Smart Data Accelerator Interface: Use cases, futures, and proof points

SNIA COMPUTE, MEMORY, AND STORAGE SUMMIT

Solutions, Architectures, and Community

VIRTUAL EVENT, MAY 21-22, 2024

Shyam Iyer

SDXI TWG Chair

SNIA Technical Council

Distinguished Engineer, Dell Technologies

sdxitwgchair@snia.org

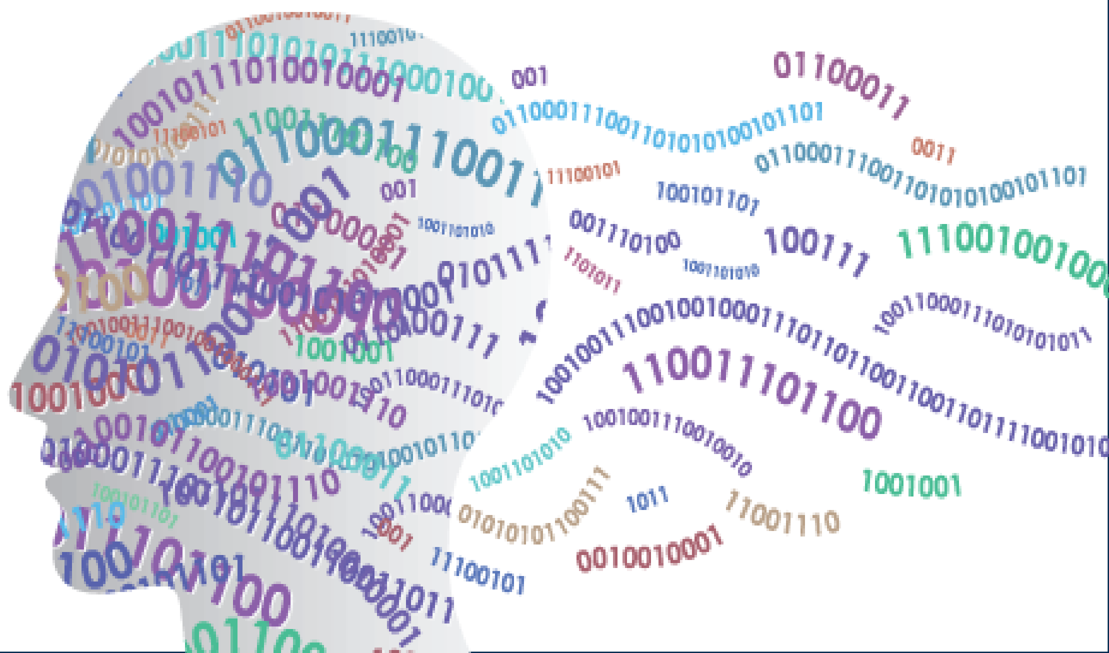

Agenda

- SNIA SDXI TWG
- SDXI Use cases
- Looking Ahead
- Proof Points
- Summary and Call to Action

