



2020 OFA Virtual Workshop

COMPUTE EXPRESS LINK™ (CXL™): A COHERENT INTERFACE FOR ULTRA-HIGH-SPEED TRANSFERS

Jim Pappas, Director, Industry Initiatives, Intel Corporation
Chairman, CXL Consortium



AGENDA

- **Industry Landscape**
- **Compute Express Link™ Overview**
- **Introducing CXL™ Consortium**
- **CXL Features and Benefits**
- **CXL Use Cases**
- **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

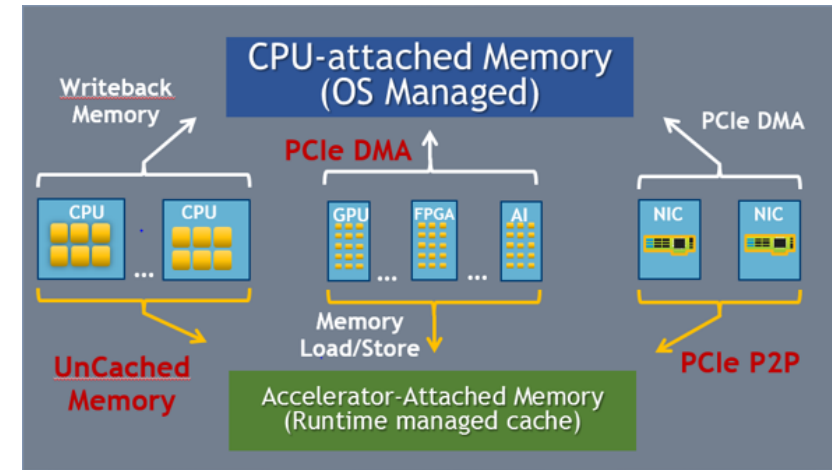


Cloudification of the
Network & Edge

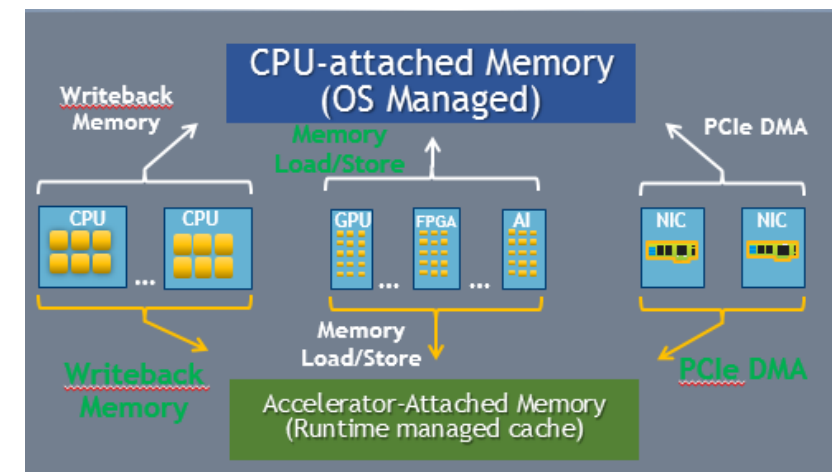


WHY THE NEED FOR A NEW CLASS OF INTERCONNECT?

- **Industry mega-trends are driving demand for faster data processing and next-generation data center performance:**
 - Proliferation of Cloud Computing
 - Growth of Artificial Intelligence and Analytics
 - Cloudification of the Network and Edge
- **Need a new class of interconnect for heterogeneous 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

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:

