CXL - Advancing Data Center Architectures with Memory Tiering

Danny Moore,
Senior Product Marketing Manager, Rambus

9-Aug-2023

Core Counts Increasing, AI Models Growing in Size

- Memory systems are becoming a critical bottleneck
- Relentless demand for more memory bandwidth and more memory capacity
- Meta DLRM: >10T parameters (Oct 2022)

Cores are underserviced by available memory bandwidth

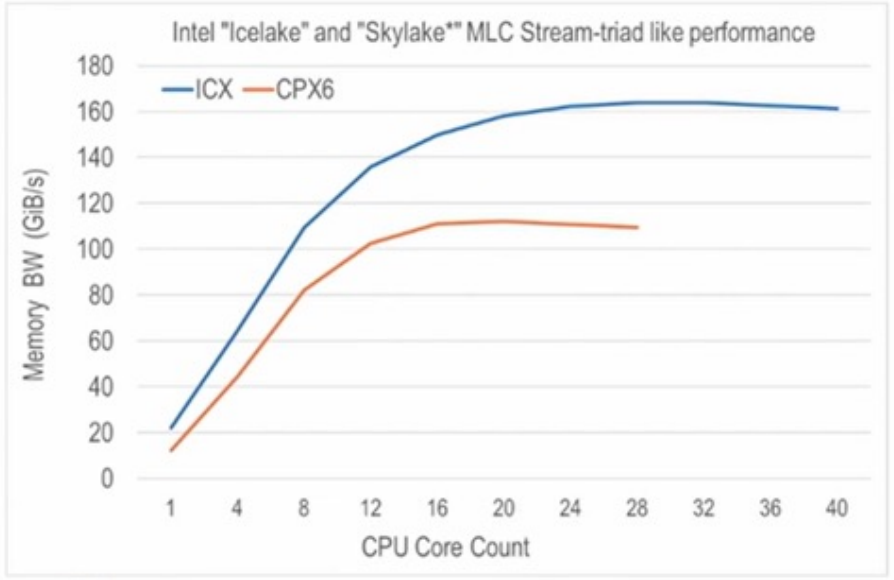

Neural network parameters increasing on the order of 10x year on year

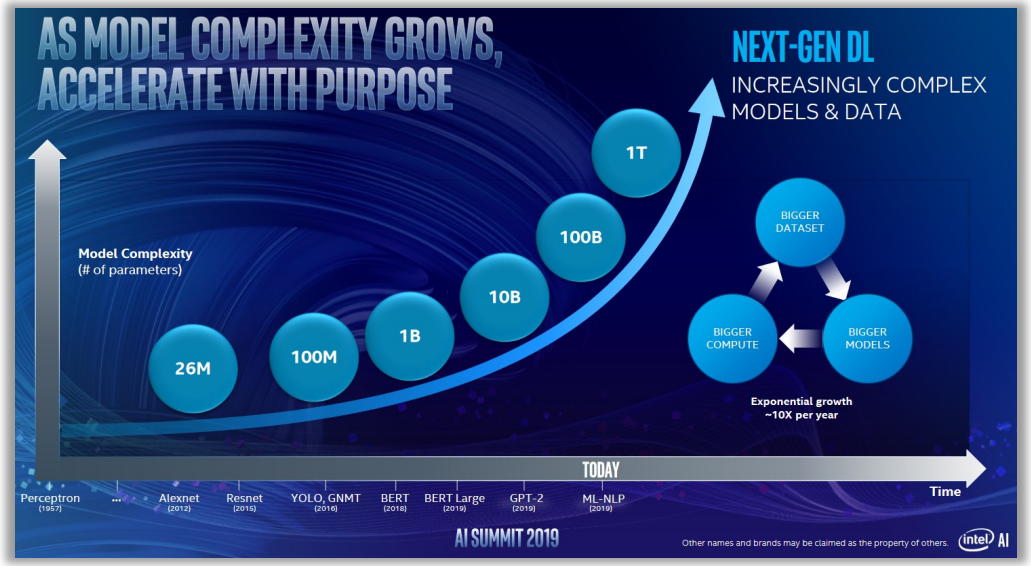

Sources:
Lenovo/VMware, https://www.vmware.com/vmworld/en/video-library/video-landing.html?sessionid=1621034264999001VNUL&videoId=6274685974001
Intel, https://newsroom.intel.com/wp-content/uploads/sites/11/2019/11/intel-ai-summit-keynote-slides.pdf

