Introducing the CXL 3.1 Specification

Presented by

Rob Blankenship and Mahesh Wagh

CXL Board of Directors

CXL Compute Express Link®

Industry Open Standard for High Speed Communications

255+ Member Companies

Industry Landscape

Proliferation of Cloud Computing

Growth of AI & Analytics

Cloudification of the Network & Edge

CXL Ecosystem

Growth of CXL ecosystem since its inception

- Ecosystem that meets the ever-increasing performance and scale requirements
- Fully backwards compatible
- Lower overall system cost
- Comprehensive compliance and testing support

Data Center: Expanding Scope of CXL

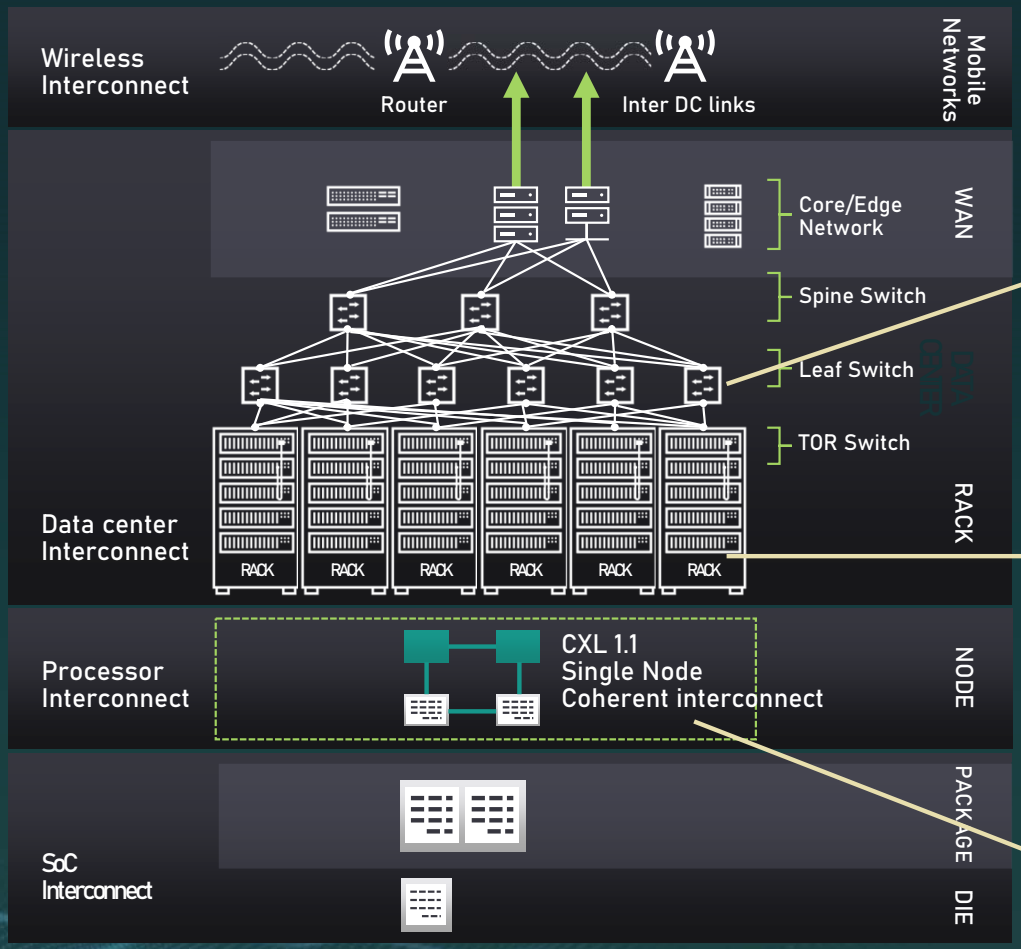

CXL 3.0/3.1

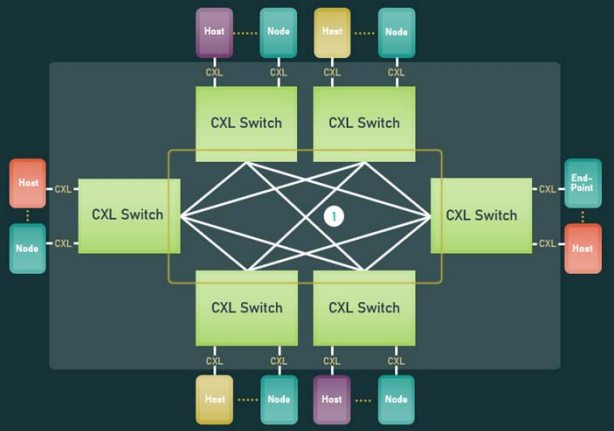

