

The logo for the Storage Networking Industry Association (SNIA), featuring the letters 'SNIA' in a bold, sans-serif font. To the left of the letters is a small square icon containing a stylized 'S' shape.

PERSISTENT MEMORY
PM SUMMIT

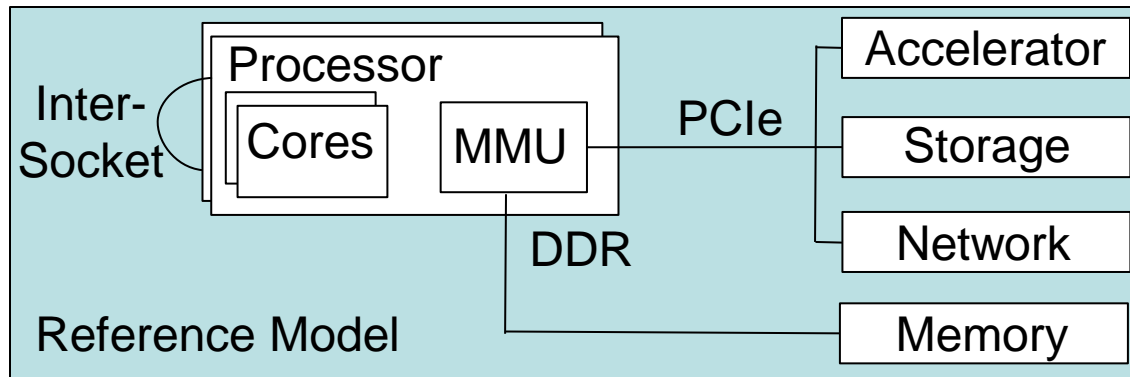
JANUARY 18, 2017 | SAN JOSE, CA

How Might Recently Formed System Interconnect Consortia Affect PM?

Doug Voigt, SNIA TC

Three Consortia Formed in Oct 2016

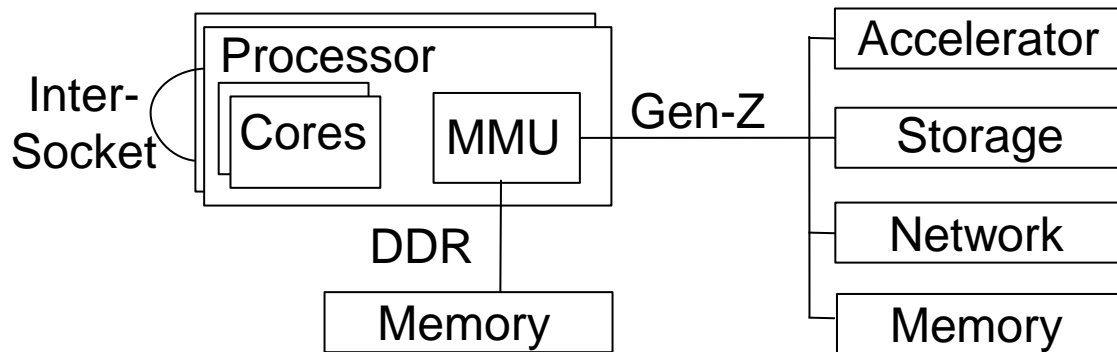
Gen-Z	Processor complex to rack scale memory fabric	❖ PM pooling
Open CAPI	Cache coherent accelerator access to processor memory	❖ Accelerator access to PM
CCIX	Cache coherency across processors and accelerators	❖ Shared PM



Gen-Z: Communication at the speed of memory

What is a Memory Semantic Fabric?

- Handles all communication as memory operations such as load/store, put/get and atomic operations typically used by a processor
- Memory semantics are optimal at sub-microsecond latencies from CPU load command to register store
- Unlike, storage accesses which are block based and managed by complex, code intensive, software stacks



Why Now?

- The emergence of low latency, Storage Class Memory (SCM) and the demand for large capacity, rack scale resource pools, and multi node architectures



Gen-Z: A New Data Access Technology

High Bandwidth Low Latency

- Memory Semantics - simple Reads and Writes
- From tens to several hundred GB/s of bandwidth
- Sub-100 ns load-to-use memory latency

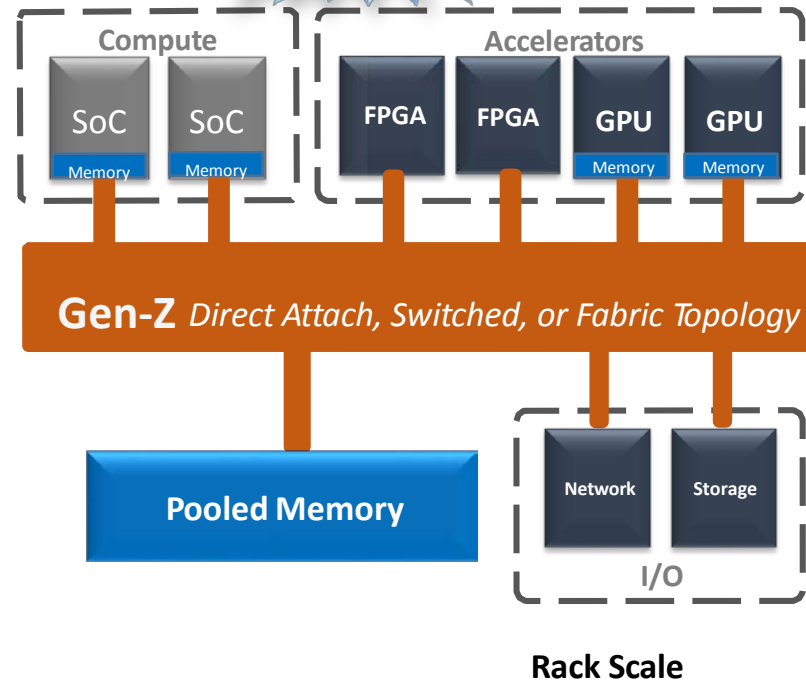
Advanced Workloads & Technologies

- Real time analytics
- Enables data centric and hybrid computing
- Scalable memory pools for in memory applications
- Abstracts media interface from SoC to unlock new media innovation

