





# Gen-Z Technology: Enabling Memory Centric Architecture

## GROWTH IN THE INTERNET OF THINGS

THE NUMBER OF CONNECTED DEVICES WILL EXCEED 50 BILLION BY 2020




# Why Gen-Z?



## Businesses' Need to 'Monetize' Data

### Worldwide Total Data Revenue by Segment (\$M) 2014-2019




Source: 451 Research Market Monitor





Big Data  
AI  
Machine Learning  
Deep Learning  
BI

## Need Answers ... FAST!




- Businesses demanding real-time insight
- Increasing amounts of data to be analyzed


# Challenge: Multiple Memory Types Need Support



Processor  
1 ns



SRAM  
5 ns



DRAM  
100 ns

- Byte Addressability
- High Bandwidth/Low Latency
- Simple & Efficient Load/Store operations



Flash  
200,000 ns



Hard Drive  
10,000,000 ns




Tape  
40,000,000,000 ns

- Data Protection methods (RAID, EC)
- Device management: Hotplug-ability
- Rich Data Services (ie deduplication, replication, thin provisioning, etc)

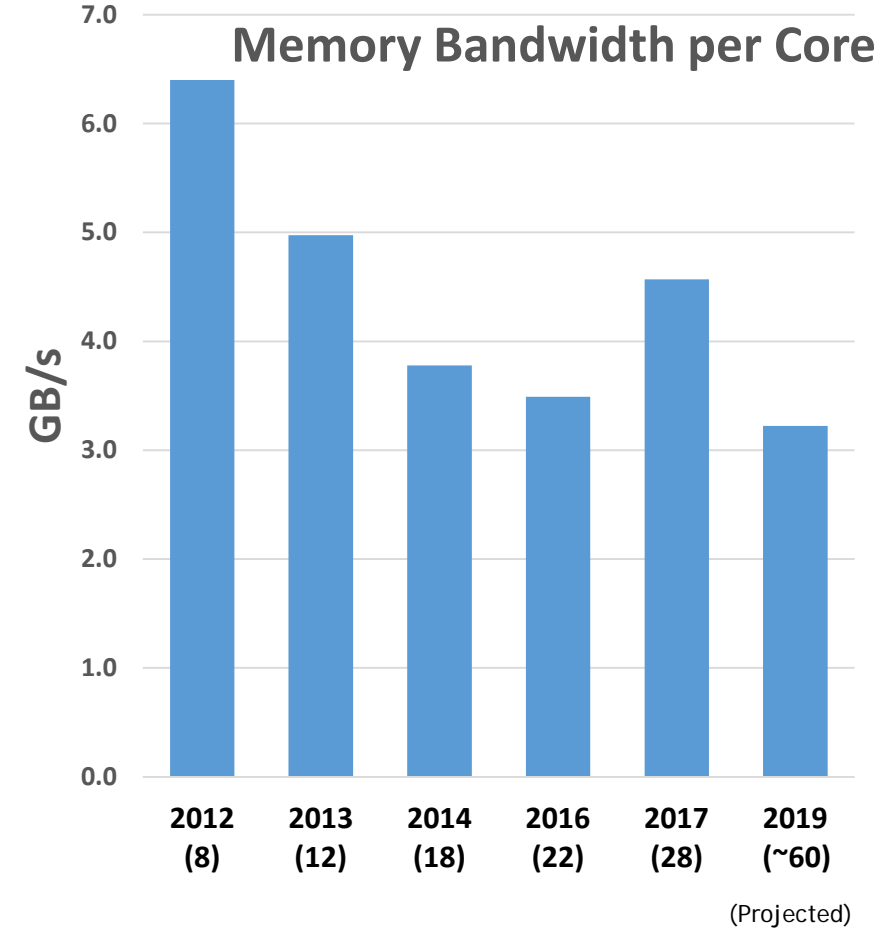
# Challenge: Memory Bandwidth Per Core Is Shrinking

## # Processor Pins & DDR Channels



Courtesy: Ideas International

## Memory Bandwidth per Core




Pin limitations lead to memory channel limitations that leads to memory bandwidth limitations.





