In-Memory Computing SUMMIT VIRTUAL EVENT 2020

Introducing Smart Data Acceleration Interface (SDXI)

Shyamkumar Iyer, Distinguished Member of Technical Staff
Dell Technologies

Interim Chair, SNIA SDXI TWG

10-28-2020

SNIA®

What is SNIA?

SNIA is a non-profit global organization dedicated to developing standards and education programs to advance storage and information technology.

Who is SNIA?

A community of storage professionals and technical experts

185
industry leading organizations

2,000
active contributing members

50,000
IT end users & storage pros worldwide

snia.org

@SNIA

Work Accomplished Through SNIA

Standards Development and Adoption

- Accepted and Ratified spec development process
- Submissions for International Standard ratification (ISO/IEC)
- Develop open source software to accelerate adoption

Technology Acceleration and Promotion

- Special Interest Groups to promote emerging technologies
- Multi-vendor collaboration to accelerate adoption
- Cross-Industry alliances and engagements

Global Vendor-Neutral Education

- Host worldwide storage developer conferences
- Organize storage technology summits
- Deliver vendor-neutral webcasts and technical podcasts
- Publish technology white papers, articles and blogs
- Vendor neutral plugfests, hack-a-thons, conformance and interoperability testing
- SNIA GitHub open source repositories

STORAGE DEVELOPER CONFERENCE SDC BY Developers FOR Developers

SNIA's Technical Work is in Eight Focus Areas

PERSISTENT MEMORY
- Non-Volatile Memory Programming Model
- Smart Data Accelerator Interface
- NVDIMMs

COMPUTATIONAL STORAGE
- Services and Products
  - Drives, Processors, Arrays

NETWORKED STORAGE
- Data Access Protocols
- Networking Technologies for Storage

CLOUD STORAGE TECHNOLOGIES
- Data into and out of the Cloud
- Data Orchestration

PHYSICAL STORAGE
- Connectors, Form Factors & Transceivers
- Hyperscaler Storage
- Object Drives
- Solid State Storage

POWER EFFICIENCY MEASUREMENT
- SNIA Emerald™ Power Efficiency

STORAGE MANAGEMENT
- Device and Environment Management
- Next Generation Storage Management

DATA GOVERNANCE & SECURITY
- Privacy and Data Protection Regulations
- Storage Security
- Integrity, Protection, Retention
- Blockchain Storage

Agenda

The problem and the need for a solution

Introducing SDXI

The problem and the need for a solution

Trends

- Core counts increasing to enable Compute scaling
- Compute density is on the rise
- Converged and Hyperconverged Storage appliances are enabling new workloads on server class systems
  - Data locality is important
- Single threaded performance is under pressure.
- I/O intensive workloads can take away compute CPU cycles available.
- Network and Storage workloads can take compute cycles
- Data Movement, Encryption, Decryption, Compression

Need for Accelerated Intra-host Data Movement

Each intra-host exchange can comprise multiple memory buffer copies (or transformations)

- Generally implemented with layers of software stacks:
- Kernel-to-I/O can leverage I/O-specific hardware memory copy
- But, SW-to-SW usually relies on per-core synchronous software (CPU-only) memory copies

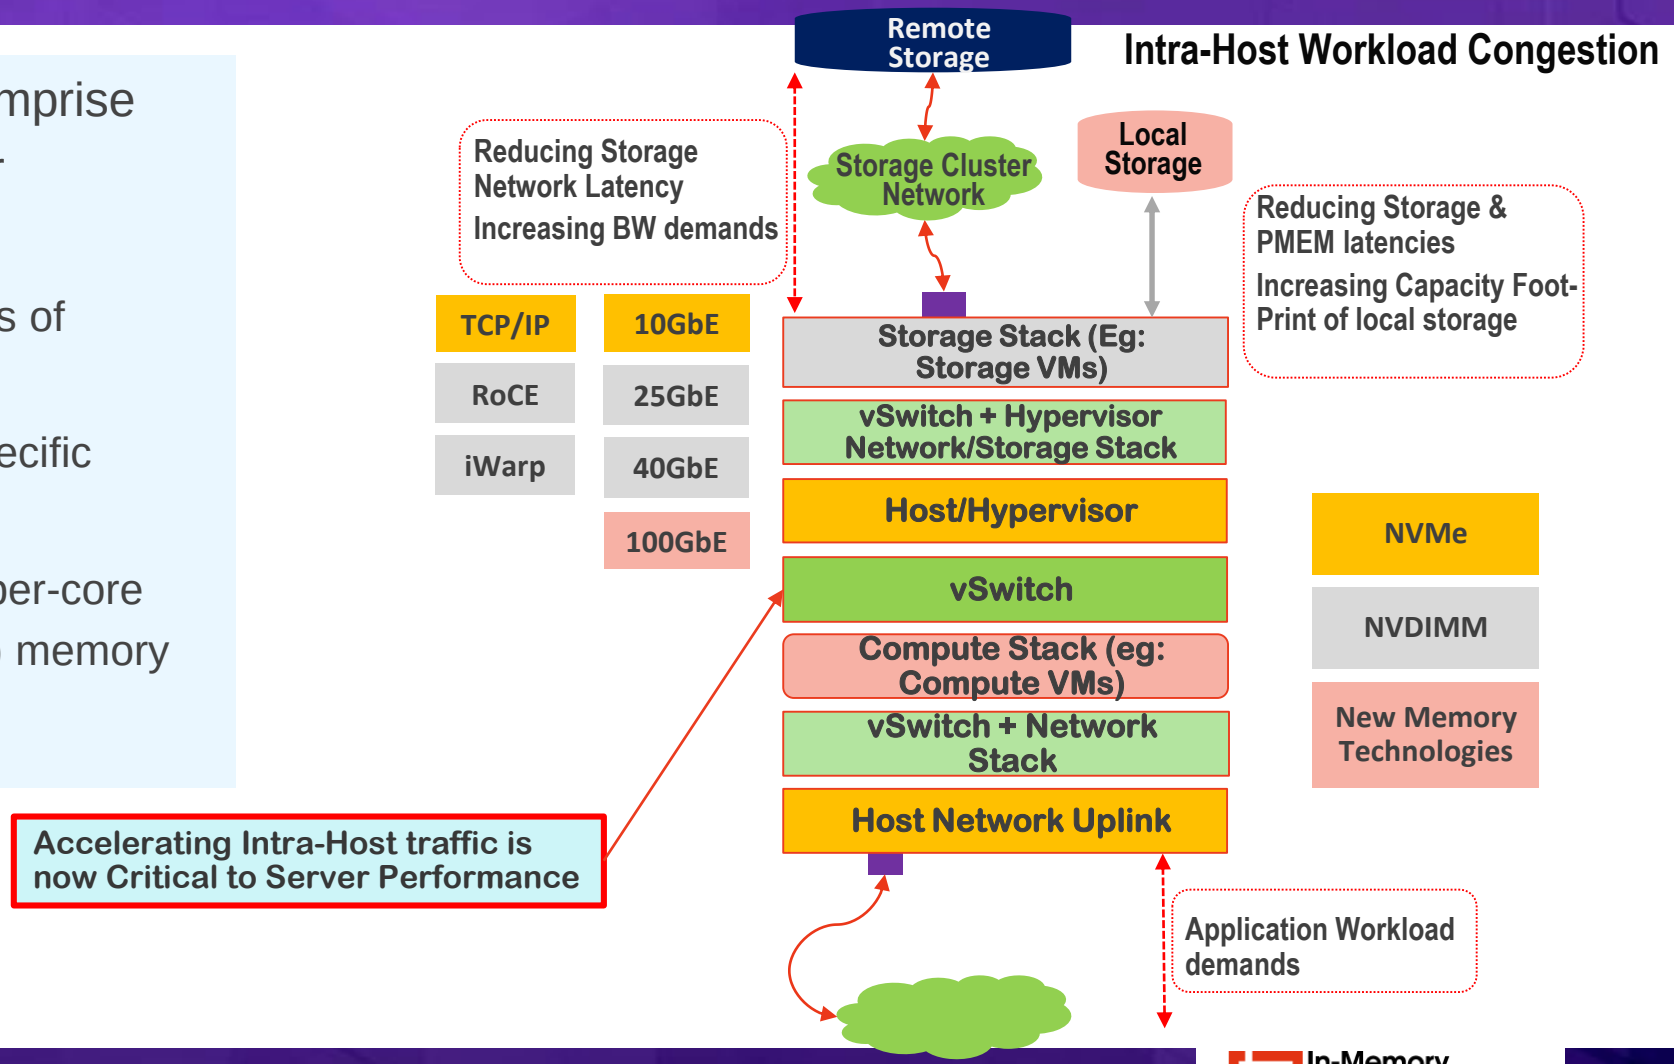

Slide Has Animation

Current data movement standard:

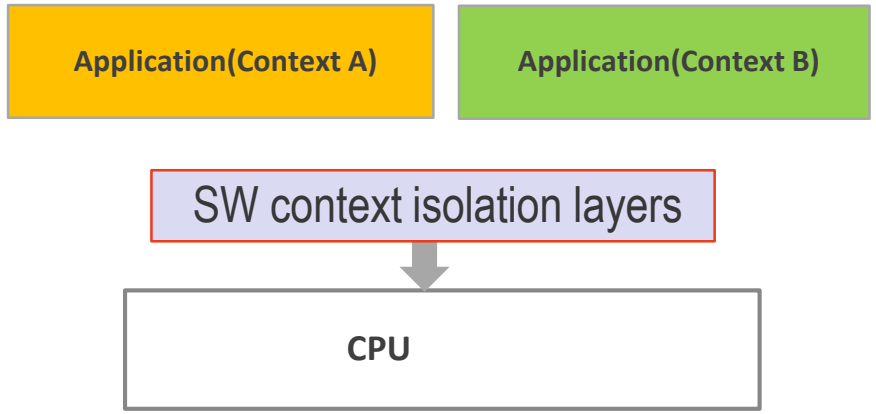

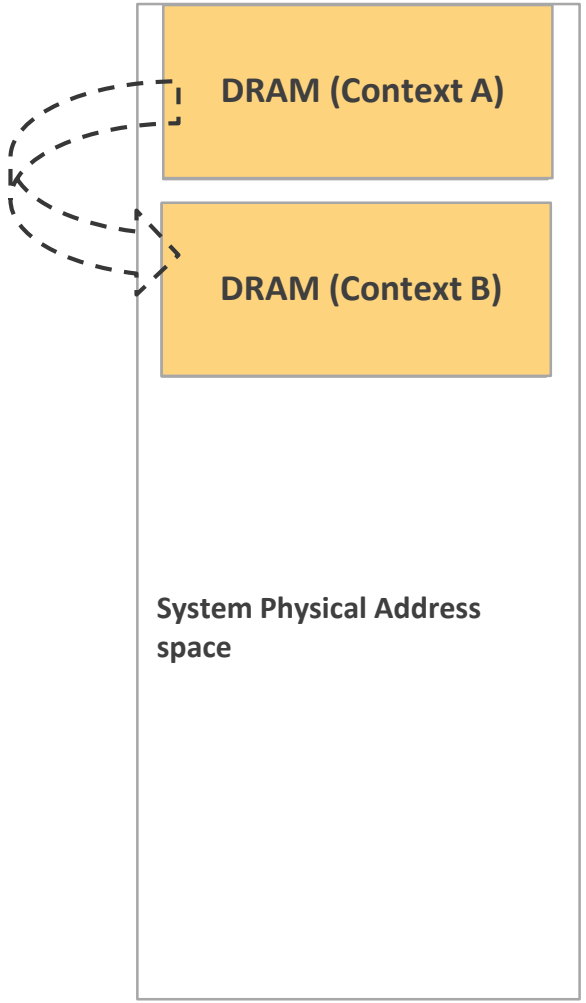

Stable CPU ISA for SW based memory copies

- Takes away from application performance
- Software overhead to provide context isolation
- Synchronous SW copies stall applications
- Less portable to different ISAs(Instruction Set Architectures)
- Finely tuned CPU data movement algorithms can break with new microarchitectures

Offload DMA engines: A new concept ?