Composable Fabric growth for disaggregation/pooling/accelerator

CXL 2.0

Memory/Accelerator Pooling with Single Logical Devices

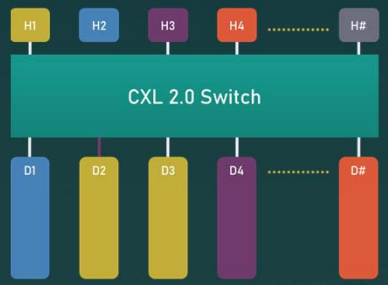

Memory Pooling with Multiple Logical Devices

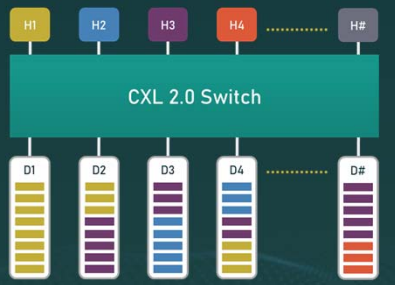

Multiple Nodes inside a Rack/ Chassis supporting pooling of resources

Representative CXL Usages

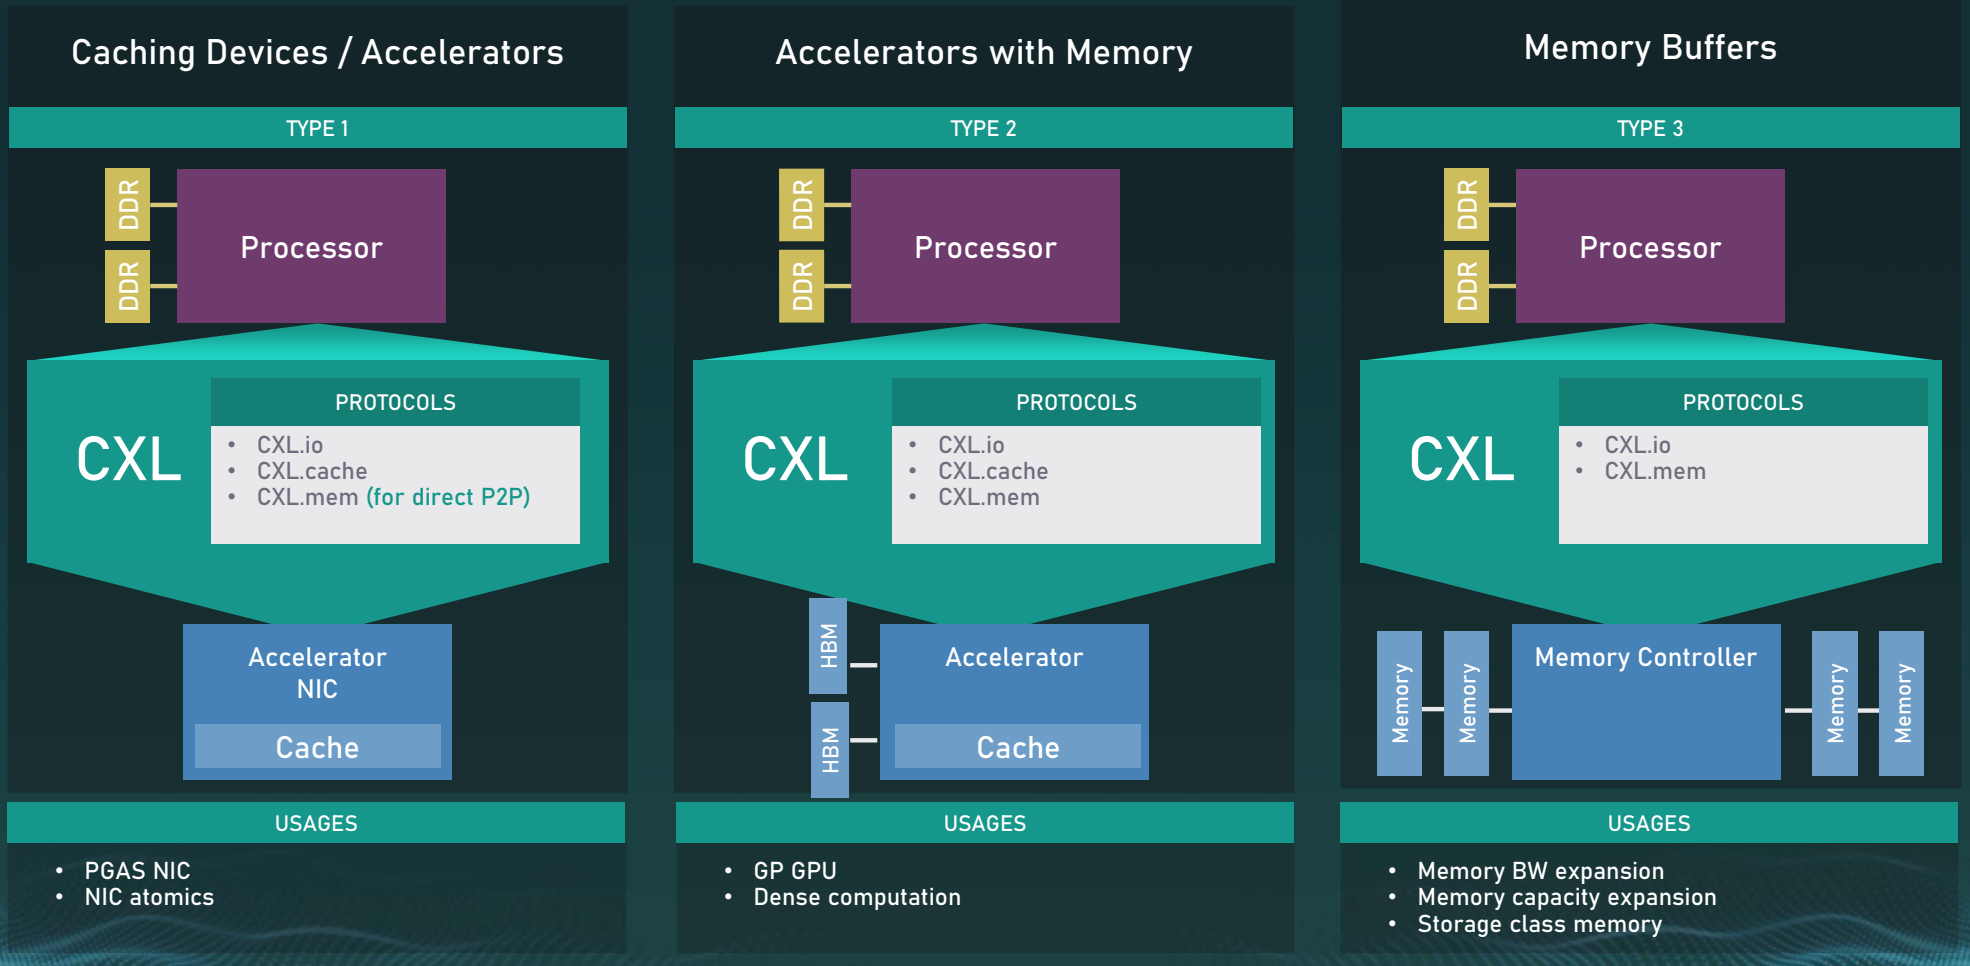

CXL Specification Release Timeline

CXL Specification Feature Summary

Features	CXL 1.0 / 1.1	CXL 2.0	CXL 3.0
Release date	2019	2020	August 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

RECAP: CXL 3.0: POOLING & SHARING

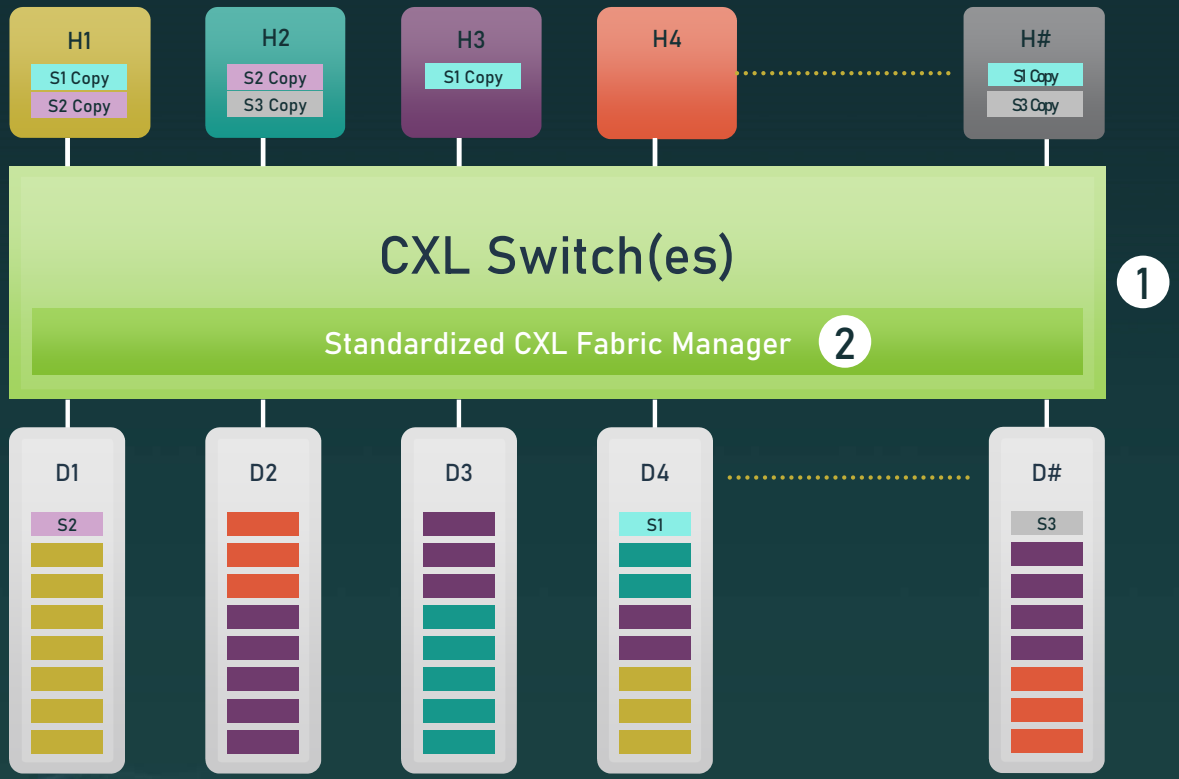

1. Expanded use case showing memory sharing and pooling
2. CXL Fabric Manager is available to setup, deploy, and modify the environment

CXL 3.1

Feature enhancements

CXL 3.1 Feature Enhancements

- The CXL specification continues to evolve to meet new usage models
- New features introduced in the CXL 3.1 specification:
  - CXL Fabric Improvements/Extensions
    - Scale-out of CXL fabrics using PBR (Port Based Routing)
  - Trusted-Execution-Environment Security Protocol (TSP)
    - Allows for Virtualization-based Trusted Execution Environments (TEEs) to host Confidential Computing Workloads
  - Memory Expander Improvements
    - Up to 32-bit of meta data and RAS capability enhancements

