SNIA NSF | NETWORKING STORAGE

What's in a Name? Memory Semantics and Data Movement with CXL™ and SDXI

Live Webcast

November 30, 2022

10:00 am PT / 1:00 pm ET

Today's Presenters

David McIntyre
SNIA Board of Directors
Samsung

Shyam Iyer
Chair, SNIA SDXI Technical Work Group
SNIA Technical Council
Distinguished Engineer, Dell

Rita Gupta
CXL Consortium Contributor Member,
Memory Systems Work Group
CXL System Architect, AMD

SNIA - By the Numbers

Industry Leading Organizations

180

Active Contributing Members

2,500

IT End Users & Storage Pros Worldwide

50,000

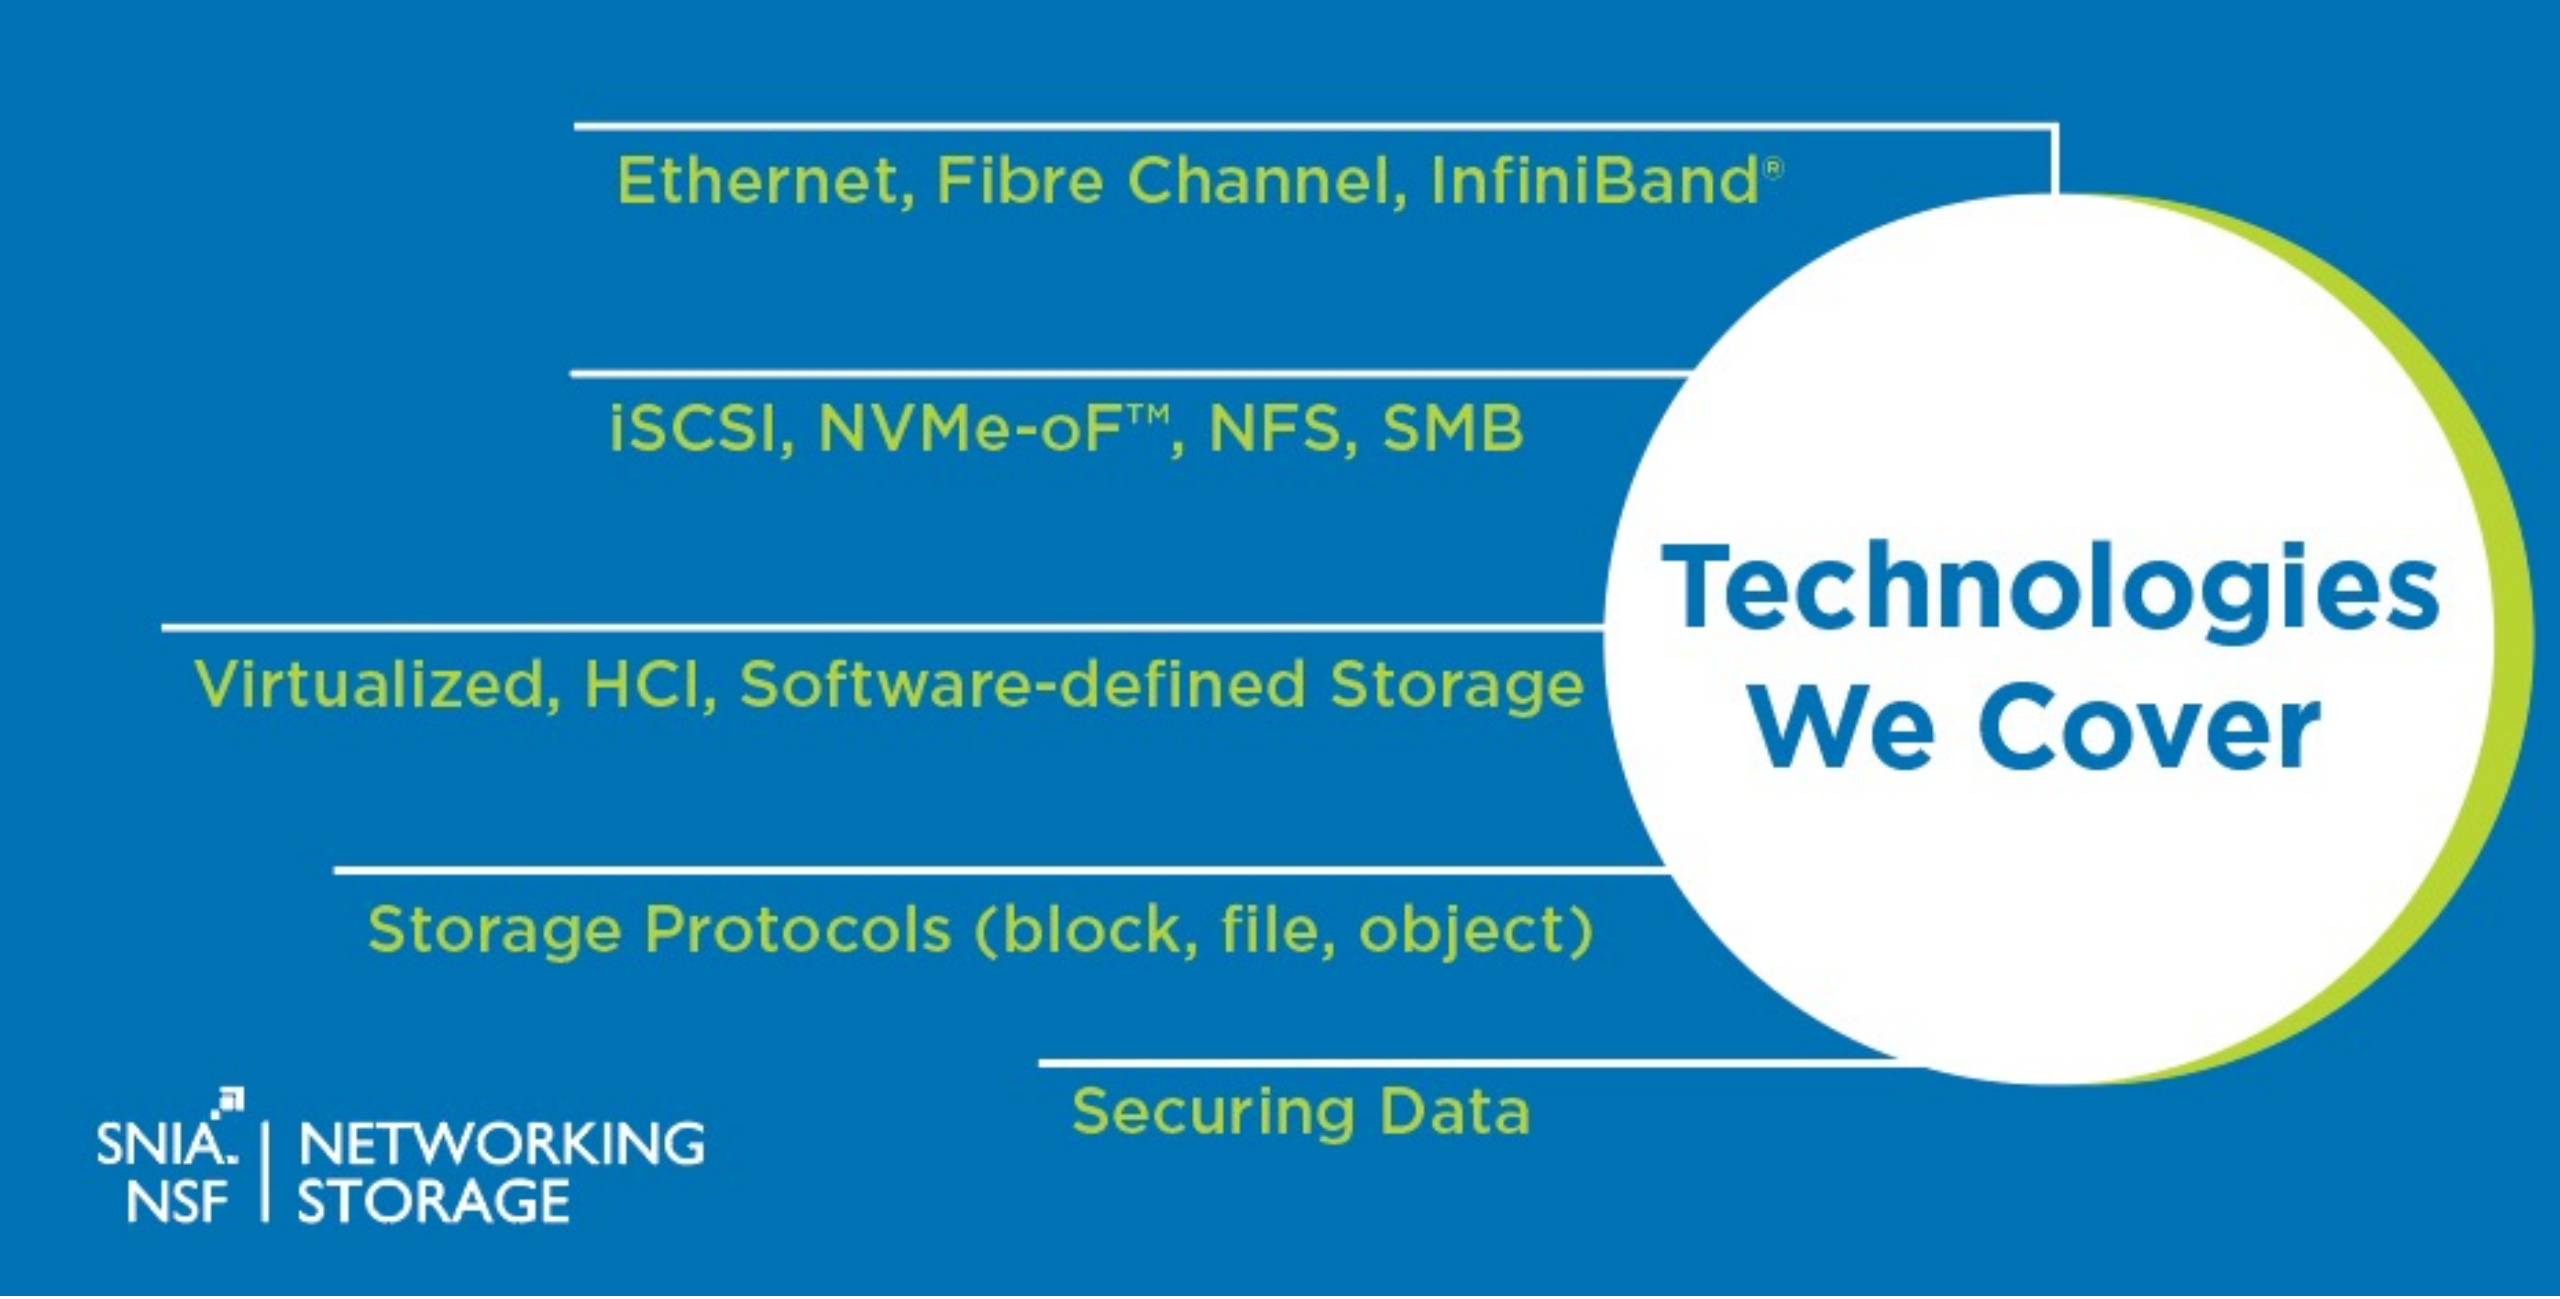

SNIA Legal Notice

- The material contained in this presentation is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Agenda

- Introduction to SDXI and CXL
- Heterogeneous Compute and Data Movement Needs with CXL
- SDXI and CXL – The Path Ahead

SDXI

Shyam Iyer, Dell

SDXI Topics

- Compute, IO, Memory Bubble
  - Current Memory to Memory Data Movement Standard
- SDXI Use Cases
  - Application Patterns and benefits of Data Movement & Acceleration
- SNIA SDXI TWG
  - Goals and Tenets
  - A brief introduction to the internals of SDXI Specification
  - SDXI Community
  - SDXI Futures
  - References, Links, and Announcements