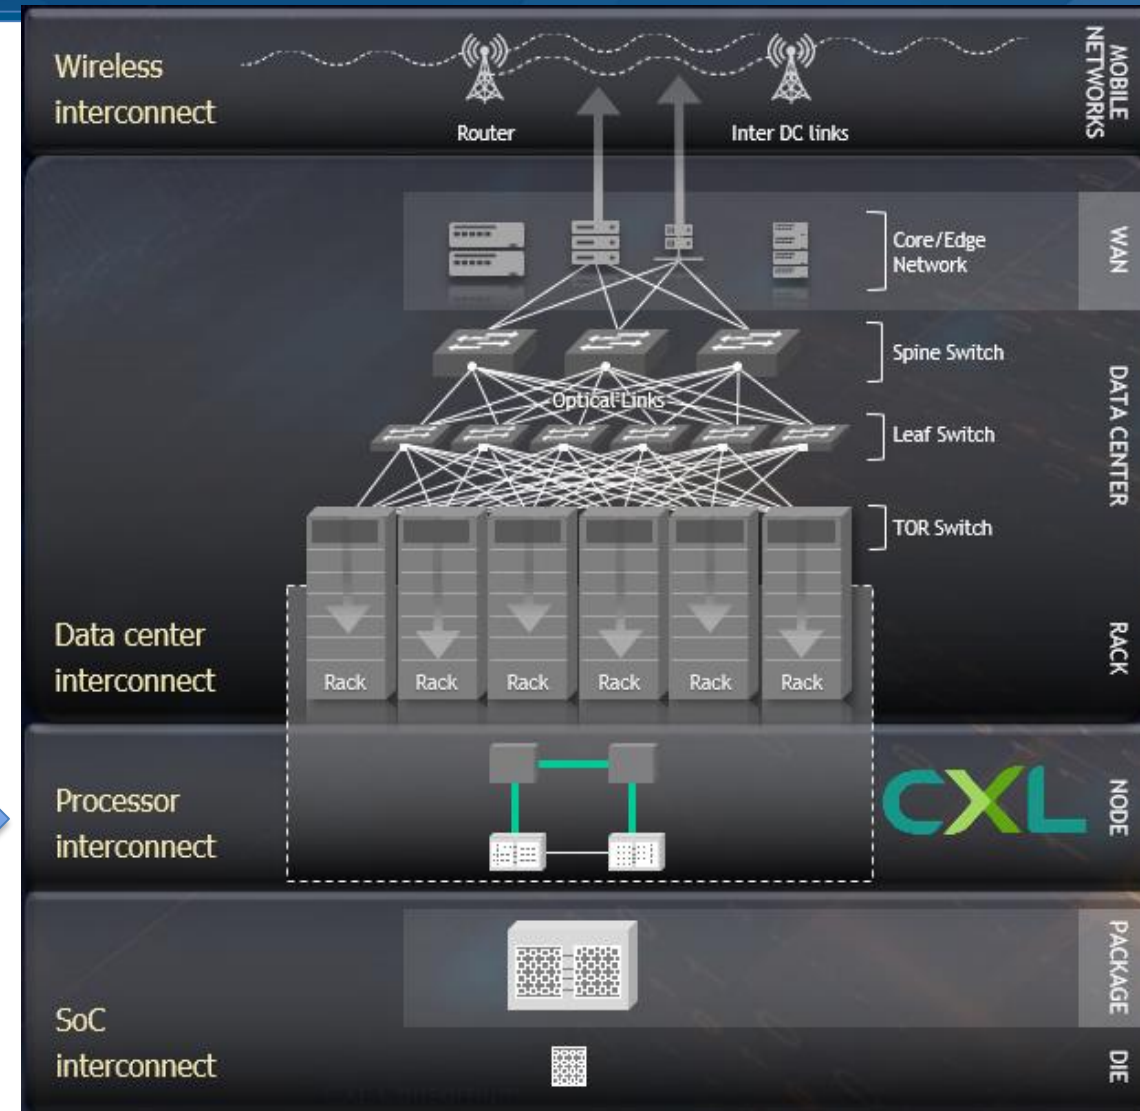



INTRODUCING CXL

■ Processor Interconnect:

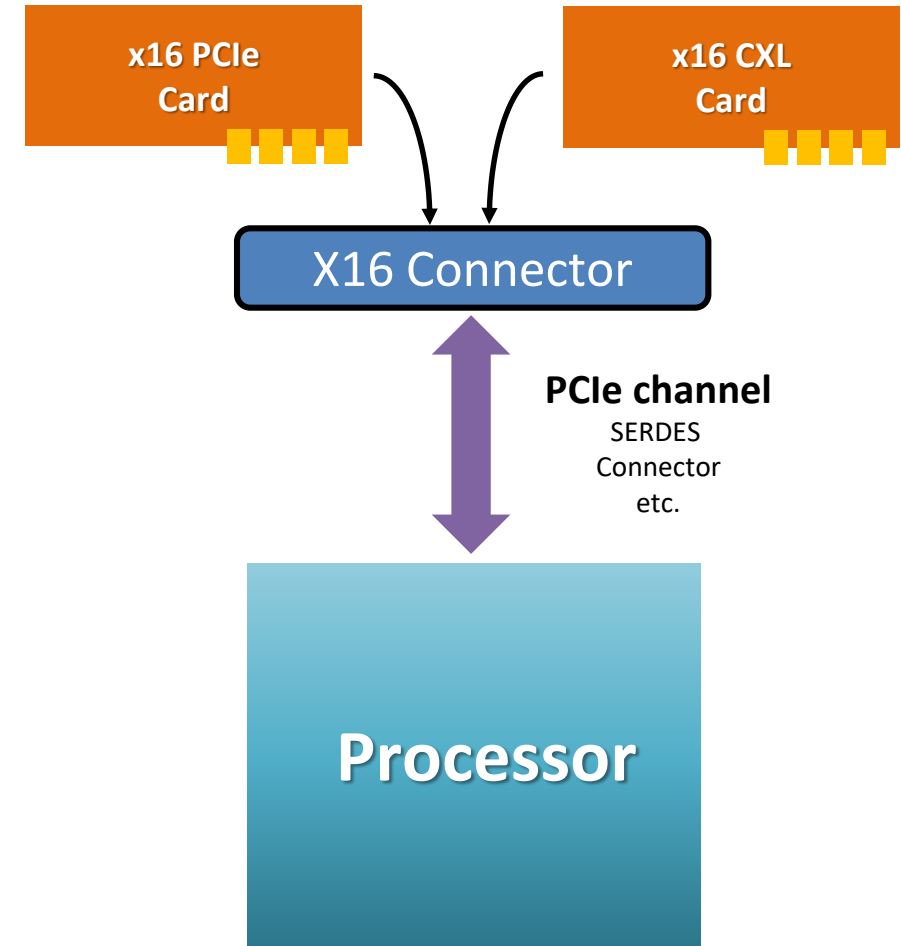
- Open industry standard
- High-bandwidth, low-latency
- Coherent interface
- Leverages PCI Express®
- Targets high-performance computational workloads
 - Artificial Intelligence
 - Machine Learning
 - HPC
 - Comms

