EIDETICOM

Accelerate Everything.

Successfully Deploying Persistent Memory and Acceleration via Compute Express Link!

Stephen Bates, Chief Technology Officer, Flash Memory Summit 2019

We've Been At This For Some Time….

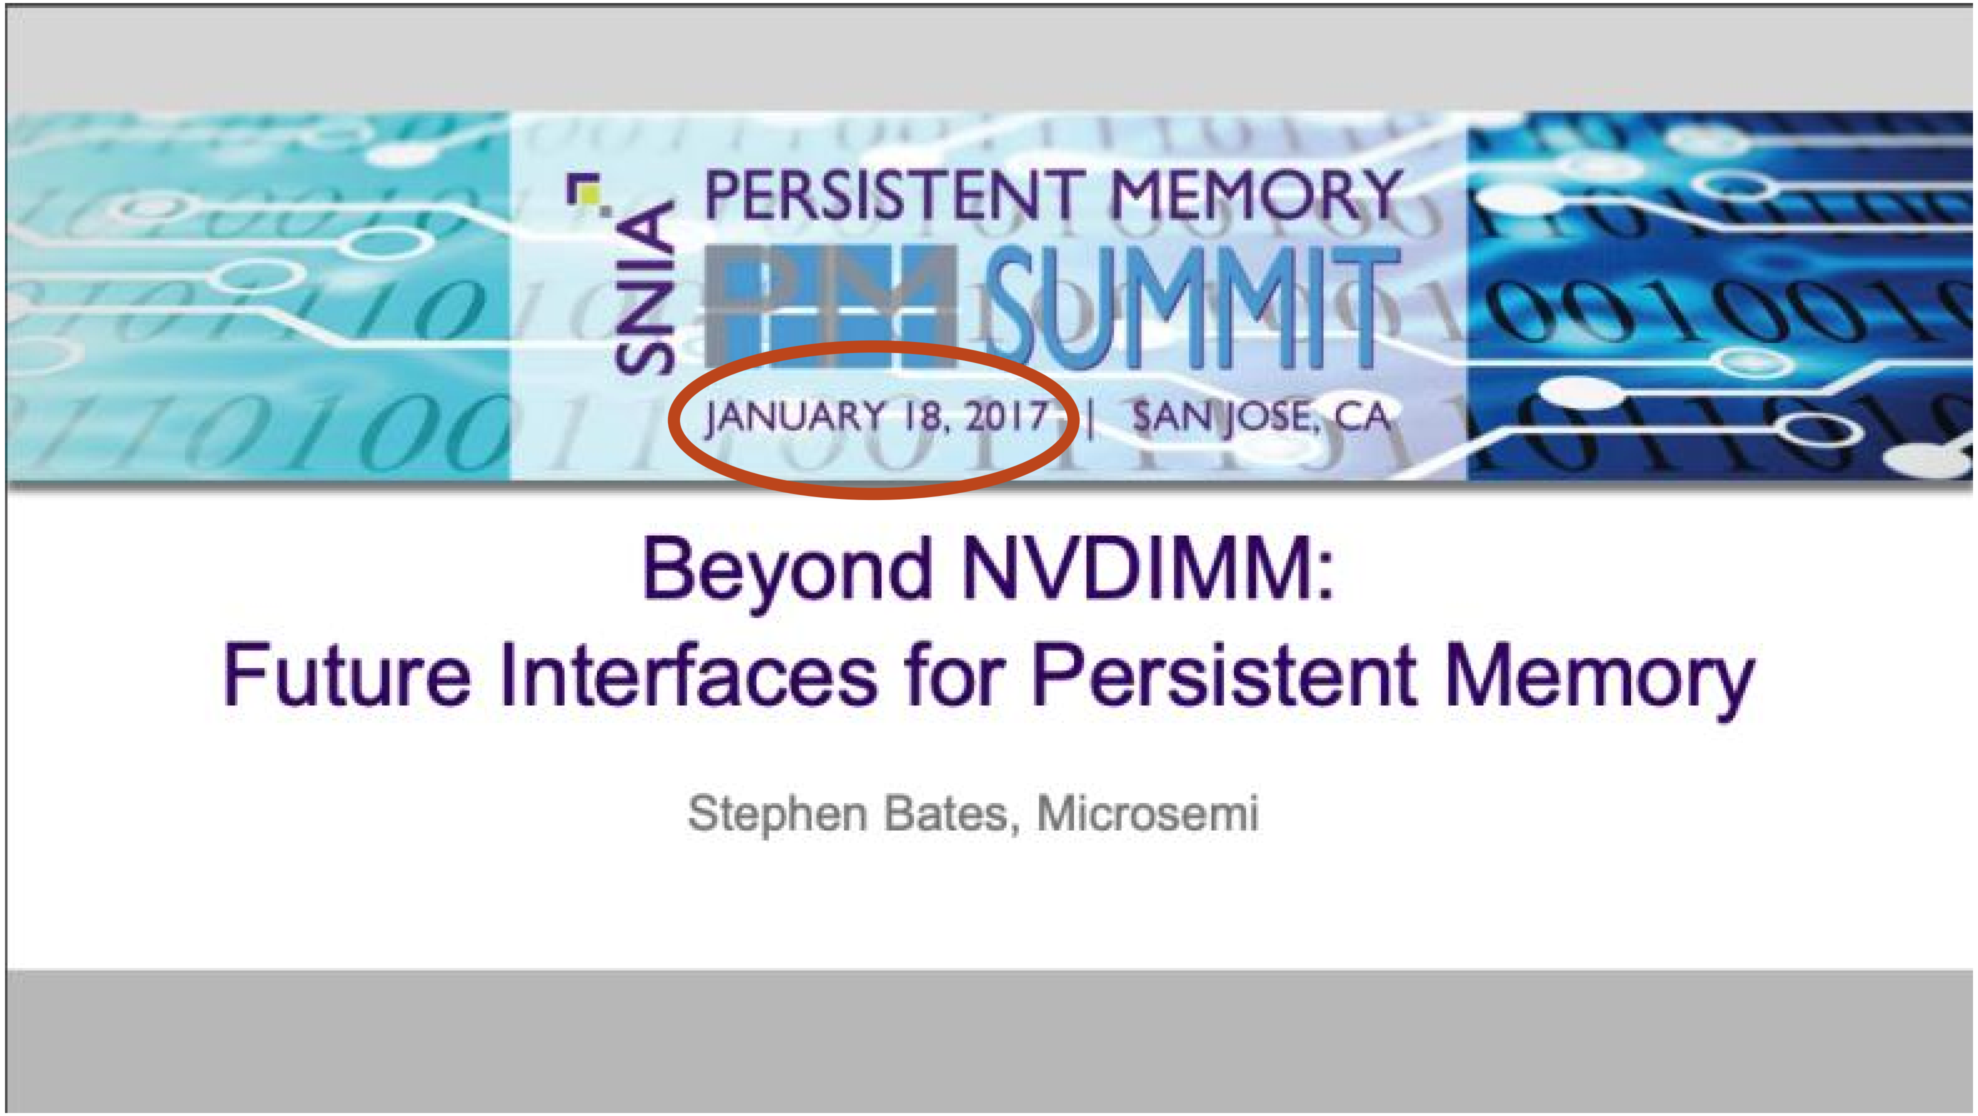

SNIA PERSISTENT MEMORY PM SUMMIT

JANUARY 18, 2017 | SAN JOSE, CA

Beyond NVDIMM:
Future Interfaces for Persistent Memory

Stephen Bates, Microsemi

Persistent Memory (PM)

