Flash Memory Summit

Addressing Memory Bottleneck with CXL™ Type 3 Memory Controllers

Presenter: Ranjit Gupte, Sr. Technical Staff Engineer

Microchip Technology Inc.

MICROCHIP

Agenda

Flash Memory Summit

- Problem Statement
- Addressing Memory Bottleneck with CXL™ High-Speed Serial Interface
- Introduction to CXL and CXL Type 3 Devices
- CXL-Based Memory Solutions
- CXL Memory Controller Requirements
- Conclusion

MICROCHIP

Problem Statement

- CPU vendors are introducing powerful CPUs with multiple cores and threads at a rapid pace
- DDR Memory is unable to keep up with the advances in the CPU space due to various reasons
- Amount of DDR memory that could be directly attached to CPU is limited due to
  - Number of parallel DDR interfaces/channels per CPU
  - Number of IO pins
  - Adding more interfaces will impact CPU footprint

MICROCHIP

Typical Host to DDR Interface

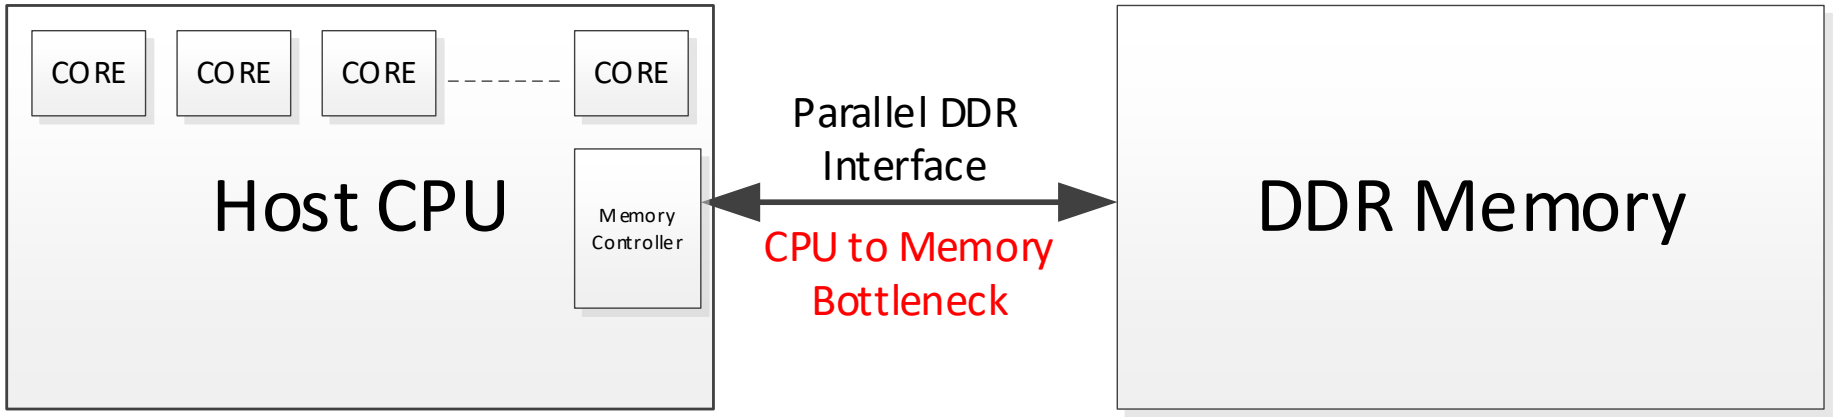

- Higher core count processors are the future
- Upgrading memory is challenging due to limited number of I/O pins
- Stranded CPU cores: memory unable to keep up with CPU's compute capability which causes bottleneck

Host to DDR Interface Improvements

- Low-latency serial-attached memory expansion with coherent memory space between the host processor's memory and any memory attached to a memory expander device alleviates CPU to memory bottleneck
- Decouples the chipset from the DDR memory type
  - For example, next-gen hosts will only support DDR5, but a hyper-scaler could expand the host memory by reusing existing DDR4 from the previous-gen systems

DDR4 DIMMs:

+

Additional memory could be DDR, HBM or any other high-bandwidth memory

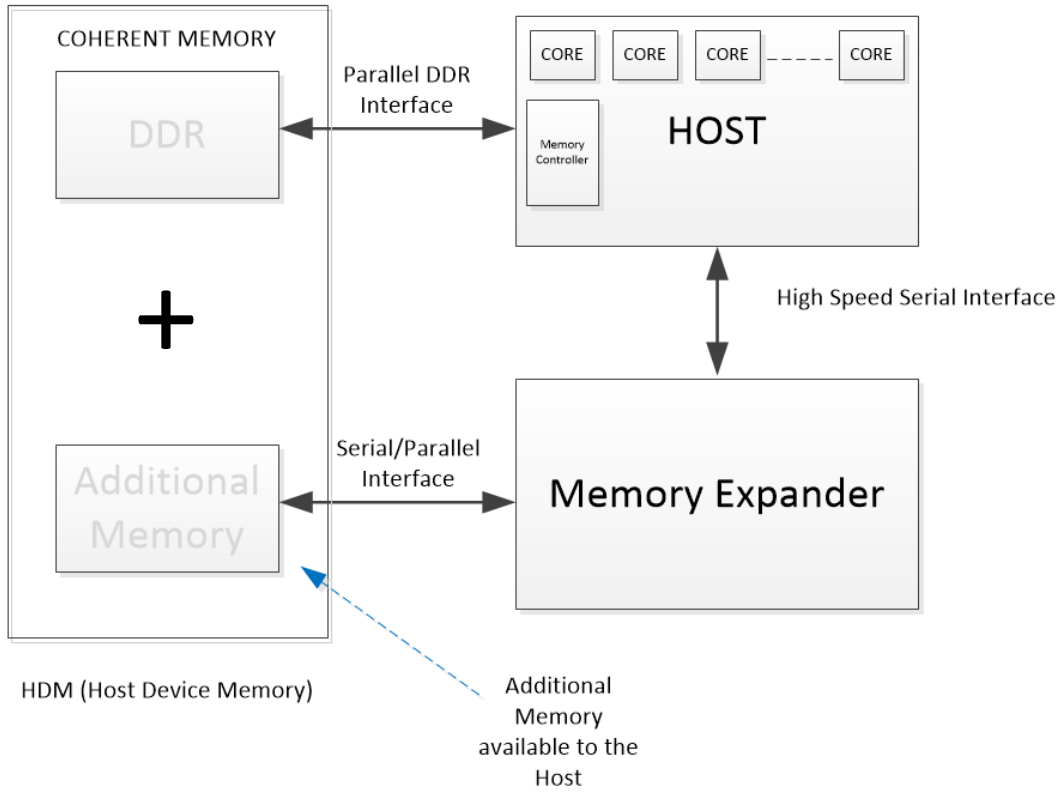

MICROCHIP

Introduction to CXL™

- CXL™ is an emerging open industry standard based on PCIe® 5.0 infrastructure
  - High bandwidth, low latency interconnect for connectivity between host processors and accelerators/memory devices/smart NIC
  - Targets high-performance computational workloads like AI, Machine Learning, Comms, etc.
  - CXL 1.0 --> CXL 1.1 --> CXL 2.0 --> CXL 3.0
- Dynamic multi-protocol capability that enables new support models
  - **CXL.io** – Mandatory PCIe-based protocol for initialization, discovery, register access, interrupts, I/O Virtualization, DMA
  - **CXL.mem** – Protocol to support memory semantics, optimized for latency
  - **CXL.cache** – Protocol to support caching semantics, optimized for latency
- CXL™ runs on PCIe 5.0® PHY. Primary Data Rate: 32 GT/s (128b/130b)
- Plug and Play: Either PCIe or CXL card can be plugged in
  - Link negotiation for PCIe or CXL occurs during link training

MICROCHIP

CXL™ Stack – Designed for Low Latency

- CXL IO transactions pass through a stack that is largely identical to PCIe® stack
- CXL cache and memory stacks are optimized for latency, use separate transaction and link layers from CXL IO
- Based on the QoS mechanism, static ARB/Mux decides which protocol to schedule

