Accelerators: SDXI, DPUs and storage

Shyam Iyer

SDXI TWG Chair

SNIA Technical Council

Distinguished Engineer, Dell Technologies

sdxitwgchair@snia.org

Legacy Compute, Memory, IO Bubbles

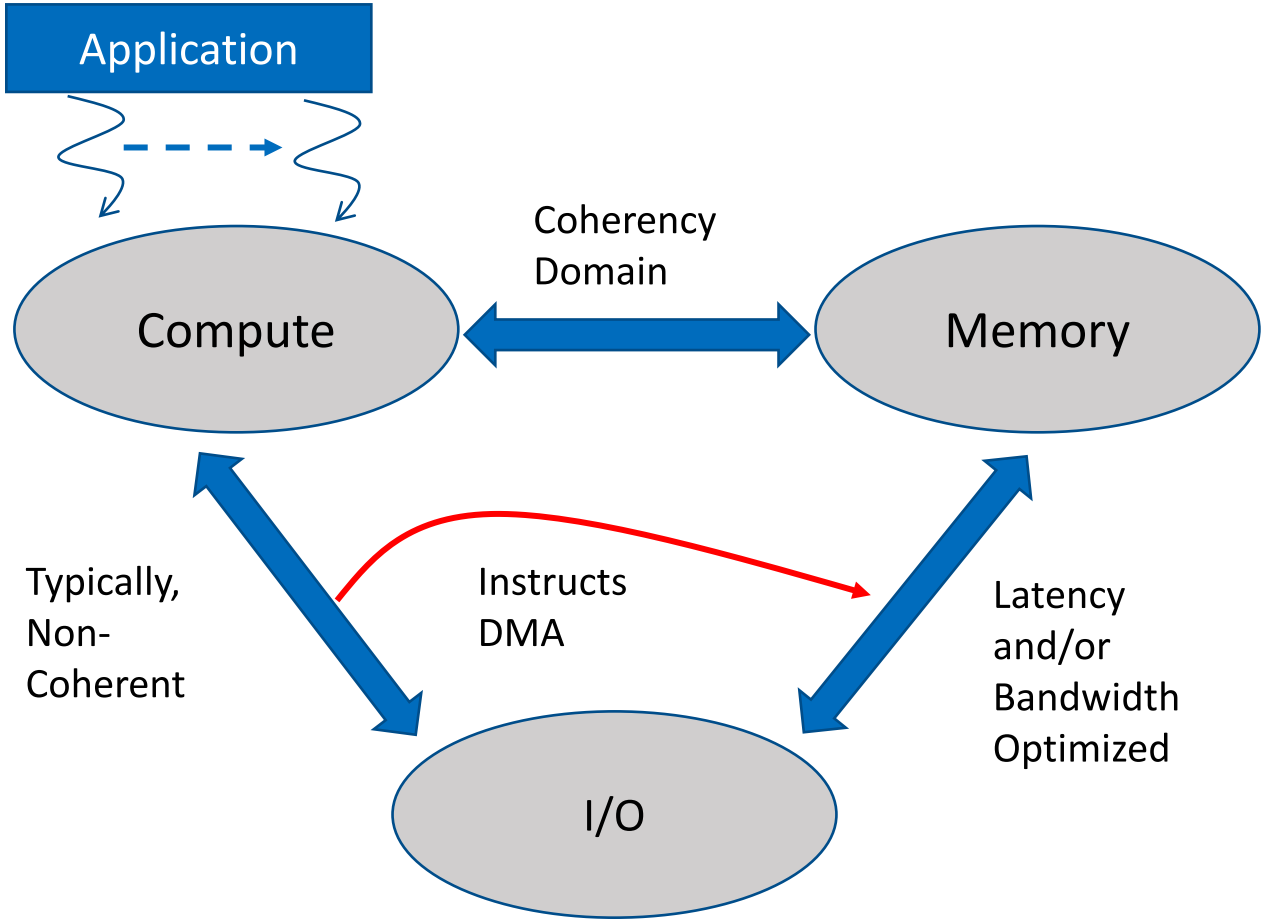

Data Storage
Data Transport

Emerging Bubbles

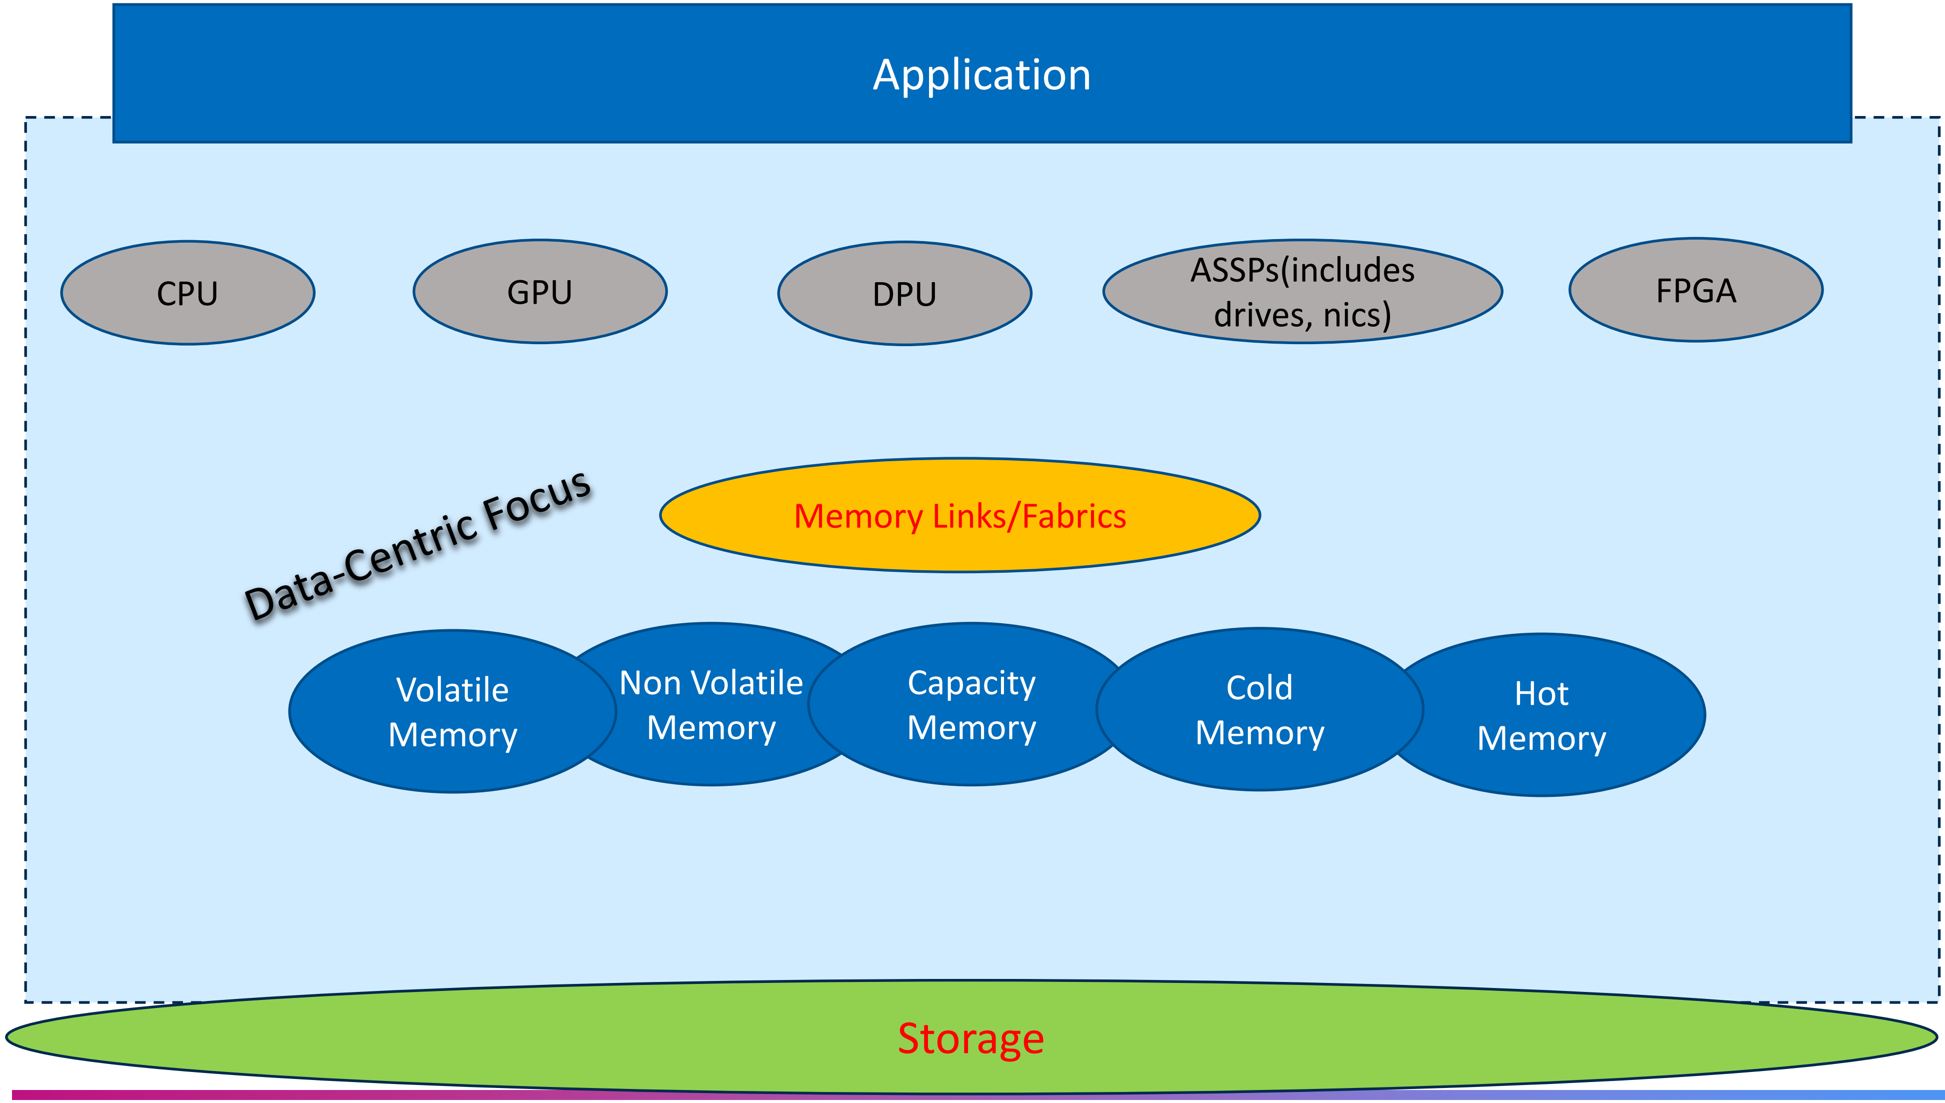

Shared Design constraints

- Latency
- Bandwidth
- Coherency
- Control

SNIA Data Focus Areas

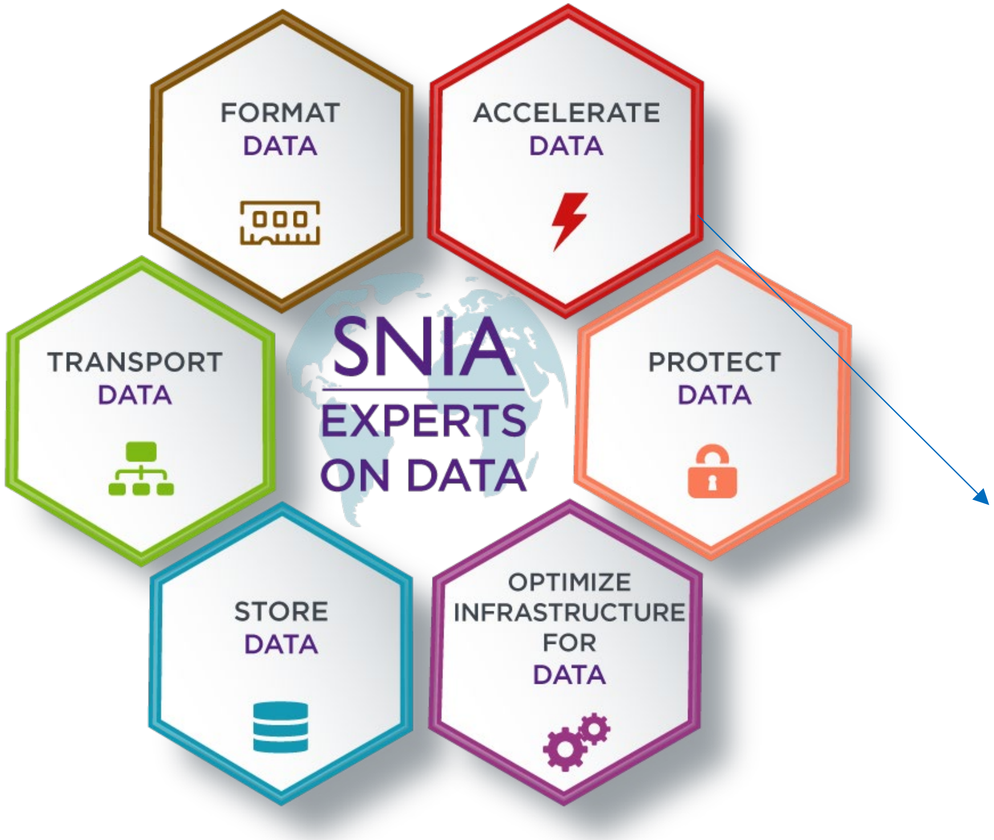

Accelerate: Move processing to the data

- Technologies that move processing closer to the data, enabling improvements in application performance and/or infrastructure efficiency through the integration of compute resources (outside of the traditional compute & memory architecture) either directly with storage or between the host and the storage.

Areas of Interest:

- SNIA SDXI (Smart Data Accelerator Interface)
- SNIA Computational Storage Architecture and Programming Model, SNIA Computational Storage API
- DPU (Data Processing Unit)