- Fast DMA offload engines are -
  - Vendor-specific HW
  - Vendor specific drivers, APIs
  - Vendor specific work submission/completion models
  - Direct access by user level software is difficult
  - Limited Usage Models
  - Vendor specific DMA states – Makes it harder to abstract/virtualize and migrate the work to other hosts

Solution Requirements

1. Need to offload I/O from Compute CPU cycles
2. Need Architectural Stability
3. Enable Application/VM acceleration but,
    - Help migration from existing SW Stacks
4. Create abstractions in Control Path for scale and management
5. Enable performance in data path with offloads

Emerging Server & Storage Architectures

Looking into the horizon …

1. Memory-centric architectures.
2. New memory interconnects.
    a. CXL
    b. Gen-Z
3. Varied memory types.
4. Heterogenous architectures are becoming main stream.
5. The need to democratize data movement.

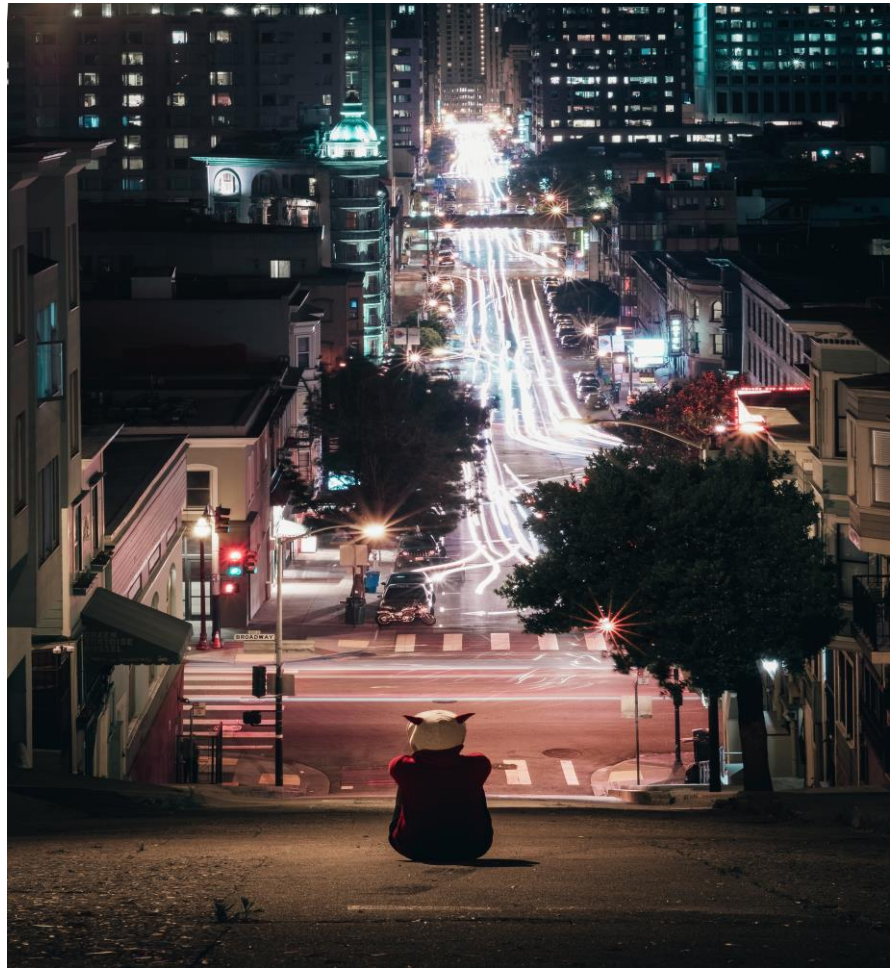

Emerging Needs: New Memory Architectures

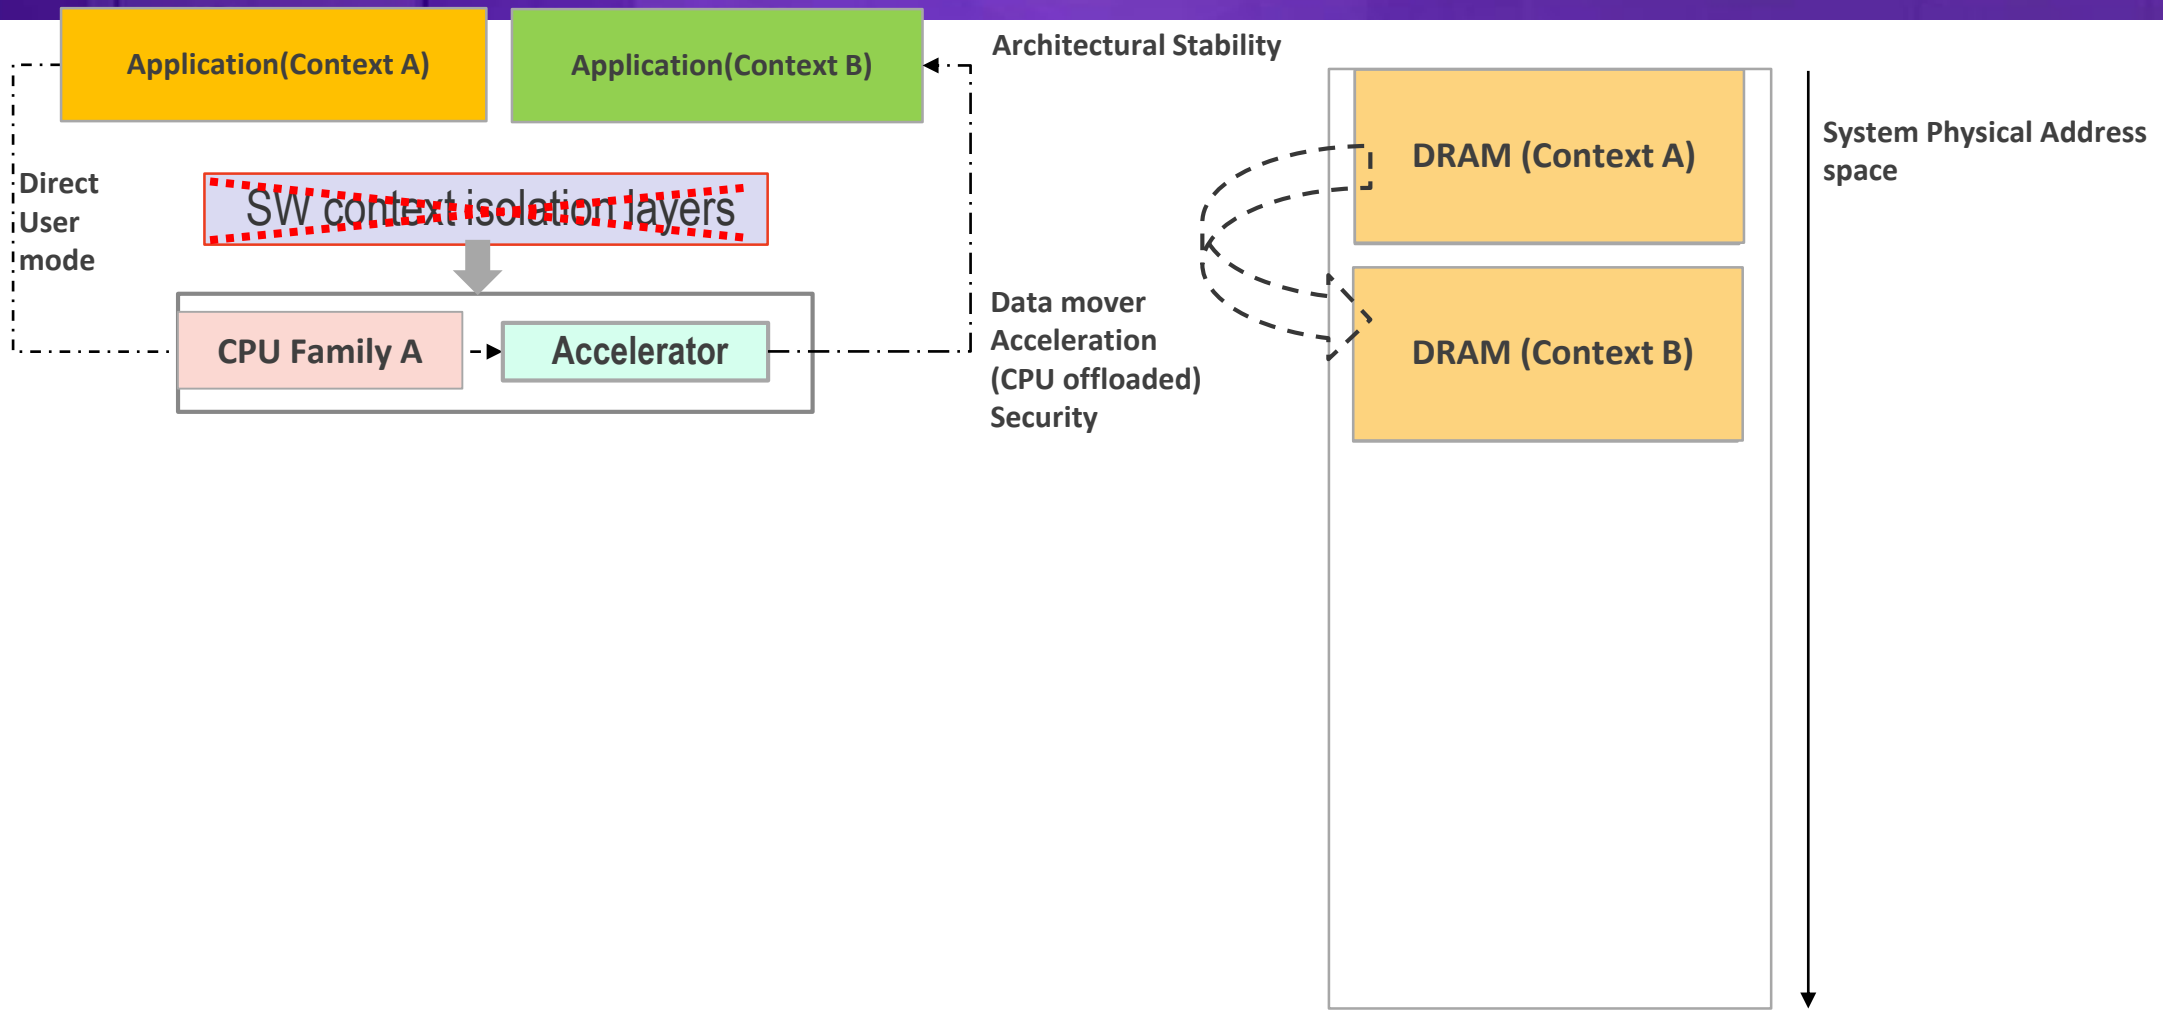

Emerging Needs: New Memory Architectures

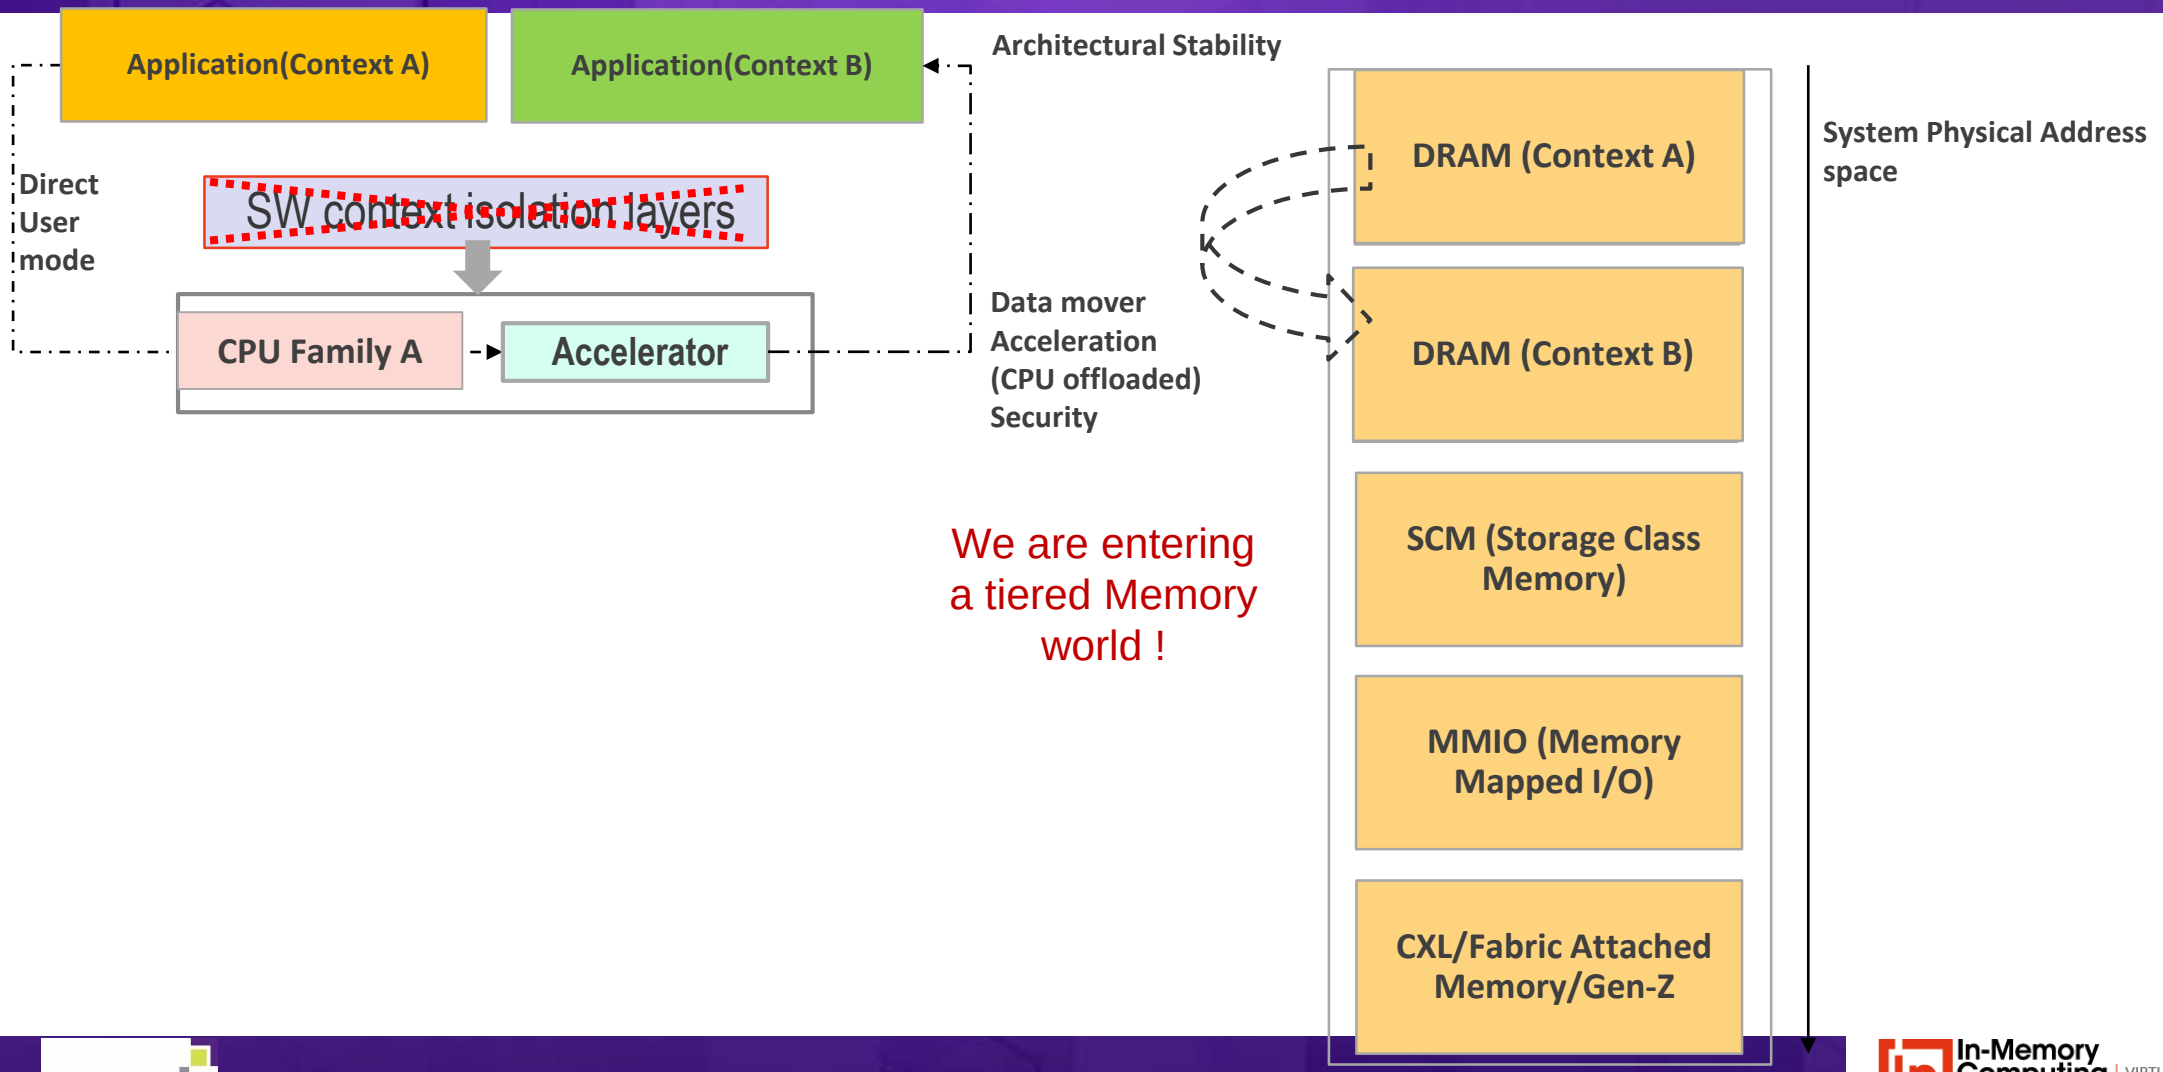

We are entering a tiered Memory world !