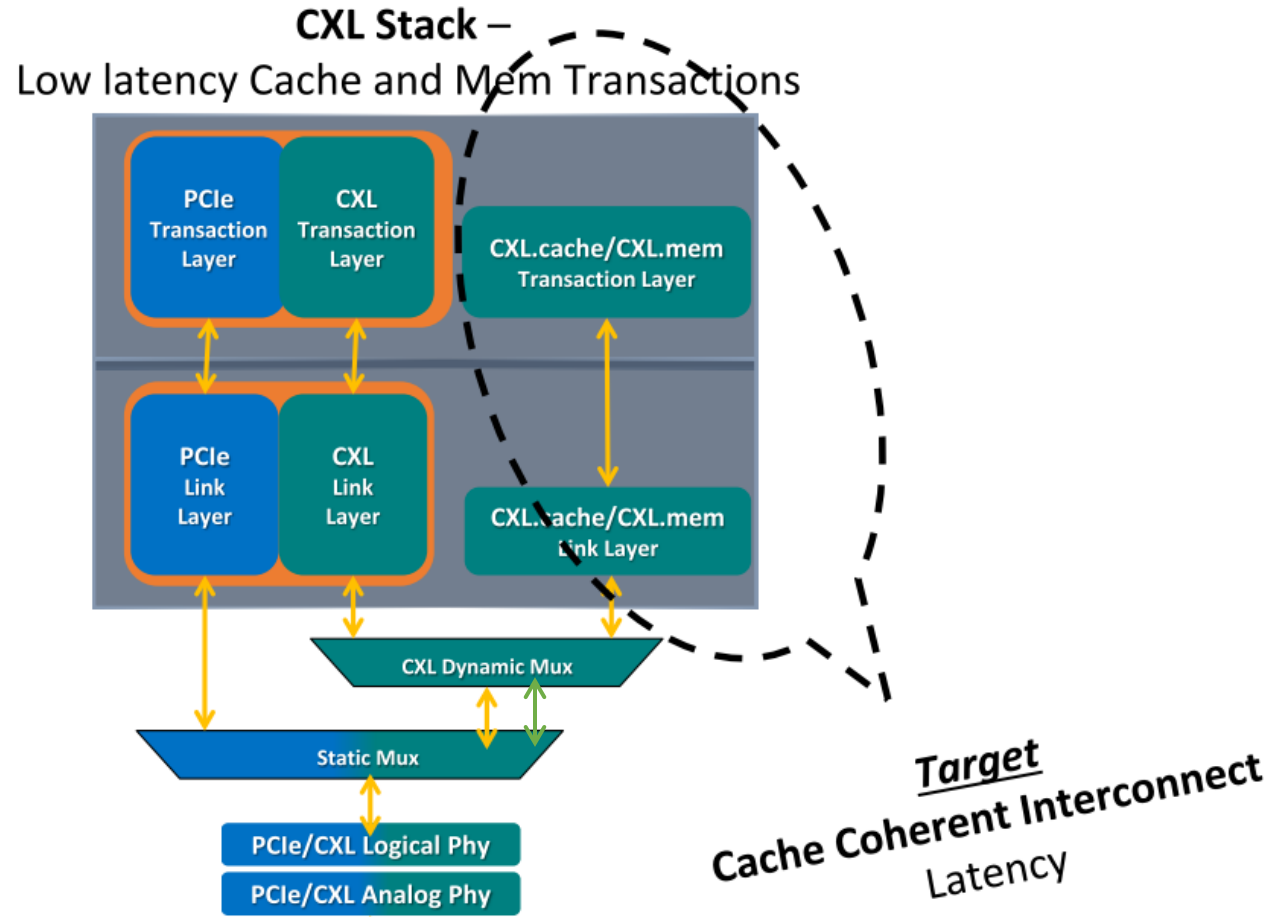

CXL™ Device Types

- CXL consortium has defined three classes of devices to support various CXL based applications