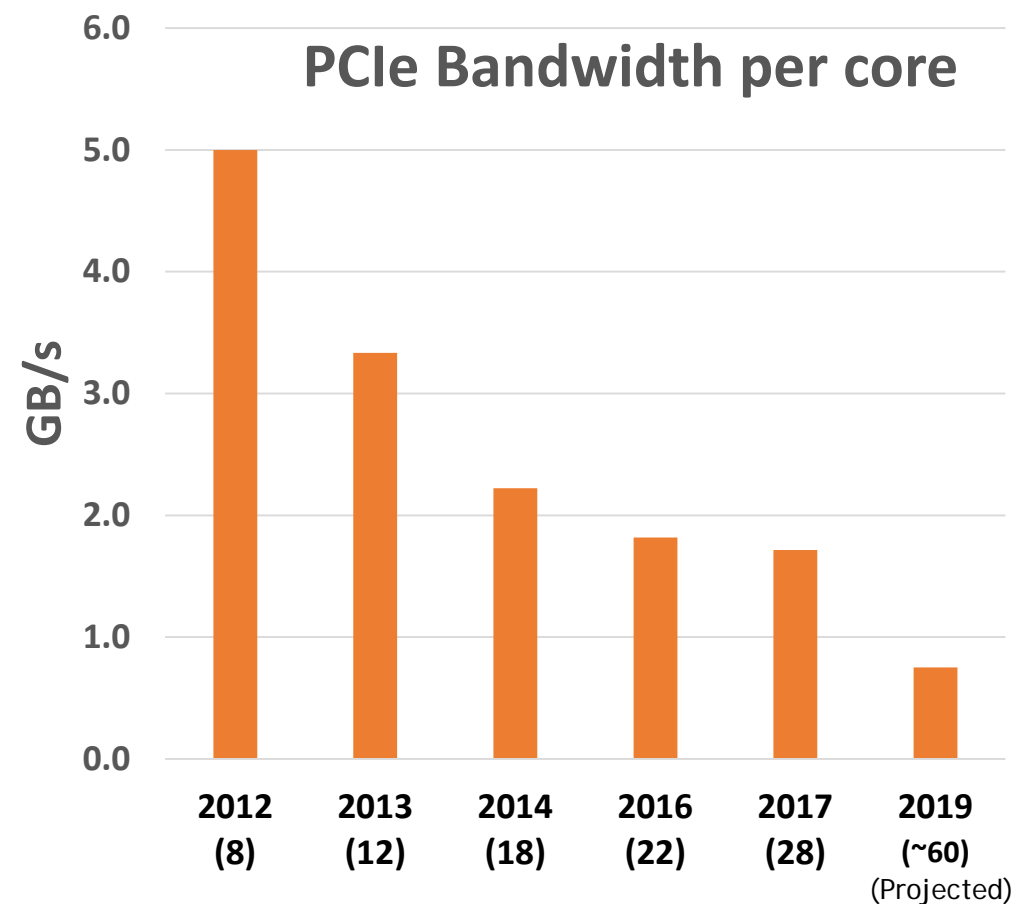
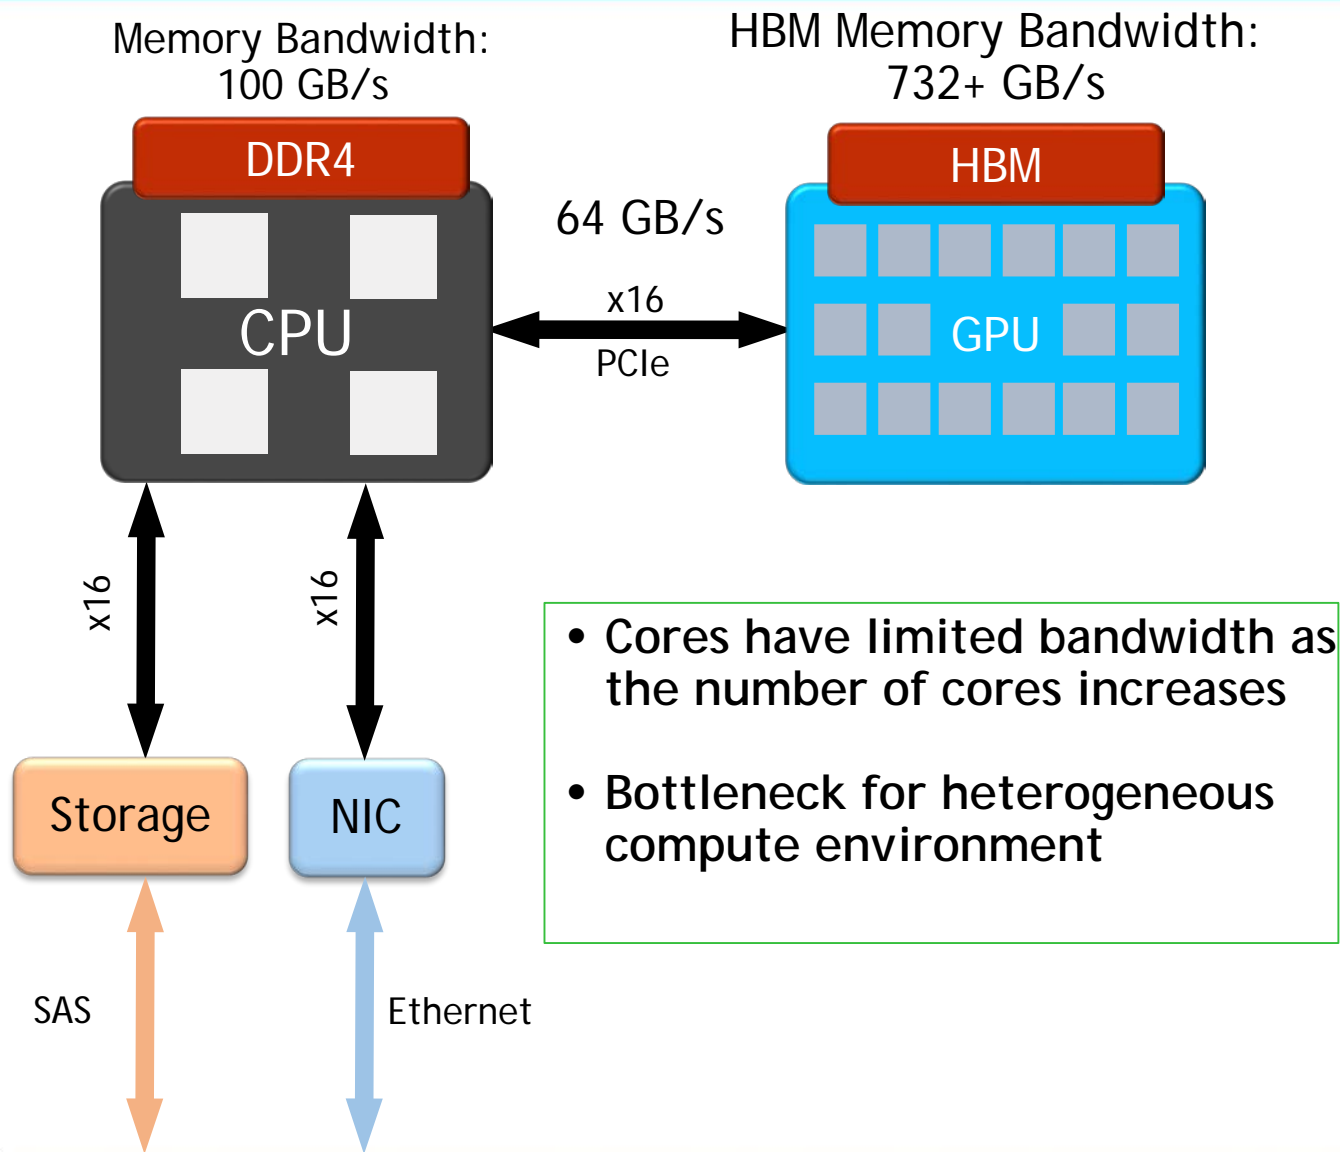
## Power, Cooling, and Physical Limitations

