



BY Developers FOR Developers

Storage Developer Conference
September 22-23, 2020

Understanding Compute Express Link™: A Cache- coherent Interconnect

Debendra Das Sharma
Intel Corporation



Industry Landscape

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

Proliferation of Cloud Computing



Growth of AI & Analytics

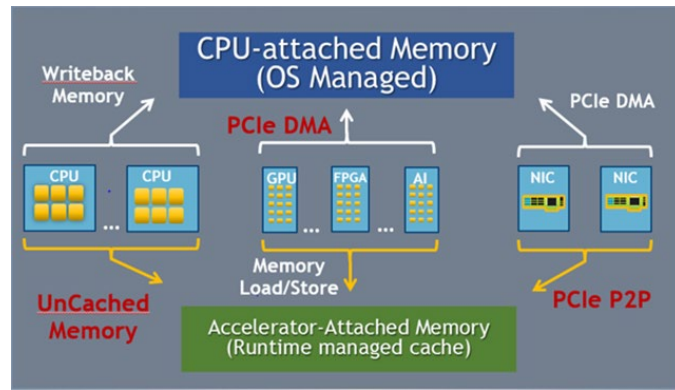


Cloudification of the Network & Edge

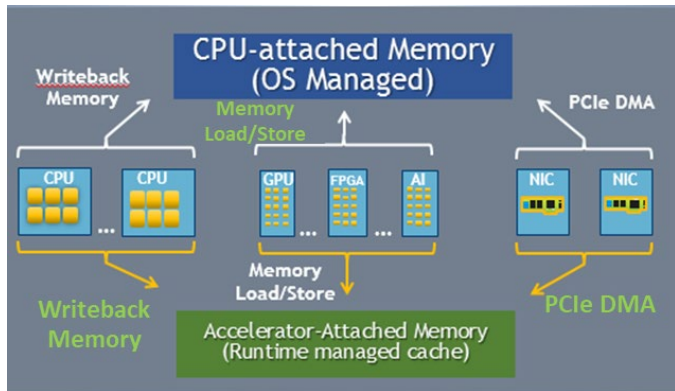


Why the need for a new class of interconnect?

- Need a new class of interconnect for heterogenous computing and disaggregation usages:
 - Efficient resource sharing
 - Shared memory pools with efficient access mechanisms
 - Enhanced movement of operands and results between accelerators and target devices
 - Significant latency reduction to enable disaggregated memory
- The industry needs open standards that can comprehensively address next-gen interconnect challenges



Today's Environment



CXL-Enabled Environment

CXL Consortium Overview

- CXL Consortium boasts 100+ member companies to date and is growing rapidly
 - Current membership ranks reflects the required industry expertise to create a robust, vibrant CXL ecosystem
- CXL Consortium Work Groups:
 - 5 Technical (Protocol, PHY, Software & Systems, Memory, Compliance) and Marketing
- CXL Board of Directors:

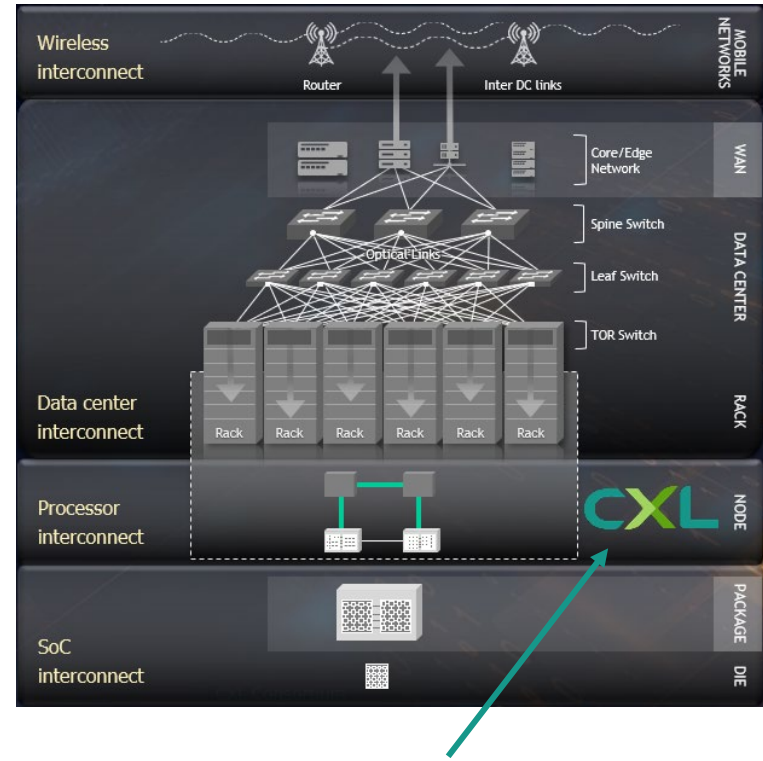




Overview of Compute Express Link™

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?