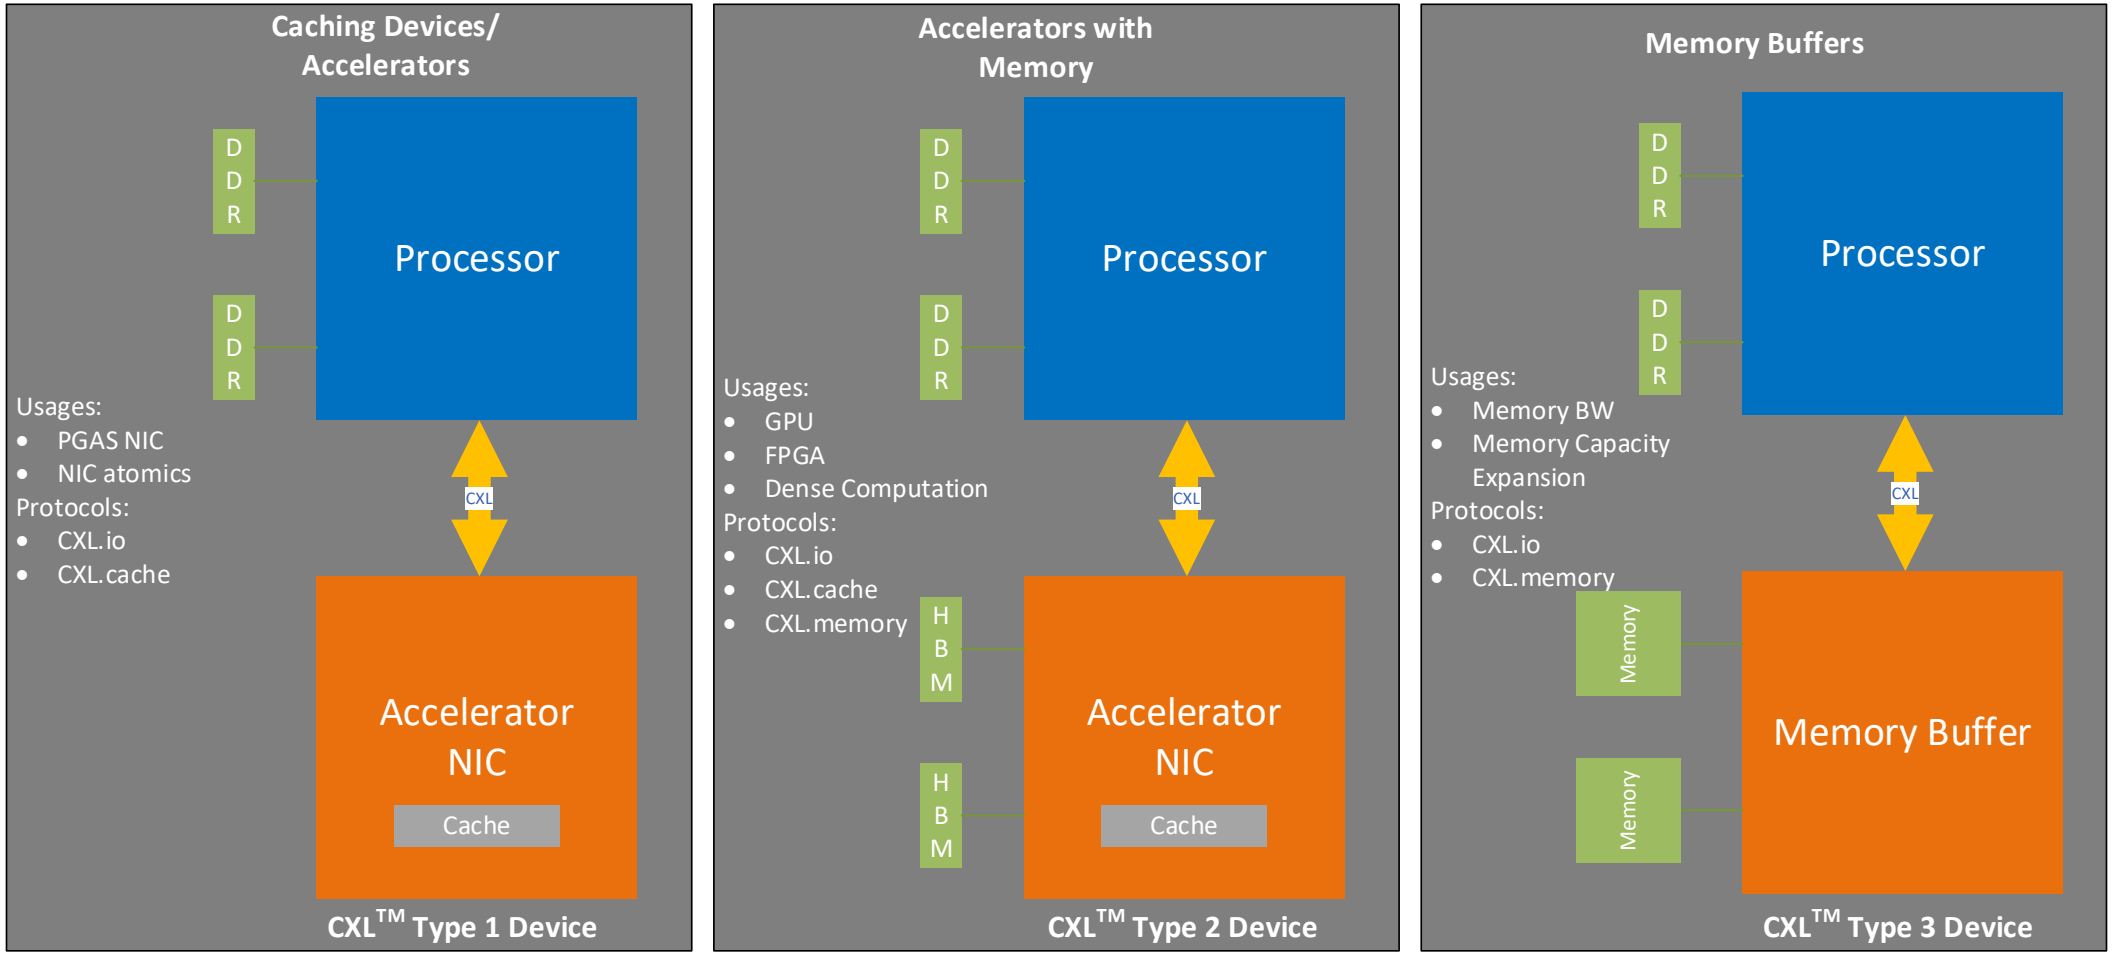

CXL™ Type 3 Devices

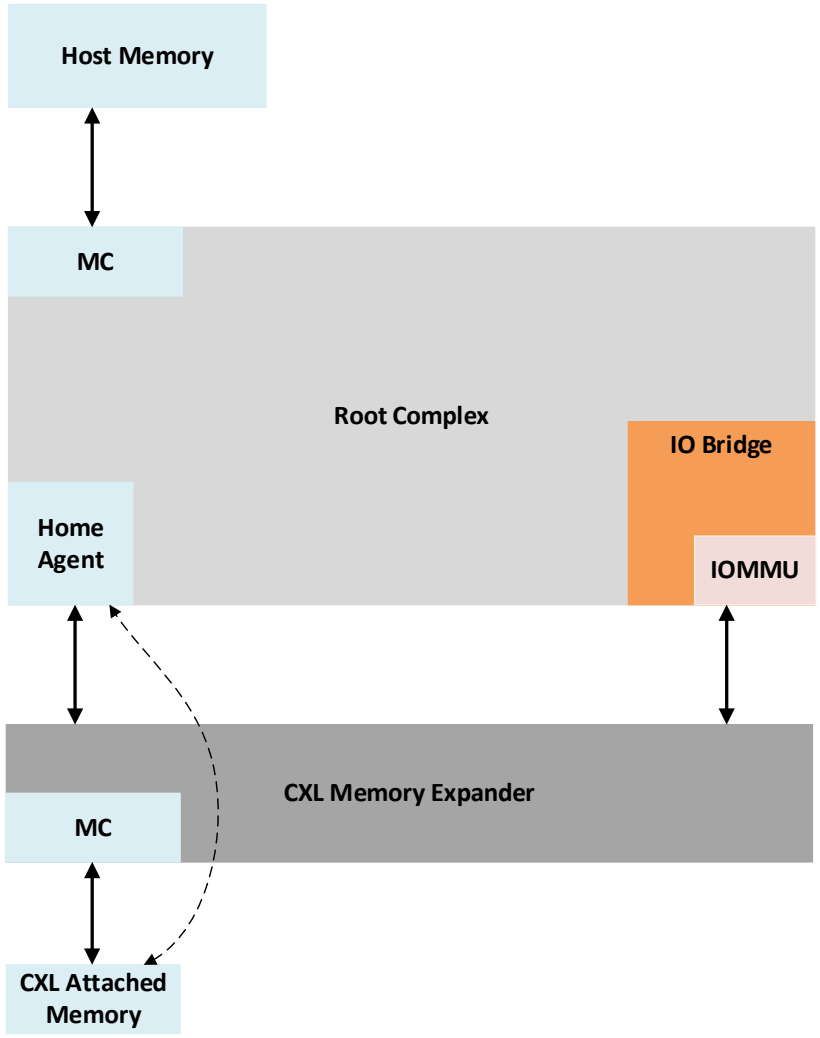

Type 3 CXL Device

- Type 3 devices are typically memory buffers or memory expanders
- Host can access and manage Device Memory coherently
- Uses CXL.io and CXL.mem

