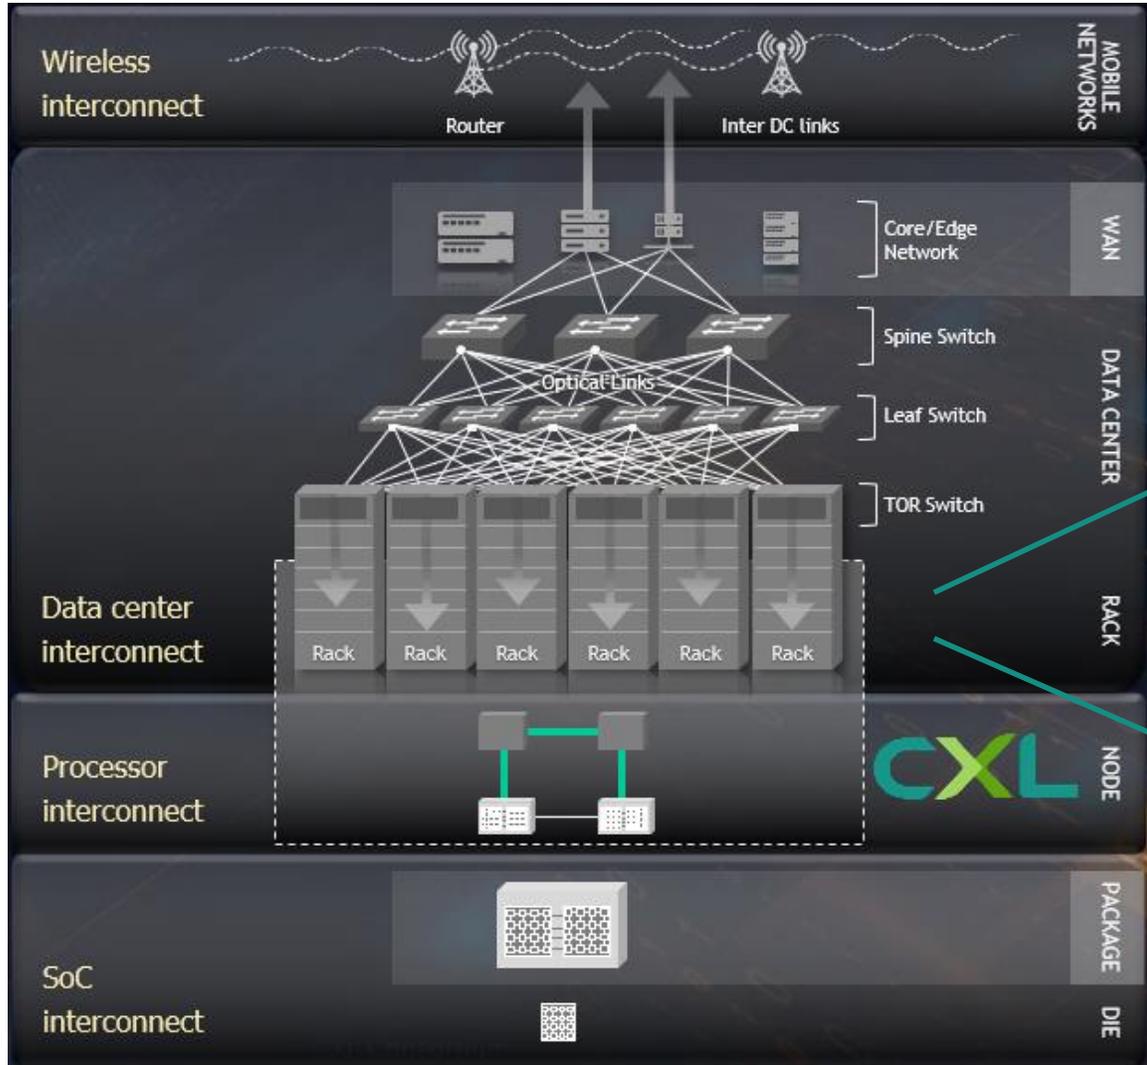
Industry trends

- Use cases driving need for higher bandwidth include: high performance accelerators, system memory, SmartNIC and leading edge networking
- CPU core-count increasing: efficiency is declining due to reduced memory capacity and bandwidth per core
- Memory bottlenecks due to CPU pin and thermal constraints
- Need efficient peer-to-peer resource sharing across multiple domains