Legacy Compute, IO, Memory Bubbles

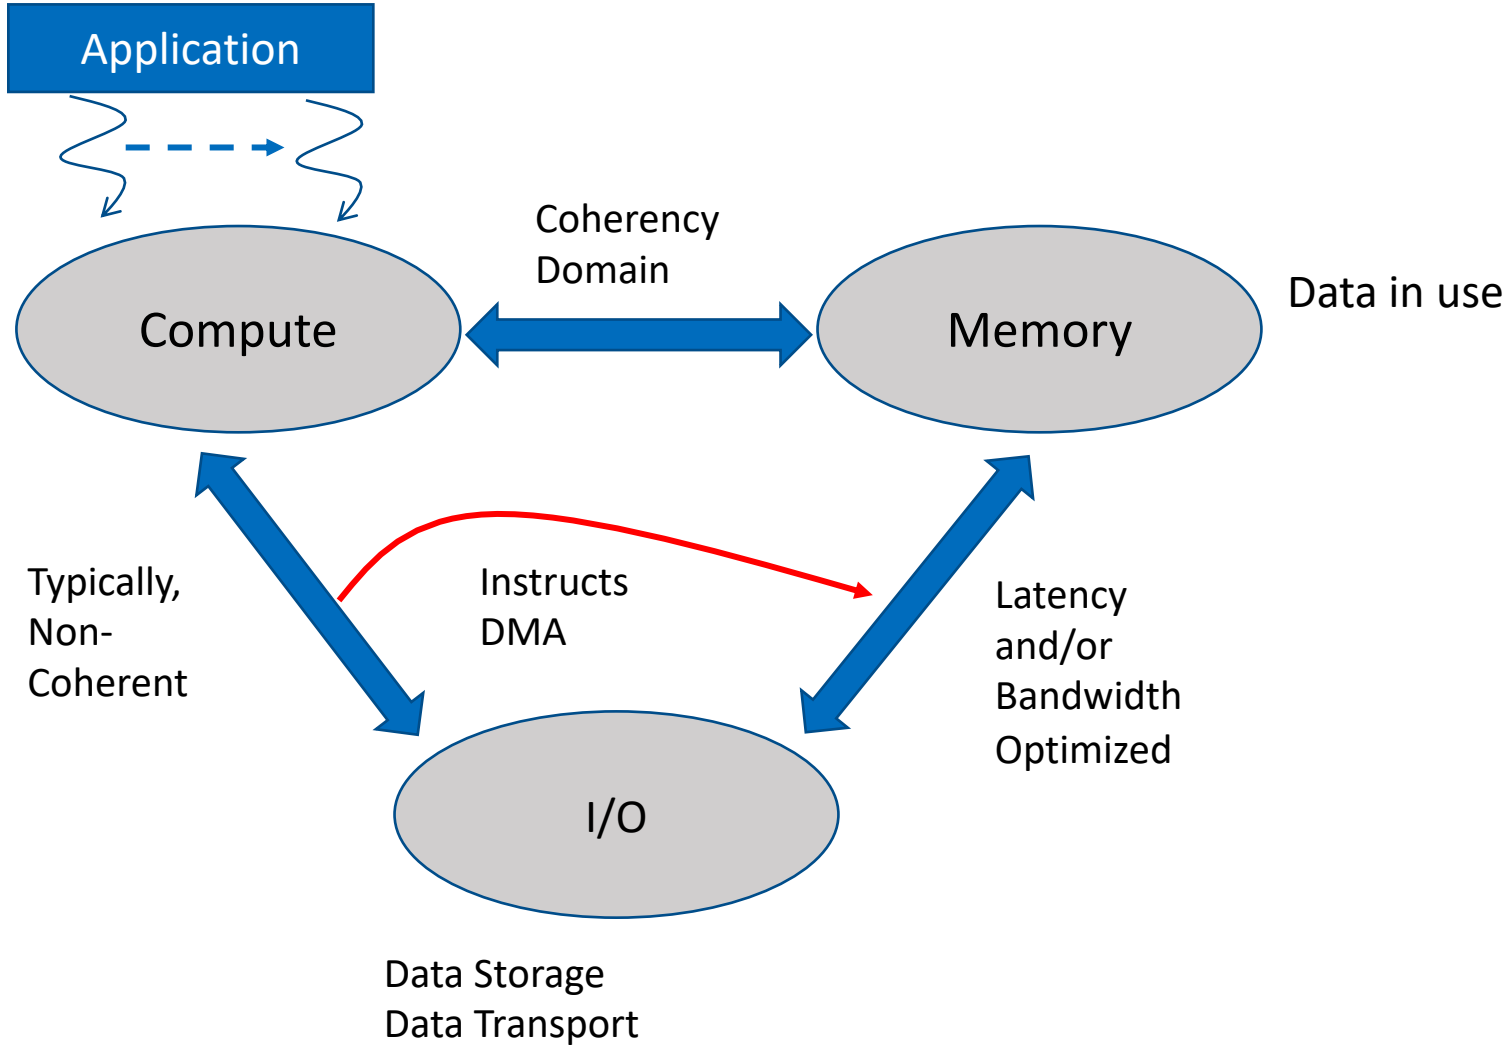

Emerging Bubbles

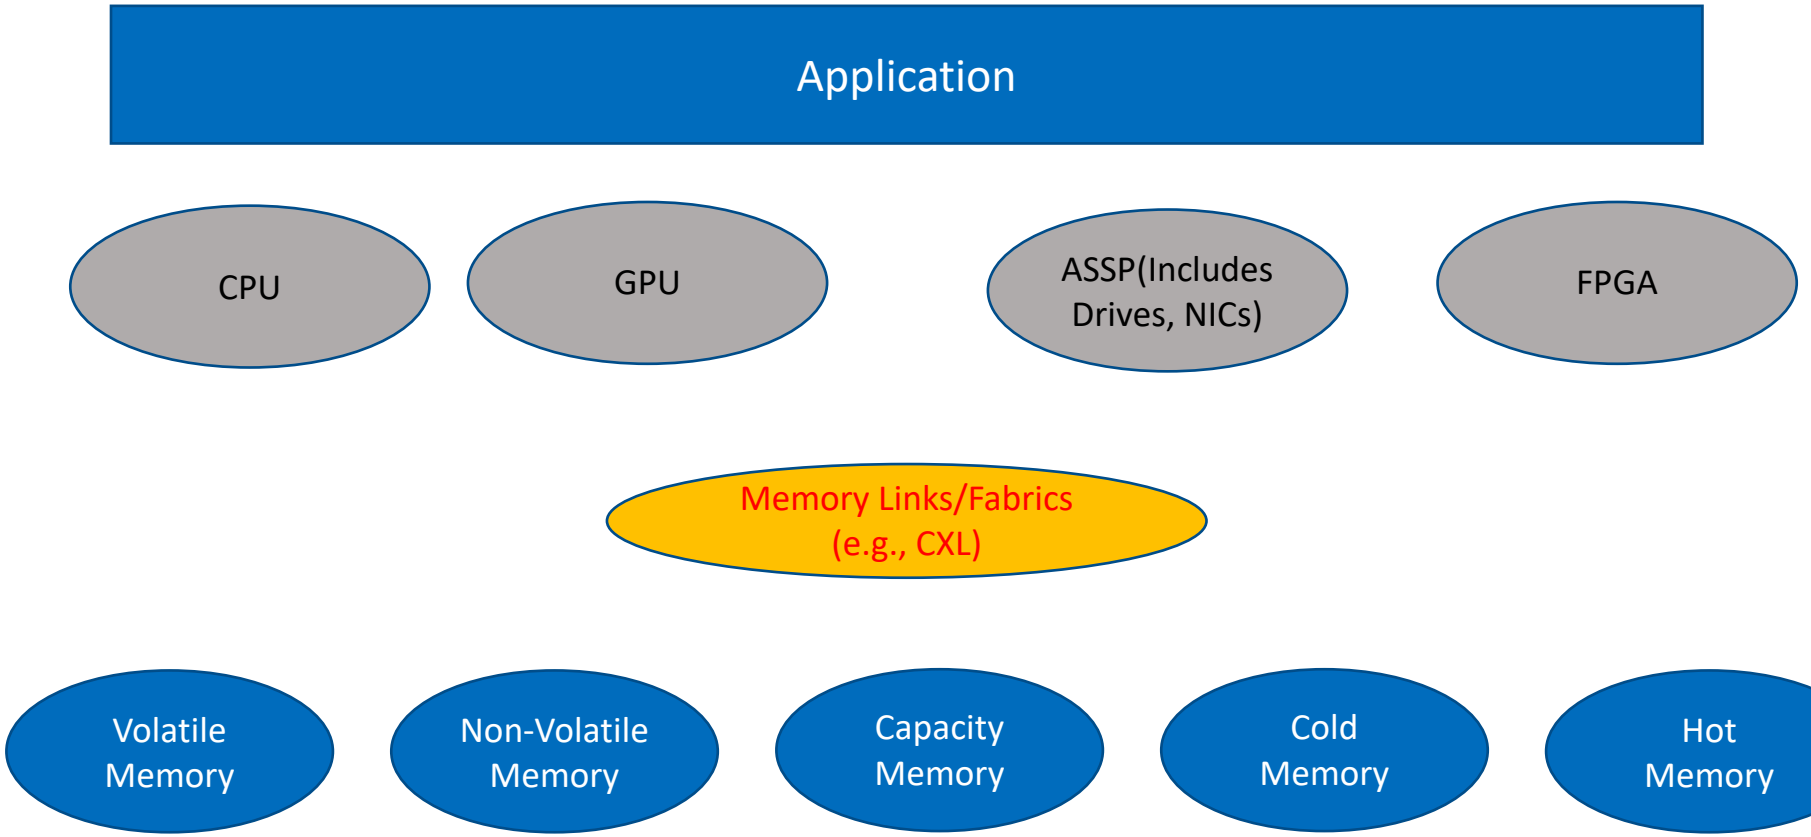

Shared Design constraints

- Latency
- Bandwidth
- Coherency
- Control

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

- Compute, IO, Memory Bubble
  - Current Memory to Memory Data Movement Standard
- Use Cases
  - Application Patterns and benefits of Data Movement & Acceleration
- SNIA SDXI TWG
  - Goals and Tenets
  - A brief introduction to the internals of SDXI Specification
  - SDXI Community
  - SDXI Futures
  - References, Links, and Announcements

Application Pattern 1 (Buffer Copies)

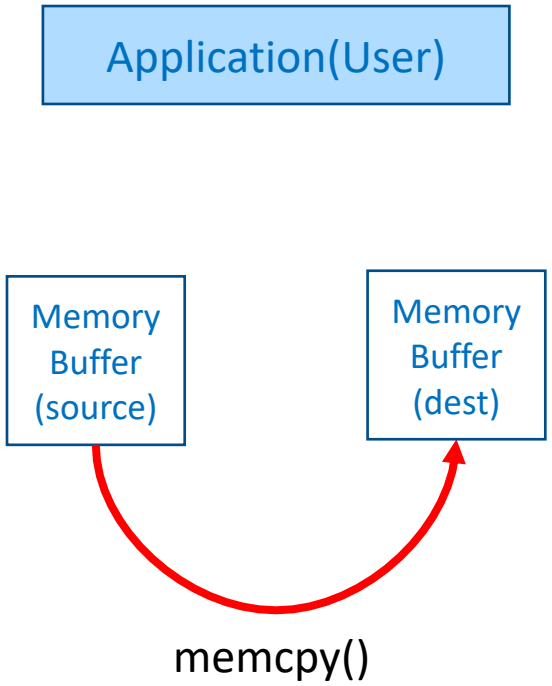

- Takes away from application performance

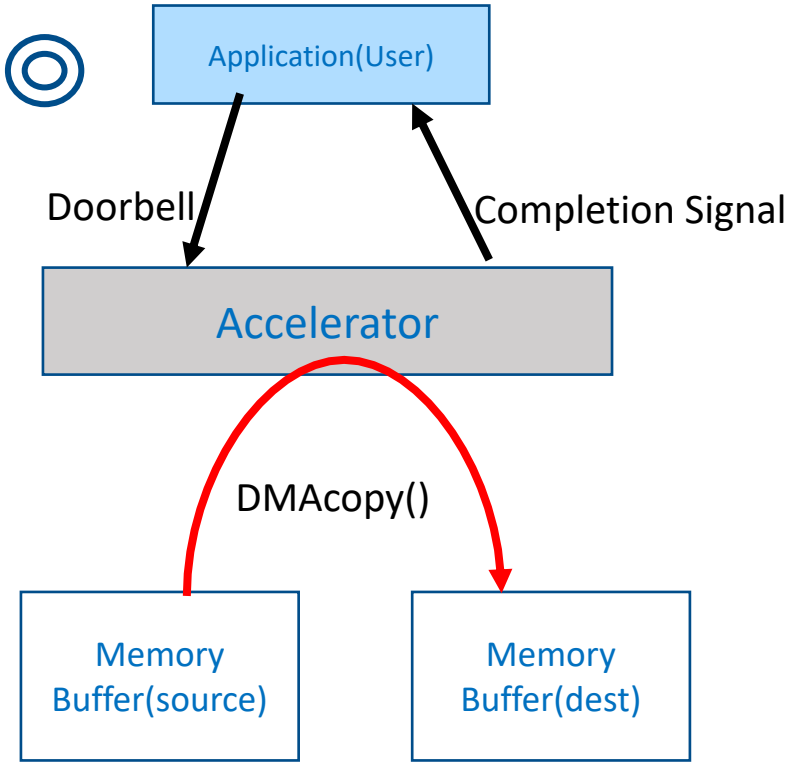

- HW based memory copies can be offloaded without affecting application performance