Architectural Stability

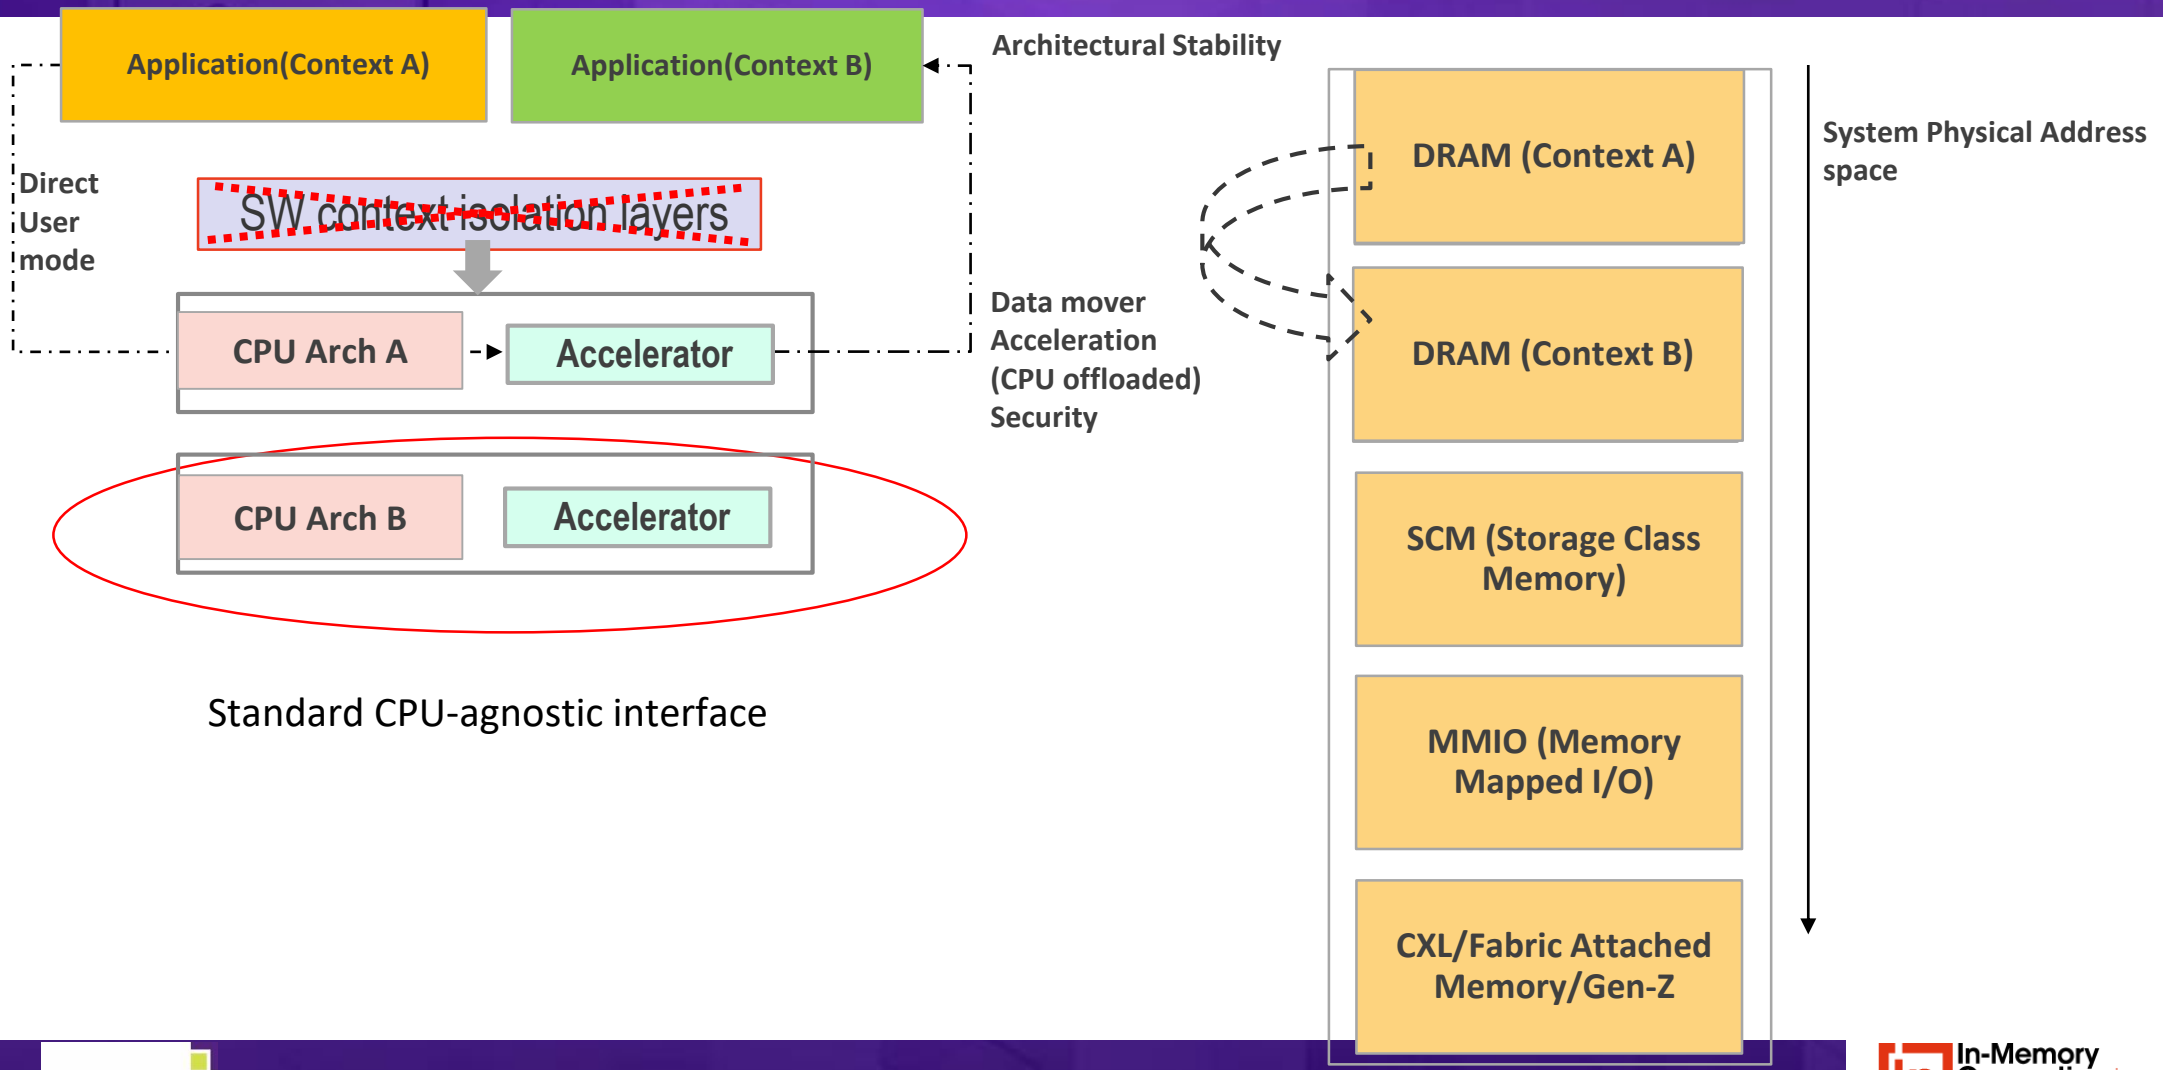

Standard CPU-agnostic interface

Enabling Accelerators

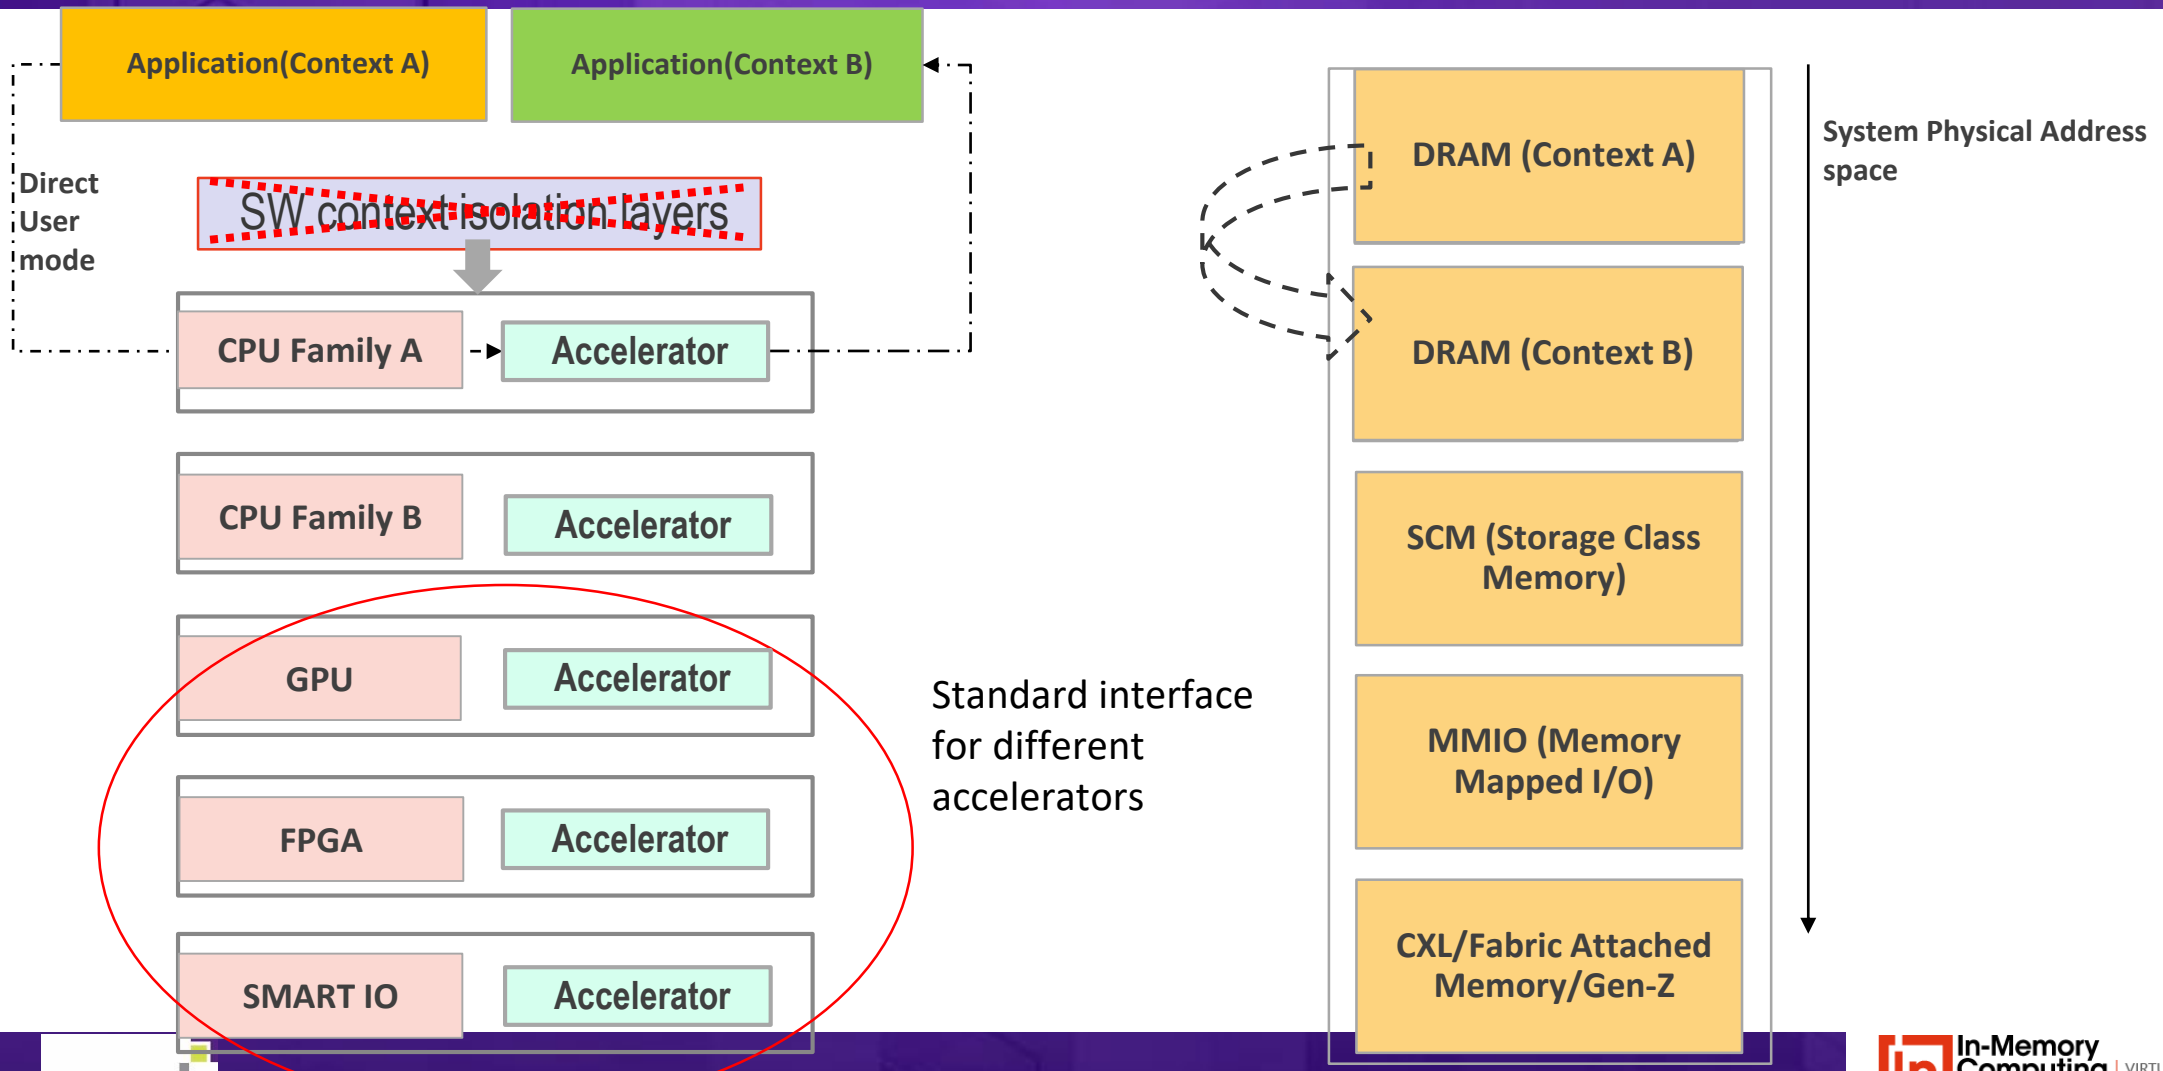

The need for an industry standard

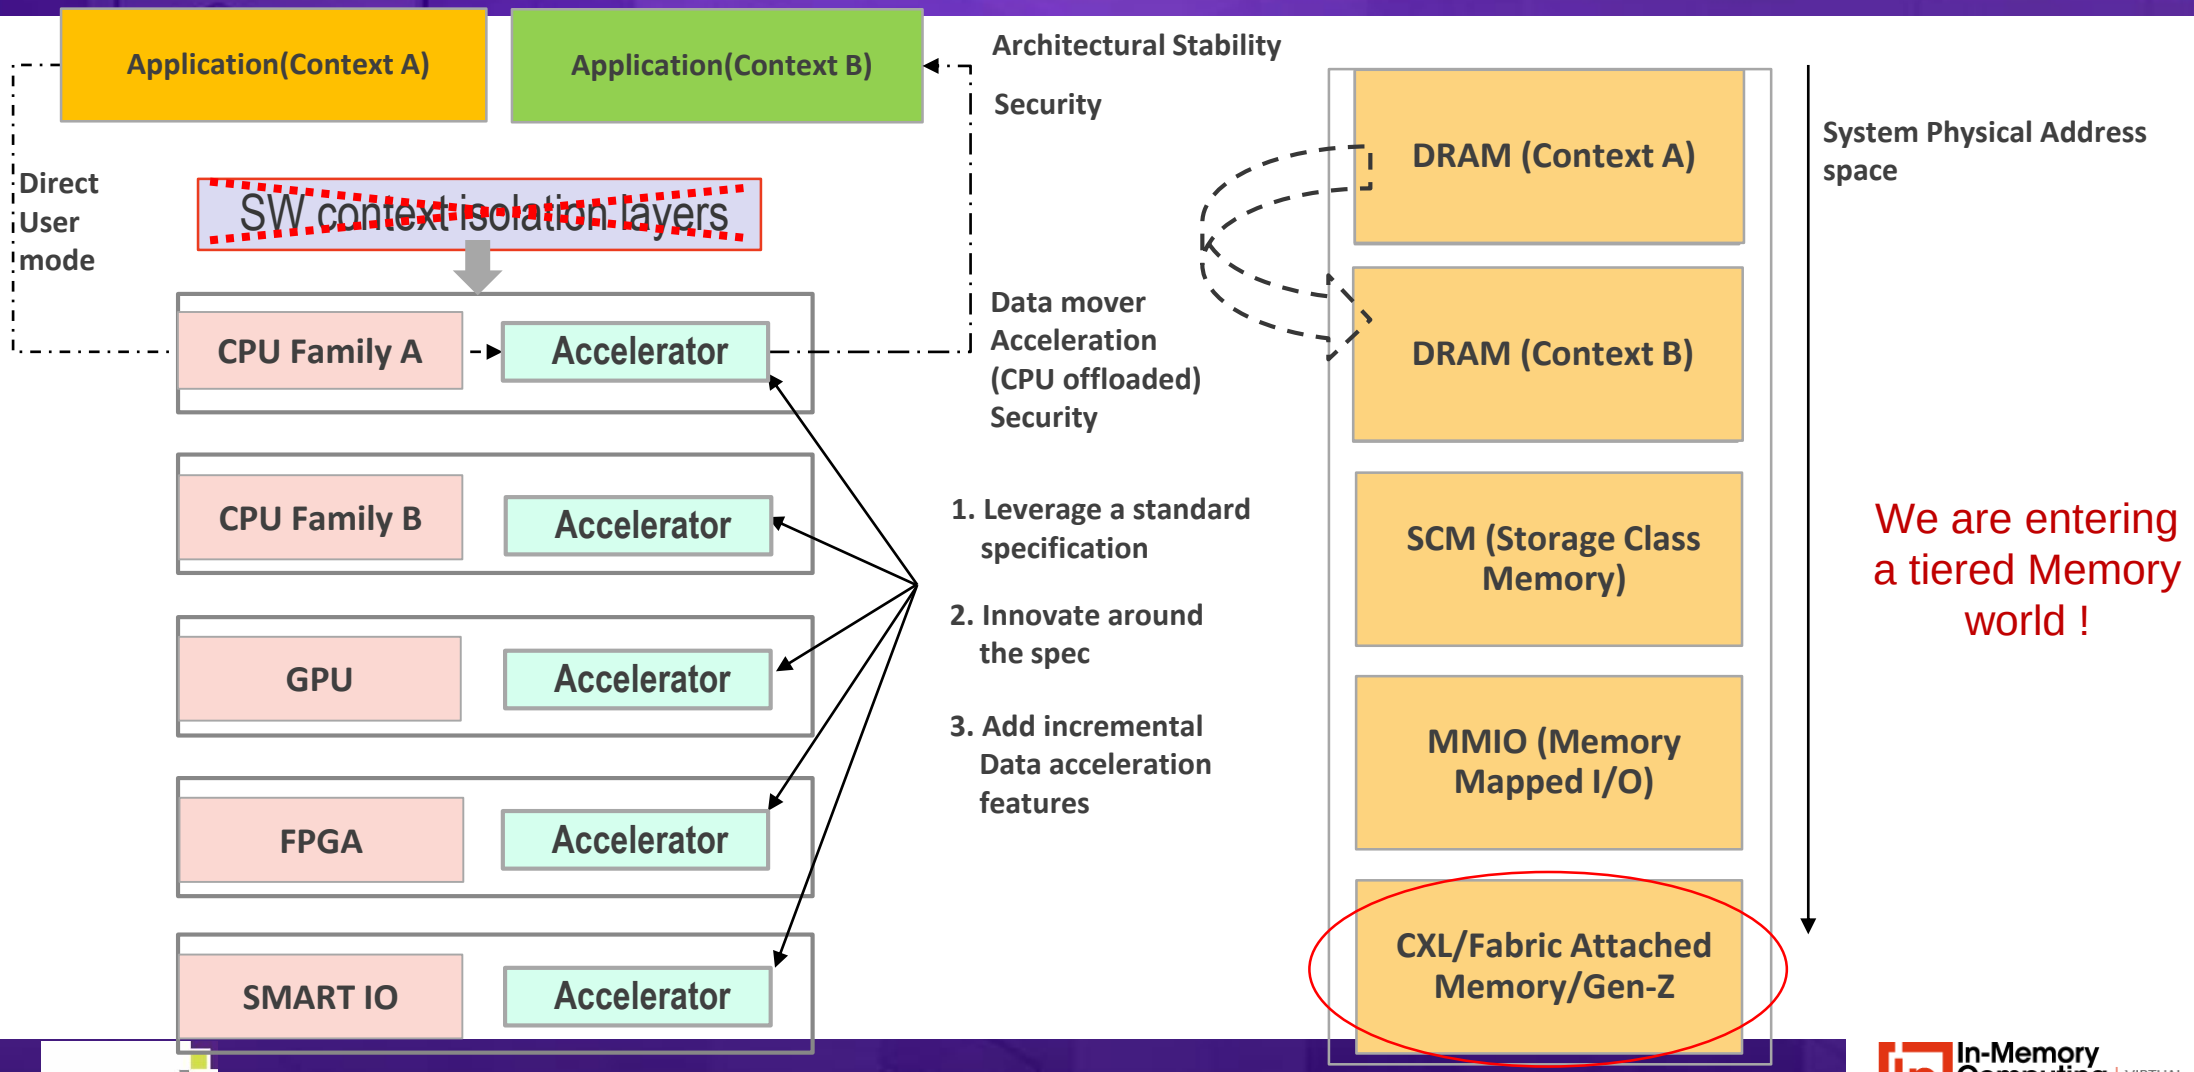

We are entering a tiered Memory world !

Agenda

The problem and the need for a solution

Introducing SDXI

Introducing SNIA SDXI

Introducing SNIA SDXI TWG

SDXI Charter

- Develop and Standardize a Memory to Memory Data Movement and Acceleration interface that is –
  - Extensible
  - Forward-compatible
  - Independent of I/O interconnect technology
- Dell, AMD, VMware contributed the starting point for the spec
- 13 TWG member companies and growing…

Design Tenets

- Data movement between different address spaces.
  - Includes user address spaces, different virtual machines
- Data movement without mediation by privileged software.
  - Once a connection has been established.
- Allows abstraction or virtualization by privileged software.
- Capability to quiesce, suspend, and resume the architectural state of a per-address-space data mover.
  - Enable “live” workload or virtual machine migration between servers.
- Enables forwards and backwards compatibility across future specification revisions.
  - Interoperability between software and hardware
- Incorporate additional offloads in the future leveraging the architectural interface.
- Concurrent DMA model.