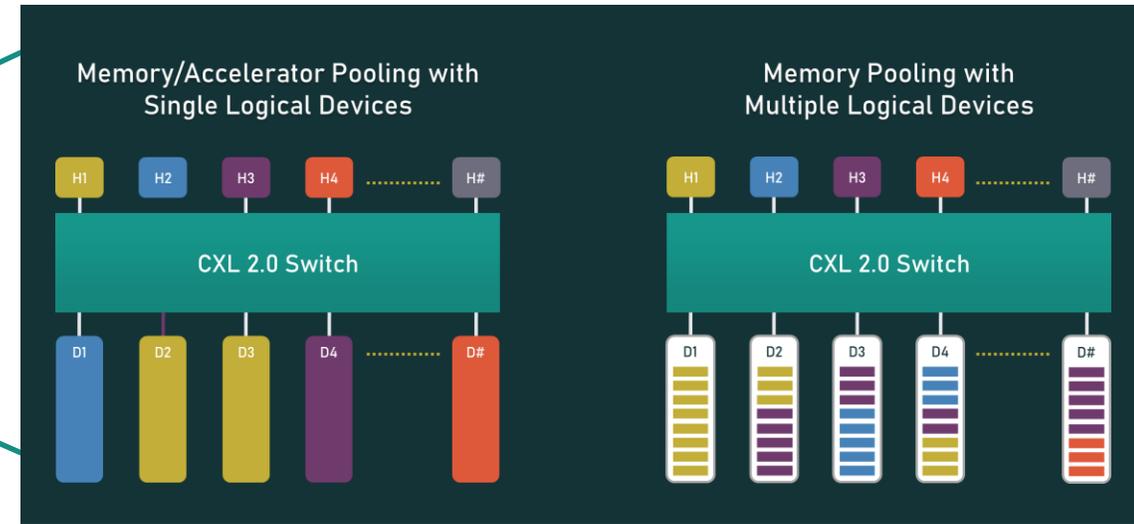
CXL 3.0 introduces

- Fabric capabilities
 - Multi-headed and fabric-attached devices
 - Enhanced fabric management
 - Composable disaggregated infrastructure
- Improved capability for better scalability and resource utilization
 - Enhanced memory pooling
 - Multi-level switching
 - New enhanced coherency capabilities
 - Improved software capabilities
- Double the bandwidth
- No added latency over CXL 2.0
- Backward-compatibility with CXL 2.0, CXL 1.1, and CXL 1.0