SDXI Memory-to-Memory Data Movement

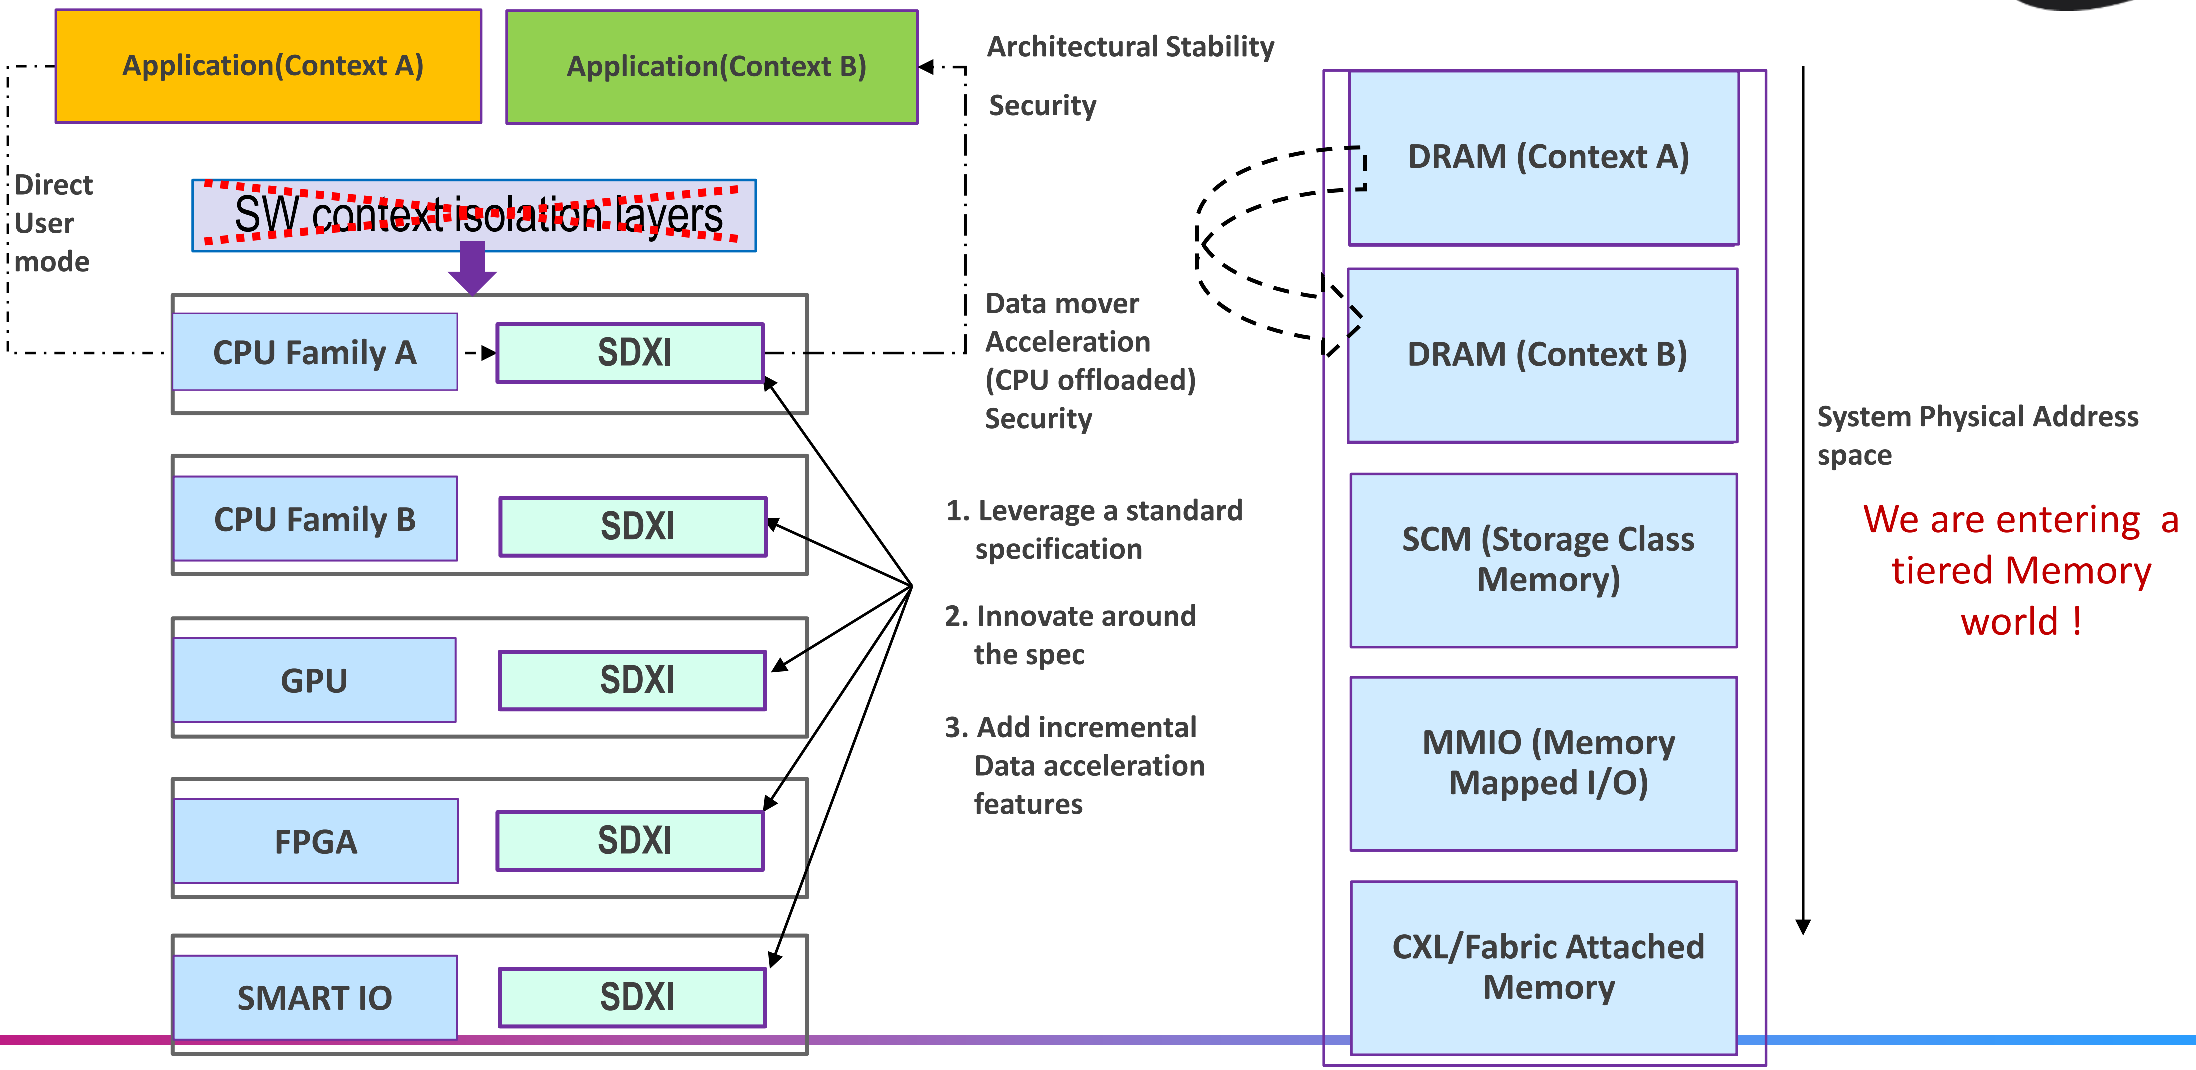

Application(Context A)

Application(Context B)

Architectural Stability

Security

Direct
User
mode

~~SW context isolation layers~~

CPU Family A — SDXI

CPU Family B — SDXI

GPU — SDXI

FPGA — SDXI

SMART IO — SDXI

Data mover
Acceleration
(CPU offloaded)
Security

1. Leverage a standard specification
2. Innovate around the spec
3. Add incremental Data acceleration features

DRAM (Context A)

DRAM (Context B)

SCM (Storage Class Memory)

MMIO (Memory Mapped I/O)

CXL/Fabric Attached Memory

System Physical Address space

We are entering a tiered Memory world !

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
- **v1.0 released!**
  - https://www.snia.org/sdxi

SDXI v1.0 Specification Contributors

AMD

arm

DELL Technologies

Hewlett Packard Enterprise

IBM

FUJITSU

MemVerge

Micron

Microsoft

NetApp

SAMSUNG

SK hynix

vmware

Western Digital

SDXI Internals

- SNIA SDXI Specification v1.0 Internals
  - https://www.youtube.com/watch?v=wjc4ZnCQibw&pp=ygUNc2RjIDIwMjMgc2R4aQ%3D%3D

