SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

Standardizing memory to memory data movement with SDXI v1.0

Shyam Iyer
Chair, SNIA SDXI TWG
Member, SNIA Technical Council
Distinguished Engineer Dell

Abstract

- For long, using software to perform memory copies has been the gold standard for applications performing memory-to-memory data movement or system memory operations. The newly released SNIA Smart Data Accelerator Interface (SDXI) Specification v1.0 attempts to change that.
- SDXI v1.0 is a standard for a memory-to-memory data mover and acceleration interface that is extensible, forward-compatible and independent of I/O interconnect technology. Among other features, SDXI standardizes an interface and architecture that can be abstracted or virtualized with a well-defined capability to quiesce, suspend, and resume the architectural state of a per-address-space data mover.
- This specification was developed by SNIA's SDXI Technical Working Group, comprising of 89 individuals representing 23 SNIA member companies.
- As new memory technologies get adopted and memory fabrics expand the use of tiered memory, SDXI Specification v1.0 will be enhanced with new accelerated data movement operations and use cases. This talk gives an overview of the SDXI's use cases, important features of SDXI specification v1.0, and what features can be expected with new versions of the specification.

SDXI

- Introduction to SNIA SDXI v1.0
- Use Cases
  - Application Patterns and benefits of Data Movement & Acceleration
- SDXI: The path ahead
  - SDXI v1.1
  - SDXI Ecosystem

Current Data Movement Standard

- Software `memcpy` is the current data movement standard
  - Stable ISA
- However,
  - Takes away from application performance
  - Incurs software overhead to provide context isolation.
  - Offload DMA engines and their interfaces are vendor-specific
  - Not standardized for user-level software.

SDXI(Smart Data Accelerator Interface)

- Smart Data Accelerator Interface (SDXI) is a SNIA standard for a memory to memory data movement and acceleration interface that is -
  - Extensible
  - Forward-compatible
  - Independent of I/O interconnect technology

- SNIA SDXI TWG was formed in June 2020 and tasked to work on this proposed standard
  - 23 member companies, 89 individual members
- **v1.0 released!**
  - https://www.snia.org/sdxi

SDXI Memory-to-Memory Data Movement

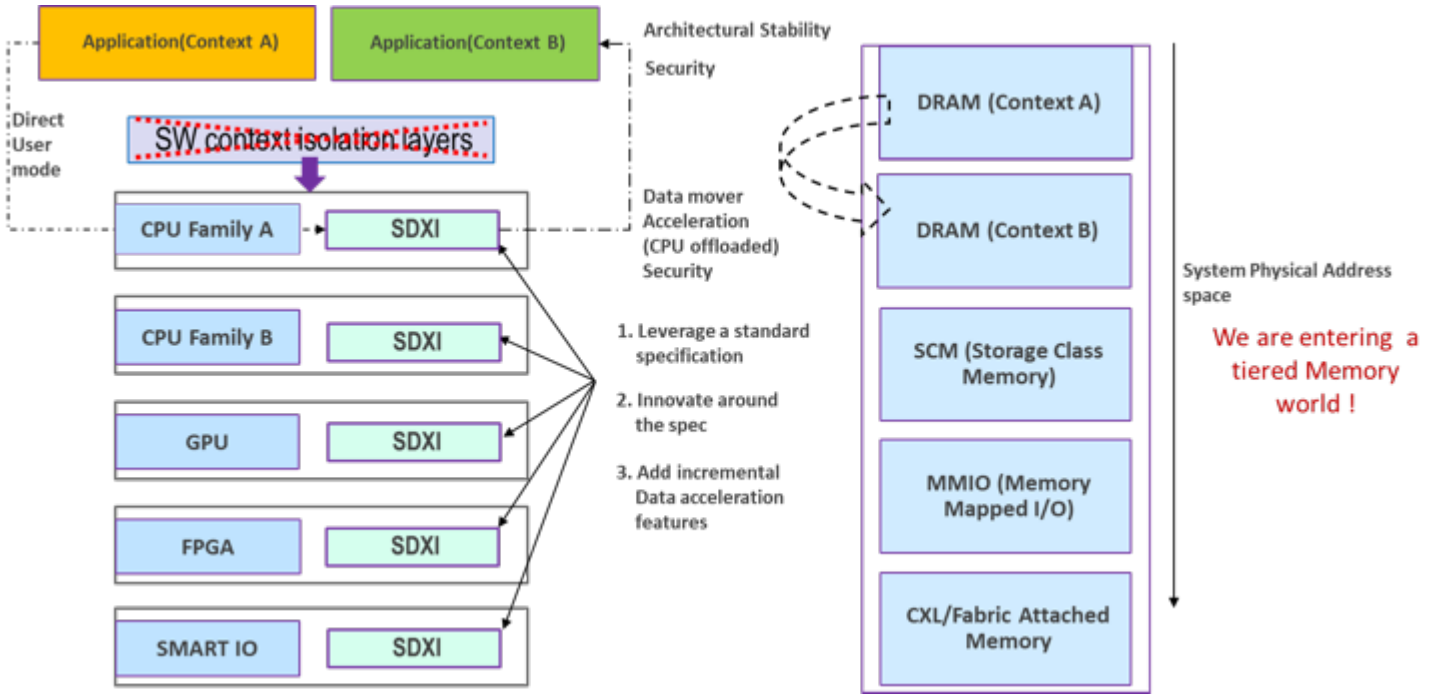

- Data movement between different address spaces.
- Data movement without mediation by privileged software.
- Allows abstraction or virtualization by privileged software.
- Capability to quiesce, suspend, and resume the architectural state of a per-address-space data mover.
- Forward and backward compatibility across future specification revisions.
- Additional offloads leveraging the architectural interface.
- Concurrent DMA model.

