Flash Memory Summit

GEN Z

The Best Interface for Emerging Memory Technologies

Presented by: Valerie Padilla, Technology Strategist in the Server CTO, Dell EMC

Gen-Z Approach

Several key factors are impacting infrastructure architectures:

- STRETCHED TO CAPACITY
- DATA EXPLOSION
- MEMORY/STORAGE CONVERGENCE
- WHAT'S THE SOLUTION?

Data Trends & Challenges

EXPLOSIVE GROWTH OF DATA

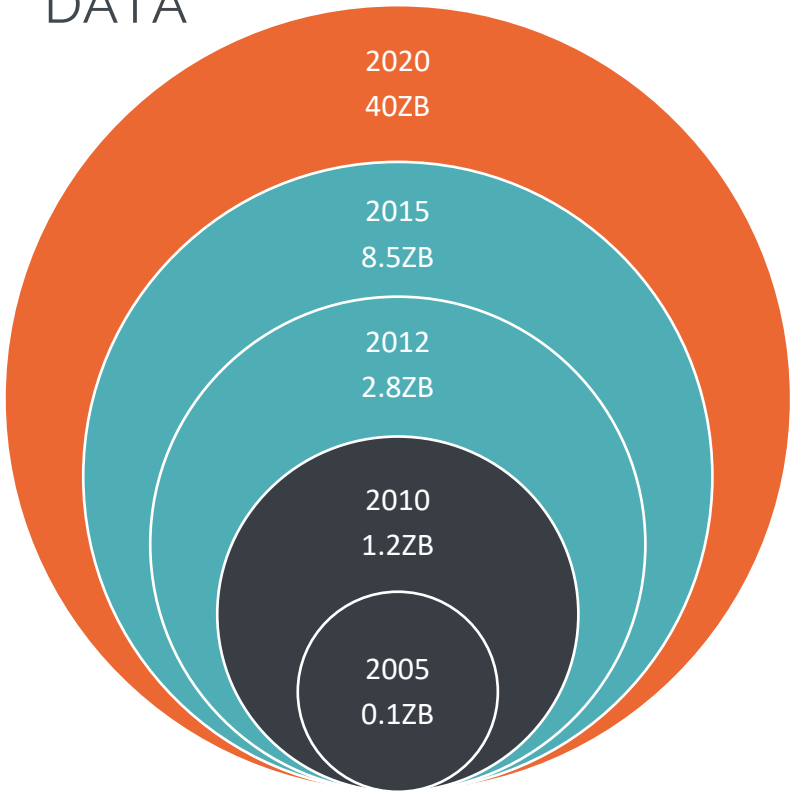

More than 37% of total data generated in 2020 (40 ZB) will have significant business value.

A SOLUTION IS NEEDED

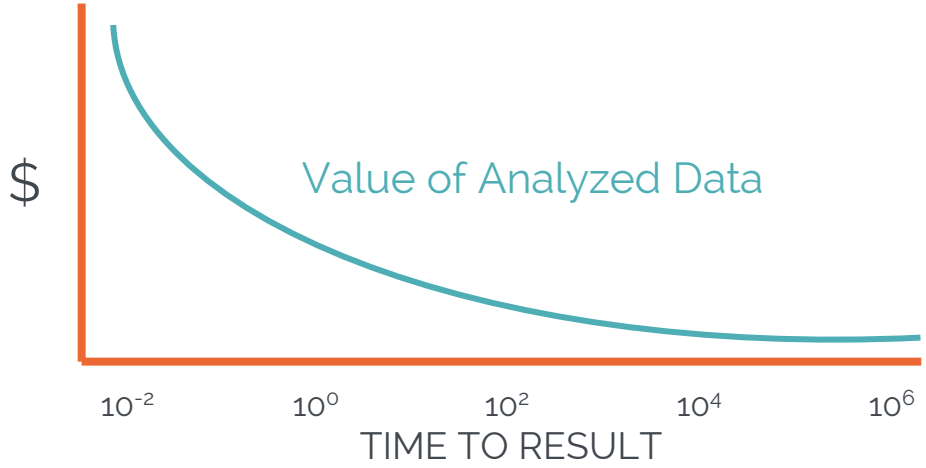

Increasing amounts of data to be analyzed & businesses demand real-time insight.

Infrastructure Maxed Out

Two-socket server architecture is stretched to its limit

- Relentless drive for performance
- Greater core counts
- Processors are more complex than ever
  - DDR data rates & DDR channel counts are increasing
  - DIMMs per channel are decreasing
  - PCIe data rates & lane counts are increasing
  - Storage data rates & lane counts are increasing
- More complex system designs drive up costs

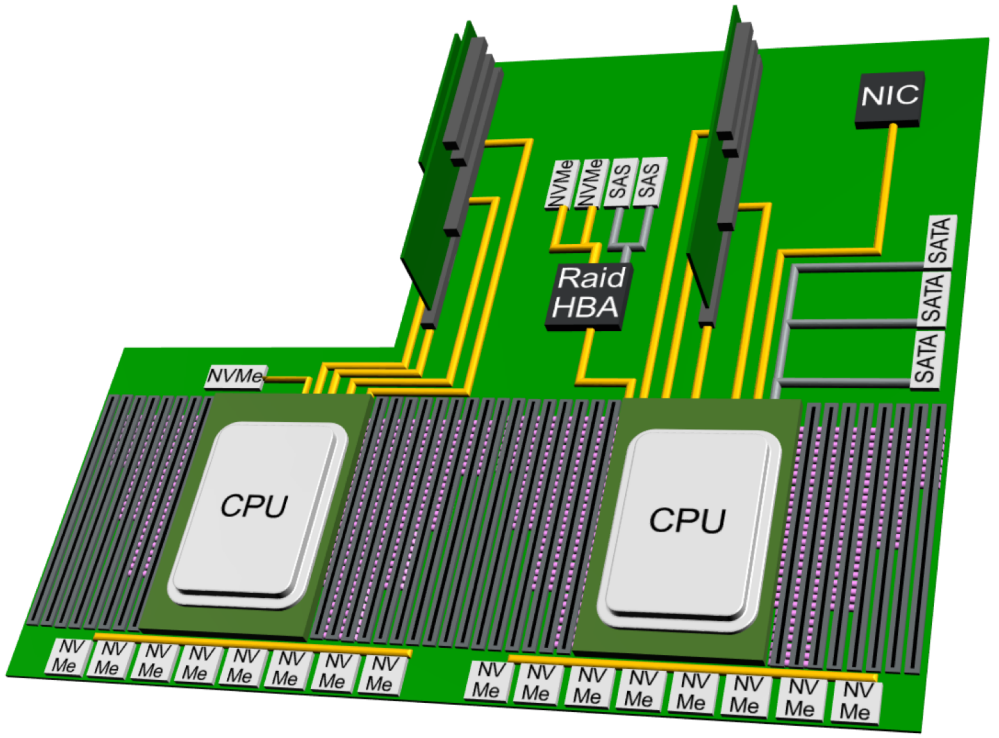

SATA GEN 4 PCIE GEN 5 CCIX OPENCAPI

DDR5 SAS GEN 4 PCIE GEN 4 CXL

SATA GEN 3 SAS GEN 3 DDR4 PCIE GEN 2

SATA GEN 2 PCIE GEN 3 SAS GEN 2

NVLINK NVME INTER-PROCESSOR LINKS

(future protocols)

Infrastructure Maxed Out

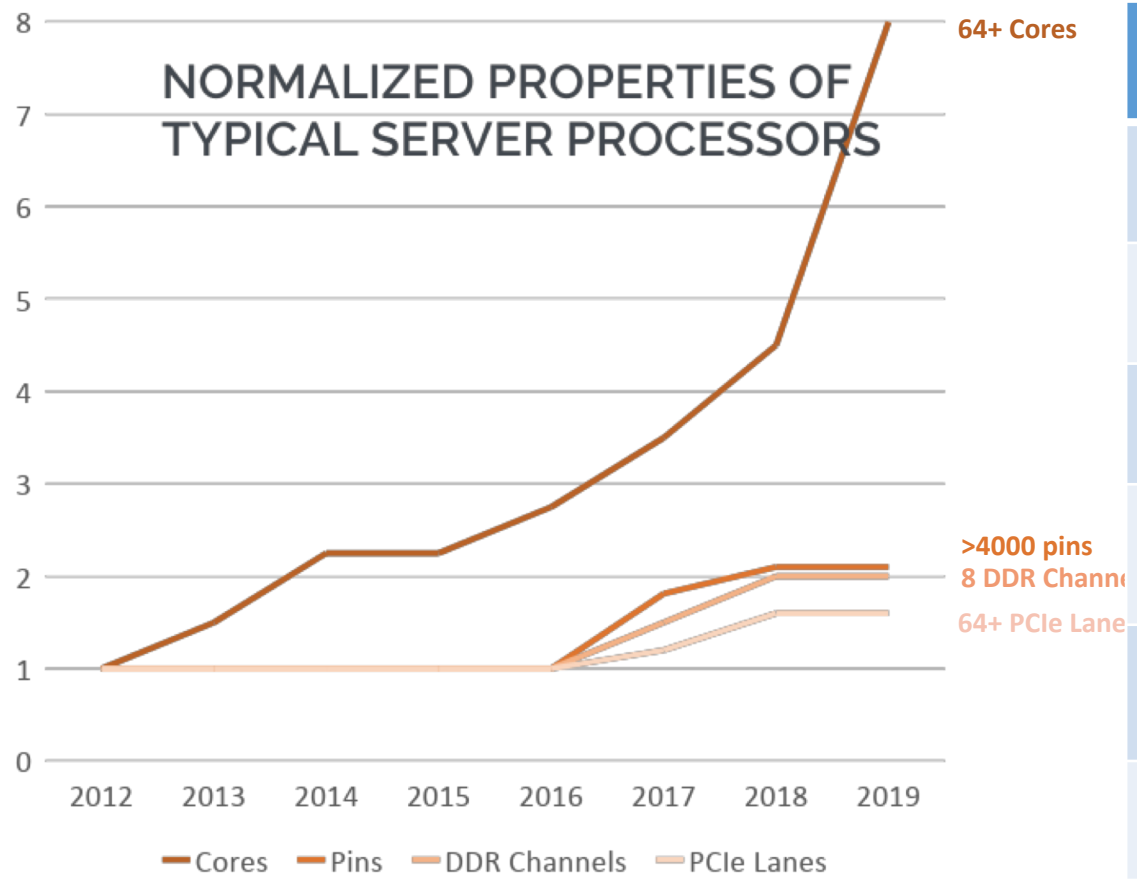

Cores per socket	
IO bandwidth per core	⇩
Memory capacity per core	⇧
Cache capacity per core	⇔
Memory bandwidth per core	⇩
Memory bandwidth per IPC per core	⇩
Local idle memory latency	⇧