CXL Fabric Improvements/ Extensions

CXL Fabric Improvements/Extensions

- Fabric Decode/Routing requirements
- Host-to-Host communication with Global Integrated Memory (GIM) concept
- Direct P2P .mem support through PBR Switches
  - Adds symmetric .mem Link Layer definition for devices
  - Enables direct caching of CXL.mem for an accelerator, which is not possible with UIO
- Fabric Manager (FM) API definition for PBR Switch

CXL 3.1: Fabric Enhancement Features

Port-Based Routing (PBR) compared to Hierarchy-Based Routing (HBR)

Support fabric topologies other than tree topologies that HBR switches offer

- Address-based, non-prescriptive routing for large memory fabrics
- Supports tree, mesh, ring, star, butterfly, and multi-dimensional topologies

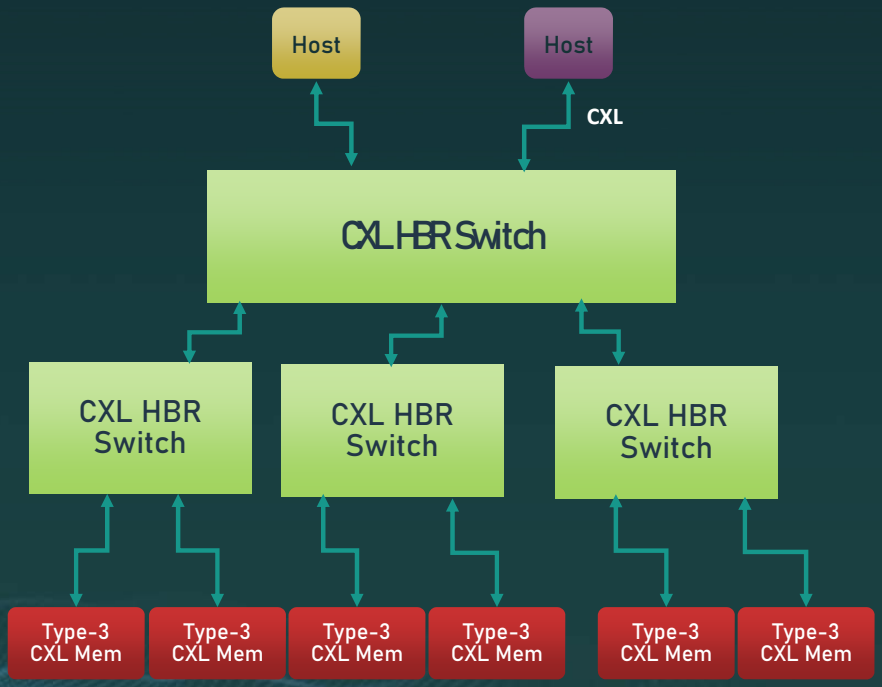

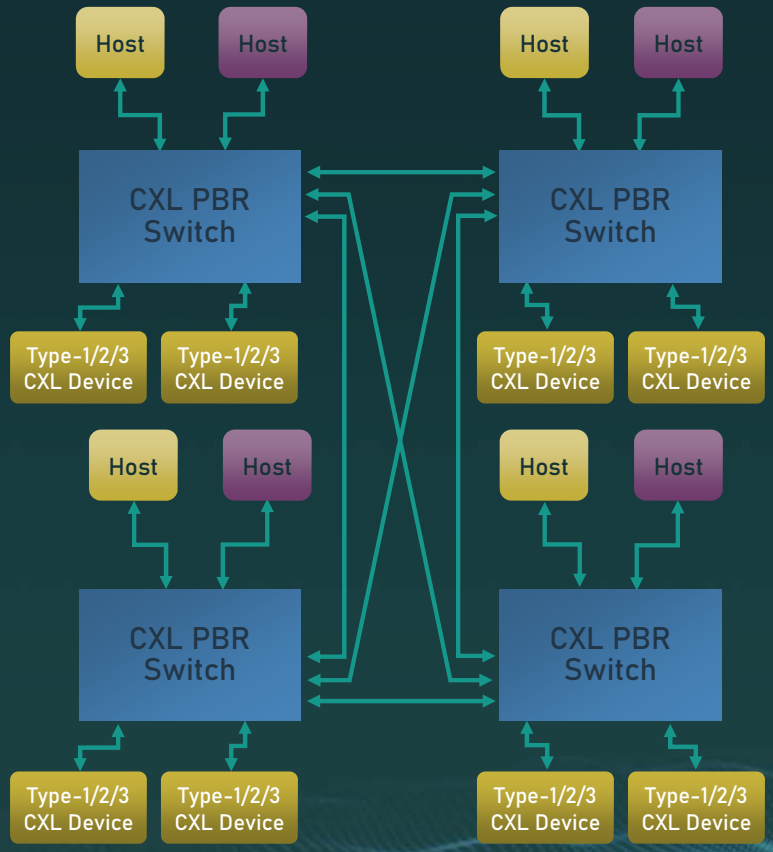

