The Perfect Trifecta. NVMe, CXL and Persistent Memory!

Millind Mittal, Fellow, Xilinx

Scott Shadley, VP Marketing, NGD Systems

We've Been At This For Some Time…

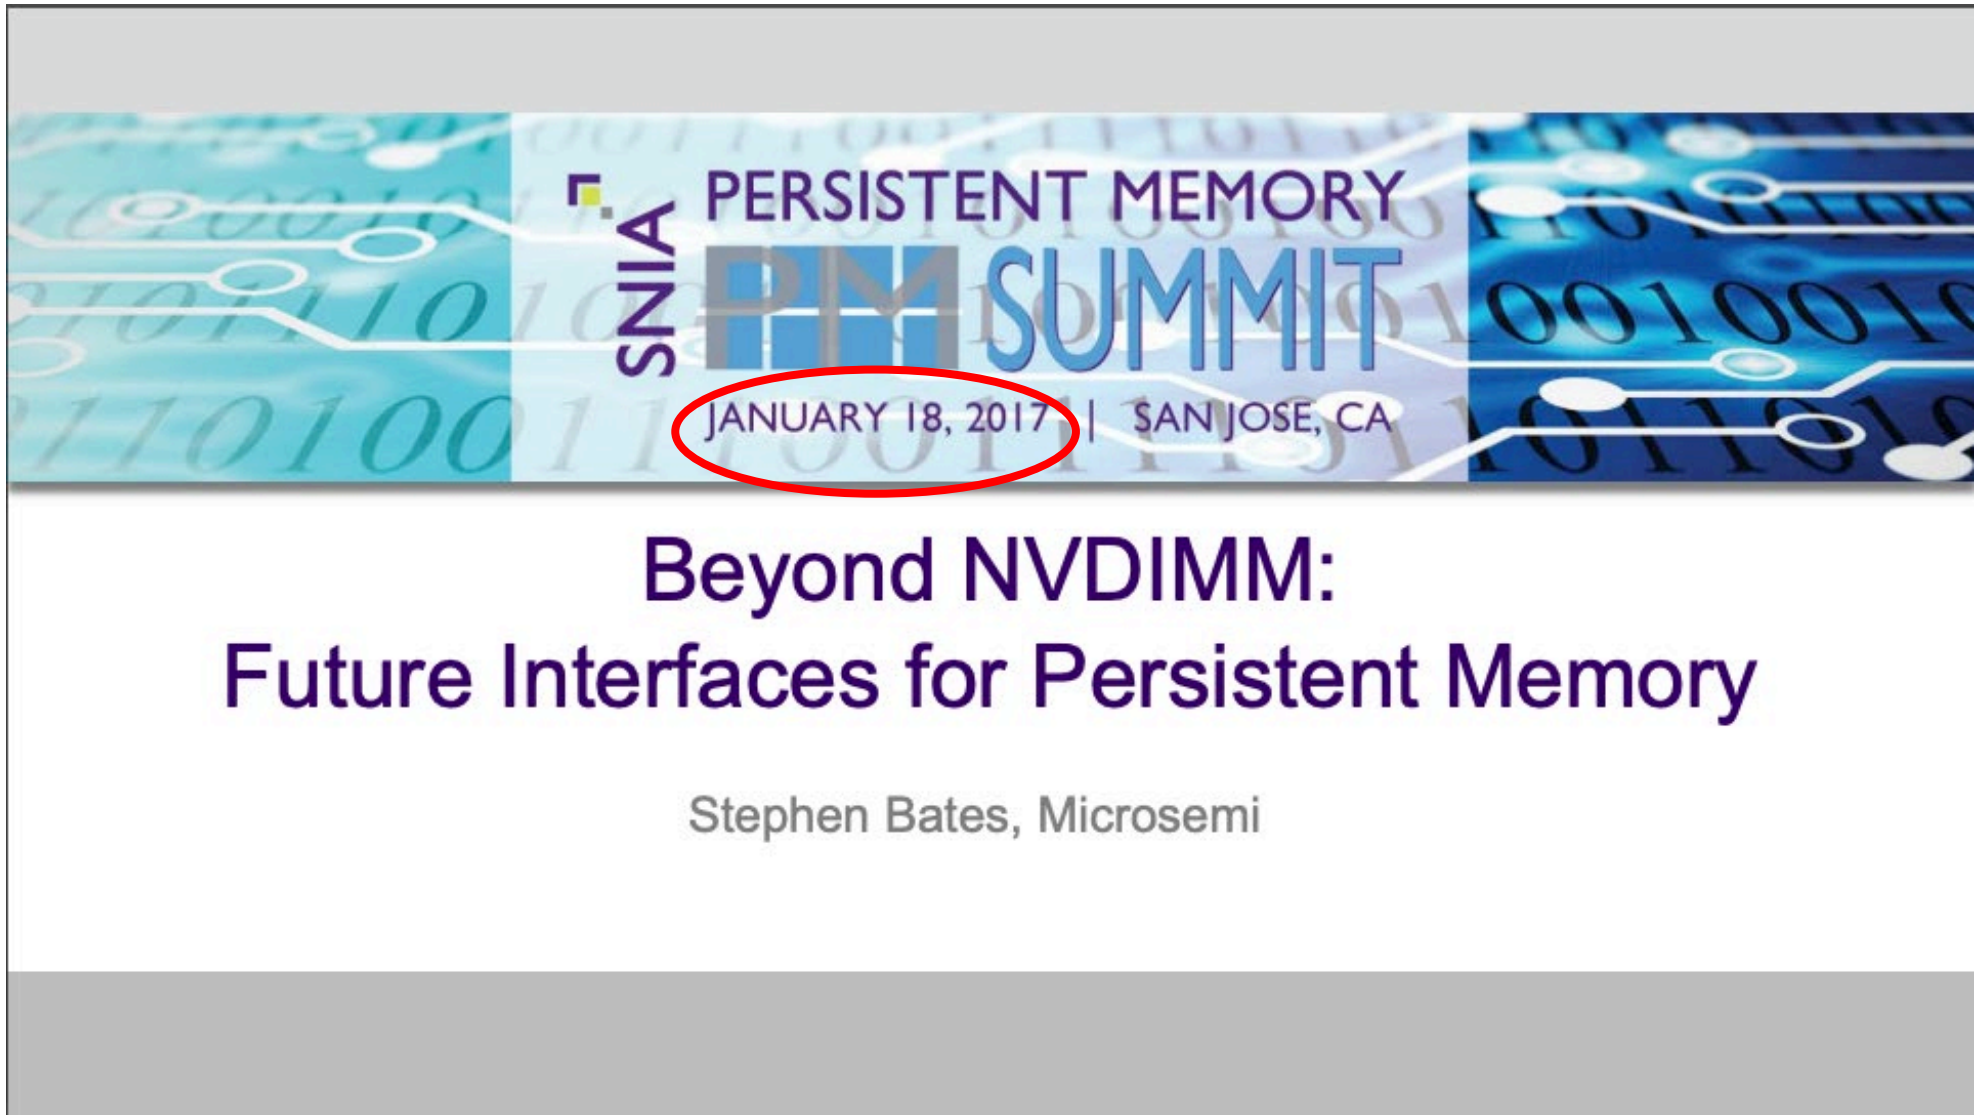

A Coherent Bus MUST Be Supported By The CPU

intel

arm

AMD

IBM

CXL Compute Express Link

CXL is the ONLY open, coherent bus standard with support from all the major CPU vendors. Ideally CPUs from all these vendors will, once day, ship with CXL IO!

What is CXL?

- CXL is an alternate protocol that runs across the standard PCIe physical layer
- CXL 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 Gen5
- CXL usages expected to be key driver for an aggressive timeline to PCIe Gen6