Next Gen High Computing Servers

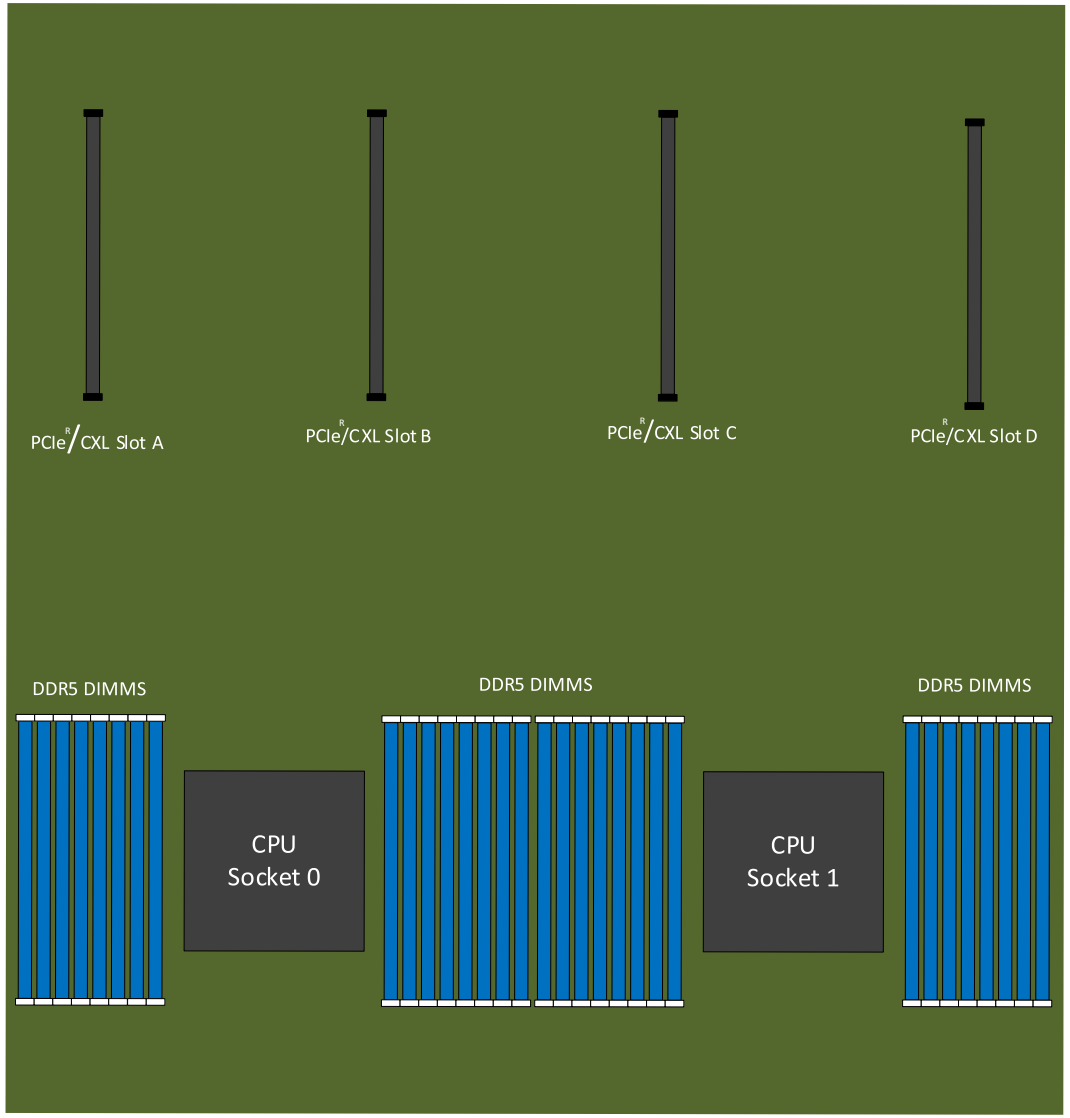

- Two CPU sockets
- 52 cores per socket
- Four memory controllers per socket
- Eight memory channels per socket
- Two DIMMS per channel (two DPC)
- 16 DDR5 DIMM slots per Socket
- 32 total DDR5 DIMMs
- 64GB x 32 = Total capacity of 2.048 TB
- Multiple PCIe® Gen 5/ CXL™ slots for CXL memory expansion