CXL 3.1: Fabric Enhancement Features

Host-to-Host communication with Global Integrated Memory (GIM) concept

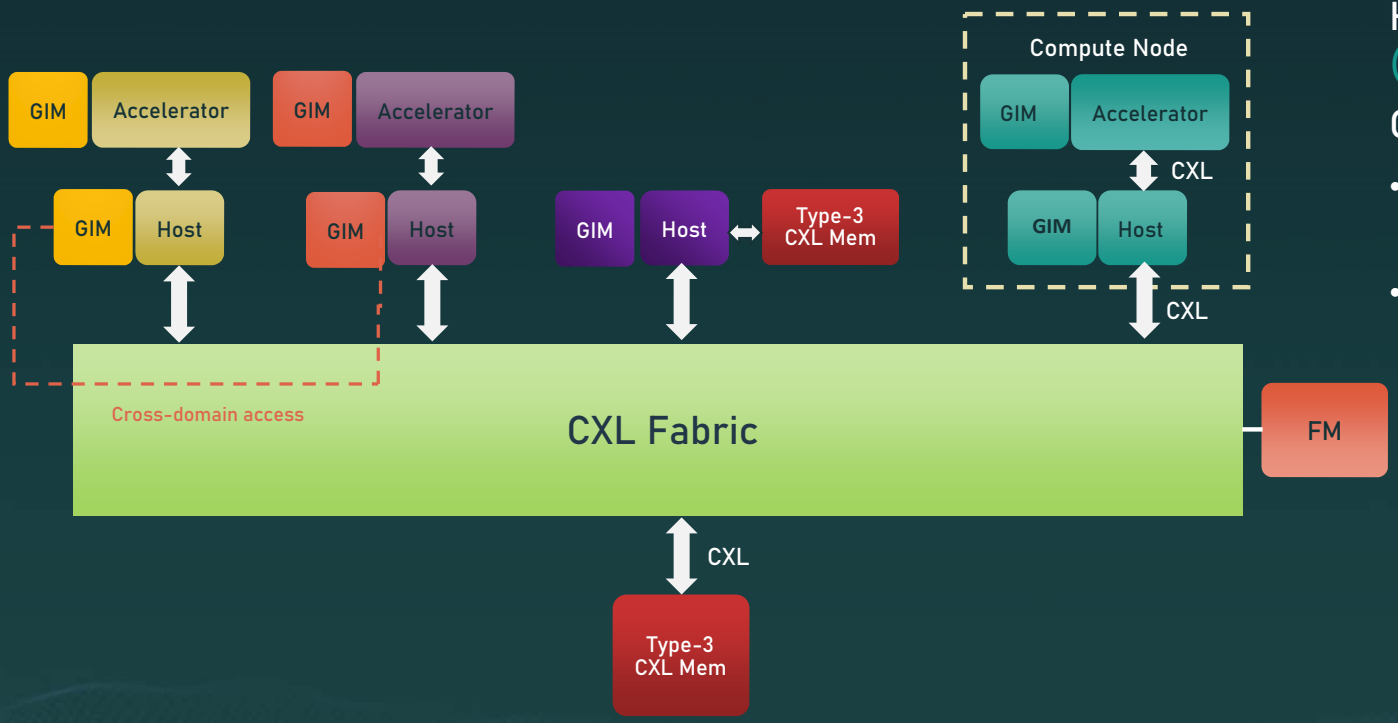

Host to Fabric-Attached Memory (FAM) communication with Global Integrated Memory (GIM)

- Multiple Hosts mapping CXL Fabric-Attached Memory devices
- Hosts and FAM devices can initiate cross-domain accesses to GIM

CXL 3.1: Fabric Enhancement Features

Direct peer-to-peer (P2P) .mem support for Accelerators through PBR Switches

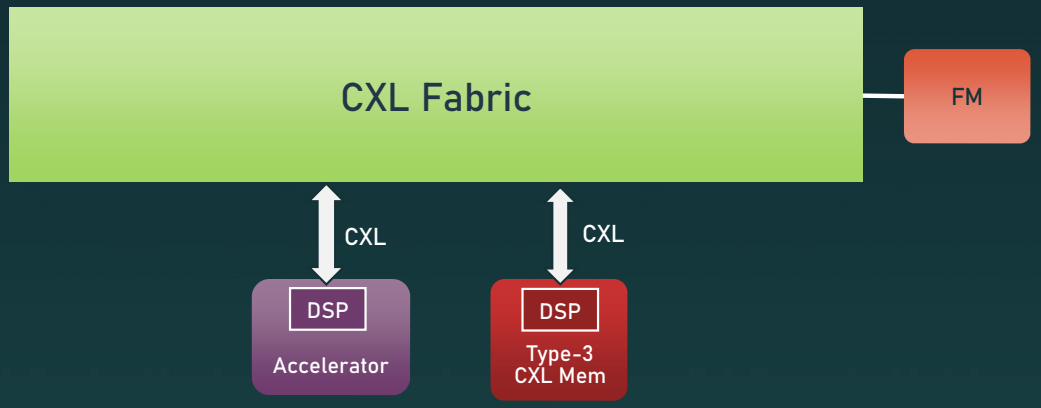

Direct P2P

Direct P2P CXL.mem for Accelerators

- Enables accelerator access to peer Type-3 memory
- The accelerator and Type-3 device must each be directly connected to an Edge Downstream Port (DSP)
- Utilizes port-based routing (PBR) for transactions
- Target memory device can be dedicated or shared with the host.

