Flash Memory Summit

Challenges of Memory Diversity in CXL™ Ecosystem

SARC-203-1

Sanketh Srinivas,

Product Manager, CXL Memory Controllers

MICROCHIP

MICROCHIP

Agenda

- Traditional System Architecture
- Introduction to CXL™ and Memory Diversity enablement
- Memory Tiering Model with CXL
- Memory Tiering with CXL direct-attached Memory Expansion
- Memory Tiering with CXL Switch-attached Memory Expansion
- Memory Tiering with CXL Fabric-attached Memory Expansion
- Memory Tiering Enablement
- Host Based CXL Memory Tiering Enablement
- Memory Controller Based CXL Tiering Enablement
- Software Based CXL Memory Tiering Enablement
- Hybrid CXL Memory Tiering Enablement
- Conclusion

Traditional Memory Tiering Model

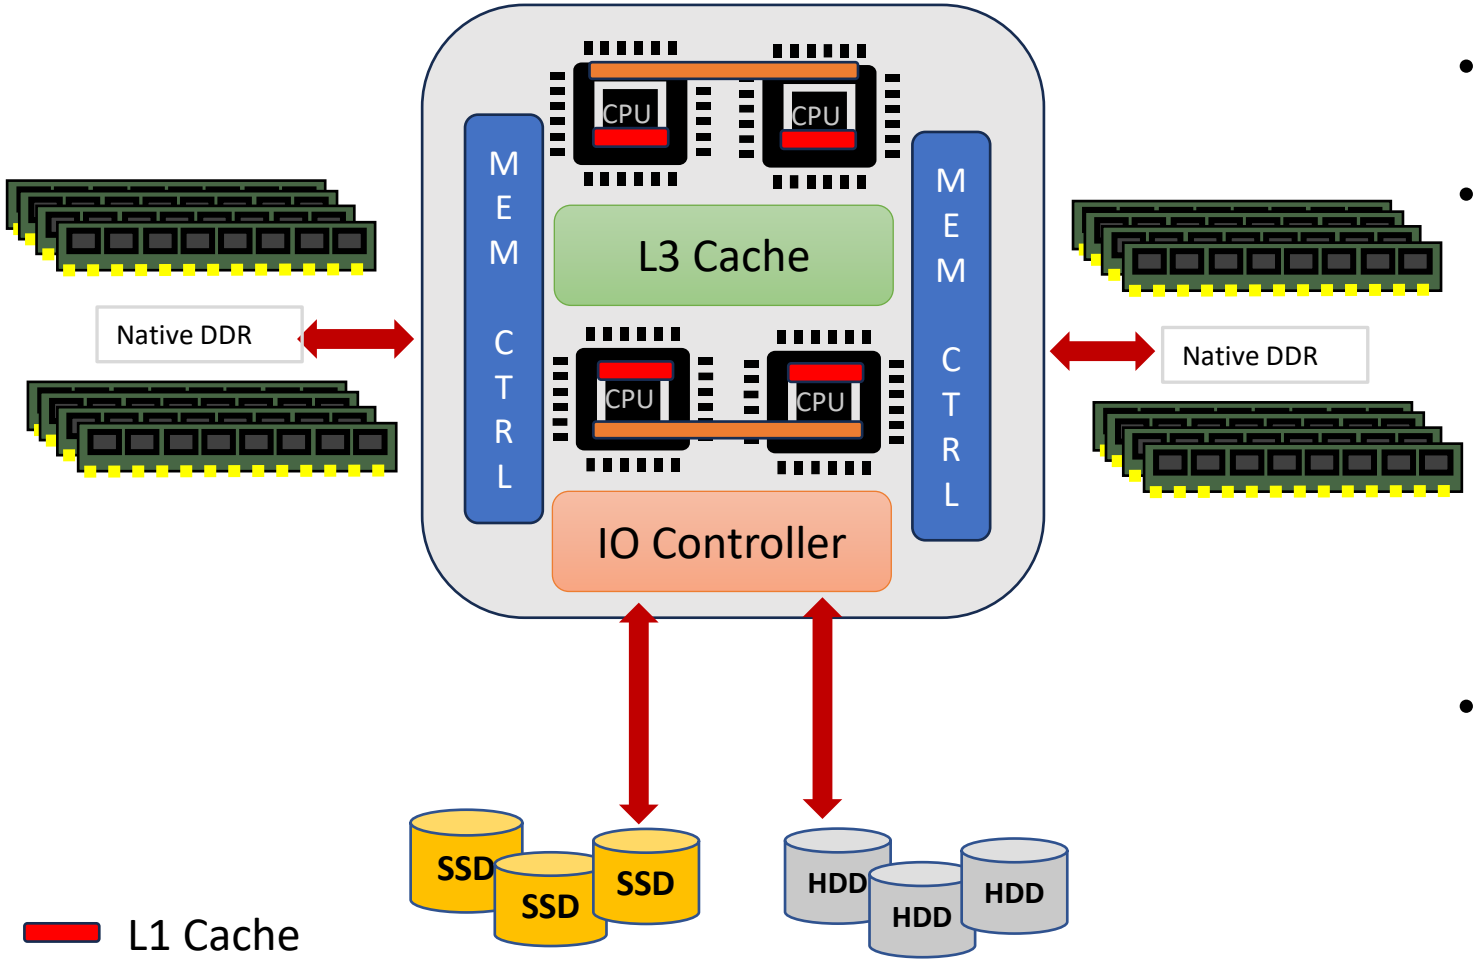

- Traditional Memory model is unable to keep up with today's workloads
- Amount of DDR memory that could be directly attached to a CPU is limited due to:
  - Limited number of parallel DDR interfaces/channels per CPU
  - Limited number of IO pins
  - Adding more interfaces will impact CPU footprint
- Growing need for data storage, analysis, processing and inference with onset of AI/ML and Heterogenous compute capabilities

Traditional Memory Tiering Model

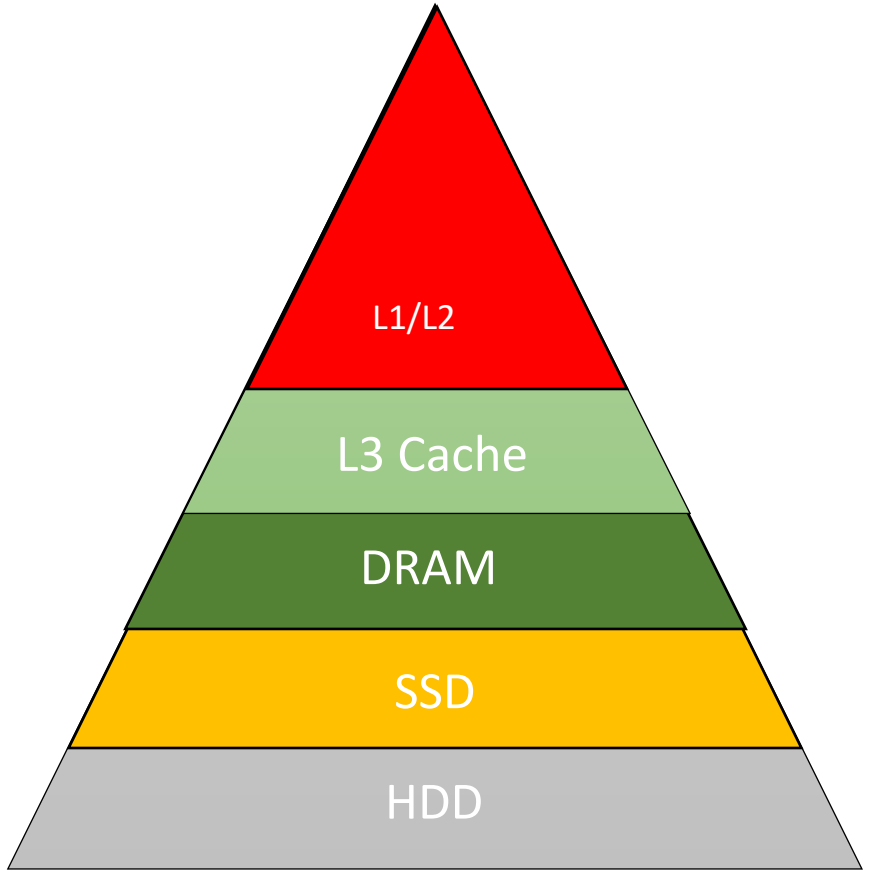

Introduction to CXL™ and Memory Tiering Enablement

- CXL is an emerging open Industry standard based on existing PCIe® 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, Communication, Hyperscale, Datacenter, etc.
- Dynamic multi-protocol capability that enables new support models