Additional Challenges

- Infrastructure networks & fabrics have challenges
- Still maintain a wide variety of networks/fabrics
  - Each has its special purposes
- None are designed to support native CPU load/store
  - To support memory/storage convergence

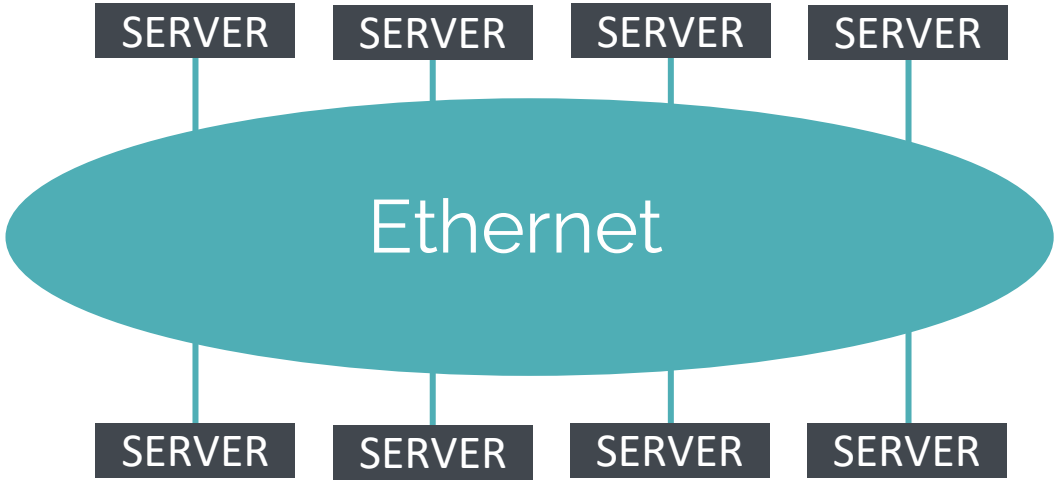

Memory/Storage Convergence

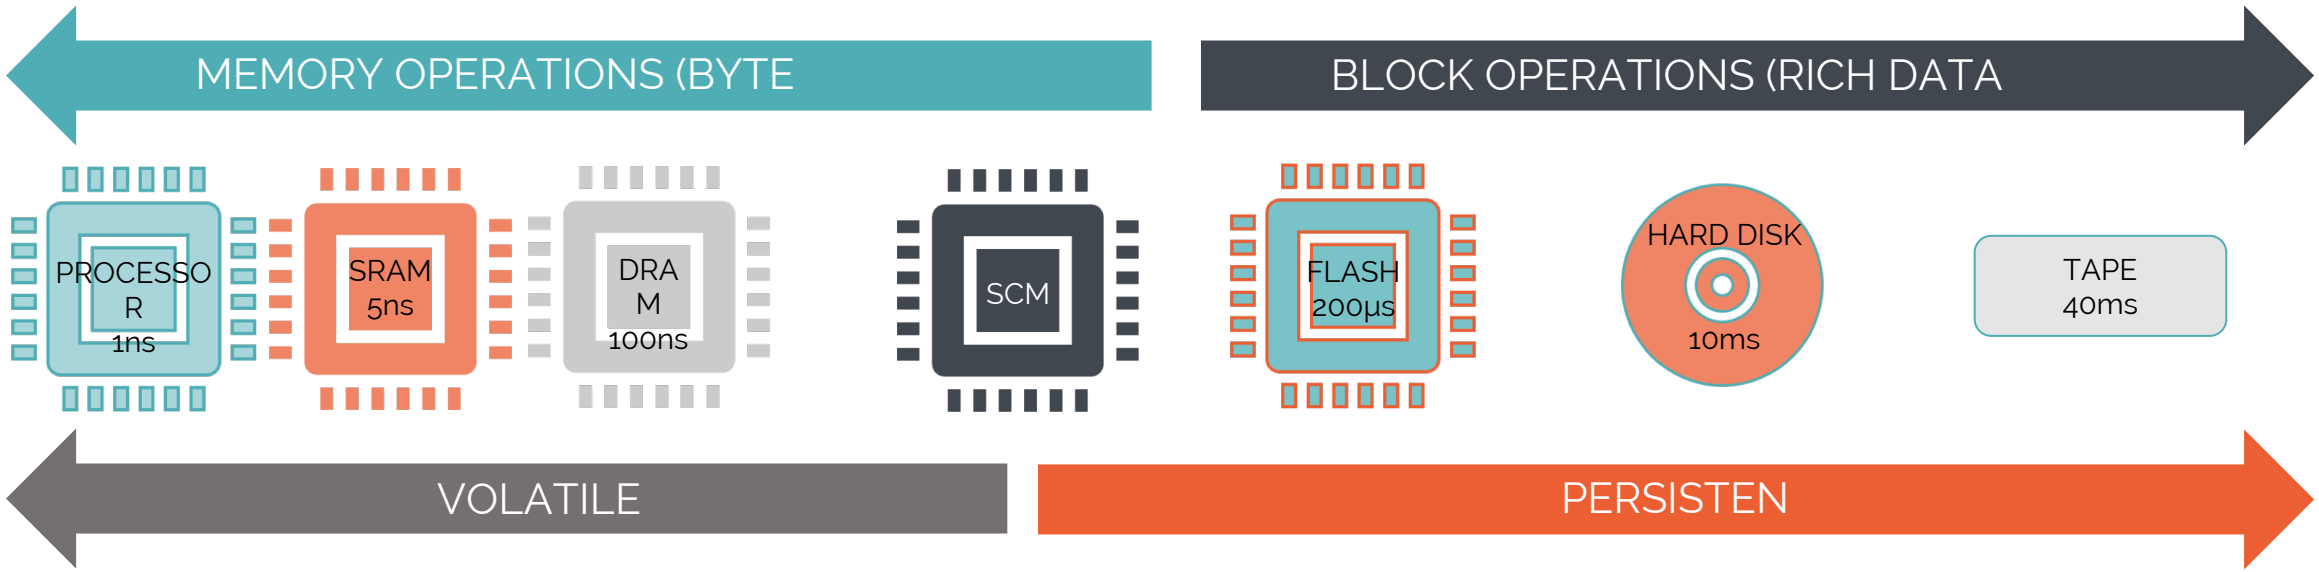

With memory/storage convergence, memory semantic operations become predominant (volatile & non-volatile).