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



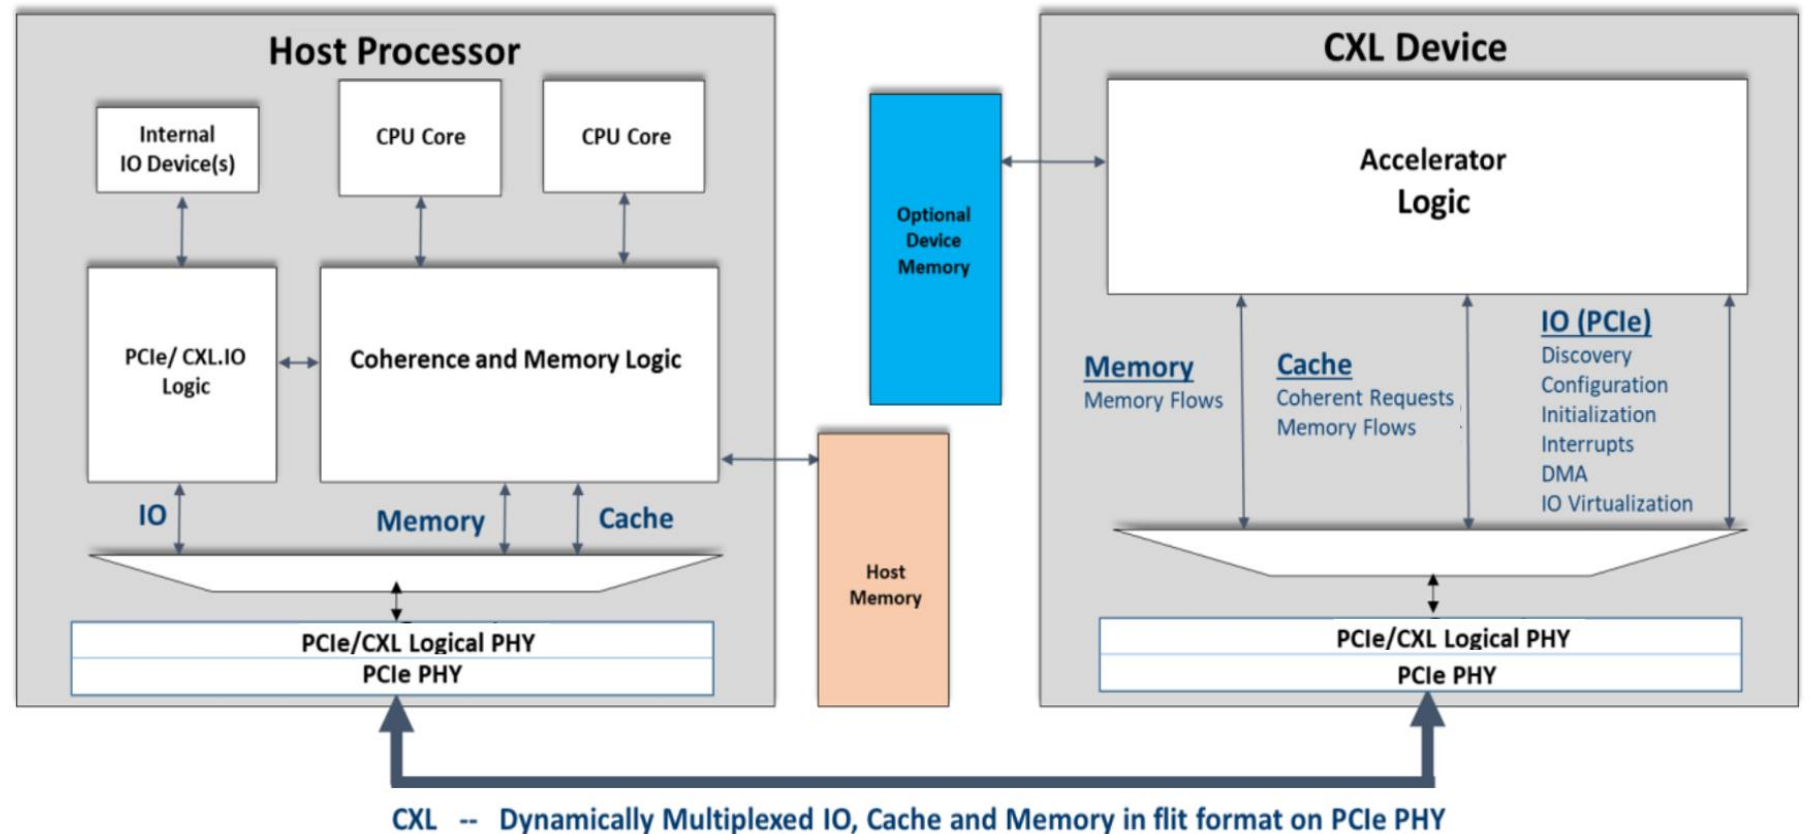
CXL PROTOCOLS

- 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



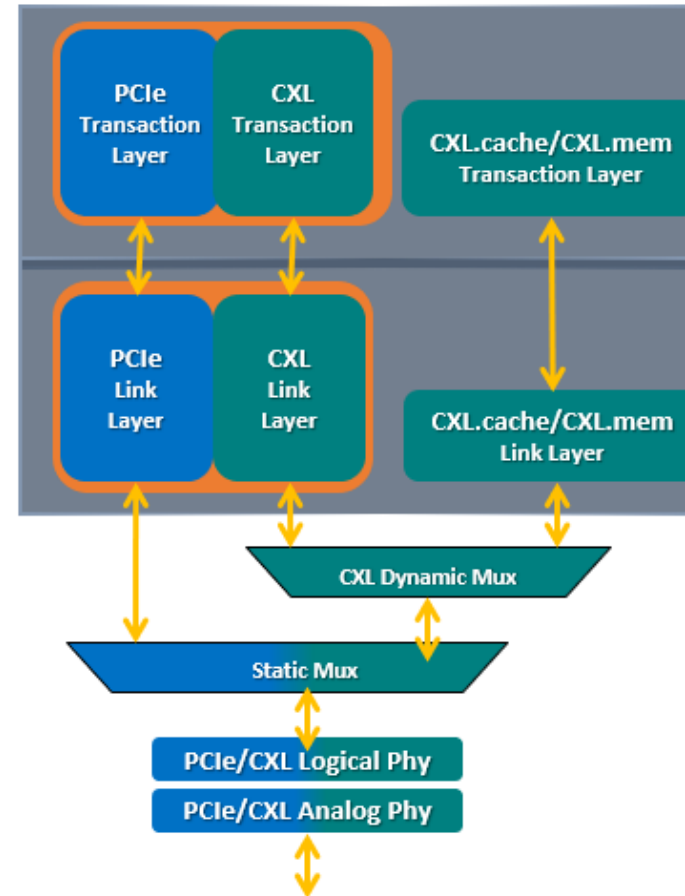