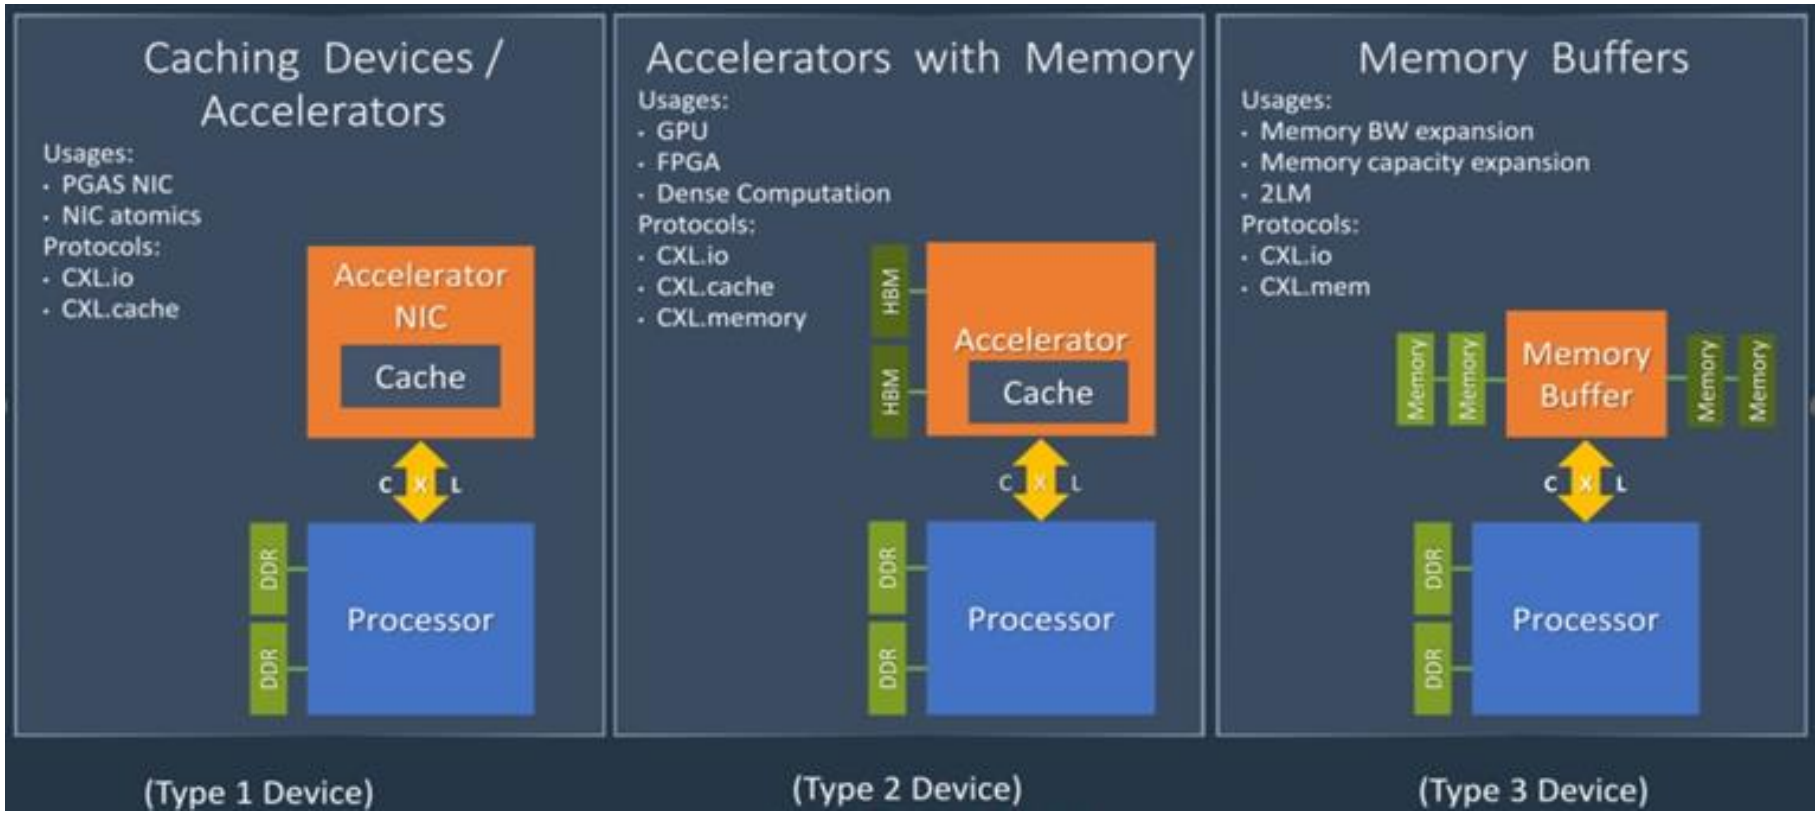

Memory Tiering with CXL™

CXL enables **Memory Diversity** with the same look and feel abstracting different memory types to be supported behind a memory controller