Processor to Memory Centric

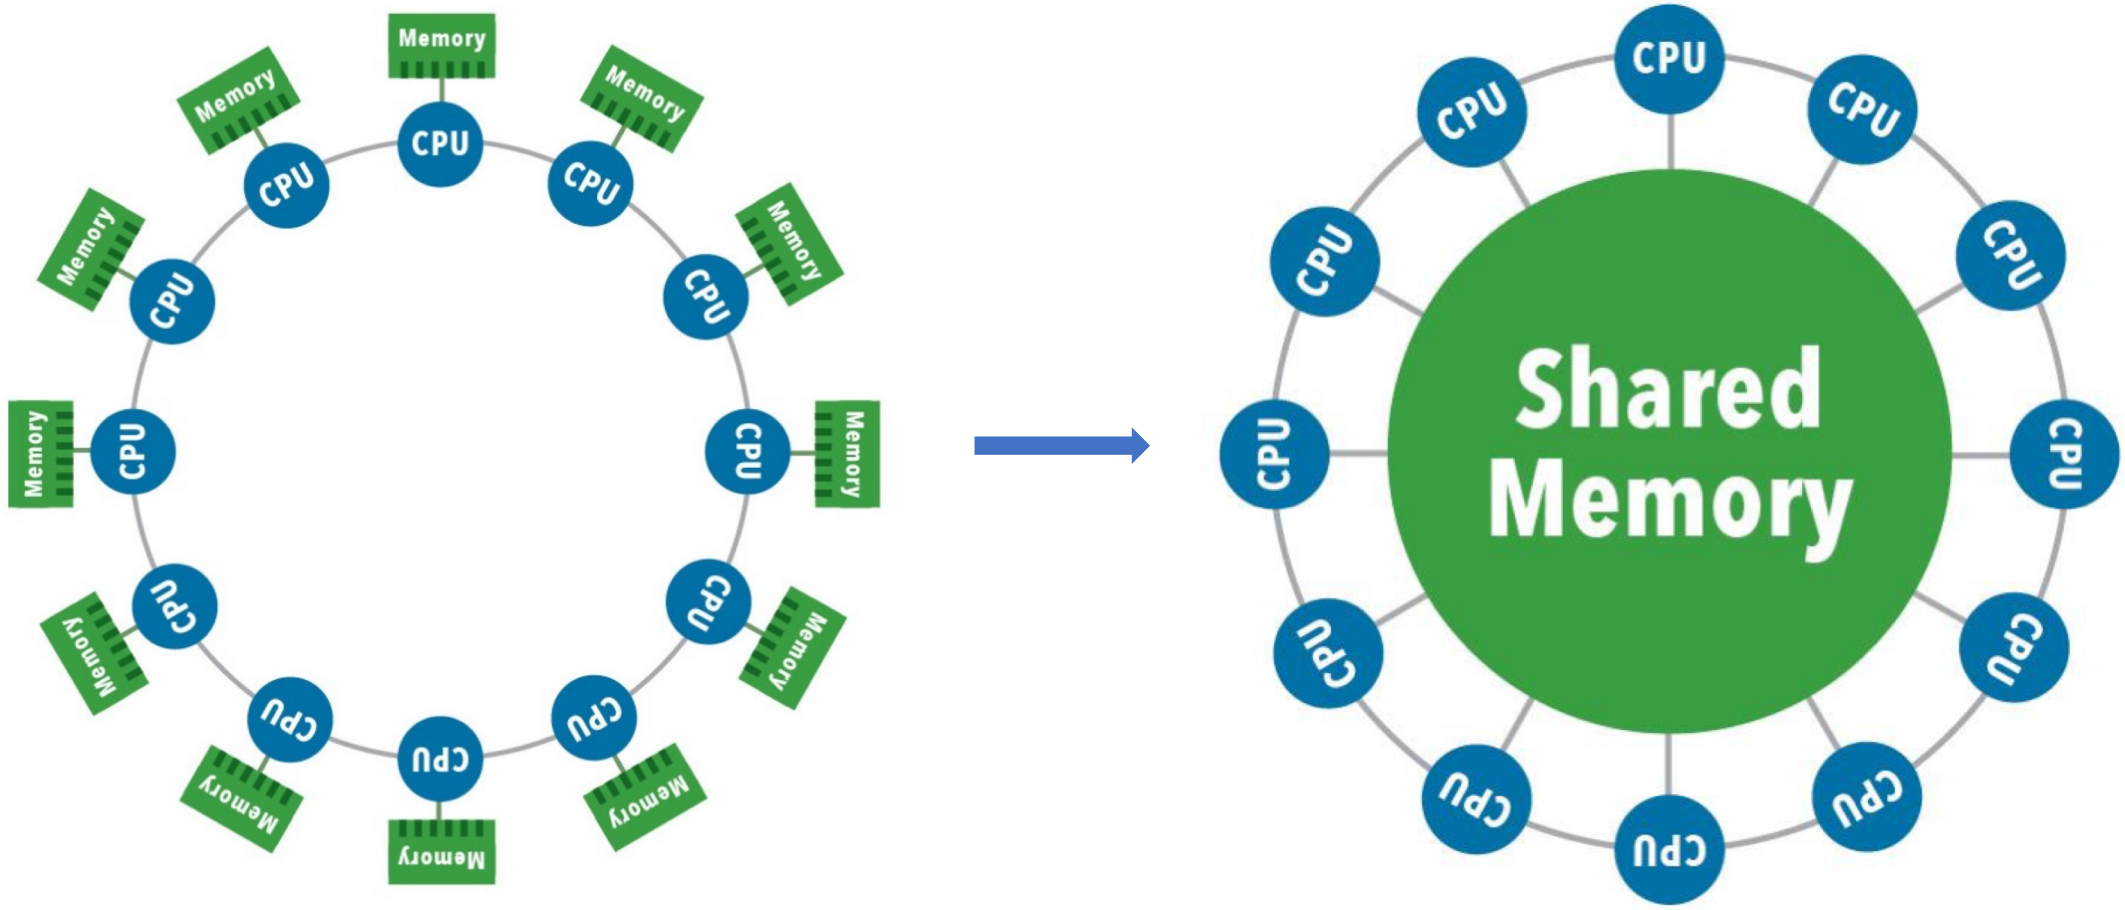

A New Data Access Technology

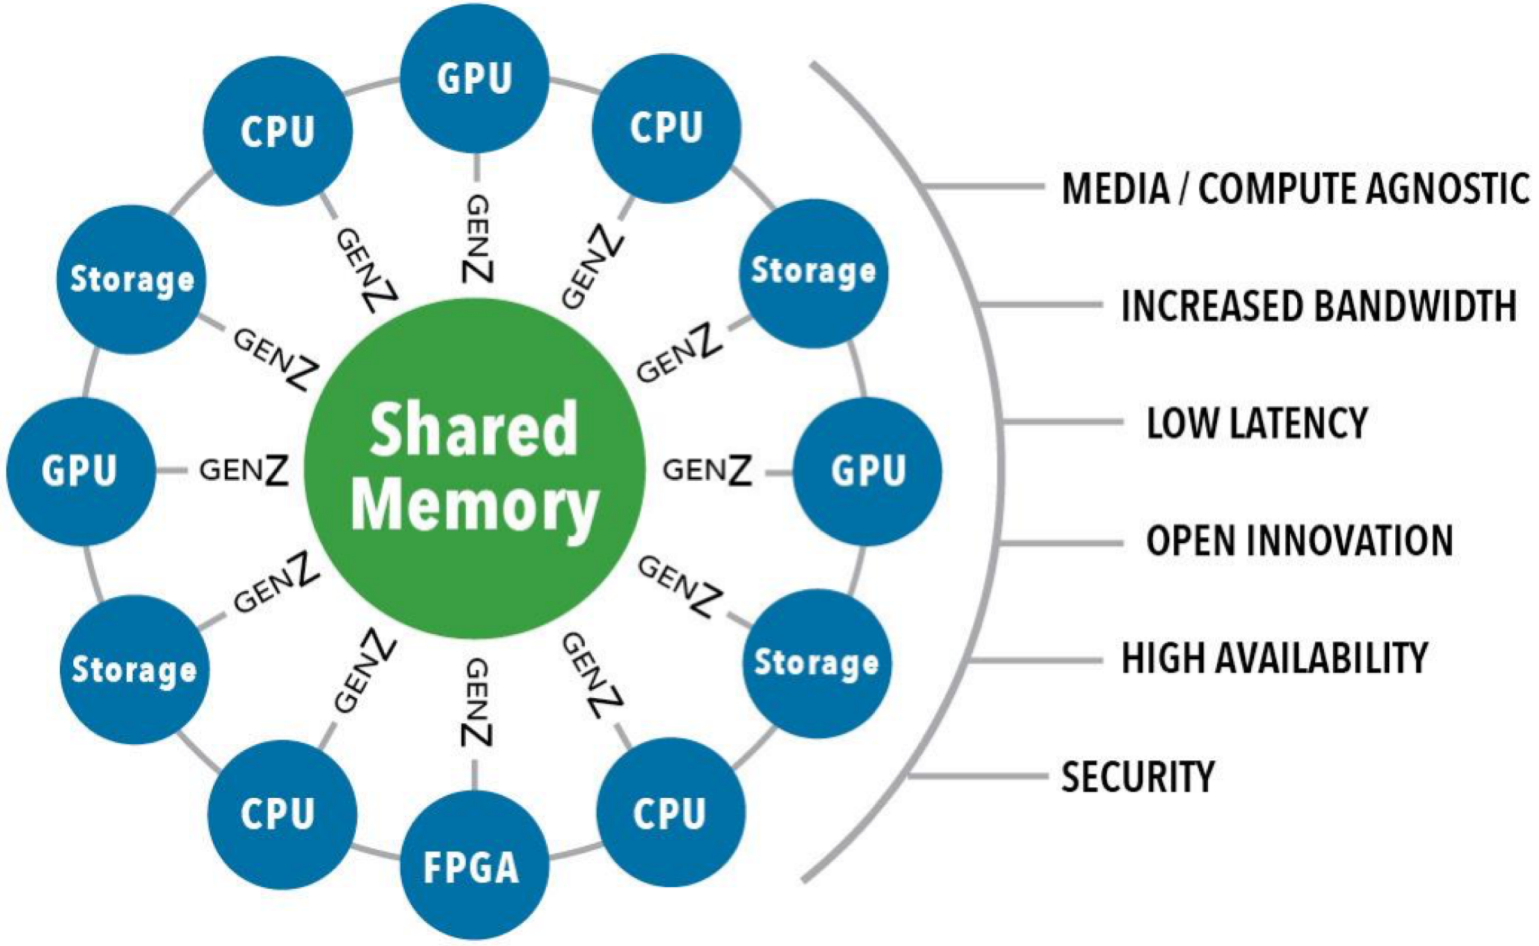

The Gen-Z Solution

Memory Semantic Fabric

- Open
- High-performance
- Reliable
- Secure
- Flexible
- Compatible
- Economic

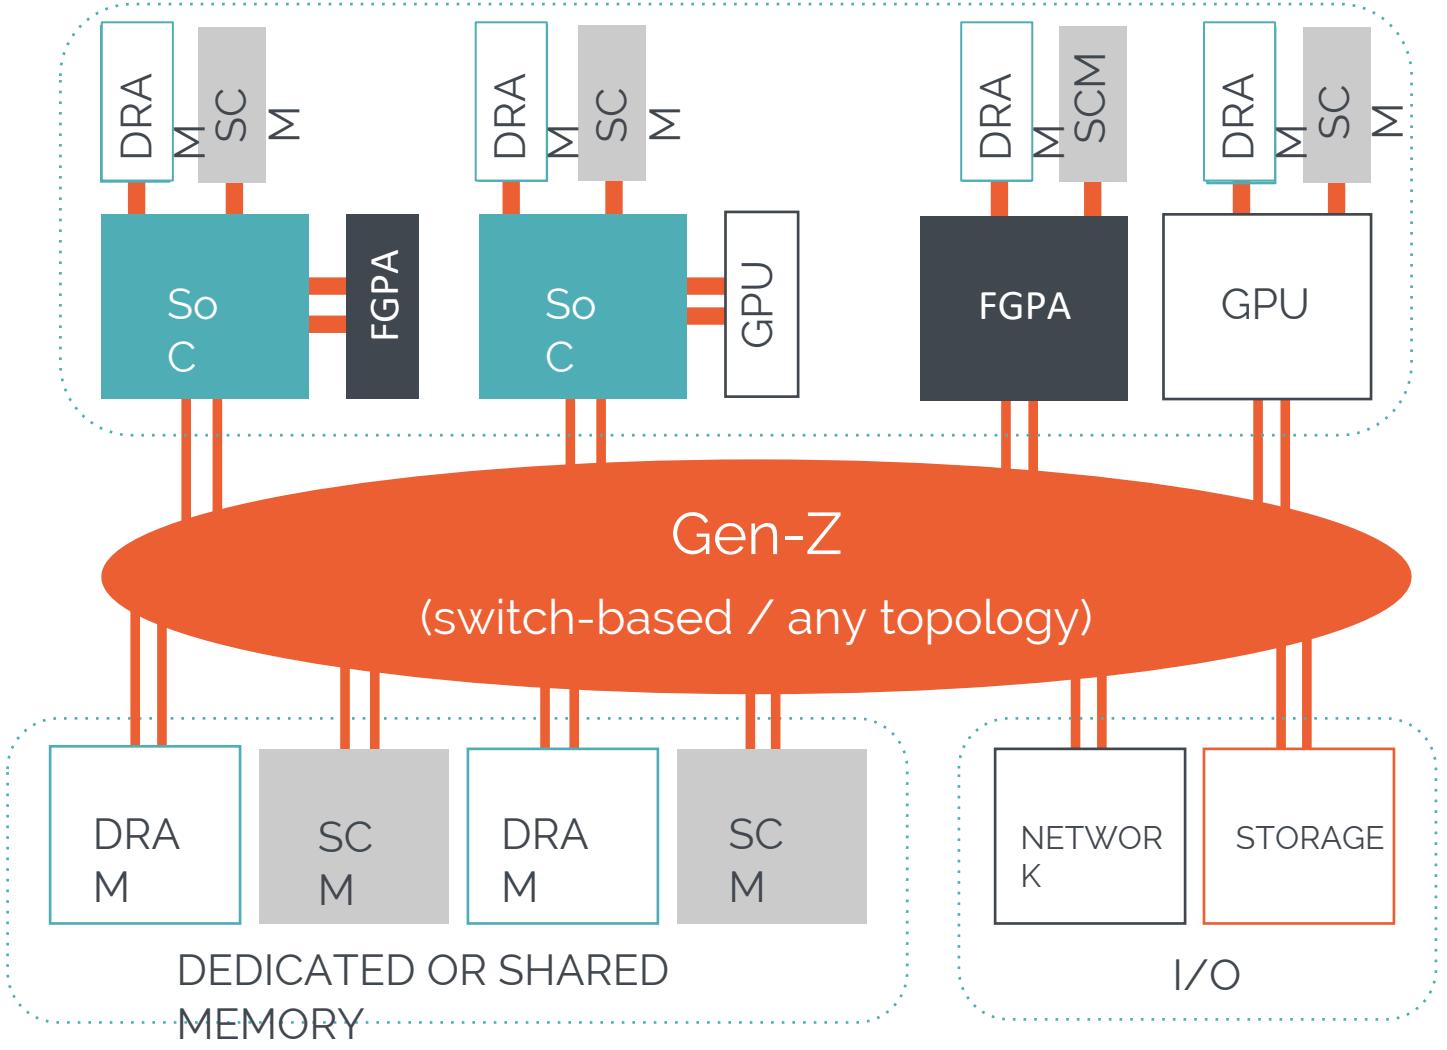

How Does Gen-Z Help Us?

