



Compute Express Link™ (CXL™): A Coherent Interface for Ultra-High-Speed Transfers

DMTF Virtual APTS

July 21, 2020

Kurt Lender, Ecosystem Enabling Manager, Intel Corporation
Co-Chair Marketing Work Group, CXL Consortium

© CXL™ Consortium 2020

Meet the Presenter



Kurt Lender
Ecosystem Enabling Manager
Intel Corporation

Agenda

- Industry Landscape
- Compute Express Link™ Overview
- CXL™ Features and Benefits
- CXL Use Cases
- CXL Consortium and Industry Liaisons
- Summary


Industry Landscape

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

Proliferation of Cloud Computing



Growth of AI & Analytics

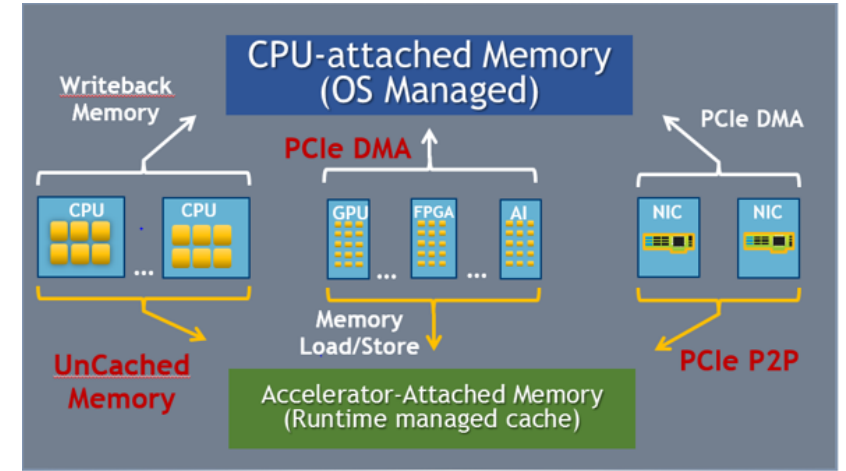


Cloudification of the Network & Edge

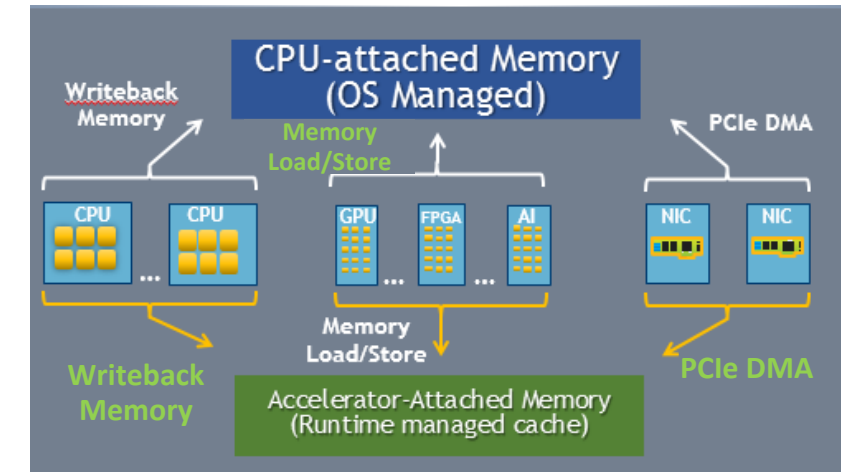


interconnect?

- **The industry needs open standards that can comprehensively address next-gen interconnect challenges**



Today's Environment



CXL-Enabled Environment

Compute Express Link™ (CXL™) Overview

- **New breakthrough high-speed CPU-to-Device interconnect**
 - Enables a high-speed, efficient interconnect between the CPU and platform enhancements and workload accelerators
 - Builds upon PCI Express® infrastructure, leveraging the PCIe® 5.0 physical and electrical interface
 - Maintains memory coherency between the CPU memory space and memory on attached devices
 - Allows resource sharing for higher performance
 - Reduced complexity and lower overall system cost
 - Permits users to focus on target workloads as opposed to redundant memory management
- **Delivered as an open industry standard**
 - CXL Specification 1.1 is available now
 - Future CXL Specification generations will continue to innovate to meet industry needs