Memory Structures(1) – Simplified view

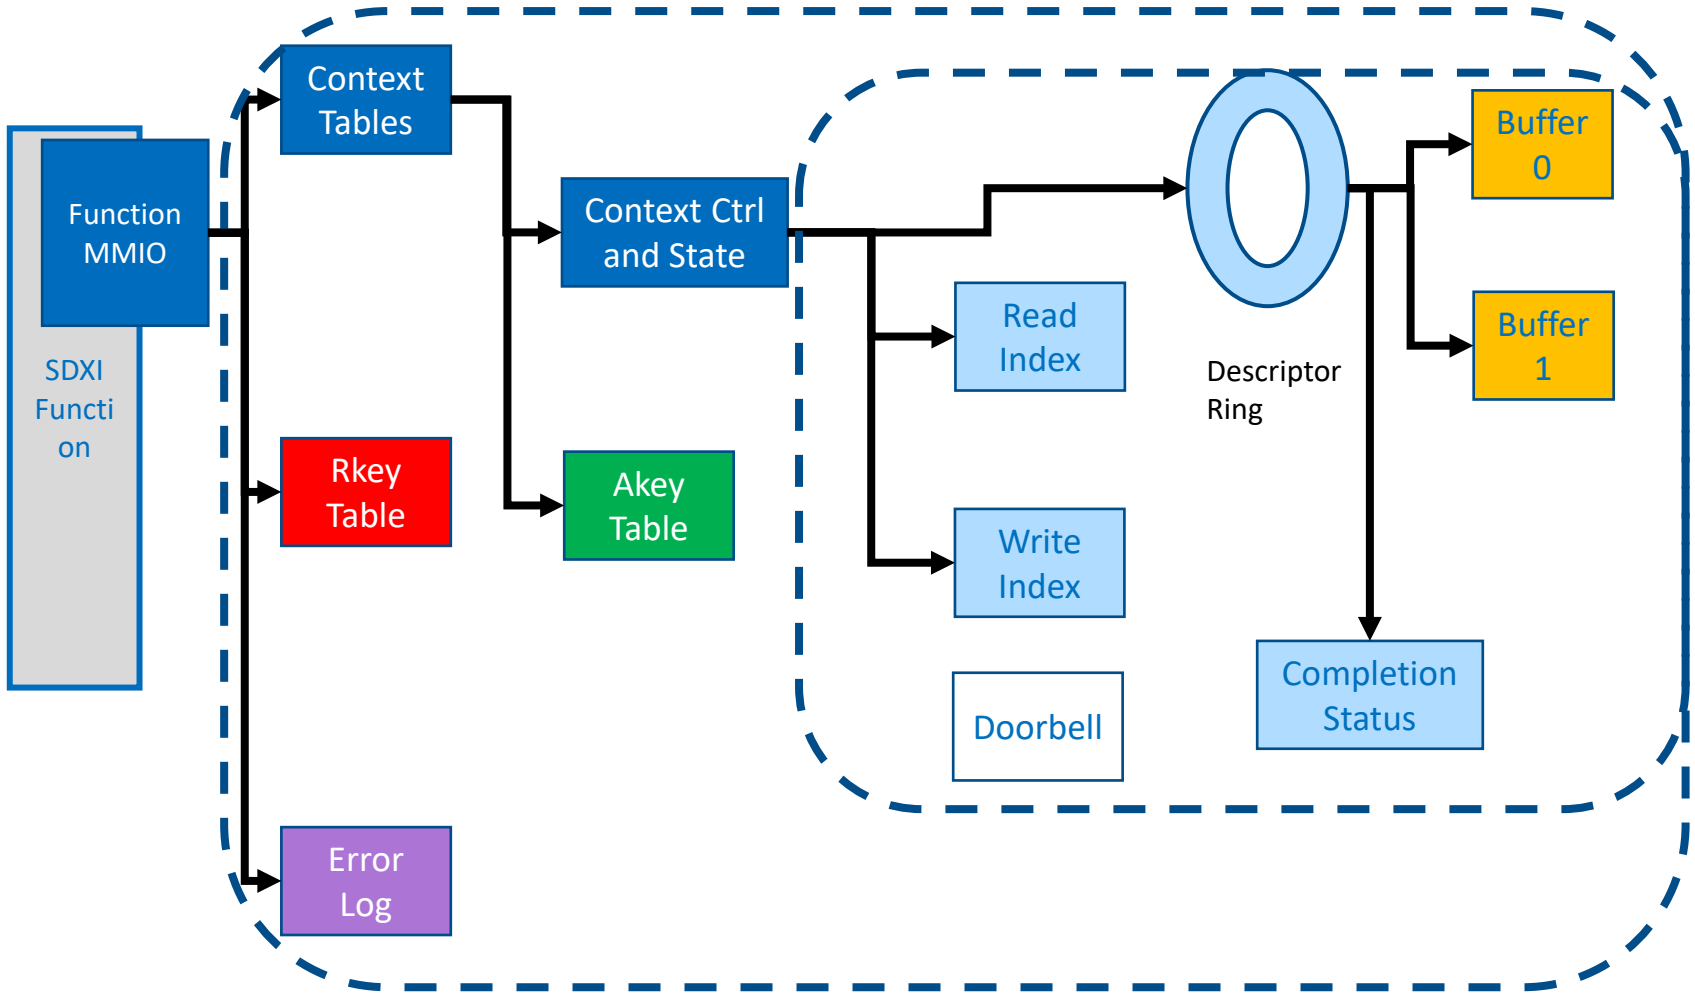

- All states in memory
- One standard descriptor format
  - Scope for future expansion
- Easy to virtualize
- Architected function setup and control
  - *layered model for interconnect specific function management
  - SDXI class code registered for PCIe implementations

Multi-Address Space Data Movement within an SDXI function group (2)