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

SNIA COMPUTE, MEMORY, AND STORAGE SUMMIT

Solutions, Architectures, and Community

VIRTUAL EVENT, MAY 21-22, 2024

Sample accelerator usage models

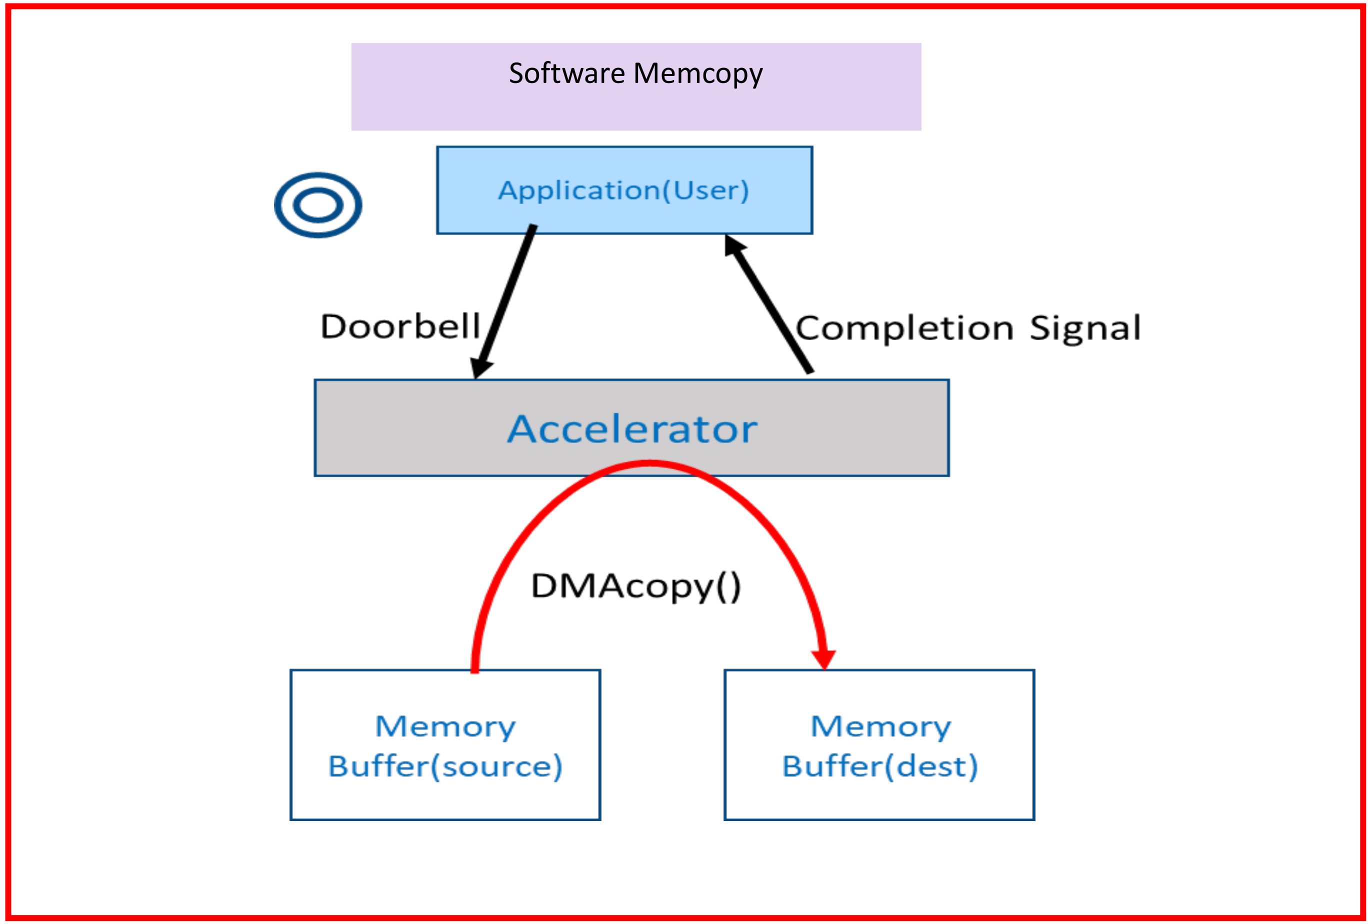

Sample accelerator usage models