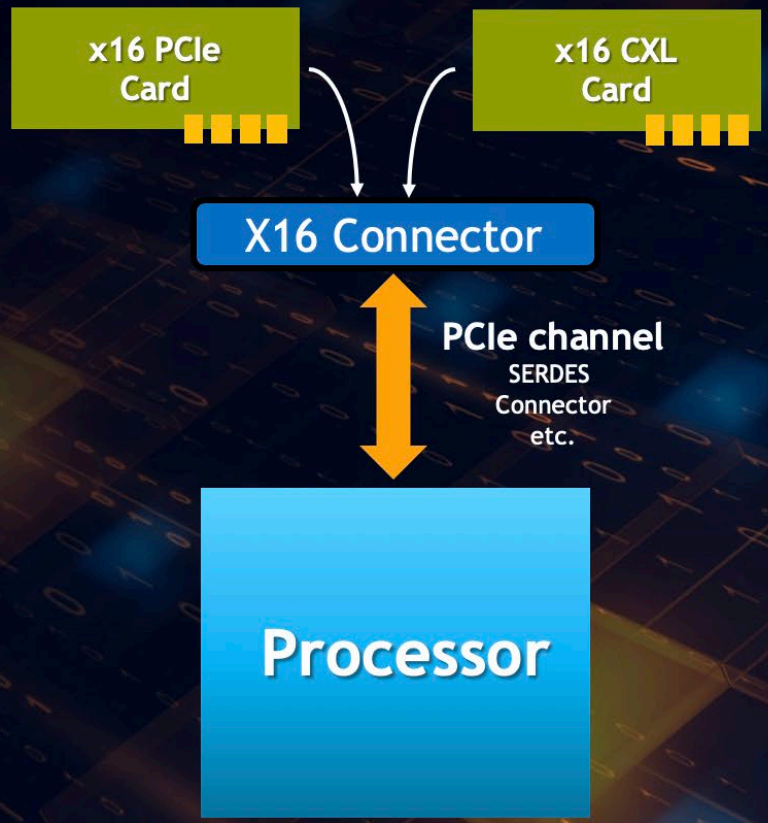

CXL Compute Express Link

Let's break that down:

- PCIe 5.0 based. Links can be connected to CXL devices or switch.
- CXL.mem exposes device memory as "conventional but Specific Purpose Memory (SPM)" via UEFI and perhaps even at run-time (MEMORY_HOTPLUG).
- PCIe connectors (and form-factors) same as CXL connectors (and form-factors). As well as Add-In-Cards we can do things with storage form-factors like U.2 and EDSFF.
- Management buses that also connect to PCIe devices can also connect to CXL devices (I2C, SMBUS). Useful for management.
- Can tie into other frameworks like ACPI (for HMAT for example) and RedFISH/SwordFISH for remote management of CXL enabled servers.