Application Pattern 2

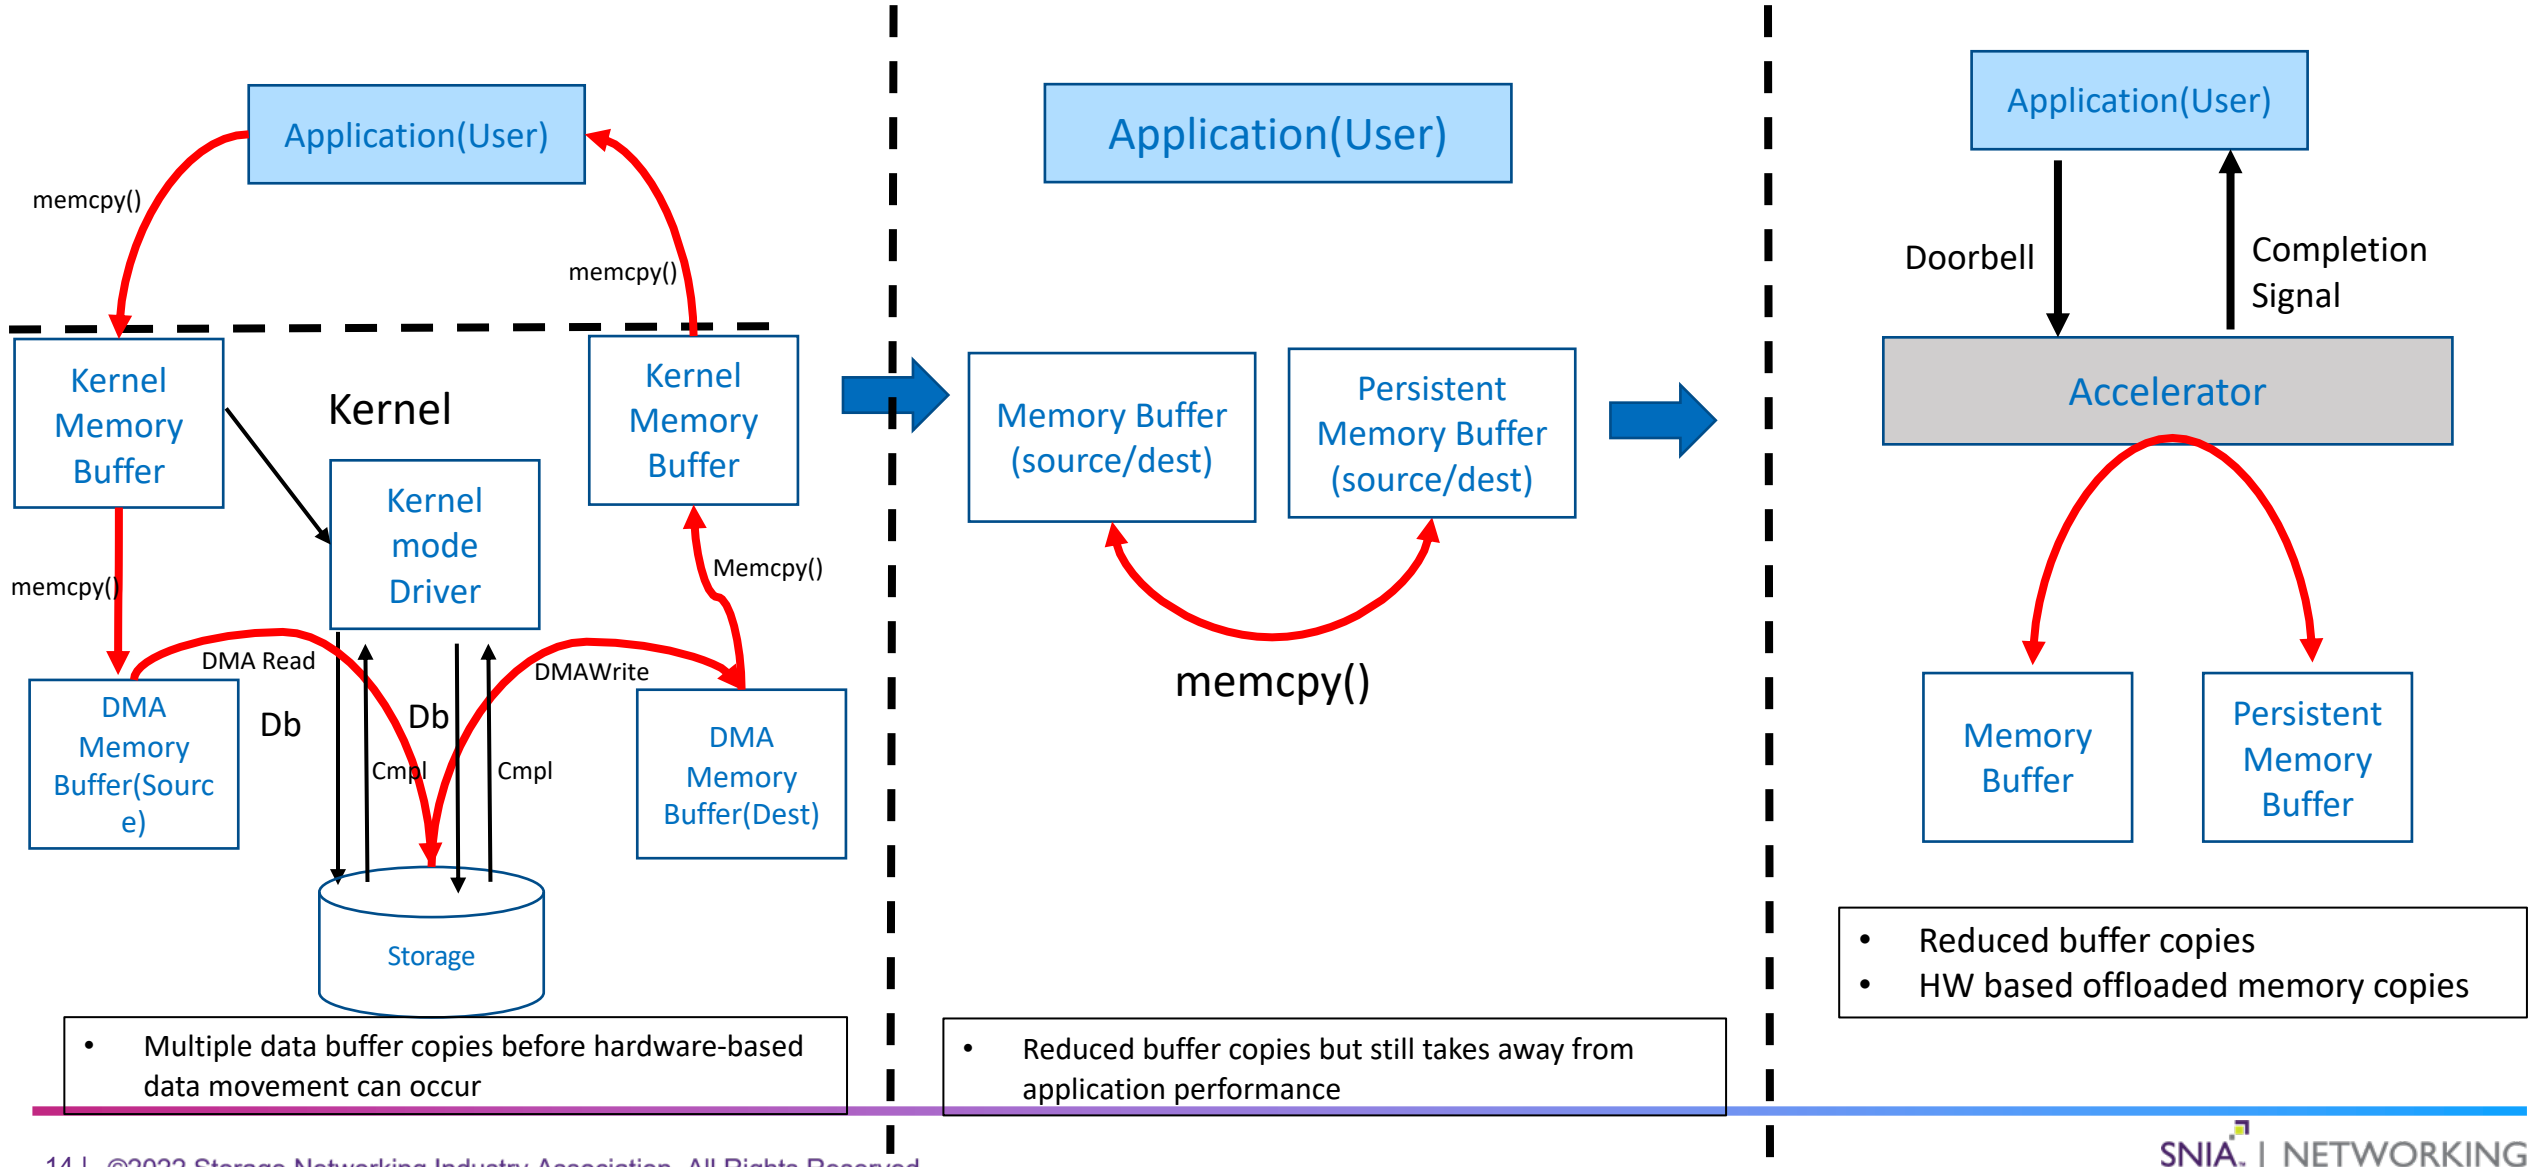

- Multiple data buffer copies before hardware-based data movement can occur

- Reduced buffer copies but still takes away from application performance

- Reduced buffer copies
- HW based offloaded memory copies

Application Pattern 3

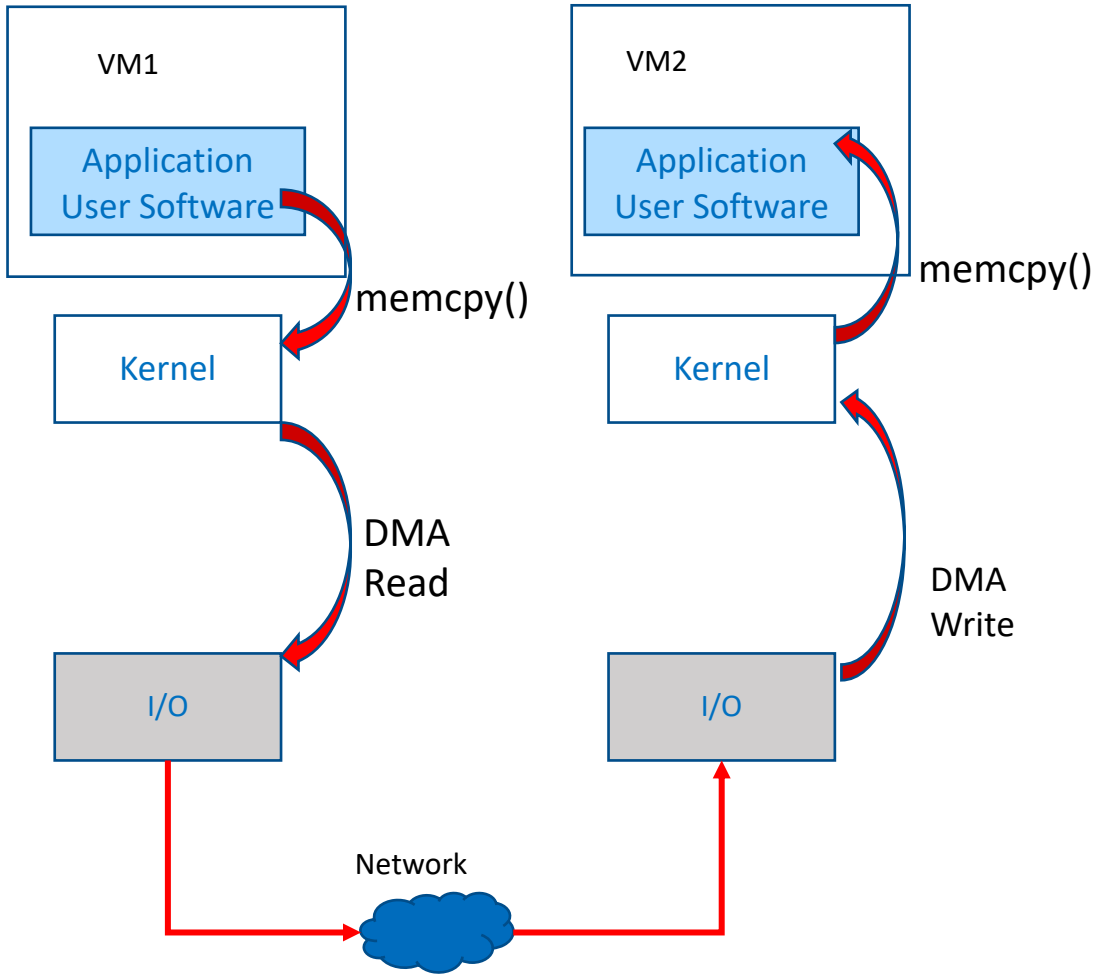

- Context isolation layers introduce multiple buffer copies

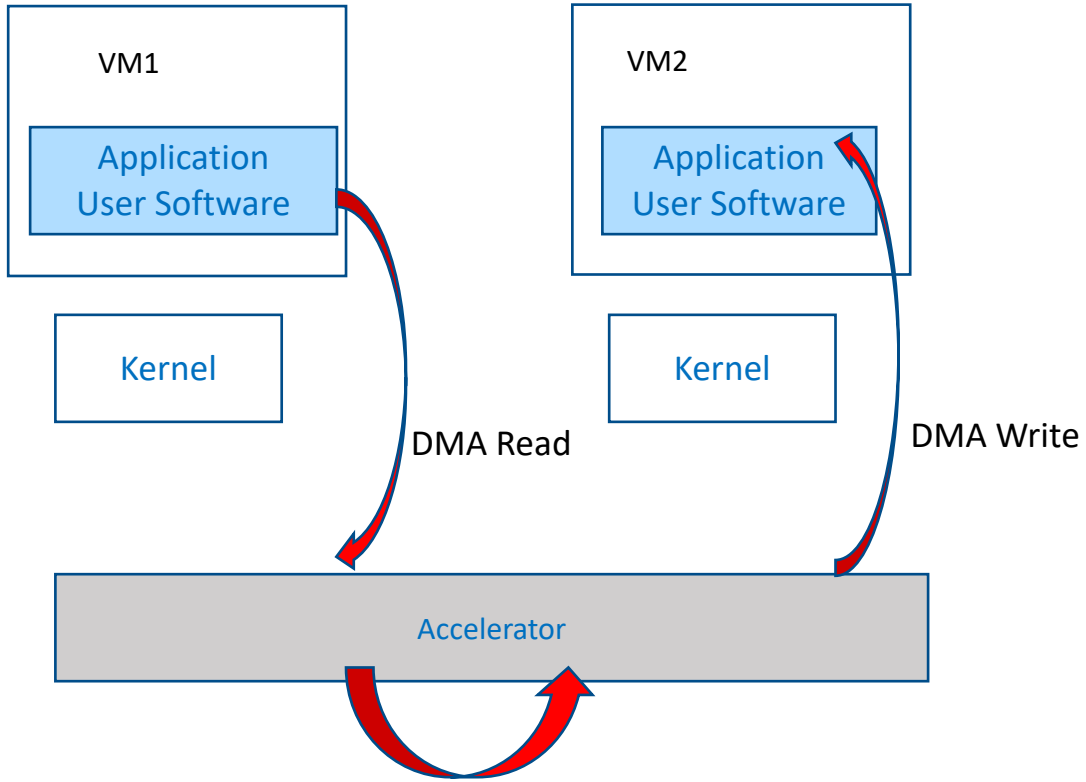

- Best of both: Context isolation layers and optimized HW based memory buffer copies

