Proprietary Interconnects and CXL

Presented by Larrie Carr
VP Engineering, Rambus

SNIA COMPUTE, MEMORY, AND STORAGE SUMMIT

Solutions, Architectures, and Community
VIRTUAL EVENT, MAY 21-22, 2024

CXL Specification Release Timeline

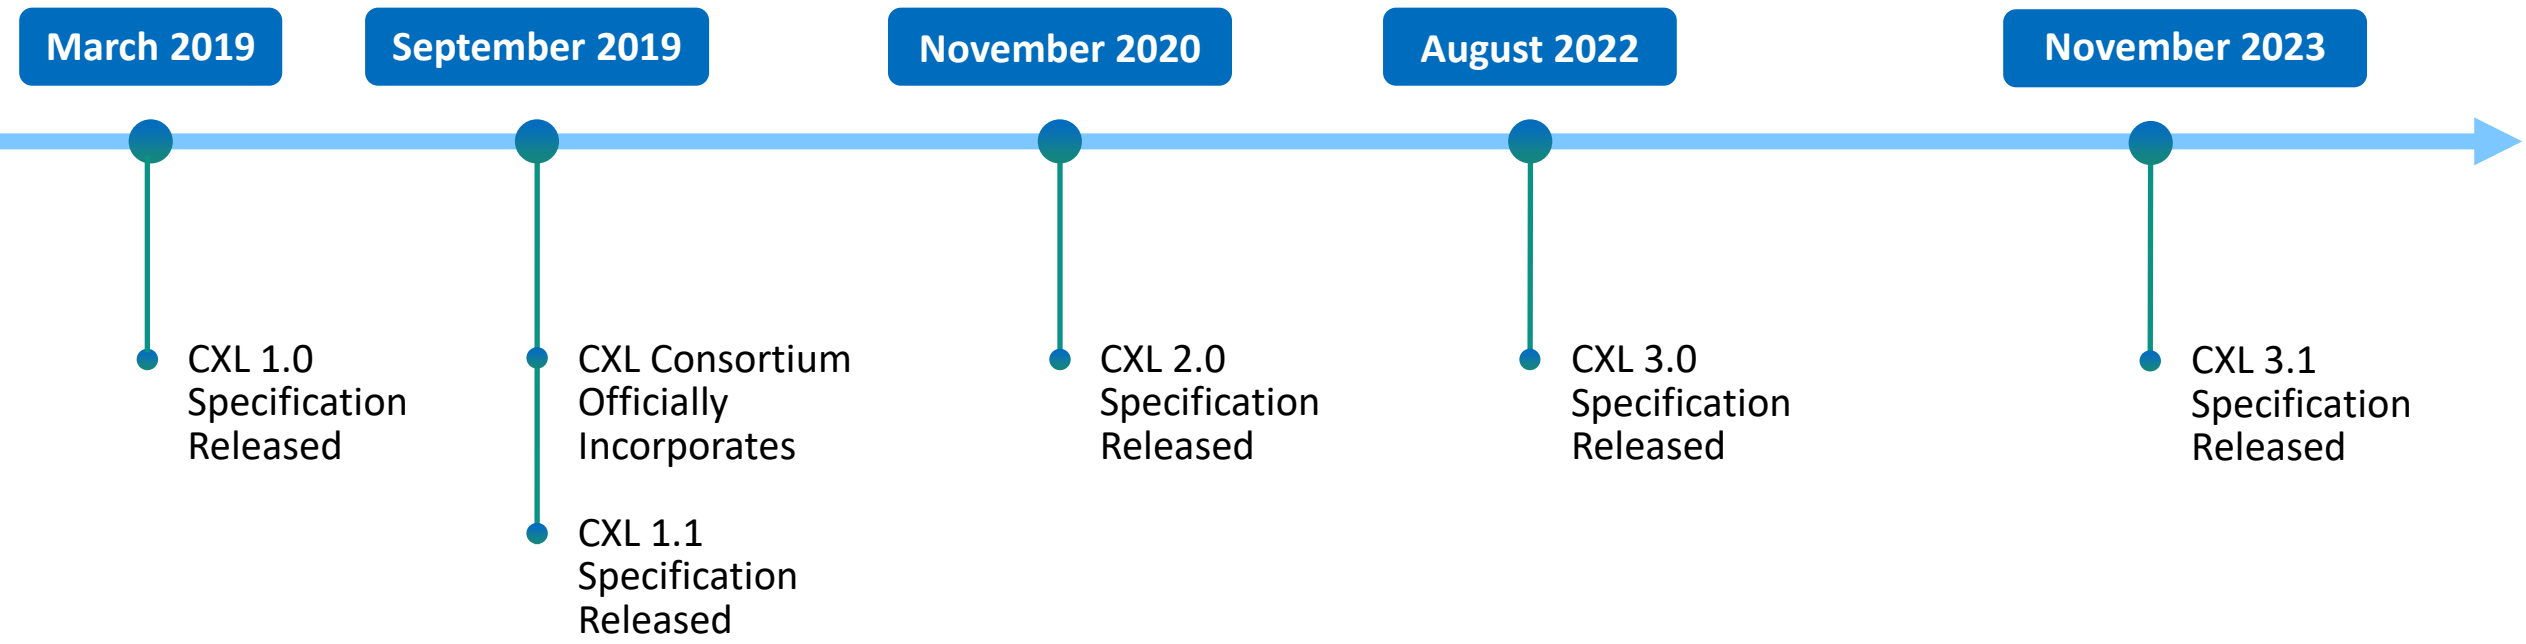

Compute Express Link™ and CXL™ are trademarks of the Compute Express Link Consortium.