- CPU power consumption is increasing
- DIMM power consumption is increasing
- Number of memory channels is slowly increasing
- Physical size of CPU is increasing
- Physical size of DIMM is increasing
- Speed of Memory Channel is increasing
- Number of Cores in each Socket is increasing
- Dataset Memory Requirements are increasing





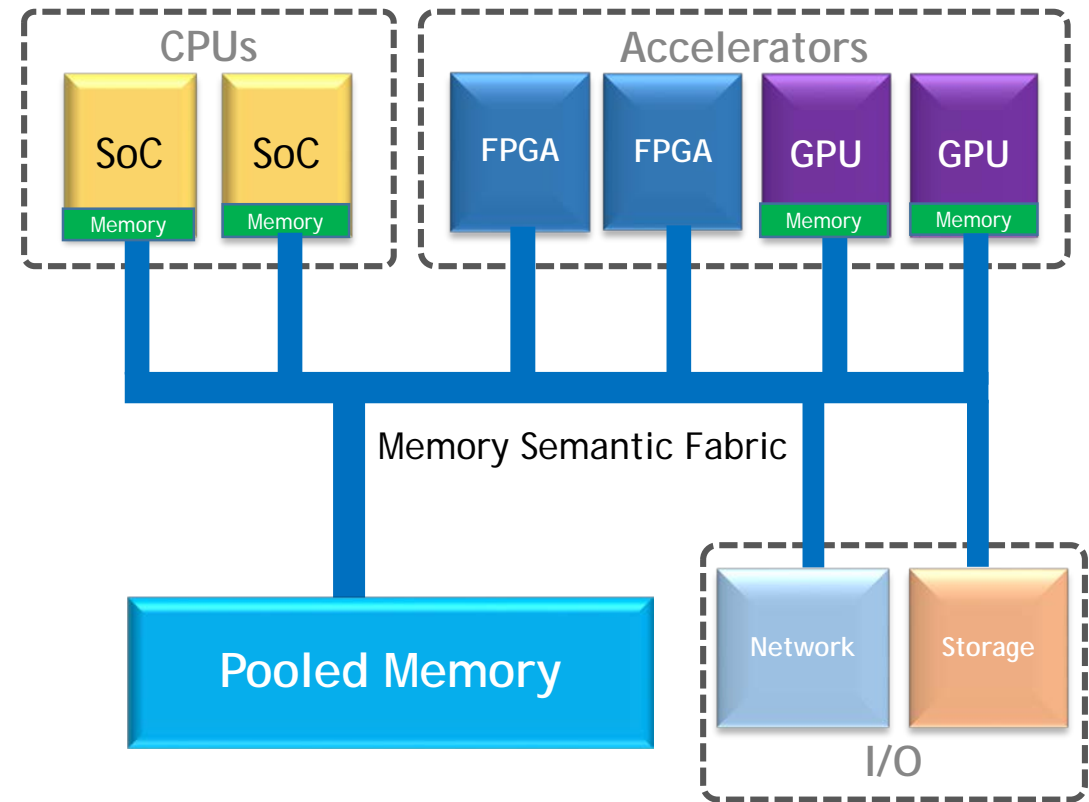
# Challenge: IO Bandwidth Per Core is Shrinking



## A Memory Semantic Fabric!

### What is a Memory Semantic Fabric?

- A communication protocol that speaks the same language the CPU speaks: load/store, put/get, and read/write
- Connectivity that extends beyond the server to the rack