CXL Consortium Board of Directors

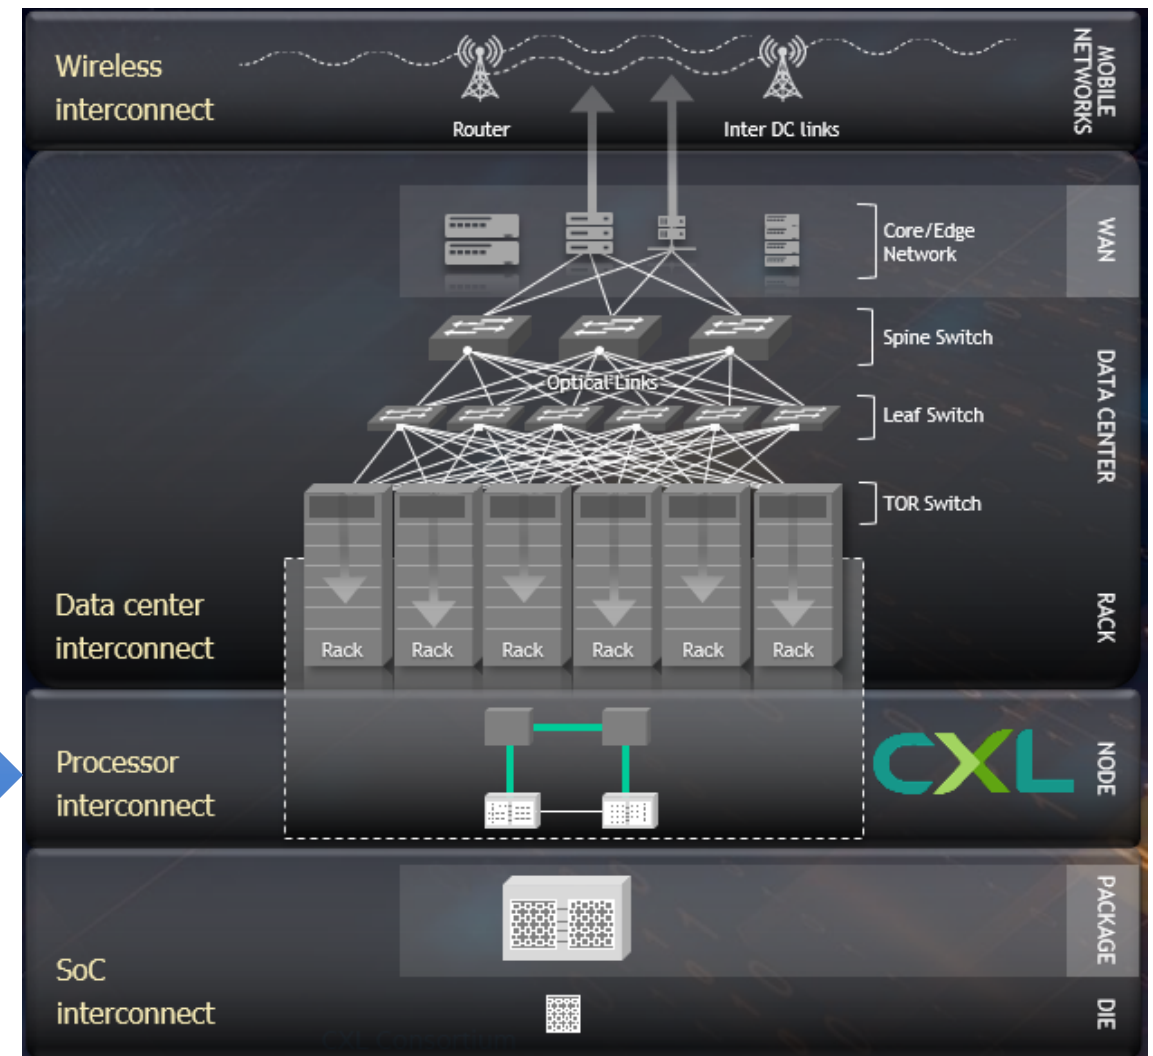
- Alibaba, Cisco, Dell EMC, Facebook, Google, Hewlett Packard Enterprise, Huawei, Intel Corporation and Microsoft [announced](#) their intent to incorporate in March 2019
- This core group [announced](#) incorporation of the Compute Express Link (CXL) Consortium on September 17, 2019 