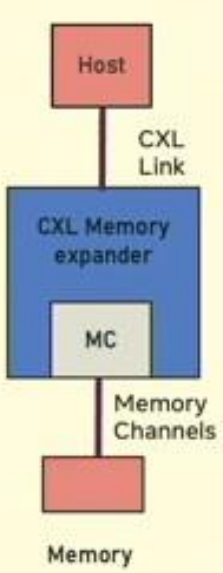

Direct attached

Add Capacity
Add Bandwidth
Slower-cheaper tier

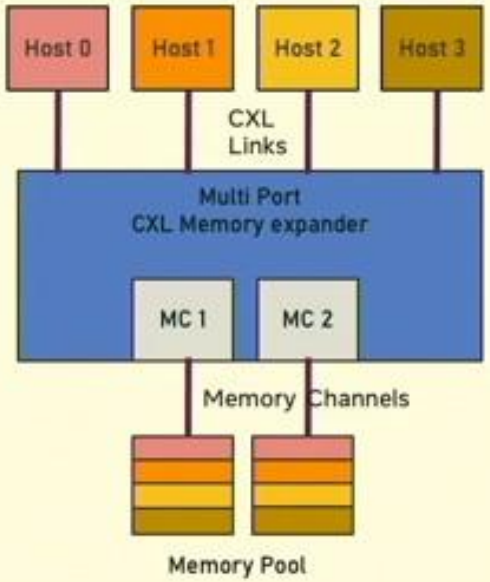

Pooled Memory

Amortize CXL infra cost
Flexible allocation

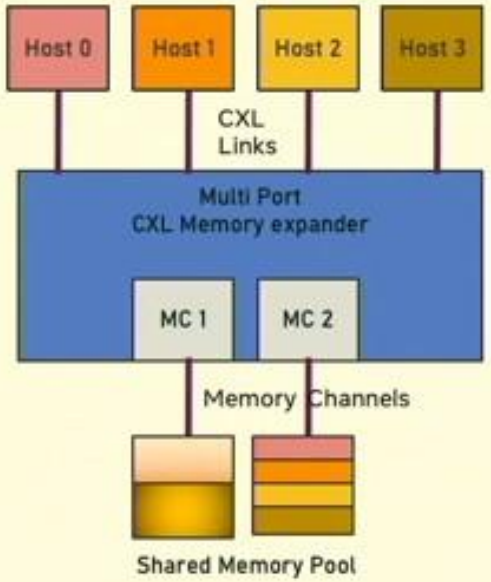

Shared Memory

Deduplication
Host2host communication
large datasets

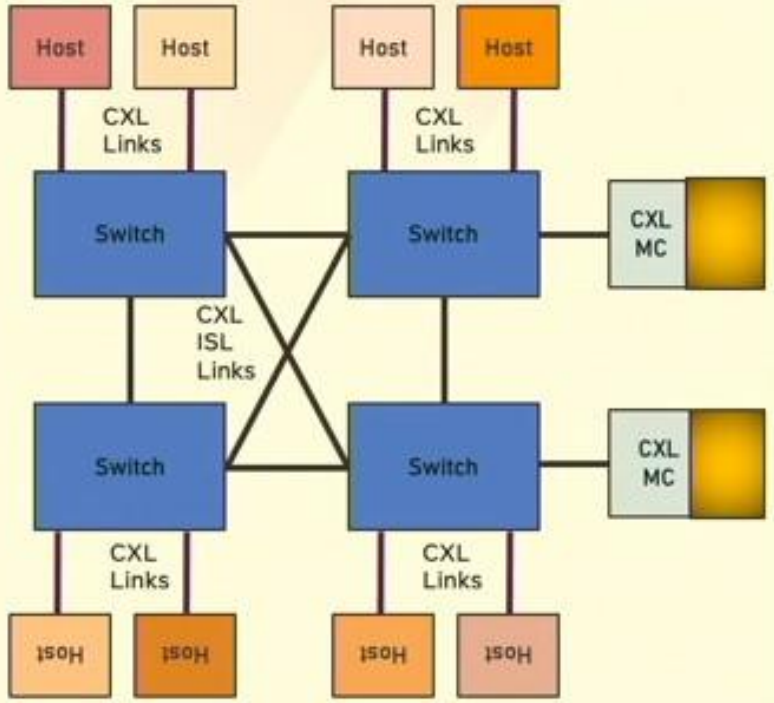

Fabric Memory

Scaling to huge datasets

Memory Tiering Model with CXL™

Flash Memory Summit

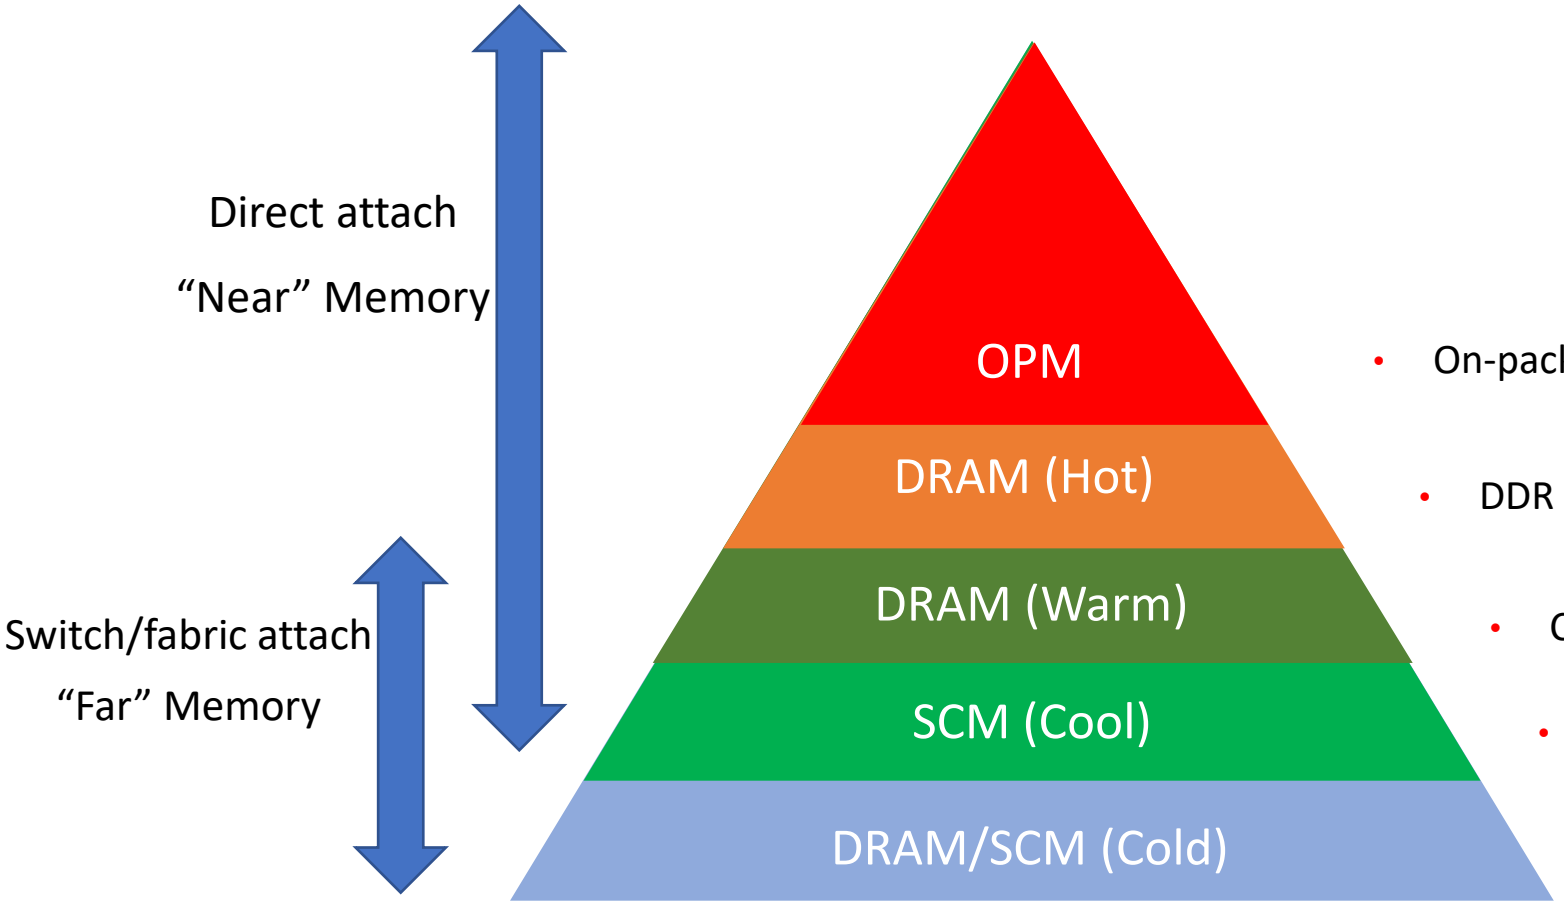

- On-package memory/High Bandwidth Memory (HBM)
- DDR direct-attached
- CXL direct-attached
- CXL Switch-attached
- CXL Fabric-attached or GFAM

Microchip Memory Infrastructure Products will enable memory tiering to optimize performance and cost for application-dependent workloads