Data Center: Expanding Scope of CXL

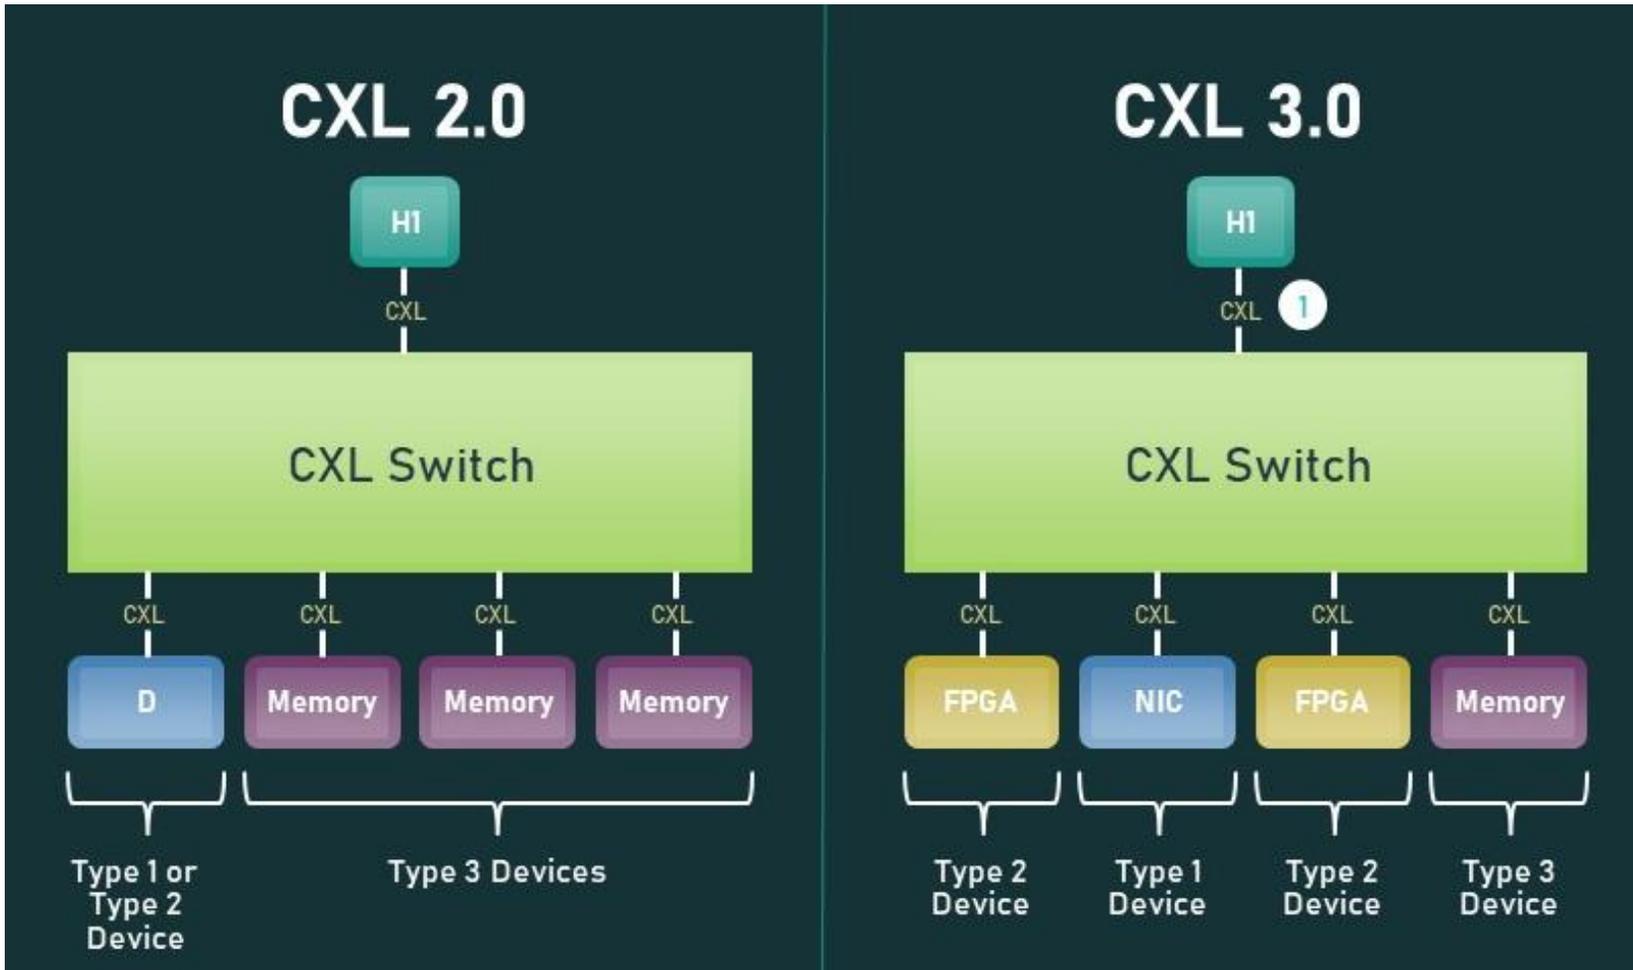


Future - CXL 3.0
Fabric growth for
disaggregation/pooling/accelerator



CXL 2.0 across Multiple Nodes inside a Rack/ Chassis supporting pooling of resources