and unveiled the names of its Board of Directors:



Introducing CXL

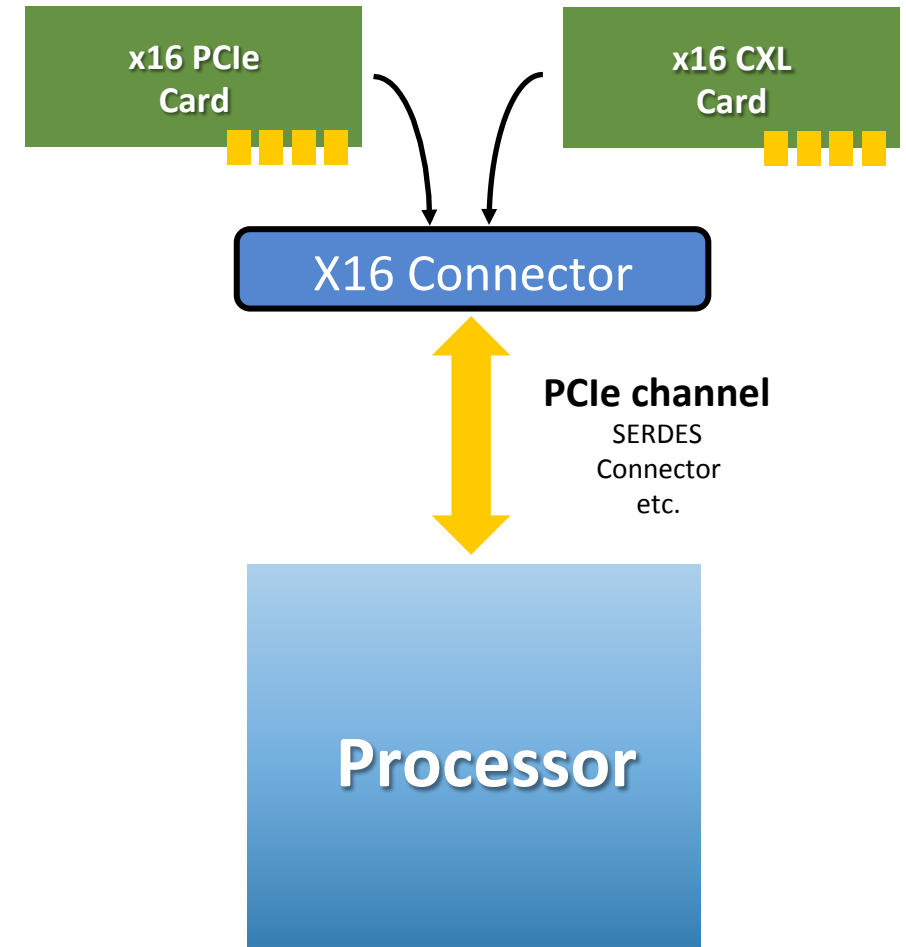
- Processor Interconnect:
 - Open industry standard
 - High-bandwidth, low-latency
 - Coherent interface
 - Leverages PCI Express®
 - Targets high-performance computational workloads
 - Artificial Intelligence
 - Machine Learning
 - HPC
 - Comms