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

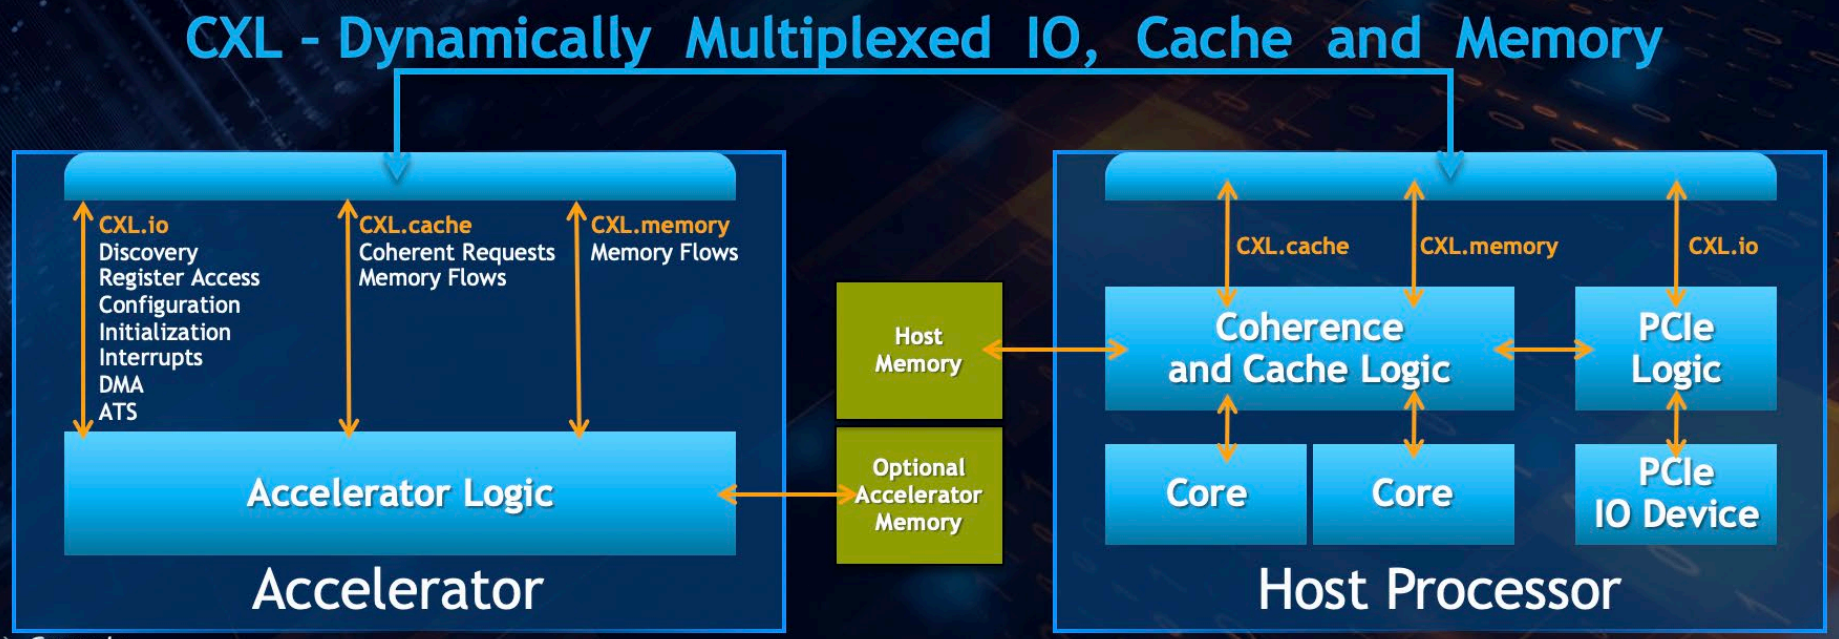

CXL Compute Express Link

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!**
- **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).