MICROCHIP

Memory Tiering with CXL™ Direct-Attached Memory Expansion

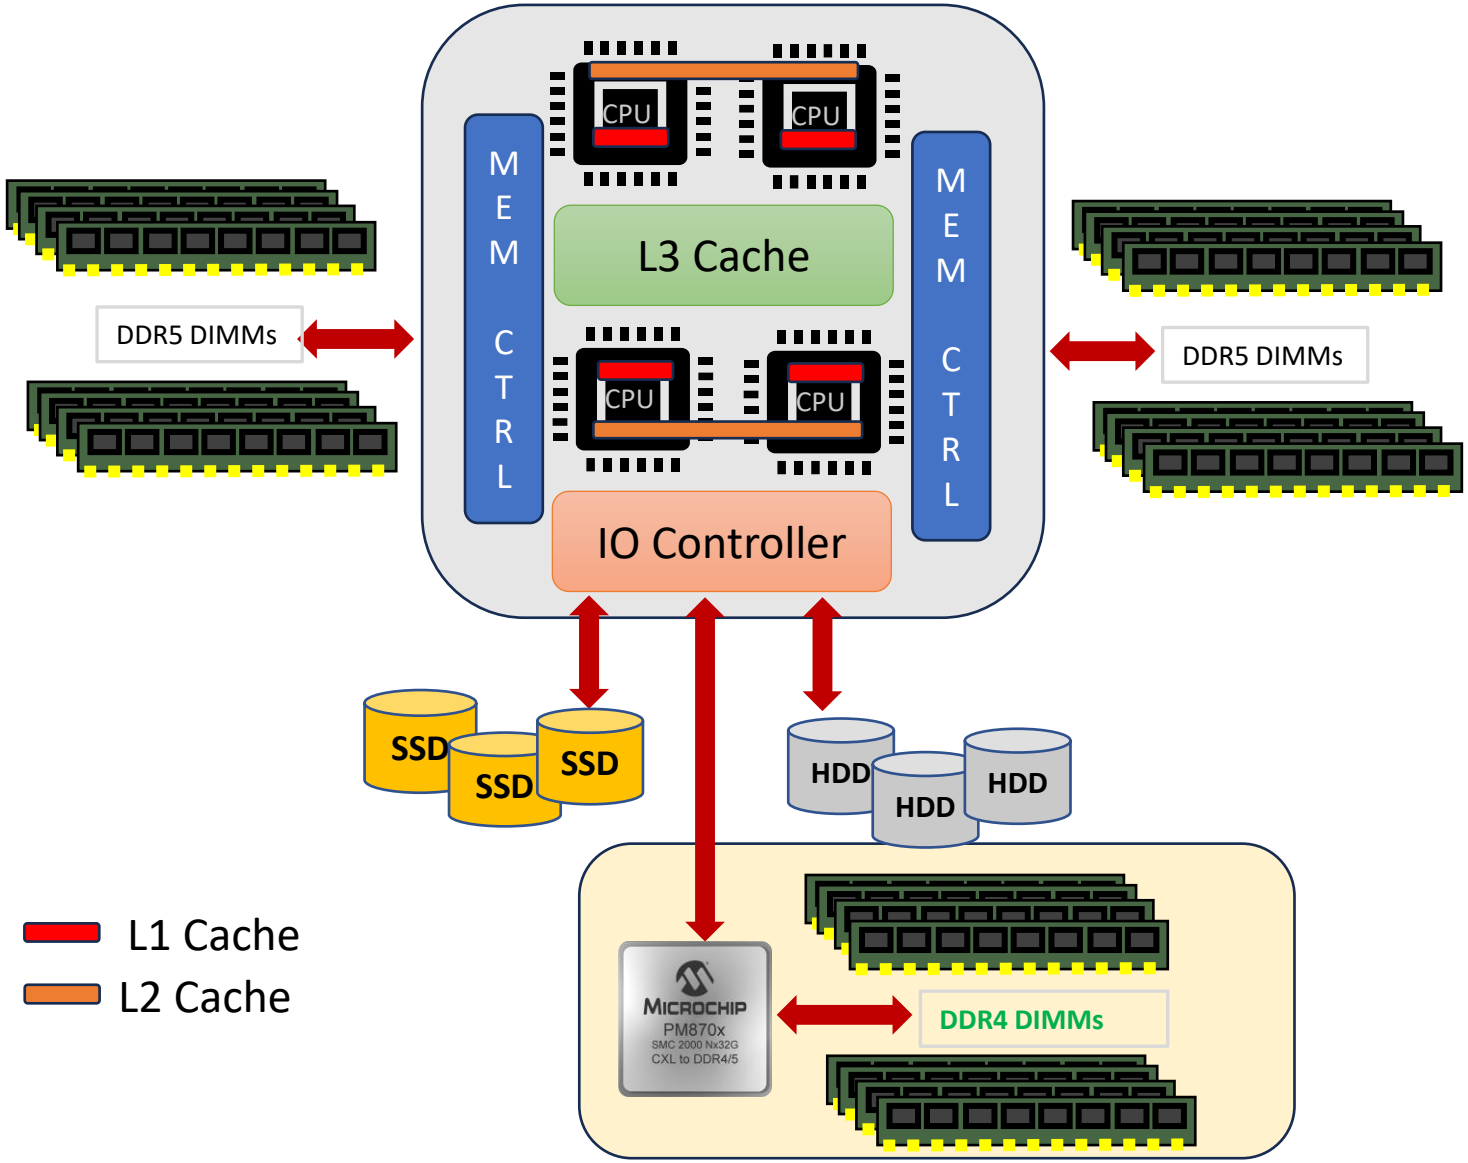

- Enables memory bandwidth and capacity expansion
- Enables re-use of decommissioned DDR DIMMs reducing TCO
- Not so latency sensitive Applications could utilize lower cost DDR4 or DDR5 memory

- Challenges
  - Memory Management
  - Reliability
  - Memory property aware Applications