Summary of Data Center Memory Challenges

Flash Memory Summit

BANDWIDTH

Decreasing memory bandwidth per core

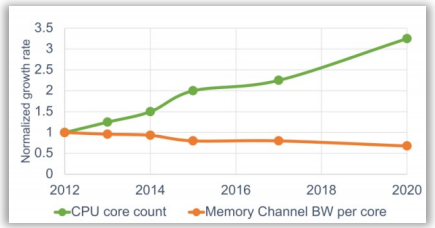

OCP GLOBAL SUMMIT

Source: Meta, OCP Summit Presentation Nov. '21

CAPACITY

Huge latency and capacity gap

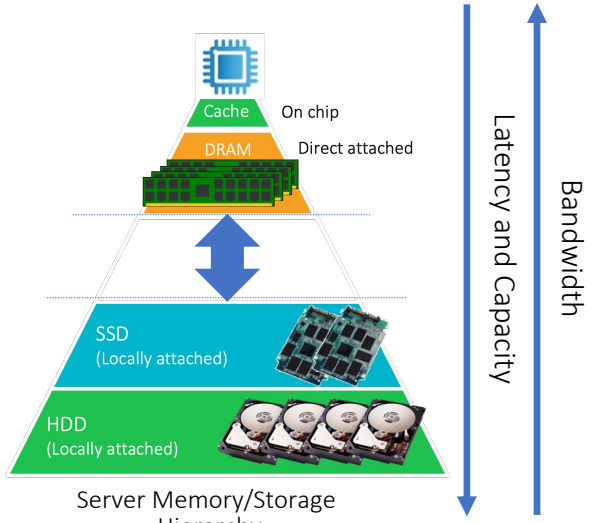

Server Memory/Storage Hierarchy

COST

Costs to achieve desired memory density and bandwidth

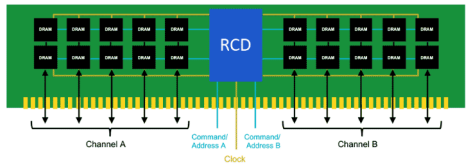

EFFICIENCY

Stranded memory resources and low utilization

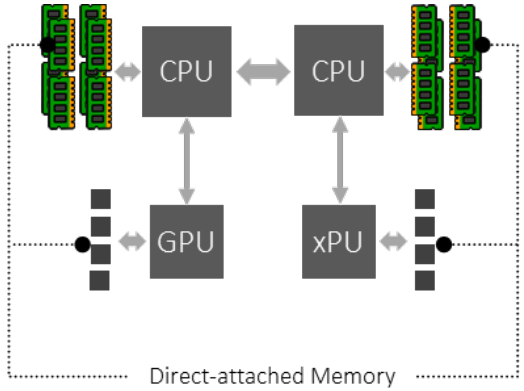

CXL Decouples the Memory Controller from the CPU and Provides Options for New Server Architectures to Address Memory Challenges