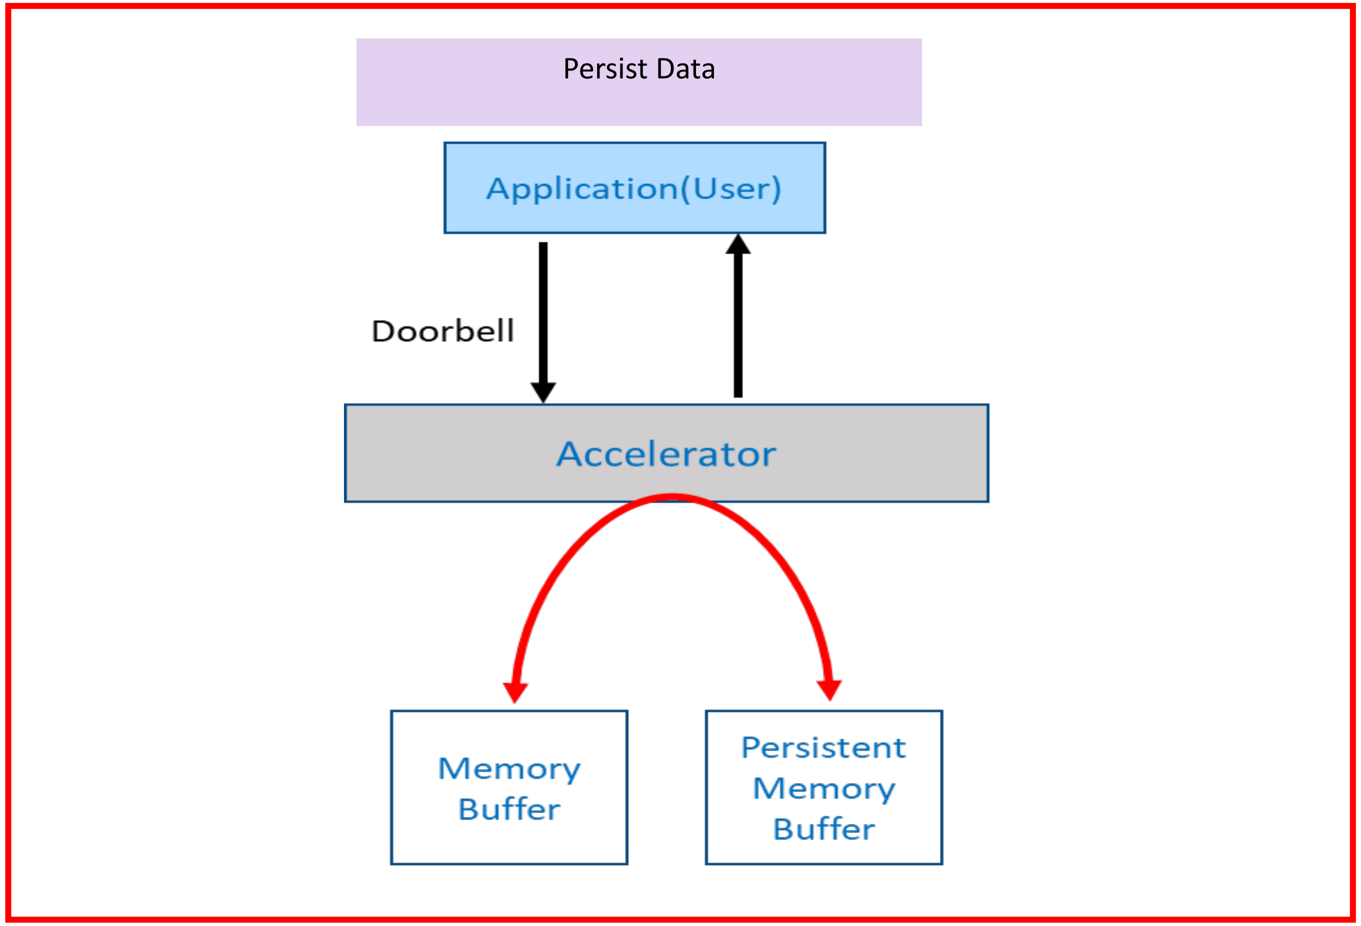

Sample accelerator usage models

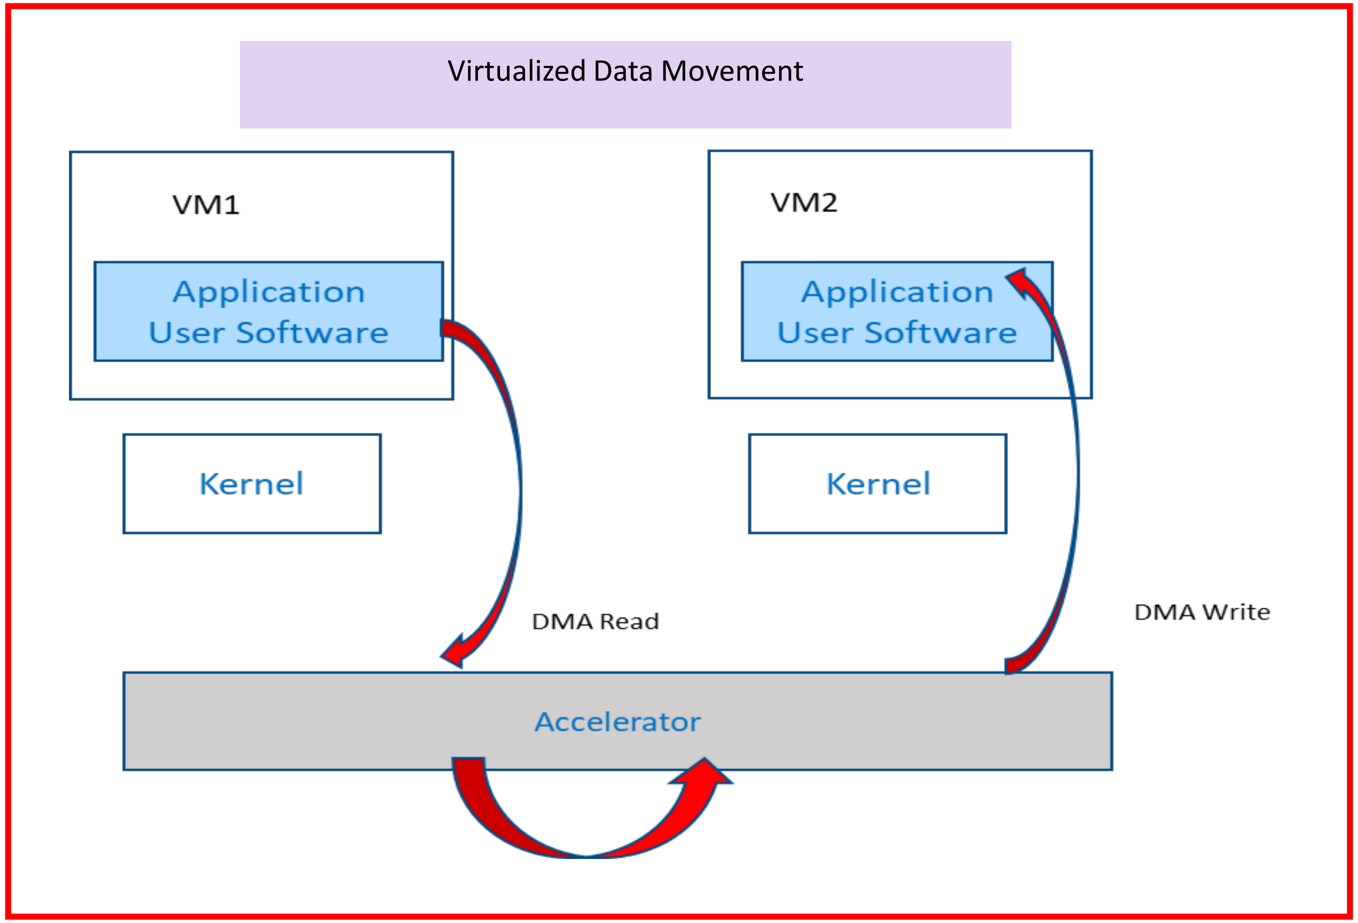
Virtualized Data Movement
VM1
Application User Software
Kernel
VM2
Application User Software
Kernel
DMA Read
DMA Write
Accelerator