Baremetal Stack View

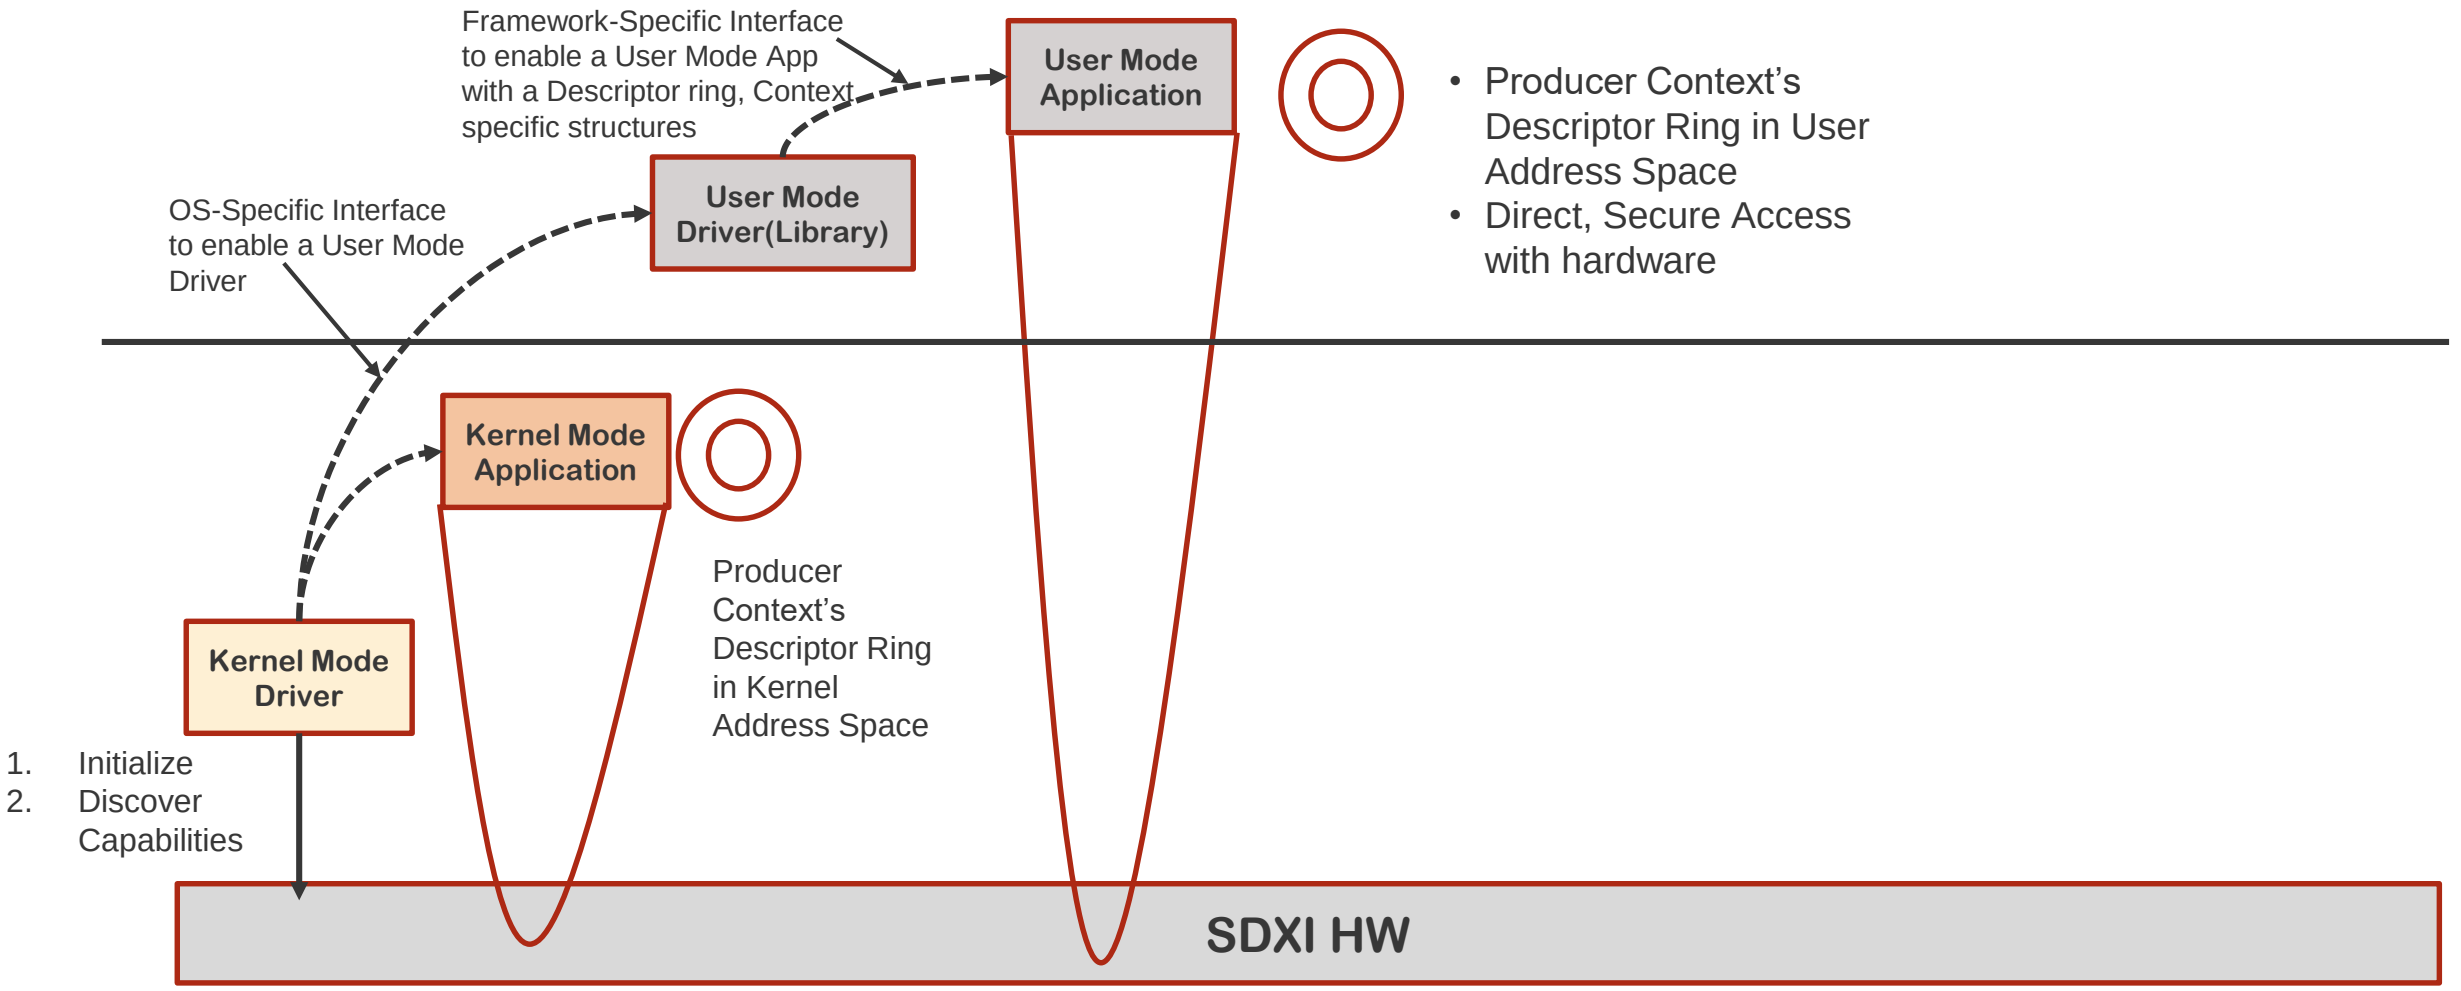

- Producer Context's Descriptor Ring in User Address Space
- Direct, Secure Access with hardware

Direct HW access, Tier across Memory Tiers

DRAM	PMEM	MMIO	Fabric Mem

Source and Destination Memory Targets for Data transfer in System Physical Address Space