- Enhanced, substitute protocol that runs across 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

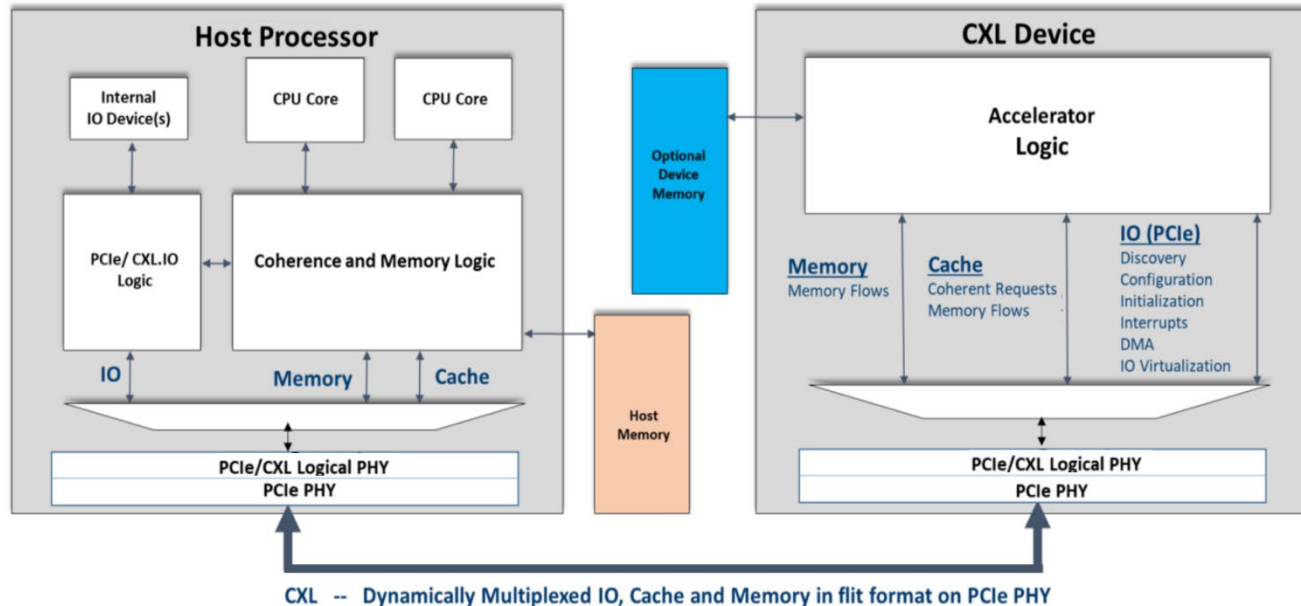
Discovery, configuration, register access, interrupts, etc.

