

SDXI Technical Workgroup Charter

SUBMITTER: Shyam Iyer, Philip Ng, Rich Brunner

DATE SUBMITTED: June 22, 2020

VERSION: 1.0

APPROVED by Technical Council: 6/22/2020

CHARTER:

SDXI is a proposed standard for a memory to memory Data Mover interface. Software memcopy is the current data movement standard for software implementation due to stable CPU ISA. However, this takes away from application performance and incurs software overhead to provide context isolation. Offload DMA engines and their interface are vendor-specific and not standardized for user-level software.

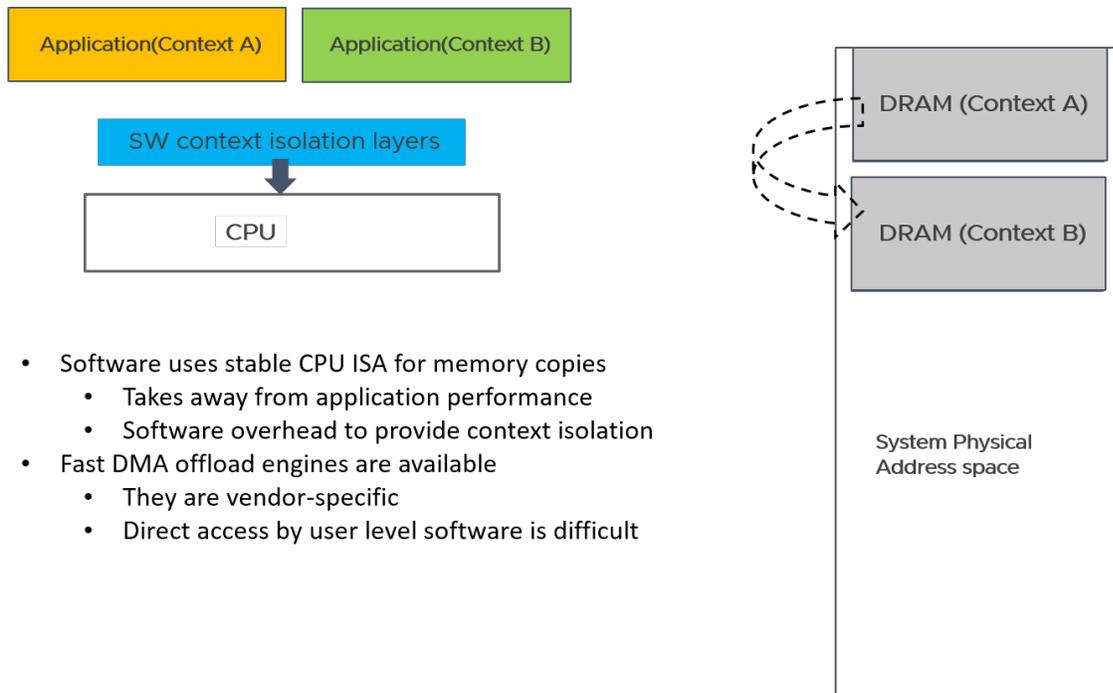


Fig.1 Current Data movement standard

The SDXI TWG -

- Develops and standardizes an extensible, forward-compatible memory to memory data mover interface that is independent of actual data mover implementations and underlying I/O interconnect technology. A PCIe device model is described. The interface can be extended to other device models.

- Proposed standard supports-
 - data movement between different address spaces including user address spaces located within different virtual machines
 - data movement without mediation by privileged software once a connection has been established.
 - an interface and architecture that can be abstracted or virtualized by privileged software to allow greater compatibility of workloads or virtual machines across different servers.
 - a well-defined capability to quiesce, suspend, and resume the architectural state of a per-address-space data mover to allow “live” workload or virtual machine migration between servers.
 - mechanisms to enable forwards and backwards compatibility across future specification revisions. Allows software and hardware designed to different specification revisions to interoperate.
 - ability to incorporate additional offloads in the future leveraging the architectural interface.
 - concurrent DMA model.
- coordinates the submission of new feature proposals to standards groups like PCISIG, CXL, OFA, UEFI, Gen-Z etc.