Secure Compatible Economical


- Provides end-to-end secure connectivity from node level to rack scale
- Supports unmodified OS for SW compatibility
- Graduated implementation from simple, low cost to highly capable and robust
- Leverages high-volume IEEE physical layers and broad, deep industry ecosystem

Open
Standard



Open CAPI (coherent accelerator processor interface)

Gives accelerators (e.g. GPU's) cache coherent access to memory

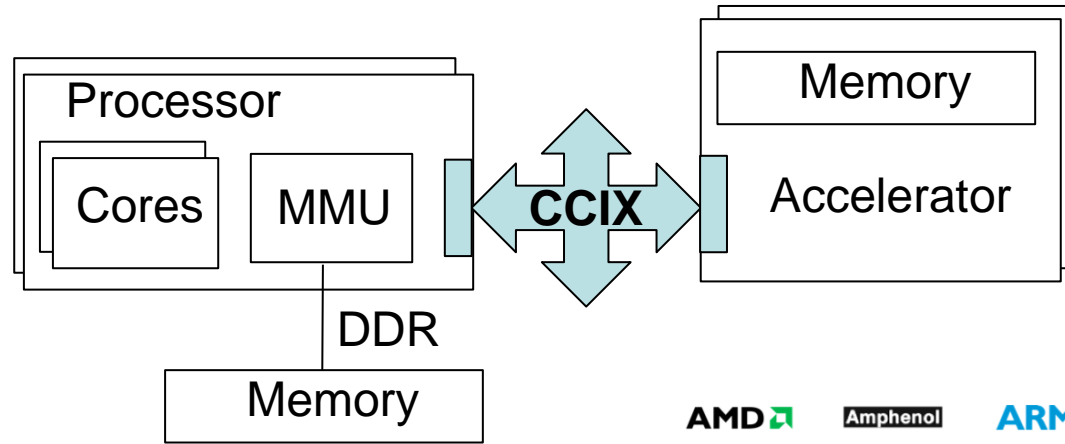


- PCIe or higher speed physical
- Accelerator operates within application address space
- Can be used for advanced IO or memory devices
- Founders: AMD, Google, IBM, Mellanox, and Micron

<http://www.opencapi.org>

- Accelerator Reads with no intent to cache, DMA write using Program Addresses
 - The accelerator is working in the same address domain as the host application (Pointer chasing, link lists are all now possible without Device Driver involvement)
 - Address translation on host (processor) with error response back to the accelerator (Efficient translation latency mechanism using host processor Address Translation Cache and MMU)
 - Non-posted writes only
 - Ability for Partial Read/Write DMAs
- Translate touch to warm up address translation caches (Allows accelerator to reduce translation latency when using a new page)
- Wake Up host thread (Efficient low latency mechanism in lieu of either interrupts or host processor polling mechanism of memory)
- Atomic Memory Operations (AMO) to Host Processor Memory (Accelerator can use atomic operations in a coherent domain just like any other host processor thread)
- Open CAPI attached memory

CCIX (Cache Coherent Interconnect for Accelerators)

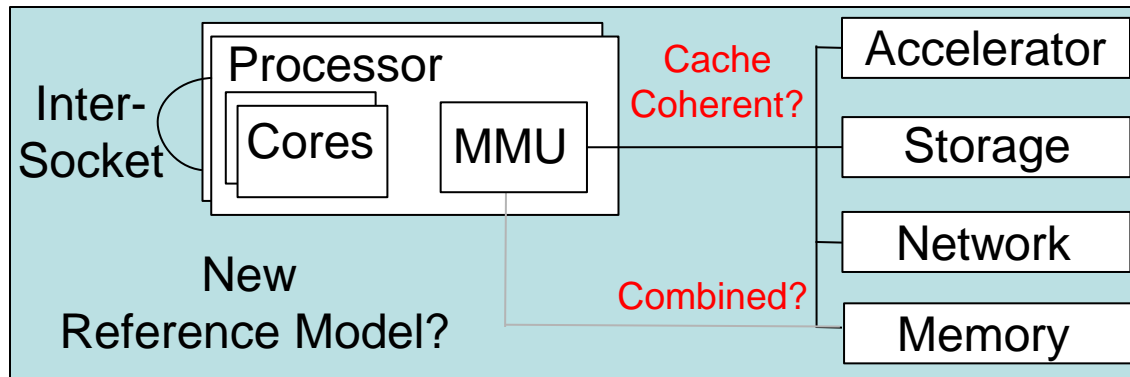


- Shared address space
- PCIe physical
- Processor and accelerator cache coherent transactions converted to CCIX
- Accelerator memory in coherency domain
- Can be used for advanced IO or memory



Three Consortia Formed in Oct 2016

Gen-Z	Processor complex to rack scale memory fabric	❖ PM pooling
Open CAPI	Cache coherent accelerator access to processor memory	❖ Accelerator access to PM
CCIX	Cache coherency across processors and accelerators	❖ Shared PM



The logo for the Storage Networking Industry Association (SNIA), featuring the letters 'SNIA' in a bold, sans-serif font. To the left of the letters is a small square icon containing a stylized 'S' shape.

PERSISTENT MEMORY
PM SUMMIT

JANUARY 18, 2017 | SAN JOSE, CA

How Might Recently Formed System Interconnect Consortia Affect PM?

Doug Voigt, SNIA TC