Low Latency

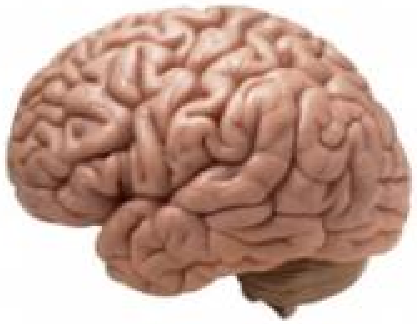

Memory Semantics

Storage Features

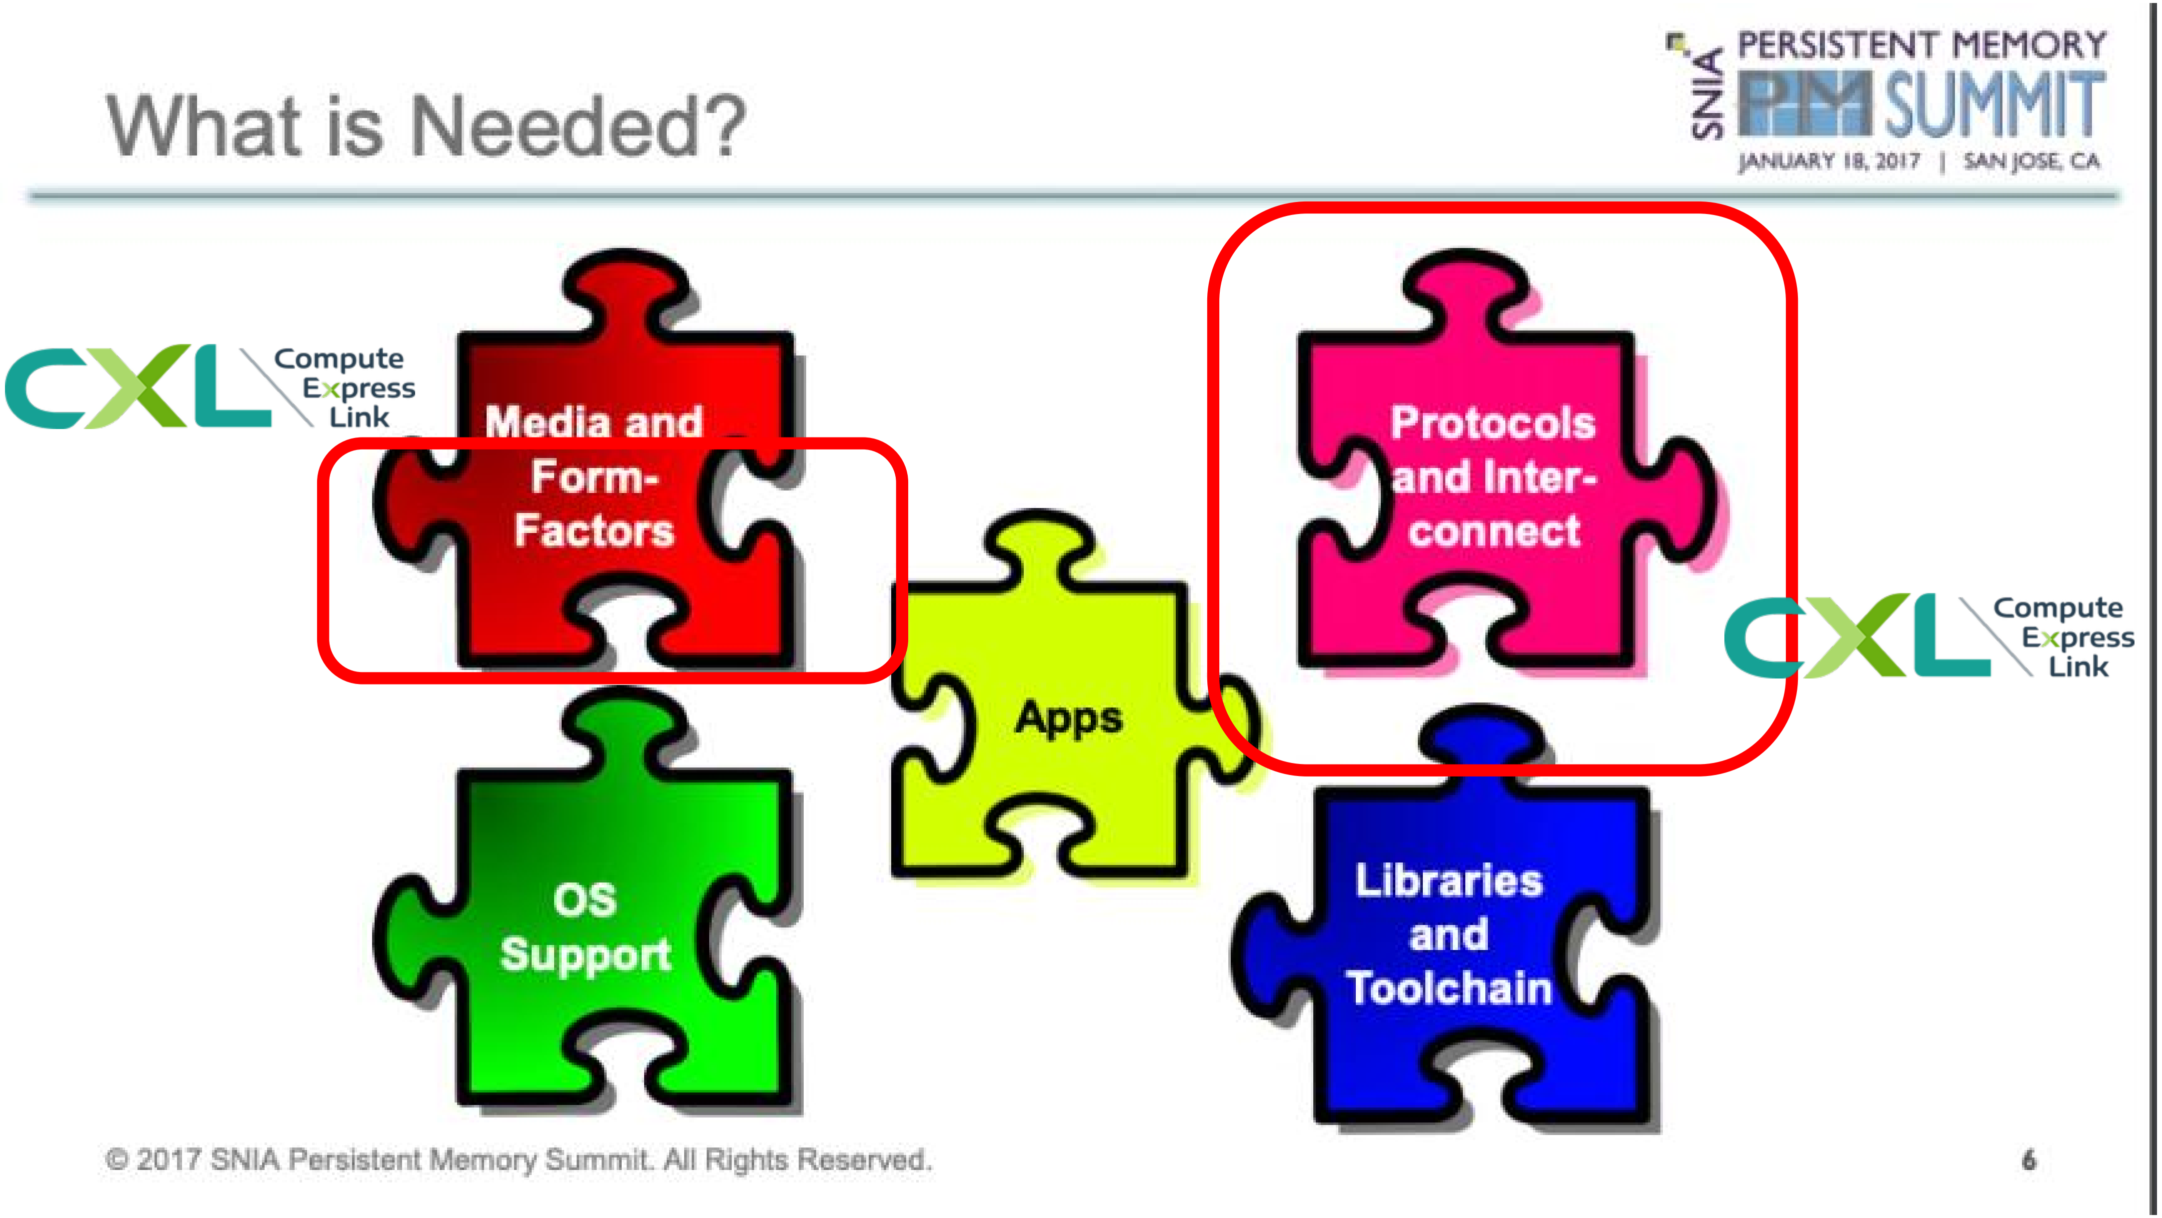

PERSISTENT MEMORY PM SUMMIT
JANUARY 18, 2017 | SAN JOSE, CA

What is Needed?

Media and Form-Factors

OS Support

Apps

Protocols and Inter-connect

Libraries and Toolchain

CXL Compute Express Link

CXL Compute Express Link

Coming Soon to a Cinema Near You!

GEN Z

A New Fabric

featuring
Optional coherency
NVMe support
Scale

Coming in 2020

CCIX

The ARMpire Strikes Back

featuring
Off the CPU bus
Accelerator support
Cache coherency
Scale?

Coming Soon??

OpenCAPI

The Return of the Big Blue

featuring
Off the CPU bus
Accelerator support
Cache coherency

Now Showing in Select Cinemas

BUT WHAT CAN I SEE TODAY???

Coming Soon to a Cinema Near You!

GEN Z

A New Fabric

featuring
Optional coherency
NVMe support
Scale

Coming in 2020

CCIX

The ARMpire Strikes Back

featuring
Off the CPU bus
Accelerator support
Cache coherency
Scale?

Coming Soon??

OpenCAPI

The Return of the Big Blue

featuring
Off the CPU bus
Accelerator support
Cache coherency

Now Showing in Select Cinemas

BUT WHAT CAN I SEE TODAY???