Representative CXL usages

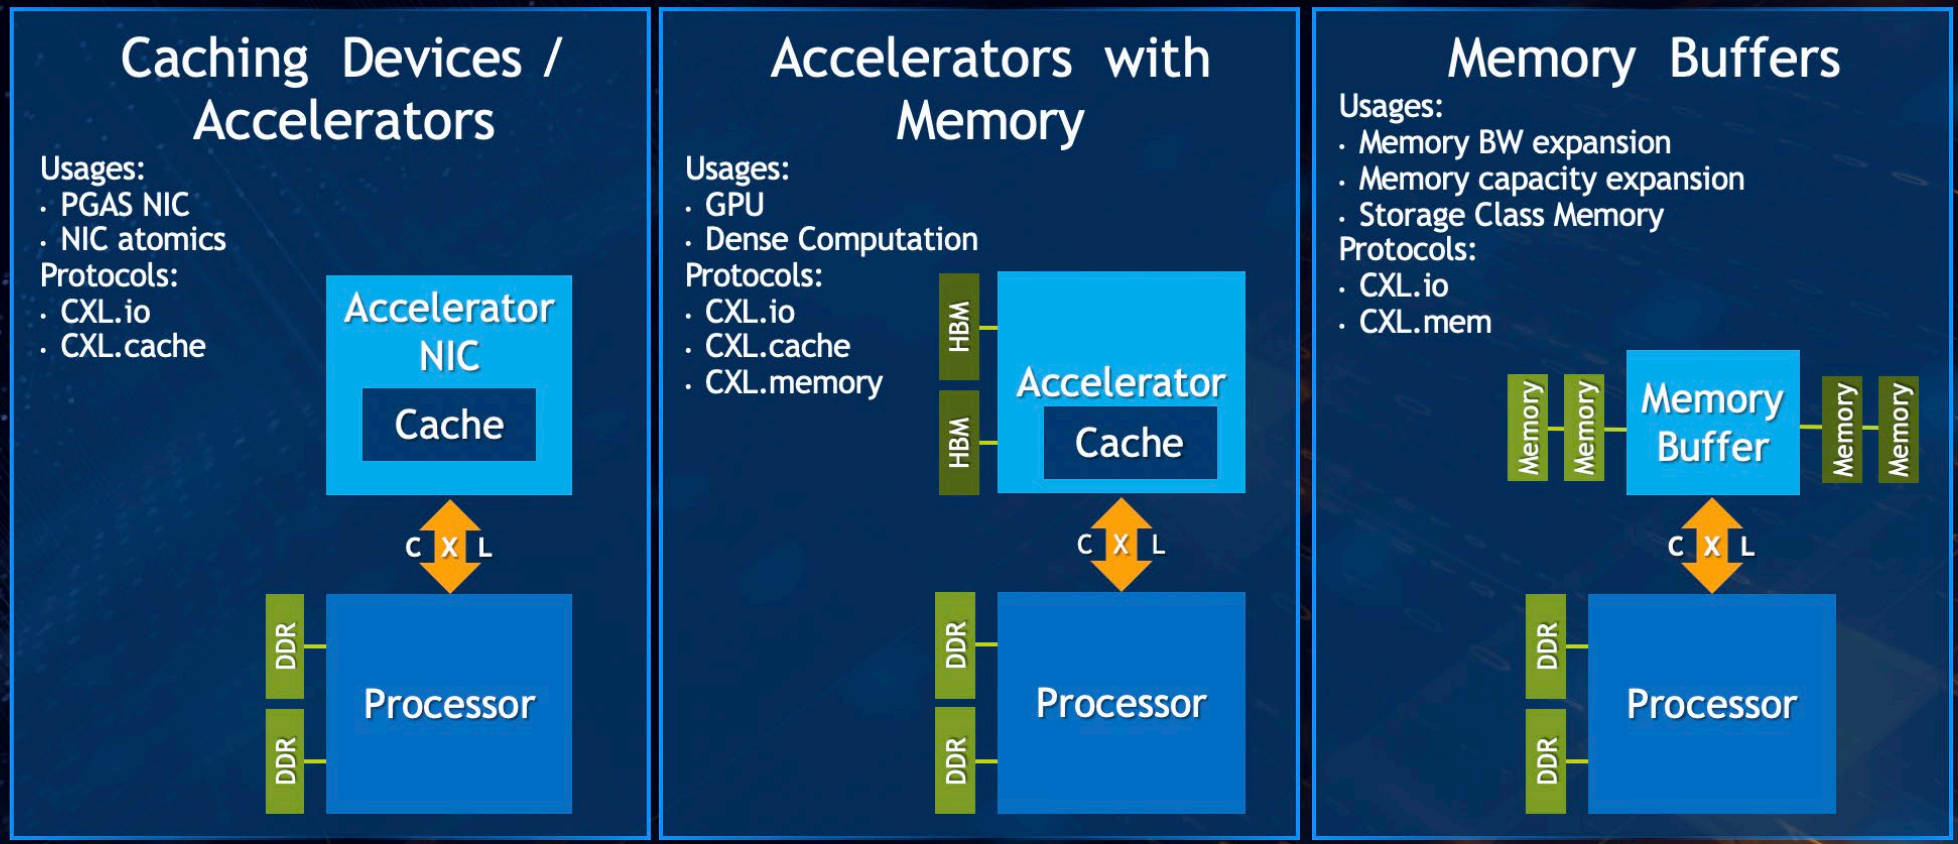

CXL Compute Express Link

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

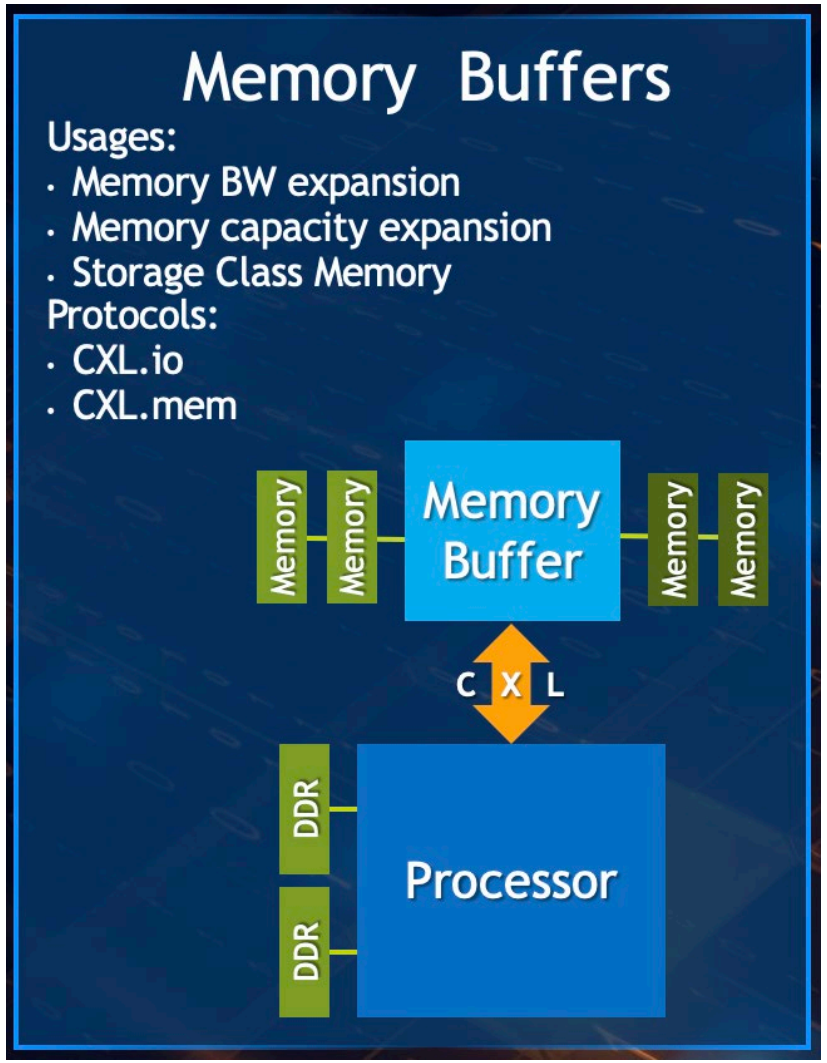

- Essentially **an NVDIMM** but **no longer constrained by the physical and electrical requirements of DDR and DIMMs!**
- Since the form-factors are PCIe we have more options around the shape, power and heat of these solutions. No longer consuming DIMM slots or channels. Save all that capacity and bandwidth for standard DRAM.
- The CXL.io allows for discovery, configuration and management (i.e. 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. **We already have a great PCIe-based protocol for doing all this!**

nvm EXPRESS

Now Add Compute!