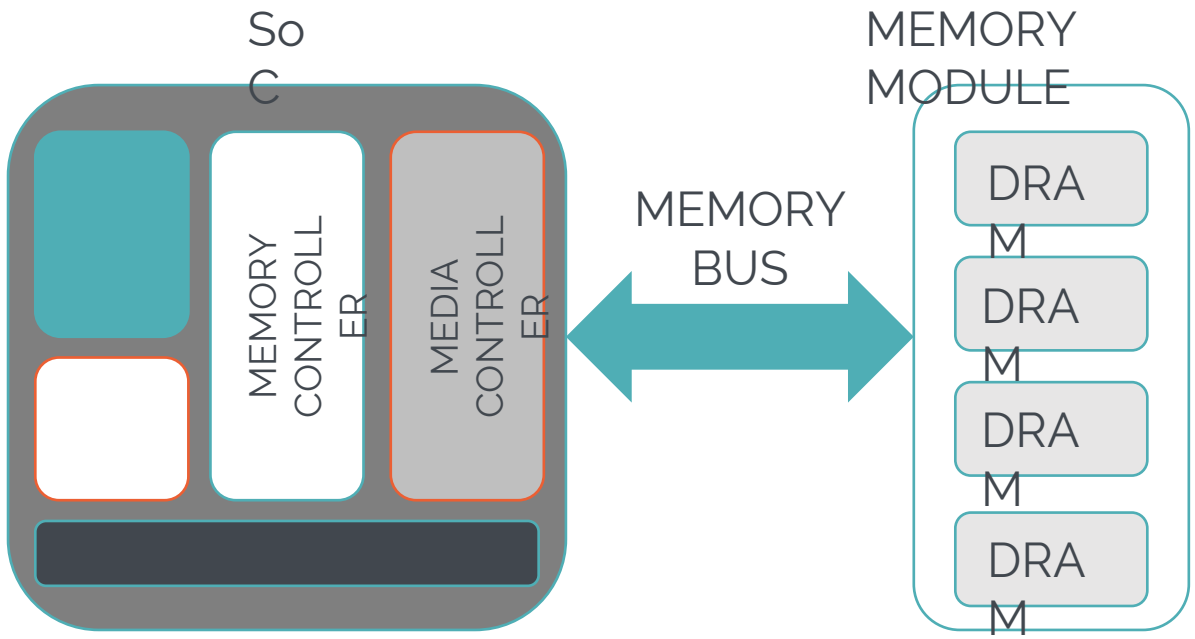

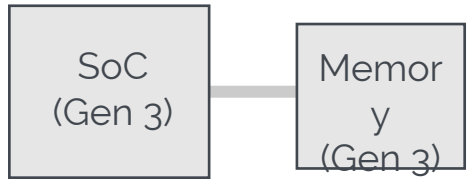

TRADITIONAL

- SoC & memory are in generational lock-step
- Limited memory types

TRADITIONAL MEMORY BUS

- Media specific logic integrated into SoC
- Tight coupling of SoC & memory technology evolution
- Limits the types of memory that can be supported

How Does Gen-Z Help Us?

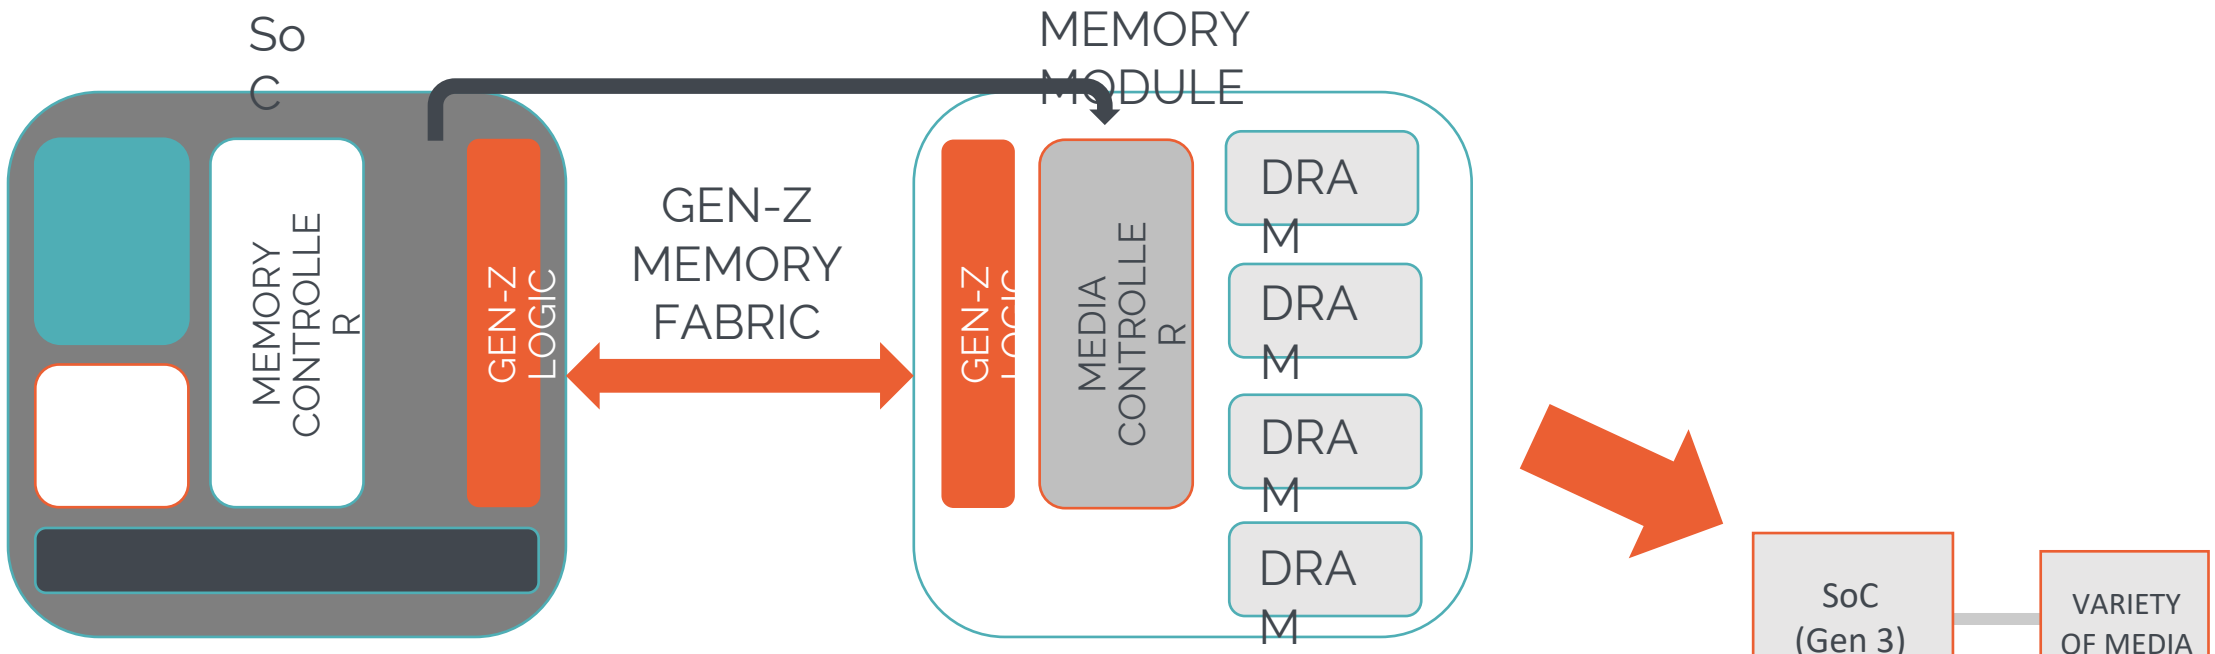

GEN-Z MEMORY INTERCONNECT

- Media specific logic integrated into memory module
- Independent SoC and memory technology evolution
- Accelerates innovation, enables variety of media support

GEN-Z

- SoC & memory evolve freely
- Type & generation independent

How Does Gen-Z Help Us?

