

SNIA SDXI Roundtable: Towards Standardizing a Memory to Memory Data Movement and Acceleration Interface



Moderator:

Shyamkumar Iyer, SDXI TWG Chair, Distinguished Engineer, Dell

Panelists(L-R):

Philip Ng, Sr. Fellow, AMD


Alexandre Romana, Principal System Architect, ARM

Jason Wohlgemuth, Partner Software Engineering Lead, Microsoft


Donald Dutile, Principal Software Engineer, Red Hat, Inc

Richard Brunner, Principal Engineer and CTO of Server Platform Technologies, VMware

Paul Hartke, Principal Engineer, Xilinx



Panelists



Philip Ng

- Philip Ng is a Sr. Fellow at AMD where he works on I/O subsystem architecture including technologies such as PCI-Express and I/O virtualization. He holds a B.A.Sc and M.Eng from the University of Toronto.



Alexandre Romana

- Alexandre Romana is an experienced, hands-on system architect who has been involved in the architecture, design, and implementation of complex system-on-chips with a strong focus on performance and energy-efficiency. He currently focuses on the future of virtualization, and was an architect for The Revere Accelerator Management Unit (Revere-AMU) System Architecture



Jason Wohlgemuth

- Jason Wohlgemuth is a Partner Software Engineering Lead in the Windows Core OS team who manages a group of engineers that aid in the implementation of OS support for emerging hardware and silicon technology. He is a key architect and technical leader for system and OS architecture, representing Microsoft in various industry forums and technical boards.



Panelists



Donald Dutile

- Donald Dutile is a Principal Software Engineer at Red Hat. Don has 40 years of experience in the computer industry. His focus over those years has primarily been in I/O subsystems. At Red Hat, he has worked in the areas of device-assignment for KVM, IOMMU, and PCI (SRIOV, RedHat's PCI-SIG representative) support. He led the RHEL RDMA subsystem for a few years, and maintained the initial RHEL kernel package for ARM64. Don moved on to the RHEL core kernel team focusing on the mm subsystem as well as dma-mapping and peer to peer DMA support. Previous to Red Hat, he provided Linux kernel support for StarGen's StarFabric and ASI fabrics. Prior to StarGen, Don worked for DEC in its UNIX group, supporting new Alpha platforms, PCI and IOMMU support. During his first 10 years, he admits to being a hardware designer on DEC's cluster interconnect and a brief stint as a MIPS System architect.



Richard A. Brunner

- Richard A. Brunner is a Principal Engineer and the CTO of Server Platform Technologies at VMware. He is responsible for defining and driving future platform features that can benefit virtualization. Richard has been focused on computer technology since 1977 as a programmer, design engineer, and then instruction set architect for IBM, DEC, Intel, and AMD. At VMware since 2006, he has been focusing on Non-Volatile Memory, composition of disaggregated resources, CPU Security Vulnerabilities, and platform offline acceleration.



Paul Hartke

- Paul Hartke is a member of the Xilinx CTO team working on architectures for heterogenous systems