Memory Tiering with CXL Switch-Attached Memory Expansion

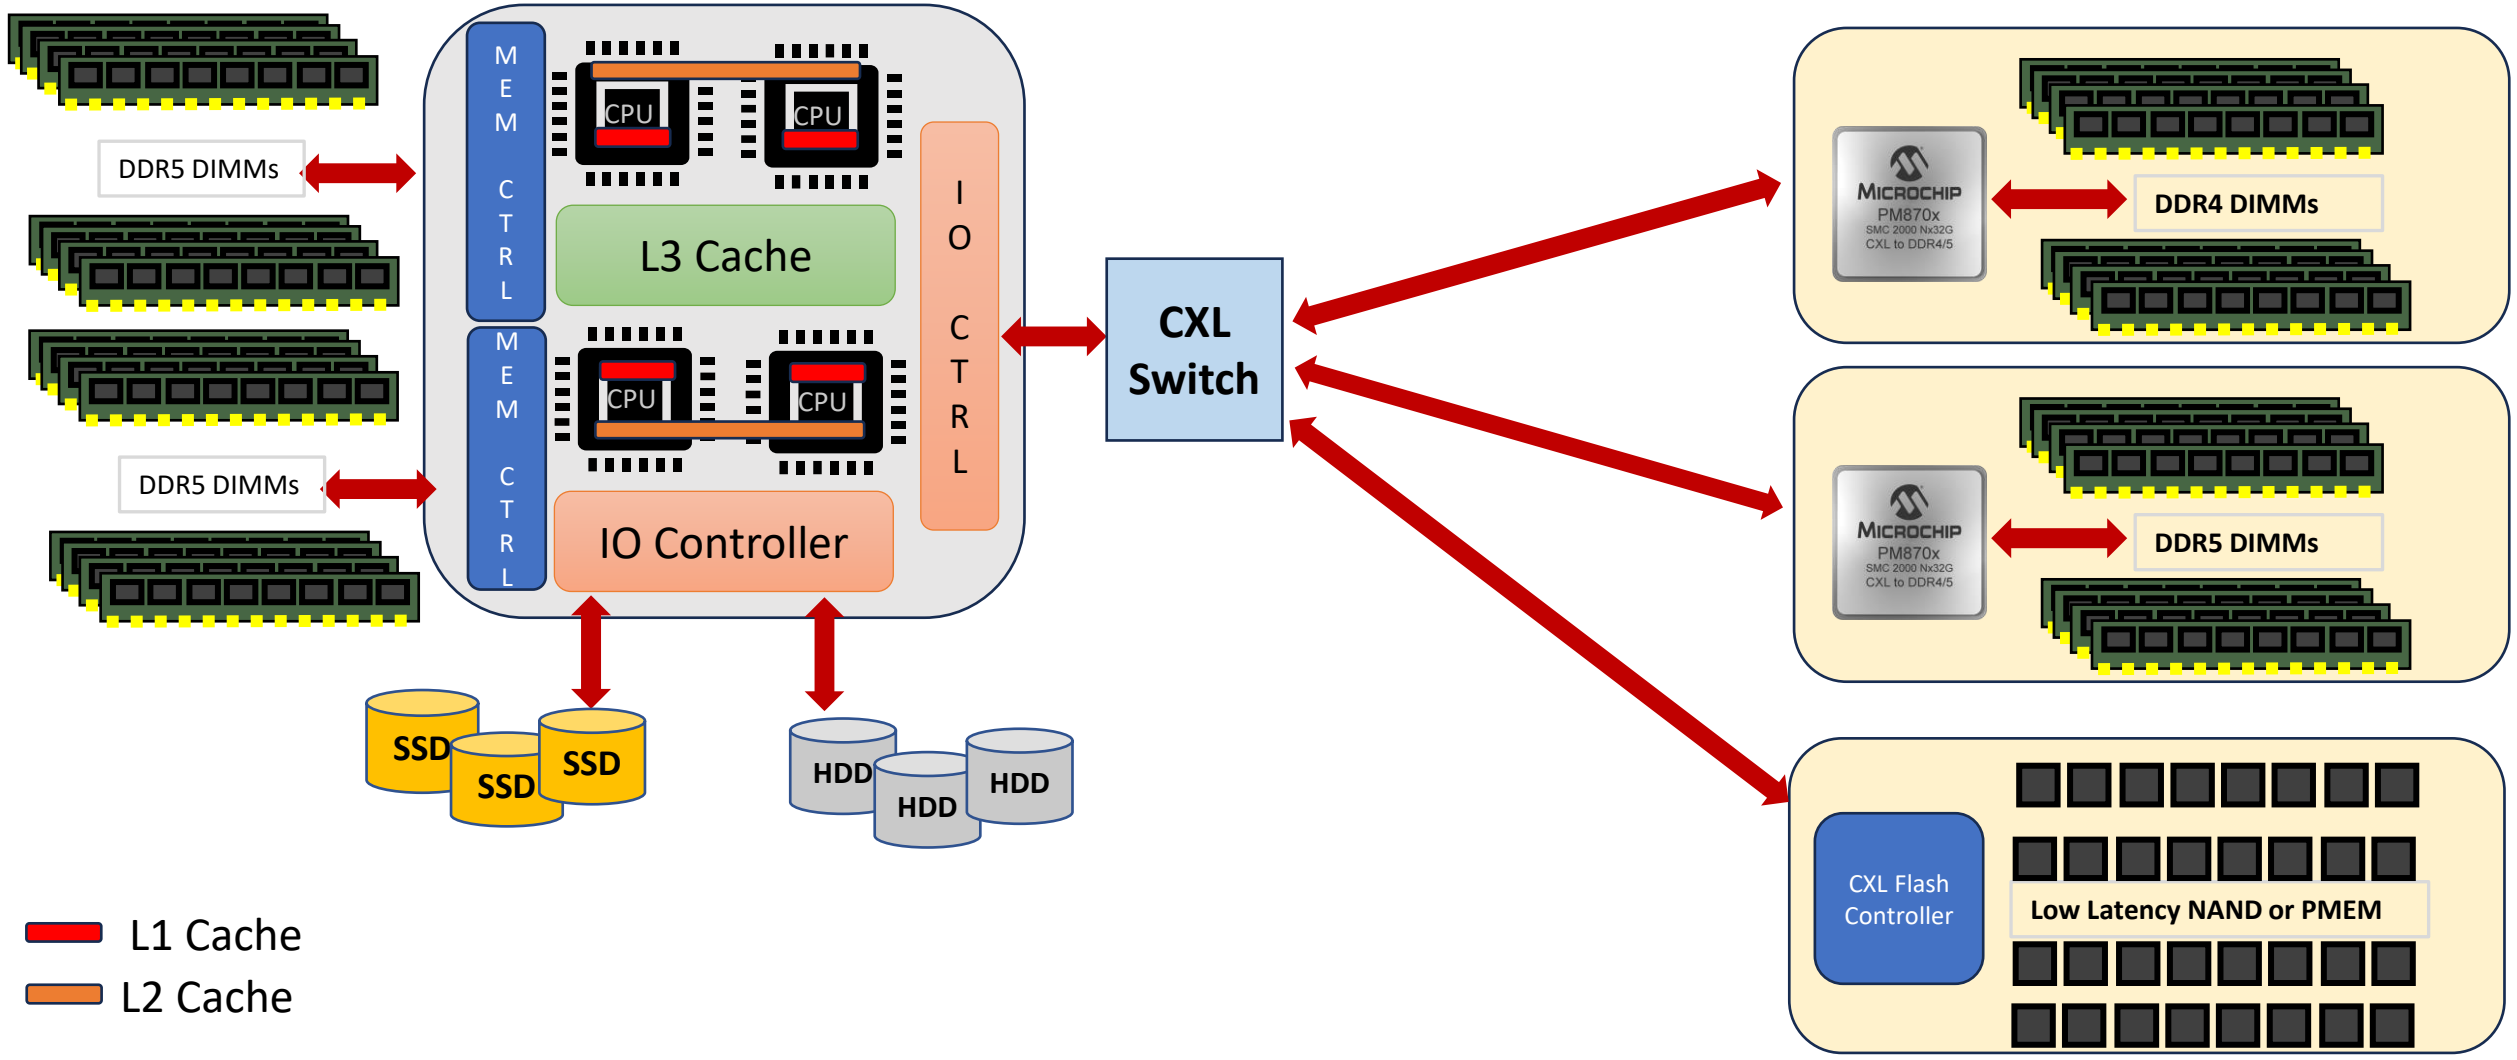

Memory Tiering with CXL Fabric-Attached Memory Expansion

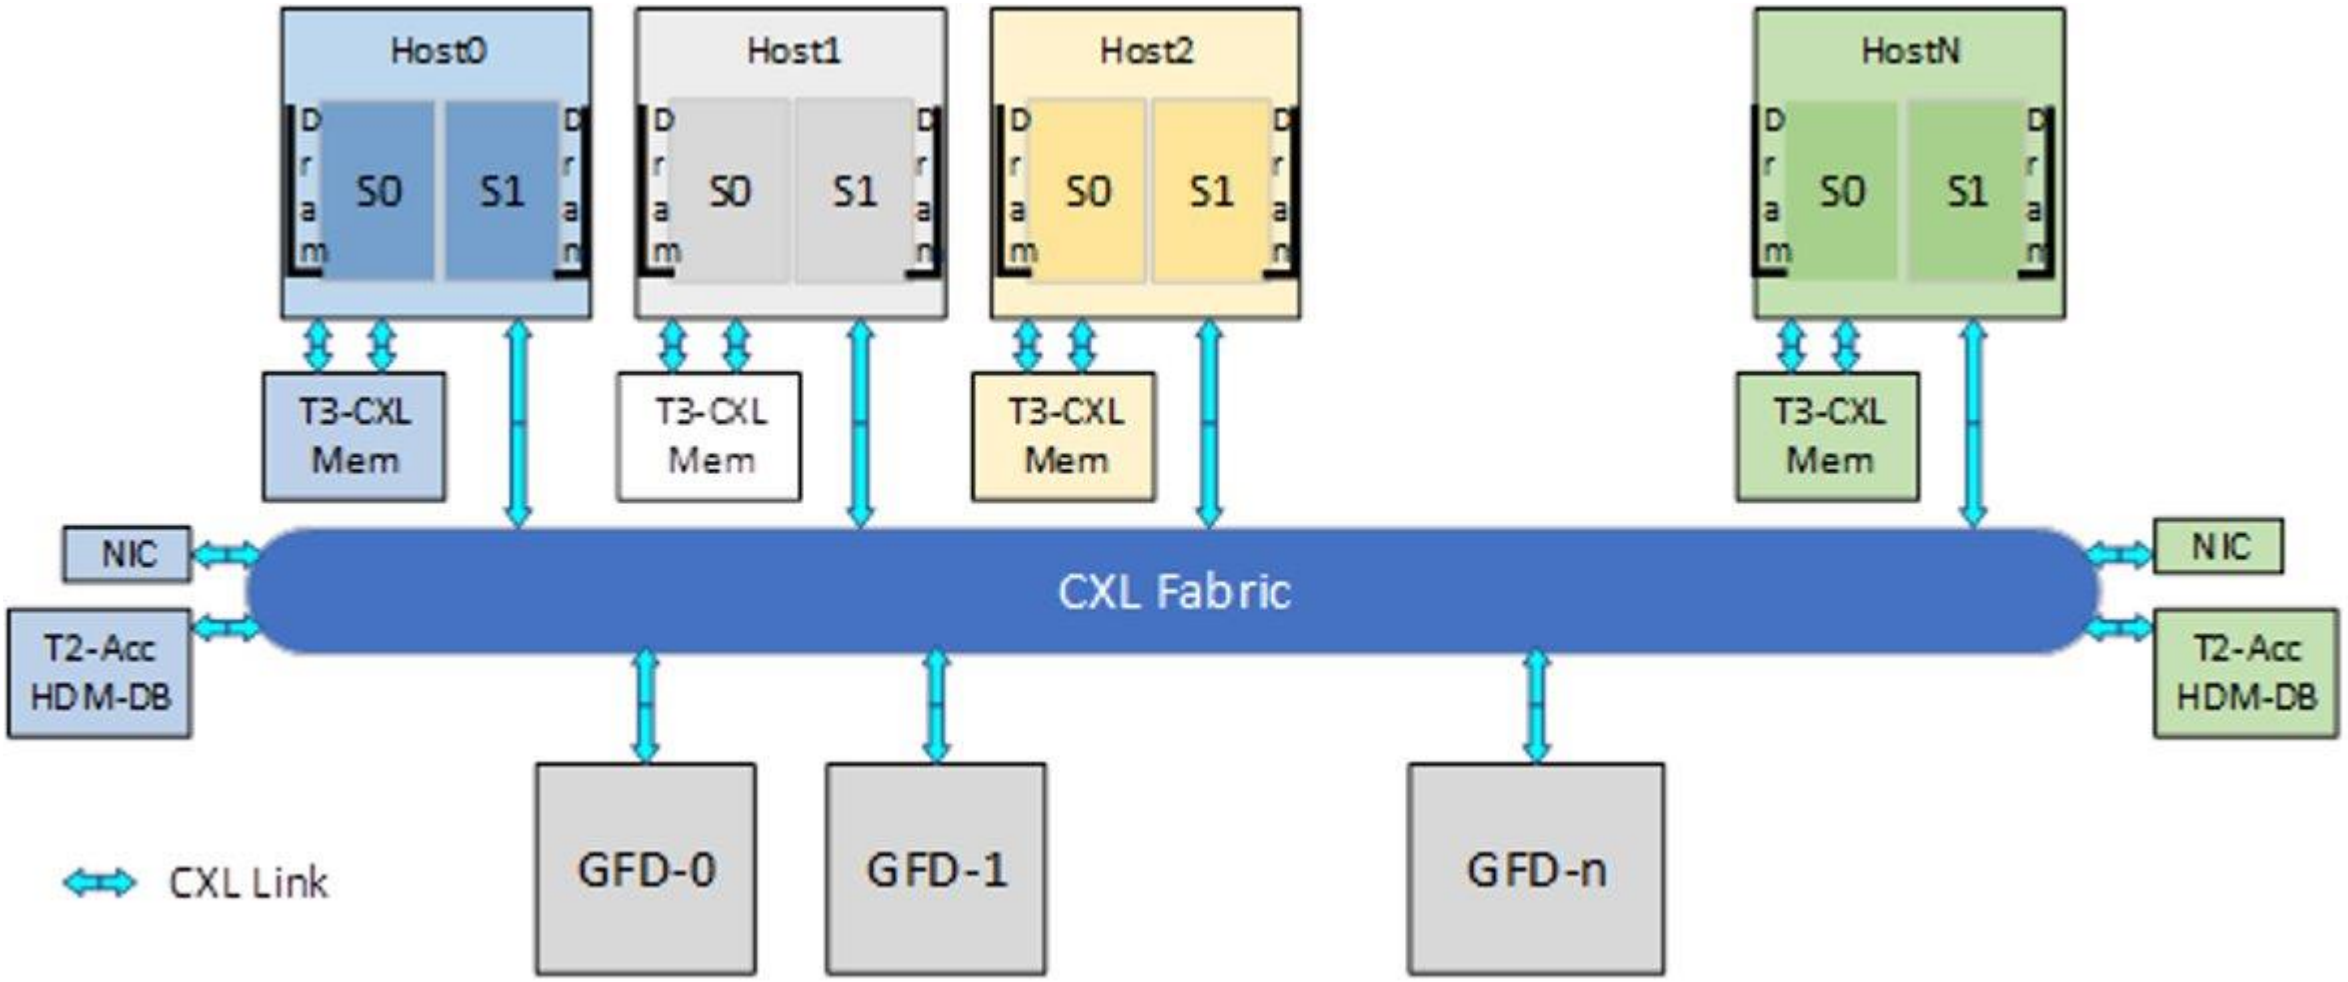

Memory Diversity Enablement

- Tiering can be used to efficiently utilize different Memory types with various properties in a system
- Basic tenets for memory tiering
  - Understand memory properties
  - Monitor data movement
  - Migrate and manage data appropriately