Sample accelerator usage models

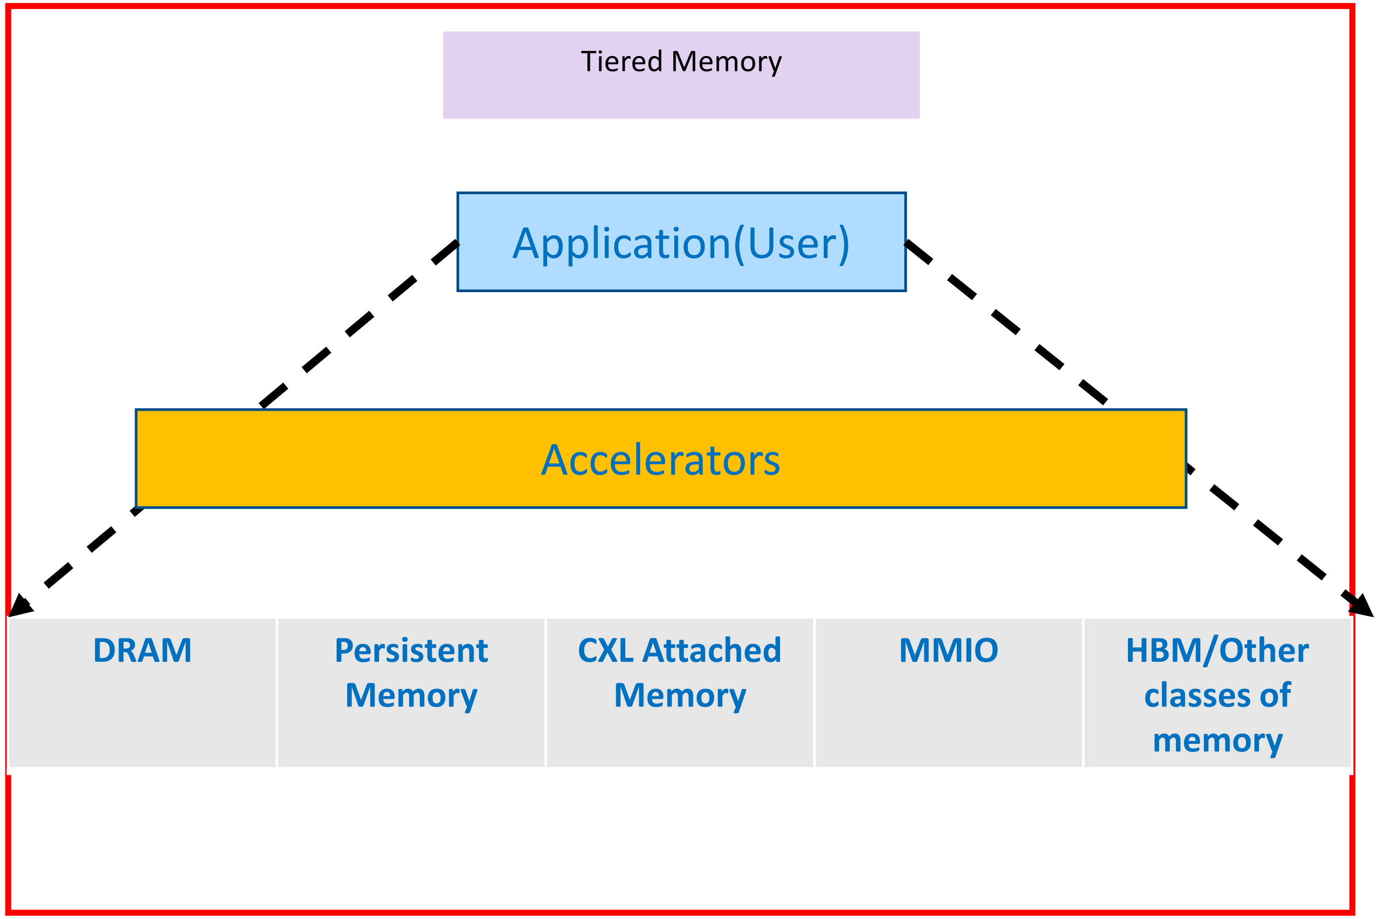

Sample accelerator usage models

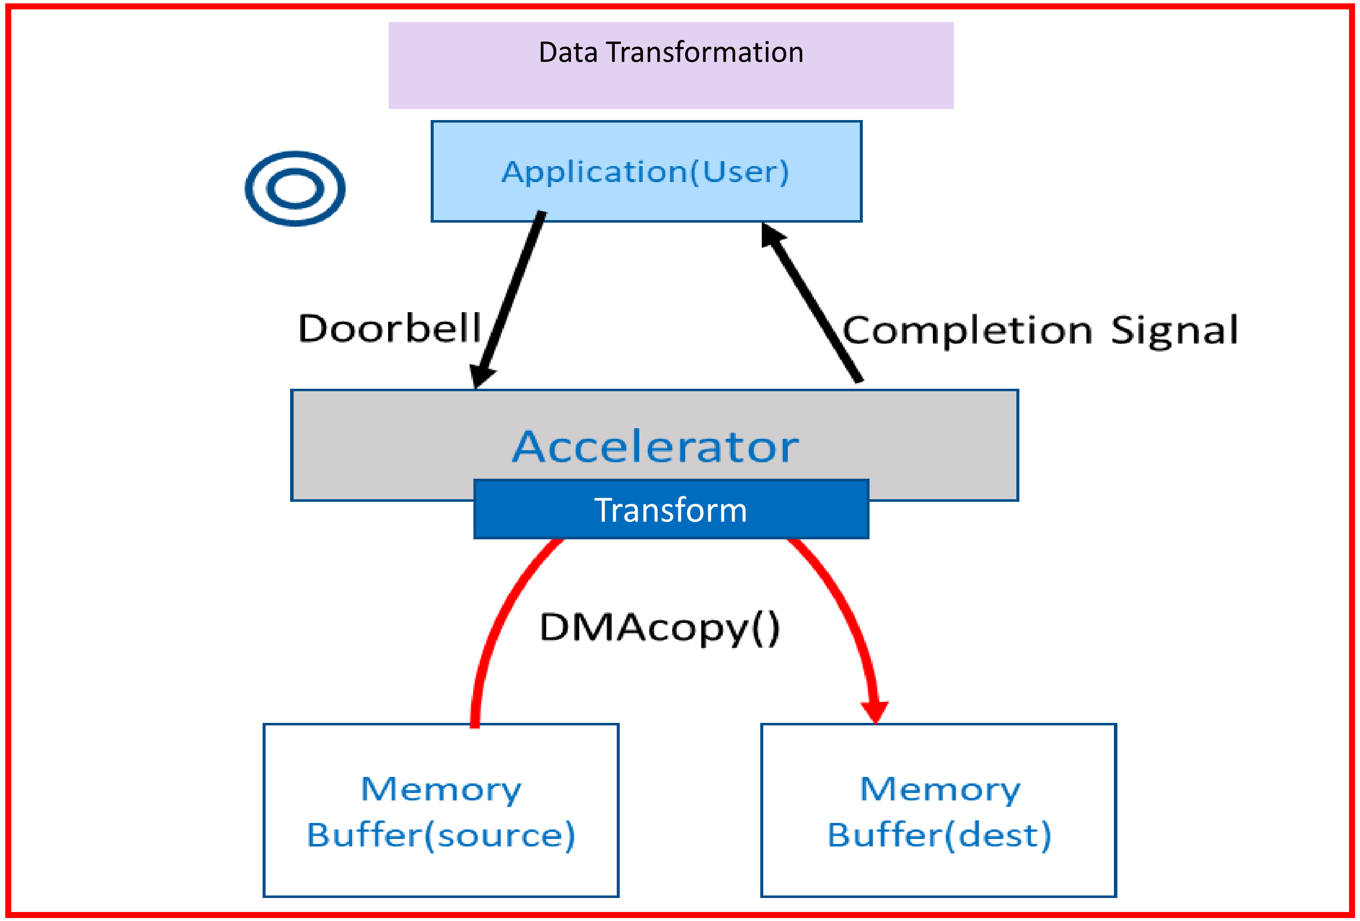

Sample accelerator usage models

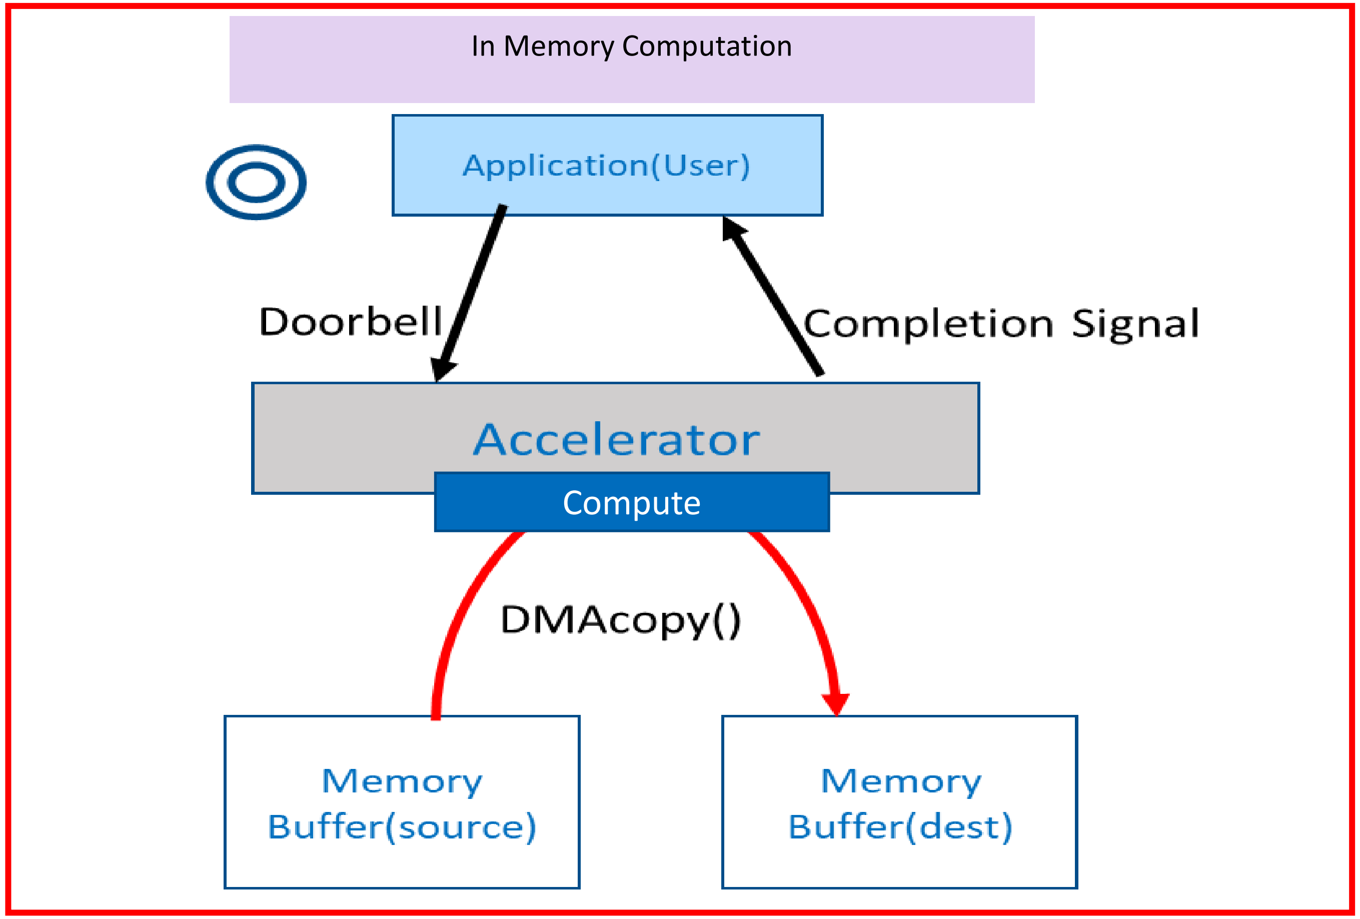

Does it apply to AI? Yes!!!

- 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 Memory, GPU Memory, etc.
- Need operations to be able to perform
  - Format Conversions
  - In memory Vector/Tensor transformations like quantization, scaling, matrix operations, etc.
  - …
- Vendor-specific accelerator operations weaken TCO
- Possible Solution: SDXI
  - 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
  - Data movement between different address spaces.
  - Standard extends to in-memory Offloads/transformations leveraging the architectural interface.