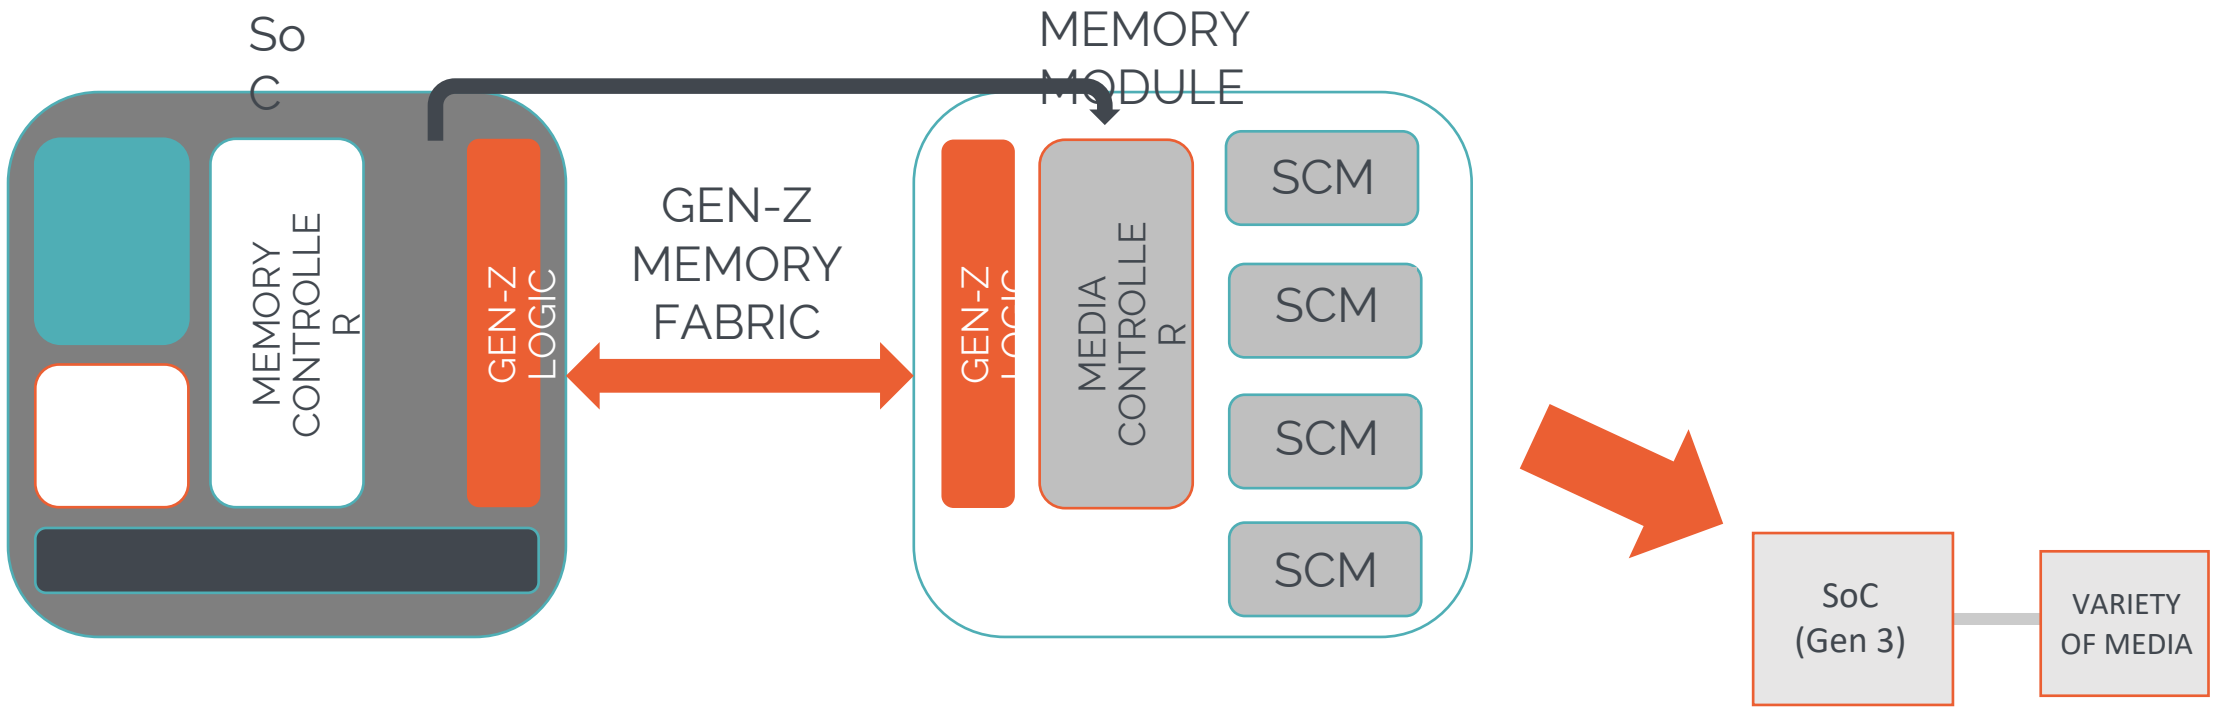

GEN-Z MEMORY INTERCONNECT

- Media specific logic integrated into memory module
- Independent SoC and memory technology evolution
- Accelerates innovation, enables variety of media support

GEN-Z

- SoC & memory evolve freely
- Type & generation independent

Gen-Z Memory Use Case Example

- Seamlessly augments DDR / HBM solutions
  - Supports unmodified OS, applications, middleware
  - Load-stores transparently translated into read-writes
- Abstracts media to break processor-memory interlock
- Very high bandwidth (16 GT/s to 112 GT/s signaling)
  - Delivers 32 GB/s to 400+ GB/s per memory module
- Supports legacy and new high-capacity form factors

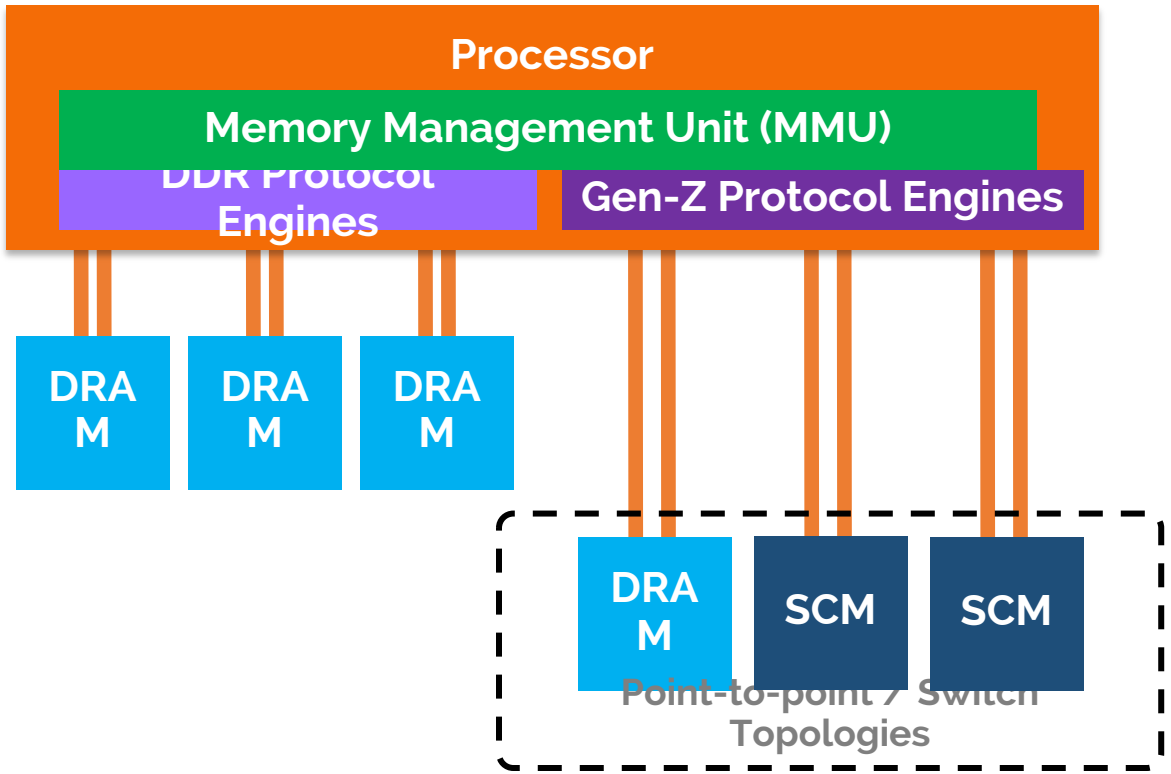

Low Latency Data Access

LOW LATENCY ACCESS TO SCM

- Load/Store byte-addressable access to SCM
- Reduced CPU utilization = more workloads per core/CPU

PERSISTENT MEMORY FILE SYSTEM

- Gen-Z utilizes PMFS developed for NVDIMMs

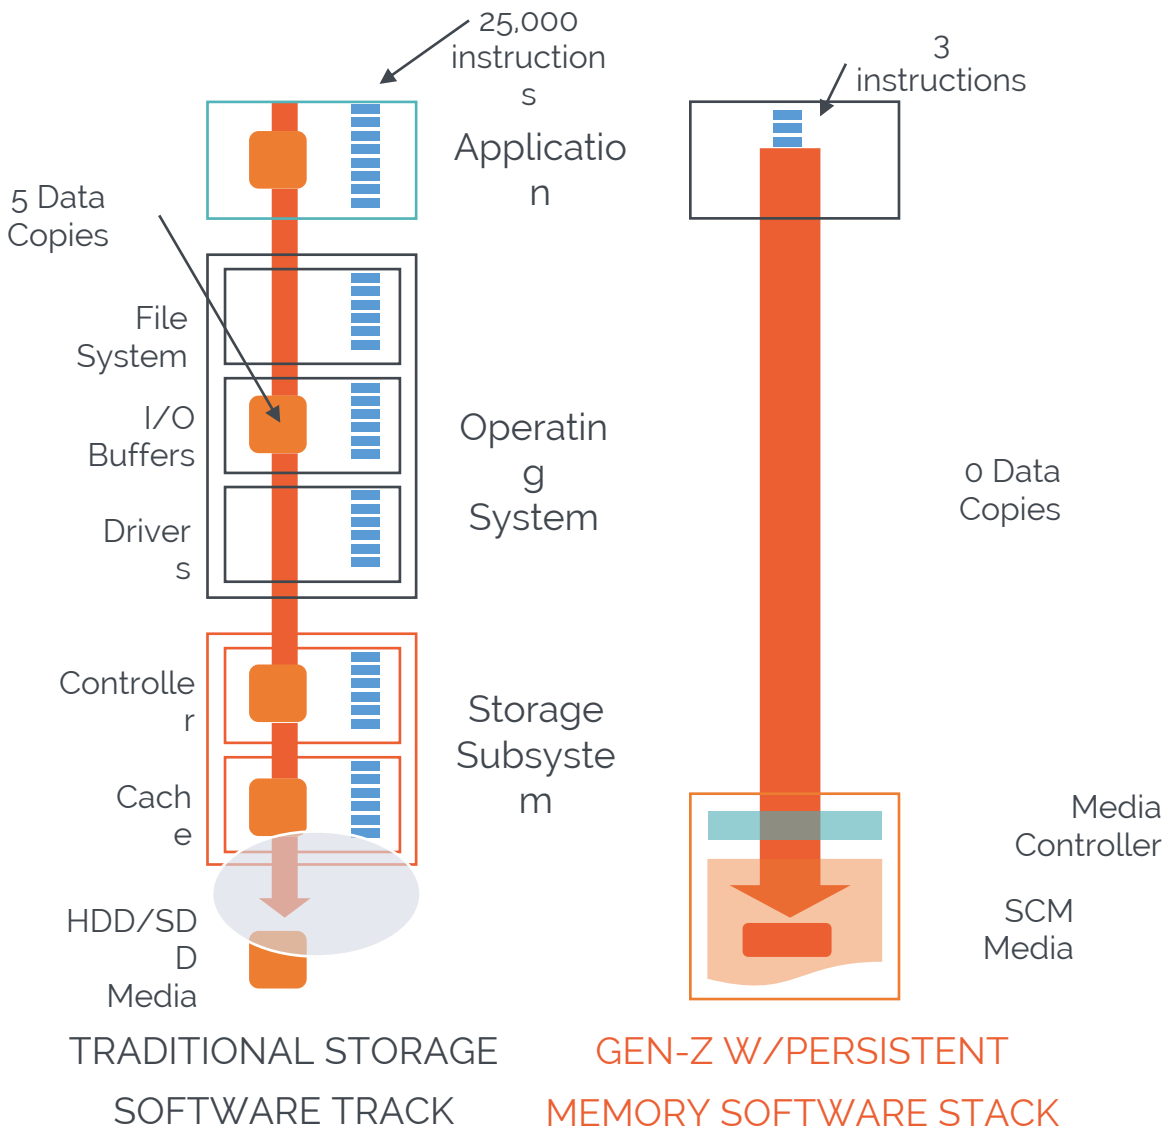

Traffic Management

TRADITIONAL NETWORKS

- Modest traffic prioritization
- Modest multi-path support
- Challenging to mix bulk data and low latency