CXL 3.1: Fabric Enhancement Features

Fabric Manager (FM) API definition for PBR Switch

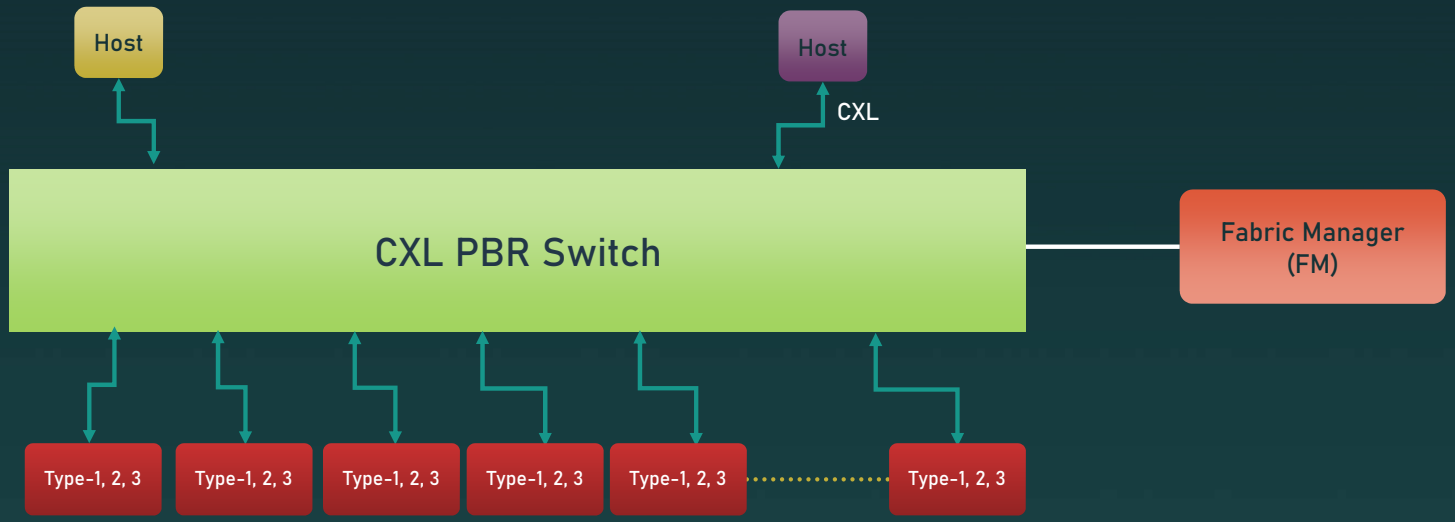

Trusted Execution Environment (TEE)
&
TEE Security Protocol (TSP)