A new class of
interconnect
for device
connectivity



What is CXL?

- Alternate protocol that runs across the standard PCIe physical layer
- Uses a flexible processor port that can auto-negotiate to either the standard PCIe transaction protocol or the alternate CXL transaction protocols
- First generation CXL aligns to 32 Gbps PCIe 5.0
- CXL usages expected to be key driver for an aggressive timeline to PCIe 6.0



CXL Protocols

- The CXL transaction layer is comprised of three dynamically multiplexed sub-protocols on a single link:

CXL.io

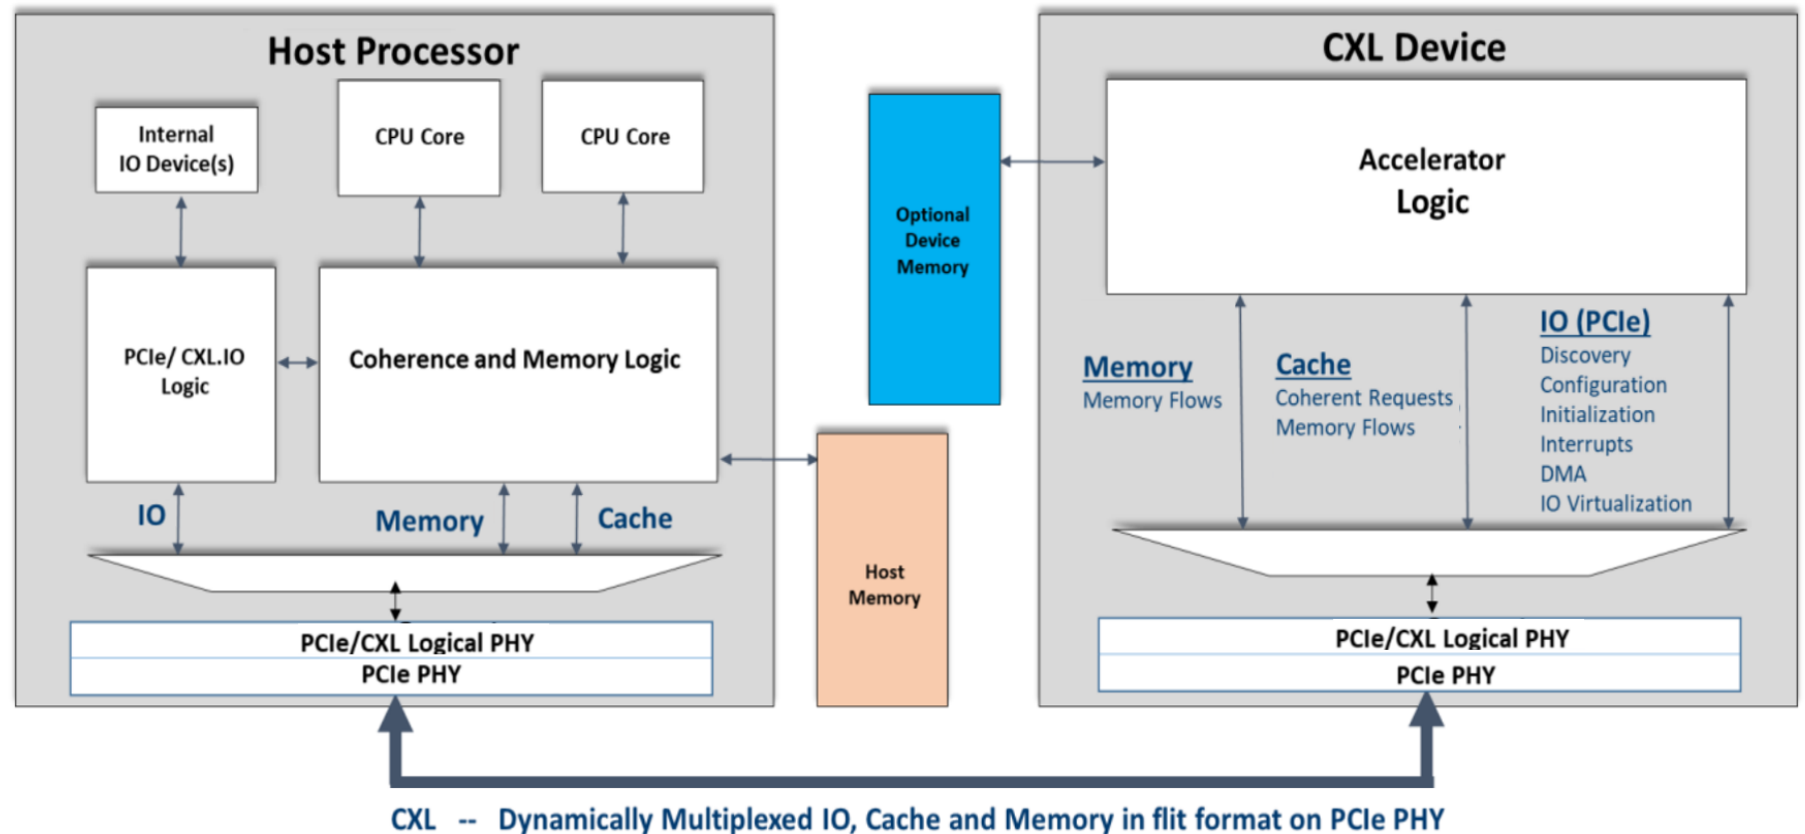
Discovery, configuration, register access, interrupts, etc.