CXL Compute Express Link

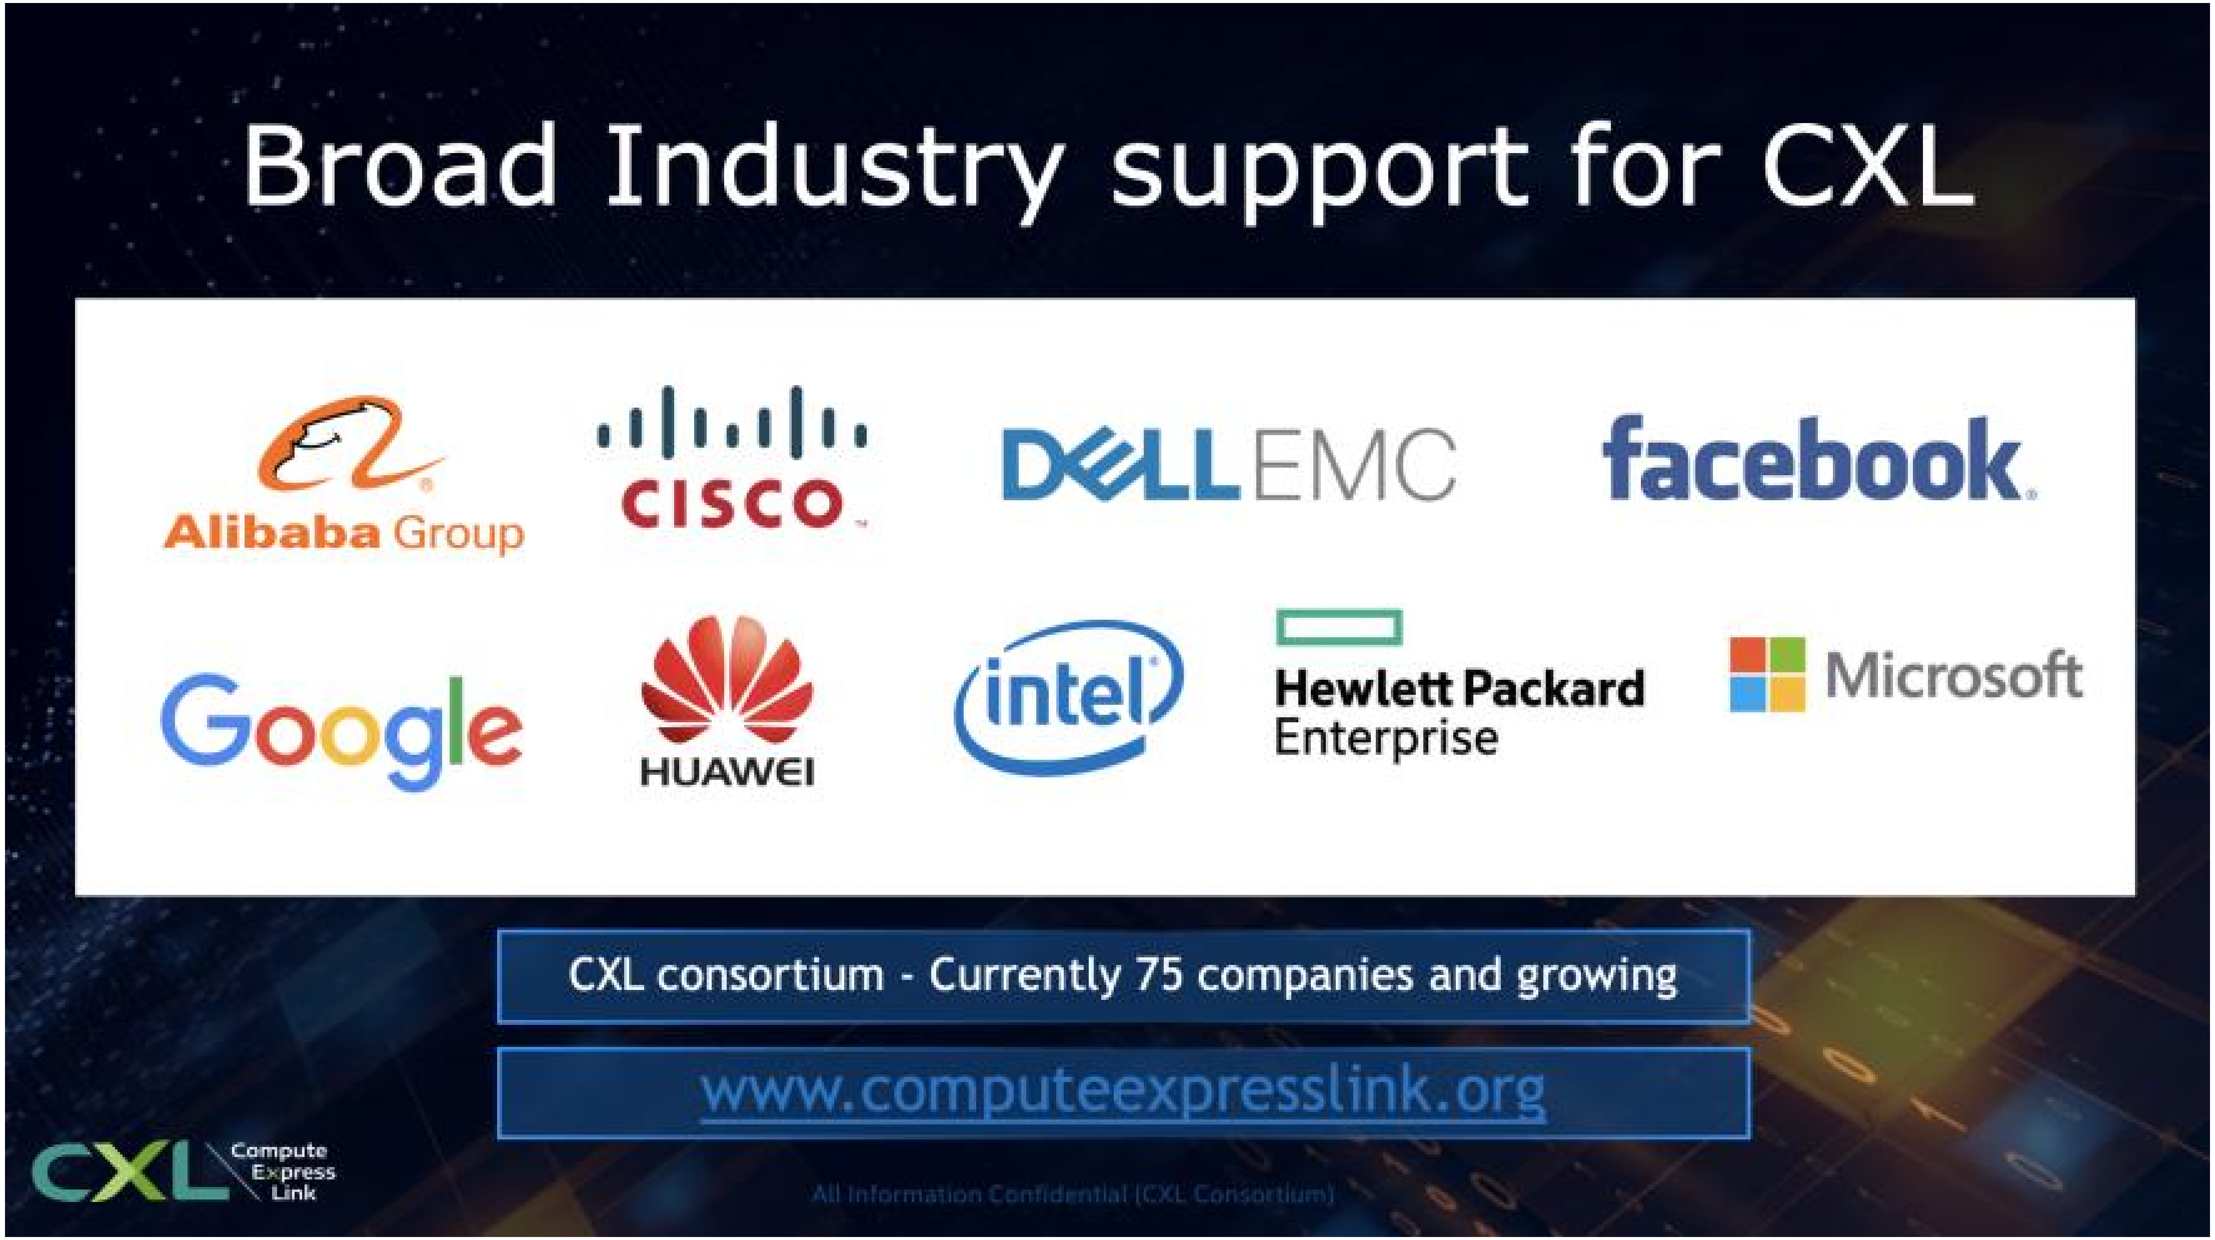

Broad Industry support for CXL

Alibaba Group

CISCO

DELL EMC

facebook

Google

HUAWEI

intel

Hewlett Packard Enterprise

Microsoft

CXL consortium - Currently 75 companies and growing

www.computeexpresslink.org

CXL Compute Express Link

All Information Confidential (CXL Consortium)

intel
arm
AMD

Last Week!

- All the CPU vendors I care about are now CXL members.
- Same cannot be said for OpenCAPI, CCIX or Genz
- Remember, coherent buses MUST come directly out of the CPU!