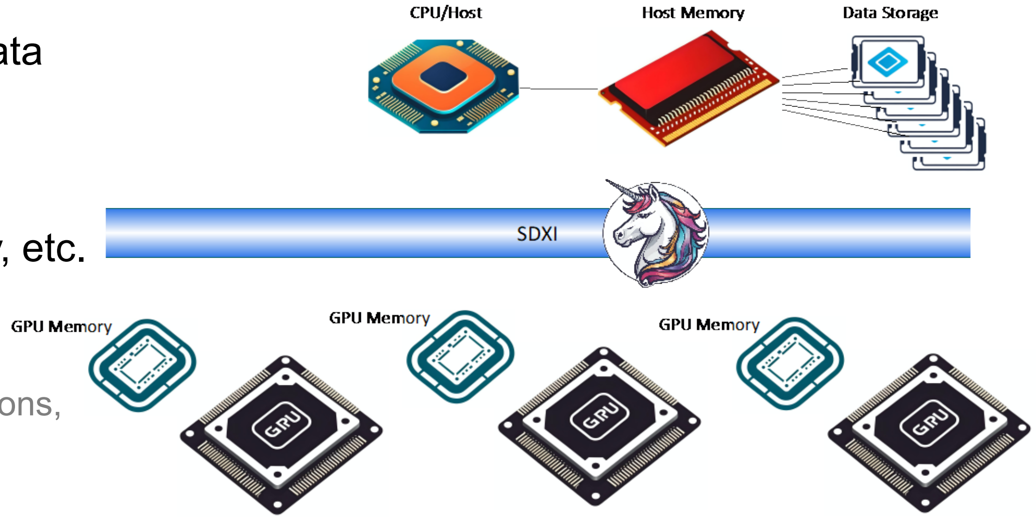

SDXI Memory-to-Memory Data Movement

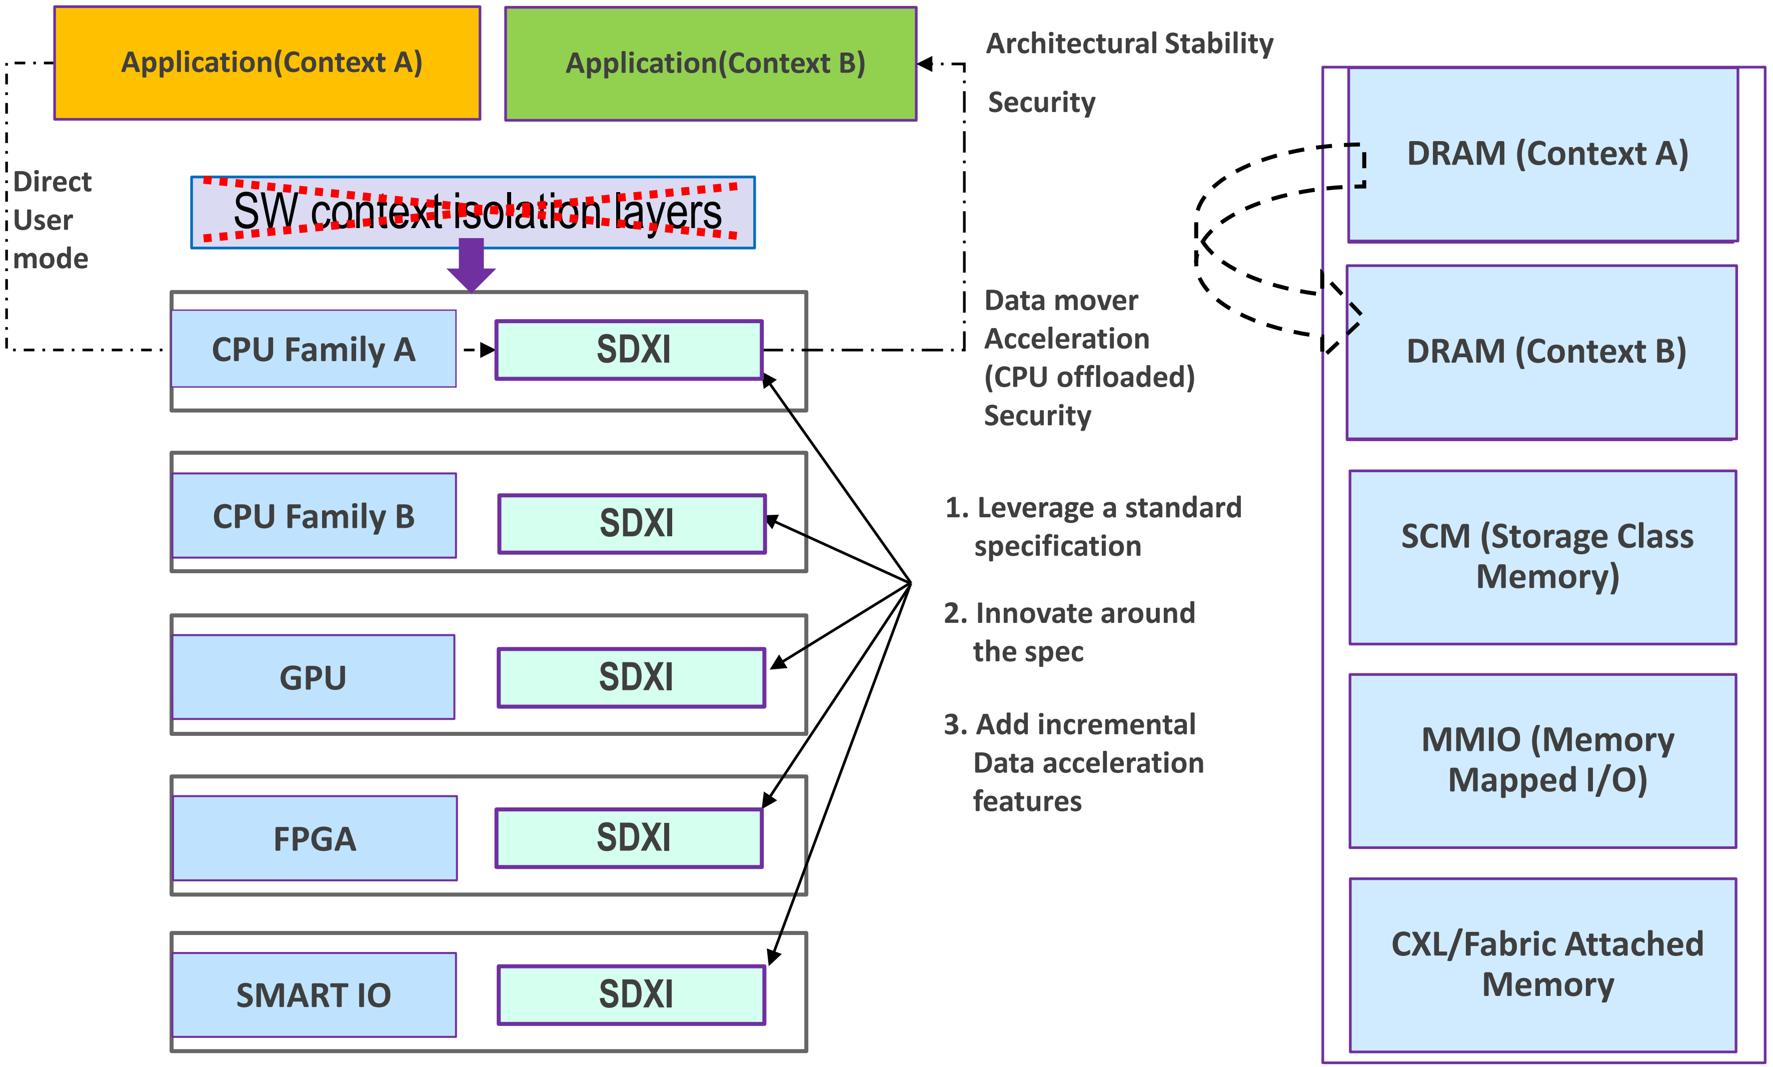

1. Leverage a standard specification
2. Innovate around the spec
3. Add incremental Data acceleration features

System Physical Address space

We are entering a tiered Memory world !

Looking ahead…

SNIA COMPUTE, MEMORY, AND STORAGE SUMMIT

Solutions, Architectures, and Community

VIRTUAL EVENT, MAY 21-22, 2024

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 enable ISVs
- SDXI specification planning and feature discussions
  - Developed framework for v1.1 features vs 2.0
  - A framework for Definable Operations
  - A connection manager for brokering connections between different address spaces