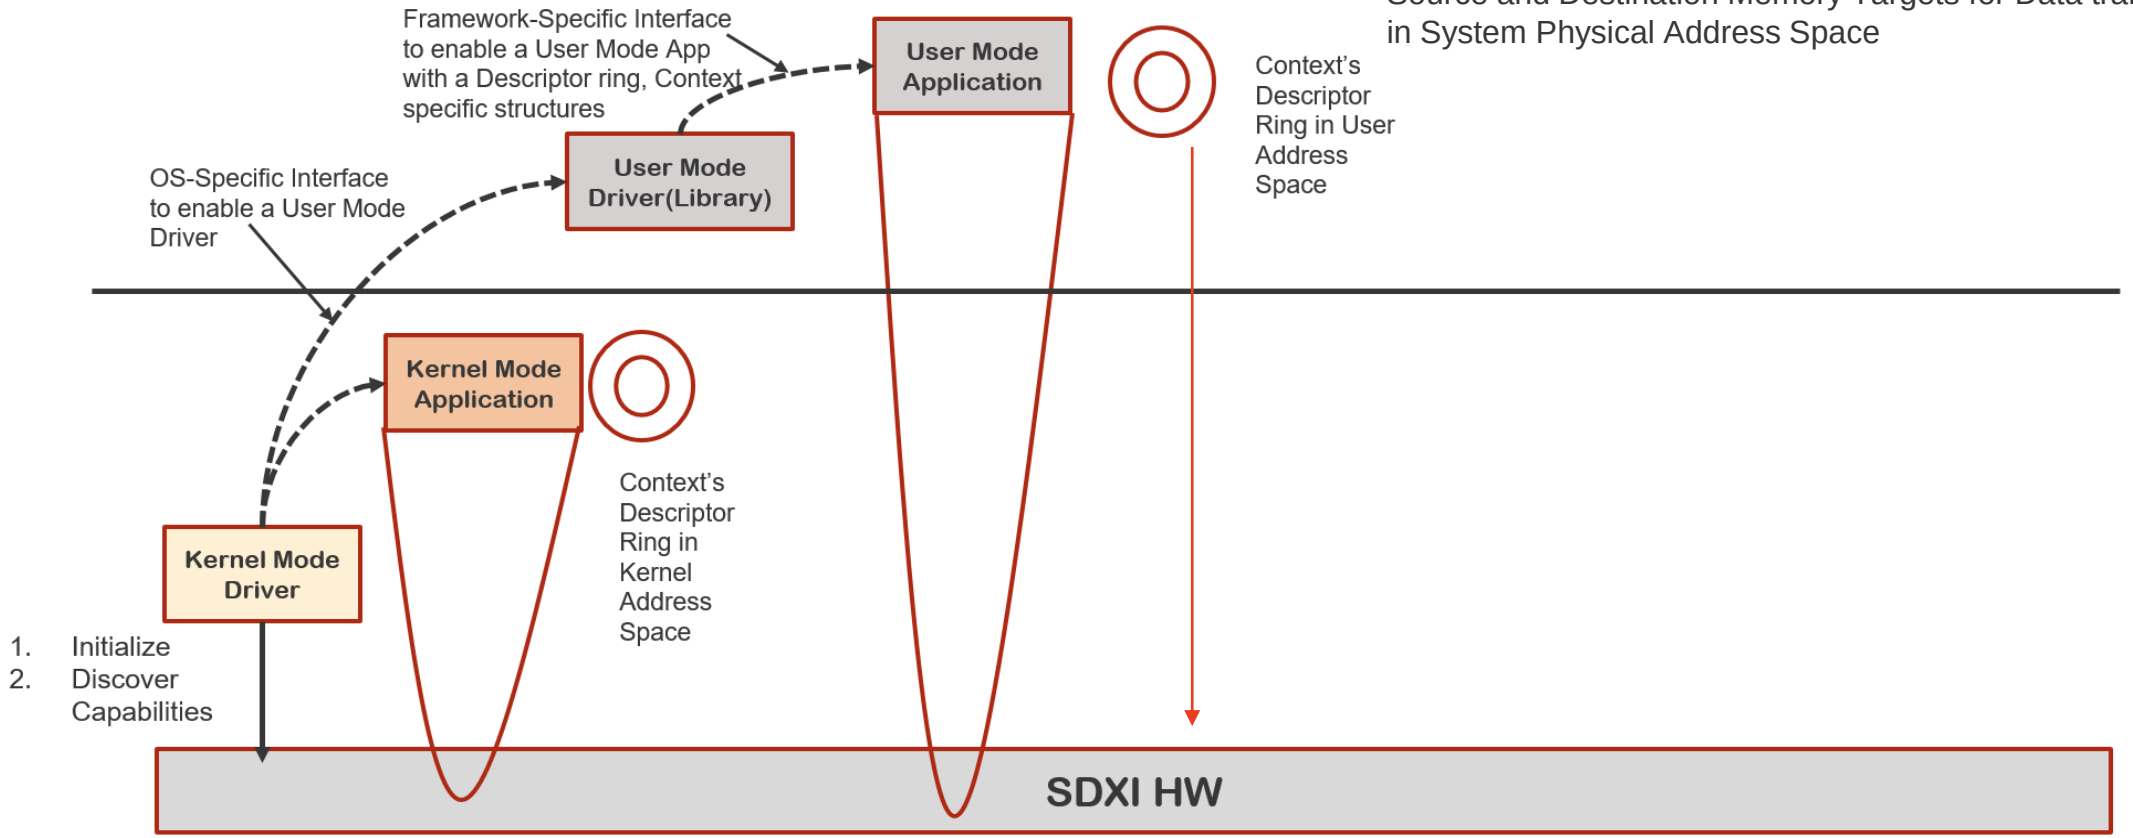

Scale Baremetal Apps – Multi-Address Space

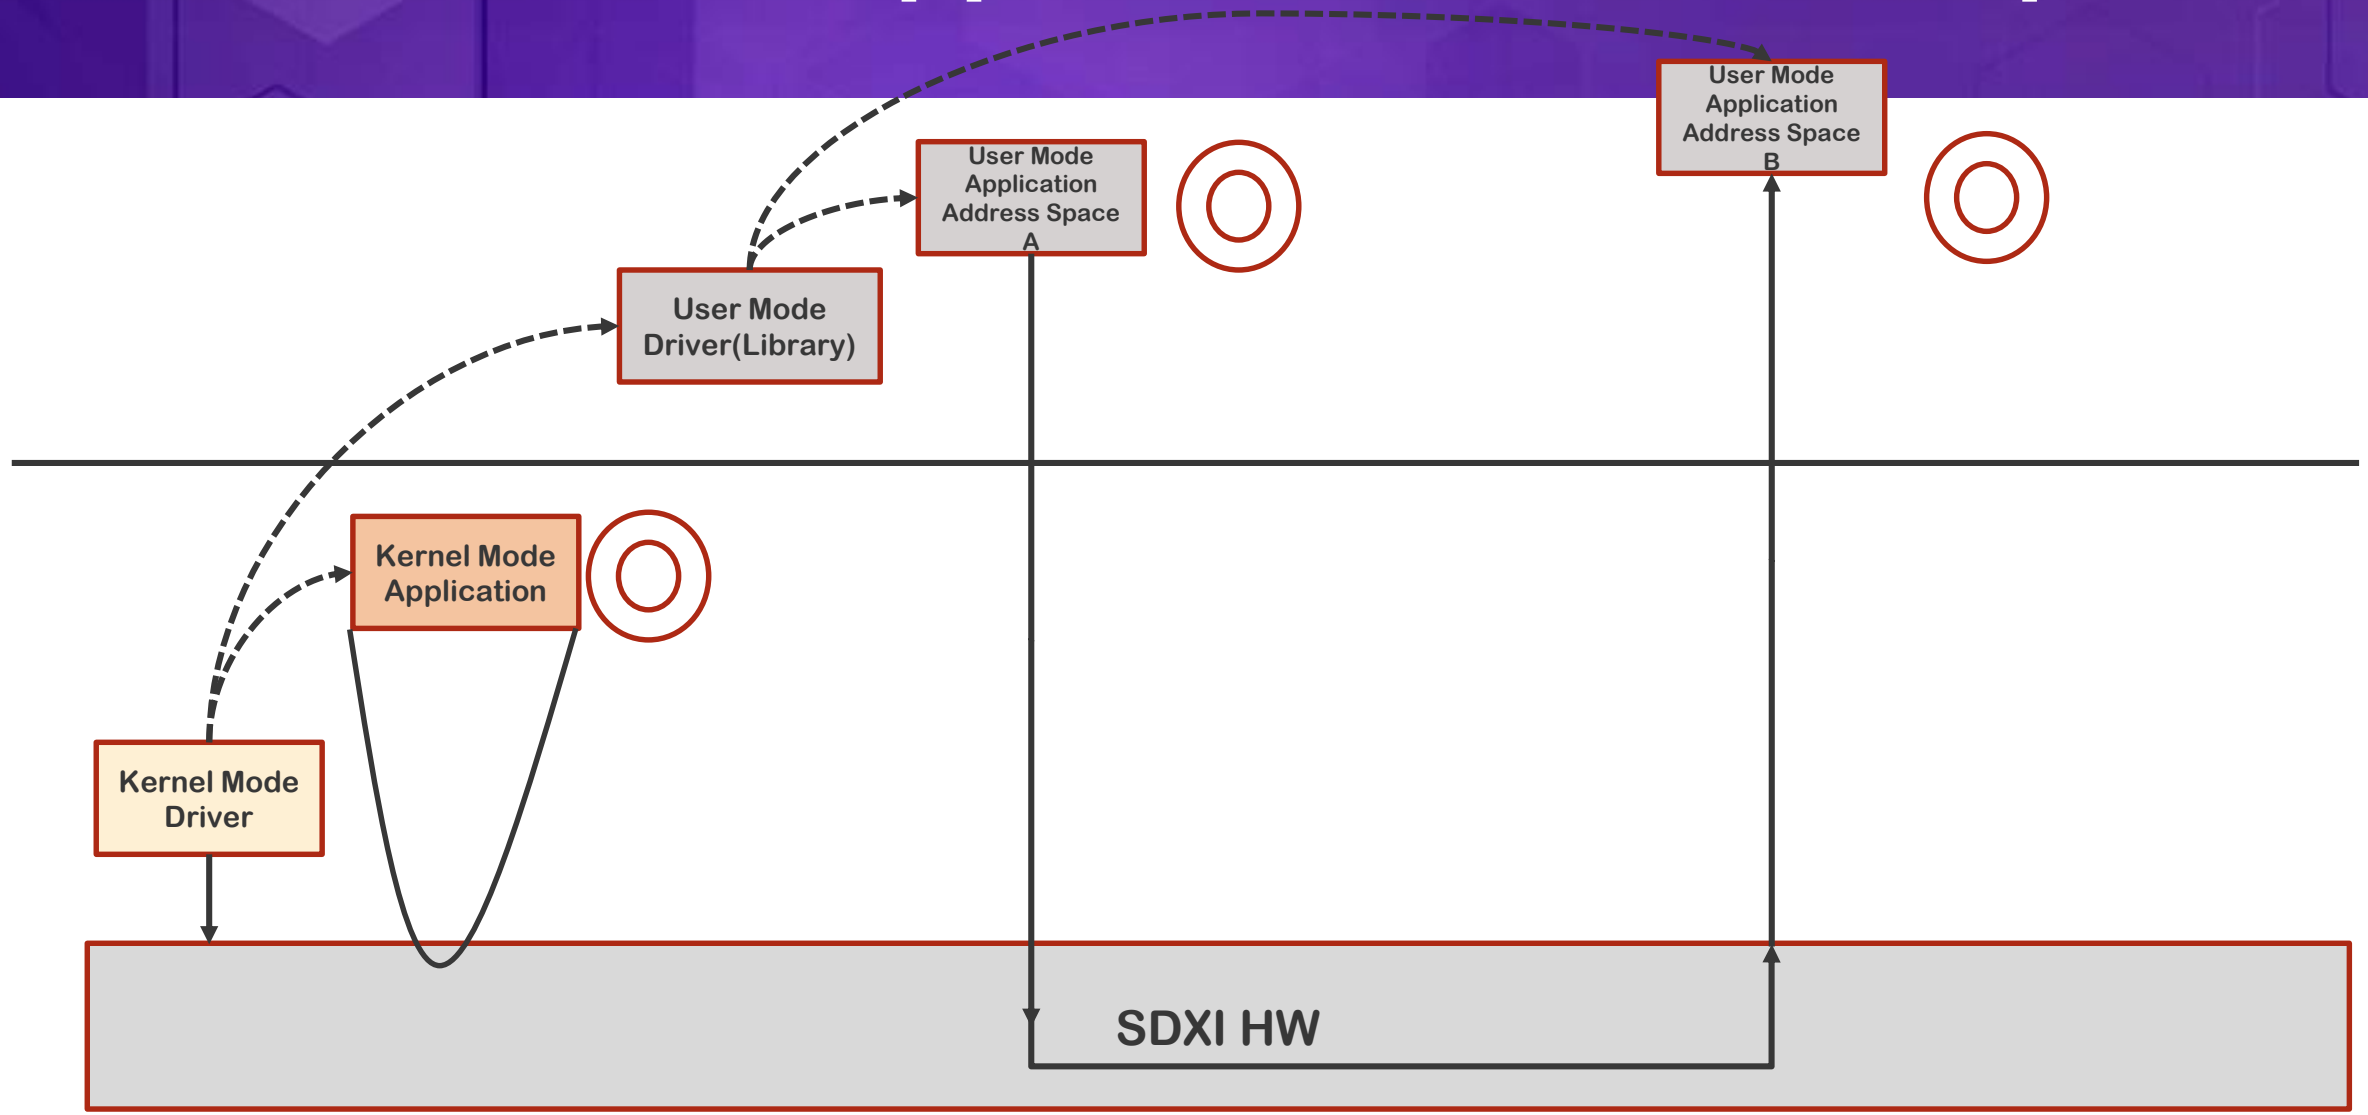

Scale with Compute Virtualization– Multi-VM address space

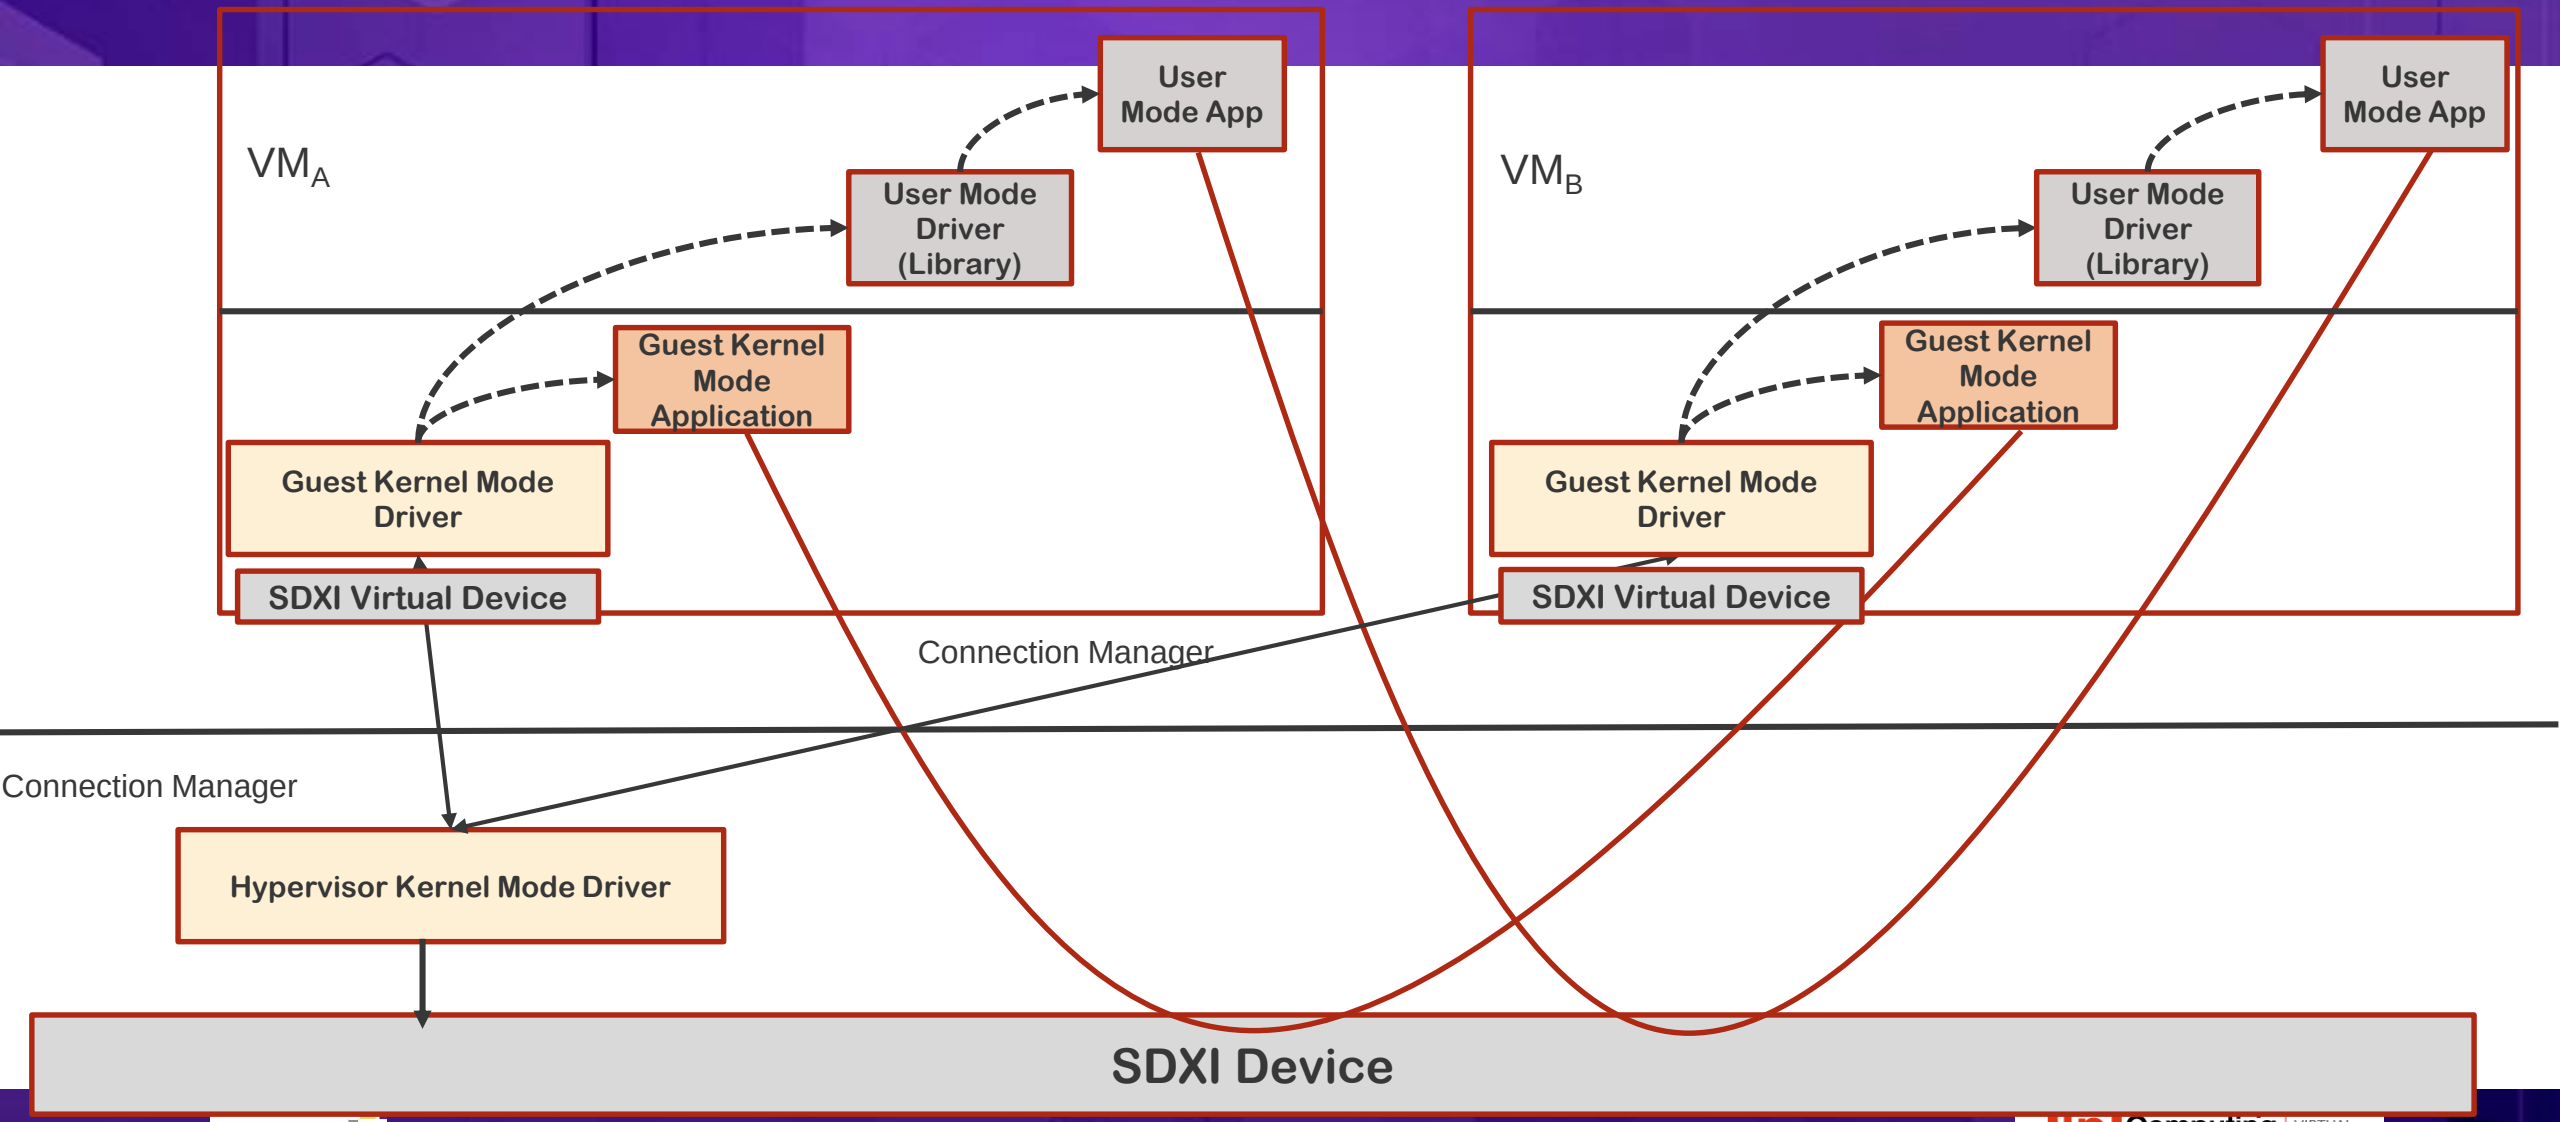

SDXI TWG's Program of Work

Advance and Standardize initial spec contribution to a v1.0 SNIA architecture standard.

Post v1.0 Focus

- 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
- Connection Management architecture for data movers

Encourage adopting companies to work towards compliant software implementations and driver models.

Educate and encourage adoption by OS, Hypervisors, OEMs, Applications and Data Acceleration vendors

Links

1. How to get more involved ?
   - https://www.snia.org/sdxi
2. Need more details ?
   - SDC 2020 Conference
   - https://www.youtube.com/watch?v=iv2GUfnxG-A
3. Questions ?
   - Linkedin - https://www.linkedin.com/in/shyam-iyer-51300ab/
   - Twitter - @kumar_iyer