CXL Memory Tiers Span the Latency Gap

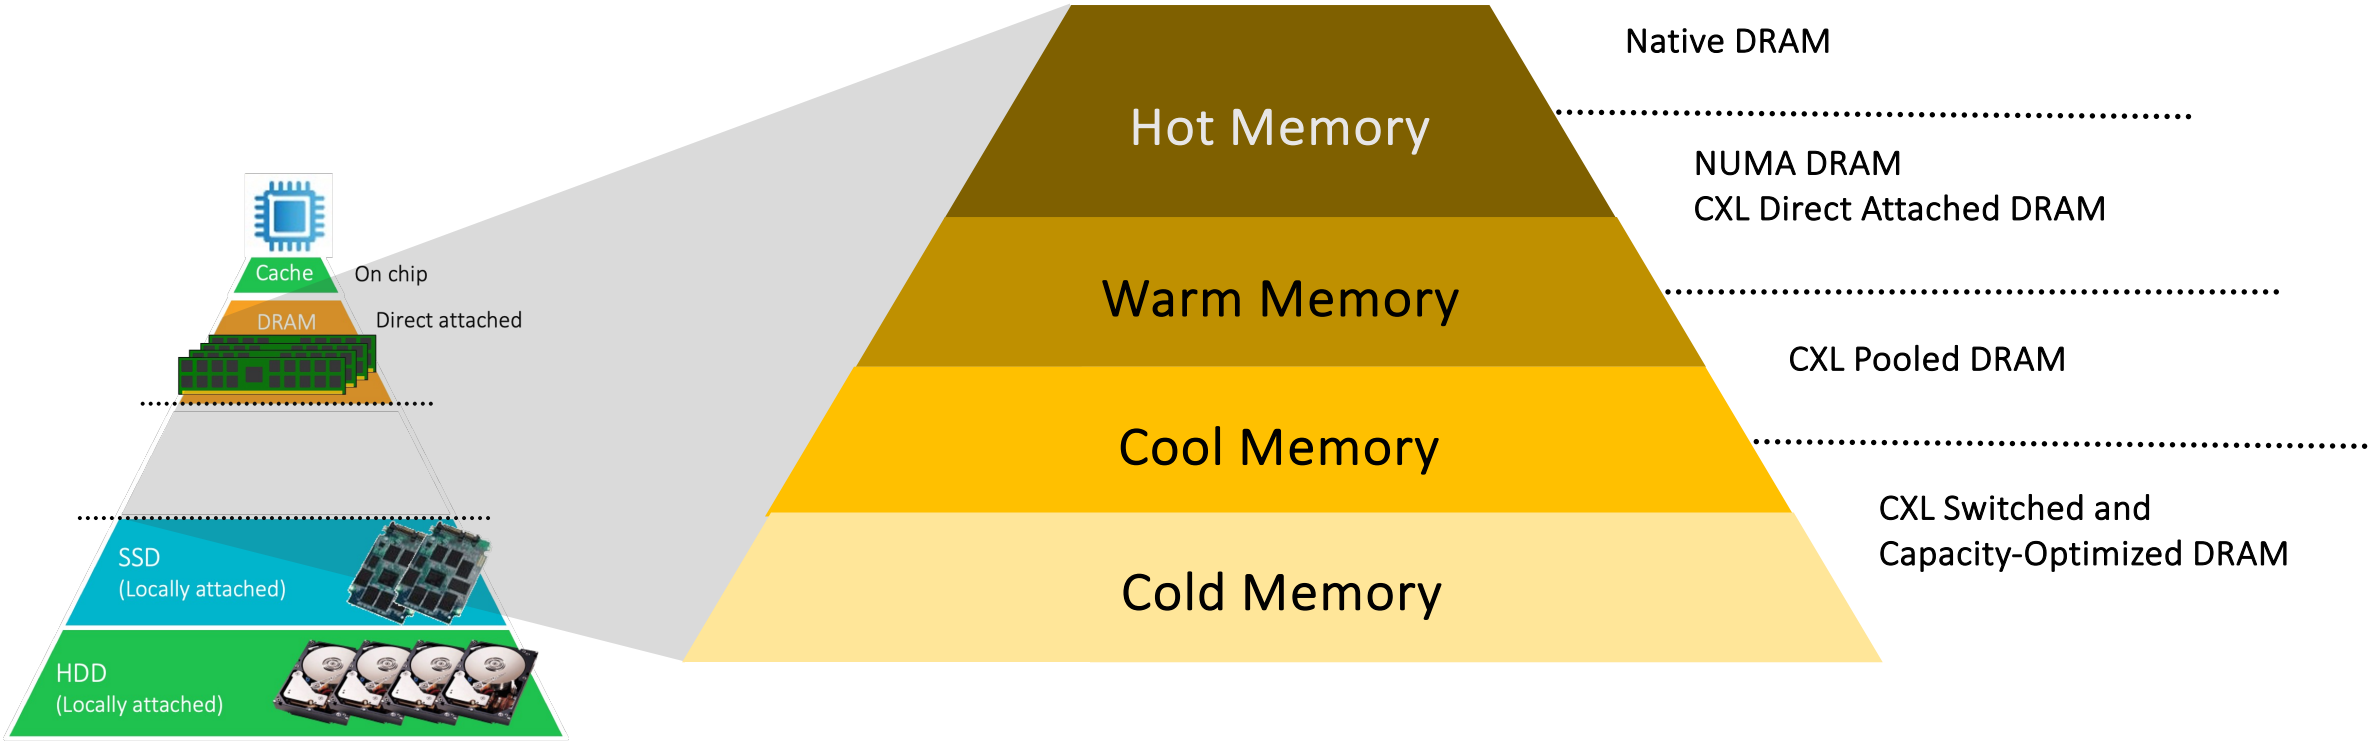

- CXL delivers new expansion options for hot DRAM, with no impact to software applications
- CXL also introduces memory tiering, to the Data Center, much like storage tiering before it
- The industry is now working on software infrastructure to take advantage of these new tiers