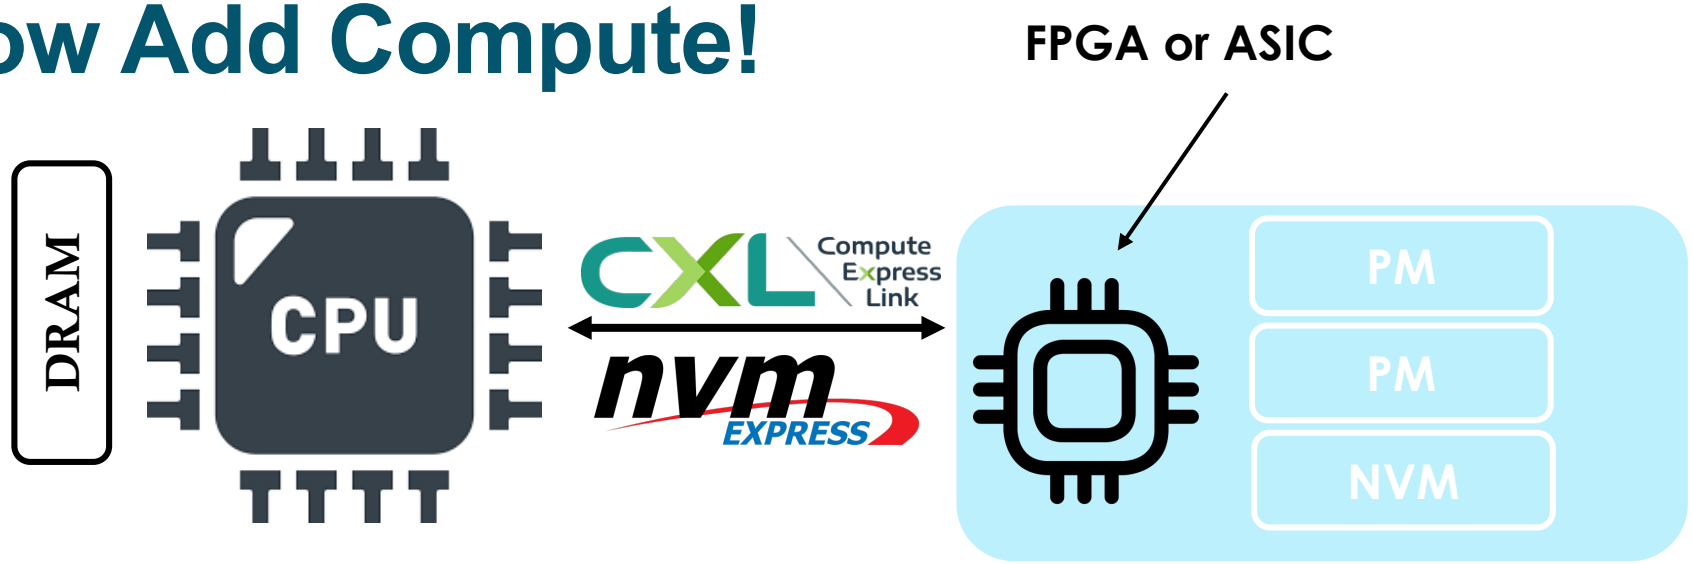

CXL-based NVDIMM+Accelerator:

- Controller chip includes compute functions (e.g. AI, search, graph database)
- Controller chip can be programmed via PCIe driver (e.g. NVMe).
- NVM and PM can still be exposed to host and accessed via CXL.mem (volatile or persistent)
- **Computation on Persistent Memory at the CXL device using standard drivers and management framework!**