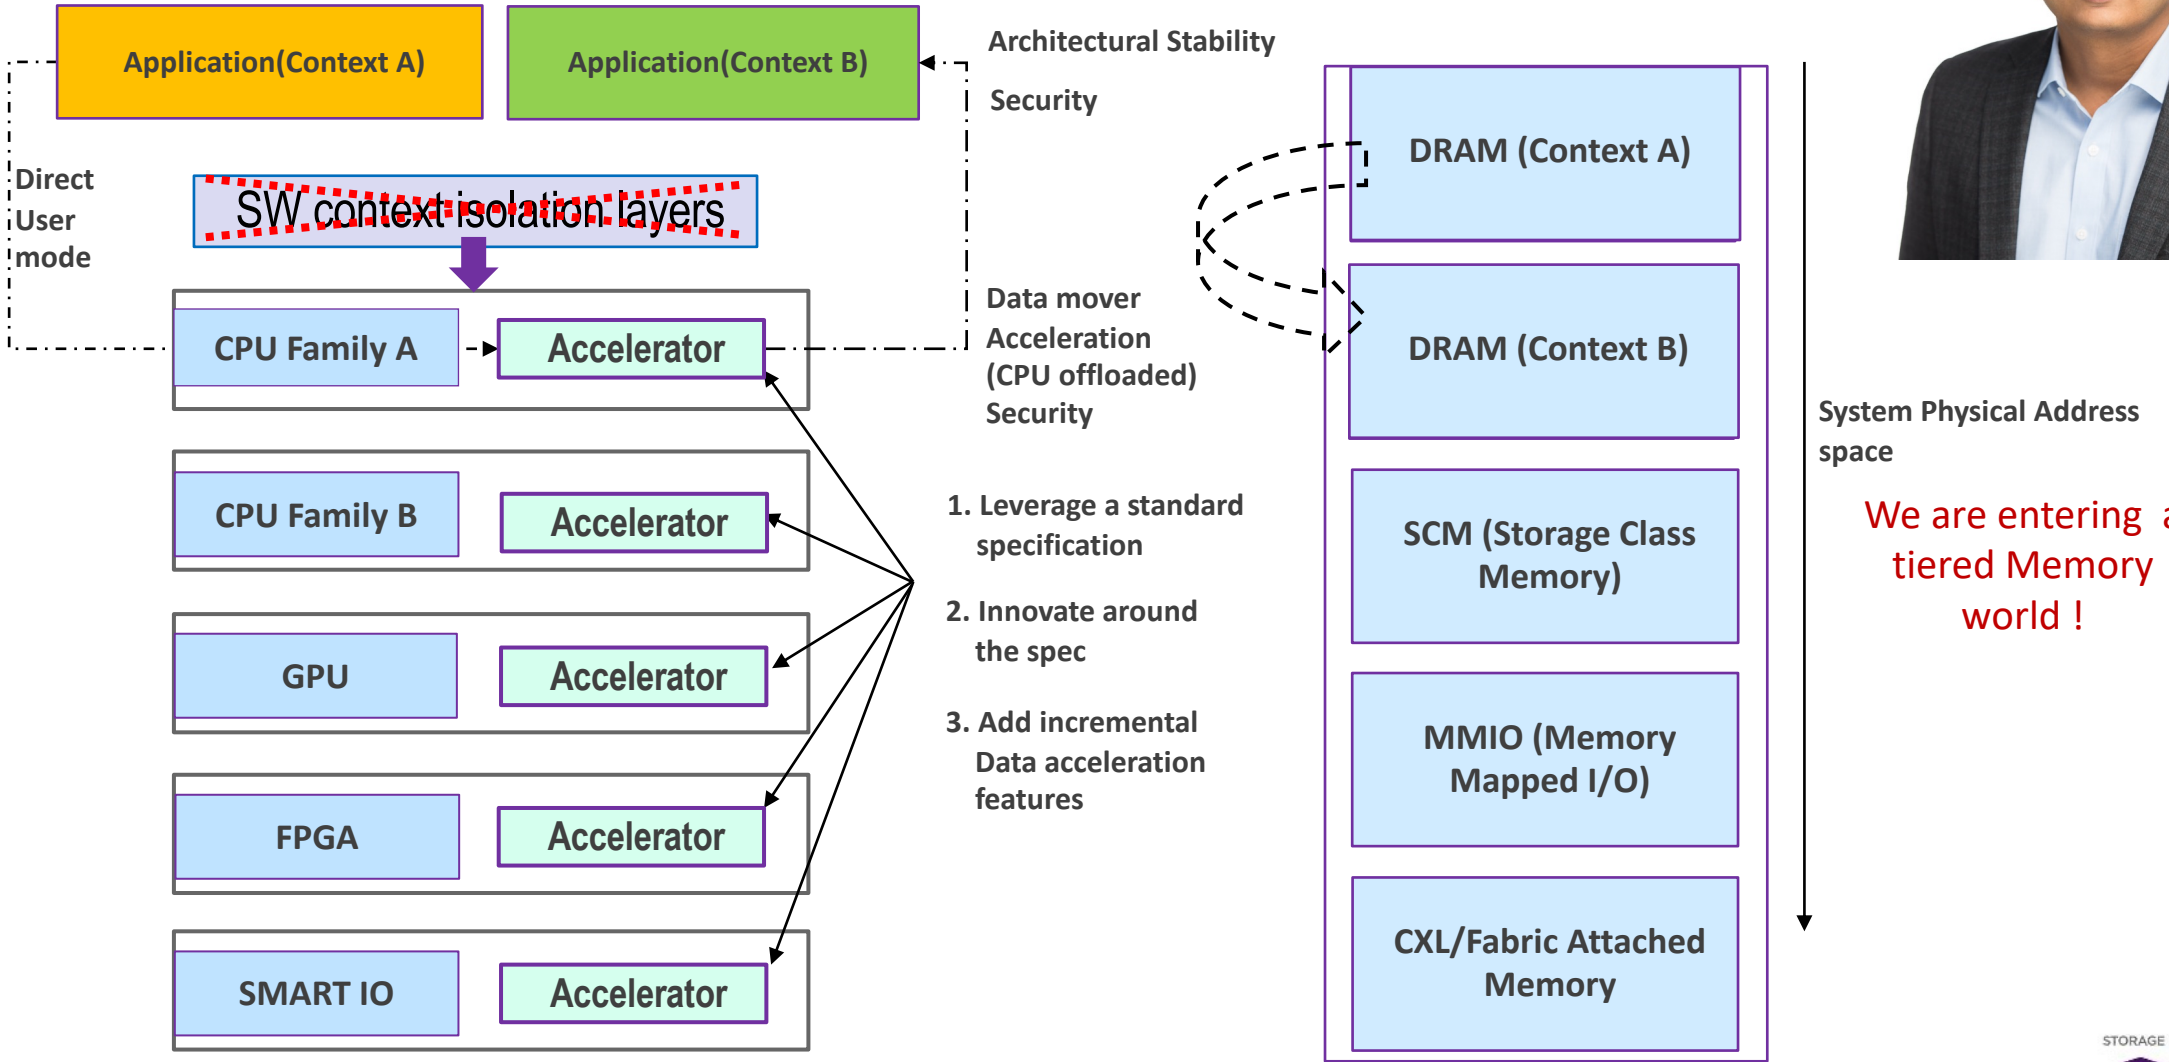


SDXI(Smart Data Accelerator Interface)

- Software memcopy is the current data movement standard
 - 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.
- 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 formed in June 2020 tasked to work on this proposed standard
 - 25 member companies, 60+ individual members



SDXI Memory-to-Memory Data Movement



SDXI 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.

Announcements

- **SDXI Specification v0.9-rev1 available for Public Review**
 - <https://www.snia.org/publicreview>
- **WEDNESDAY, SEPTEMBER 29TH, Birds of a Feather (BoF) Session**
 - 3:00 - 4:00 PM SDXI (Smart Data Accelerator Interface) discussion with SDXI TWG members
- **SNIA SDXI Page**
 - <https://www.snia.org/sdxi>