CXL FEATURES AND BENEFITS

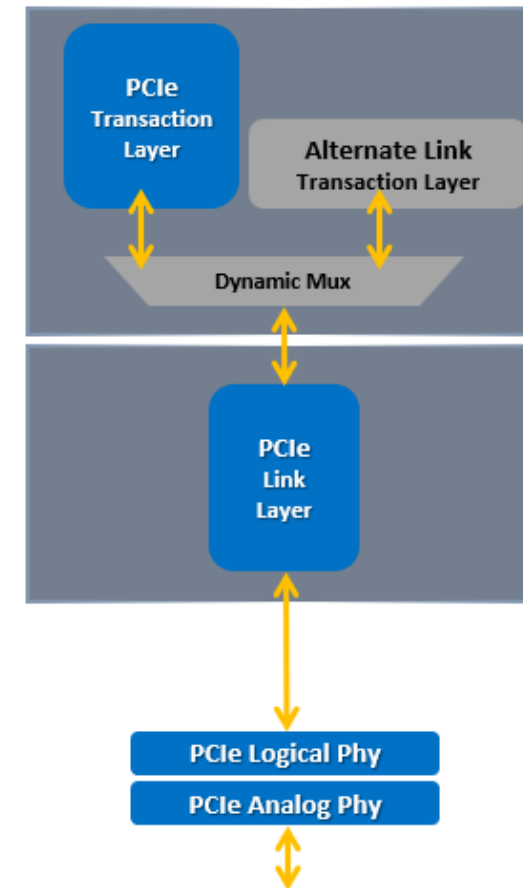
CXL STACK – DESIGNED FOR LOW LATENCY

- **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 (TLP)/Data Link Layer Packet (DLLP) encapsulated in CXL flits