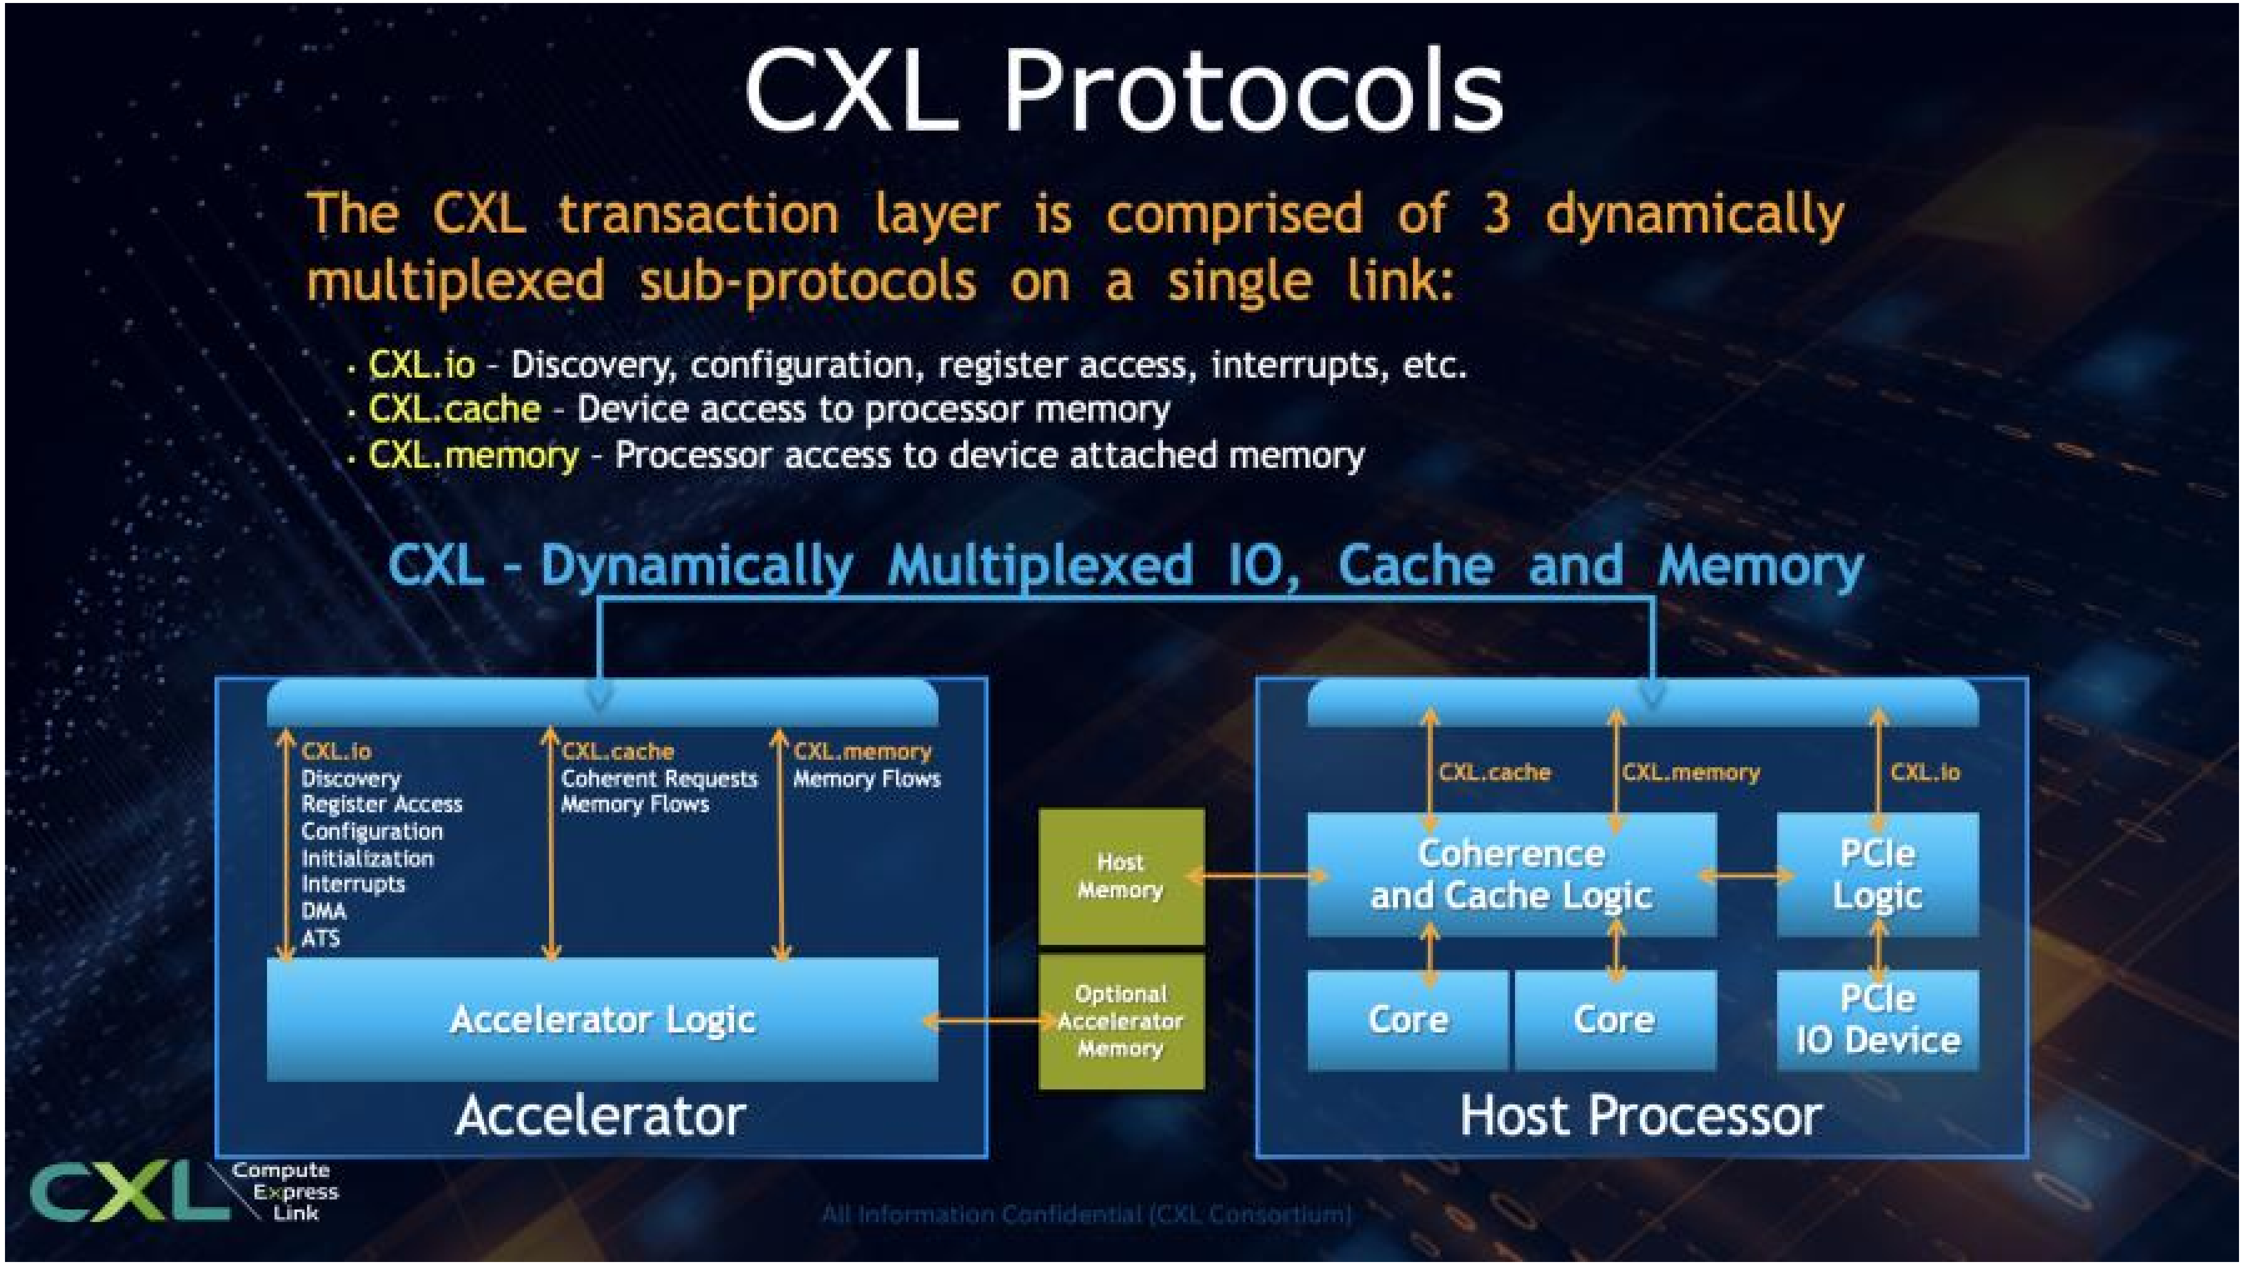

CXL Protocols

The CXL transaction layer is comprised of 3 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 - Dynamically Multiplexed IO, Cache and Memory

All Information Confidential (CXL Consortium)

Let’s break that down. Three protocols on one physical layer:

- **CXL.io**: This is PCIe Gen 5.0. All PCIe services will just work!
  - DMA
  - Interrupts (MSI/MSIX)
  - SR-IOV, ACS, ATS etc. for virtualization
  - NVM Express!!!??? – We will come back to this
- **CXL.mem**: This is the protocol by which the host CPU accesses (persistent) memory on the CXL device.
- **CXL.cache**: The is the protocol by which the CXL device accesses host memory (useful for accelerators, not covered here today).