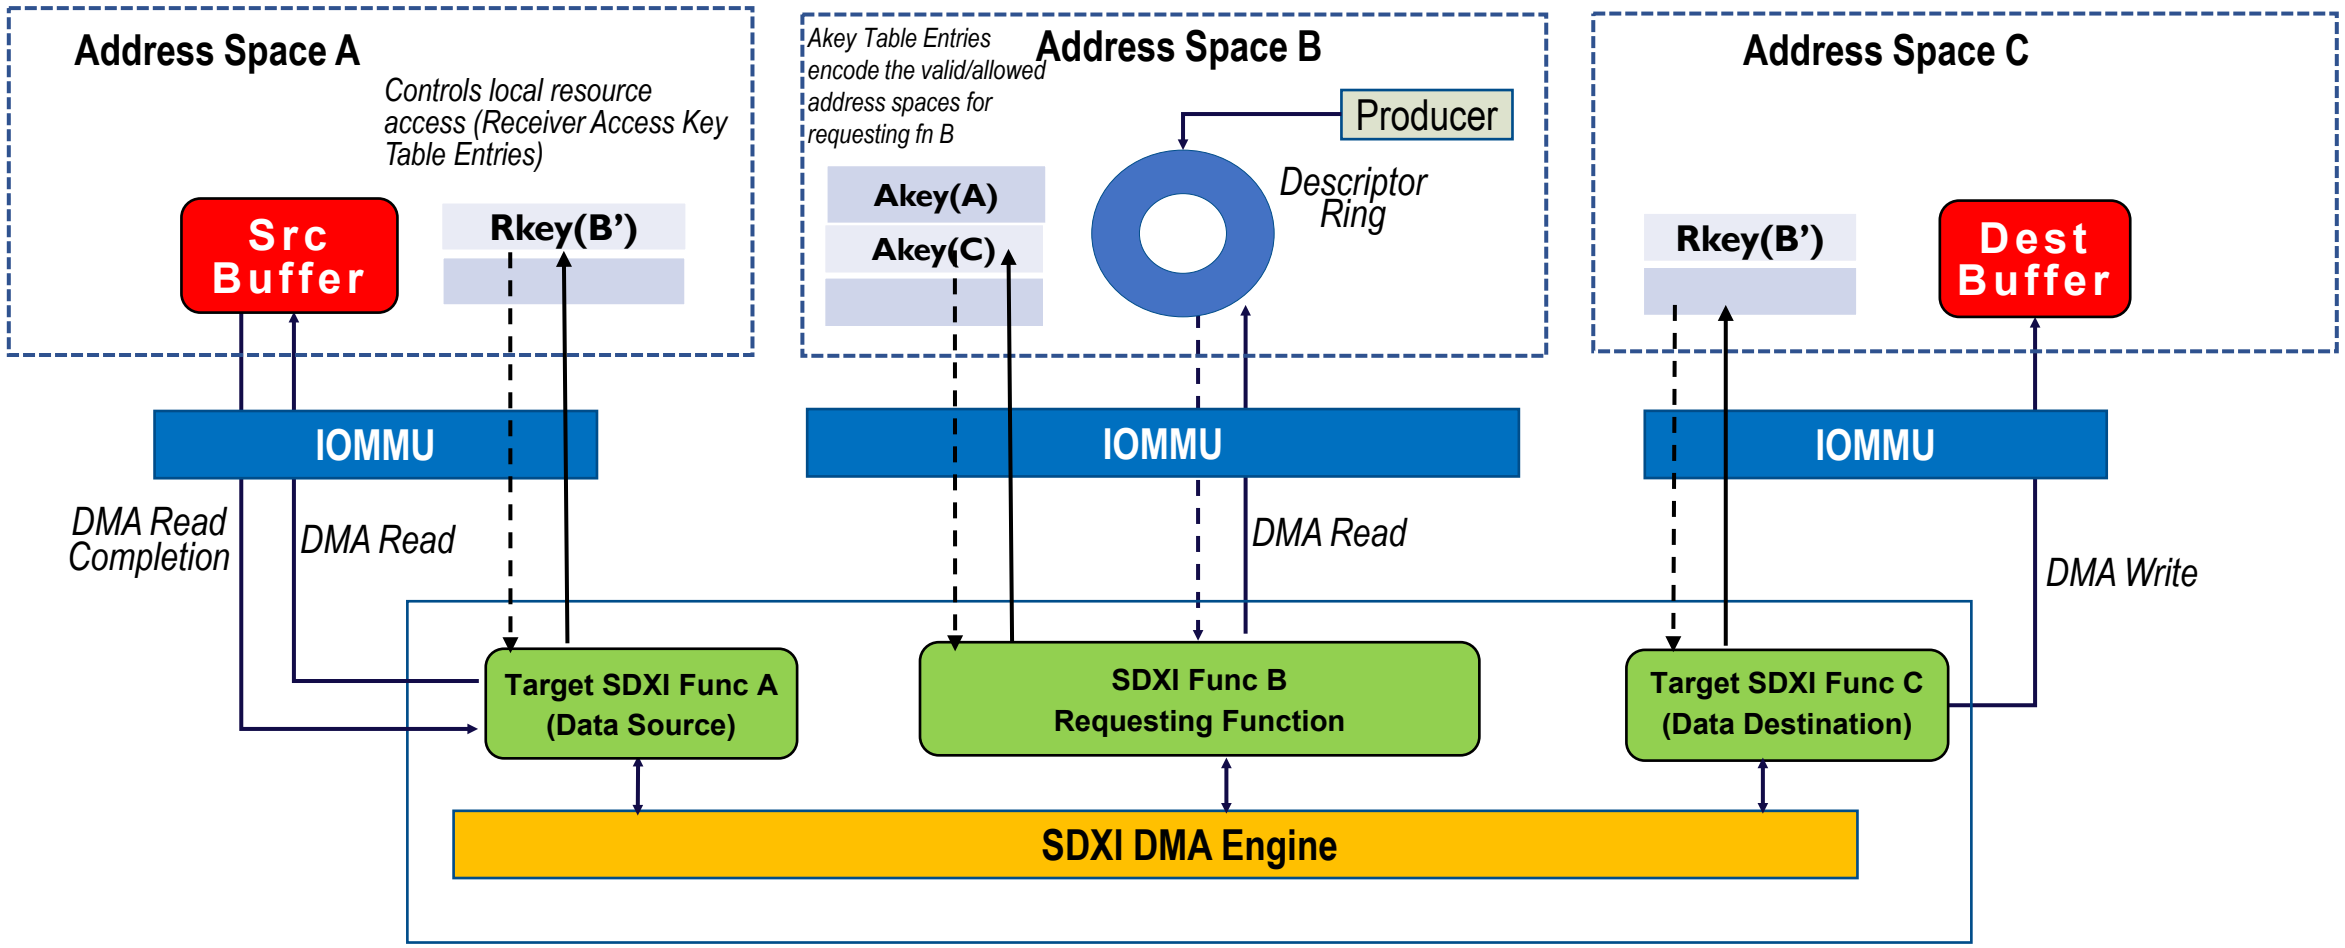

Active Contributors and growing…

SDXI

- Introduction to SNIA SDXI v1.0
- Use Cases
  - Application Patterns and benefits of Data Movement & Acceleration
- SDXI: The path ahead
  - SDXI v1.1
  - Software Ecosystem

Application Pattern 1 (Buffer Copies)