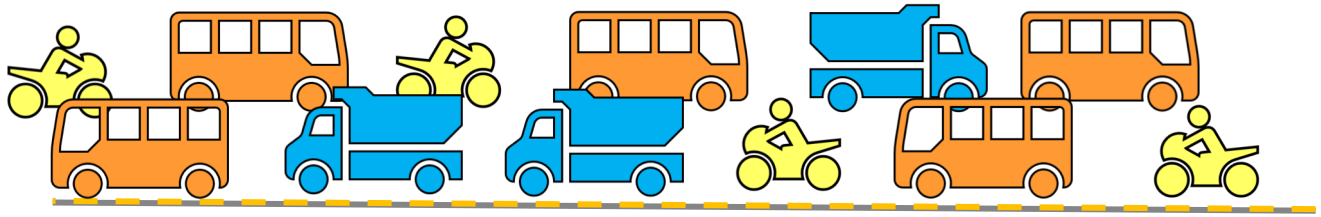

GEN-Z

+ Hardware driven Multi-Link
+ Hardware driven Multi-Path
+ 32 Virtual Channels (VCs) per link
+ 256 byte max packet size
= Industry leading performance

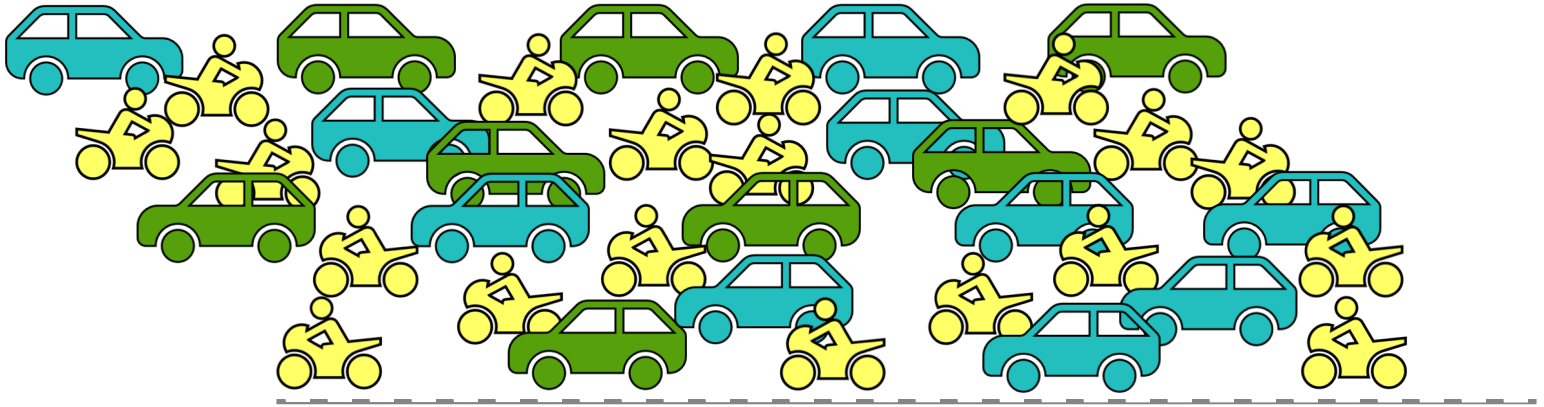

Universal System Interconnect

- Any device connected into any topology
- No dedicated memory, I/O, or storage links
- Enables fluid deployments
- Enables construction of "right-sized" infrastructure

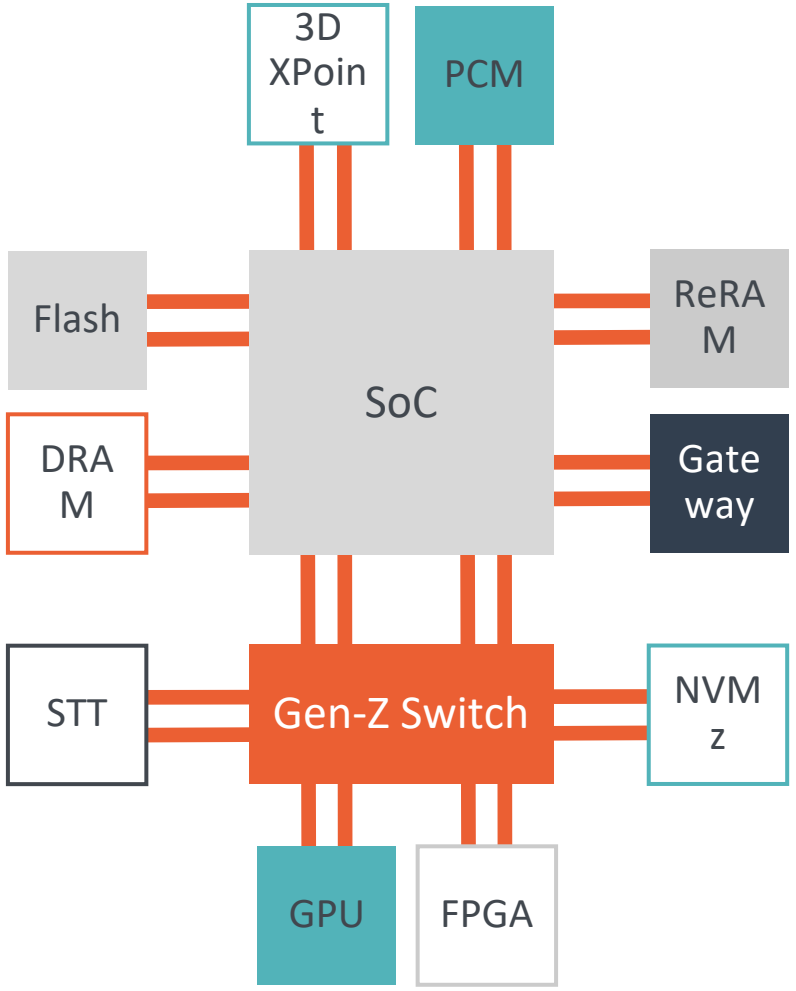

Leading Edge Reliability

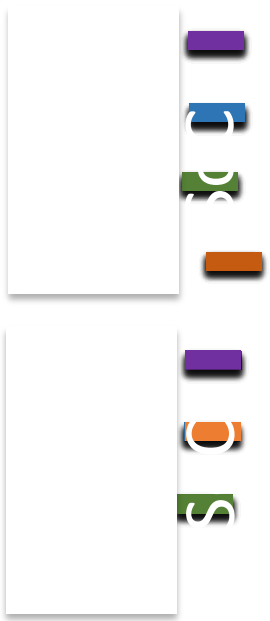

X

MULTI-LINK

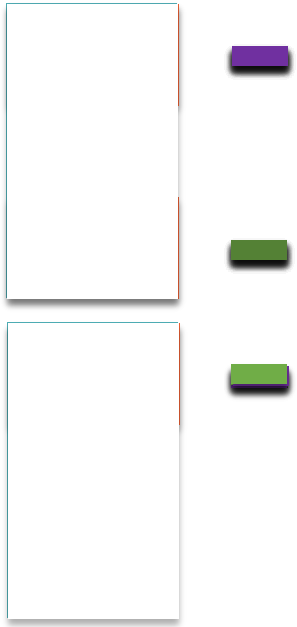

X

FABRIC MULTI-PATH

Industry Leading System RELIABILITY, AVAILABILITY, and SERVICEABILITY

- Multiple links and paths between components with H/W load-balancing, automatic fail-over/fail-back
- Applicable to CPUs, GPUs, FPGAs, DRAM, SCM, storage, etc.
- No OS or software teaming drivers necessary

Gen-Z is Designed to be Secure

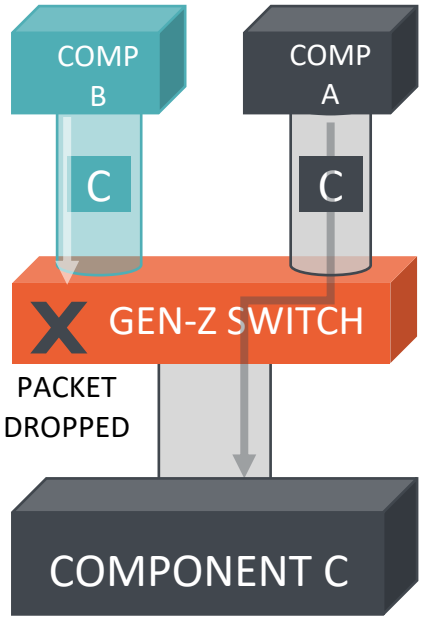

ACCESS KEYS

- Link & device partitioning
- Similar to VLANs or Zones
- Enforced by Gen-Z switches