Data *in use* Memory Expansion

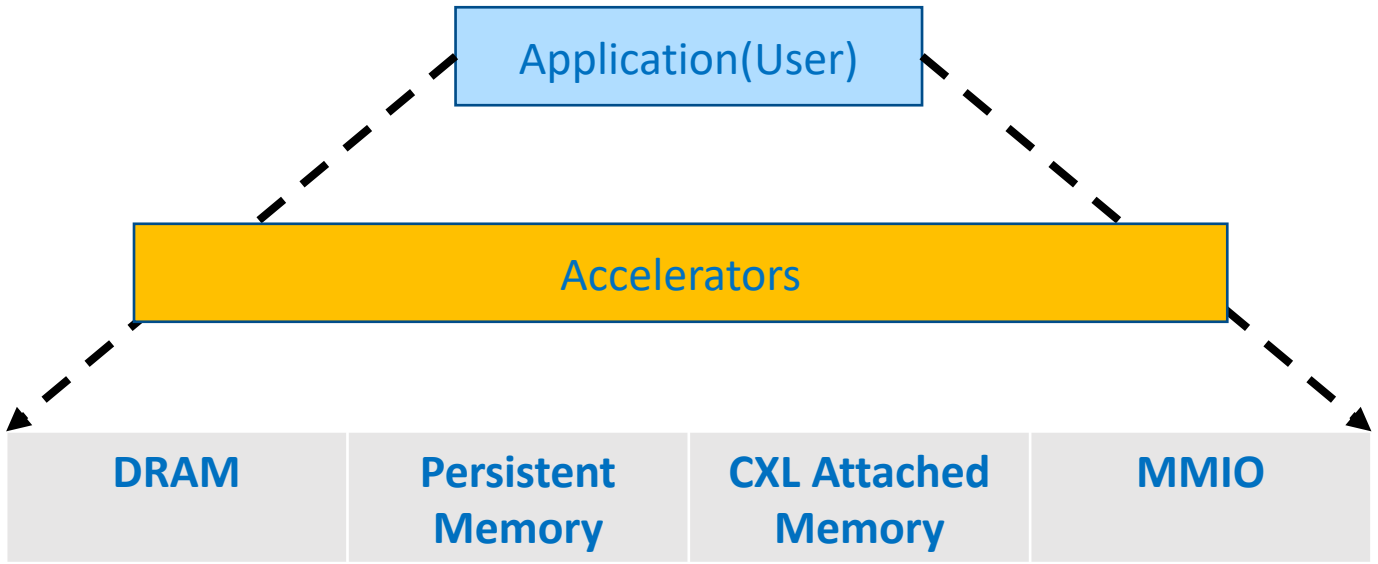

- Memory expansion expands the memory target surface area for accelerators
- Different tiers of memory
- Diversity in accelerator programming methods

Stack View

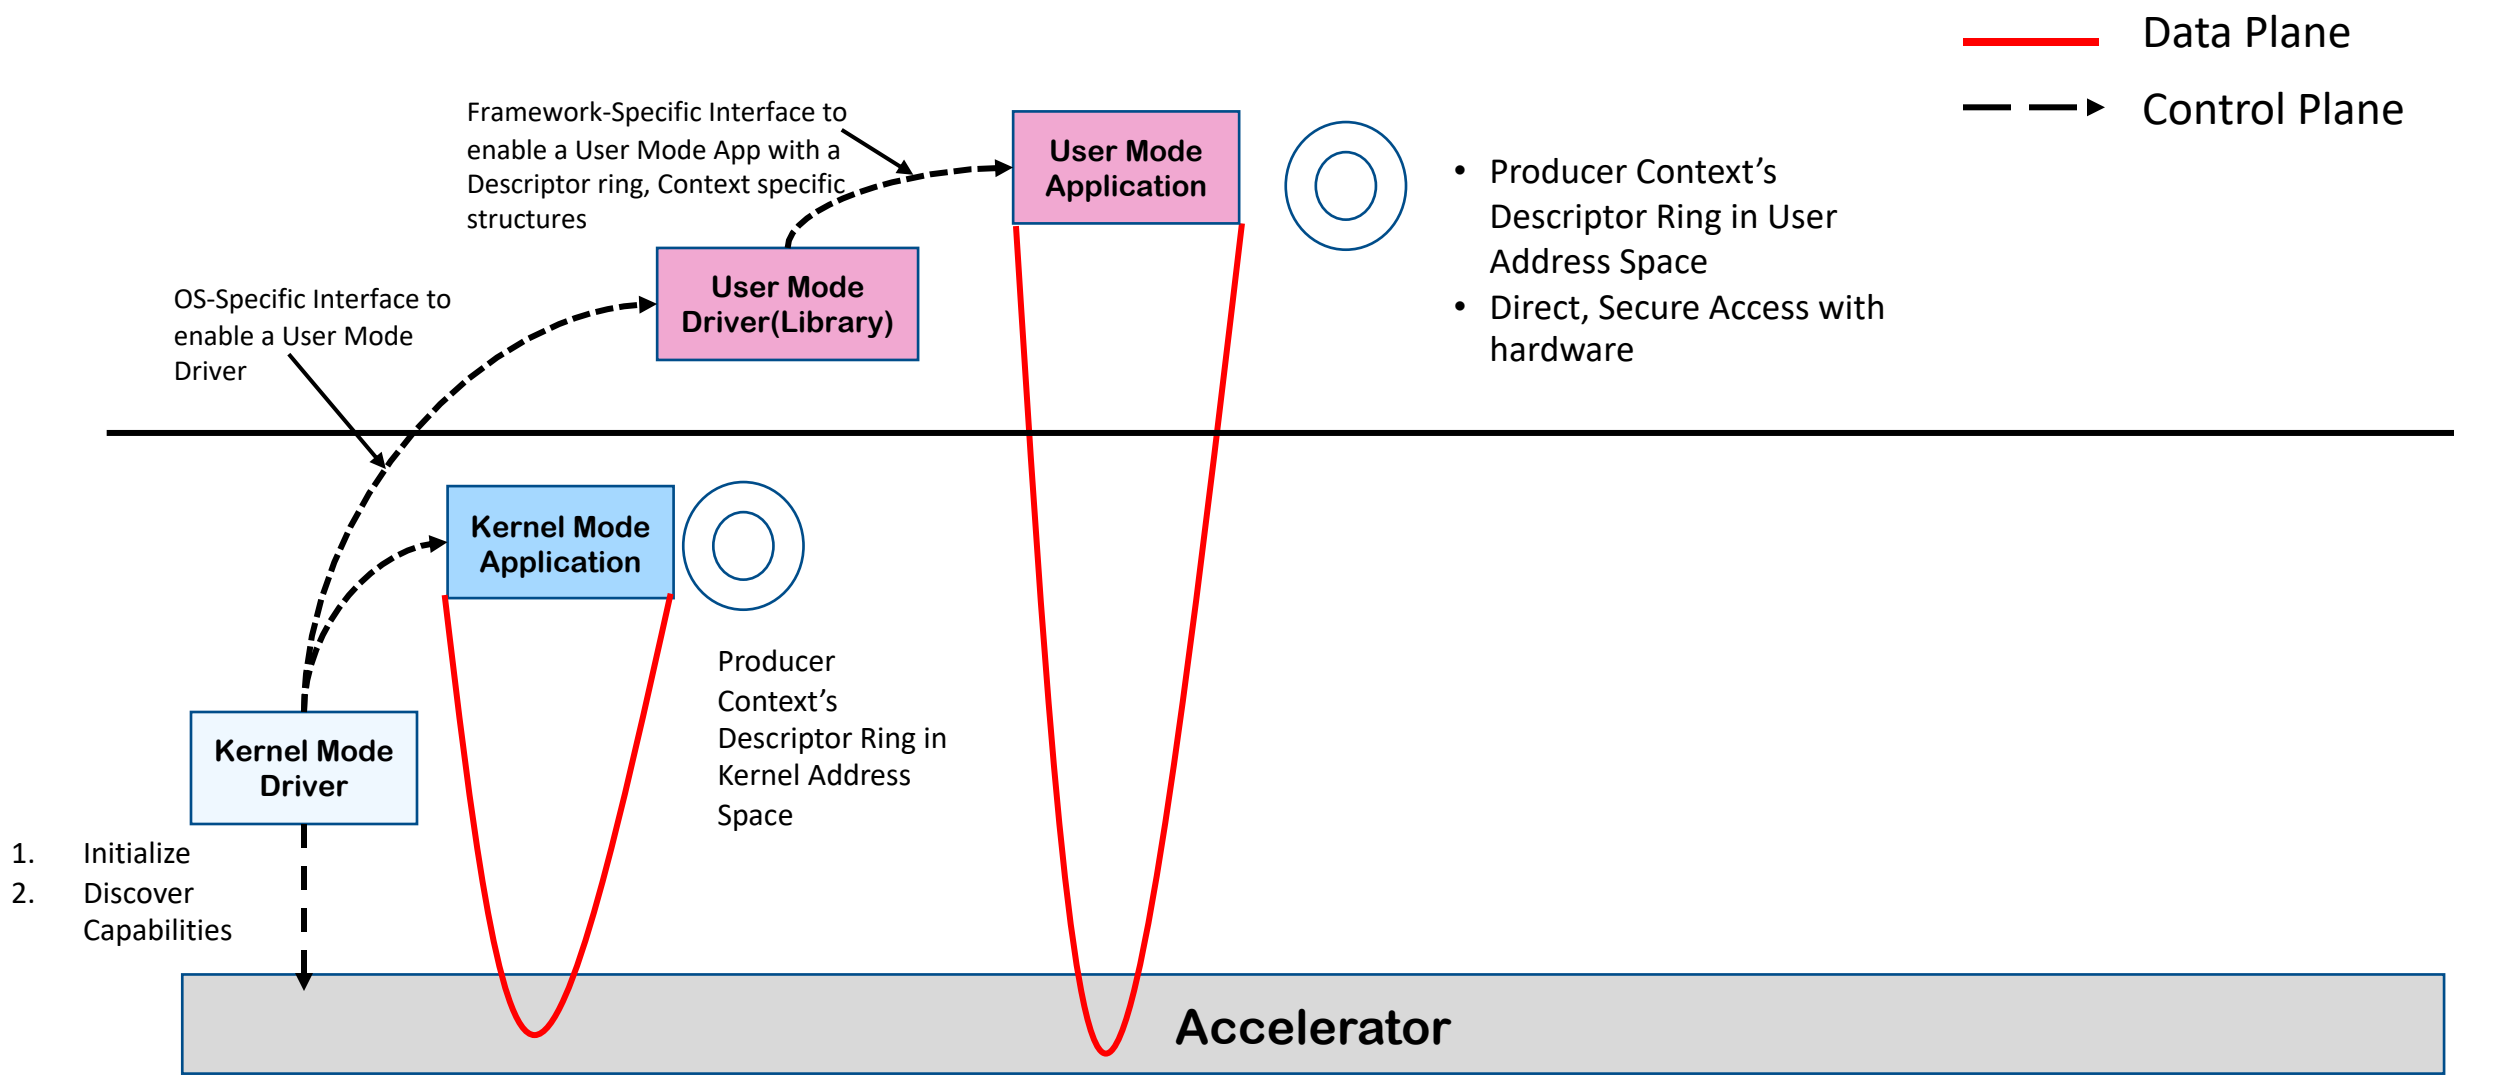

- Producer Context's Descriptor Ring in User Address Space
- Direct, Secure Access with hardware

1. Initialize
2. Discover Capabilities

SDXI

- Compute, IO, Memory Bubble
  - Current Memory to Memory Data Movement Standard
- Use Cases
  - Application Patterns and benefits of Data Movement & Acceleration
- SNIA SDXI TWG
  - Goals and Tenets
  - A brief introduction to the internals of SDXI Specification
  - SDXI Community
  - SDXI Futures
  - References, Links, and Announcements

SDXI(Smart Data Accelerator Interface)

- Smart Data Accelerator Interface (SDXI) is a proposed 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
  - 28 member companies, 80+ individual members