Multiple Devices of all Types per Root Port

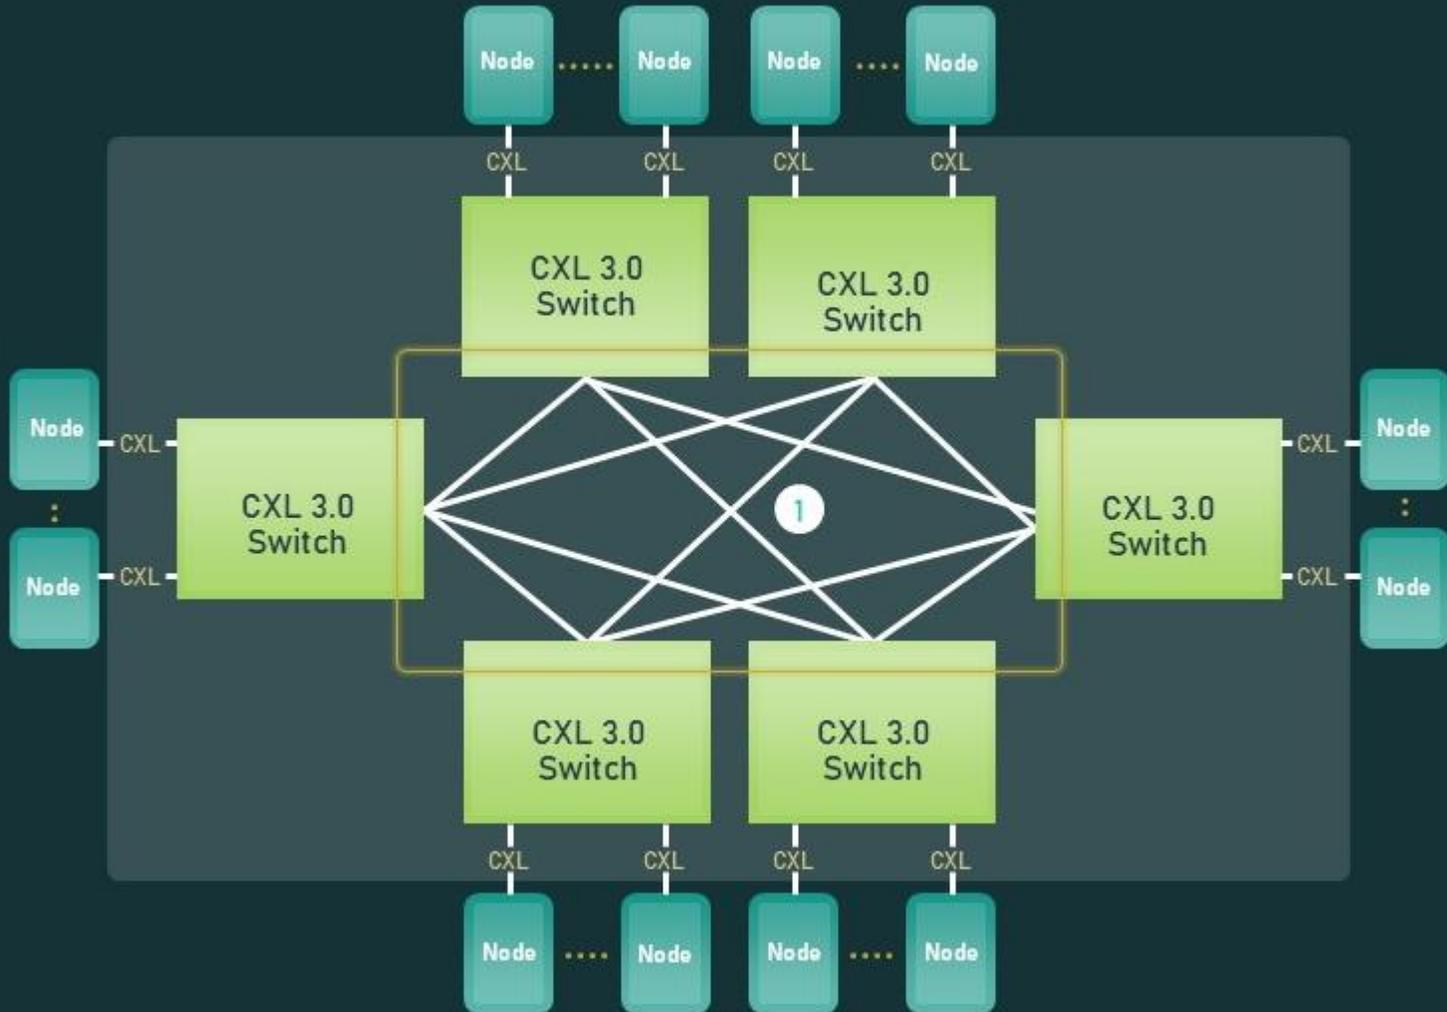


Fabric capabilities

- Multi-headed and fabric-attached devices
- Enhanced fabric management
- Composable disaggregated infrastructure