In the Beginning

- PCI Express Gen 1.0 was introduced in 2003
  - Serial ATA in 2003 and Serial Attached SCSI in 2004
- Coherent Accelerator Processor Interface (CAPI) in 2014
  - Focused on attaching specialized accelerators to the IBM Power processors
  - Low-latency, coherency connectivity between processor and accelerator memories
  - Leveraged PCIe Gen3 and later PCIe Gen4
  - OpenCAPI consortium created in 2016 to support all ISA's
- Cache Coherent Interconnect for Accelerators (CCIX) in 2018
  - Open cache-coherent interconnect extension to PCIe
  - Simplify the communication between the processor memory and accelerators
  - Leveraged PCIe Gen4, but supported faster non-standard link rates

Stranded Resources in 2014

- NERSC analyzed the workloads on their HPC systems
  - Hopper HPC (2010) : 32GB DRAM per node
  - Edison HPC (2014) : 64 GB DRAM per node
  - Cori HPC (2016) : 96GB DRAM, 16GB HBM per node
- 8% of Edison workloads uses more than 80% of available memory per node
- 16% of Edison workloads would not run on Hoppers 32GB nodes
- 71% of Edison's workloads will fit within Cori's HBM

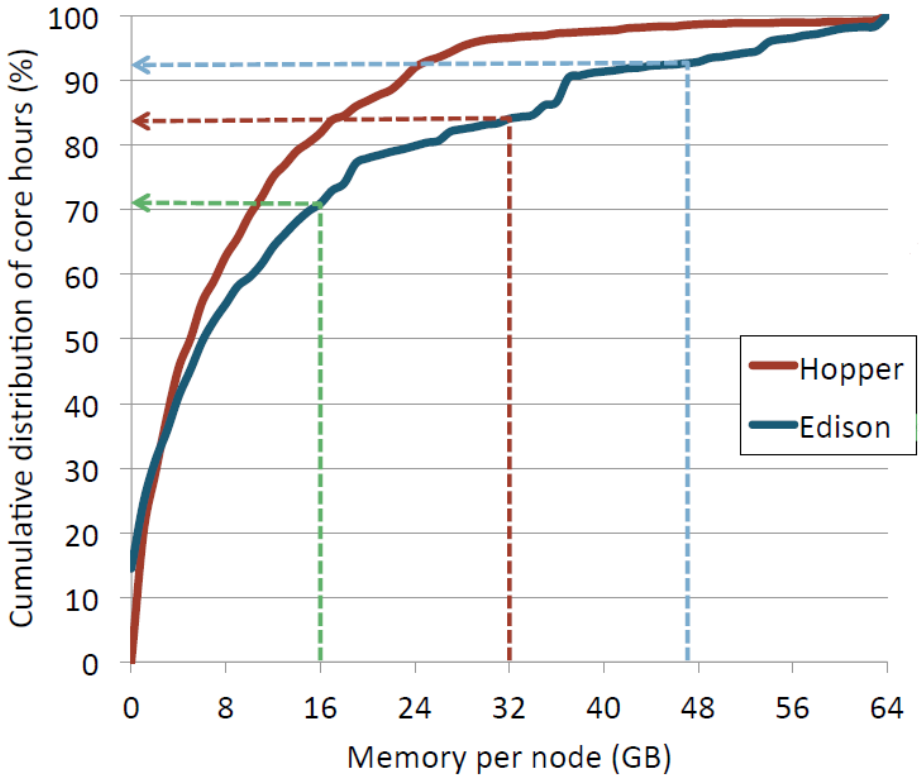

Exascale Program

- DOE launched the Exascale Computing program in 2016 to accelerate the development of exascale HPC
  - Prepare the US industry for exascale systems required for Nuclear Stockpile, Energy, Healthcare, Material Science
  - Acquire 3 HPC systems code named Aurora, Frontier and El Capitan in the 2021 to 2023 timeframe
- The PathForward funding in 2016 focused on HPC connectivity (fabric)
- Gen-Z consortium formed as part of PathForward's efforts
  - Open general-purpose system interconnect
  - Low-latency load/store (shared/pooled) memory, messaging, PCIe
  - New connectors, formfactors, optical