CXL Stack –
Low latency Cache and Mem Transactions

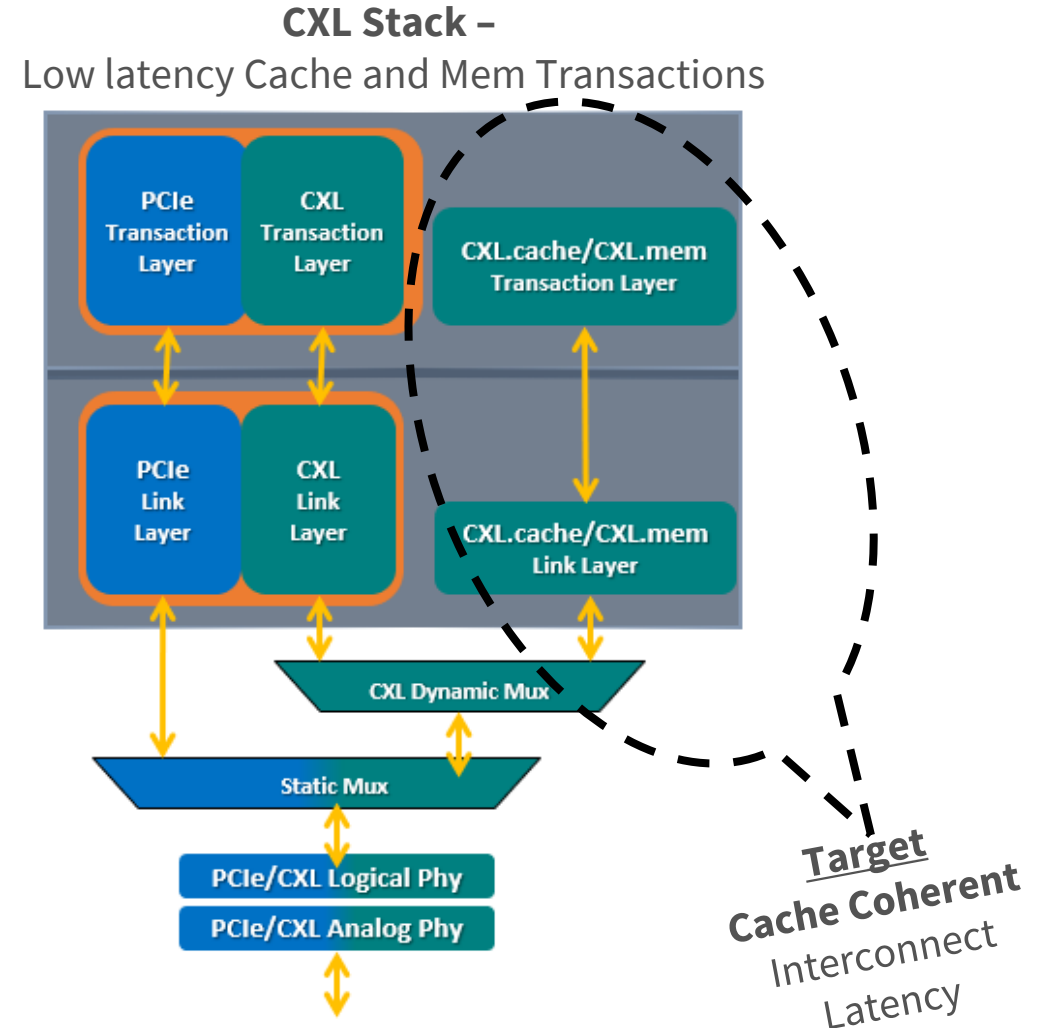


Alternate Stack –
for contrast



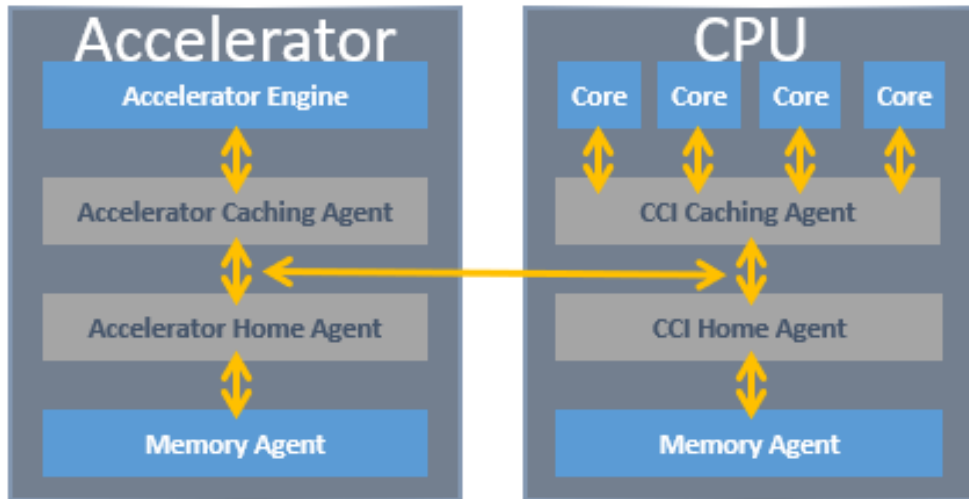
CXL STACK – DESIGNED FOR LOW LATENCY