- Multiple tiering types available
  - HW based management
  - SW based management
  - Hybrid Management

Host Based CXL™ Memory Tiering Enablement

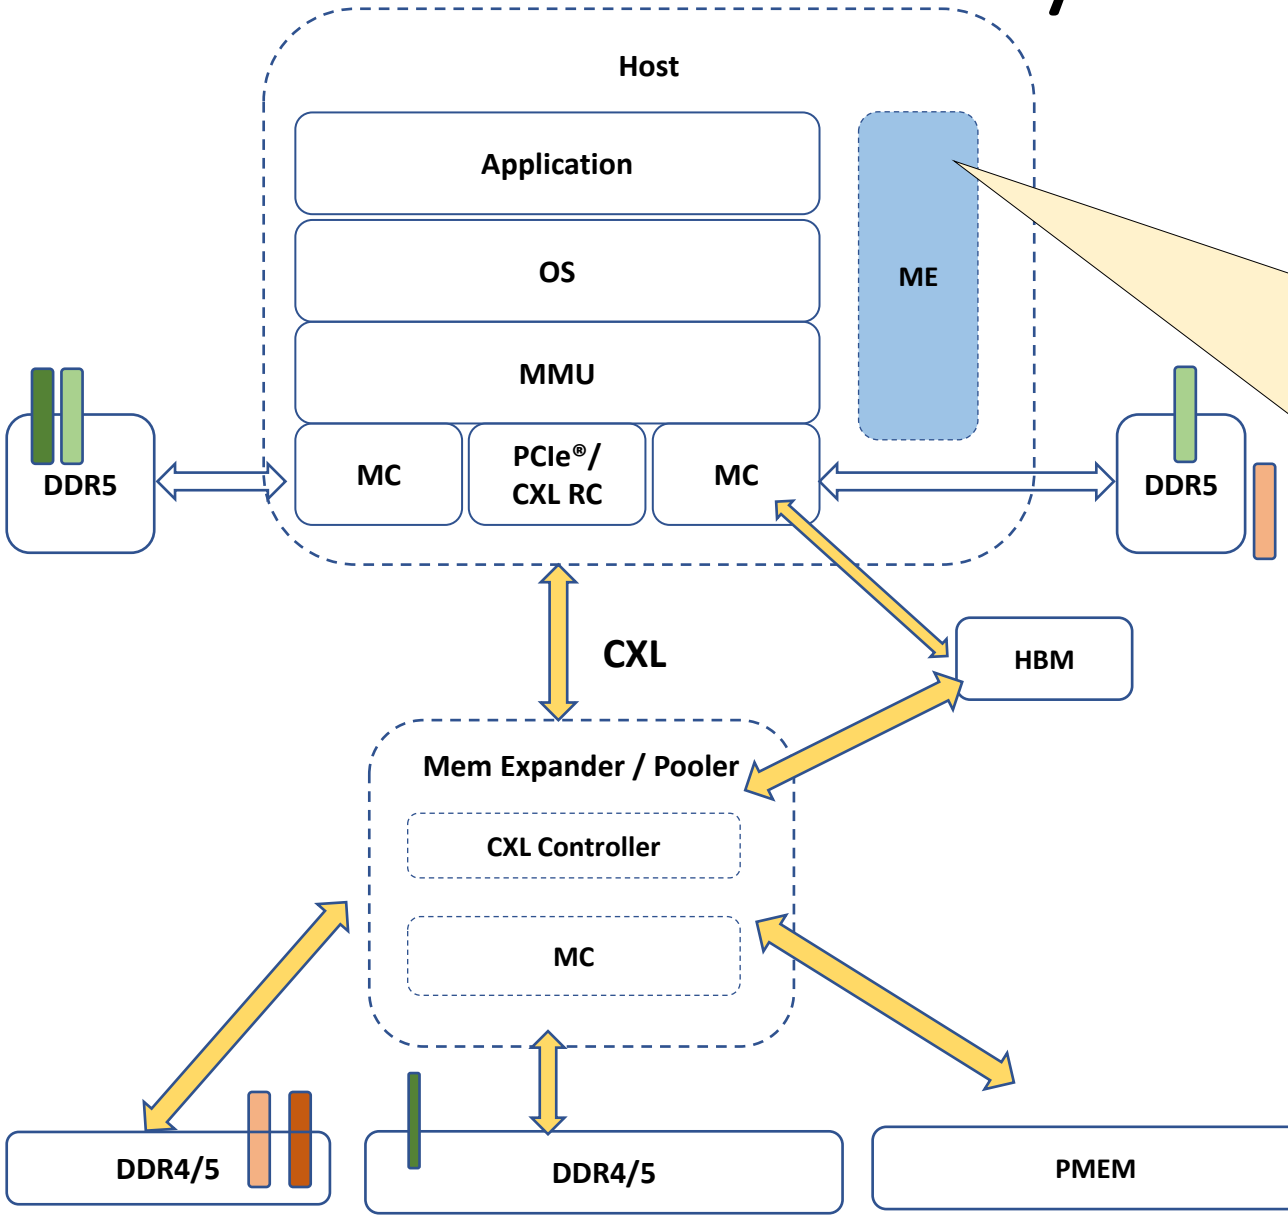

Migration Engine Functionalities

1. Builds statistics/histogram of page access in fabric memory, and publishes it to OS
2. Migrates pages to/from socket memory and informs OS upon completion
3. Migrates pages between different memory class (DDR4, DDR5, PMEM or HBM) and informs OS upon completion