Use Cases

Sample accelerator usage models

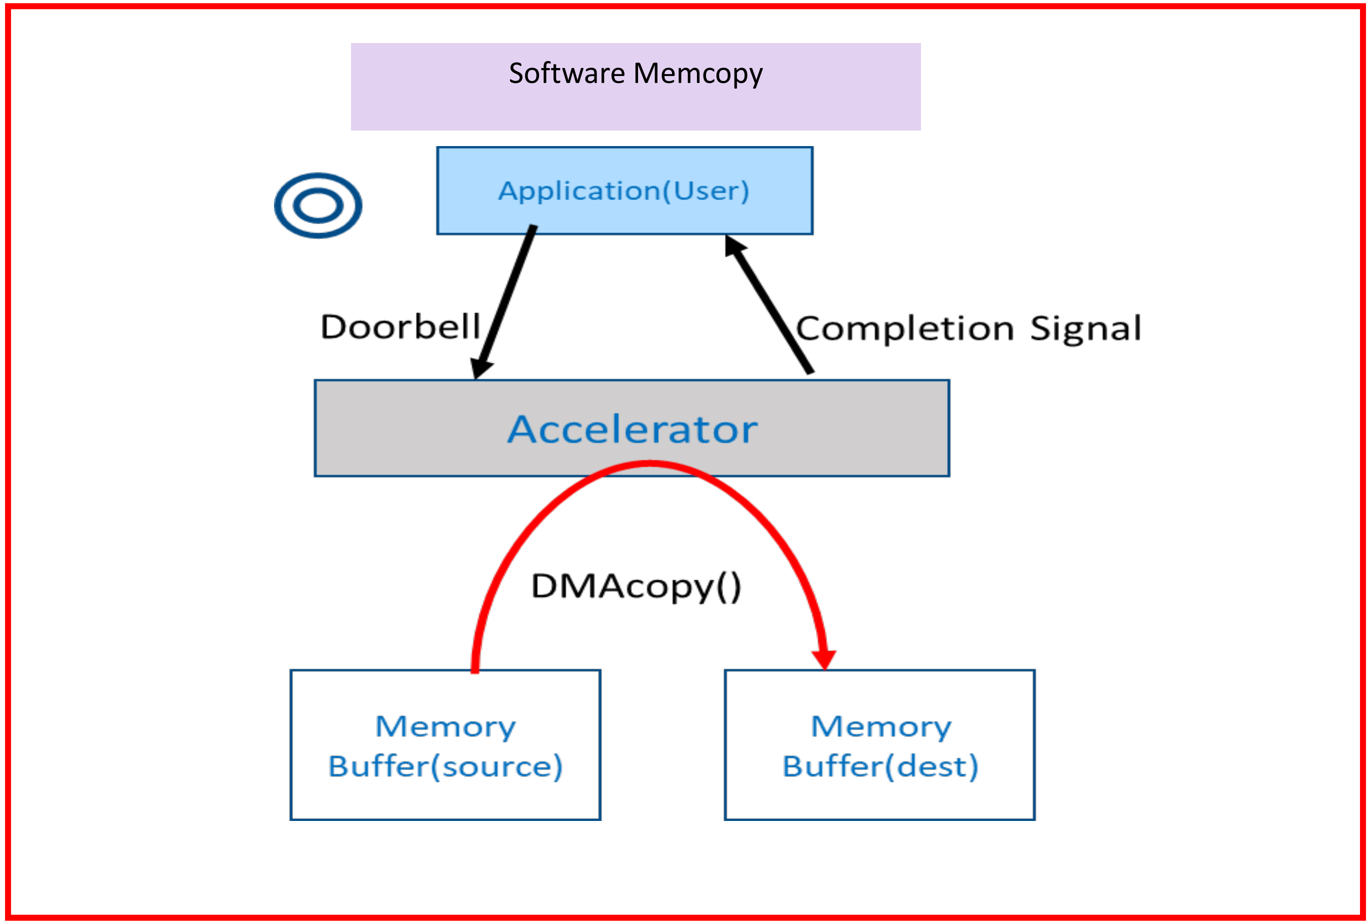

Sample accelerator usage models

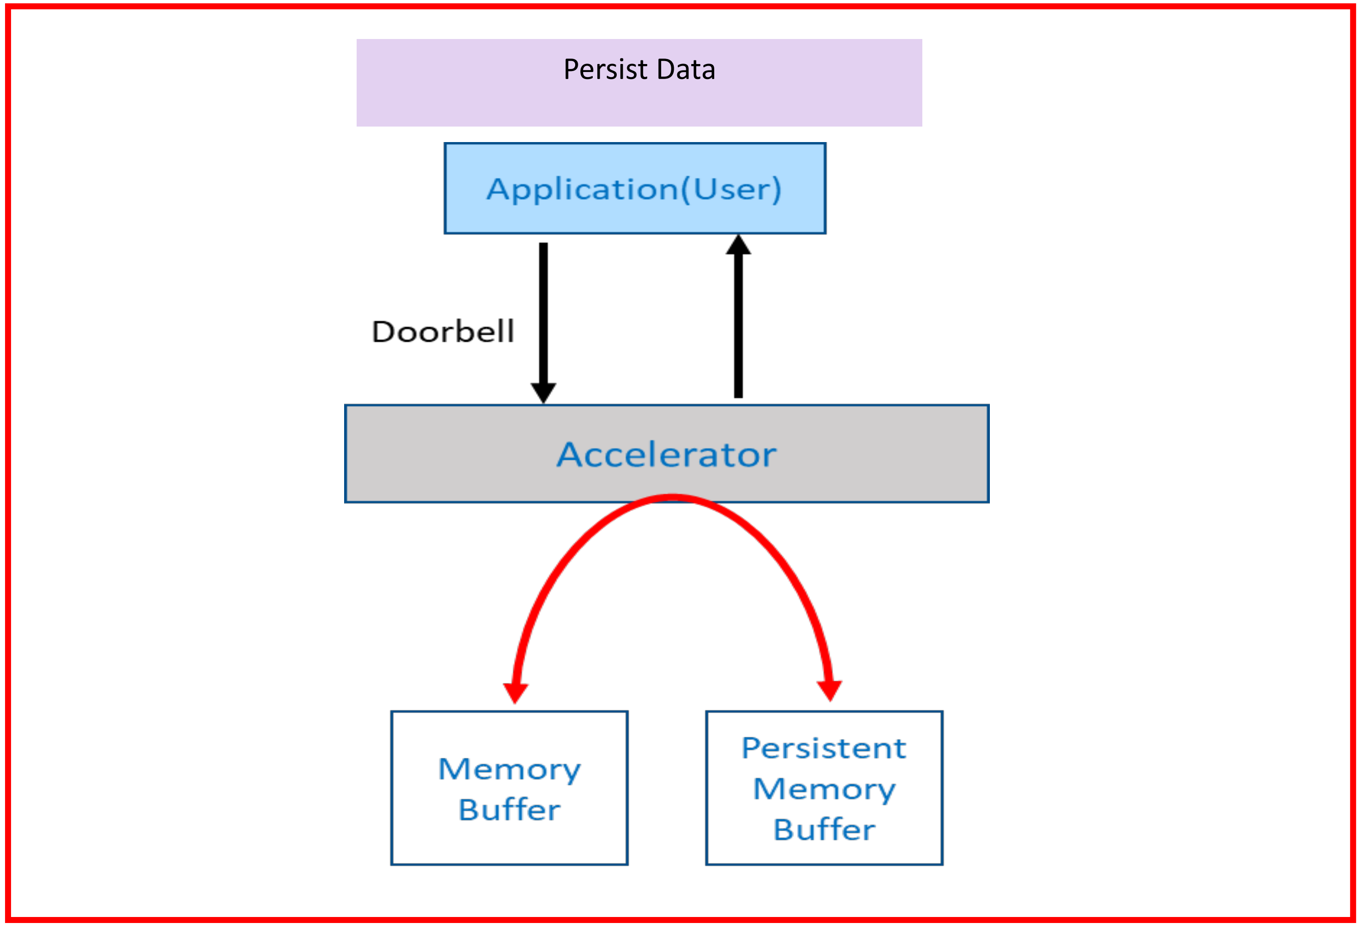

Sample accelerator usage models

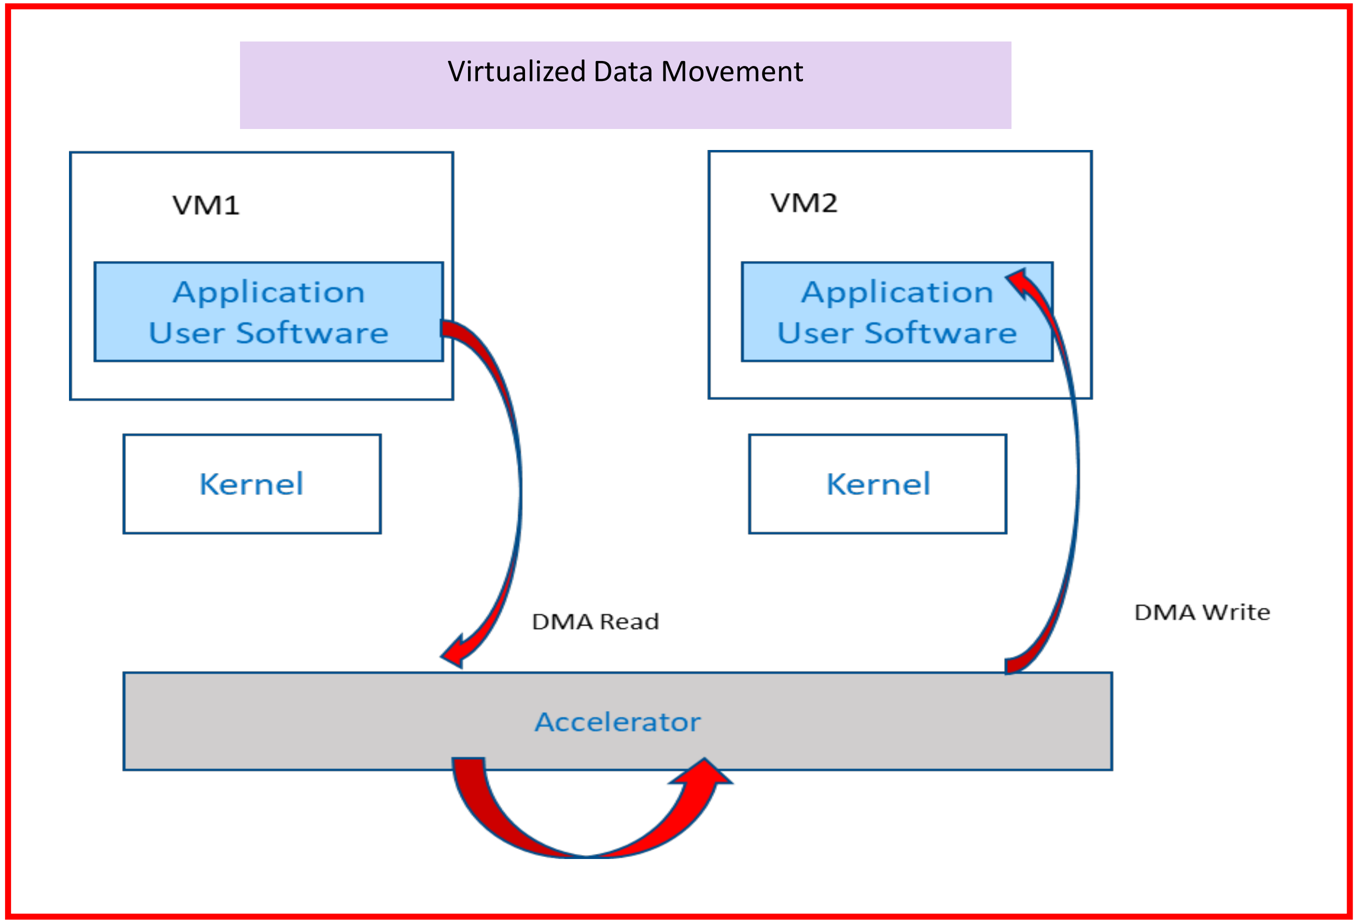
Virtualized Data Movement
VM1
Application User Software
Kernel
VM2
Application User Software
Kernel
DMA Read
DMA Write
Accelerator