  - New data mover operations for smart acceleration
    - E.g., POSIX style memory operations, CRC, DIF, compression, etc.
  - Security Features involving data movers
    - SDXI devices and Confidential Computing, IDE
    - Threat modeling
  - Other considerations
    - RAS, QoS, Latency improvements, CXL-related discussions, SDXI Host to Host investigations, Heterogenous environments
- SDXI+CS subgroup, a collaboration with Computational Storage TWG in SNIA.

Proof Points

SNIA COMPUTE, MEMORY, AND STORAGE SUMMIT

Solutions, Architectures, and Community

VIRTUAL EVENT, MAY 21-22, 2024

SDXI in news..

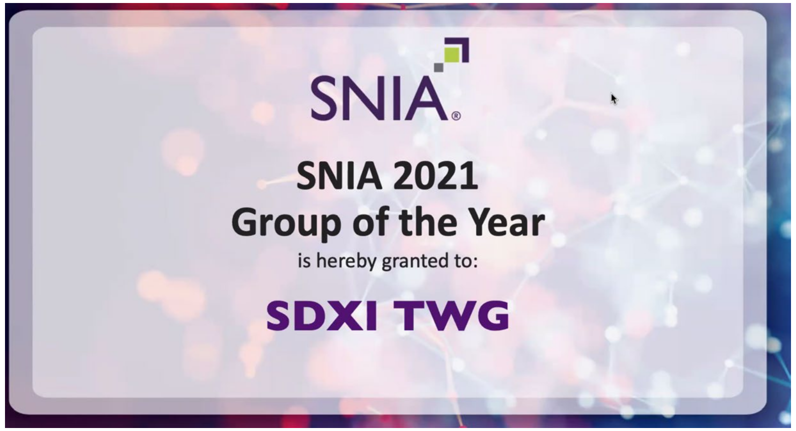

SNIA

SNIA 2021
Group of the Year

is hereby granted to:

SDXI TWG

Most Innovative Technology

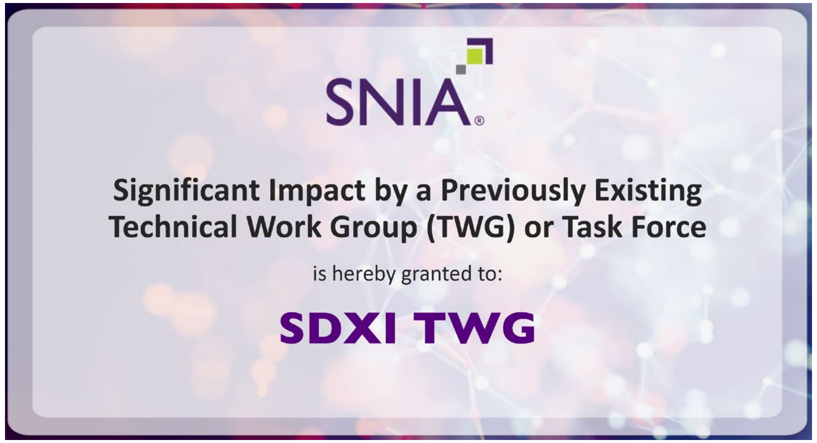

SNIA

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

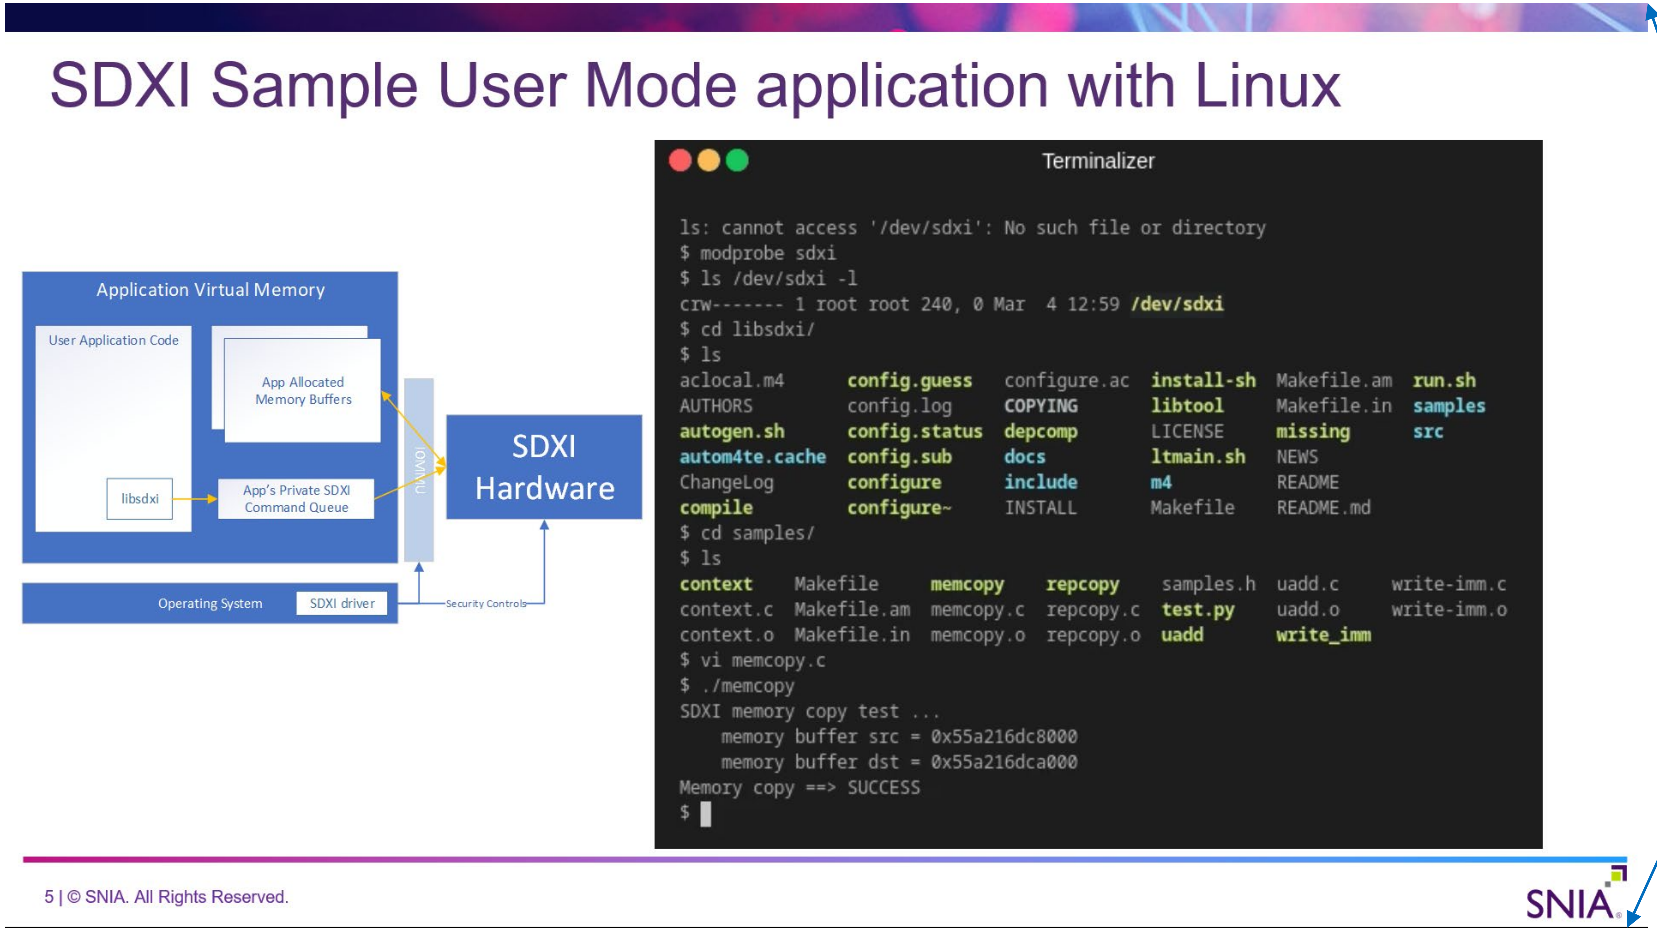

SDXI Sample User Mode application with Linux

```
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
$ 
```

5 | © SNIA. All Rights Reserved.

Summary and call to action

- SDXI TWG is working on SDXI v1.1 and beyond
- SDXI Software ecosystem is emerging
- SDXI-based accelerators have a wide variety of applicable use cases including emerging areas like AI
- Come join the group and participate in the growing ecosystem