SDXI Memory-to-Memory Data Movement

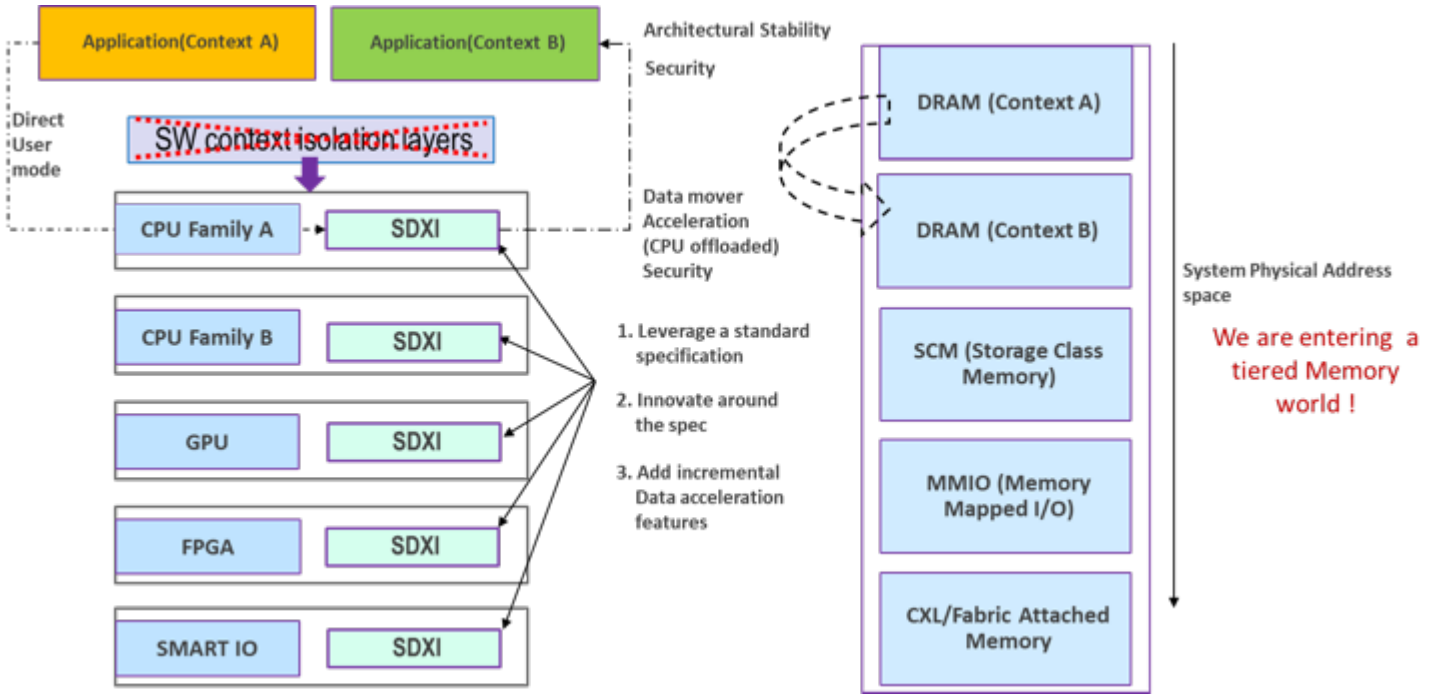

- Data movement between different address spaces.
- Data movement without mediation by privileged software.
- Allows abstraction or virtualization by privileged software.
- Capability to quiesce, suspend, and resume the architectural state of a per-address-space data mover.
- Forward and backward compatibility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

Memory Structures(1) – Simplified View

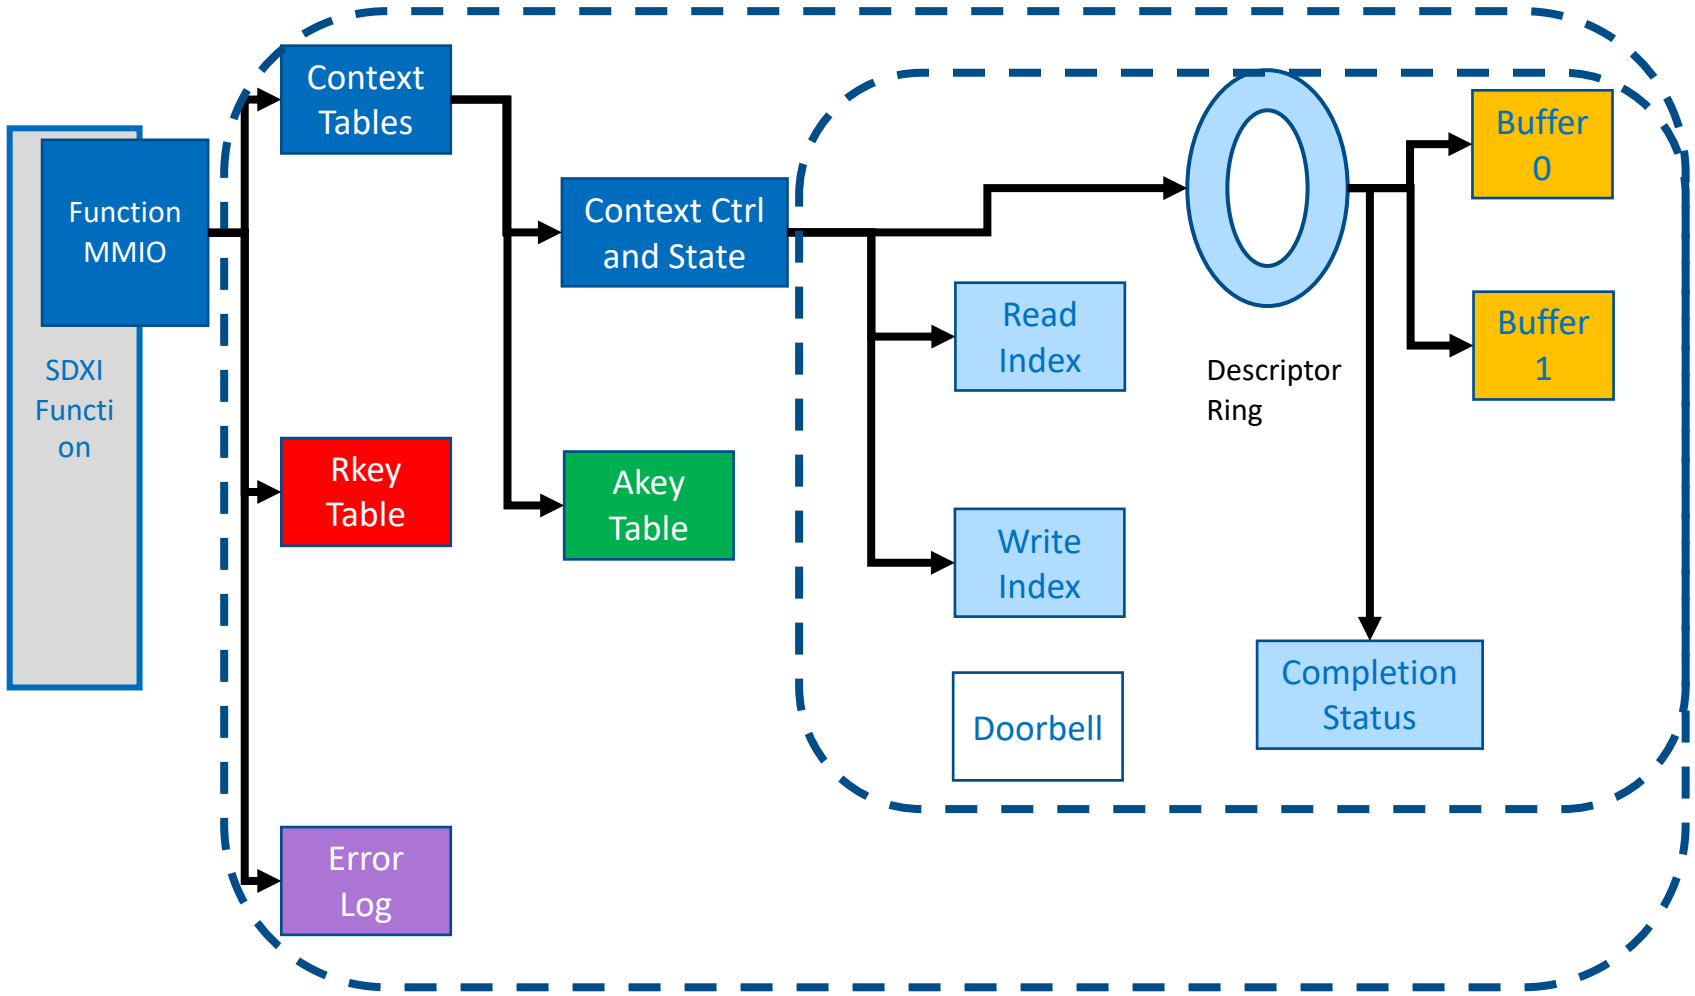

- All states in memory
- One standard descriptor format
  - Scope for future expansion
- Easy to virtualize
- Architected function setup and control
  - *layered model for interconnect specific function management
  - SDXI class code registered for PCIe implementations