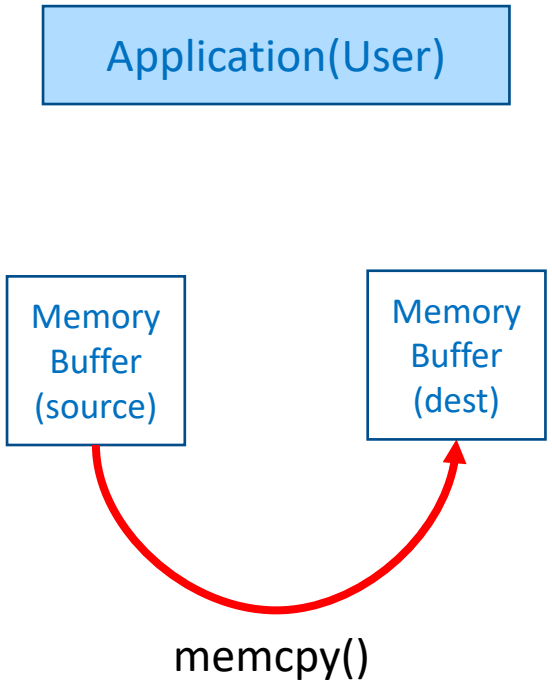

- Takes away from application performance

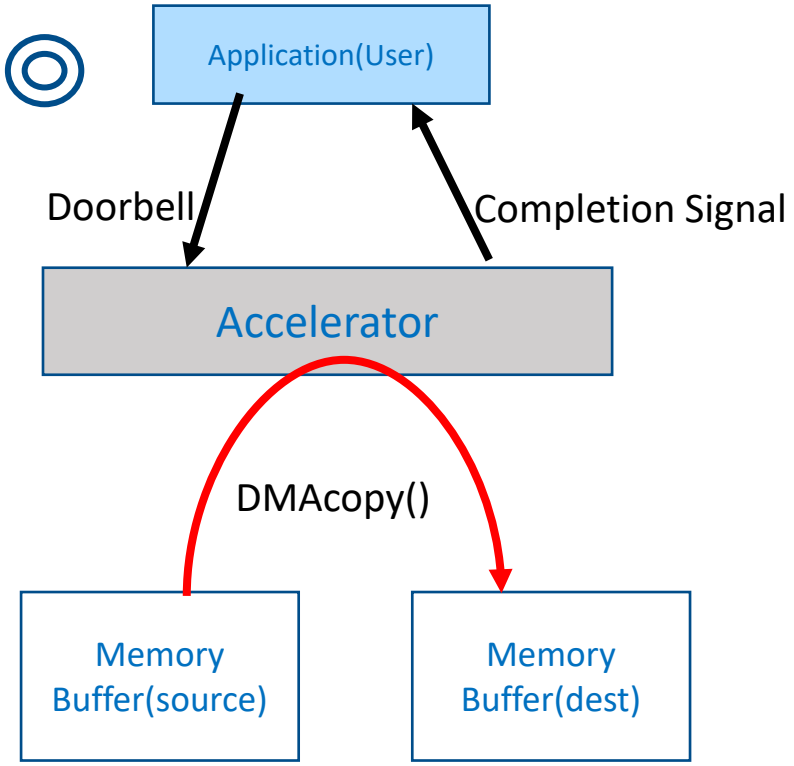

- HW based memory copies can be offloaded without affecting application performance