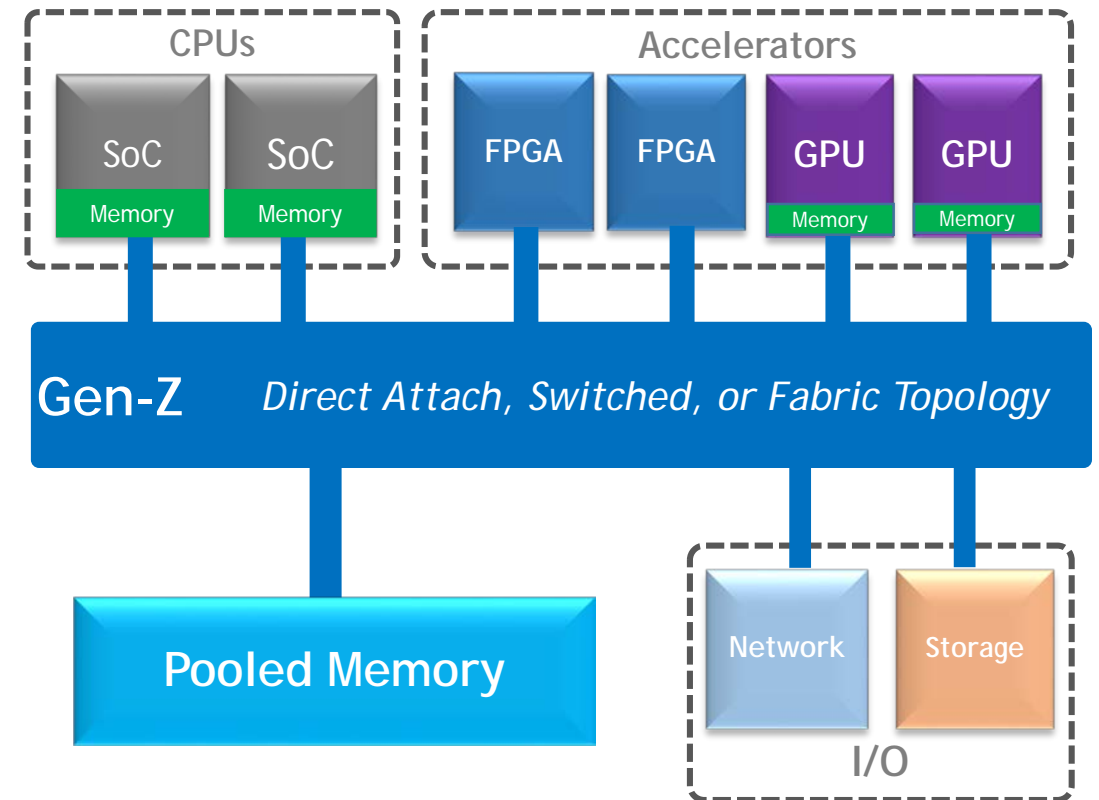


Communication at the speed of memory

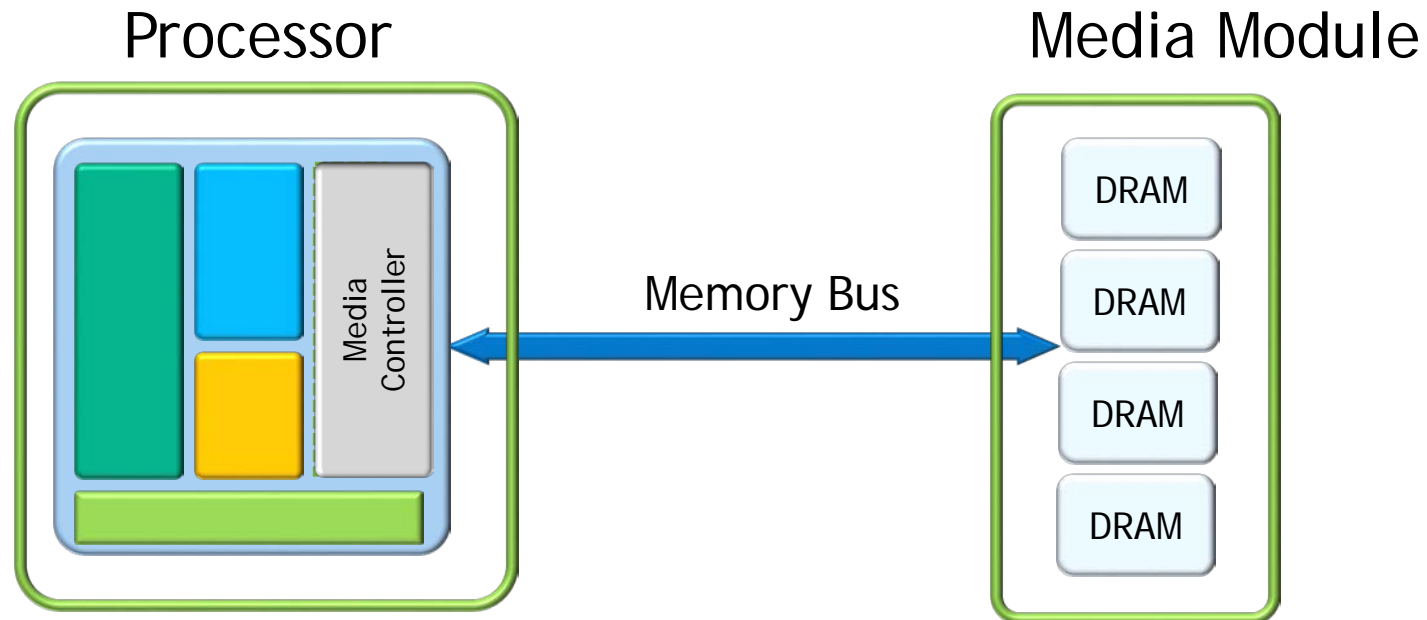
# Gen-Z: A New Data Access Technology



- High Bandwidth, Low Latency
- Packetized Memory Semantic Protocol
- Abstracted Hardware Interface
- Simplify HW/SW Boundary
- Compatible & Economical
- Universal democratized communication




# Gen-Z Allows Memory Innovation



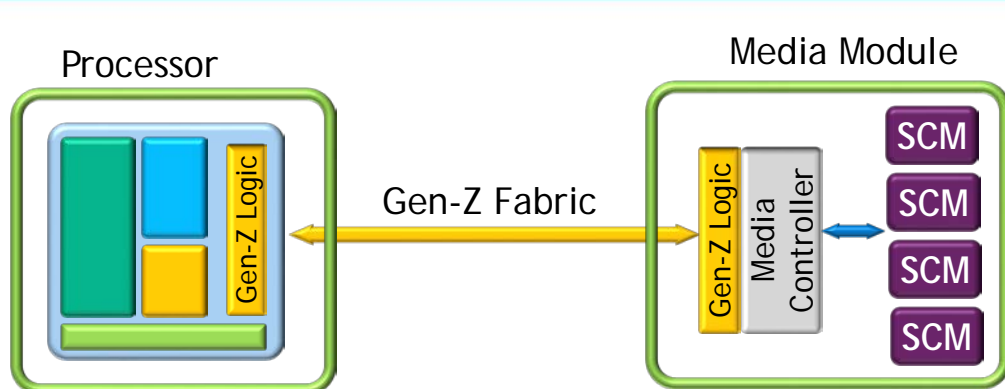
4-8 Memory Channels  
17-25 GB/s / Channel  
288pins / DIMM  
Synchronous Interface