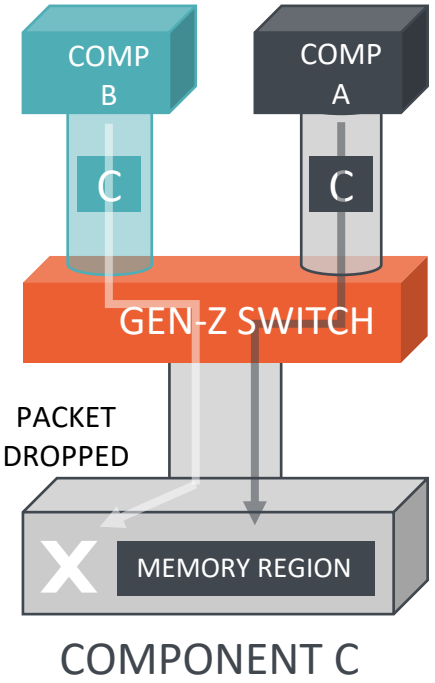

MEMORY REGION KEYS

- Component memory partitioning
- Enables multiple requesters
- Enforced by target component

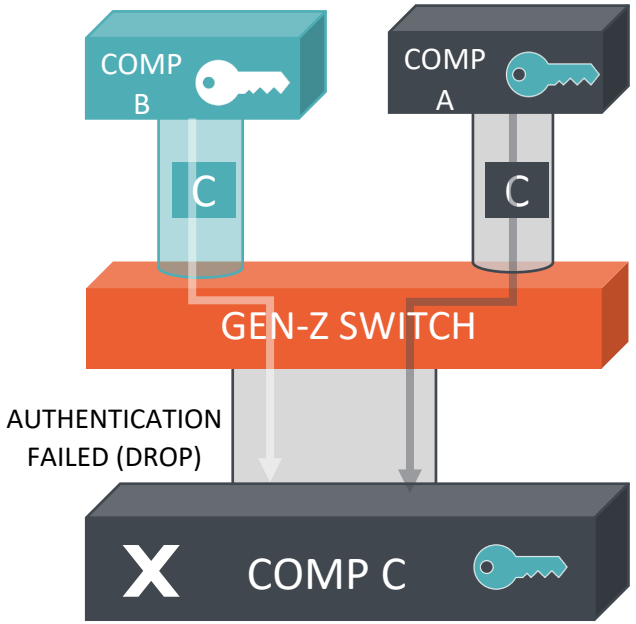

CRYPTO AUTHENTICATION

- Crypto-signed packets (HMAC)
- Multi-tenant, Gov, Financial
- Enforced by target component

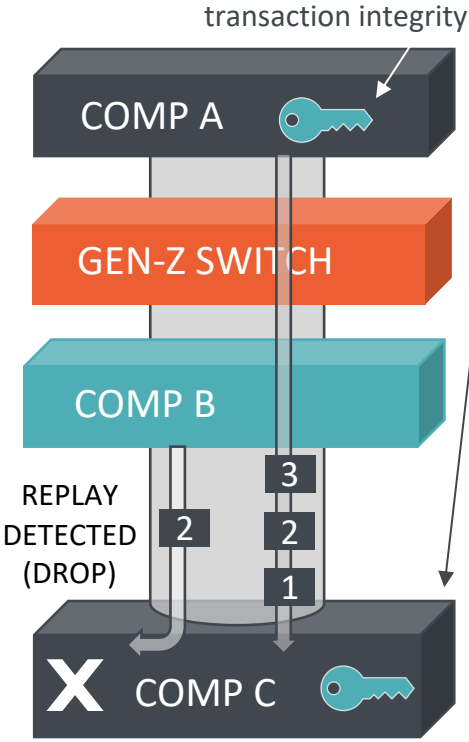

ANTI-REPLAY TAGS

- Prevents "in-the-middle" attacks
- HMAC protected sequence #'s
- Enforced by target

X Dropped packets are reported to fabric management (to identify misconfigurations or malicious attacks)

Unmodified OS Support

Storage Class Memory (SCM)

- SCM driver for block & persistent memory stacks
  - Use Persistent Memory File System developed for NVDIMMs

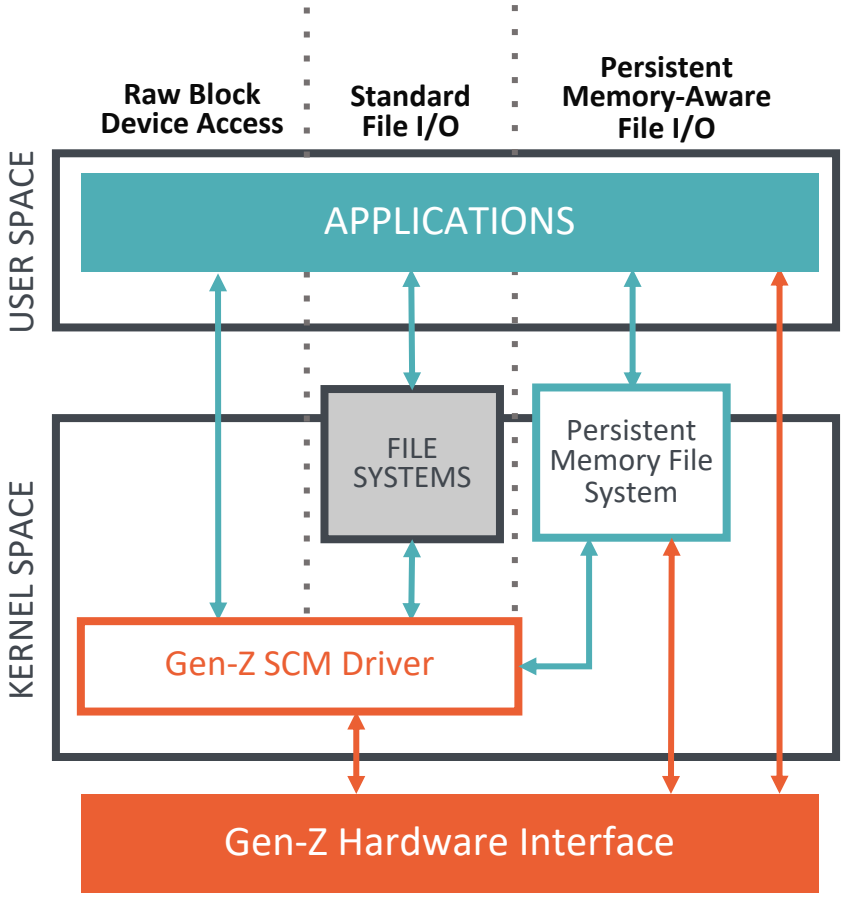

Benefit from the rich set of Gen-Z features while maintaining compatibility with legacy software.

Right-Sized Solutions

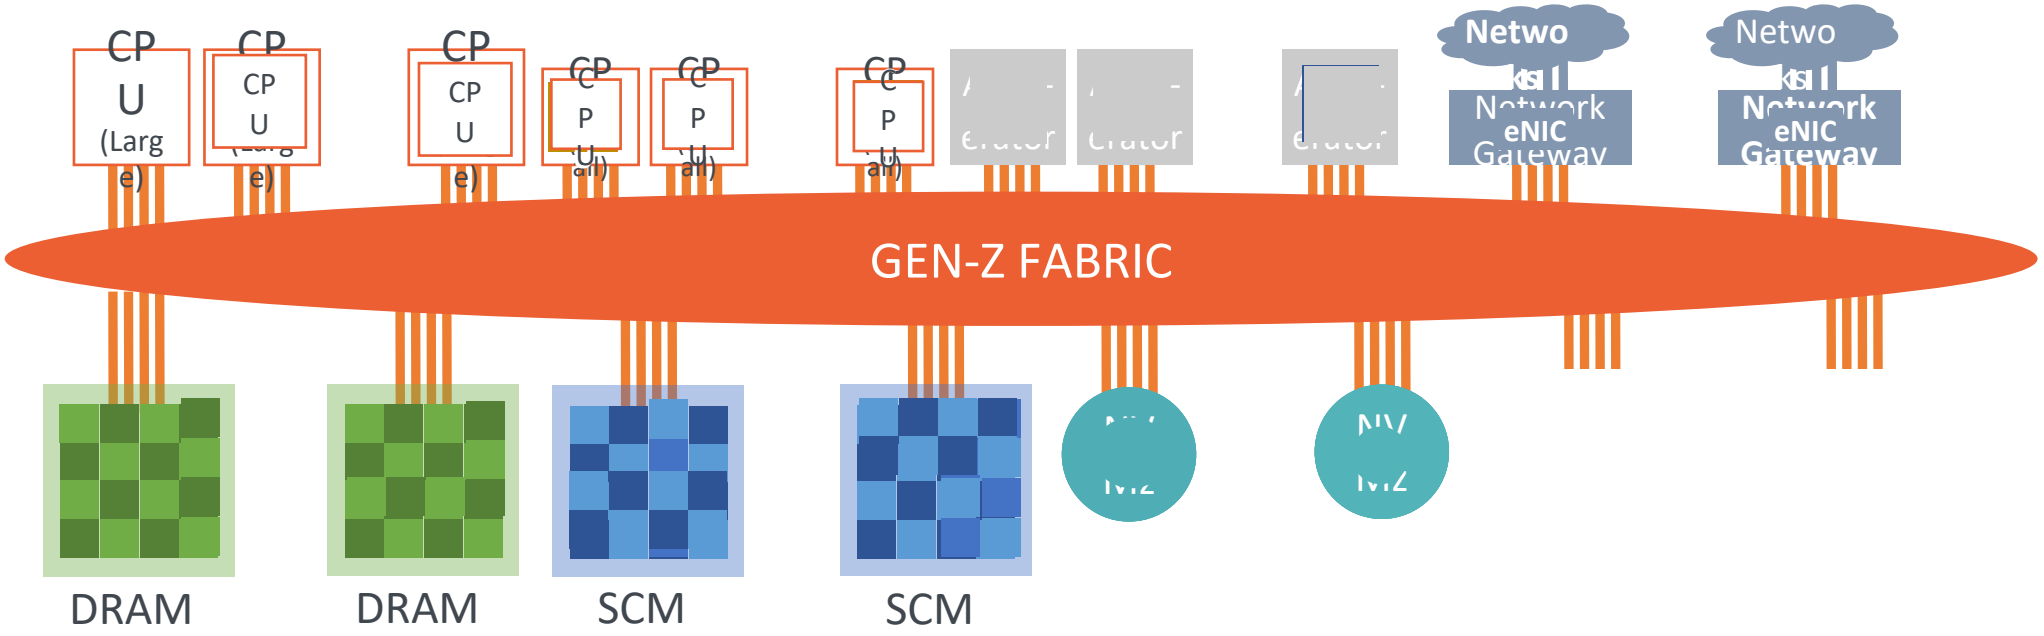

Logical systems composed of physical components

- Or subparts or subregions of components (e.g. memory/storage)