- **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



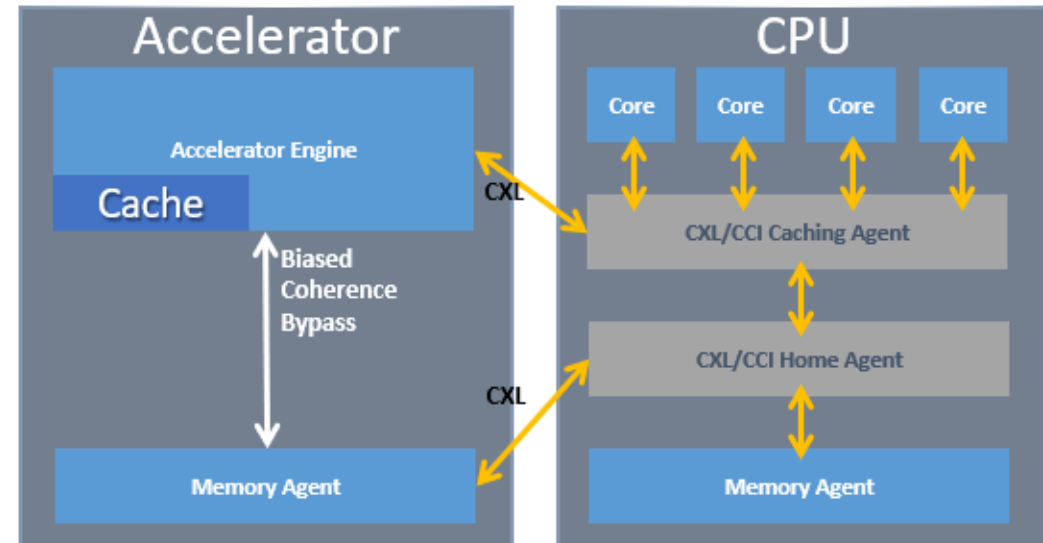
ASYMMETRIC COMPLEXITY

CCI* Model – Symmetric CCI Protocol



*Cache Coherent Interface

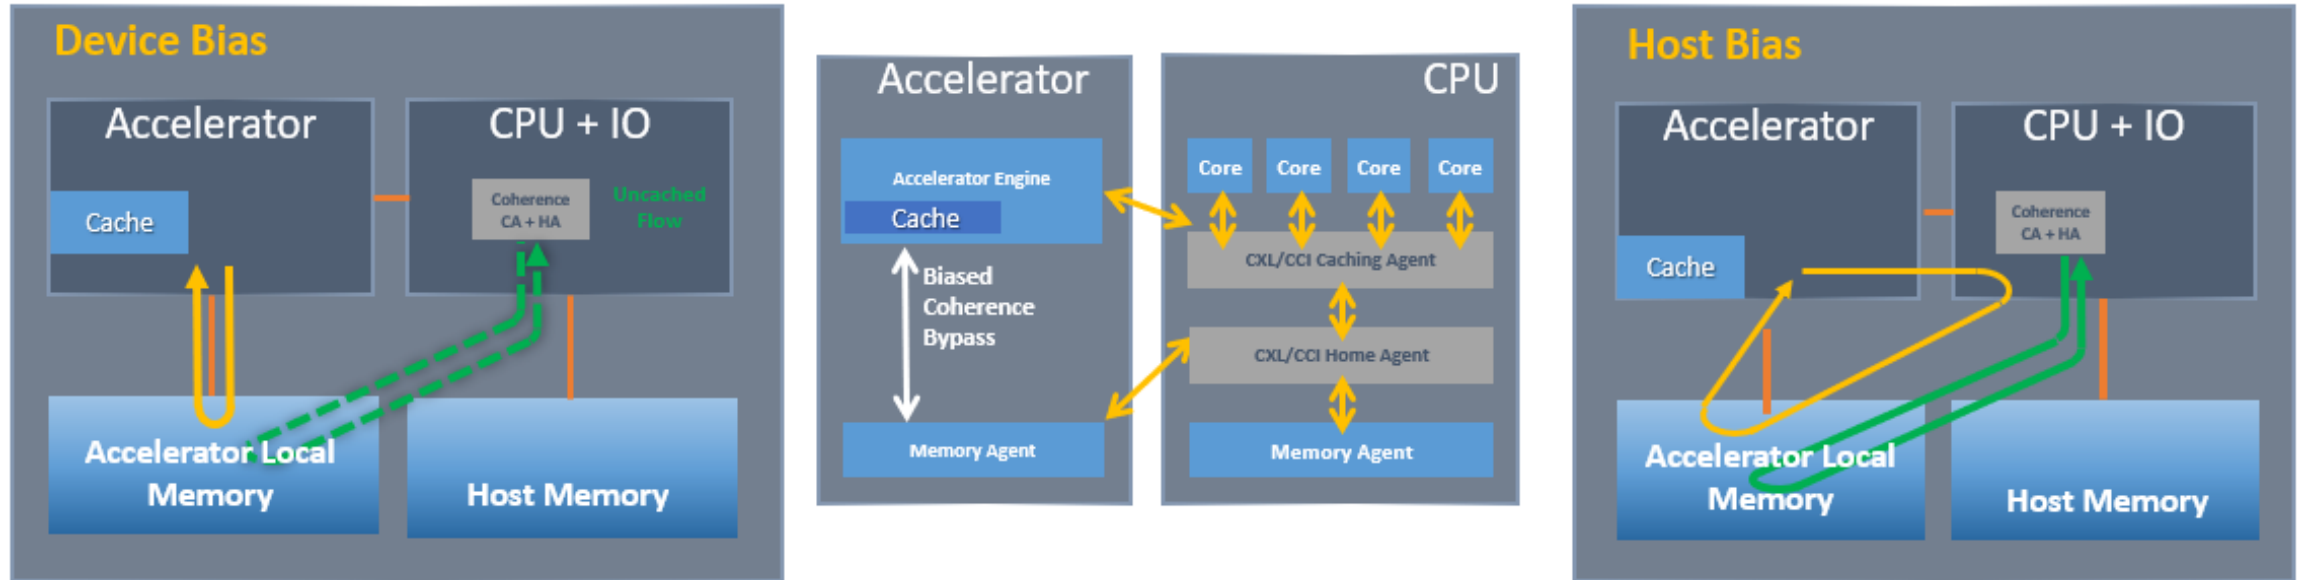
CXL Model – Asymmetric Protocol