The CXL-Enabled Server

Flash Memory Summit

Traditional Server

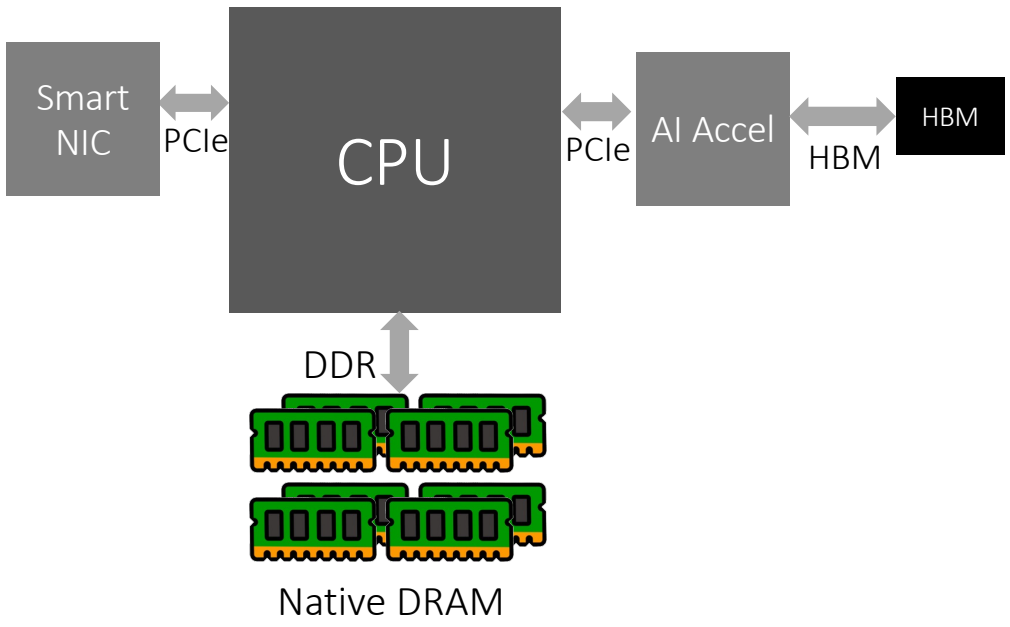

Maximum configuration of 12 DDR5 DRAM channels

CXL-Enabled Server

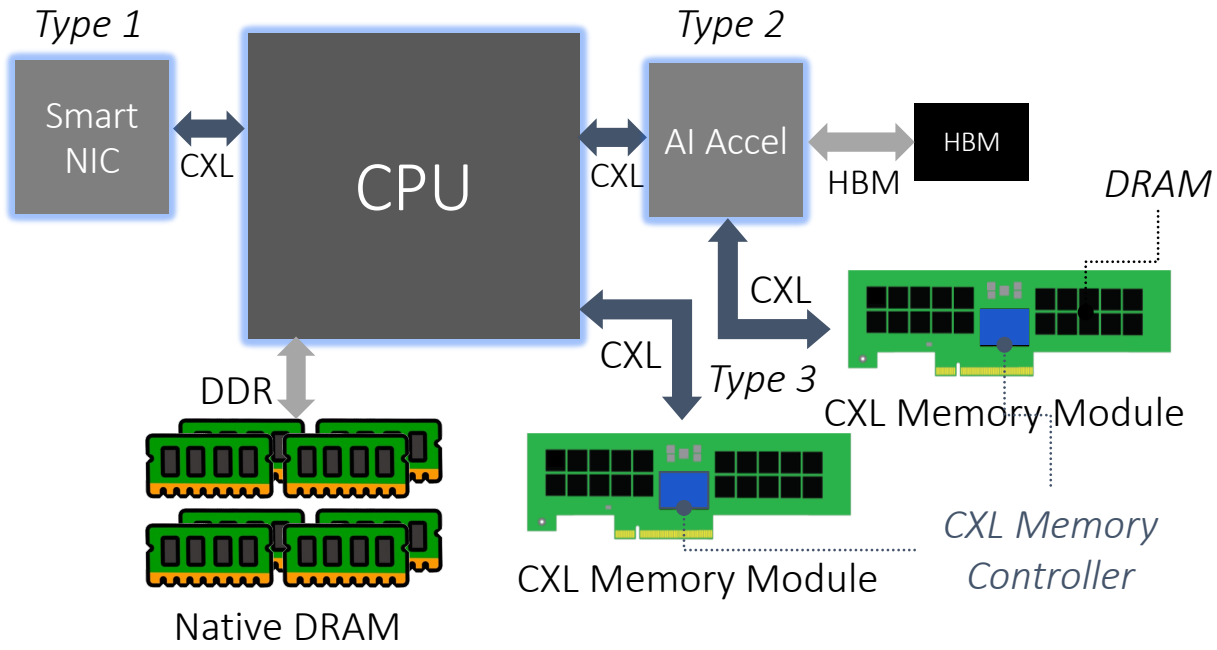

Addition of CXL DRAM can provide >2x the memory capacity and 50% more memory bandwidth