Application Pattern 2

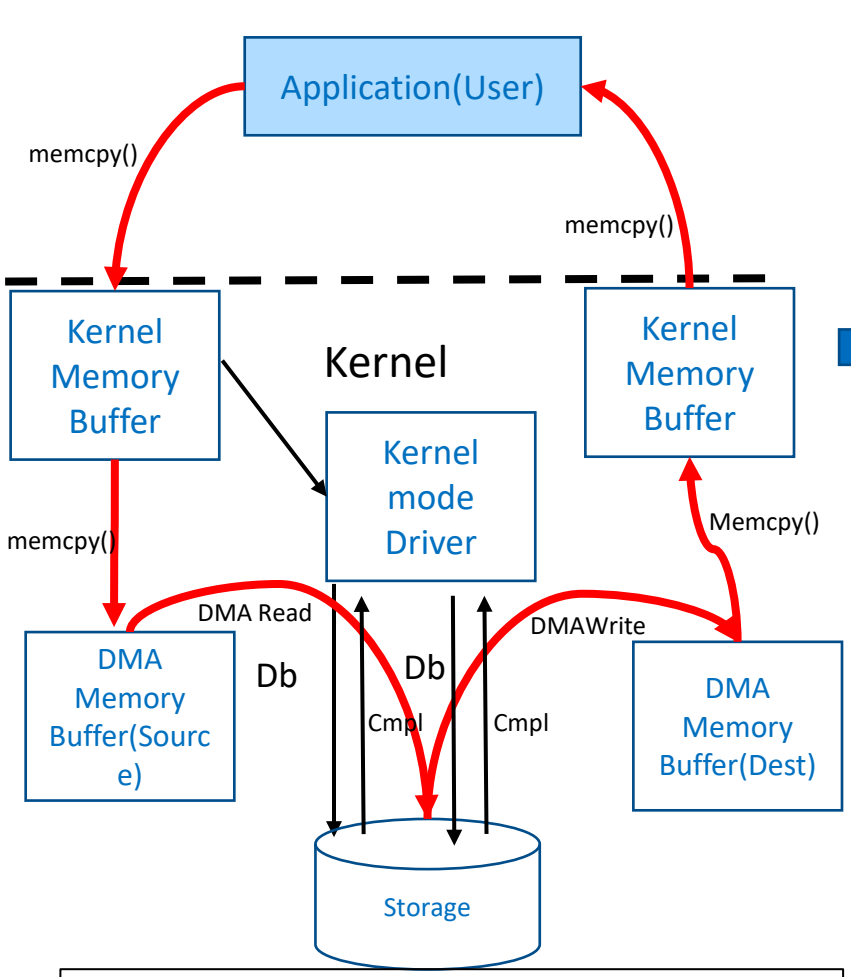

- Multiple data buffer copies before hardware-based data movement can occur

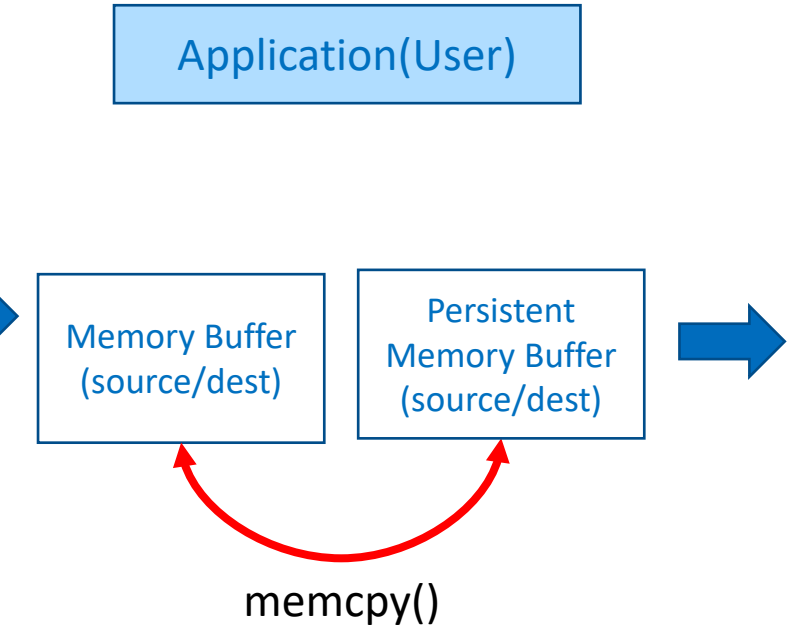

- Reduced buffer copies but still takes away from application performance

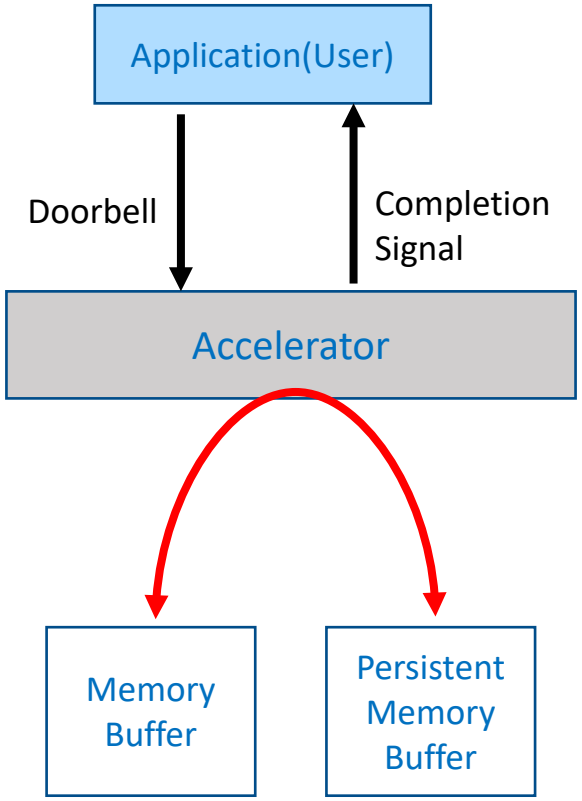

- Reduced buffer copies
- HW based offloaded memory copies

Application Pattern 3

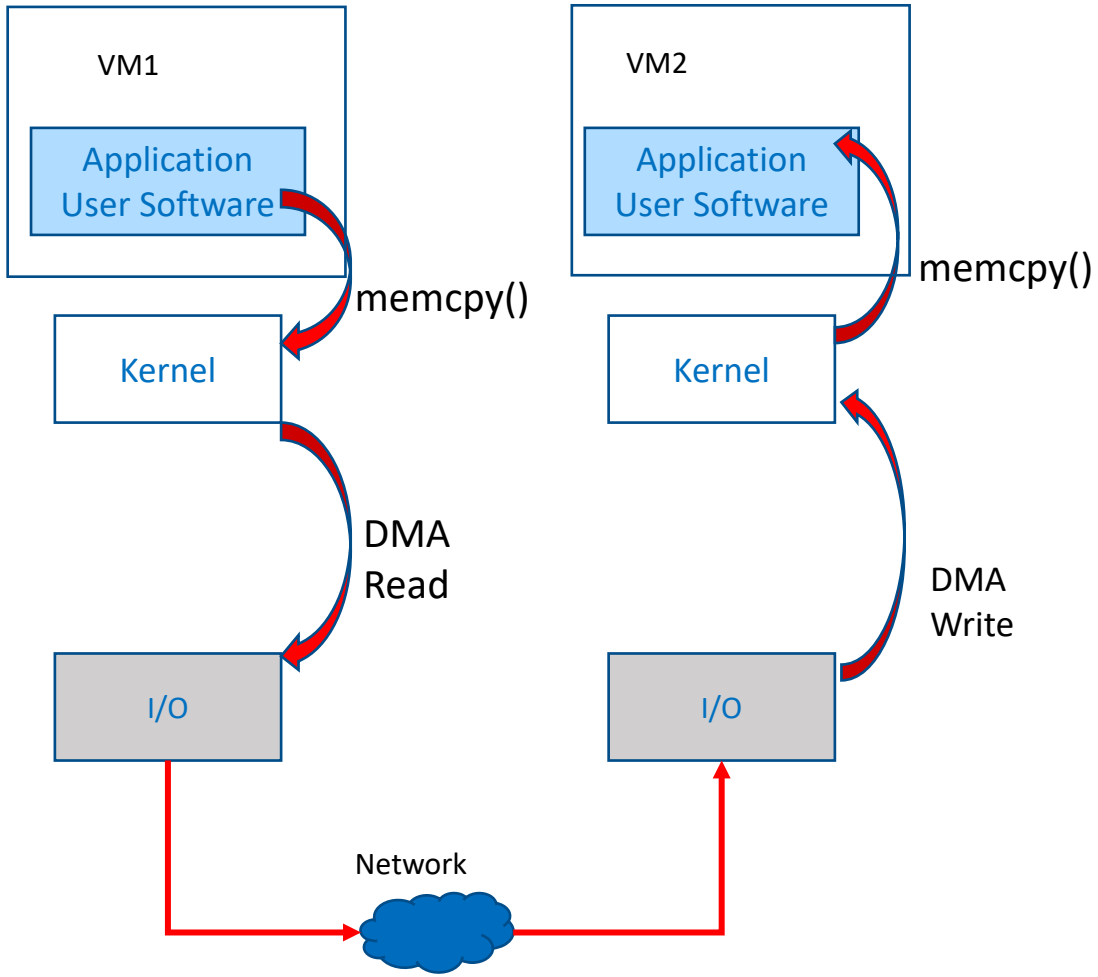

- Context isolation layers introduce multiple buffer copies

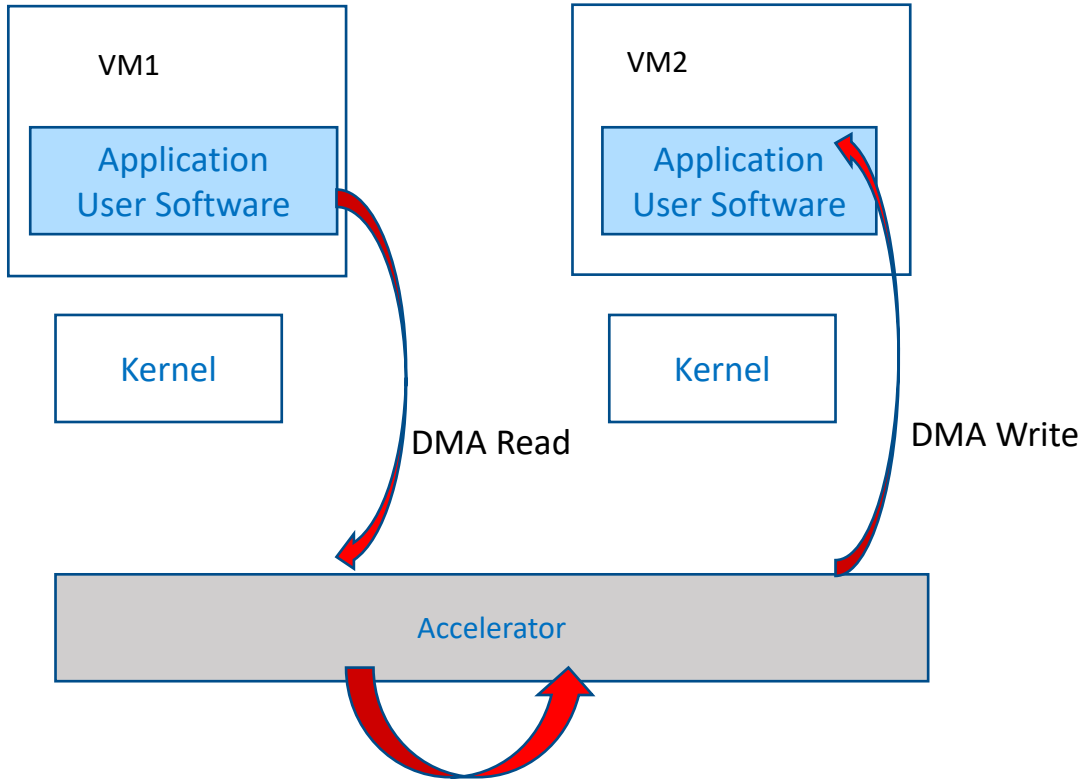

- Best of both: Context isolation layers and optimized HW based memory buffer copies