Memory Controller Based CXL™ Tiering Enablement

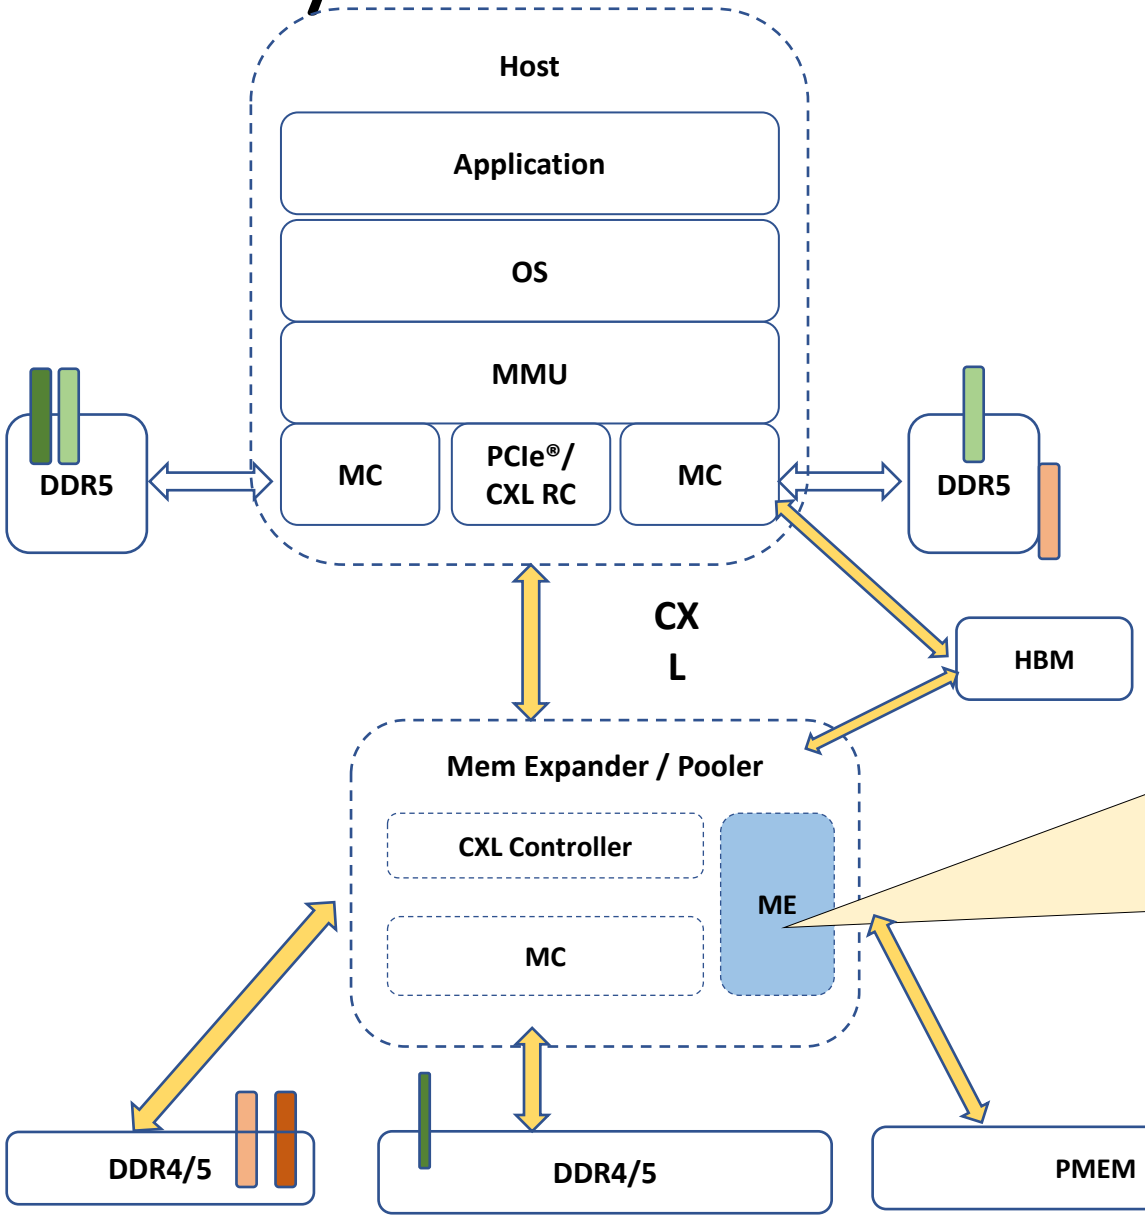

Migration Engine Functionalities

1. Builds statistics/histogram of page access in fabric memory, and publishes it to OS
2. Migrates pages to/from socket memory and informs OS upon completion
3. Migrates pages between different memory class (DDR4, DDR5, PMEM or HBM) and informs OS upon completion