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

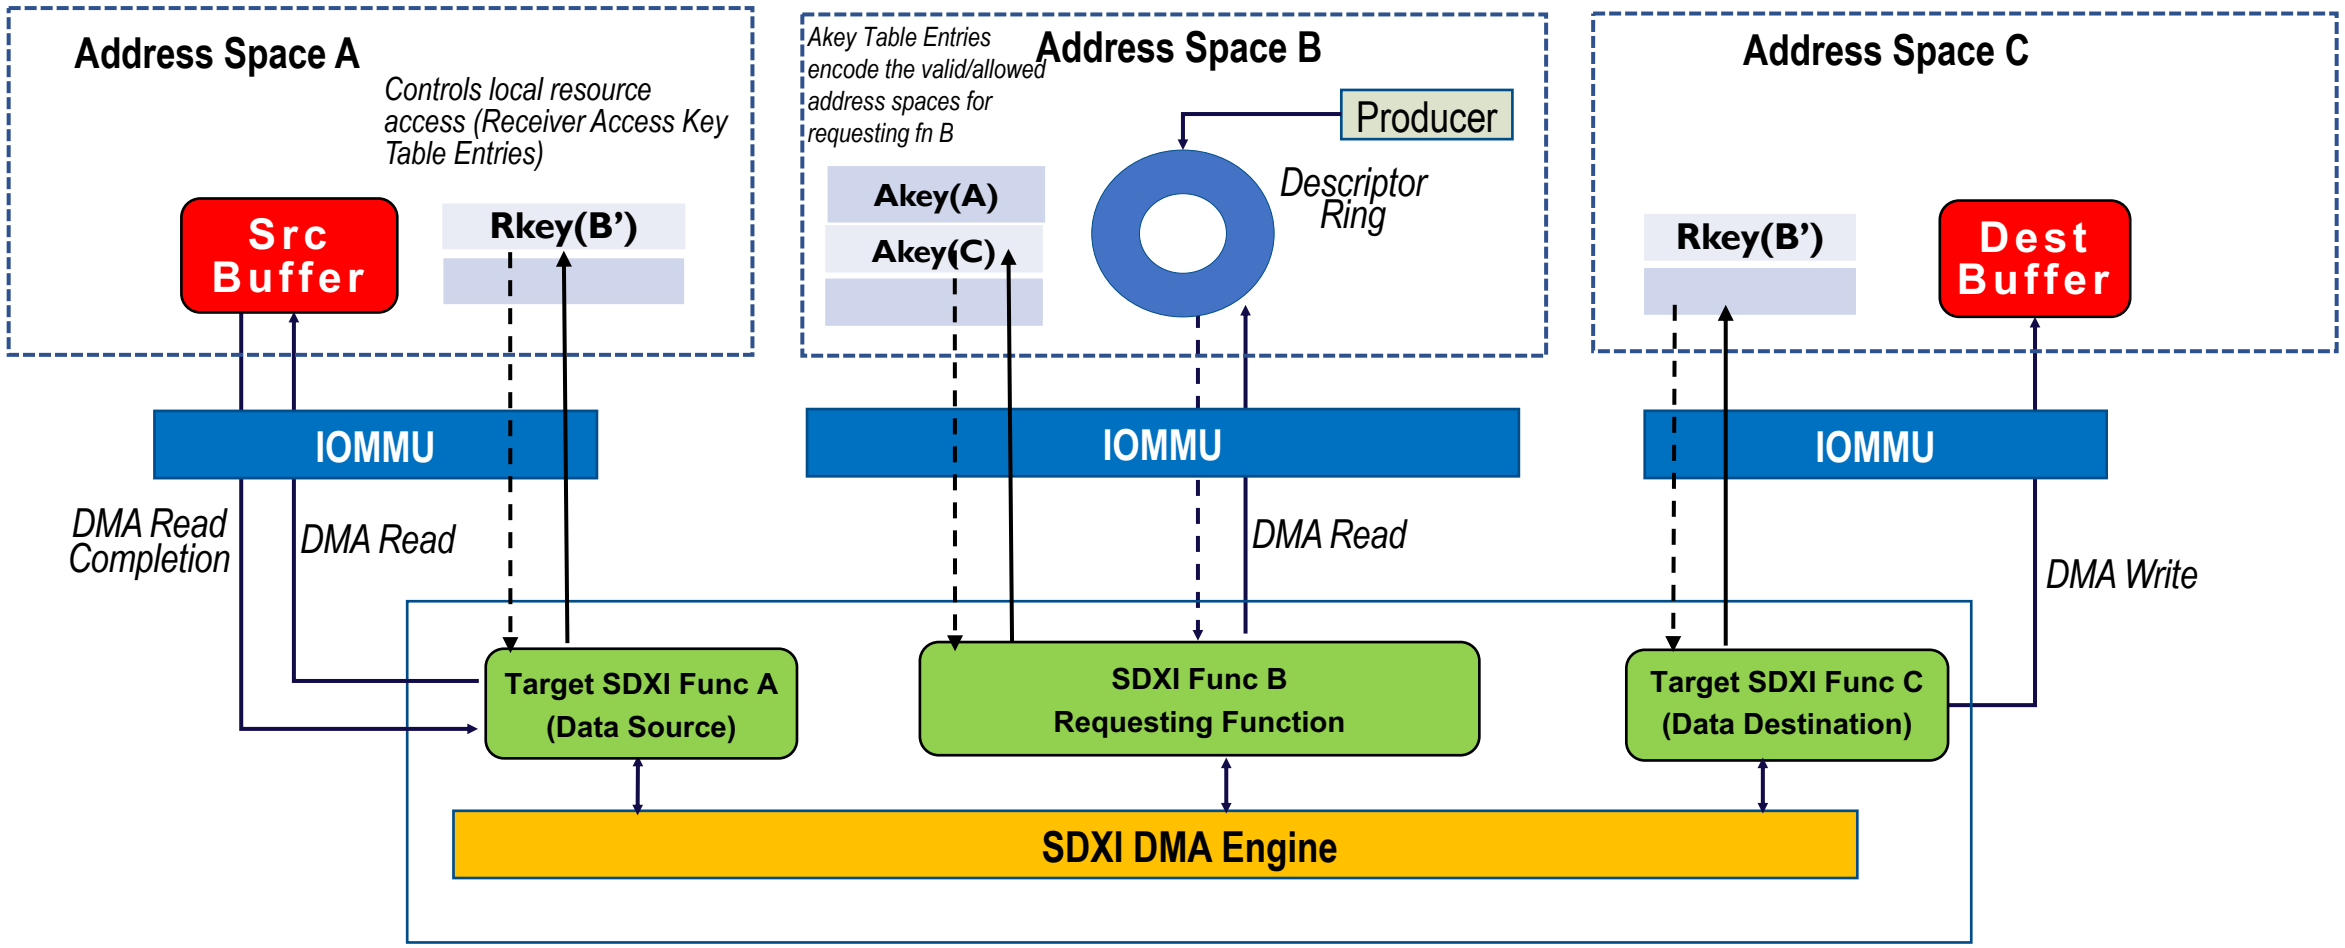

Active Contributors as of November 2022

SDXI v1.0 and Beyond

- **v1.0 released!**
  - https://www.snia.org/sdxi
- Post v1.0 activities. The current charter includes:
  - New data mover operations for smart acceleration
  - Data mover operations involving persistent memory targets
  - Cache coherency models for data movers
  - Security Features involving data movers
  - Management architecture for data movers (includes connection manager)
- Some additional discussion topics being considered post-v1.0
  - QoS improvements
  - Latency improvements
  - RAS improvements
  - CXL related discussions
  - Heterogenous environments

CXL

Rita Gupta, AMD

Agenda

- CXL™ Consortium Update
- Compute Express Link™ Overview
- CXL Features Gen to Gen
- CXL Usage models
- Heterogeneous Compute and Data Movement

CXL Board of Directors

CXL Compute Express Link™

Industry Open Standard for High Speed Communications

225+ Member Companies

Industry Focal Point

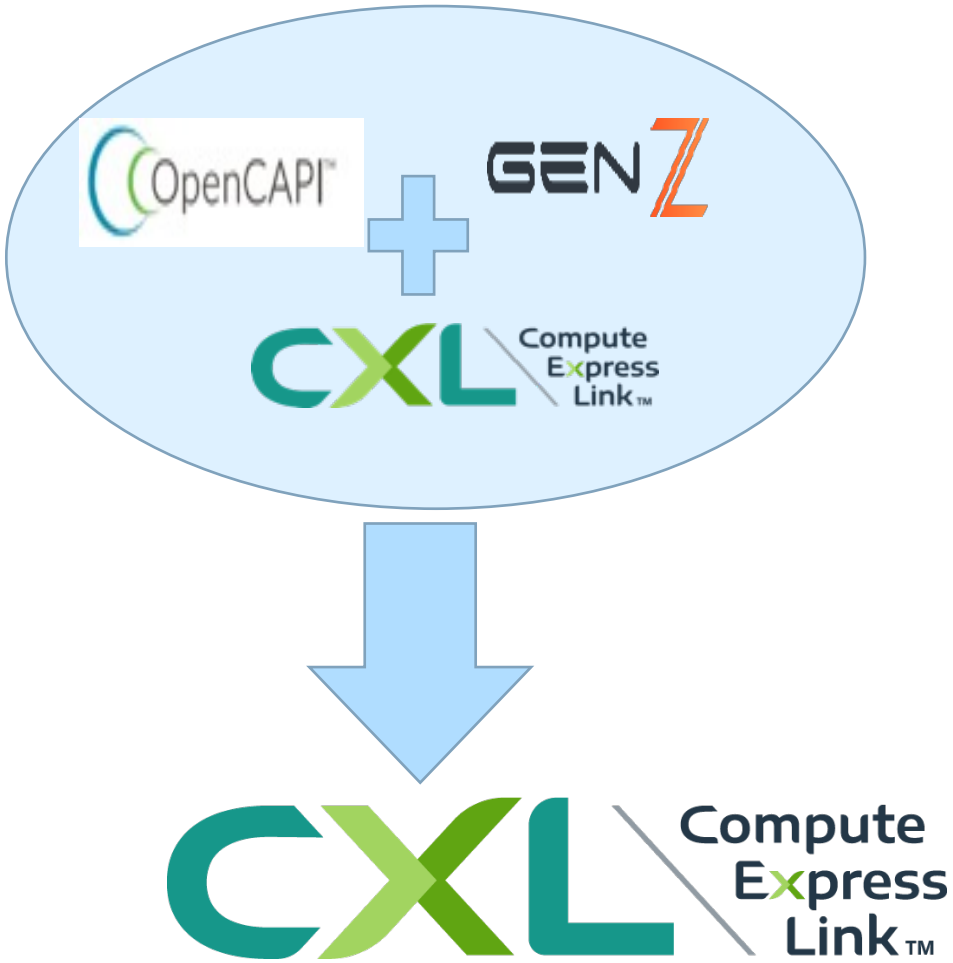

CXL is emerging as the industry focal point for coherent IO

- CXL Consortium and OpenCAPI sign letter of intent to transfer OpenCAPI specification and assets to the CXL Consortium

CXL Compute Express Link™ OpenCAPI™

August 2, 2022, Flash Memory Summit
CXL Consortium and OpenCAPI Consortium Sign Letter of Intent to Transfer OpenCAPI Assets to CXL

- In February 2022, CXL Consortium and Gen-Z Consortium signed agreement to transfer Gen-Z specification and assets to CXL Consortium

CXL Specification Release Timeline

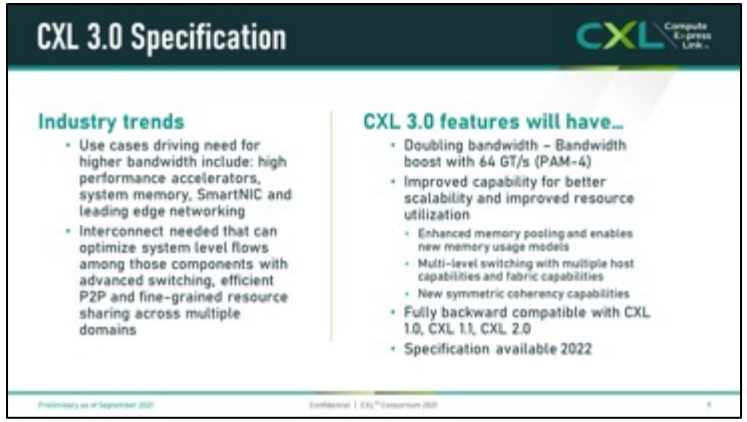

March 2019

CXL 1.0 Specification Released

September 2019

CXL Consortium Officially Incorporates CXL 1.1 Specification Released

November 2020

CXL 2.0 Specification Released

August 2022

CXL 3.0 Specification Released

CXL™ Features and Benefits

CXL Delivers the Right Features & Architecture

Challenges

Industry trends driving demand for faster data processing and next-gen data center performance

Increasing demand for heterogeneous computing and server disaggregation

Need for increased memory capacity and bandwidth

Lack of open industry standard to address next-gen interconnect challenges

CXL

An open industry-supported cache-coherent interconnect for processors, memory expansion and accelerators

Coherent Interface

Leverages PCIe with 3 mix-and-match protocols: CXL.IO, CXL.CACHE, CXL.MEMORY

Low Latency

.Cache and .Memory targeted at near CPU cache latency

Asymmetric Complexity

Eases burdens of cache coherent interface designs

What is CXL?

- High bandwidth, Low latency, Cache- coherent interconnect
- Uses a flexible processor port that can auto-negotiate to either the standard PCIe transaction protocol or the alternate CXL transaction protocols
- Based on PCIe® 5.0 PHY infrastructure and leverages channel, retimers, PHY, Logical, Protocols
- Dynamic multiplexing of 3 protocols – CXL.io, CXL.cache, CXL.mem