Data *in use* Memory Expansion

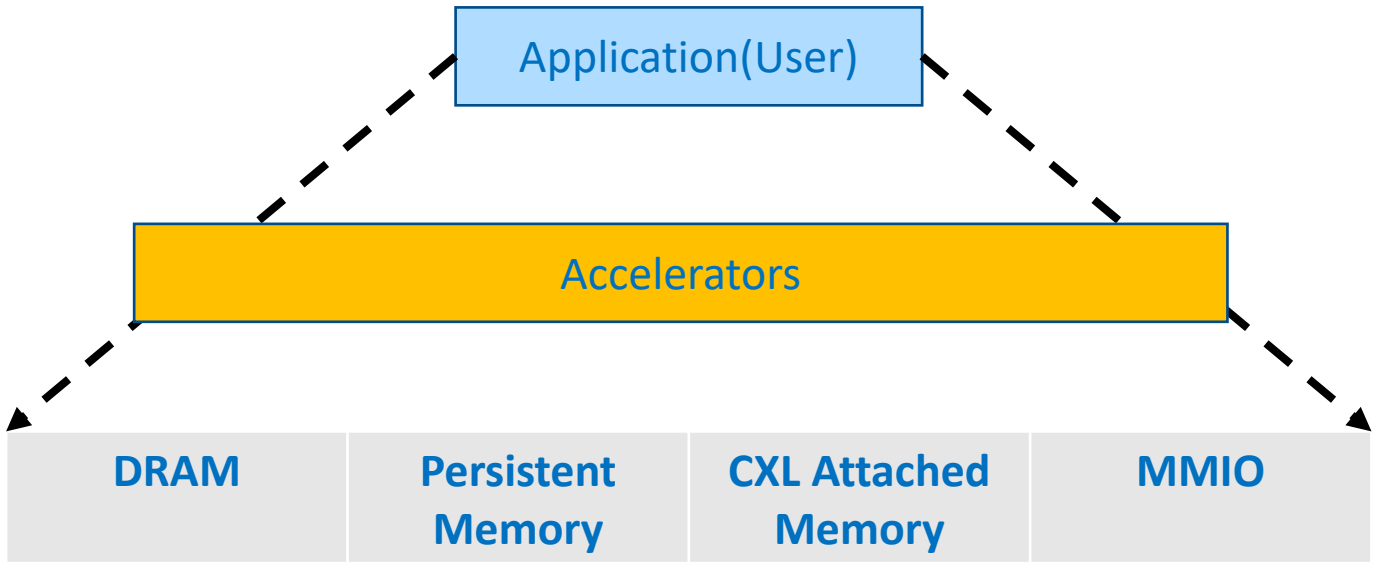

- Memory expansion expands the memory target surface area for accelerators
- Different tiers of memory
- Diversity in accelerator programming methods