CXL.cache

Device access to processor memory

CXL.Memory

Processor access to device attached memory



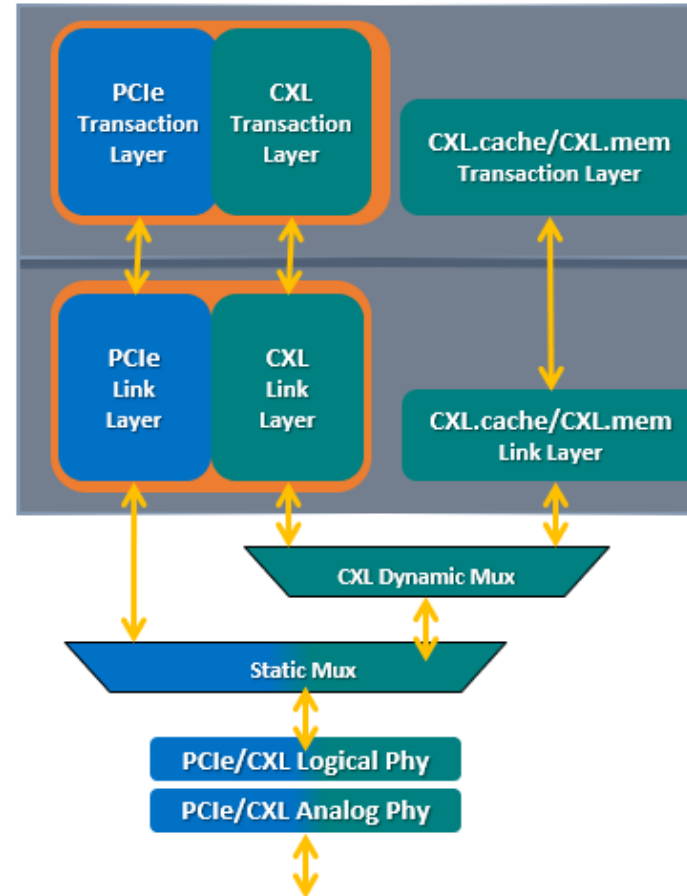


CXL™ Features and Benefits

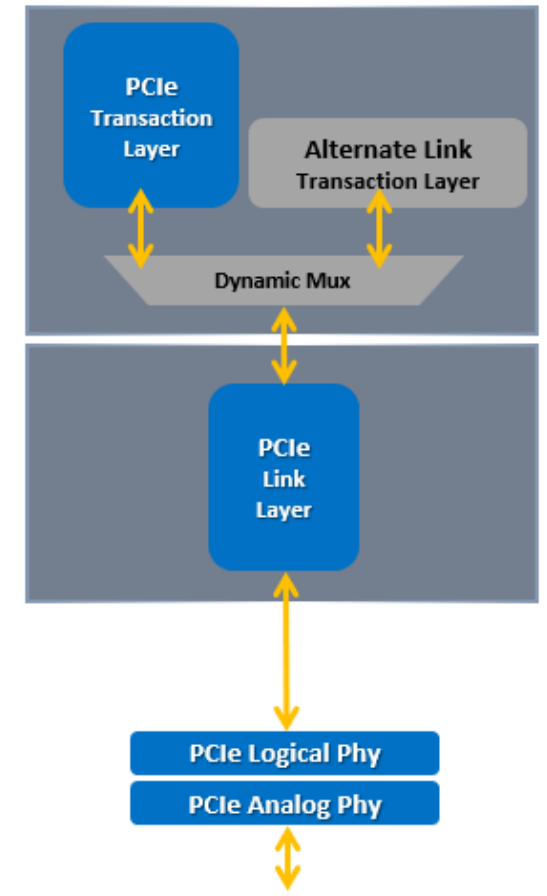
CXL Stack – Designed for Low Latency

- All 3 representative usages have latency critical elements:
 - CXL.Cache
 - CXL.Memory
 - CXL.io
- CXL cache and memory stack is optimized for latency:
 - Separate transaction and link layer from IO
 - Fixed message framing
- CXL io flows pass through a stack that is largely identical a standard PCIe stack:
 - Dynamic framing
 - Transaction Layer Packet (TL)/Data Link Layer Packet (DLPP) encapsulated in CXL flits