① Each host's root port can connect to **more than one device type**

Fabrics Overview



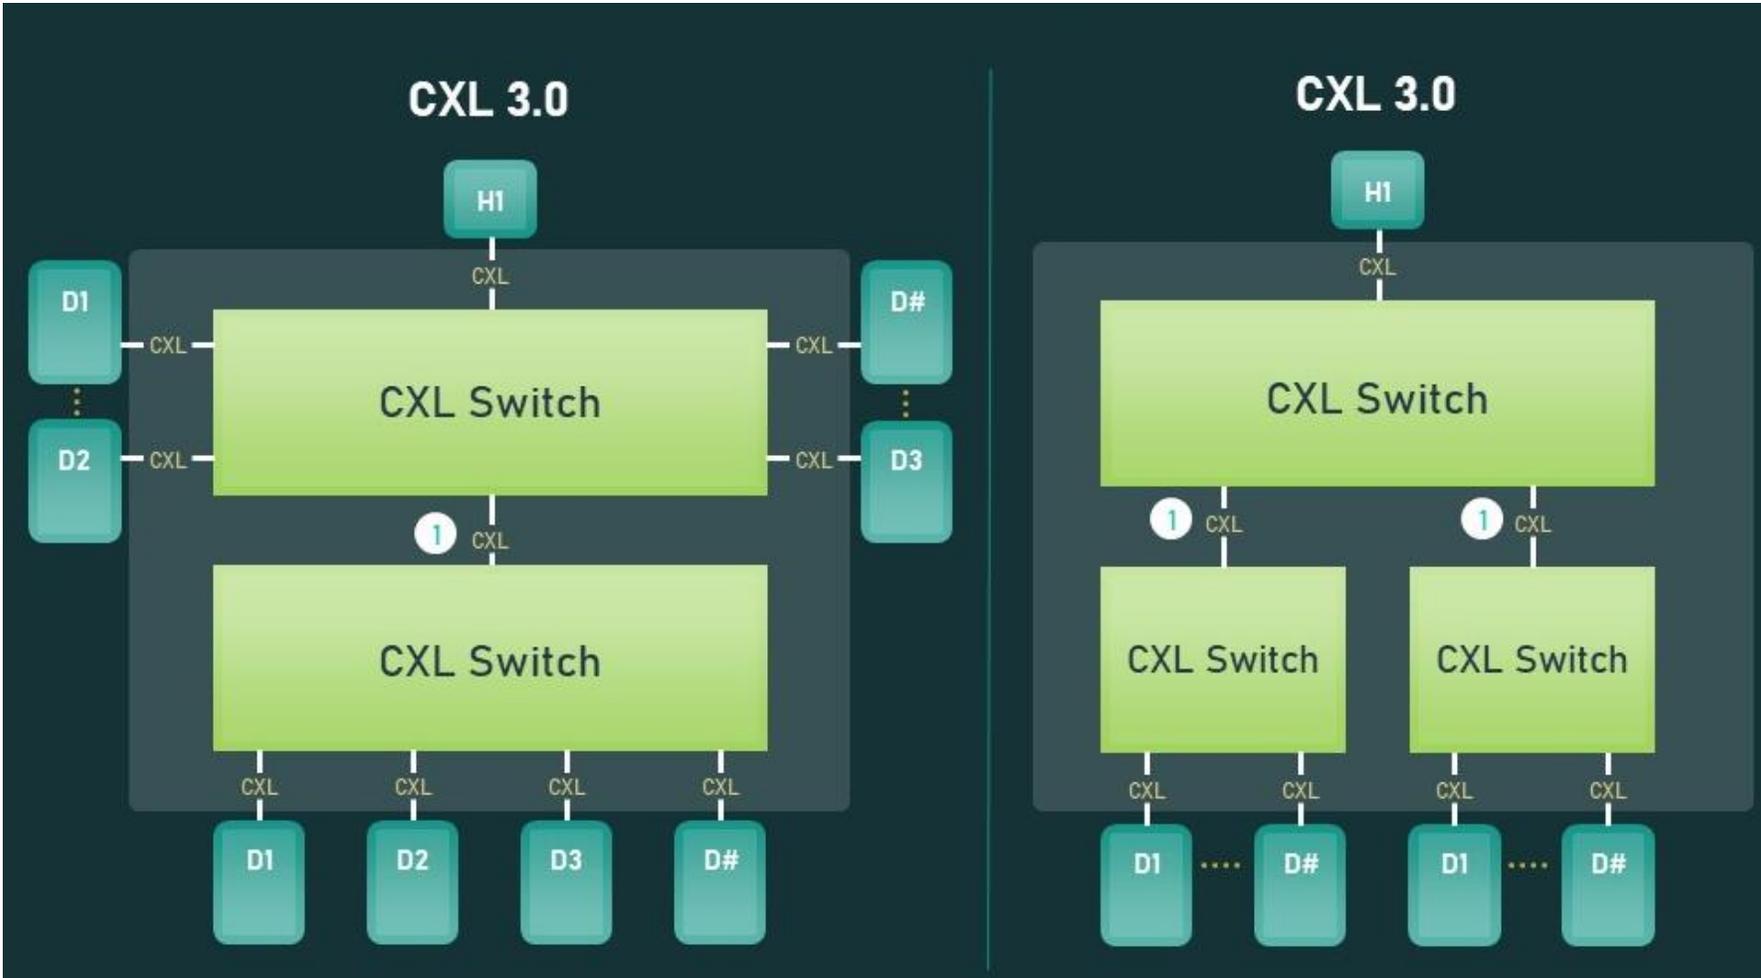
Improved capability for better scalability and resource utilization