CXL allows for significant memory capacity and bandwidth expansion within a server, leveraging existing PCIe electrical interfaces

Benefits of CXL-Attached Memory

1. Increase memory bandwidth & capacity
2. Improve bandwidth per unit of capacity
3. Media independence
   - For the first time a CPU will be able to utilize a prior generation of DDR memory
4. Improve support for persistent memory technology
5. Lower solution costs
   - Less expensive DRAM
   - 1/3 the pins for the same memory bandwidth

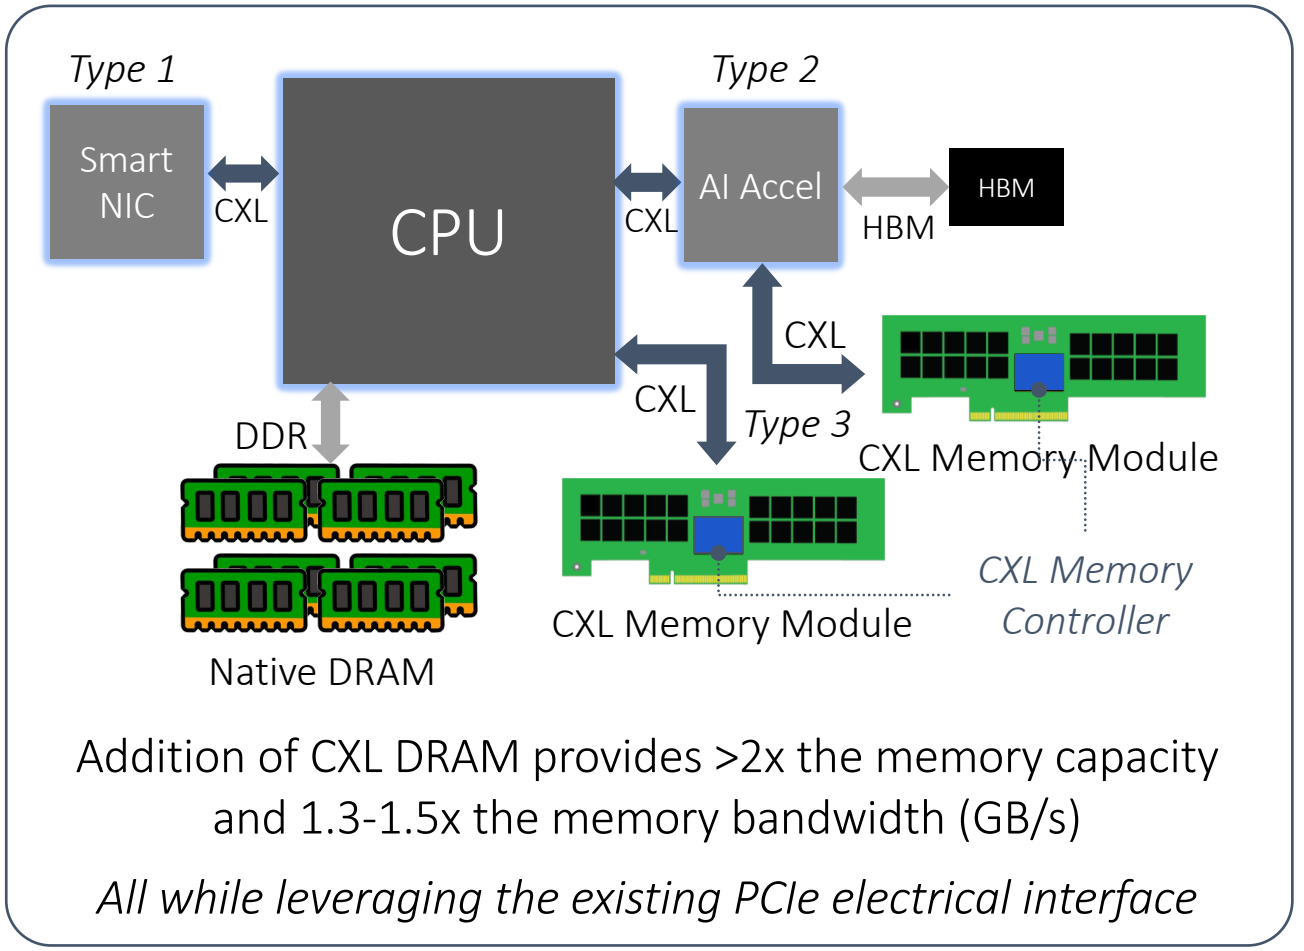

CXL enables new memory alternatives and lower solution costs

What Will these Memory Modules Look Like?