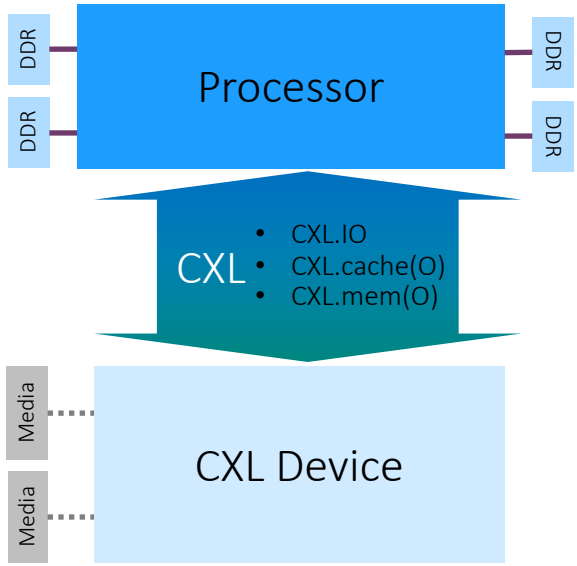

Representative CXL Usages

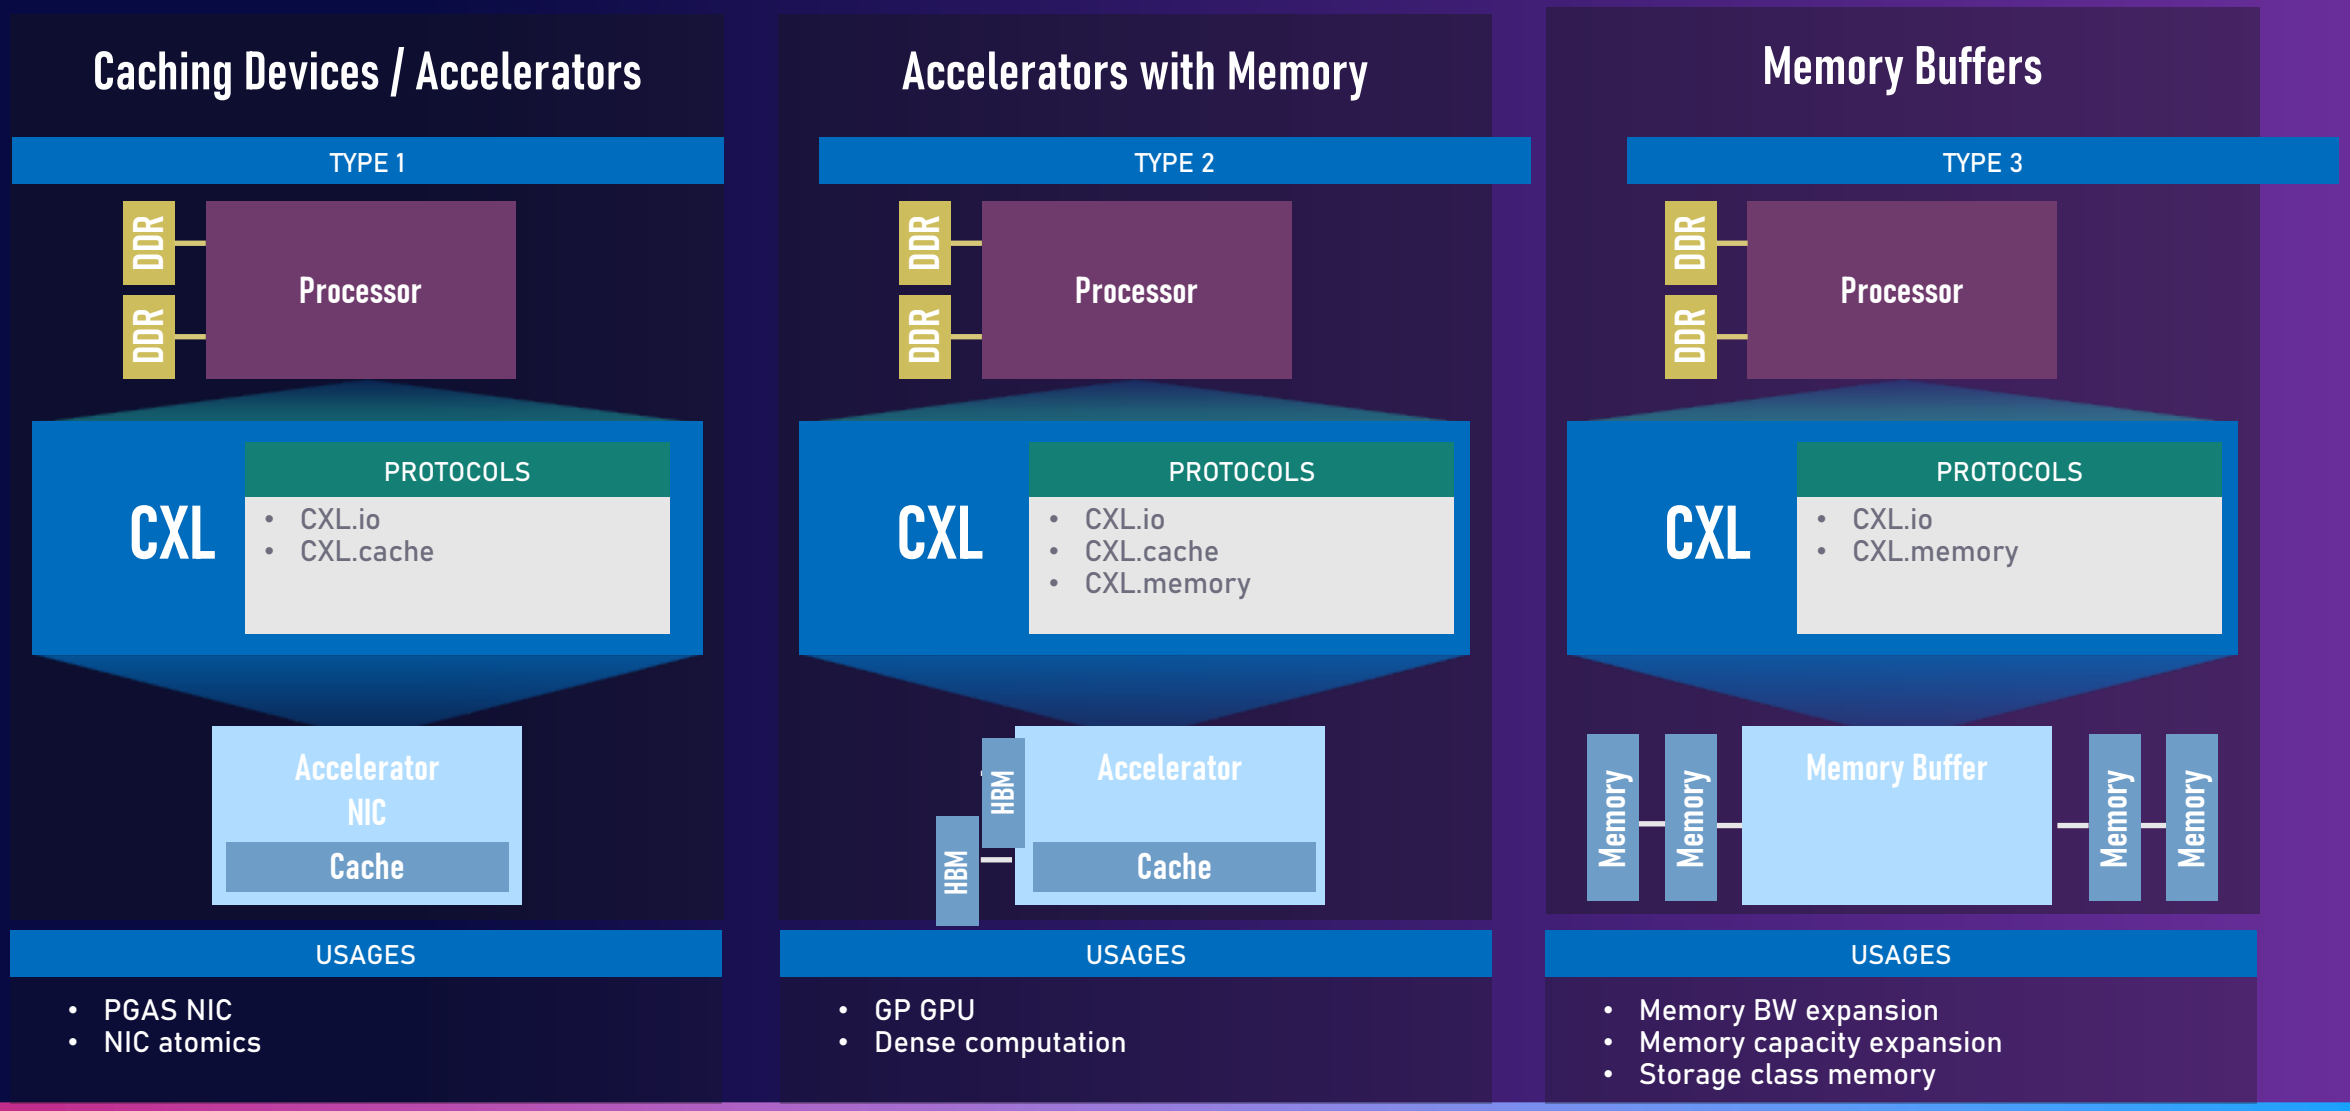

CXL Spec Feature Summary

Features	CXL 1.0 / 1.1	CXL 2.0	CXL 3.0
Release date	2019	2020	1H 2022
Max link rate	32GTs	32GTs	64GTs
Flit 68 byte (up to 32 GTs)	✓	✓	✓
Flit 256 byte (up to 64 GTs)			✓
Type 1, Type 2 and Type 3 Devices	✓	✓	✓
Memory Pooling w/ MLDs		✓	✓
Global Persistent Flush		✓	✓
CXL IDE		✓	✓
Switching (Single-level)		✓	✓
Switching (Multi-level)			✓
Direct memory access for peer-to-peer			✓
Enhanced coherency (256 byte flit)			✓
Memory sharing (256 byte flit)			✓
Multiple Type 1/Type 2 devices per root port			✓
Fabric capabilities (256 byte flit)			✓

	Not supported
✓	Supported

CXL Use Cases

Current Use Cases

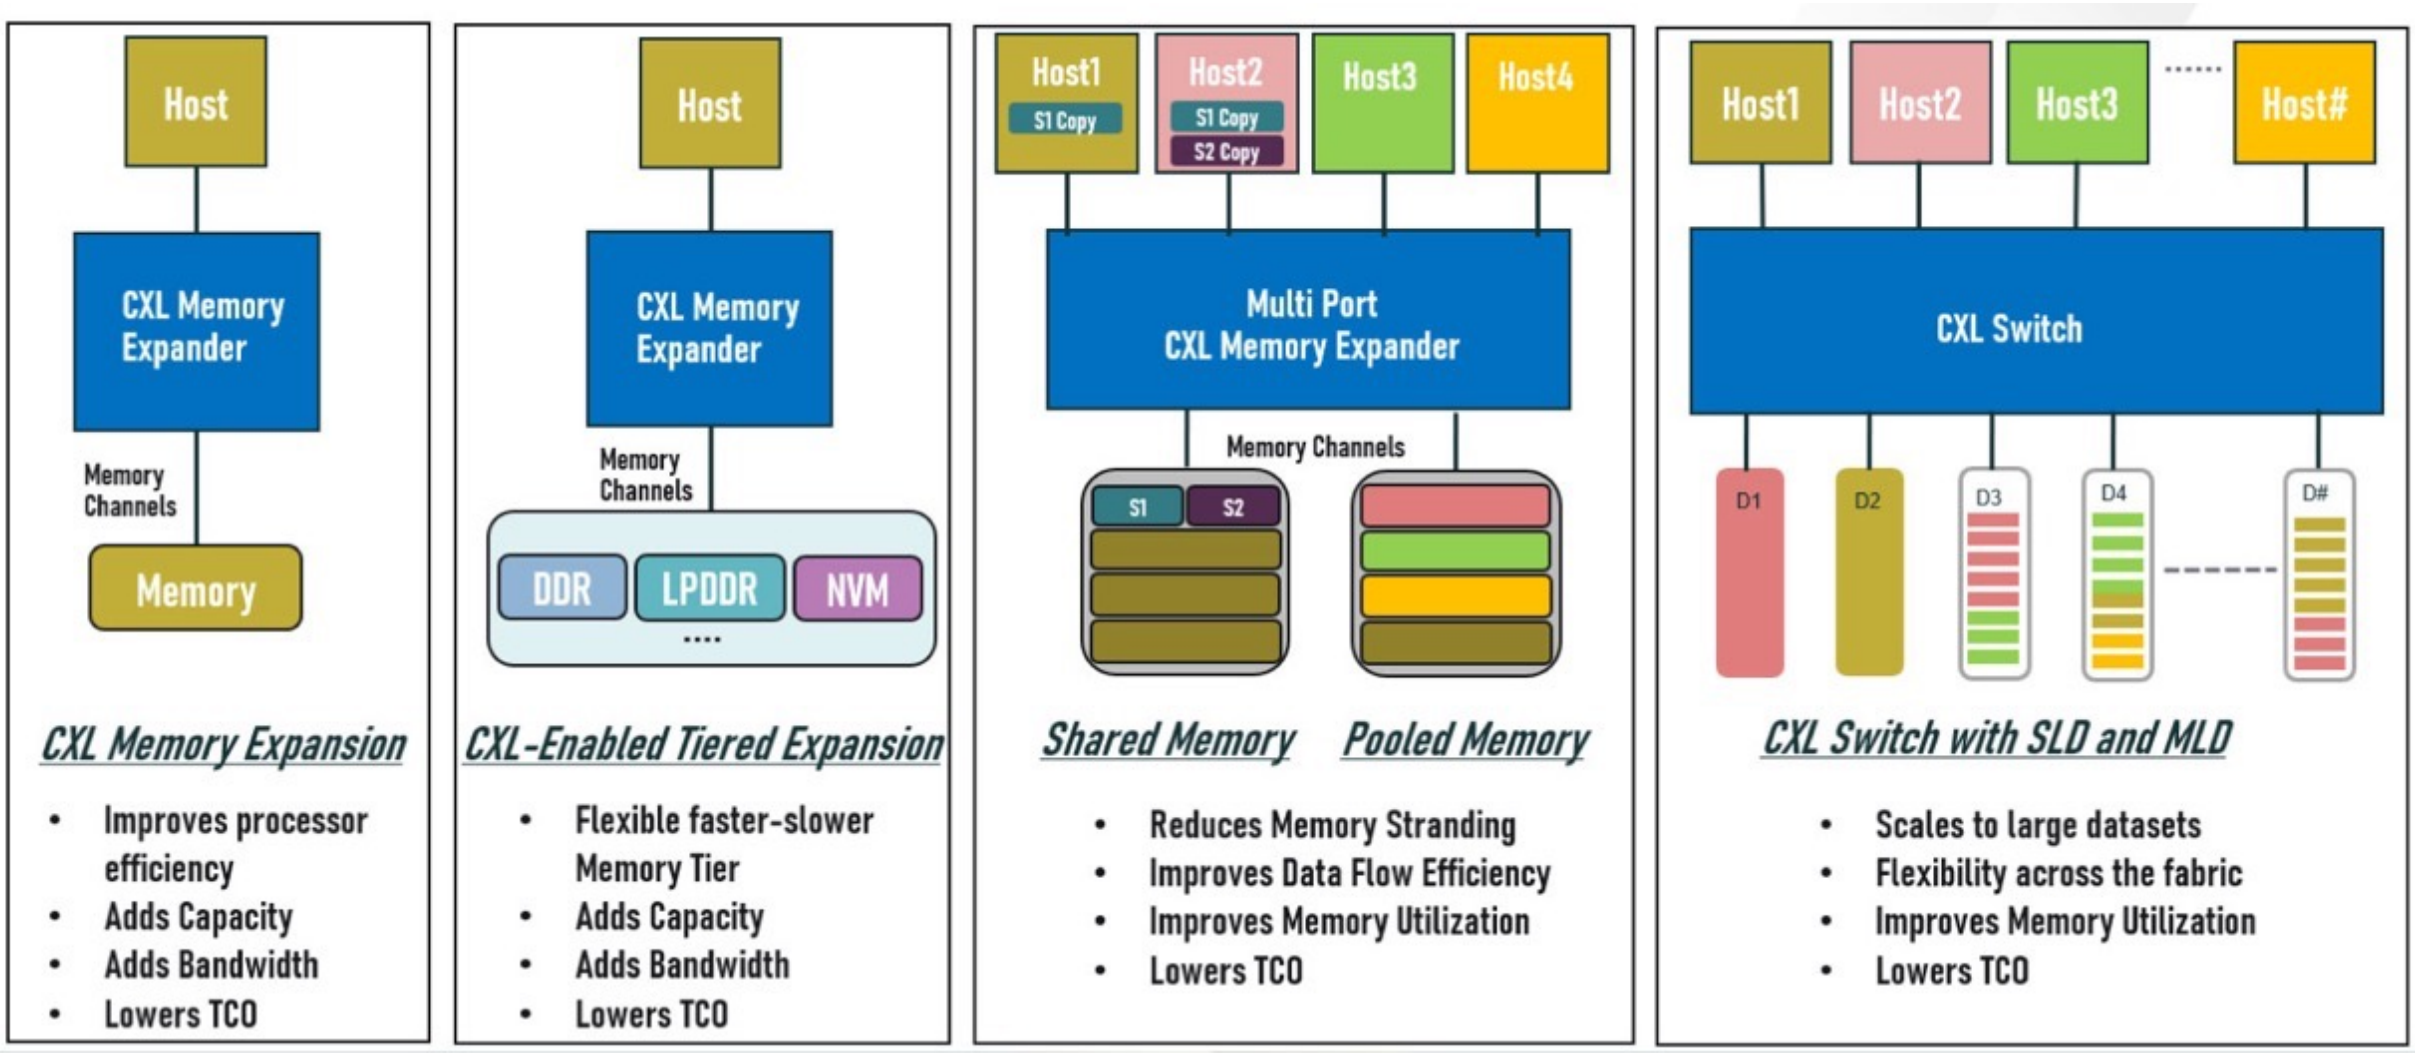

Heterogeneous Compute and Data Movement

Heterogeneous Compute Revisited

- CXL enables a more fluid and flexible memory model
- Single, common, memory address space across processors and devices