Sample accelerator usage models

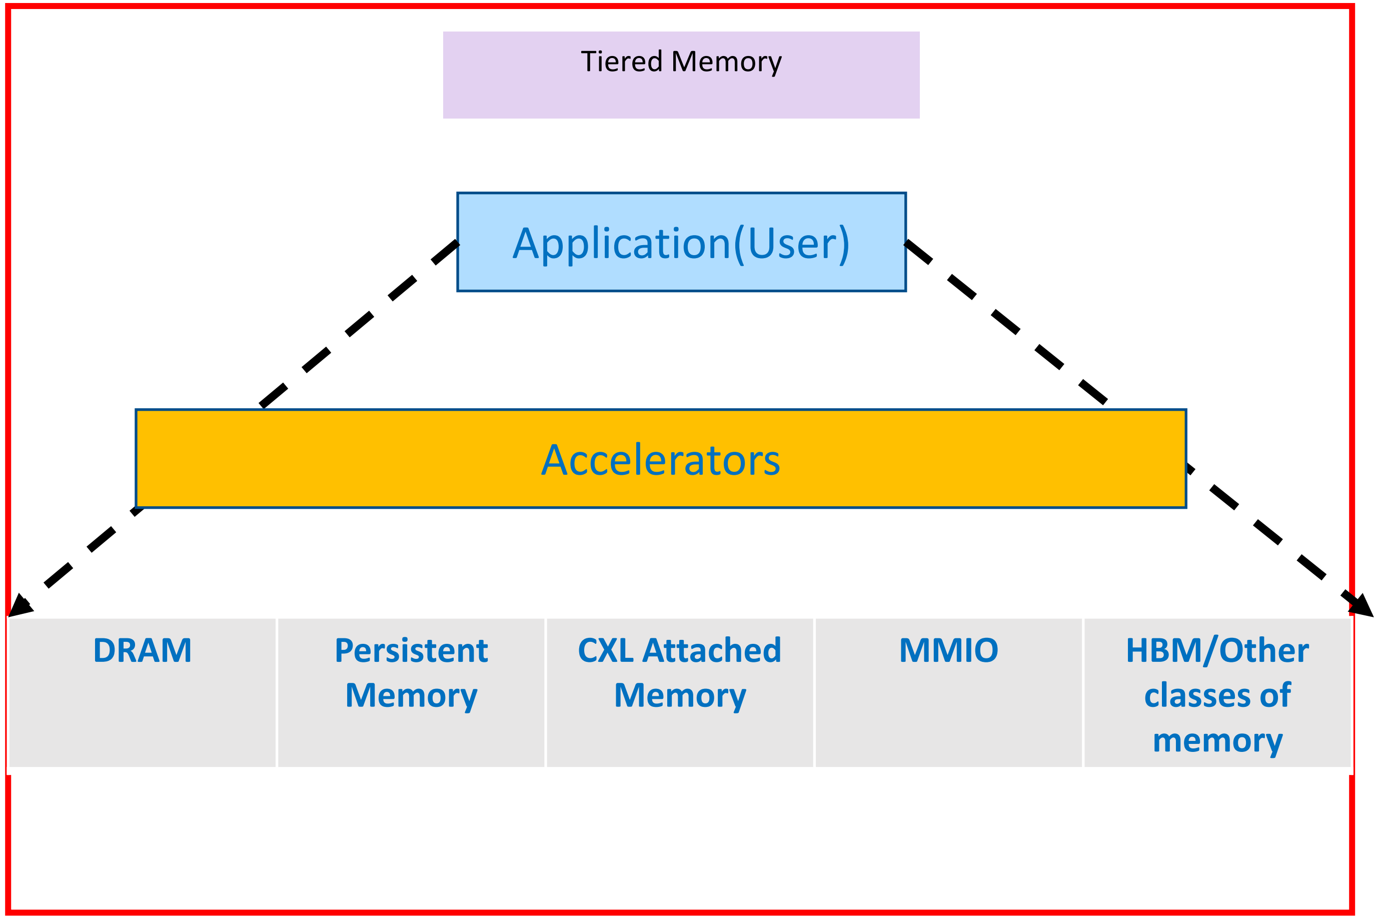

Sample accelerator usage models

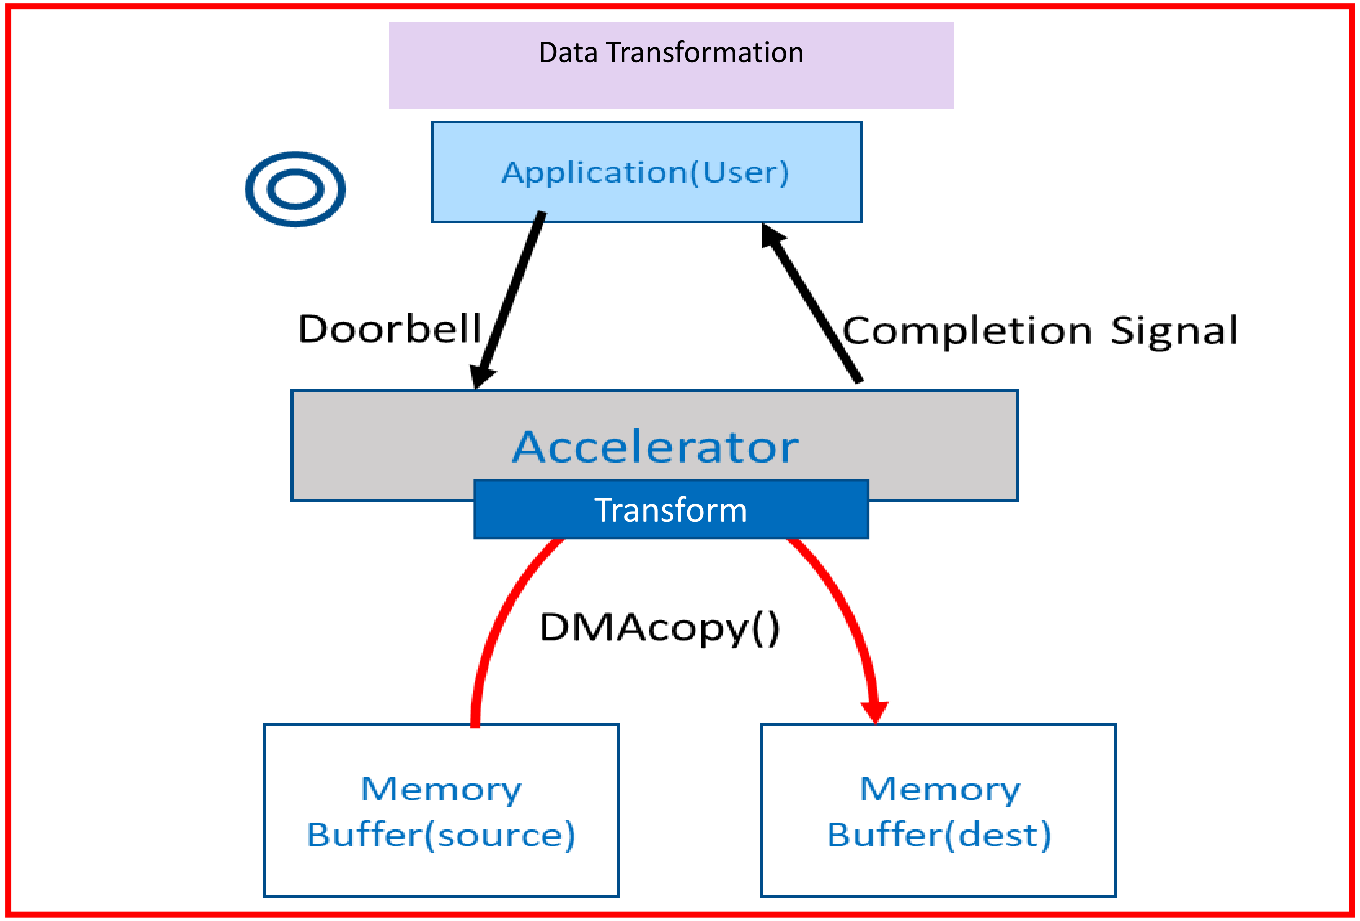

Sample accelerator usage models

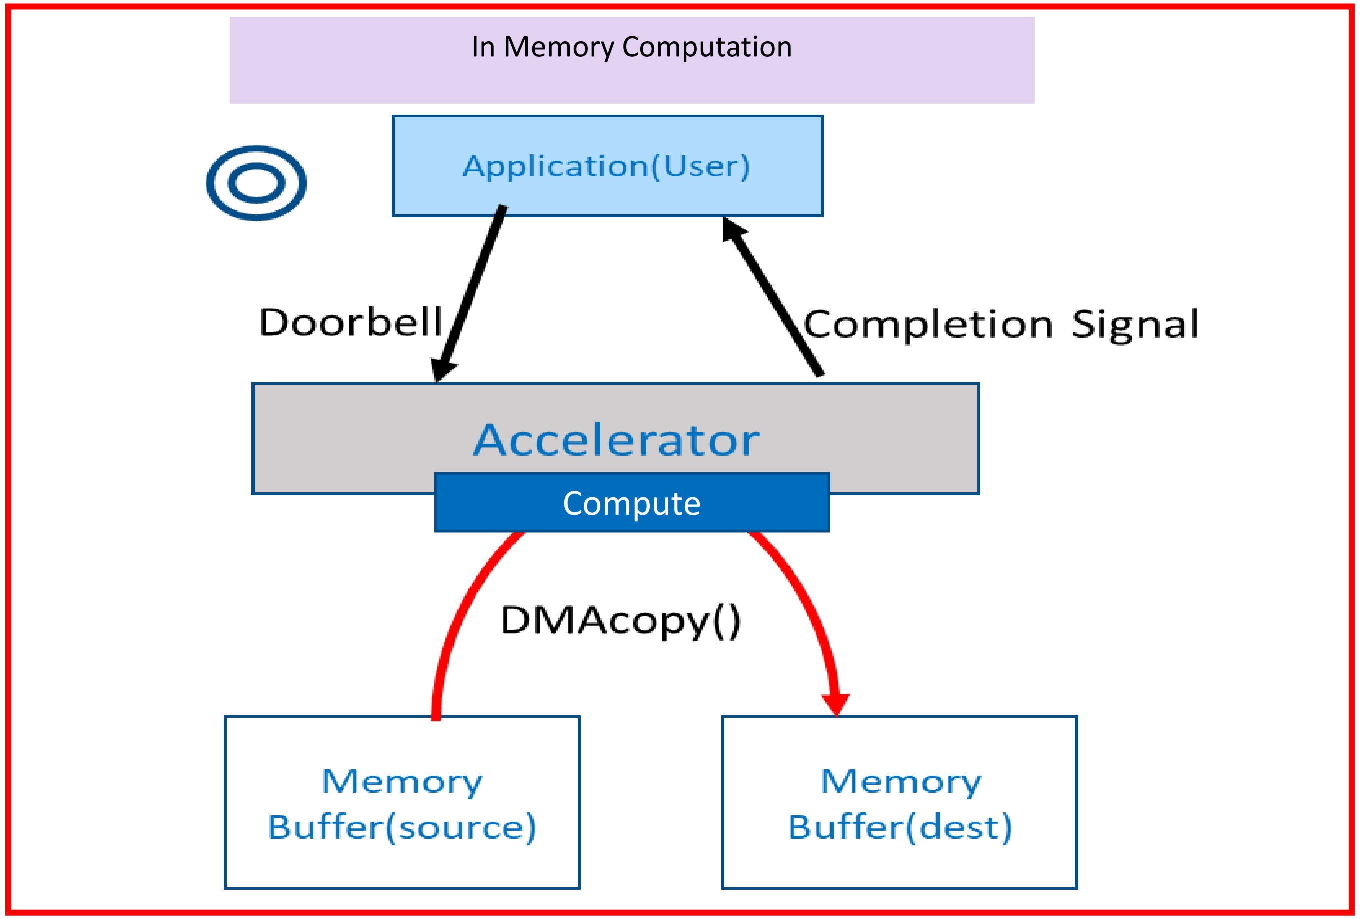

Does it apply to AI? Yes!!!

- Varying data formats and intermediate data representations used in AI/ML data pipelines
  - E.g., file, Columnar, Binary, Text, Tabular, Nested, Array-based, Hierarchical
- Training/inferencing operations involve tensors in memory
- Tensors may be in different address spaces like Host Memor ry, etc.
- Need operations to be able to perform
  - Format Conversions
  - In memory Vector/Tensor transformations li ing, matrix operations, etc.
  - …
- Vendor-specific accelerato en TCO
- Possible Solution: SD
  - Smart Dat ce (SDXI) is a SNIA standard for a memory to memory data m eration interface that is
    - compatible
    - dependent of I/O interconnect technology
  - ata movement between different address spaces.
  - Standard extends to in-memory Offloads/transformations leveraging the architectural interface.