Flash Memory Summit

- CXL EDSFF Memory Modules
  - A standard “front of server” form factor for data center and enterprise servers
  - e.g., E3.S 1T (seen in NVMe SSDs today)
- CXL Add-in Cards (AIC)
  - PCI-SIG CEM or custom AIC variants
  - Allows deployment of standard RDIMMs for memory supplier flexibility
- New Form Factors
  - The industry should expect new standard form factors to emerge which combine some elements of both EDSFF and AIC

CXL EDSFF E3.S

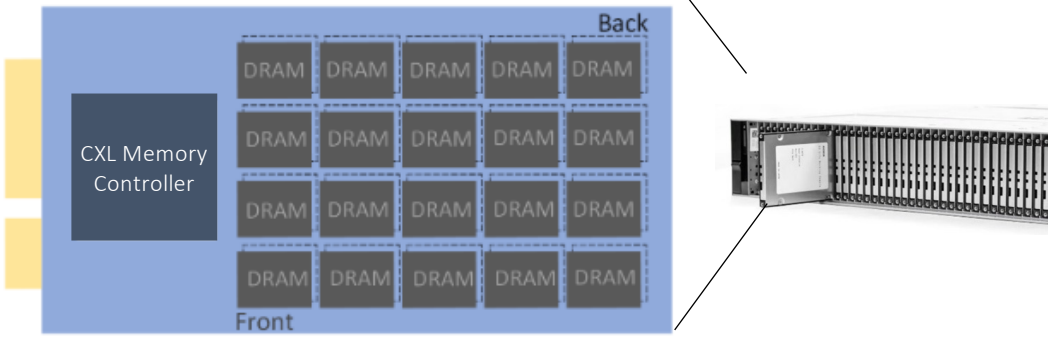

CXL RDIMM ADD-IN-CARD (AIC)

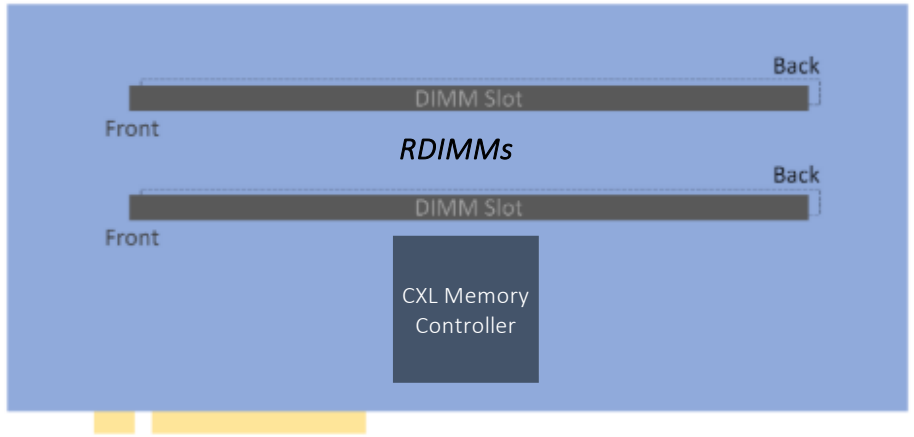

CXL memory modules will come in a variety of form factors, including standard EDSFF