As new memory technologies get adopted and memory fabrics expand the use of tiered memory, data mover acceleration and its uses will increase. This TWG will encourage adoption and extensions to this data mover interface.

This TWG will assist and cooperate with other SNIA TWGs such as the Persistent Memory work group and Computational Storage work group to leverage data mover technology. This TWG will also consider other SNIA TWG and Alliance partner work for leveraging.

This TWG will work for an architectural vision (See Fig.2) that supports a standard data mover interface for a variety of memory technologies and implementation choices.

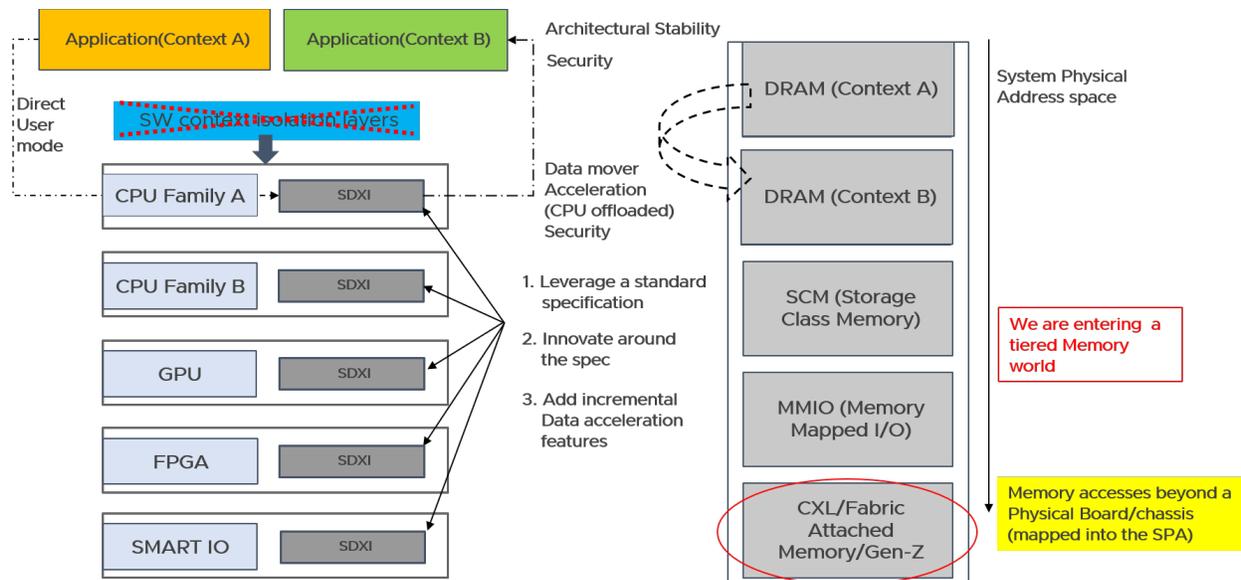


Fig.2 SDXI Architectural vision

PROPOSED PROGRAM OF WORK:

1. The sponsoring group will submit a draft spec as a starting point for a SNIA Architecture.
2. The TWG will develop a SNIA Architecture for a memory to memory data mover based on the submission.
3. Following completion of version 1.0 the TWG will commence activities on new focus areas related to data movement, such as:
 - 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
4. The TWG will encourage adopting companies to work towards compliant software implementations and driver models.
5. The TWG will work with SNIA Marketing to educate and encourage adoption by OS, Hypervisors, OEMs, Applications and Data Acceleration vendors.

RESOURCES:

- Dell, AMD and VMware have committed resources and support for this activity
 - Shyam Iyer (Dell), Philip Ng (AMD), Rich Brunner (VMware)
- Update 6-22-2020:
 - Microsoft, NetApp, Xilinx, Arm, Fujitsu, SK Hynix, Western Digital, NGD Systems have committed support to participate and join the SDXI TWG
- More companies are expected to join this effort