- More efficient population and update of operands
- More efficient extraction of results
- Memory resource "borrowing"
- User/Kernel level data access and data movement
- Low latency to memory, host to device and device to host

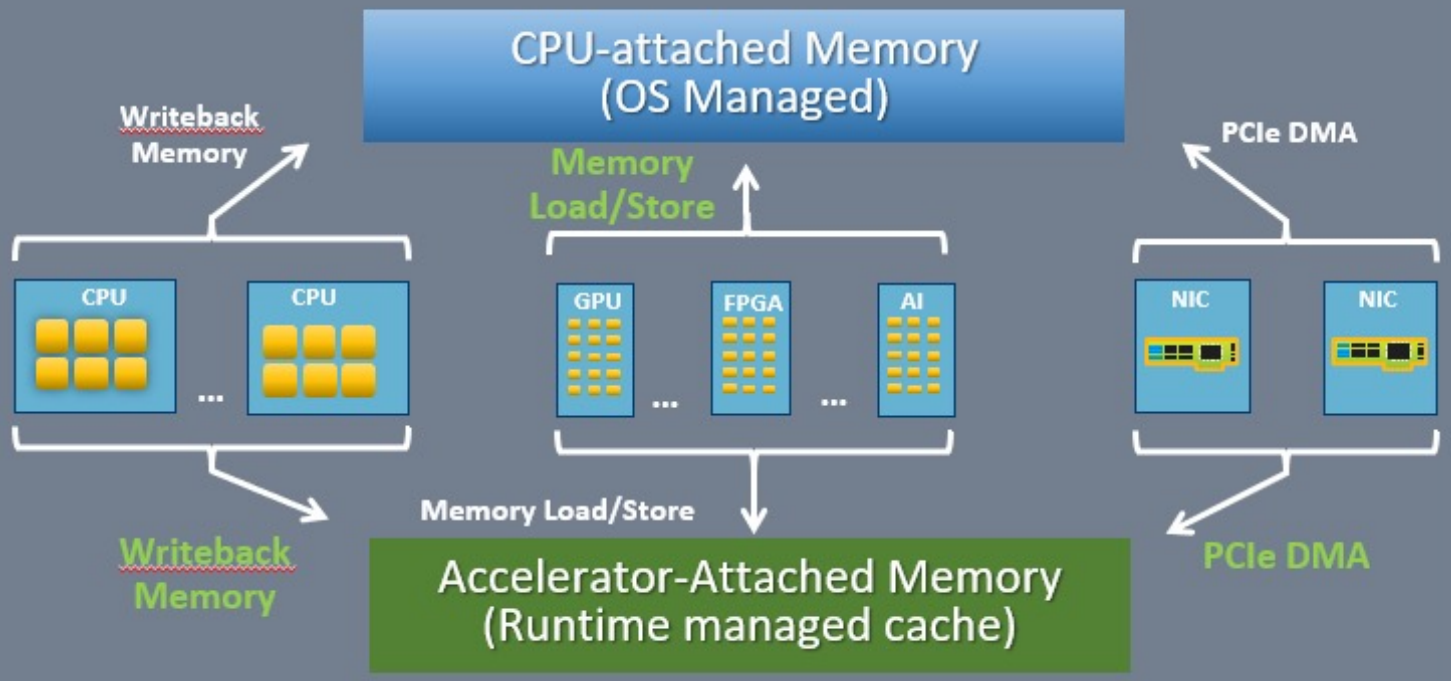

SDXI Assisted Data Movement in a CXL Architecture

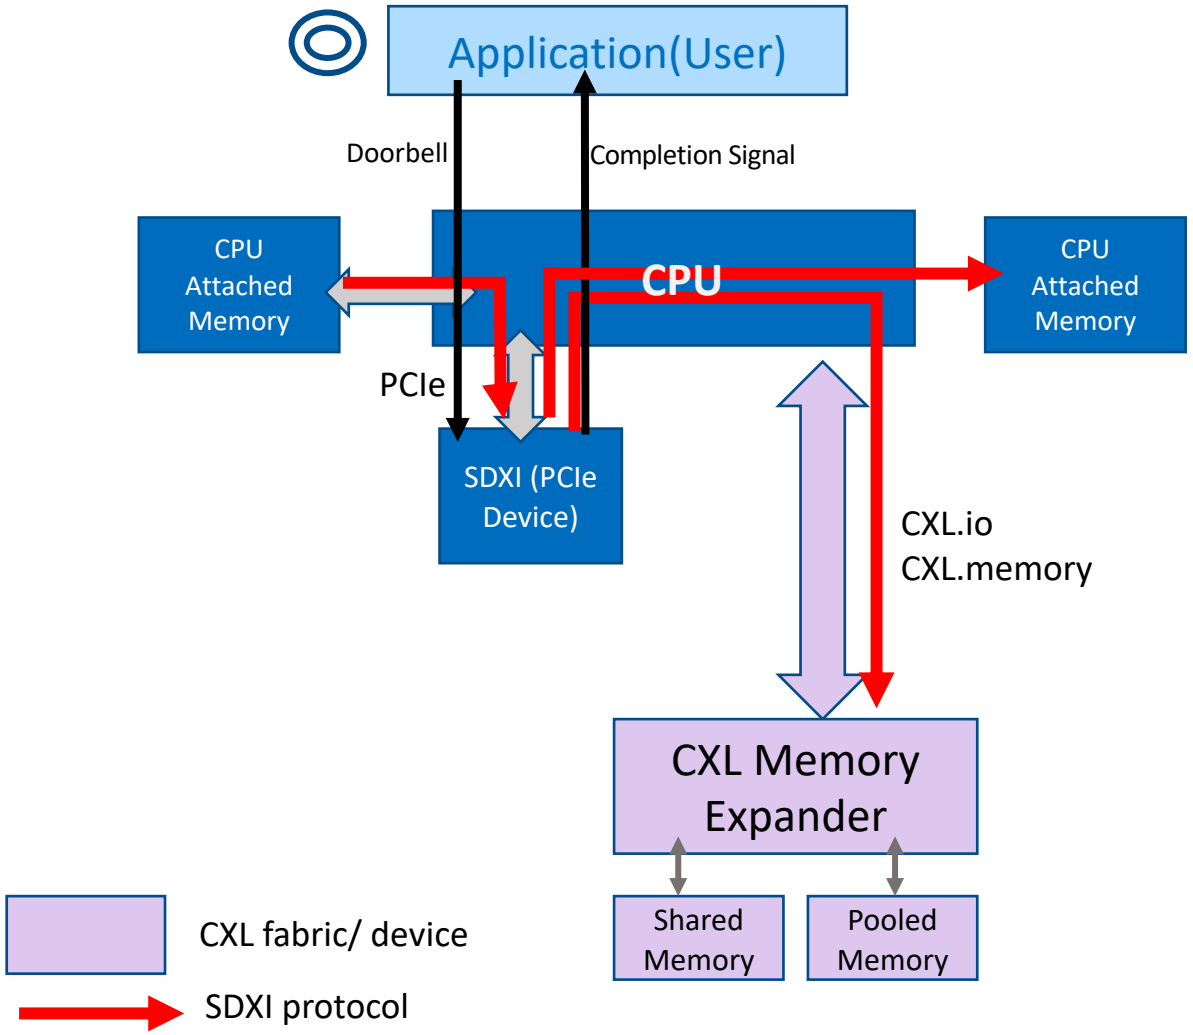

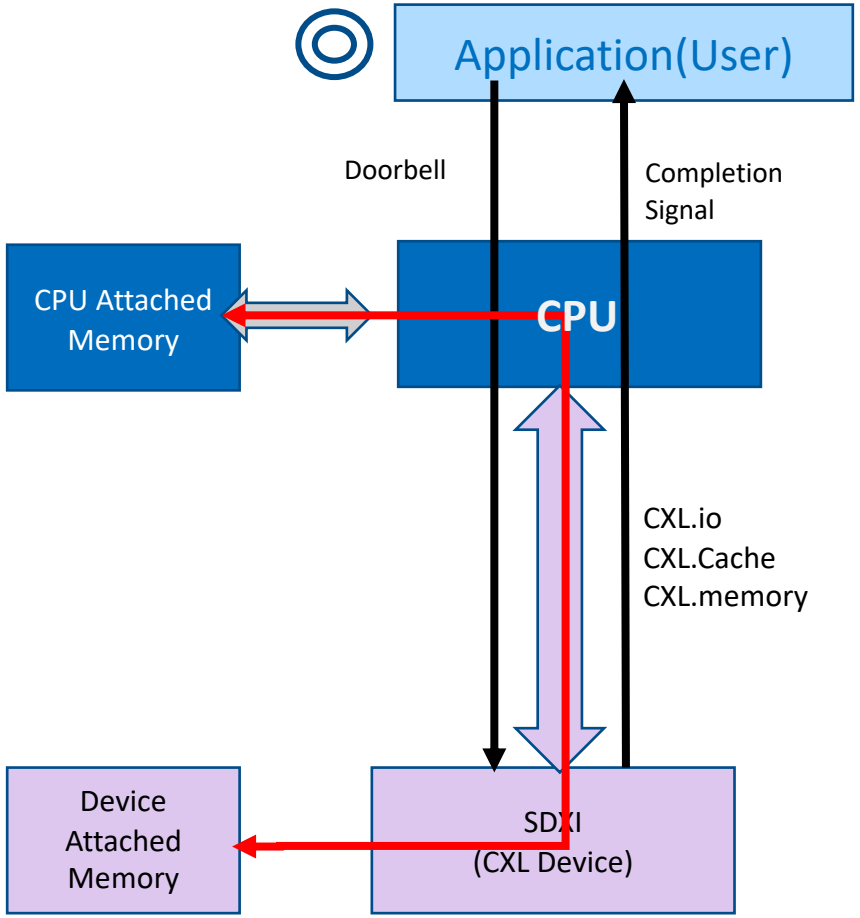

CXL and SDXI: The Path Ahead

- Memory tiering, memory borrowing with heterogeneous memories, and composable systems need efficient data movement
- CXL enables low latency, high bandwidth fabric for efficient data movement
- SDXI enables standard application interface and consistency for data mover usage models
- Community call to action: Align and drive SDXI standardized data movement in CXL-enabled system architecture

Q&A

After this Webcast

- Please rate this webcast and provide us with your feedback
- This webcast and a copy of the slides are available at the SNIA Educational Library https://www.snia.org/educational-library
- A Q&A from this webcast, including answers to questions we couldn't get to today, will be posted on our blog at https://sniansfblog.org/
- Follow us on Twitter @SNIANSF

Thank You

SNIA NSF | NETWORKING STORAGE

Backup

Call to Action

- To join the CXL Consortium, visit www.computeexpresslink.org/join
- View CXL technology demos: www.computeexpresslink.org/videos
- Download an evaluation copy of the CXL 3.0 specification
- Engage with us on social media

@ComputeExLink

www.linkedin.com/company/cxl-consortium/

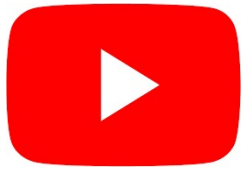

CXL Consortium Channel

CXL Technology Demonstrations

- CXL Consortium showcases first public demonstrations of CXL technology at SC'21
- View virtual and live demos from CXL Consortium members here: https://www.computeexpresslink.org/videos
  - Demos showcase CXL usages, including memory development, memory expansion and memory disaggregation

CXL Resources

Webinars:

- Webinar: Compute Express Link™ (CXL™): Supporting Persistent Memory
- Webinar: Compute Express Link™ 2.0 Specification: Memory Pooling
- Webinar: Introducing the Compute Express Link™ 2.0 Specification
- Webinar Archive: https://www.computeexpresslink.org/webinars

Blogs:

- CXL™ Consortium Member Spotlight: Elastics.cloud
- Questions from the Compute Express Link™ (CXL™): Supporting Persistent Memory Webinar
- CXL™ Consortium Member Spotlight: Synopsys
- CXL™ Consortium and Gen-Z Consortium™ MoU Update: A Path to Protocol
- Blog Archive: https://www.computeexpresslink.org/blog

White Papers:

- An Overview of Reliability, Availability, and Serviceability (RAS) in Compute Express Link™ 2.0
- Compute Express Link 2.0™ White Paper
- Introduction to Compute Express Link™

CXL 2.0 Technical Trainings

- Download videos and presentations: https://www.computeexpresslink.org/cxl-2-technical-training