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

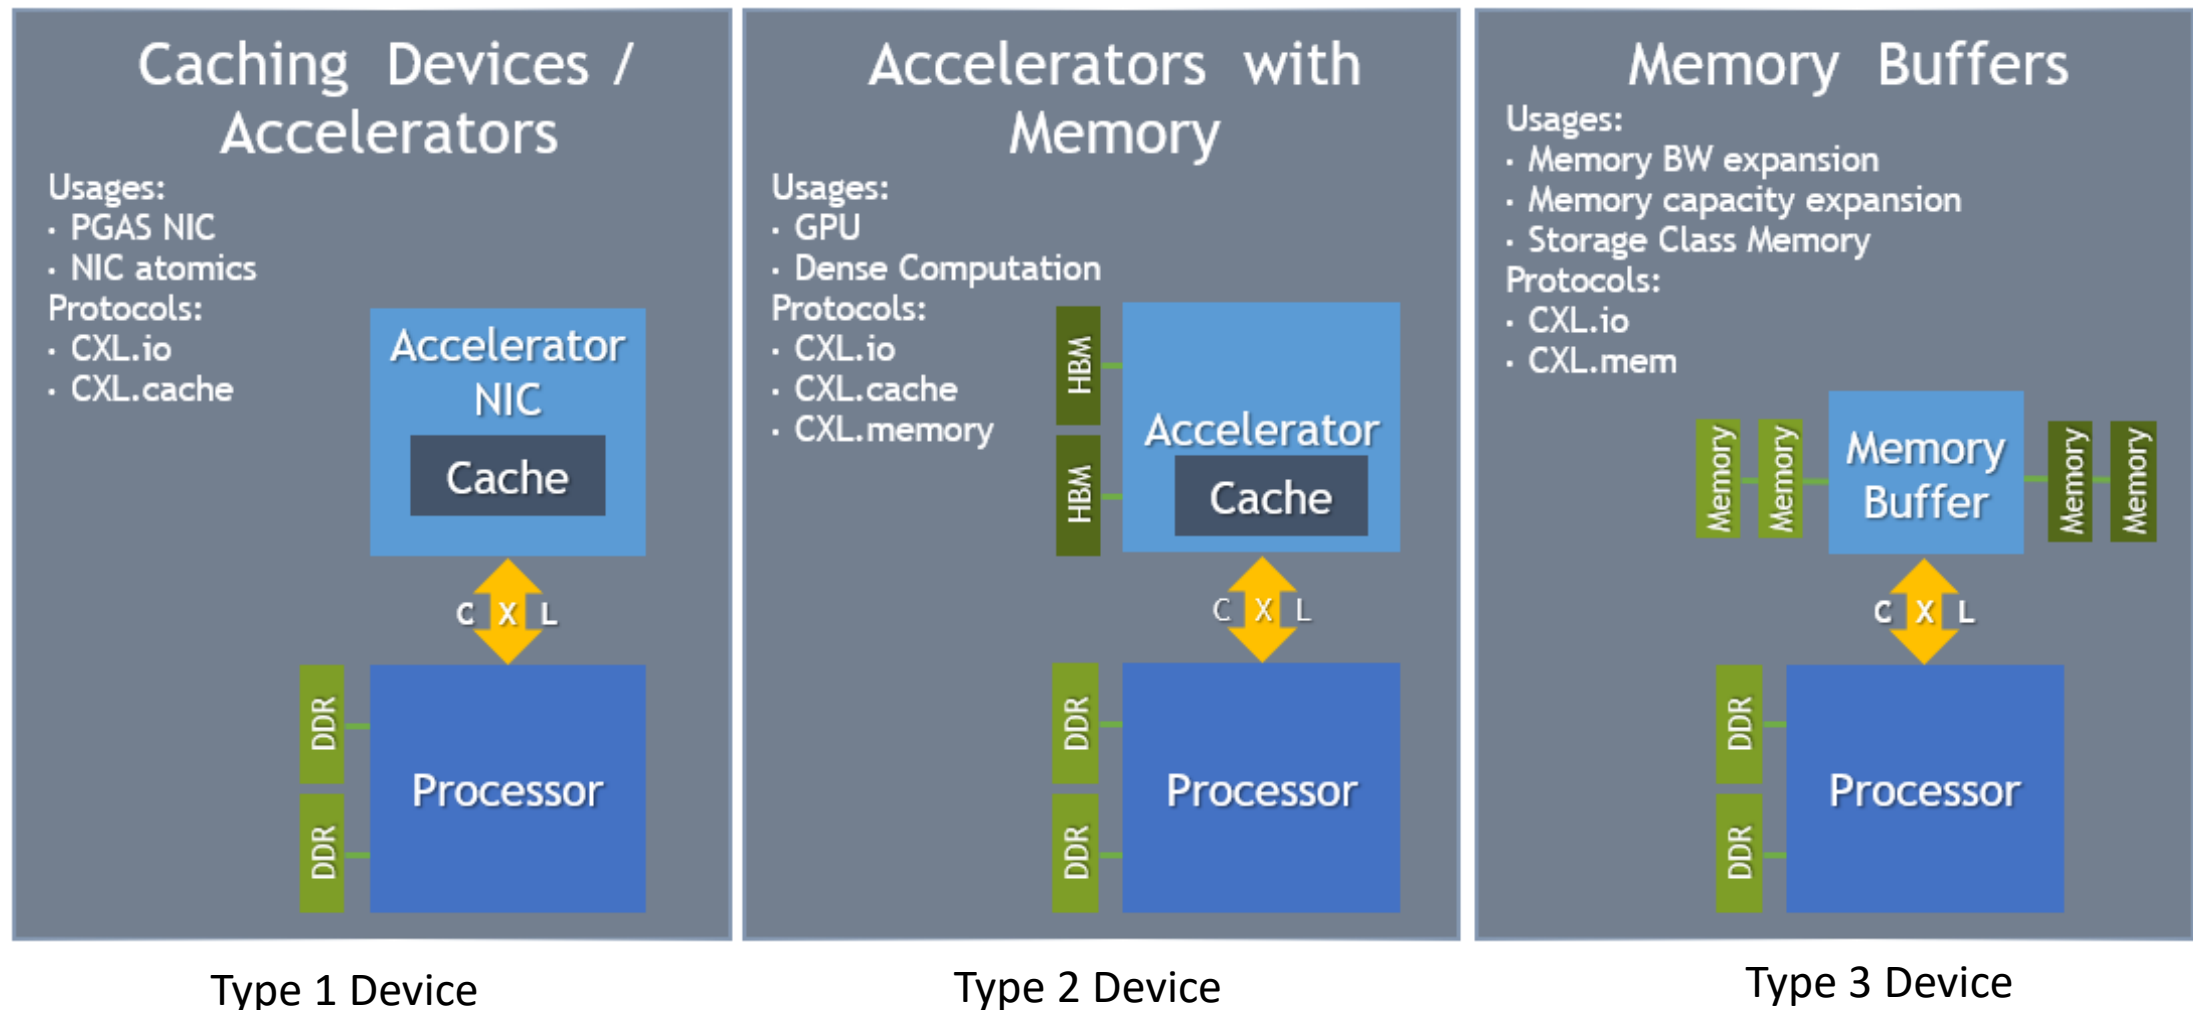
Both biases guaranteed correct/coherent

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



CXL USE CASES