Scaling CXL-Attached Memory

Flash Memory Summit

CXL Memory Pooling

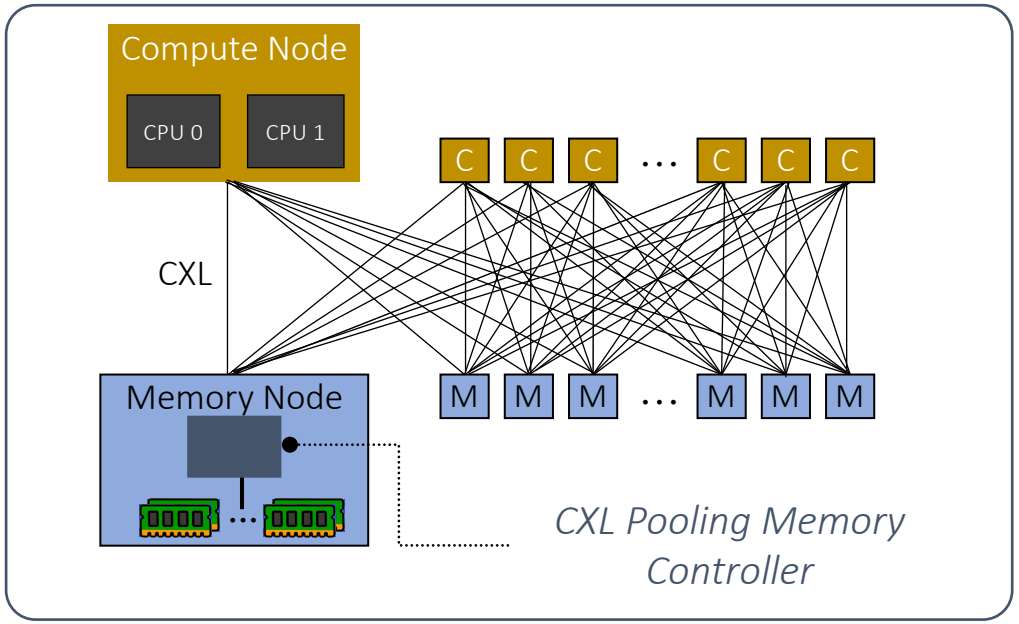

CXL Switch/Fabric-Attached Memory

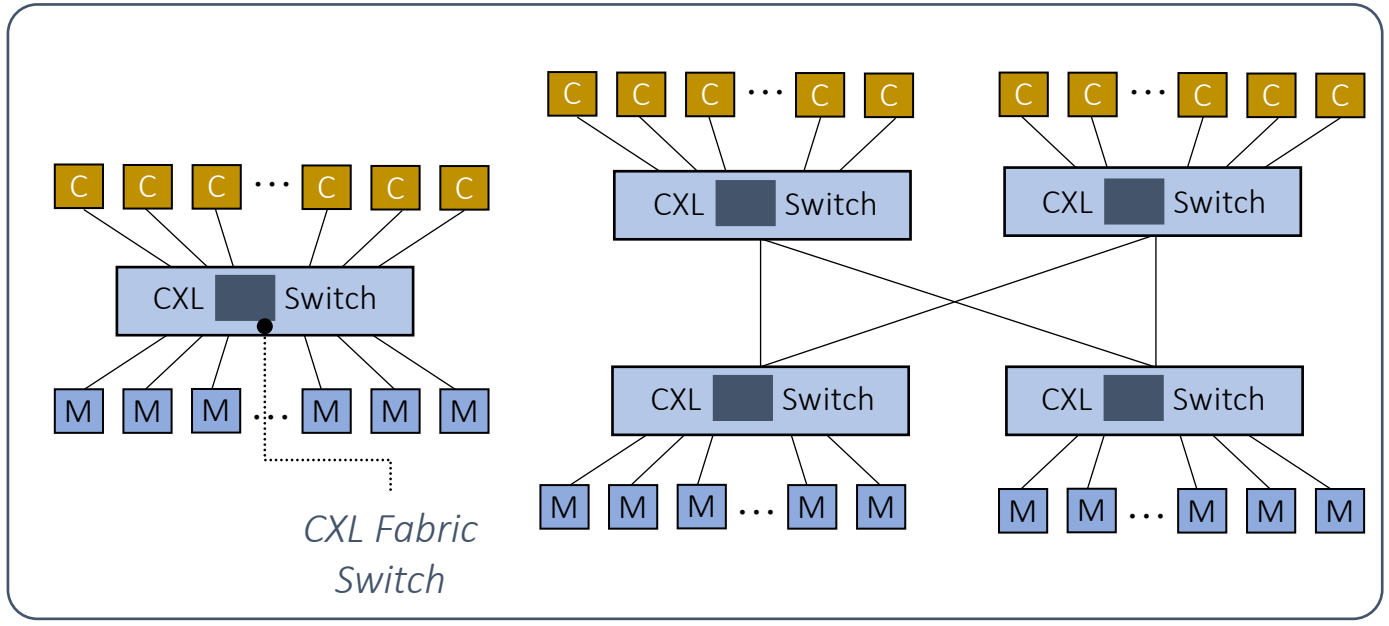

- Reducing over-provisioning, and hence memory stranding, are key objectives of large data centers
- CXL Pooling Memory Controllers → Lowest latency, moderate scale, specific memory types
- CXL Fabric Switches → Highest scale, varied memory media, adds GPUs and NICs, latency penalty