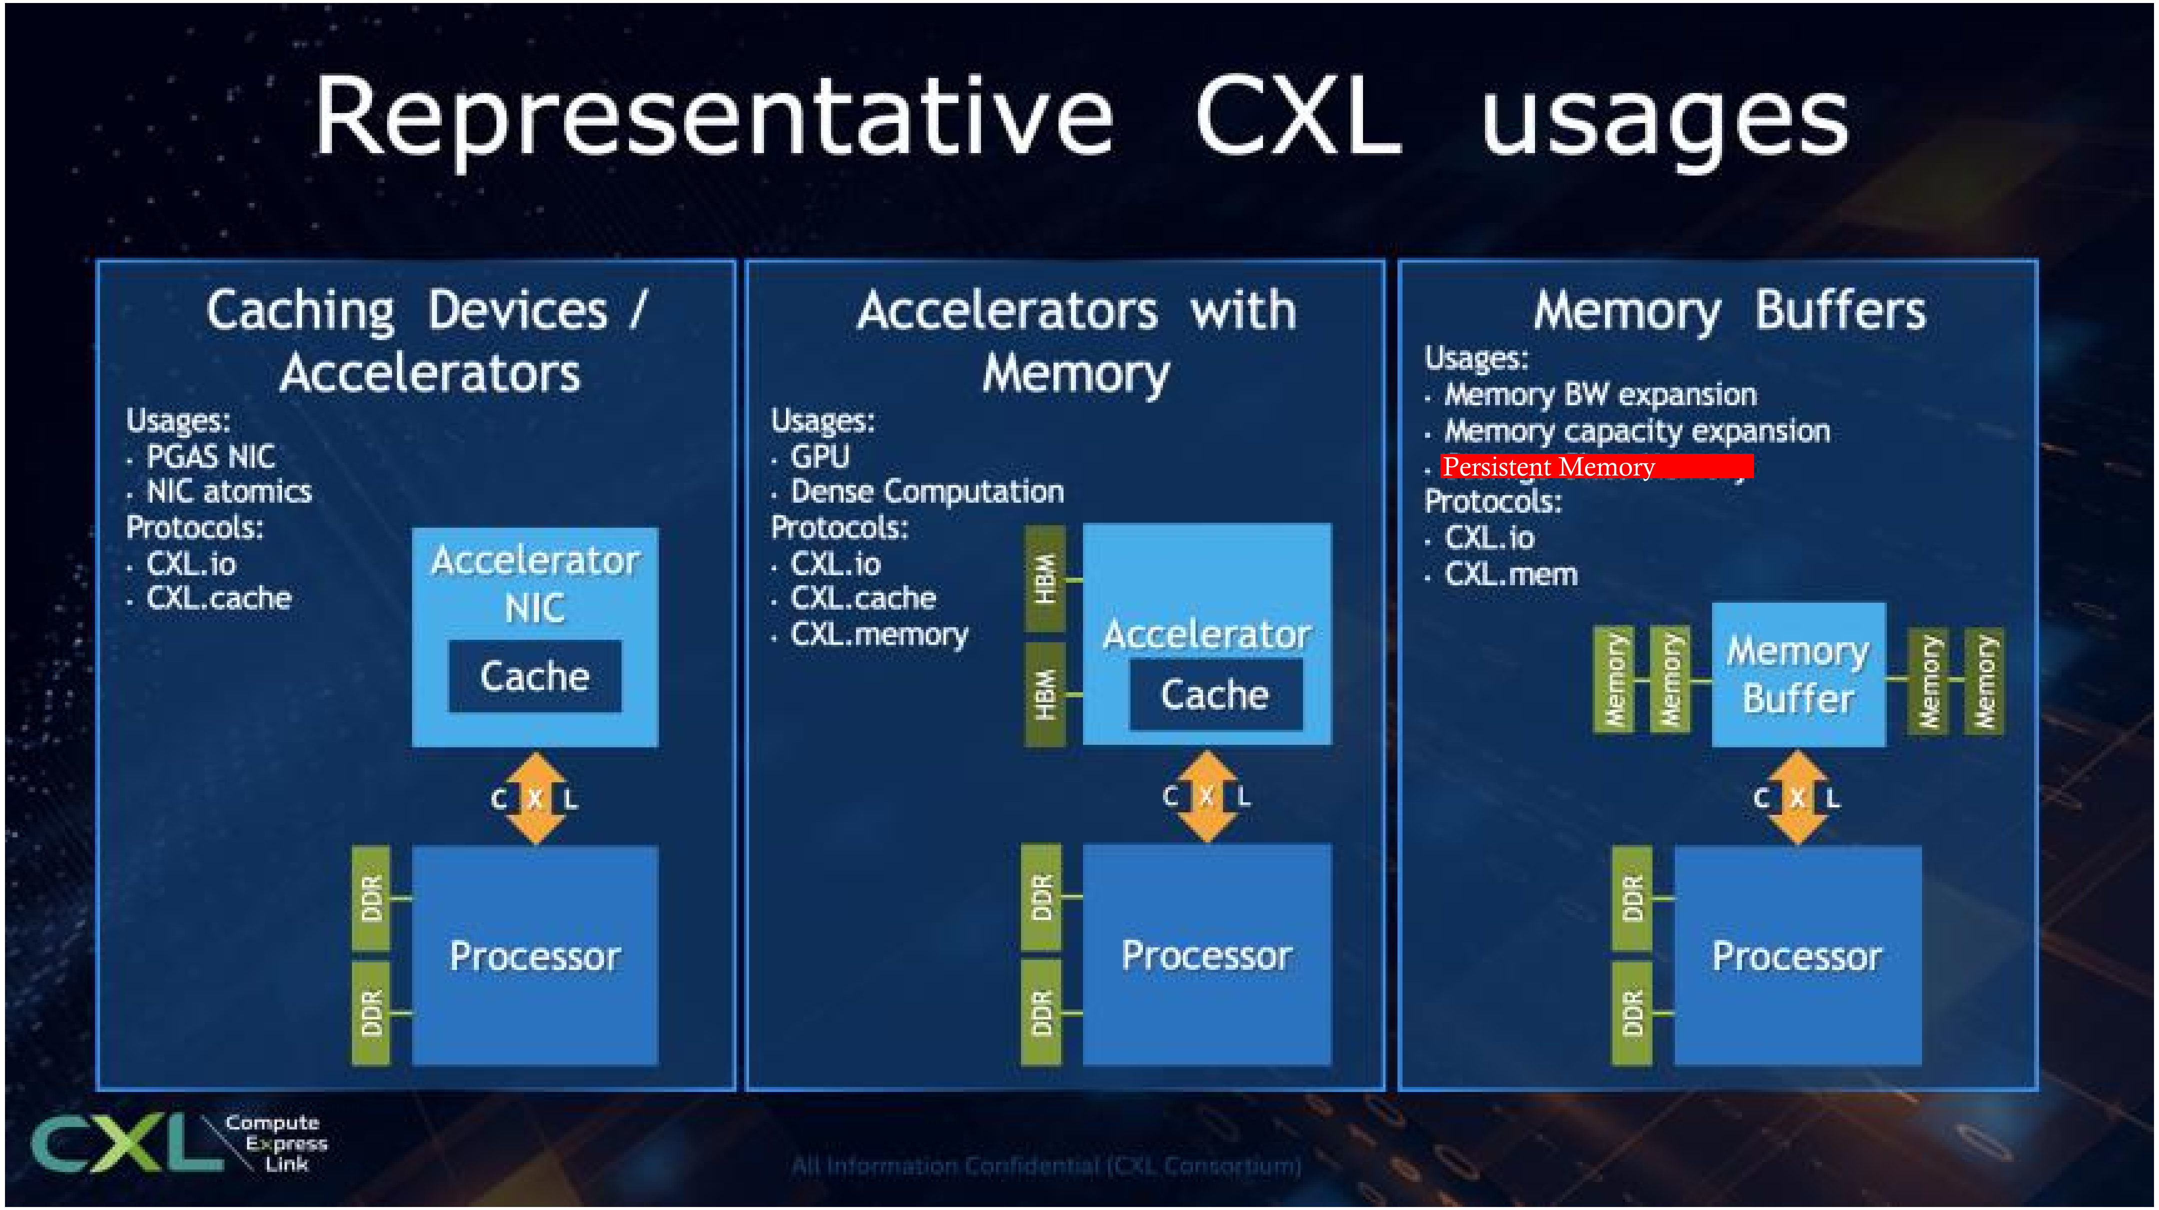

Representative CXL usages

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
- Persistent Memory

Protocols:

- CXL.io
- CXL.mem

All Information Confidential (CXL Consortium)

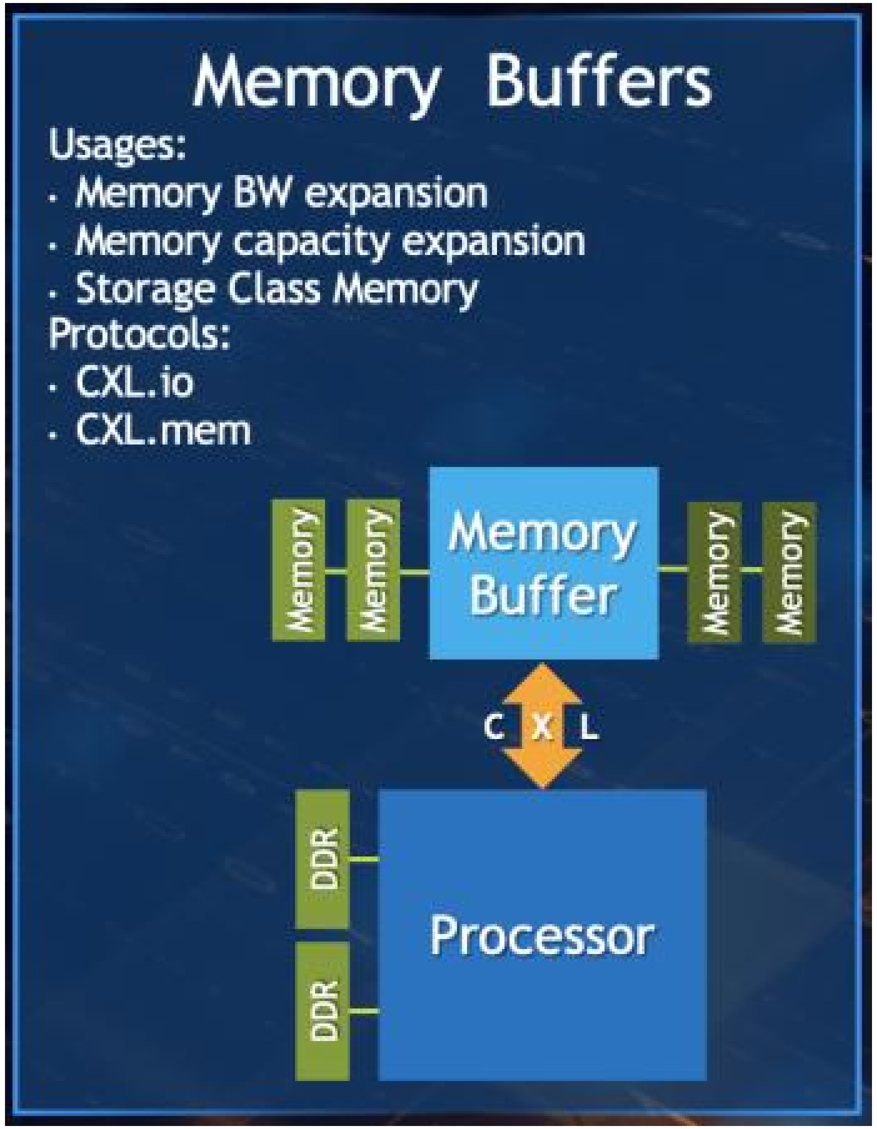

Memory Buffers

Usages:
- Memory BW expansion
- Memory capacity expansion
- Storage Class Memory

Protocols:
- CXL.io
- CXL.mem

Let’s break that down. Consider the right-most model:

- Essentially a NVDIMM but no longer constrained by the physical and electrical requirements of DDR and DIMMs.
- Since the form-factors are PCIe we have more options around the shape, power and heat of these solutions.
- The CXL.io allows for discovery, configuration and management (we can write a PCIe driver for these devices).
- We can put a DMA engine on the Memory Buffer and program that via PCIe to do data movement for us.
- No longer consuming DIMM slots or channels. Save all that capacity and bandwidth for standard DRAM.

DDR NVDIMM vs CXL NVDIMM

Attribute	DDR	CXL	Comment
Form-factor	DIMM	Many	CXL has many form-factor options
DMA	No	Yes	CXL allows placement on DMA engine on device. Can be programmed via PCIe driver.
HW Virtualization	No	SR-IOV	NVDIMM can be virtualized via software which impacts performance.
Management	SMBus and MMIO	SMBus and MMIO and CXL.io	If we adopt NVMe for CXL devices we can use NVMe Management Interface (NVMe-MI).
Latency	Very Low	Low	Until we get hardware it is hard to get comparative numbers for NVDIMM vs CXL.mem to the same memory types (e.g. 3DXP).
Throughput	19GB/s	64GB/s	NVDIMM is 64 bits @ 2400MT/s/channel. CXL is (upto) 16 lanes of PCIe Gen 5 in each direction.

Linux Support for CXL

- (Persistent) Memory discovery will be done via ACPI. This can include Heterogenous Memory Attribute Tables (HMAT) to describe properties of the memory.
- The discovered memory will be added to the physical memory pool.
- We can control how and who this memory is used by to some extent by the numactl framework.
- *If* the CXL device has a DMA engine and accelerator(s) these can be programmed via a PCIe driver (perhaps NVM Express).