Solutions For CXL™ Memory Expansion (AIC)

Flash Memory Summit

Add-in Cards

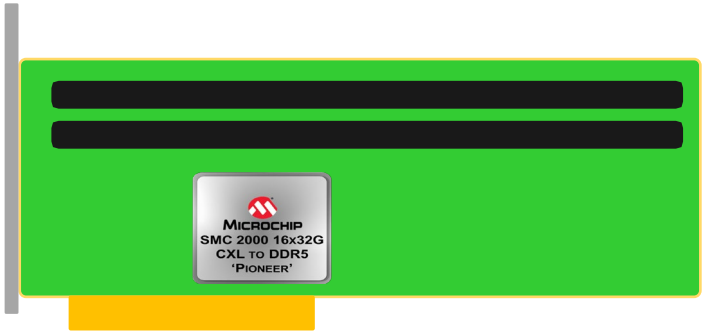

Add-in Card with DIMM slots

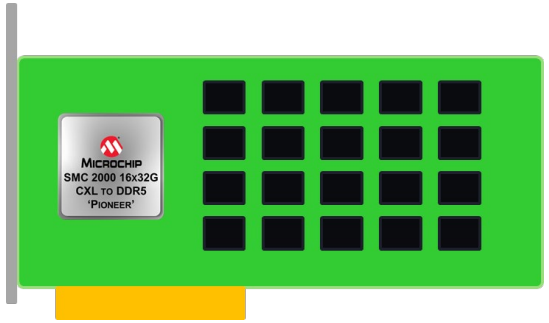

Add-in Card with discrete DRAM devices

- AIC Form Factor for CXL Memory Expansion
- There are two potential variations for AIC solutions:
  - AIC with DIMM slots
  - AIC with discrete DRAM devices
  - These AIC can be plugged into CXL slots to seamlessly expand the memory on as needed basis

MICROCHIP

Solutions For CXL™ Memory Expansion (EDSFF)

Flash Memory Summit

EDSFF Form Factor

EDSFF E1.S

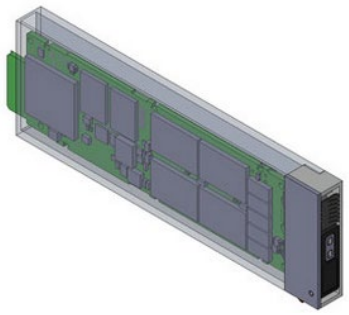

EDSFF E3.S

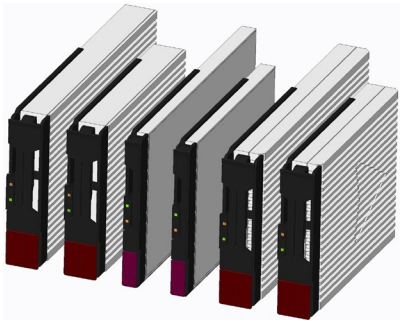

- Solid State Drives (SSDs)

are commonly used in enterprise and hyperscale data centers for storage memory

- Newer SSDs can be built in various form factors now
  - M.2, 2.5-inch U.2, EDSFF
    - EDSFF displacing M.2 and U.2
  - Apart from being used as storage devices, these form factors are also being used for emerging CXL Memory Expansion Technology

MICROCHIP

Server with E1.L CXL™ Memory Modules Fitted in 1U Chassis

CXL™ Attached Memory Controller Requirements

- Low latency interface
- Power-efficient
- Cost-effective
- Data integrity
- Security
- Rich feature set like direct attached DDR
  - Error correction (SECDED, Chipkill)
  - Patrol scrubbing and demand scrubbing
  - Thermal throttling capability
  - sPPR, hPPR

MICROCHIP

Conclusion

- Compute Express Link™ (CXL™) leverages the existing PCI Express ® (PCIe®) electrical structure to enable the expansion of memory resources in high-performance computing applications like Artificial Intelligence (AI) and Machine Learning (ML).

- With CXL, memory expansion can be achieved by connecting additional memory devices (such as DDR DRAM, HBM or persistent memory) to the CPU via CXL memory controllers. The CXL memory controllers allow the system to access the additional memory as if it were part of the CPU's native memory pool, resulting in increased memory capacity and bandwidth.

MICROCHIP