Use Cases

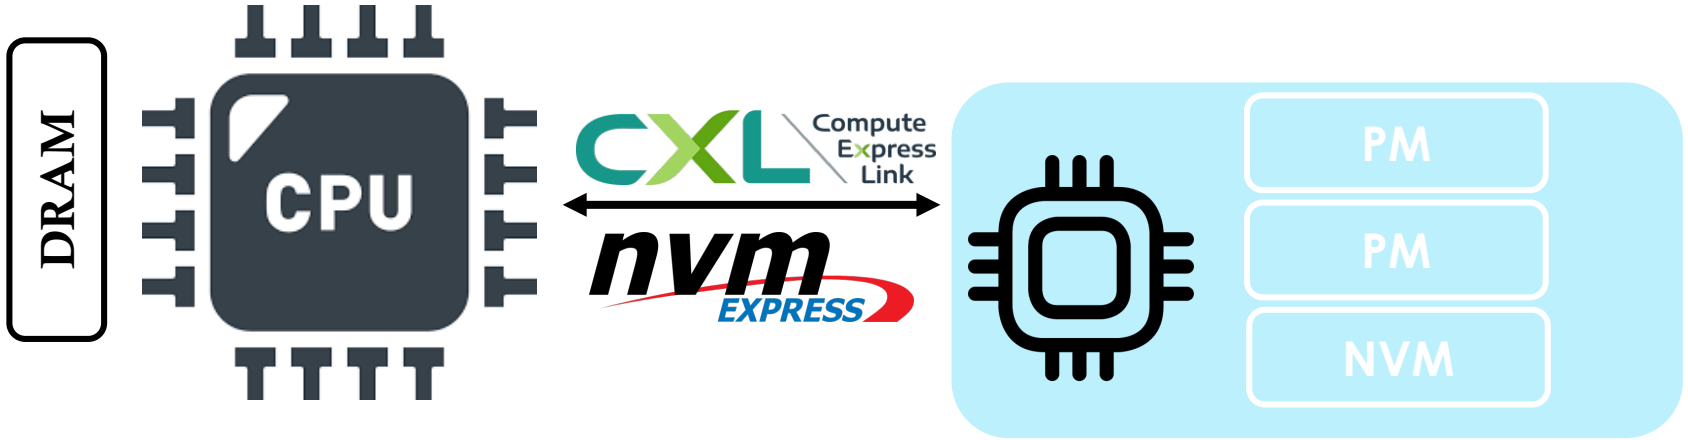

CXL/NVMe-Based NVDIMM with DMA and Computation

- **Storage centric computation** can reduce cost (compression), provide data protection (perform EC or RAID) and provide data integrity (checksums, positions etc).
- **Analytic computation** can offload applications (compaction for a KV database or RegEx for a search engine or vector processing for AI/ML.
- The **DMA capabilities** allow us to move data into and out of the device to CPU memory or other PCIe or CXL IO devices (like GPUs, Ethernet NICs or NVMe SSDs).
- The **NVMe interface** allows us to easily program those DMAs, expose memory as a CMB (for peer-2-peer DMAs), manage the device (via NVMe admin queue), access virtualization (via SR-IOV).

Get Involved!

- The Solid State Storage Initiative (SSSI) is now the one-stop shop for all things NVMe, CXL, Persistent Memory and Computational Storage:
  - SSSI contains both the Computational Storage and the Persistent Memory Special Interest Groups (SIGs).
  - The Computational Storage SIG has a liaison in place with NVM Express and its Technical Working Group.
  - SSSI has a liaison with the CXL Consortium and topics will bridge CXL, Persistent Memory and Computational Storage (and Memory) will be covered in this.

SNIA

https://www.snia.org/forums/sssi

Conclusions

- CXL is the first open, coherent bus standard with support from all the major CPU vendors – and coherency needs support from all the CPU vendors.
- CXL.io offers an “on ramp” from PCIe to CXL and allows us to use legacy PCIe protocols (like NVMe) for DMA, management etc.
- CXL.mem offers support for NVDIMM in PCIe power envelopes and form-factors allowing us to deploy new products which are not constrained by DRAM and DDR.
- CXL.cache allows accelerators on CXL devices to share the address space with the CPU and provide a path to coherent, performat Domain Specific Architecture accelerators.
- **Combining CXL with NVMe Computation and Persistent Memory creates a superb platform for disruption and innovation. The Perfect Trifecta!**

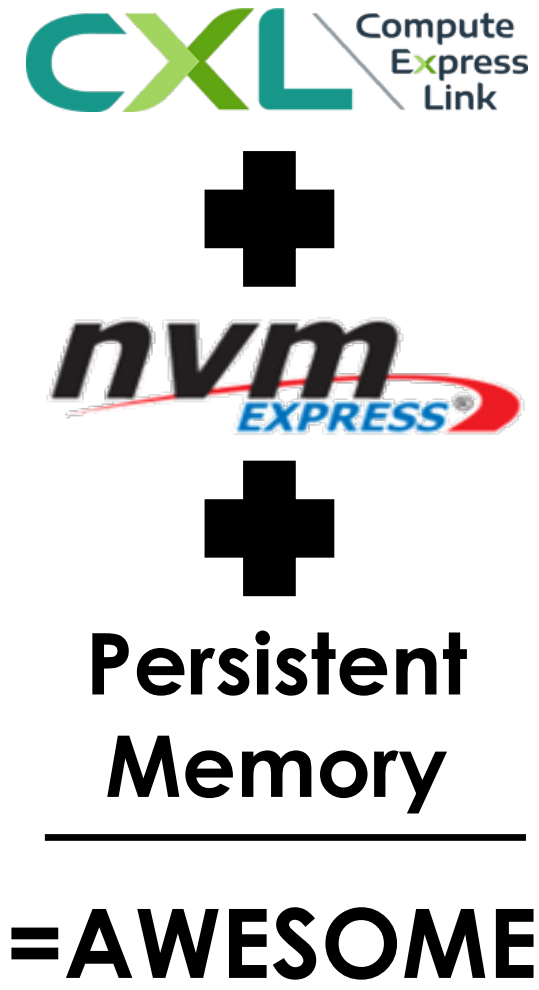

Thank you

Please visit www.snia.org/pmsummit for presentations