RECAP: CXL 2.0 Security Benefits

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

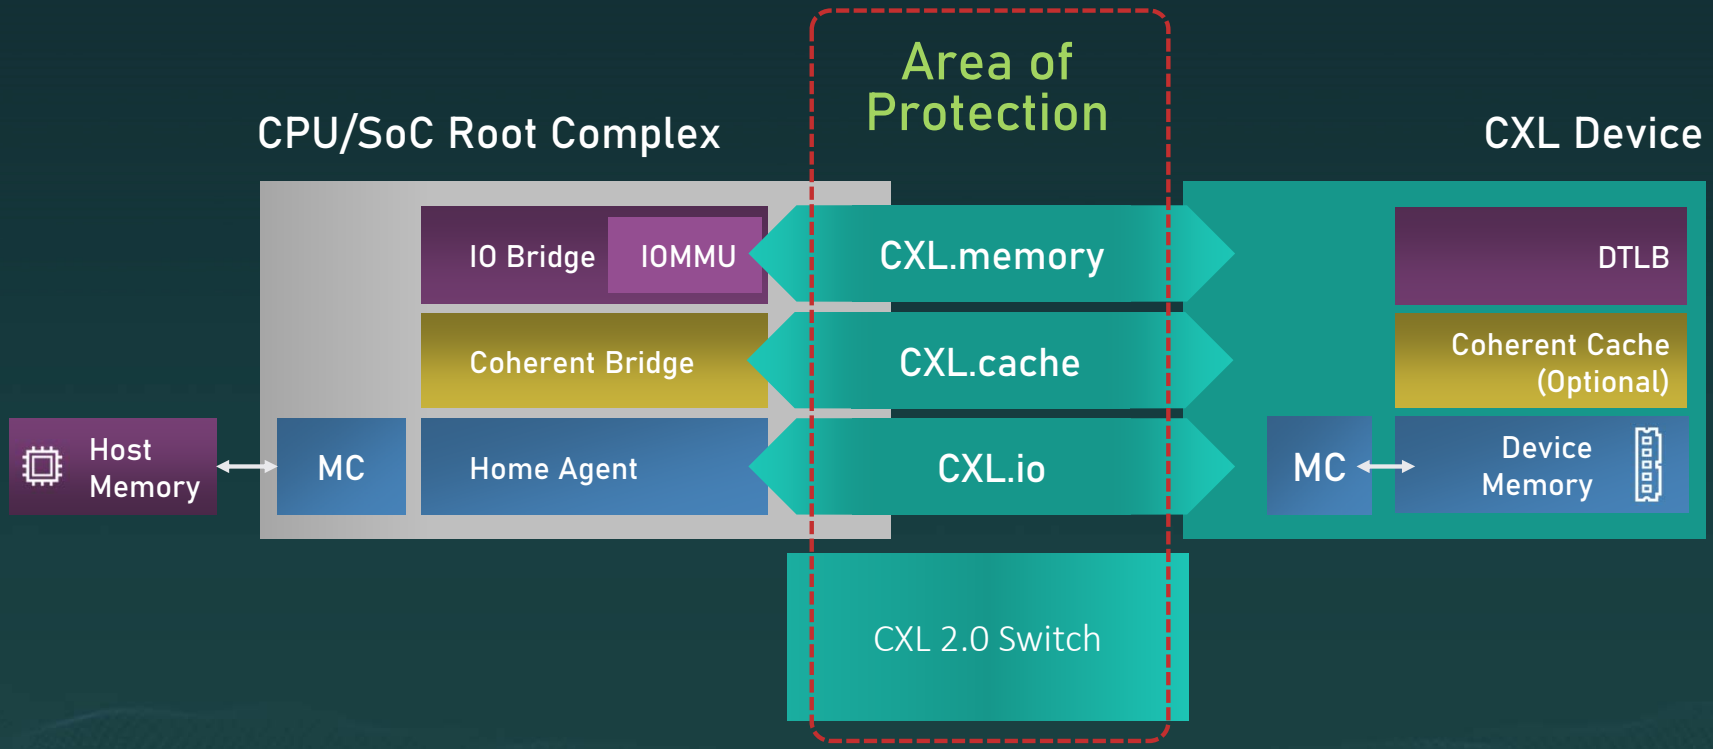

CXL 3.1 Trusted Security Protocol (TSP)

Allows for Virtualization-based, Trusted Execution Environments (TEEs) to host Confidential Computing Workloads (CC WL)

Key Capabilities:

- Separation between TVM* & CSP's infrastructure (VMM)
- Configuration of CXL device
- Encryption of sensitive data in both Host/Device memory
- Cryptographically verify correct configuration of trusted computing environment

Benefits:

- Freedom to migrate sensitive WLs to TSP-enabled Clouds
- Collaboration with multiple parties for sharing data
- Conform to Compliance & Data sovereignty programs
- Strengthen Application security & Software IP protection

*TVM = Trusted VM

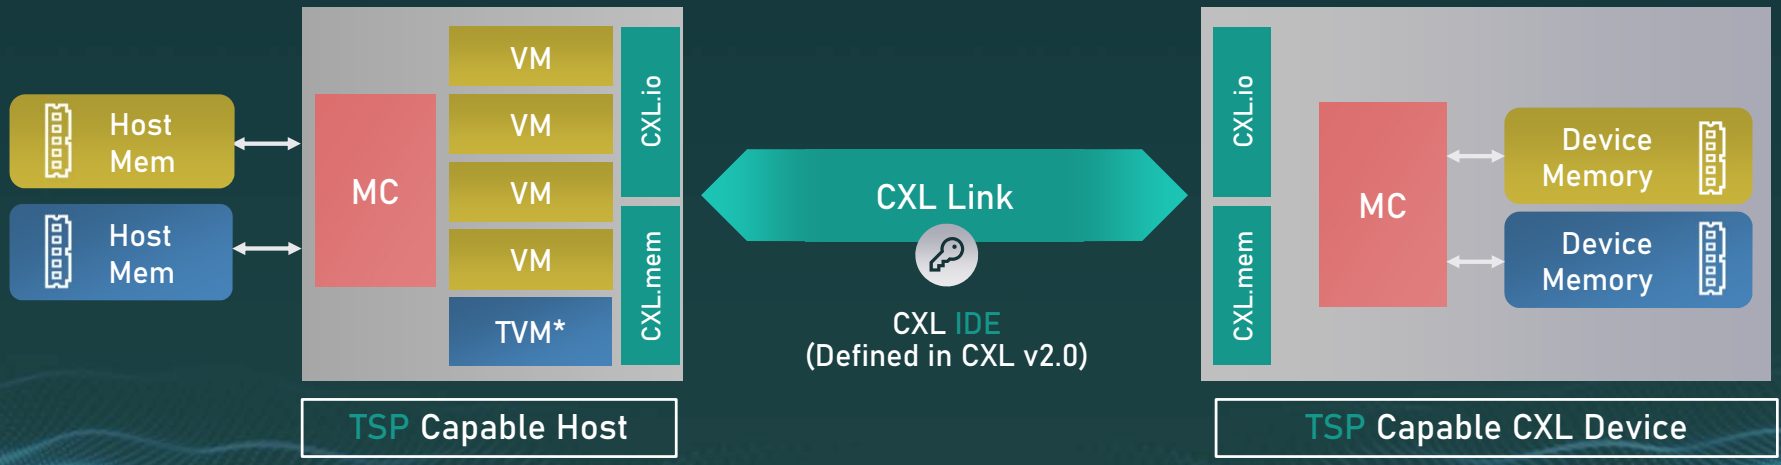

Elements of TSP / TSP Overview

TSP Components for Confidential Computing

- Trusted Execution State & Access Control
  - How access to memory is controlled
- Configuration
  - Ability to determine the supported security features on the device, enable required features, and lock the configuration
- Attestation & Authentication
  - Trusting who you are talking to
- Memory Encryption
  - Encrypting data-at-rest
- Transport Security
  - Encrypting the link to protect data-in-flight and detect/prevent physical attacks