# Gen-Z Allows Memory Innovation




2-8 High-speed Serial Links  
Low Latency, High performance  
Split Memory Controller  
Asynchronous Interface  
Processor is media agnostic

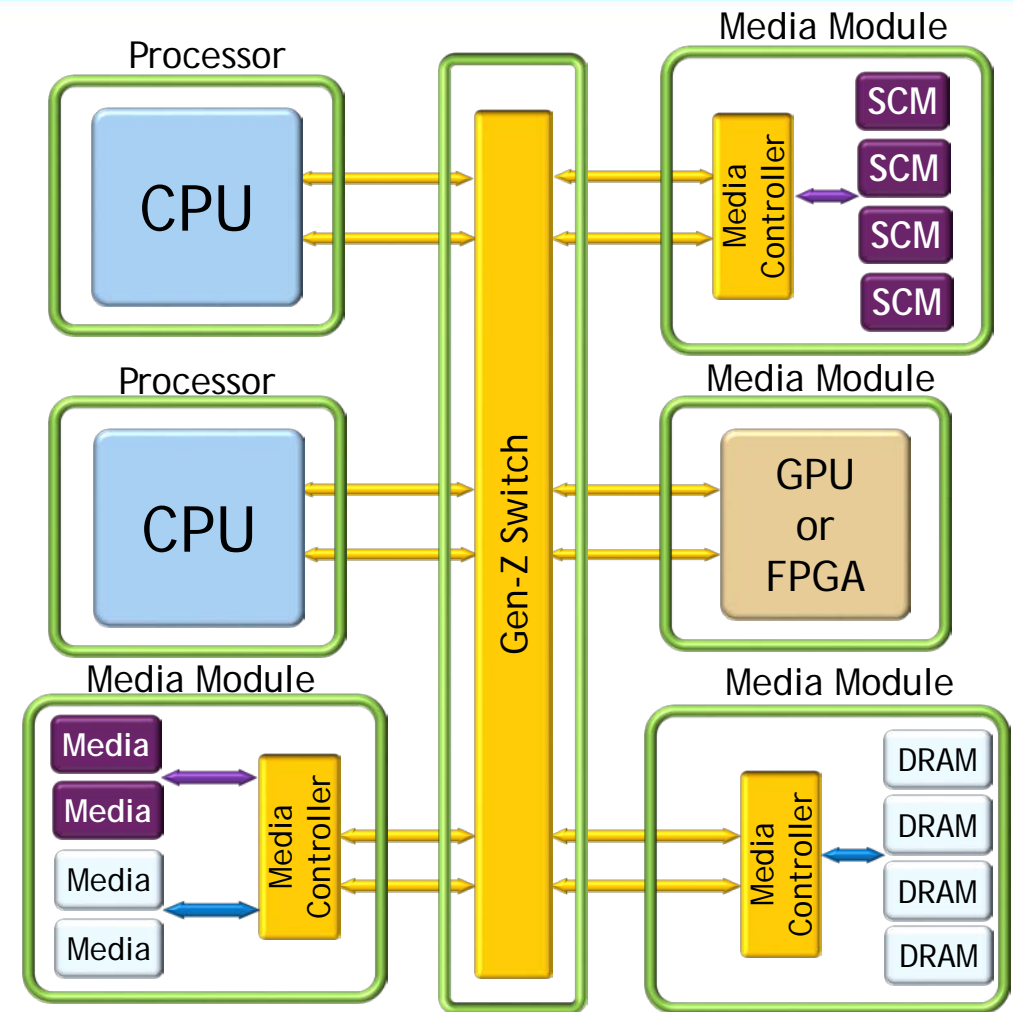