- Enhanced memory pooling
- Multi-level switching
- New enhanced coherency capabilities
- Improved software capabilities

- ① CXL 3.0 enables non-tree architectures
- Each node can be a CXL Host, CXL device or PCIe device

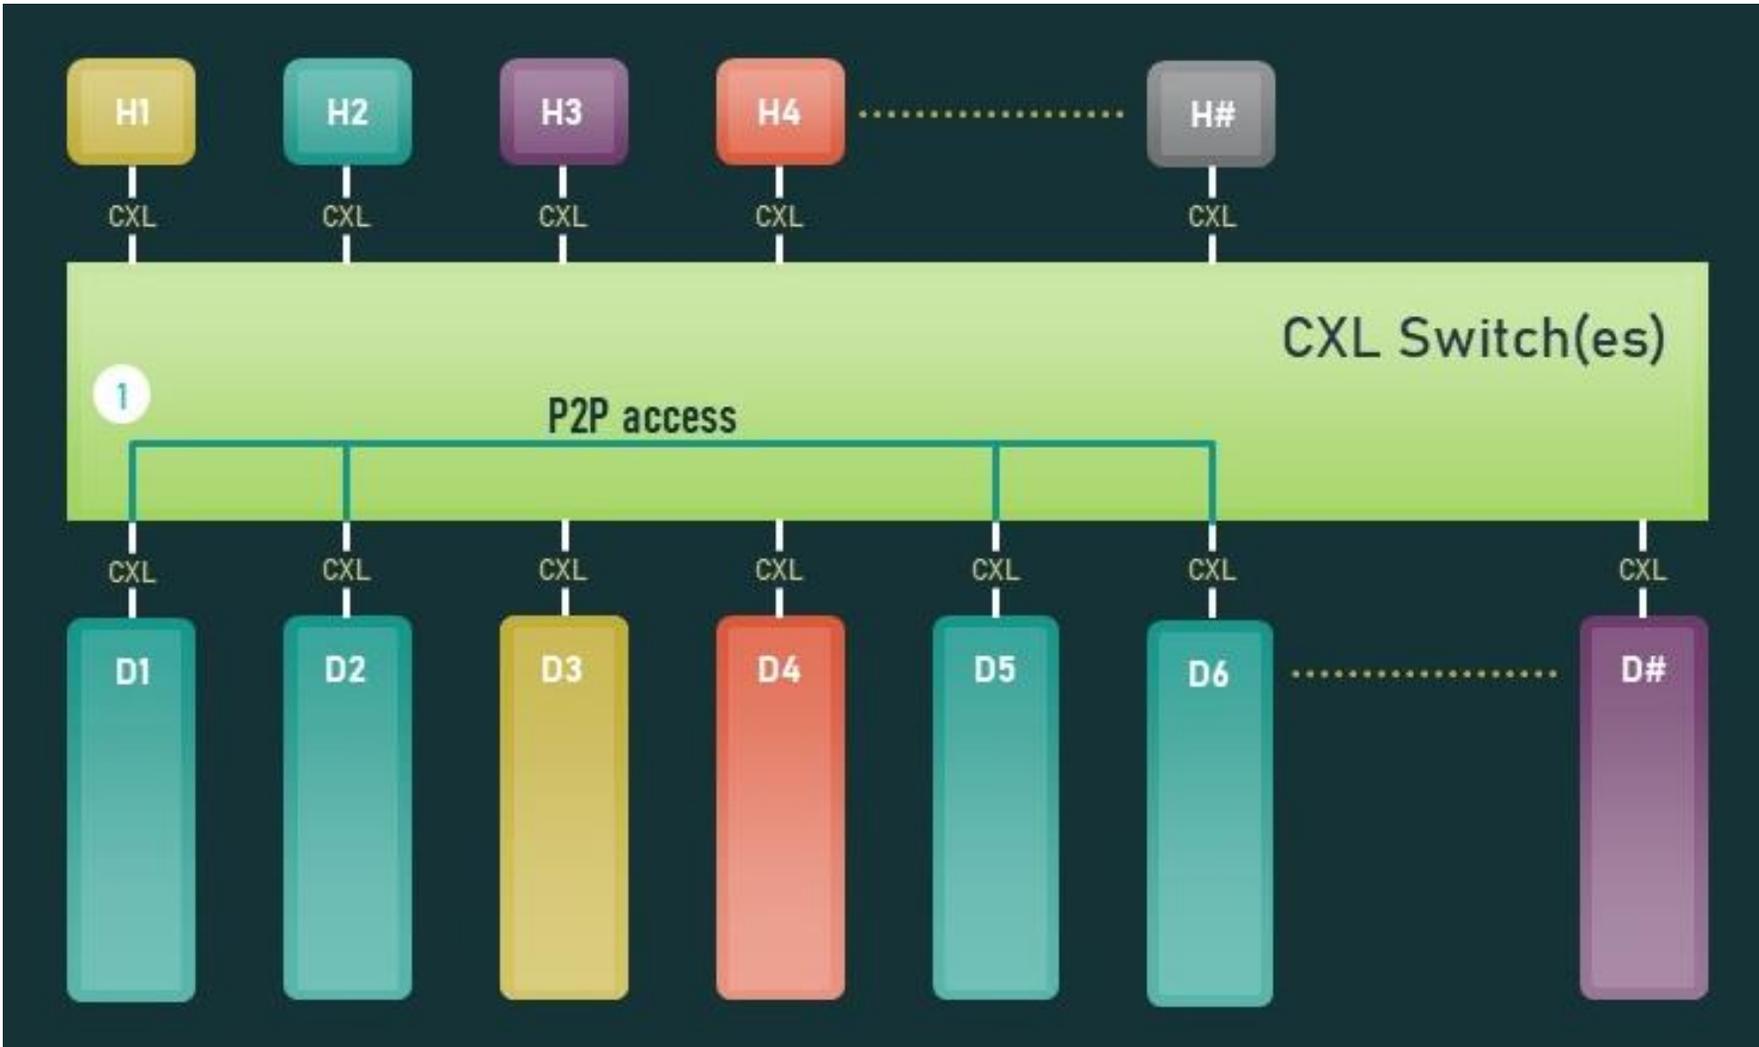
Switch Cascade/Fanout

Supporting vast array of switch topologies



- ① Multiple switch levels (aka cascade)
 - Supports fanout of all device types

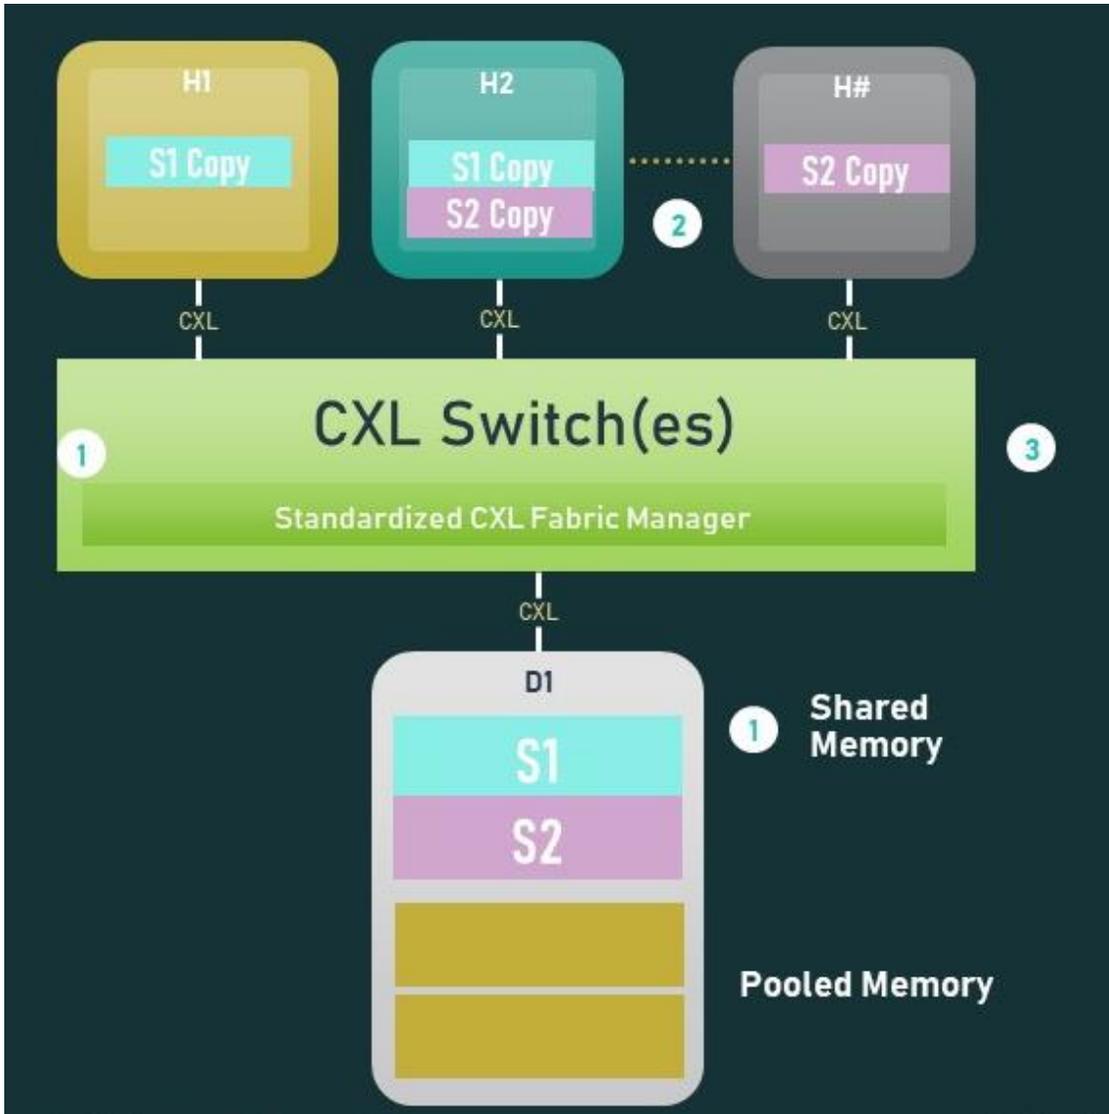
Device to Device Comms



1

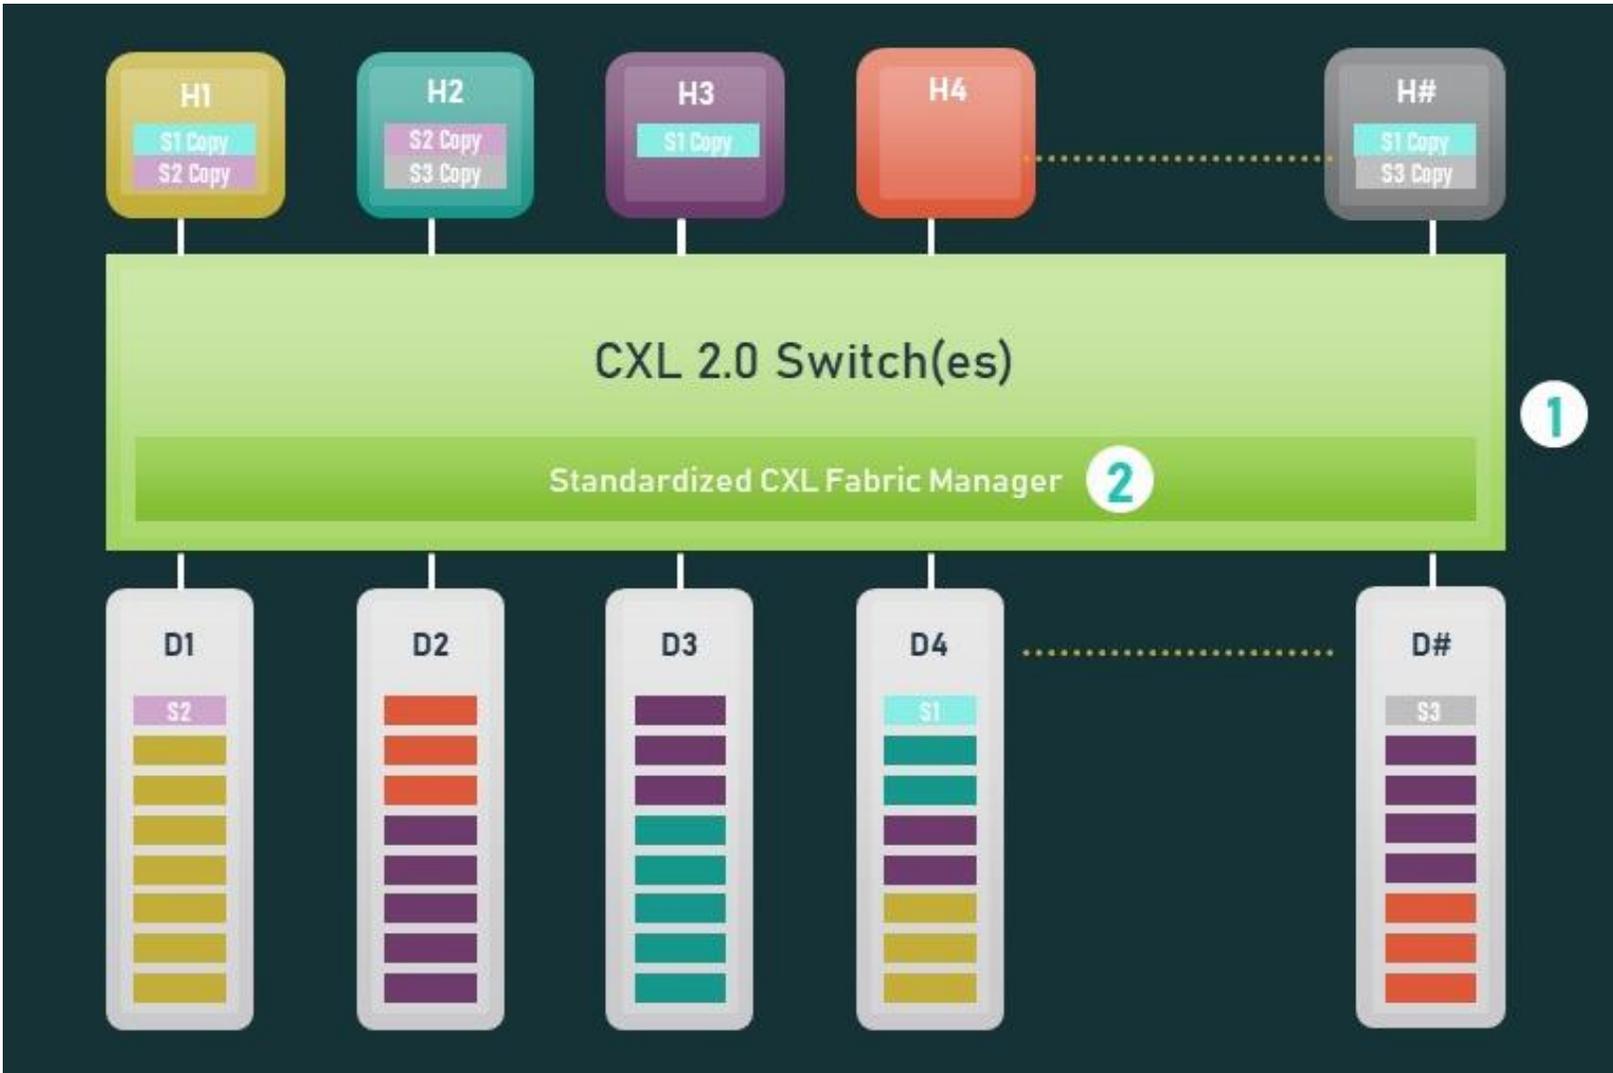
- CXL 3.0 enables **peer-to-peer communication (P2P)** within a virtual hierarchy of devices
- Virtual hierarchies are associations of devices that maintains a coherency domain

Coherent Memory Sharing