Software Based CXL™ Memory Tiering Enablement

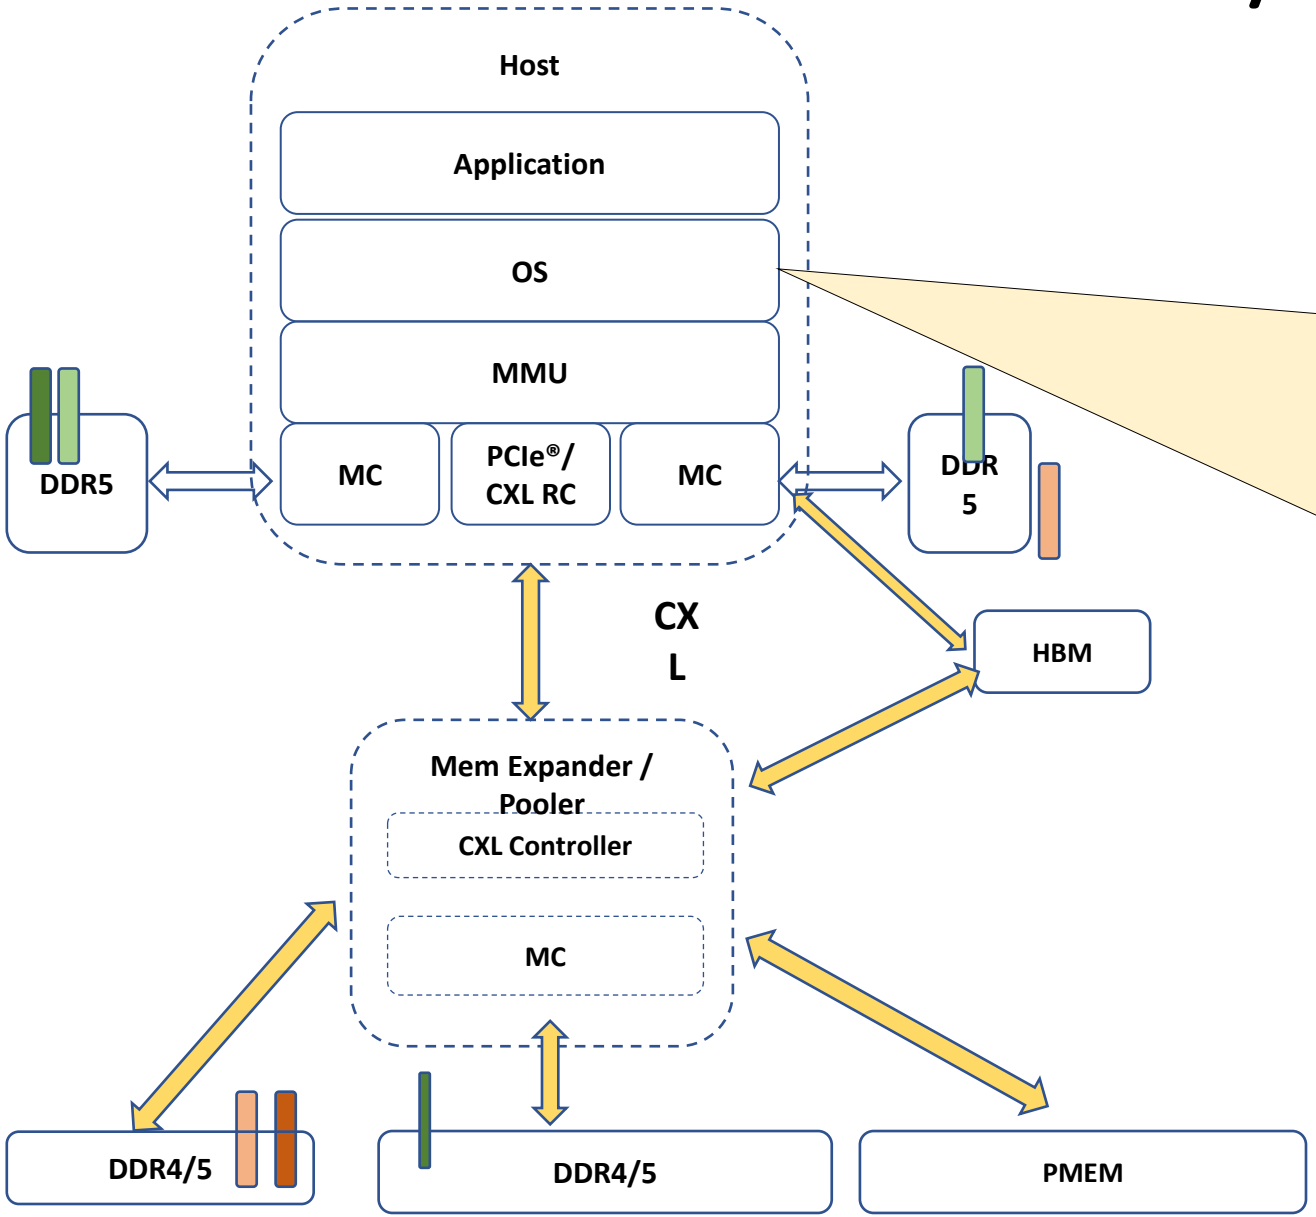

OS or Kernel level Migration algorithm implementation

1. Builds statistics/histogram of page access in fabric memory
2. Migrates pages to/from socket memory and informs OS upon completion
3. Migrates pages between different memory class (DDR4, DDR5, PMEM or HBM)

Additional processing overhead required to enable monitoring memory access and data move which may not be efficient

Hybrid CXL™ Memory Tiering Enablement

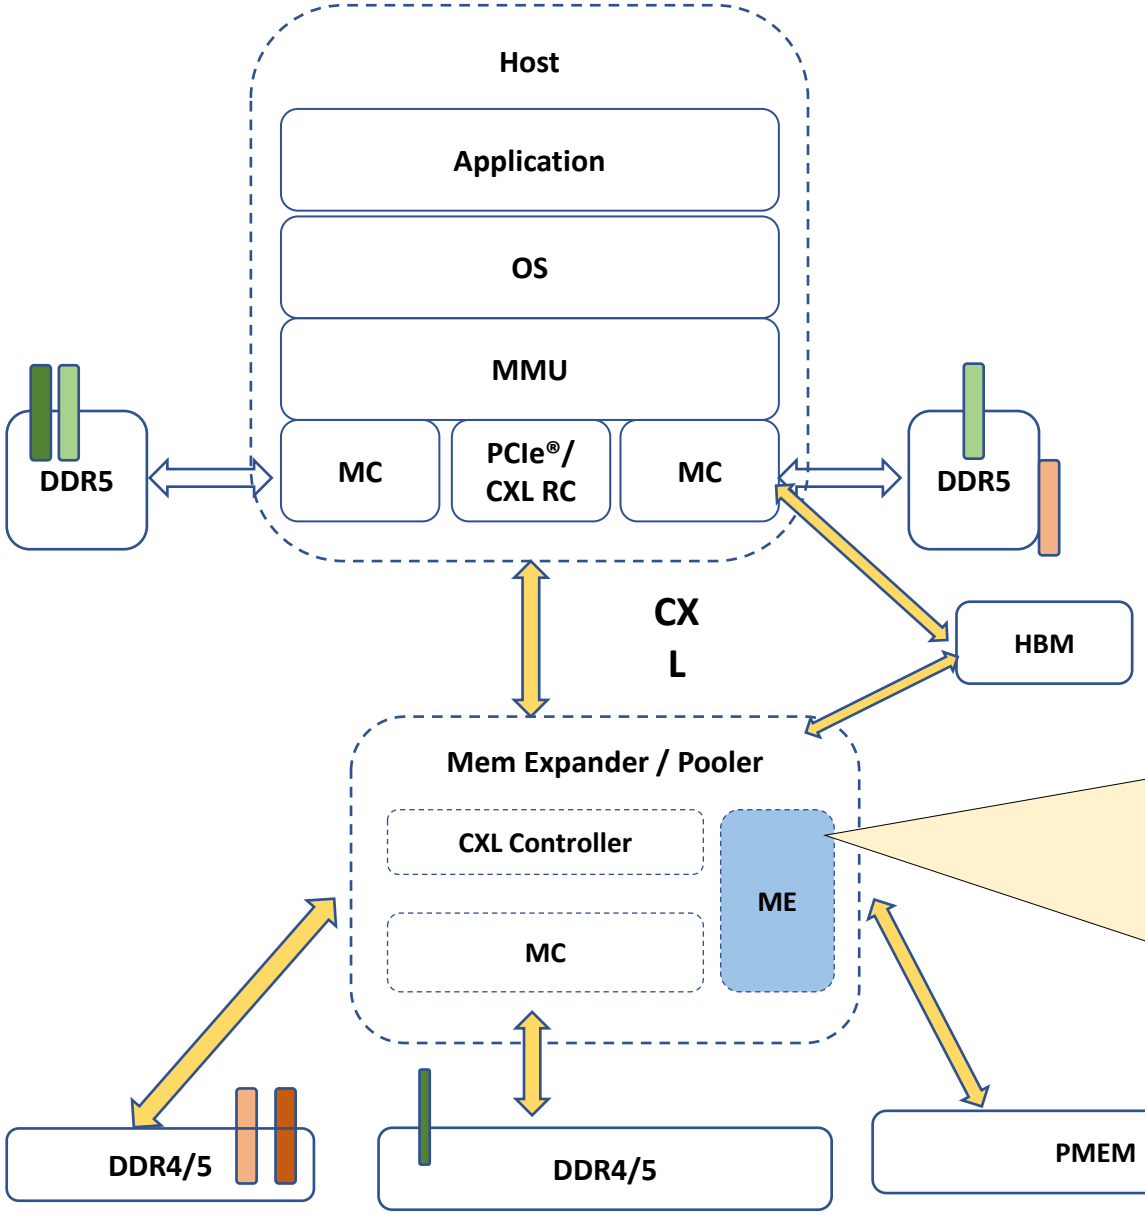

Migration Engine Functionalities

1. Migration Engine Builds statistics/histogram of page access in fabric memory, and publishes it to OS

2. OS Migrates pages to/from socket memory

3. OS Migrates pages between different memory class (DDR4, DDR5, PMEM or HBM) and informs OS upon completion

- Enables building quick statistics and histogram of data movement across memory and profiling
- Enables customization on per application basis

Conclusion

- CXL™ ecosystem enables memory diversity in order to utilize memory with different properties
  - A standardized approach to deal with different types of memories
  - Improves TCO by reducing the requirement of high cost fast memory
  - Increased data placement efficiency based on application or system requirements making data available just in time
  - Enables faster back ups and restores
- Memory tiering when done efficiently can reduce the burden on today's system architectures bridging the compute memory gap.