CXL Stack –
Low latency Cache and Mem Transactions

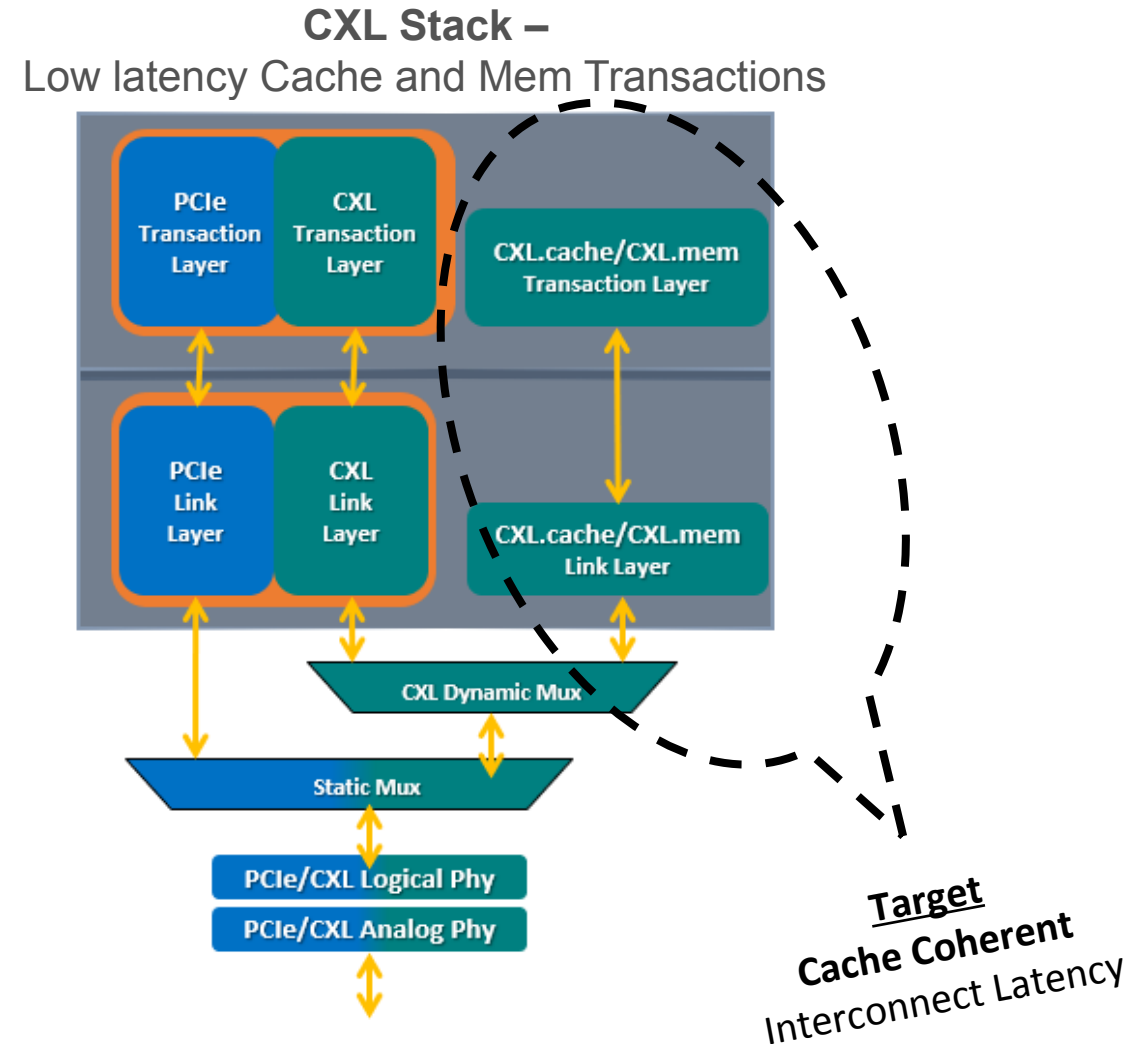


Alternate Stack –
for contrast



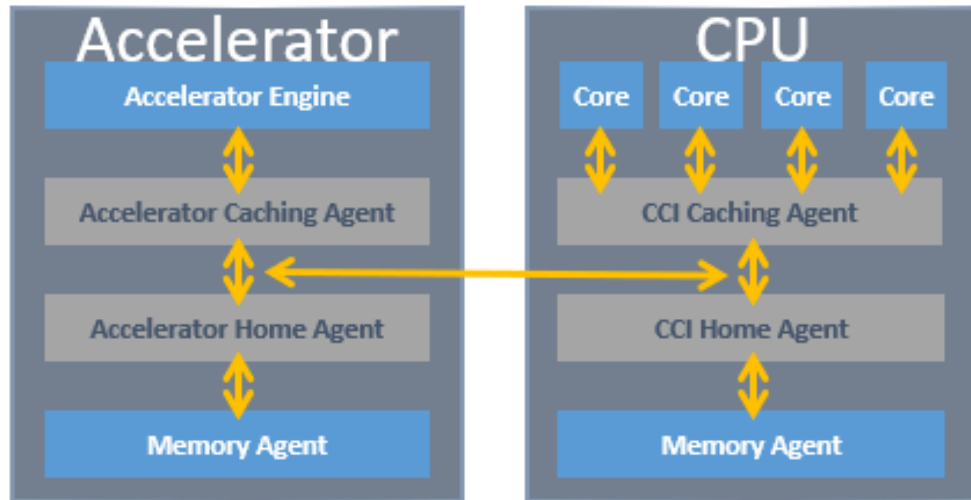
CXL Stack – Designed for Low Latency

- All 3 representative usages have latency critical elements:
 - CXL.Cache
 - CXL.Memory
 - CXL.io
- CXL cache and memory stack is optimized for latency:
 - Separate transaction and link layer from IO
 - Fixed message framing
- CXL io flows pass through a stack that is largely identical a standard PCIe stack:
 - Dynamic framing
 - Transaction Layer Packet (TLP)/Data Link Layer Packet (DLLP) encapsulated in CXL flits