- ① Device **memory can be shared by all hosts** to increase data flow efficiency and improve memory utilization
- ② Host can have a **coherent copy of the shared region** or portions of shared region in host cache
- ③ **CXL 3.0 defined mechanisms to enforce hardware cache coherency** between copies

Memory Pooling and Sharing



- ① Expanded use case showing **memory sharing and pooling**
- ② CXL Fabric Manager is available to setup, deploy, and modify the environment

CXL 3.0 Specification Feature Summary

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			✓
Enhanced coherency (256 byte flit)			✓
Memory sharing (256 byte flit)			✓
Multiple Type 1/Type 2 devices per root port			✓
Fabric capabilities (256 byte flit)			✓

Not supported

✓ Supported

CXL 3.0: Expanding CXL Use Cases

- **Enabling new usage models**

- Memory sharing between hosts and peer devices
- Support for multi-headed devices
- Expanded support for Type-1 and Type-2 devices
- GFAM provides expansion capabilities for current and future memory

Download the CXL 3.0 specification on www.ComputeExpressLink.org

Call to Action

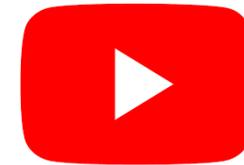
- Join the CXL Consortium, visit www.computeexpresslink.org/join
- Attend CXL Forum at OCP Global Summit on Thursday, October 20th for a deep-dive into the CXL 3.0 specification
- Engage with us on social media



[@ComputeExLink](https://twitter.com/ComputeExLink)



www.linkedin.com/company/cxl-consortium/



[CXL Consortium Channel](https://www.youtube.com/channel/UC...)

Thank you!