Emerging use cases: SDXI Assisted Data Movement in a CXL Architecture

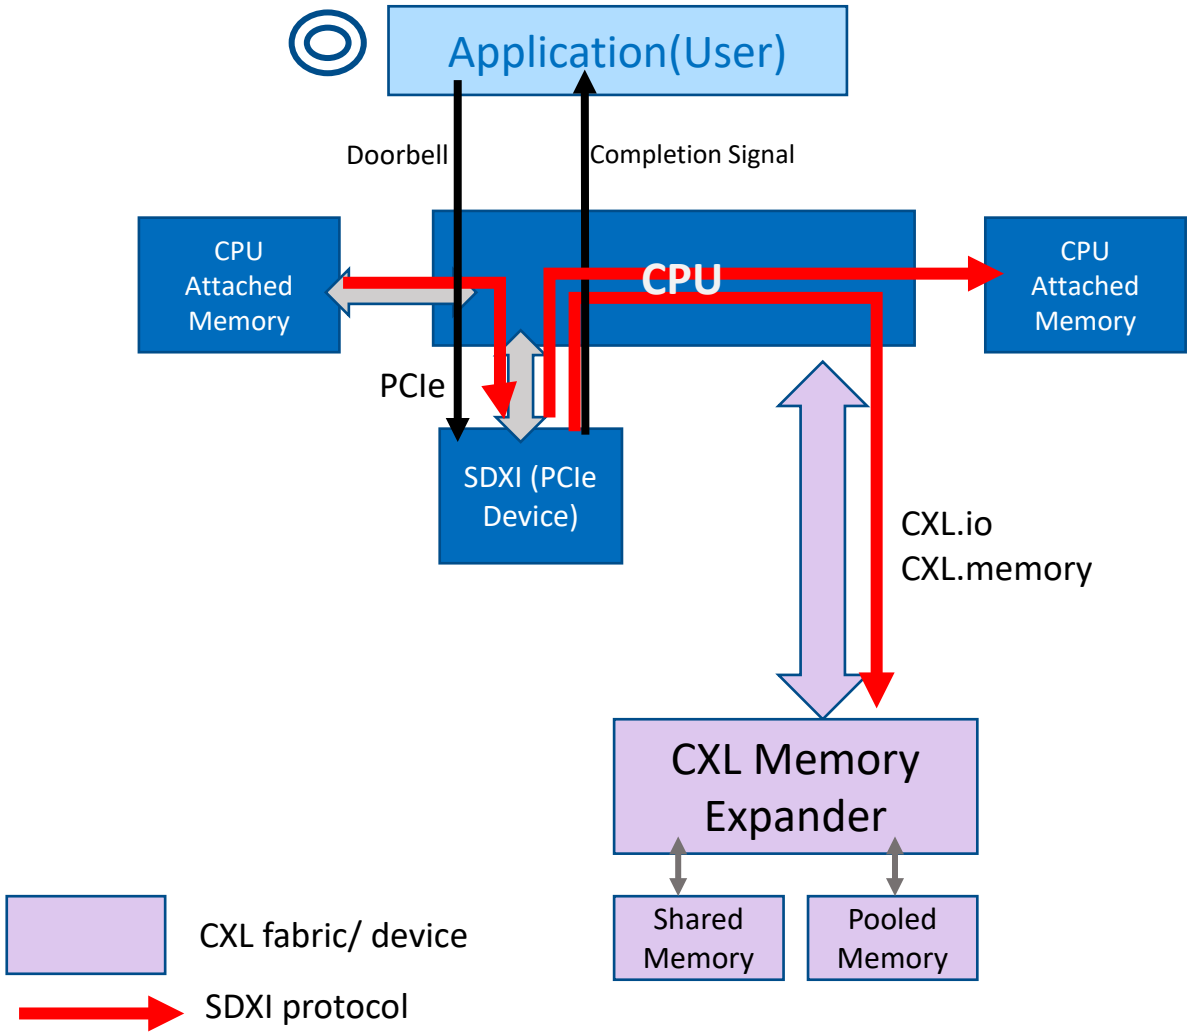

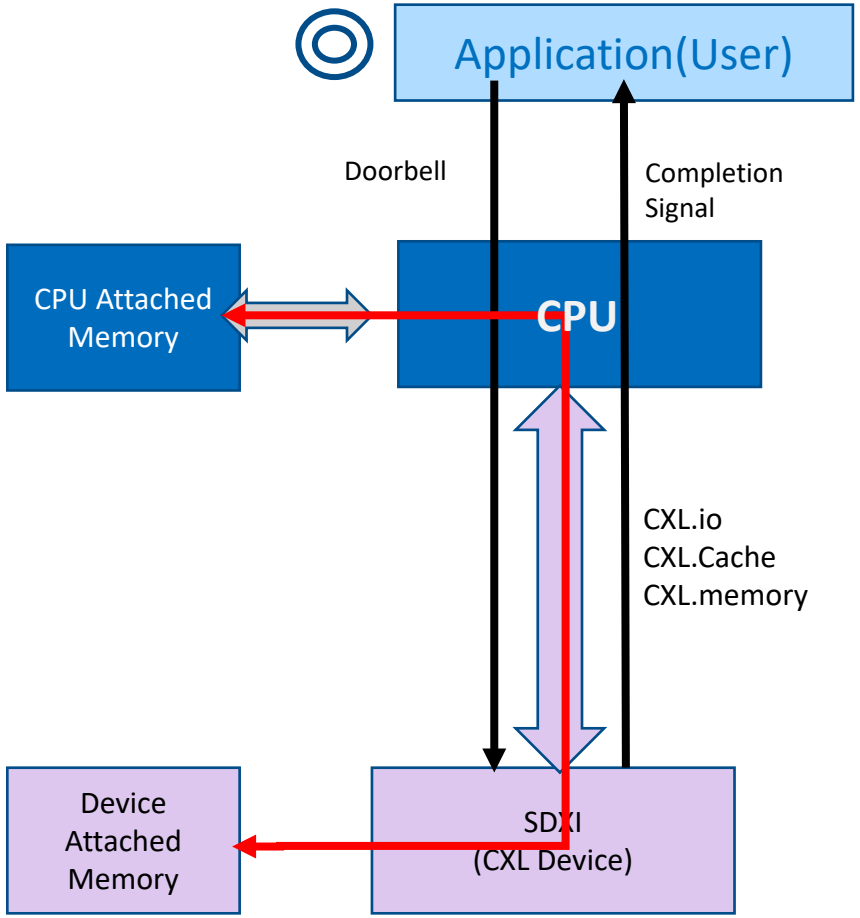

SDXI

- Introduction to SNIA SDXI v1.0
- Use Cases
  - Application Patterns and benefits of Data Movement & Acceleration
- SDXI: The path ahead
  - SDXI v1.1
  - Software Ecosystem

SDXI v1.1 investigations

- Management architecture for data movers(includes connection manager)
- New data mover operations for smart acceleration
  - DIF, DIX, Compression, Encryption, and other operations
- SDXI Host to Host investigations
- Scalability & Latency improvements
  - SIOV, Direct Work Submission, and other investigations
- Cache coherency models for data movers
- Security Features involving data movers
  - SDXI devices and Confidential Computing
  - IDE
  - Threat modeling
- Data mover operations involving persistent memory targets
- Other considerations
  - QoS
  - CXL-related discussions,
  - Heterogenous environments

Additional SDXI Ecosystem activities

- Libsdxi project work in SNIA
  - OS agnostic user space library development
- Linux Upstream driver efforts
  - SDXI TWG members are supporting this effort outside SNIA as a community
- SDXI emulation project investigation for ecosystem development
- Investigations to enable SDXI compliance for SW and HW interoperability
- SNIA's CS+SDXI Subgroup is kicking off activities to:
  - Envision SDXI in a Computational Storage Architecture
  - Implement features in SDXI to support Computational Storage use cases

Call to Action

- Join the TWG to influence the next version of the specification
- v1.0 is available for implementation
  - Feedback via SNIA feedback portal
  - https://www.snia.org/feedback
- Participate in the SDXI Ecosystem

SNIA COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

Please take a moment to rate this session.

Your feedback is important to us.