# Gen-Z is Media Agnostic & Composable



Storage Class Memory



Accelerator

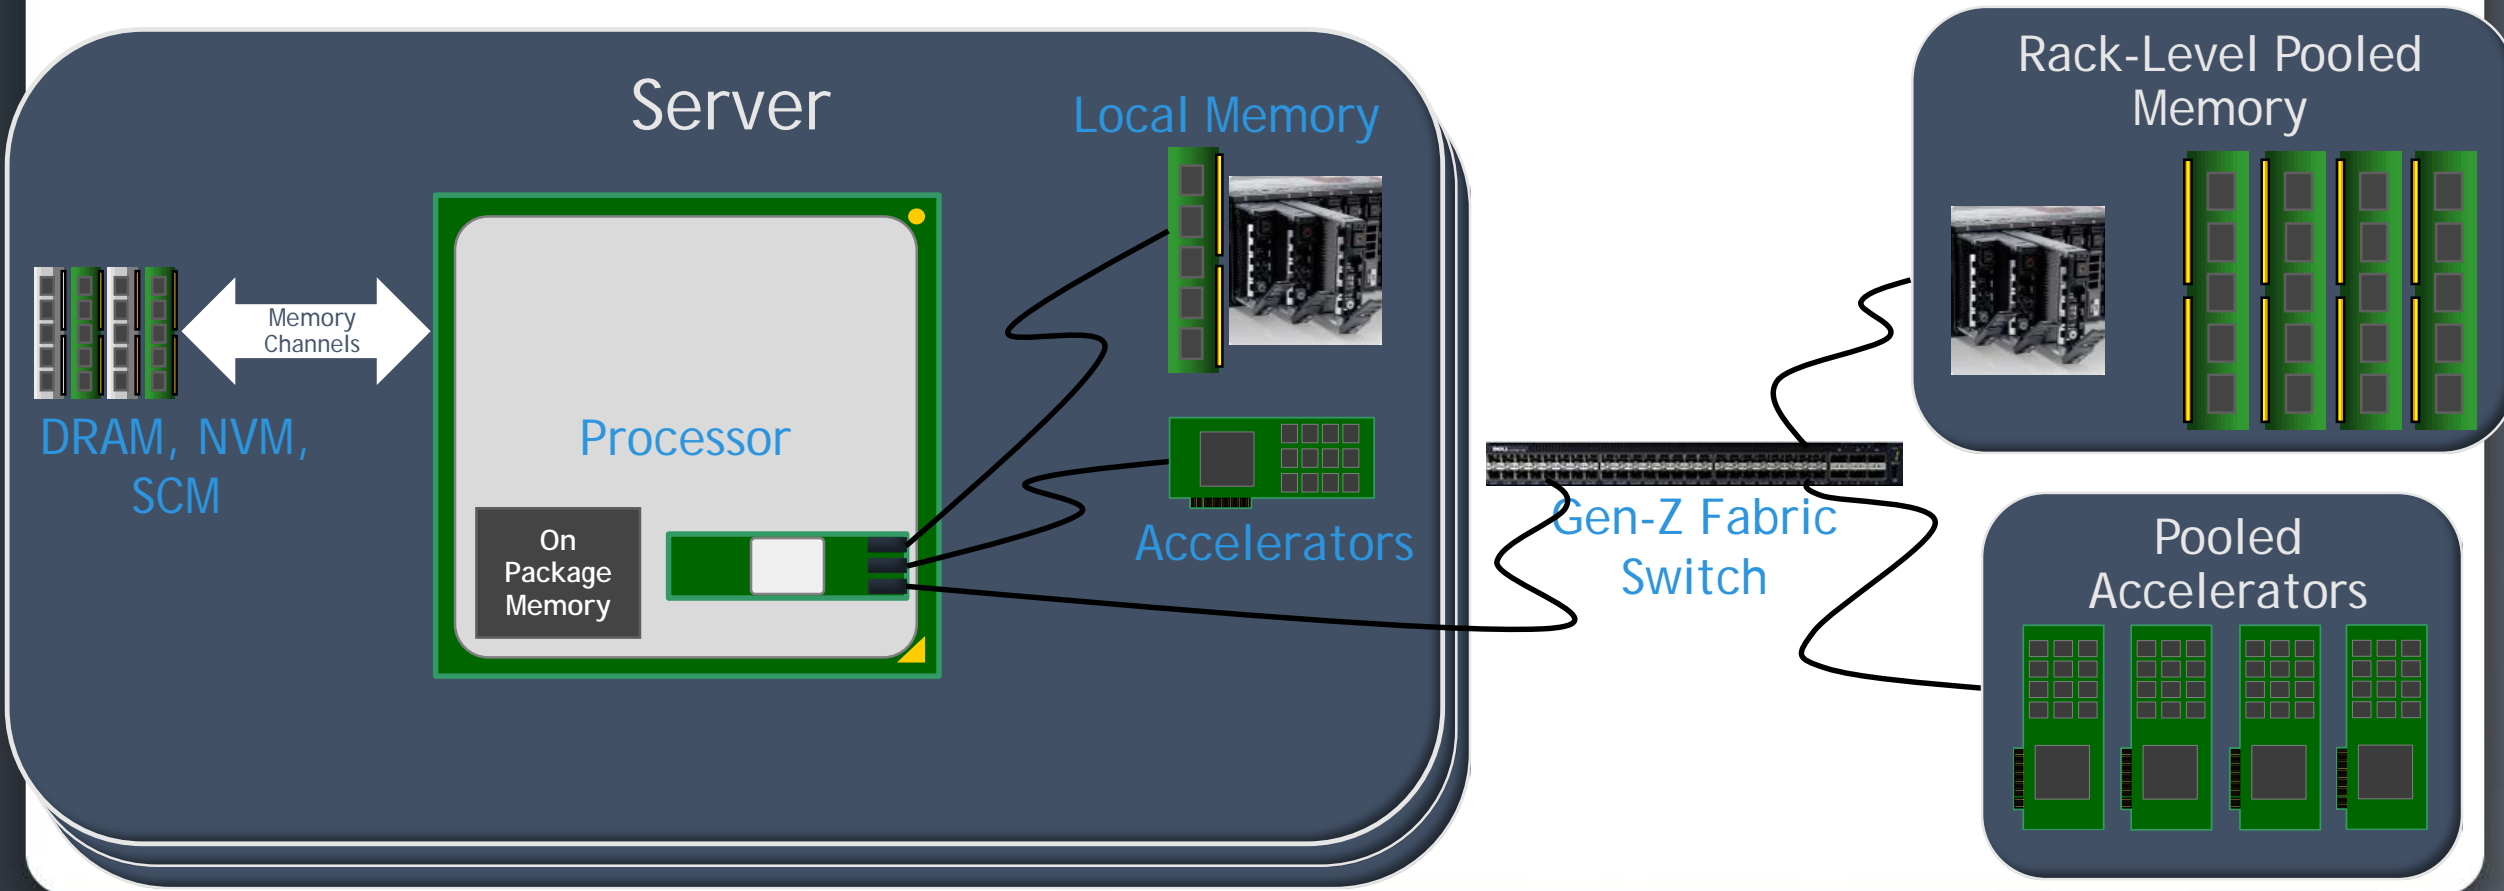


Multiple resources enabled by Universal Interconnect



# Gen-Z Creates More Real Estate

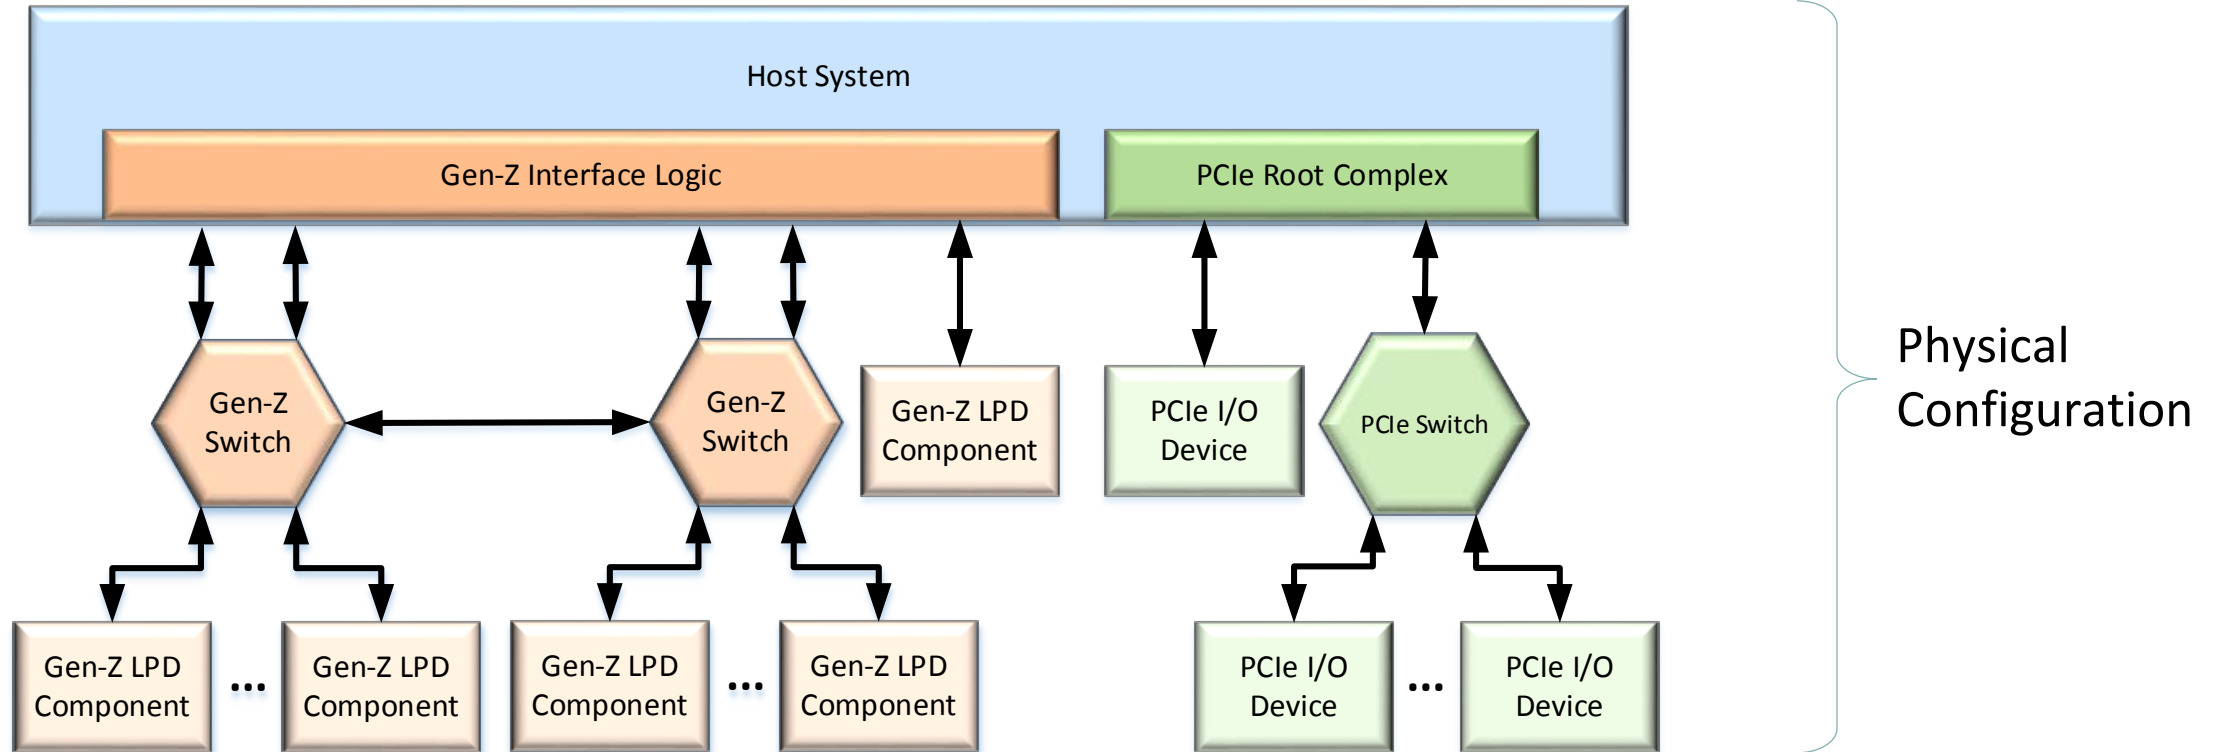
- Gen-Z can be used to connect to devices inside and outside the server enabling a composable system
- Gen-Z's pin efficient interface and ability to multiplex with PCIe pins allows for an overall increase in memory bandwidth