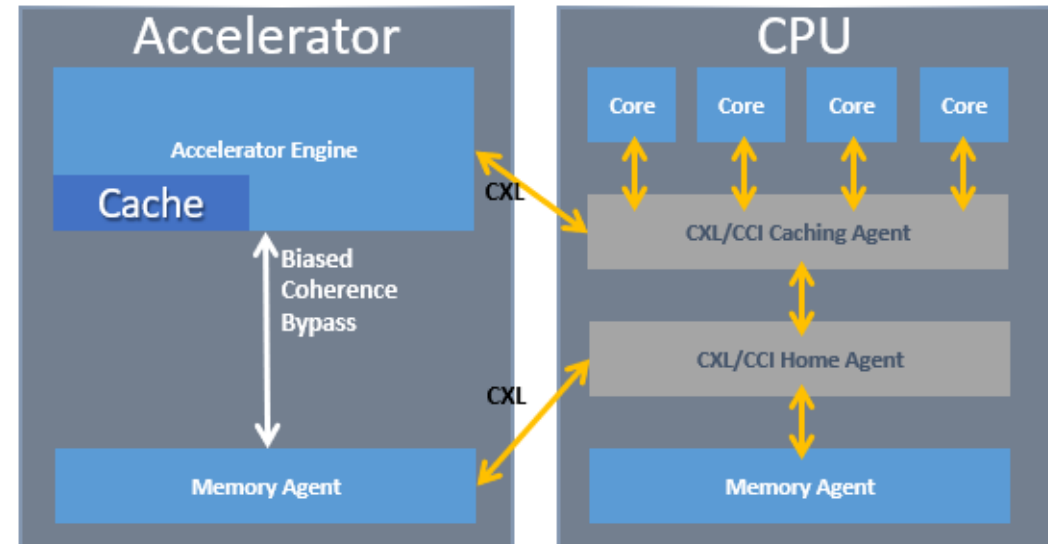
Asymmetric Complexity

CCI* Model – Symmetric CCI Protocol



*Cache Coherent Interface

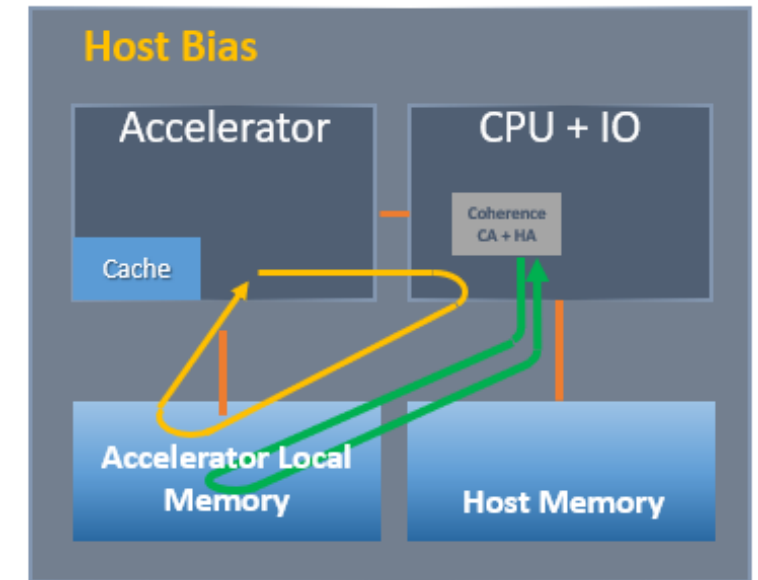
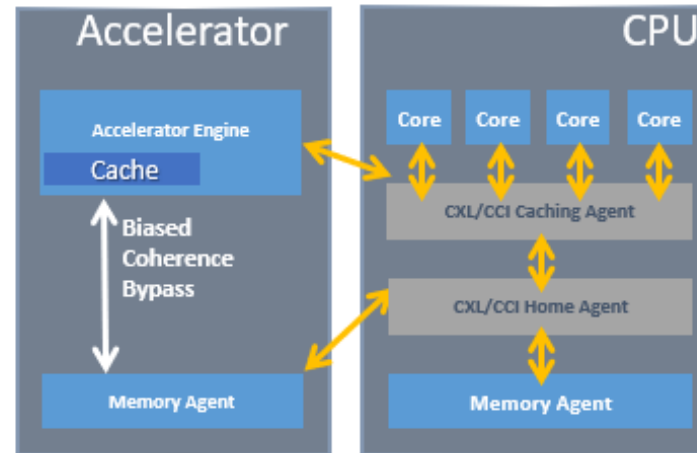
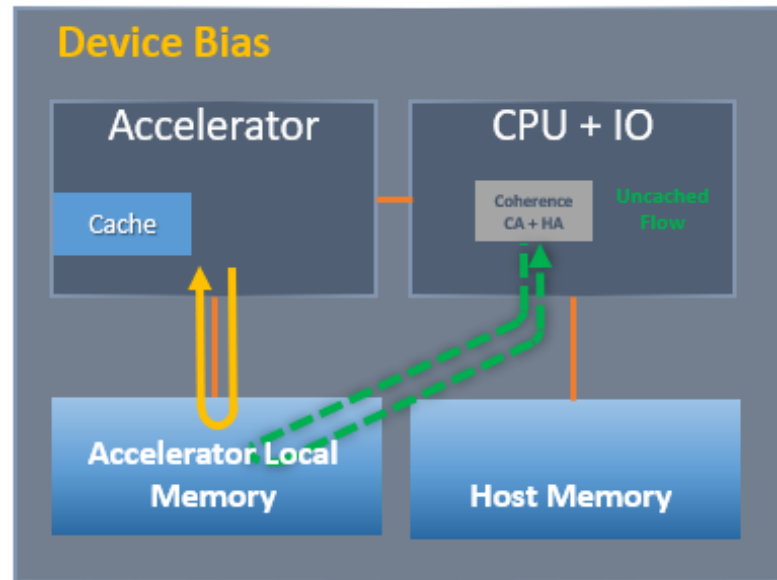
CXL Model – Asymmetric Protocol



CXL Key Advantages:

- Avoid protocol interoperability hurdles/roadblocks
- Enable devices across multiple segments (e.g. client / server)
- Enable Memory buffer with no coherency burden
- Simpler, processor independent device development

CXL's Coherence Bias



Critical access class for accelerators is “device engine to device memory”

“Coherence Bias” allows a device engine to access its memory coherently without visiting the processor

Two driver managed modes or “Biases”

HOST BIAS: pages being used by the host or shared between host and device

DEVICE BIAS: pages being used exclusively by the device

Both biases guaranteed correct/coherent

Guarantee applies even when software bugs or speculative accesses unexpectedly access device memory in the “Device Bias” state.

CXLTM Use Cases

Representative CXL Usages


Caching Devices / Accelerators

Usages:

- PGAS NIC
- NIC atomics

Protocols:

- CXL.io
- CXL.cache



Accelerators with Memory

Usages:

- GPU
- Dense Computation

Protocols:

- CXL.io
- CXL.cache
- CXL.memory




Memory Buffers

Usages:

- Memory BW expansion
- Memory capacity expansion
- Storage Class Memory

Protocols:

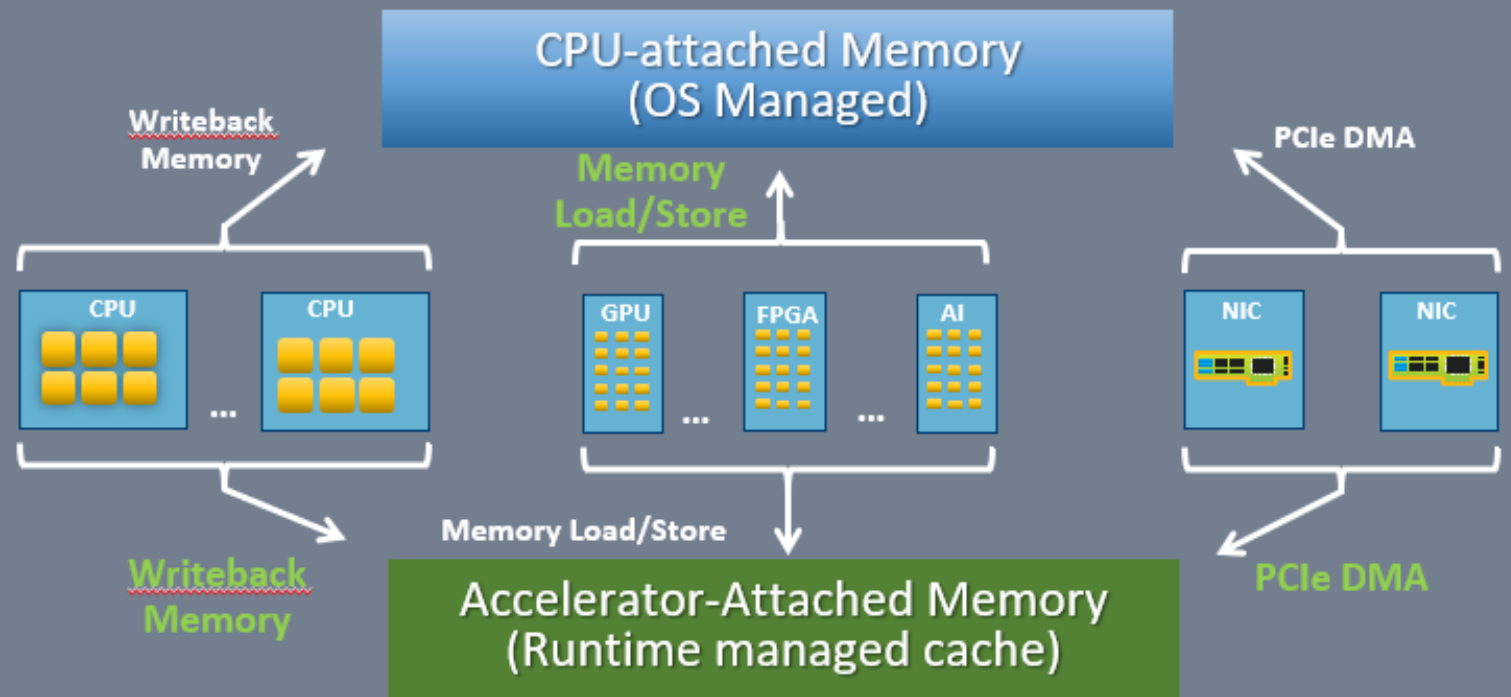
- CXL.io
- CXL.mem



Heterogeneous Computing Revisited – with CXL

- CXL enables a more fluid and flexible memory model
- Single, common, memory address space across processors and devices

- More efficient population and update of operands
- More efficient extraction of results
- Memory resource “borrowing”
- User/Kernel level data access and data movement
- Low latency to memory, host to device and device to host





CXL Consortium

CXL Consortium™ Membership

- CXL Consortium boasts 100+ member companies to date and is growing rapidly
 - Membership reflects required industry expertise to create a robust, vibrant CXL ecosystem
 - View the [List of Members](#)
- Members have immediate access to the CXL Specification 1.1
 - Both the Host and Target side of the interface is published, allowing it to be implemented on any type of system and with any type of target device
 - All members can implement the spec under the Consortium's IP protection policy
 - Evaluation Copy of the CXL 1.1 Specification is available for [download](#)
- The CXL Consortium will continue to define and deliver future generations of the CXL Specification
 - Contributor level members and above can participate in the definition and promotion of future specifications in the following CXL Working Groups:
 - Compliance WG, Marketing WG, PHY WG, Protocol WG, System WG and Software WG
 - Will maintain backwards compatibility with prior generations to protect member investments

Industry Liaison – DMTF



- Work register established between CXL Consortium and DMTF
- Areas of Technical Collaboration:
 - CXL is built on top of PCIe infrastructure, so CXL can leverage all PMCI standards that apply to PCIe adapters. CXL will adopt DMTF defined SPDM standard
 - Assist in extending the Redfish standard to include CXL management
 - Provide CXL expertise and assistance to the Redfish Forum
 - Current definition of Redfish data model is rich enough to describe CXL accelerators and memory expanders for the most parts, CXL will leverage.
 - Ensure CXL management support in standards developed by DMTF's Platform Management Components Intercommunication working group
 - CXL consortium will collaborate with DMTF to define extensions (e.g. new MCTP message type, new properties) as needed

CXL Summary

- CXL has the right features and architecture to enable a broad, open ecosystem for heterogeneous computing and server disaggregation:

Coherent Interface:

Leverages PCIe® with 3 mix-and-match protocols

Low Latency:

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

Asymmetric Complexity:

Eases burdens of cache coherent interface designs

Open Industry Standard:

With growing broad industry support

Call to Action

- To join the CXL Consortium, visit www.computeexpresslink.org/join
- If your company is a member, consider joining various workgroups and contribute to future generation of CXL.
- Download an evaluation copy of the CXL 1.1 specification
- Engage with us on social media



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



CXL Consortium Channel

CXL Resources



Webinars:

Upcoming webinar: [Memory Challenges and CXL Solutions](#) (August 6, 8 AM PT)

Webinar: [Exploring Coherent Memory and Innovative Use Cases](#)

Webinar: [Introduction to Compute Express Link \(CXL\) Webinar Presentation](#)



CXL Blogs:

- [Compliance and Interoperability: Critical Indicators of Technology Success](#)
- [The Benefits of Serial-Attached Memory with Compute Express Link™](#)
- [Questions from the Compute Express Link™ Exploring Coherent Memory and Innovative Use Cases Webinar](#)



Whitepaper:

- [Introduction to Compute Express Link™](#)

www.ComputeExpressLink.org



Thank you!