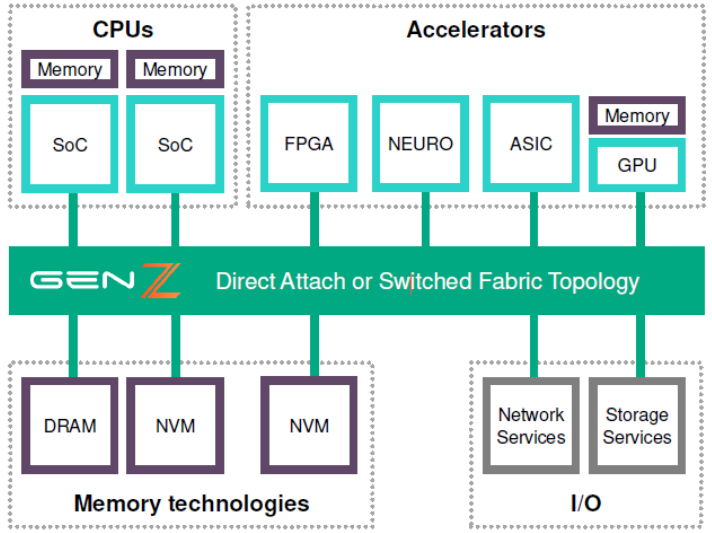

More HPC Milestones

- In 2018, the first CPU/GPU hybrid HPC system, Summit, goes live
  - 9,216 Power9 processors (221K cores) and 27,648 Nvidia Tesla GPU's over 4356 nodes
  - Nvidia NVLink and Mellanox InfiniBand for connectivity
- IBM's next-generation Power9 in 2019 used 16 OpenCAPI serial memory interfaces (OMI) instead of DDR interfaces
  - 2x the memory bandwidth to the core vs DDR
  - Less than 5ns read latency impact
- JEDEC standardized the OMI DIMM (DDIMM) providing greater memory density and improved thermals

CXL Consortium Goes Live

- During the 2016 to 2018 timeframe, numerous companies were proposing new connectivity standards (>10)
- The CXL Founding Members railed around Intel to donated their internal specification to create a public open standard CXL 1.0 and with editorial changes CXL 1.1
- By CXL 1.1 timeframe (2019), all major processor vendors pledged support in their roadmaps

Exascale HPC Achieved

- Frontier first to achieve 1.6 (FP64) exaflops in 2022
  - 9,472 AMD HPC EPYC CPU for 606K cores
  - 37,88 AMD Instinct GPU's for 8.3M cores
  - AMD Infinity and HPE Slingshot connectivity
- Most efficient HPC at 63 GigaFlops/Watts, but still over 20MW total consumption
- New Exascale Computing initiatives aiming for 10x improvement in power consumption for next generation systems

“Proprietary Interconnect”

- Custom connectivity within (rapidly) evolving architectures will always exist
- Two dominant accelerator fabrics are now open to partners
  - Nvidia’s NVLink
  - AMD’s (XGMI) Infinity Fabric
- Industry much more open to finding common ground now…
  - Ultra Ethernet

CXL 3.0/3.1

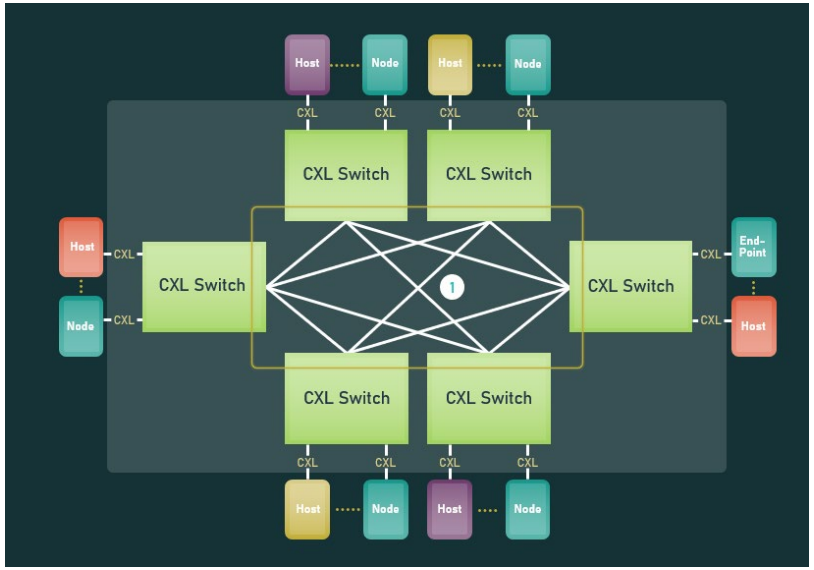

Composable Fabric growth for disaggregation/pooling/accelerator

Final Thoughts

- Believe the rise of Generative AI in 2021 and 2022 disrupted certain CXL opportunities & deployments
- Improving AI fidelity will push the boundaries of compute, memory and power consumption
  - Nvidia Venado 93 AI (FP8) exaflops at ~20MW announced 2023
- “Economies of scale” will drive standardization within the ecosystem long term
  - Including with JEDEC, SNIA and OCP efforts

Please take a moment to rate this session.

Your feedback is important to us.