CXL provides mechanisms for CPUs to allocate/deallocate memory from a common pool

Things We Need to Think About as an Industry

Flash Memory Summit

BANDWIDTH

Bandwidth matching between CXL interfaces and memory media

LATENCY

Lowest latency “load to use” memory access

POWER EFFICIENCY

Aligned to server “drive slot” thermals

Save power for memory

SCALABILITY

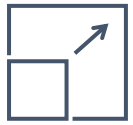

Maximum capacity density

High radix fanout Compression

Quality of service

SECURITY

Encrypted links

Encrypted data at rest

Confidential compute

Root of trust

RELIABILITY

Advanced error detection and correction

Ecosystem suppliers will be differentiating offerings in numerous dimensions

New Reliability Challenges & Opportunities

- Components require Data Center, Enterprise-grade RAS capabilities
  - ECC
  - No silent data corruption
  - Performance monitoring
- Flexibility comes at a price
  - Additional connectors, cabling over direct-attached DRAM
- Live serviceability now possible
  - FRUs must be reported and identifiable

Maintaining or improving user-experience is critical

New Management Challenges and Opportunities

- Leverage existing DRAM management
  - Minimize SW/FW development requirements
  - UEFI CDAT
  - SPD, PMIC, etc.
- New component architectures require extended topology reporting
  - PLDM Type 2 PDRs
  - Redfish models
- Standardization is key to deployment
  - Parallel efforts from CXL Consortium, JEDEC, and DMTF

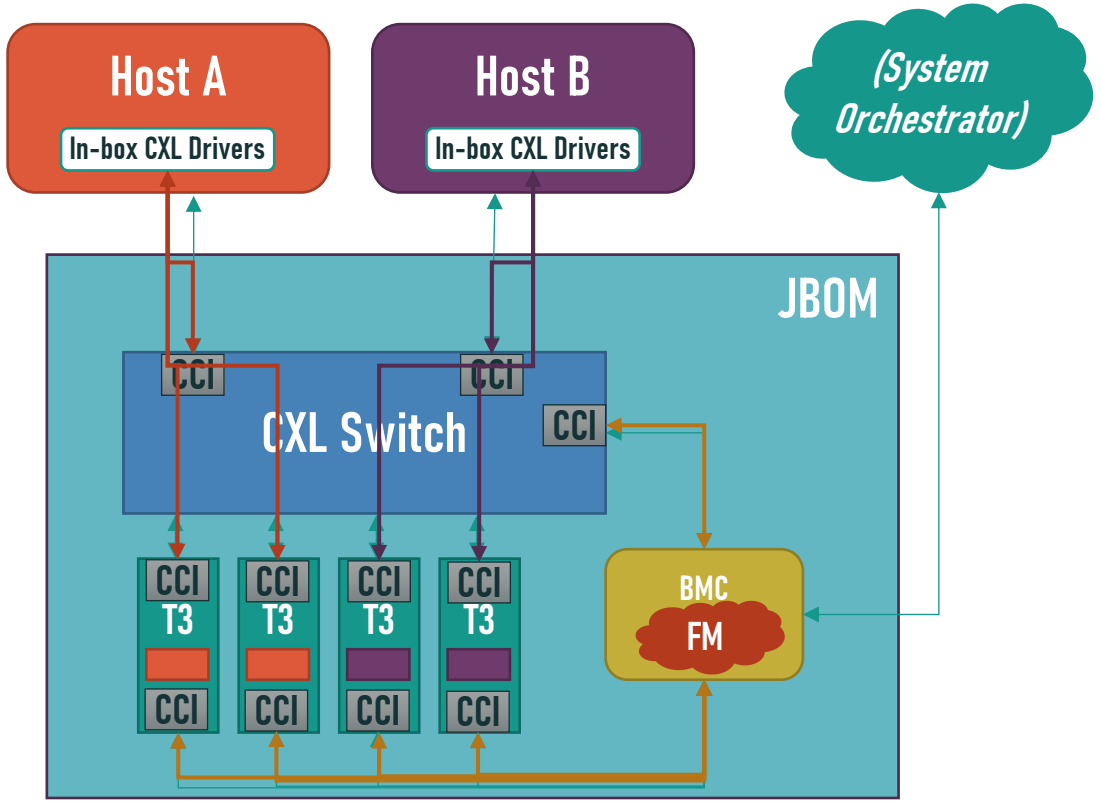

Thank you