Q&A

Backup

Legacy Compute, IO, Memory Bubbles

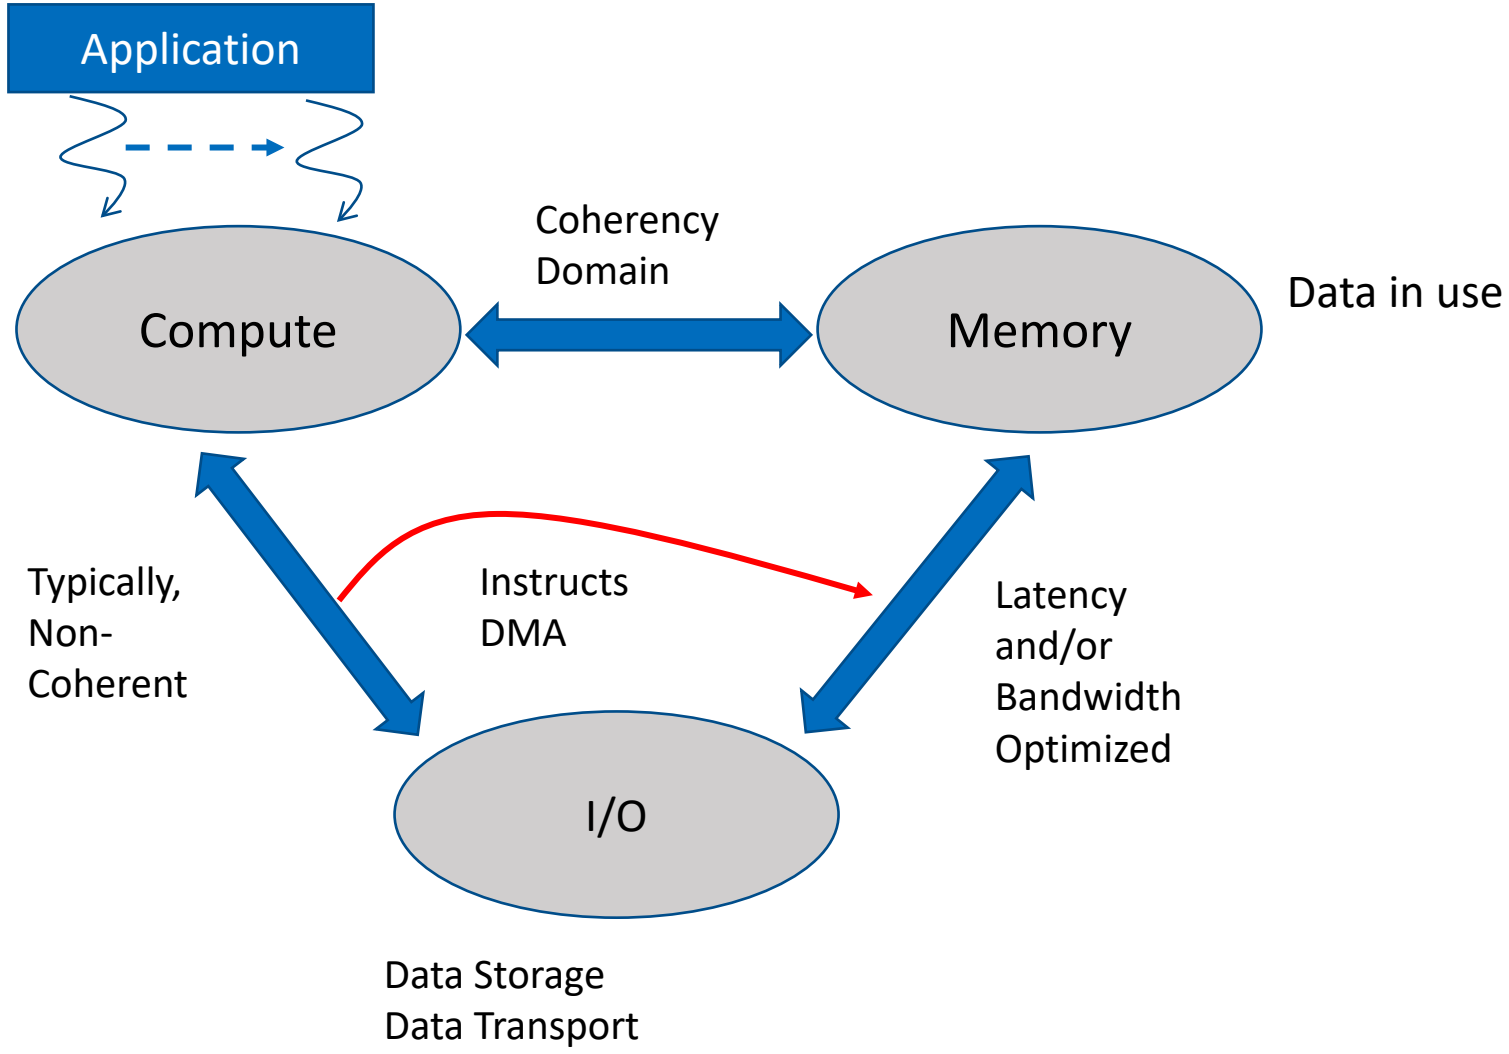

Emerging Bubbles

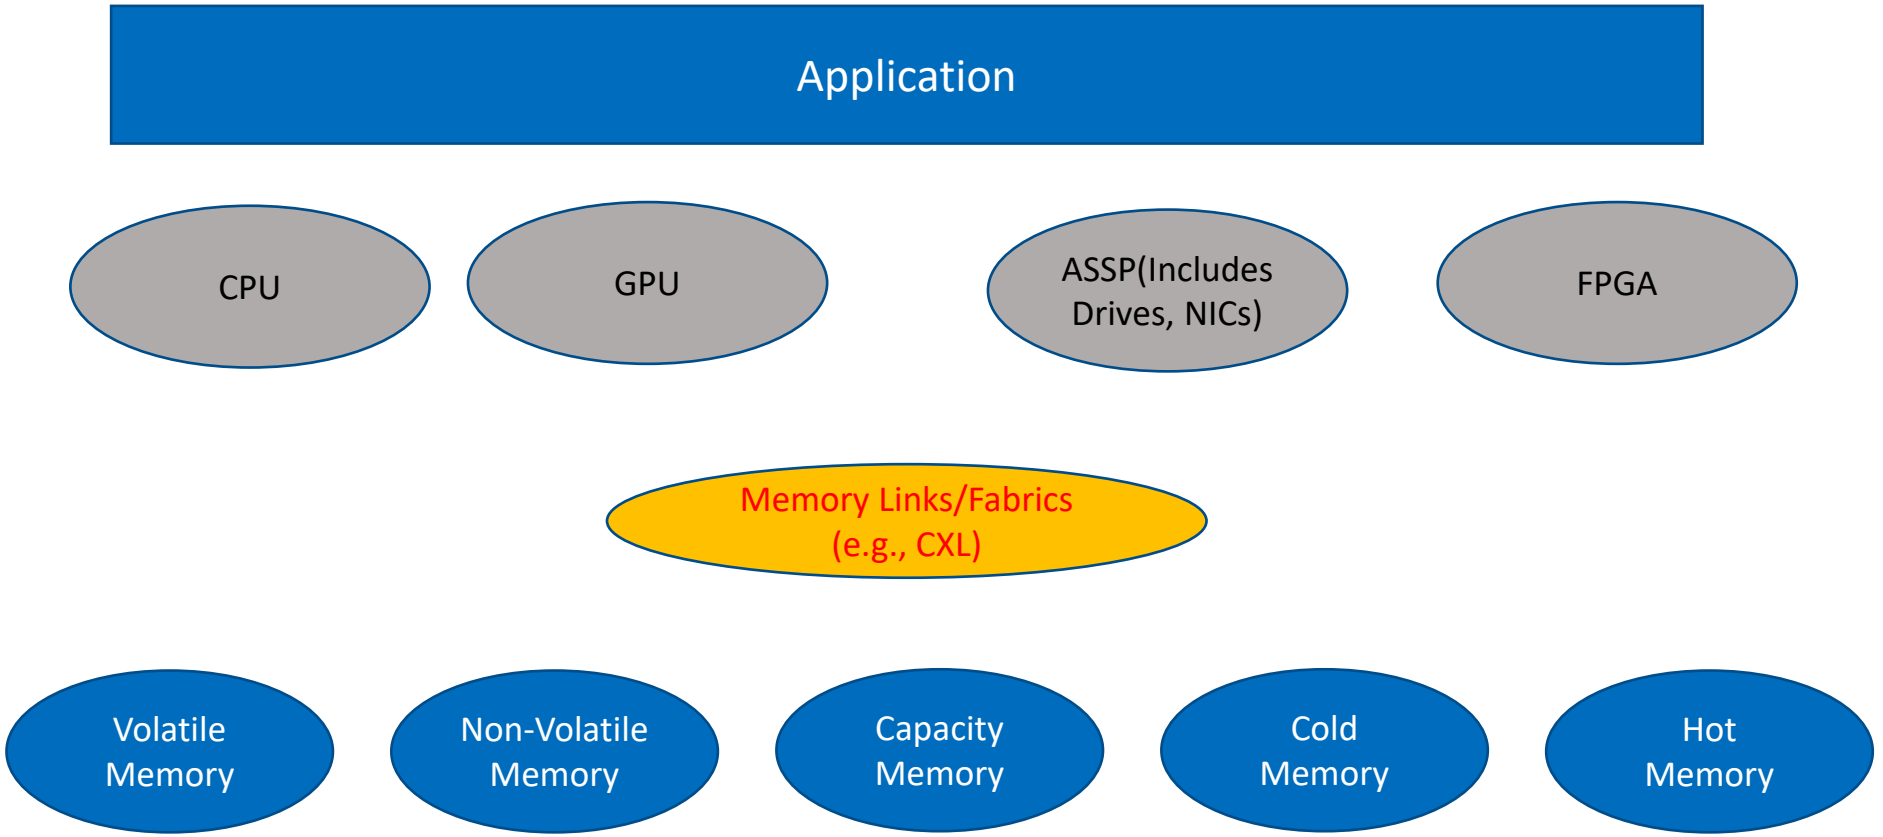

Shared Design constraints

- Latency
- Bandwidth
- Coherency
- Control

Stack View