Use Case: CXL-Based NVDIMM

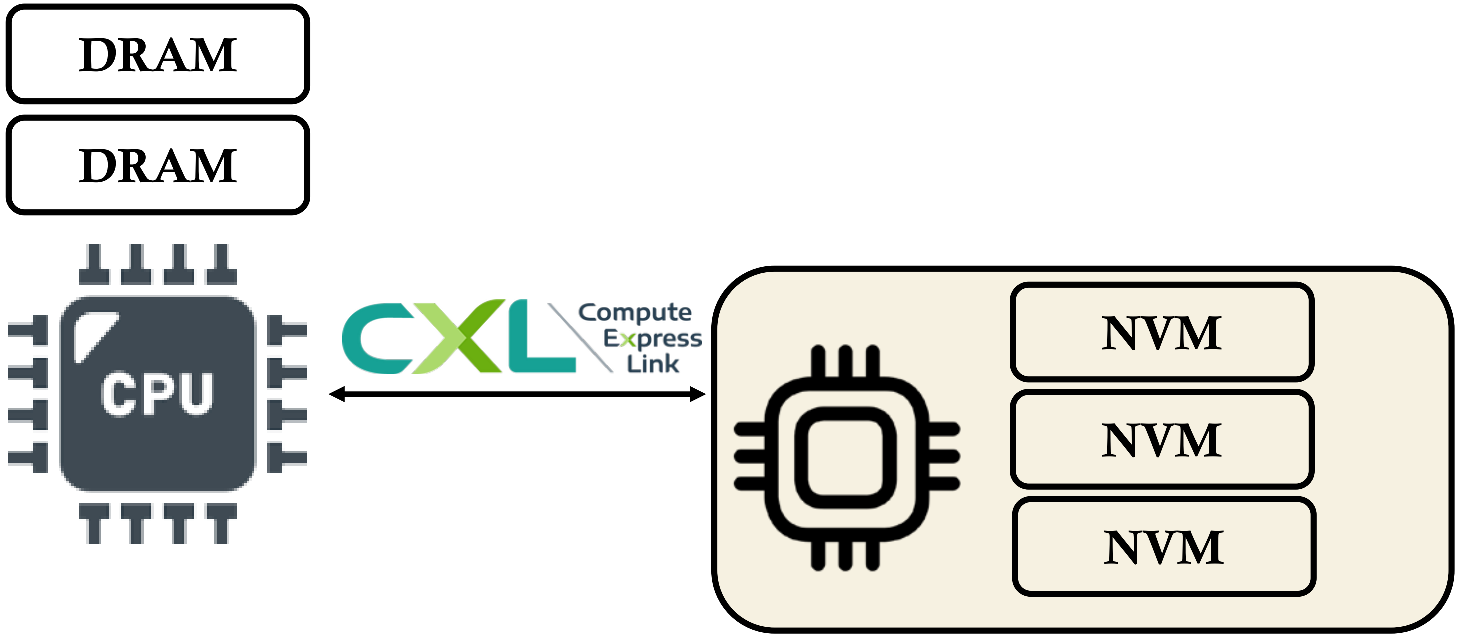

CXL-based NVDIMM:

- Use all the DIMM slots for DRAM, not NVDIMM.
- NVM can be managed by controller chip if needed.
- A lot more flexibility on form-factor, power etc than DDR based NVDIMM.
- A DMA engine on CXL device could assist with data movement.
- Can also be used just to expand volatile memory capacity.

NoLoad® NVMe-based Computational Storage Processor

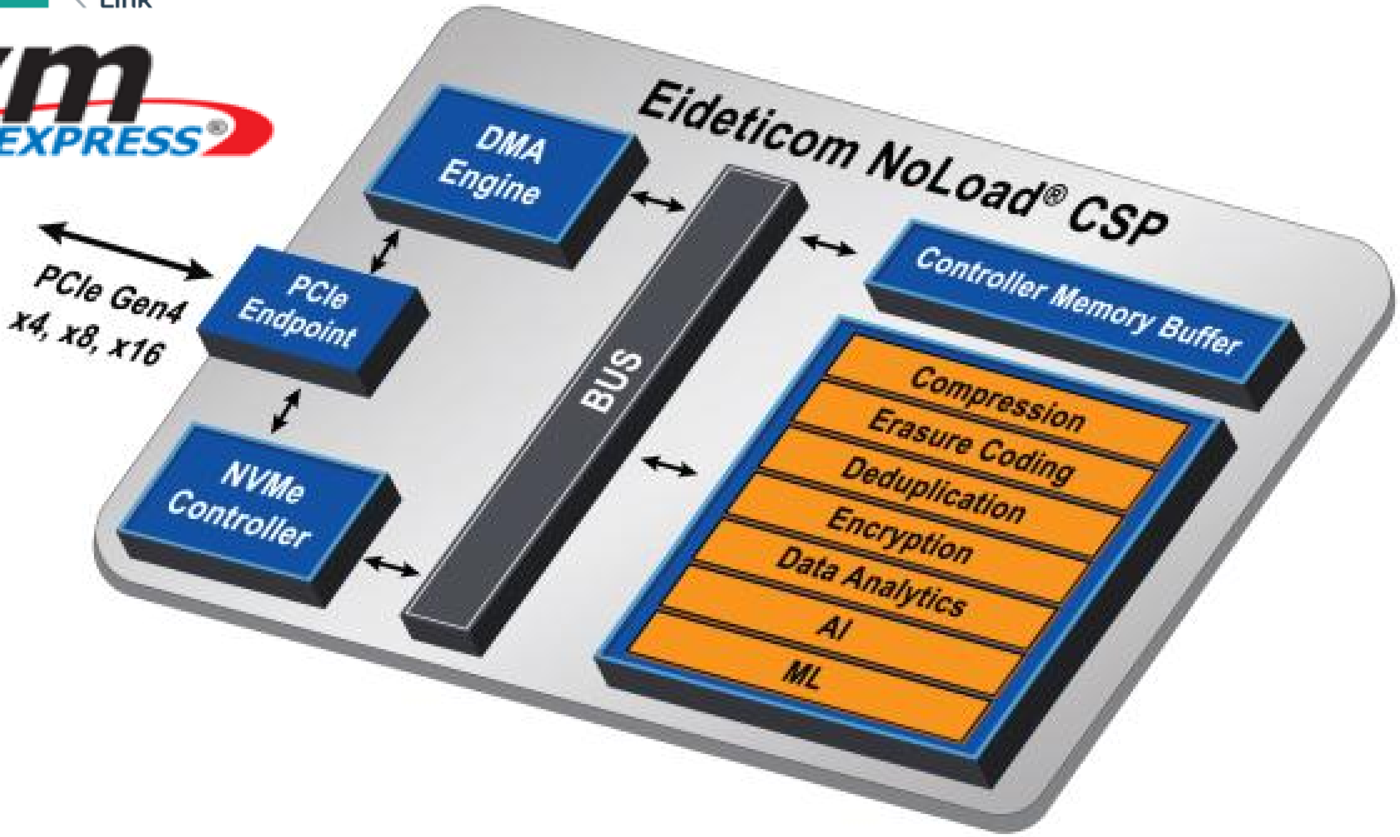

Best-In-Class Storage and Analytic Acceleration delivered via an NVMe-based Computational Storage Processor.

CXL features for 2.0:

- Improved throughput and latency (PCIe Gen6).
- Switching via enhanced PCIe switches
- Memory pooling (allowing multiple hosts to connect to a pool of (persistent) memory.

Conclusions

- CXL may finally be bringing some clarity to the "Star Wars" of open, coherent buses.
- Minimal software changes needed to deploy (persistent) memory on CXL.
- Adding acceleration and remote PM both possible.
- We all get a pony!

Eideticom HQ
3553 31st NW,
Calgary, AB,
Canada T2L 2K7

Eideticom (Bay Area)
168 South Park,
San Francisco, CA 94107
USA

www.eideticom.com

Contact: sales@eideticom.com