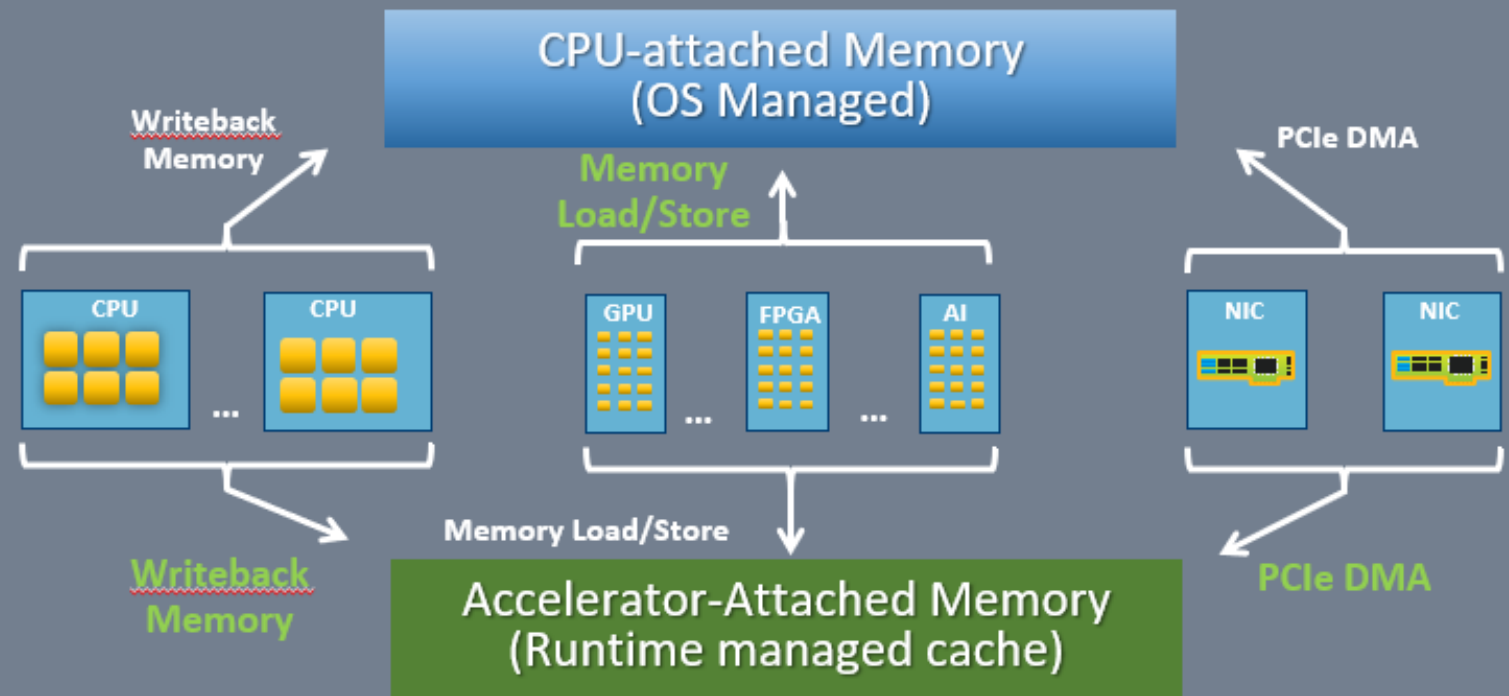
REPRESENTATIVE CXL USAGES



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



2020 OFA Virtual Workshop

THANK YOU

Jim Pappas, Director, Industry Initiatives, Intel Corporation
Chairman, CXL Consortium