# Fully Composable Through Disaggregation Using Gen-Z



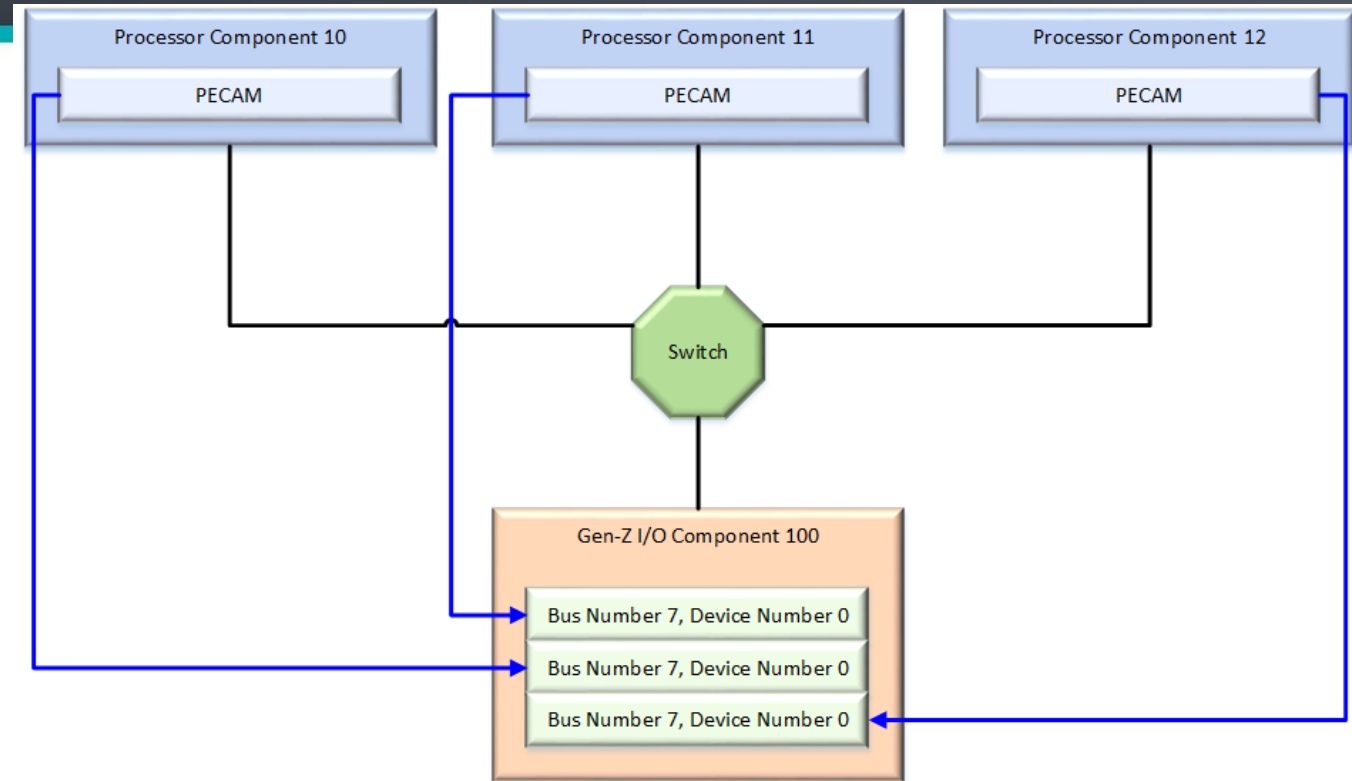
# Example Gen-Z System with Gen-Z & PCIe® I/O



- Gen-Z systems may contain a mix of Gen-Z I/O components and PCI Express® I/O devices
- *PCIe worked immediately with unmodified OSs by appearing to SW as conventional PCI devices & bridges*
- Gen-Z is using a similar strategy via the concept of **logical PCI devices (LPDs)**

# I/O Component Sharing

- With suitable Gen-Z fabric management SW, a single Gen-Z I/O component can be shared among multiple Gen-Z systems
- Component sharing offers significant availability benefits, e.g., system failover for continued operation and avoidance of stranded resources
- Each Gen-Z system running an unmodified OS only “sees” LPDs assigned to it by Gen-Z fabric management SW
- Each Gen-Z system’s BDF space for LPDs is orthogonal to the BDF space in other systems
- A Gen-Z I/O component designed to be shared by multiple systems must ensure that the virtual device instances do not interfere with each other
- Gen-Z can deliver the major benefits of PCIe MR-IOV\* with less cost & complexity



\*MR-IOV: Multi-Root I/O Virtualization and Sharing



- Gen-Z's 112GT/s Connector

- The connector supports vertical, horizontal, right angle and cabled solutions
- Same connectors for memory, accelerators, I/O, etc.
- Gen-Z defined the connector and is donating it to SFF

- Leverage High Volume Industry PHYs


- Derivative of the IEEE 802.3 PHY
- PCIe PHY
- Lower development and product costs

- Unmodified OS Support

- Minimize cost & effort to adopt Gen-Z solutions

- Enables Right Sized Solutions

- Enables the right amount of compute, memory, storage, acceleration, and networking to be applied to each workload



# Open With Broad Industry Support



## GEN Z Consortium Members

- Alpha Data
- AMD
- Amphenol
- ARM
- Avery Design Systems
- Broadcom
- Cadence
- Cavium
- Cray
- Dell EMC
- Everspin
- FIT
- HP Enterprise
- Huawei
- IBM
- IDT
- IntelliProp
- Jabil
- Lenovo
- Lotes
- Luxshare-ICT
- Mellanox
- Mentor Graphics
- Micron
- Microsemi
- Molex
- NetApp
- Nokia
- Numascale
- PLDA Group
- Red Hat
- Samsung
- Seagate
- SK hynix
- Smart Modular Technologies
- Spin Transfer Technologies
- TE
- Tyco Electronics
- VMware
- Western Digital
- Xilinx
- Yadro





- Scalable system interconnect and protocol
- Optimized for memory-semantic communications
- Breaks processor-memory innovation dependency
- Opportunity to simplify software overhead and complexity
- Unmodified OS support
- Common modular connector and mechanical form factors
- Visit the Gen-Z booth, #739, for more information and see the demonstrations
- Attend “Chat with the Experts” tonight to speak with Mike Krause about Gen-Z at the Fabrics Table
- Visit us at [genzconsortium.org](http://genzconsortium.org)