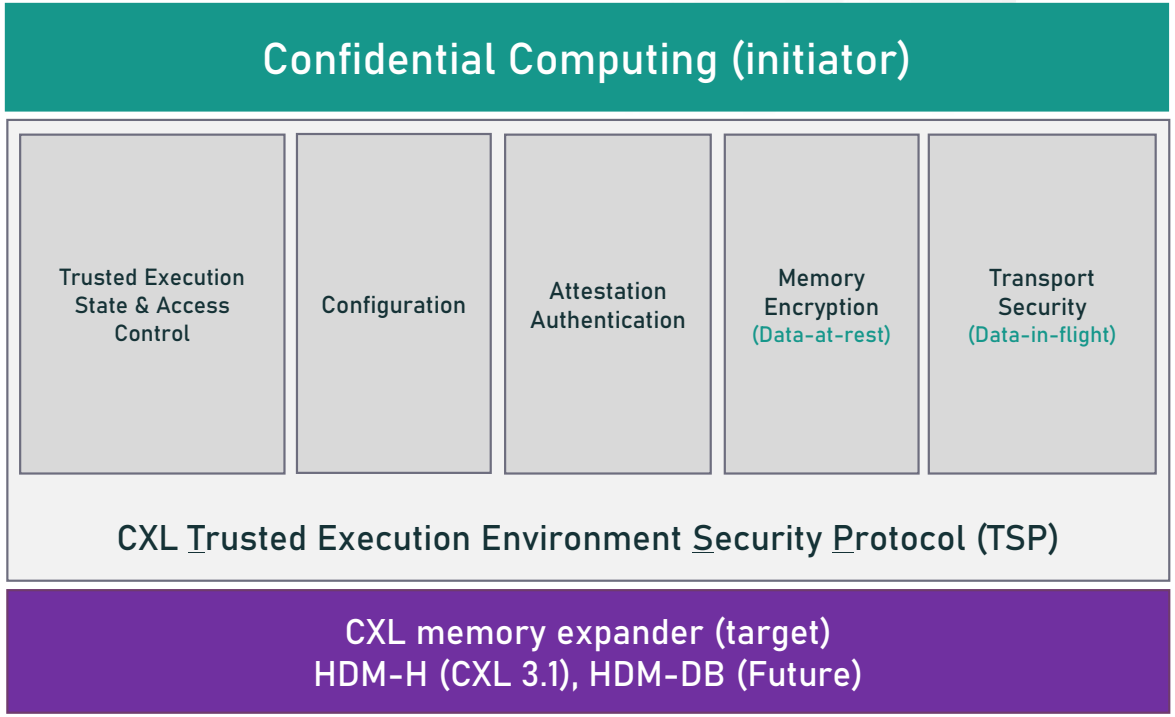

Memory Expander Improvements

Enhancements to CXL-Attached Memory

The CXL 3.1 specification provides for up to 32-bits of meta data per cacheline

CXL 3.1 supports 2-bits of meta-data for cache coherency that are used to describe if a cache line is shared, exclusive, modified, or invalid and up to 32-bits for additional meta-data

The extended meta data may be used with Host coherent and Device coherent memory

Extended meta data allows more information about the line to be available

Possible use cases

- Access control, data type tagging, and memory-tiering algorithms
- Supports proposed DDR6 feature that adds 16-32 bits of meta data for each 64B of data

Improved visibility into CXL memory device errors

Additional information on errors such as correctable error limits

New information on the source of the errors and the transaction occurring during the error condition

Added control over memory device RAS

Memory-sparing

DDR5 error-check-scrubbing

Media testing

Patrol-scrub

Capacity or performance degradation

Direct peer-to-peer CXL memory access for accelerators

CXL Specification Feature Summary

Features	CXL 1.0 / 1.1	CXL 2.0	CXL 3.0	CXL 3.1
Release date	2019	2020	August 2022	November 2023
Max link rate	32GTs	32GTs	64GTs	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)			✓	✓
Fabric Manager API definition for PBR Switch				✓
Host-to-Host communication with Global Integrated Memory (GIM) concept				✓
Trusted-Execution-Environment (TEE) Security Protocol				✓
Memory expander enhancements (up to 32-bit of meta data, RAS capability enhancements)				✓

Not Supported ✓ Supported

CXL 3.1 Summary

- The CXL specification continues to evolve to meet the usage models
- New features introduced in the CXL 3.1 specification:
  - CXL Fabric Improvements/Extensions
    - Scale-out of CXL fabrics using PBR (Port Based Routing)
  - Trusted-Execution-Environment Security Protocol (TSP)
    - Allows for Virtualization-based, Trusted Execution Environments (TEEs) to host Confidential Computing Workloads
  - Memory Expander Improvements
    - Up to 32-bit of meta-data and RAS capability enhancements

Call to Action

- Support future specification development by joining the CXL Consortium
- Download the CXL 3.1 Specification
- Follow us on social media for updates!

www.ComputeExpressLink.org

Q&A

Please share your questions in the Question Box

Thank You

www.ComputeExpressLink.org