CXL.cache

Device access to processor memory

CXL.Memory

Processor access to device attached memory



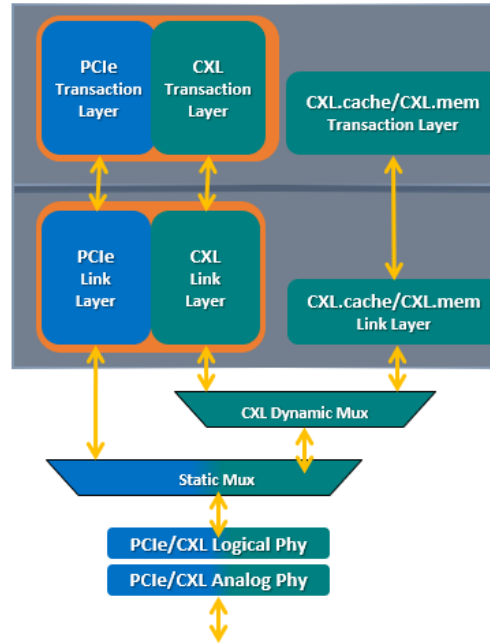


CXL™ Features and Benefits

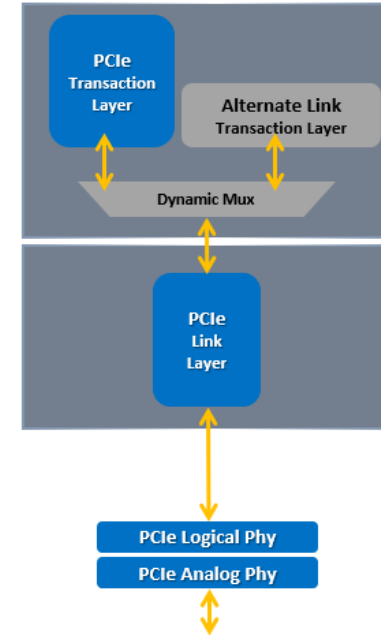
CXL Stack – Designed for Low Latency

- 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

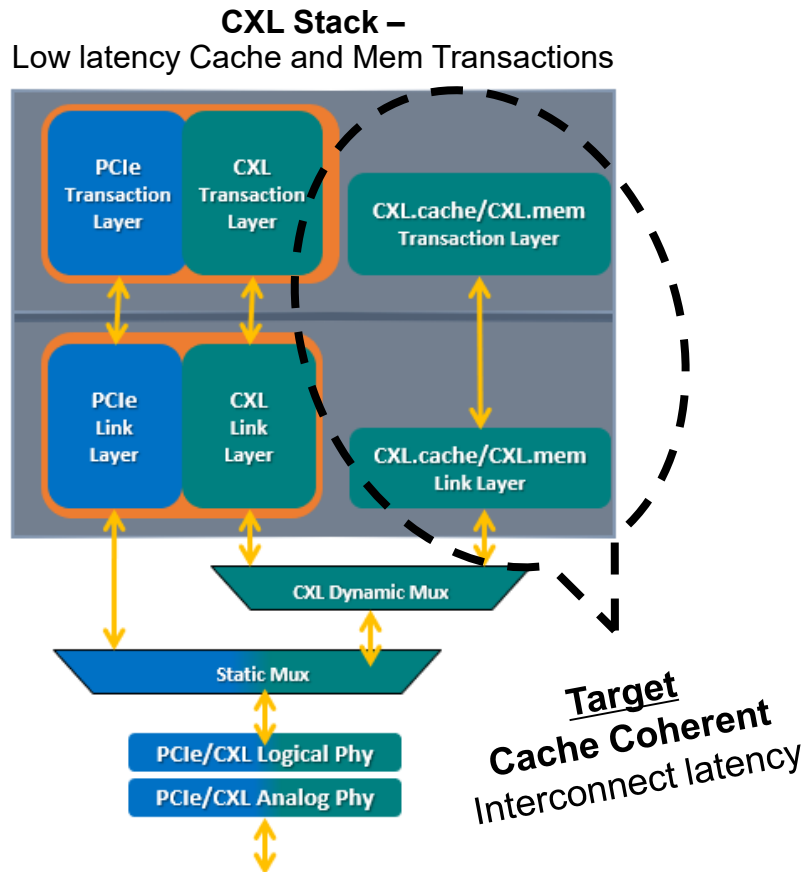


Alternate Stack –
for contrast



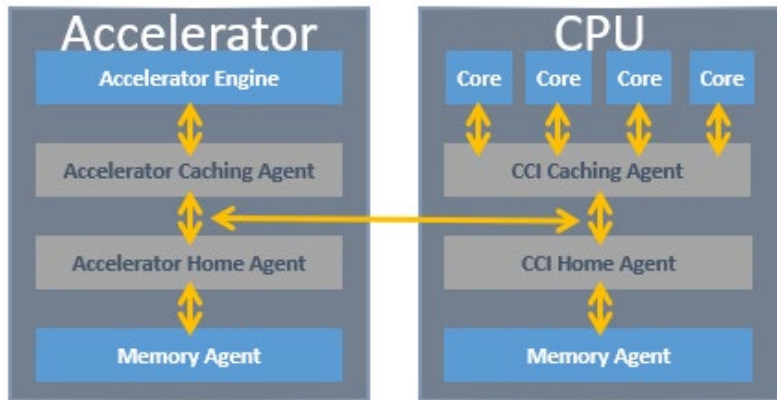
CXL Stack – Designed for Low Latency

- 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:



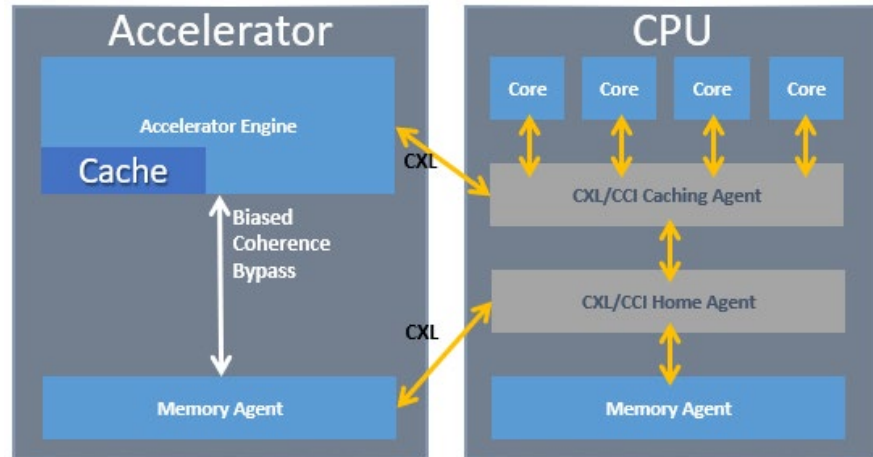
Asymmetric Complexity

CCI* Model – Symmetric CCI Protocol



*Cache Coherent Interface

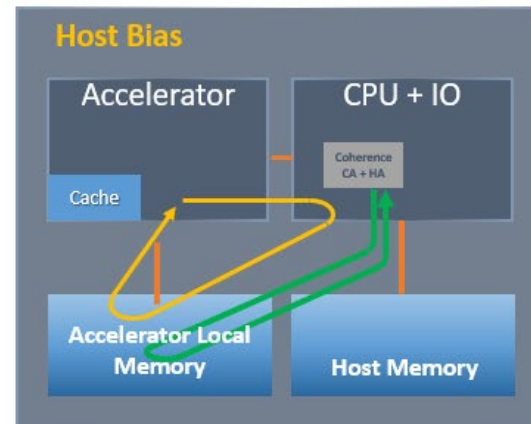
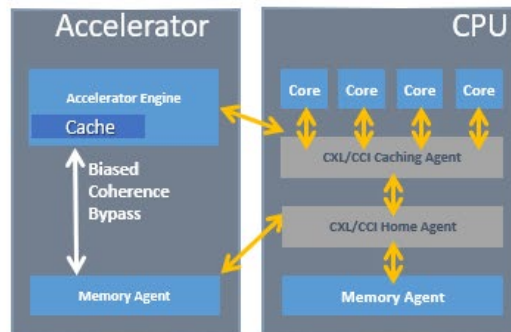

CXL Model – Asymmetric Protocol



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 Coherency 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.



CXL™ 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:

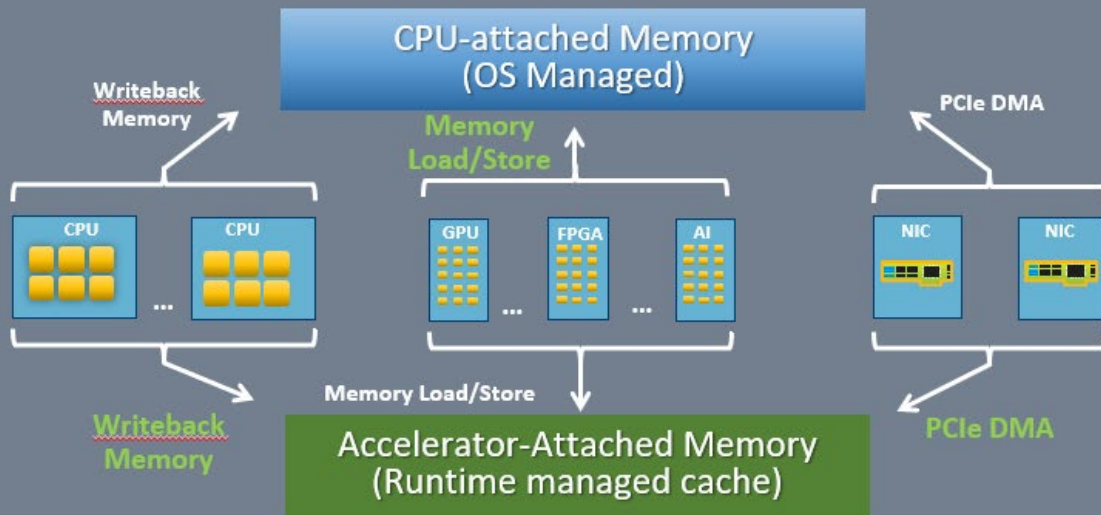
- CXL.io
- CXL.mem



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 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



Thank you!



**Please take a moment
to rate this session.**

Your feedback matters to us.