SDXI has emerging AI use cases

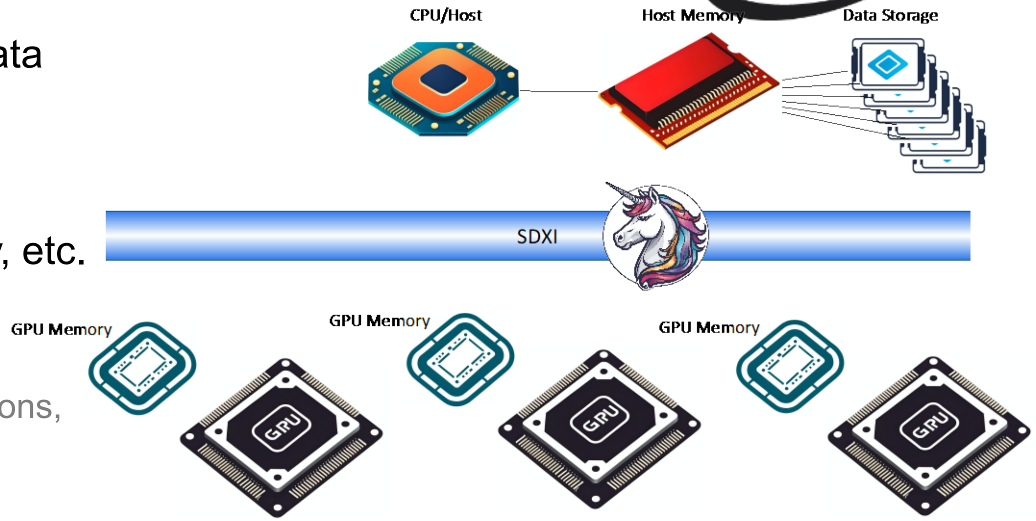

SDXI TWG activities beyond v1.0

- TWG is working on an OS-independent user space software library **libsdxi** for applications.
- TWG members are enabling efforts on SDXI driver work in various Operating Systems
- TWG is discussing efforts to enable SDXI emulation to en SVs
- SDXI specification planning and feature discussion
  - Developed framework for v1.1 features vs 2.0
  - A framework for Definable Operations
  - A connection manager for brokerin een different address spaces
  - New data mover operations tion
    - E.g., POSIX style m , compression, etc.
  - Security Feat movers
    - S al Computing, IDE
  - erations
    - AS, QoS, Latency improvements, CXL-related discussions, SDXI Host to Host investigations, Heterogenous environments
- SDXI+CS subgroup, a collaboration with Computational Storage TWG in SNIA.

The specification continues to evolve beyond v1.0

SDXI in news..

FMS

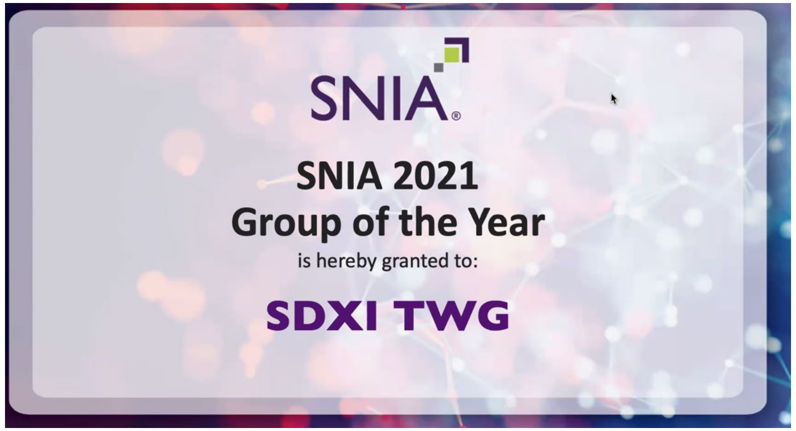

SNIA

SNIA 2021
Group of the Year
is hereby granted to:

SDXI TWG

FLASH MEMORY SUMMIT
BEST OF SHOW

Most Innovative Technology

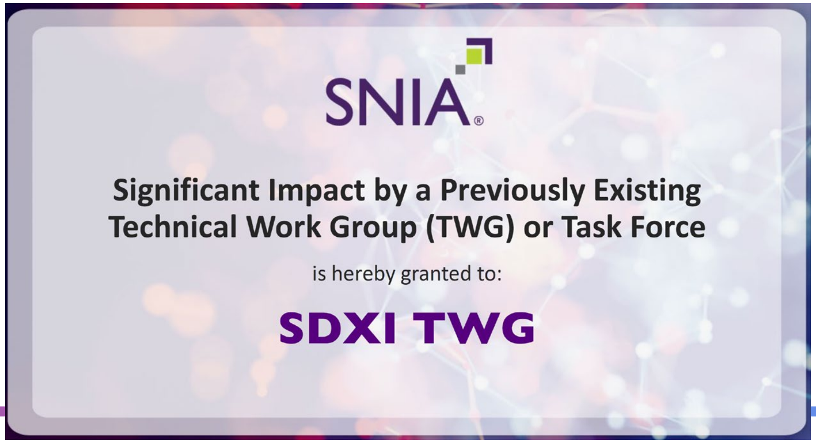

SNIA

Significant Impact by a Previously Existing Technical Work Group (TWG) or Task Force
is hereby granted to:

SDXI TWG

SDXI PoC Demo at Memcon 2024

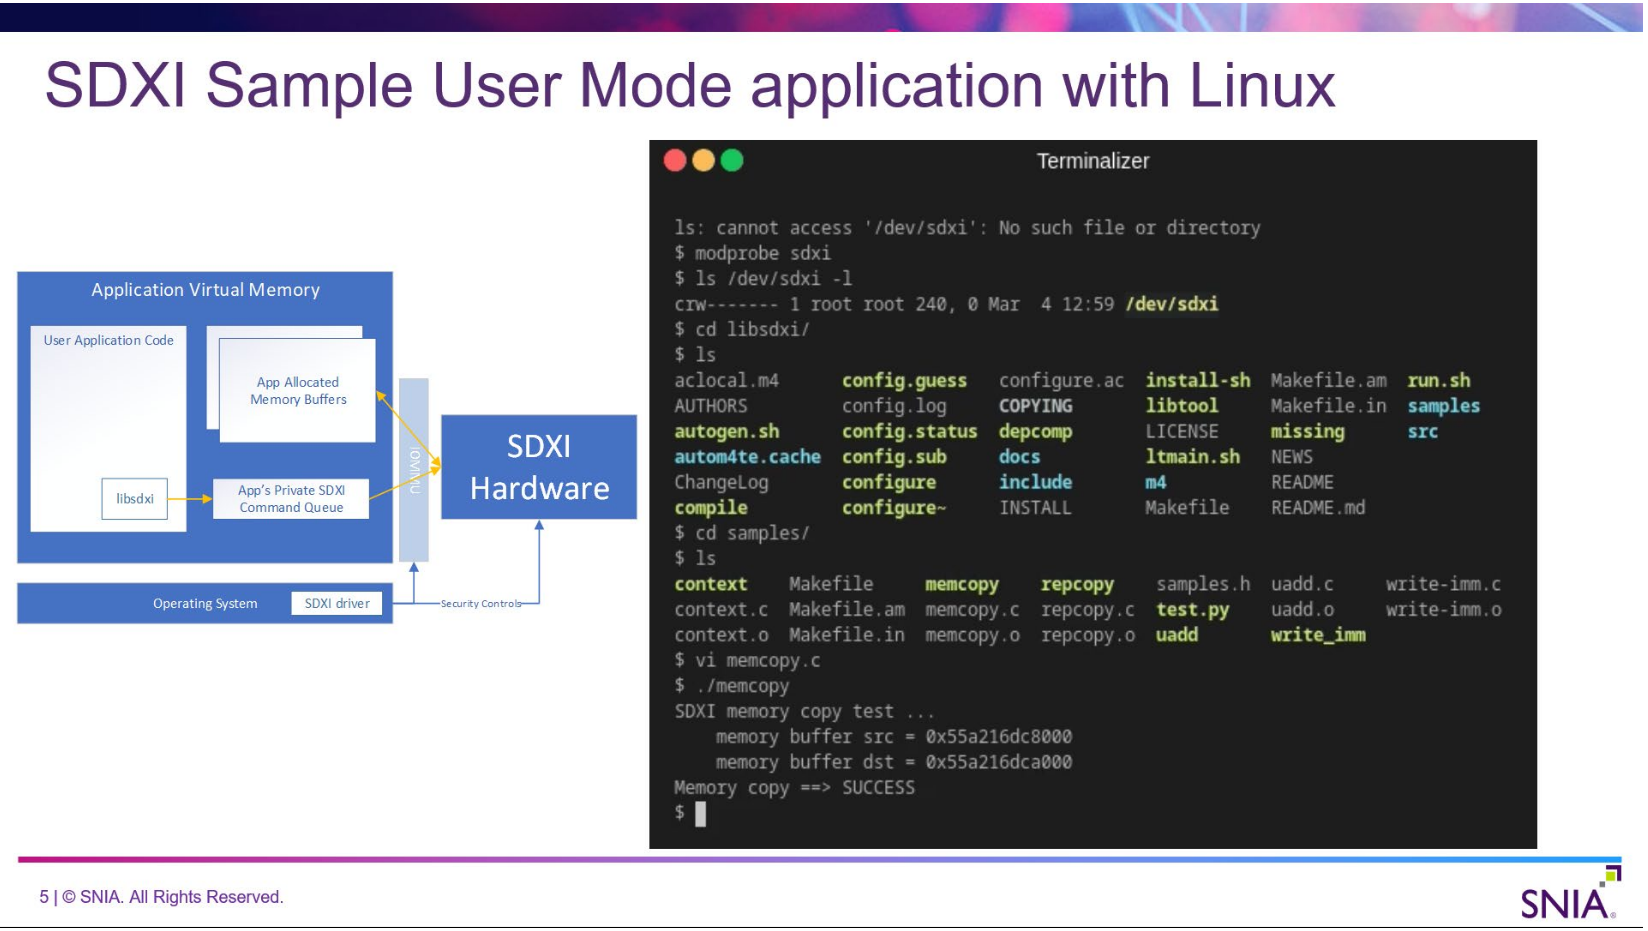

SDXI Sample User Mode application with Linux

```
ls: cannot access '/dev/sdxi': No such file or directory
$ modprobe sdxi
$ ls /dev/sdxi -l
crw------- 1 root root 240, 0 Mar  4 12:59 /dev/sdxi
$ cd libsdxi/
$ ls
aclocal.m4      config.guess   configure.ac  install-sh  Makefile.am  run.sh
AUTHORS         config.log     COPYING       libtool     Makefile.in  samples
autogen.sh      config.status  depcomp       LICENSE     missing      src
autom4te.cache  config.sub     docs          ltmain.sh   NEWS
ChangeLog       configure      include       m4          README
compile         configure~     INSTALL       Makefile    README.md
$ cd samples/
$ ls
context    Makefile     memcopy    repcopy    samples.h  uadd.c     write-imm.c
context.c  Makefile.am  memcopy.c  repcopy.c  test.py    uadd.o     write-imm.o
context.o  Makefile.in  memcopy.o  repcopy.o  uadd       write_imm
$ vi memcopy.c
$ ./memcopy
SDXI memory copy test ...
    memory buffer src = 0x55a216dc8000
    memory buffer dst = 0x55a216dca000
Memory copy ==> SUCCESS
$ 
```

SNIA’s SDXI standard is a great start…

SNIA Data Focus Areas

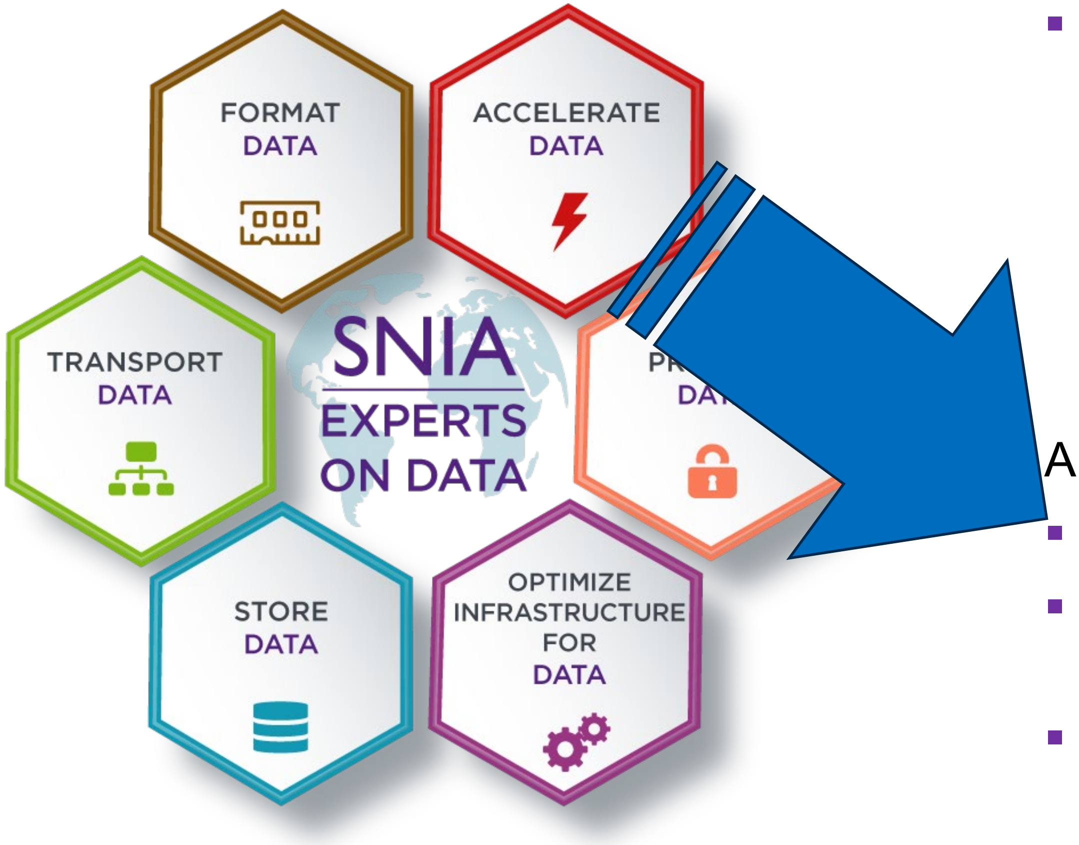

Accelerate: Move processing to the data

- Technologies that move processing closer to the data, enabling improvements in application performance and/or infrastructure efficiency through the integration of compute resources (outside of the traditional compute & memory architecture) either directly with storage or between the host and the storage.

Areas of Interest:

- SNIA SDXI (Smart Data Accelerator Interface)
- SNIA Computational Storage Architecture and Programming Model, SNIA Computational Storage API
- DPU (Data Processing Unit)

And there is more to do…

- TCO evaluation
  - Workload specific usage, deeper analysis
  - Leverage common framework for stack normalization
- Accelerators: To specialize or Generalize
- Security
  - Baseline platform security with/without accelerators
  - New standardization requirements
  - Security cannot be viewed in isolation
- Management
  - Newer modeling of Systems management profiles including accelerators in system topology
- Diversity in accelerator types and their specific standardization requirements
  - CPU-integrated accelerators
  - Discrete Accelerators
    - GPUs
    - DPUs
    - Look aside accelerators
  - In-Memory accelerators
  - Network Inline accelerators

Summary and Call to Action

- Accelerators becoming essential to various workloads
- SNIA is championing a data-centric focus area with accelerators
- SNIA SDXI TWG has taken first steps in standardizing memory data movement and transformation accelerators with SDXI v1.0
- New accelerator form factors like DPUs complementing CPUs and GPUs
- Join SNIA's focus groups to impact the industry!