Logical systems match exact workload requirements

- No stranded resources overprovisioned to workloads

Facilitates data-centric computing via shared memory

- Eliminates data movement: Do more with less, reduces cost

LOGICAL WEB SERVER

LOGICAL DATABASE SERVER

LOGICAL ANALYTICS SERVER

LOGICAL KEY-VALUE STORE

Universal Connector System

- Vertical, horizontal, right angle, straddle mount
- Same connectors for memory, I/O, storage, etc.
- Cabled solutions: for copper & optical
- Eliminates "hard choices"
  - Universal connector eliminates industry fragmentation
  - Any component, any slot, any time
  - Multi-connector option to provide added scalability
  - Supports internal and external cable applications
  - Multipath—can bifurcate connector into multiple links
  - Supports multiple interconnect technologies

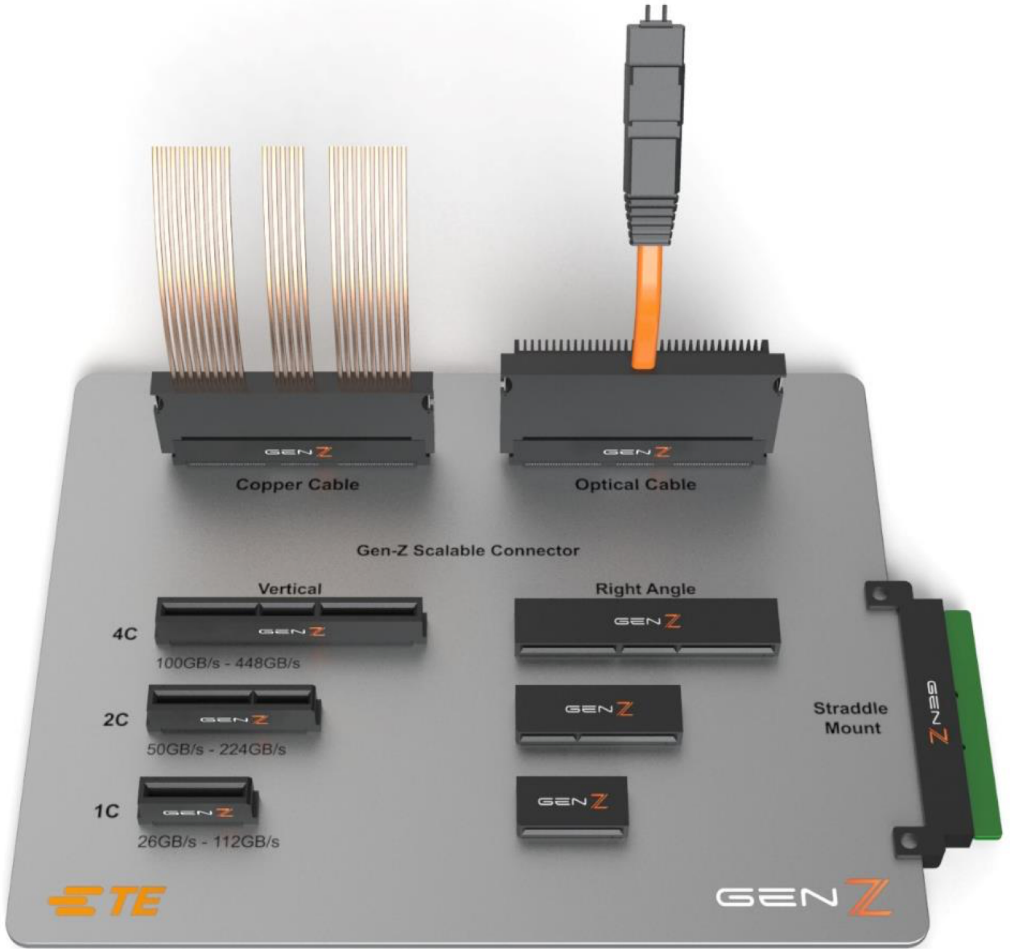

Scalable Resource Enclosures

RESOURCE ENCLOSURES SUPPORT MULTIPLE DRAWERS

- Fixed or hot-plug drawers
- Supports any mix of component types
- Supports memory-centric & data-centric architectures
- Eliminates / reduces need for TOR switches
- Supports multiple interconnect technologies

BENEFITS

- Reduces customer CAPEX / OPEX
- Fully enable / exploit composable infrastructure
- Eliminates stranded media, stranded communication capacity (e.g., no stranded memory channels / I/O links), etc.
- Unlocks solution innovation and agility

Modular, Scalable Form Factors

Supports any component type

- Memory: Flash, SCM, DRAM, etc.

Single/double-wide: scales in x-y-z directions

- Double-wide fits into pairwise single slots

Flexibility

- Media types, power, performance, thermal, capacity

Consistent, low-cost customer experience

- Similar look and feel as SSD drives
- Agile deployment for right sized solutions
- Leverage all Gen-Z benefits (reliability, no SPOF, etc.)

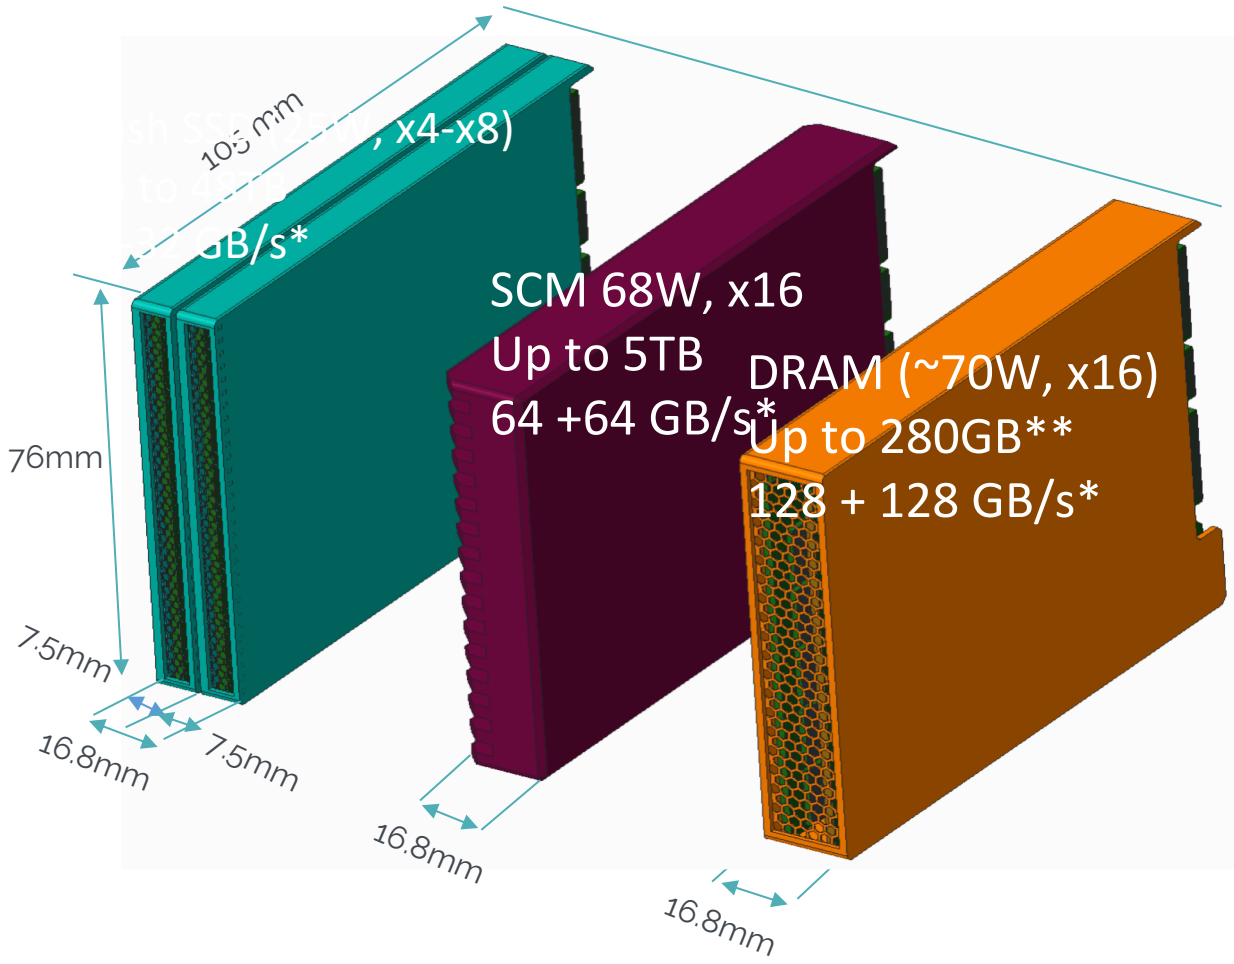

* Bandwidth calculated using 32 GT/s Signaling
** DRAM module provides 3.5x the highest-capacity DDR5 DIMM

Memory-Composability

CURRENT

- Memory is captive of the host device (processor)
- Can't scale memory independently of processing
- All accesses must traverse host processor

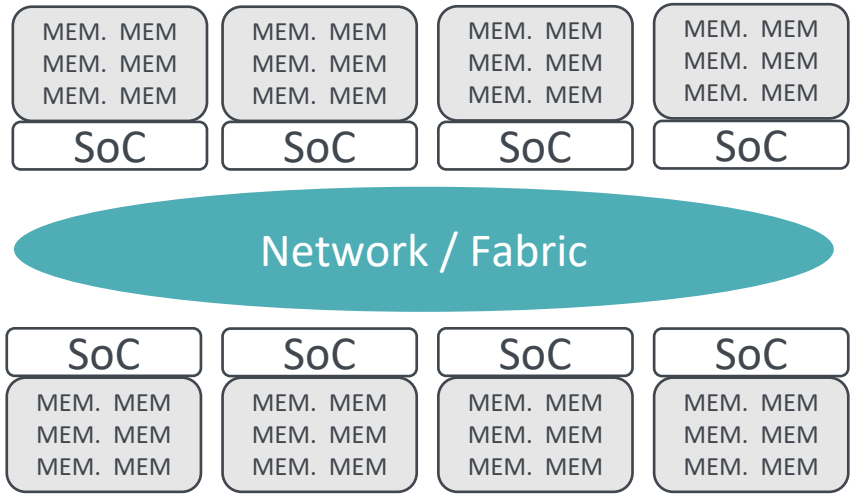

GEN-Z

- Memory and processing scale independently
- Heterogeneous compute & memory deployments
- Direct access to memory devices across fabric
- Memory can be dedicated or shared by processors
- Any topology can be supported

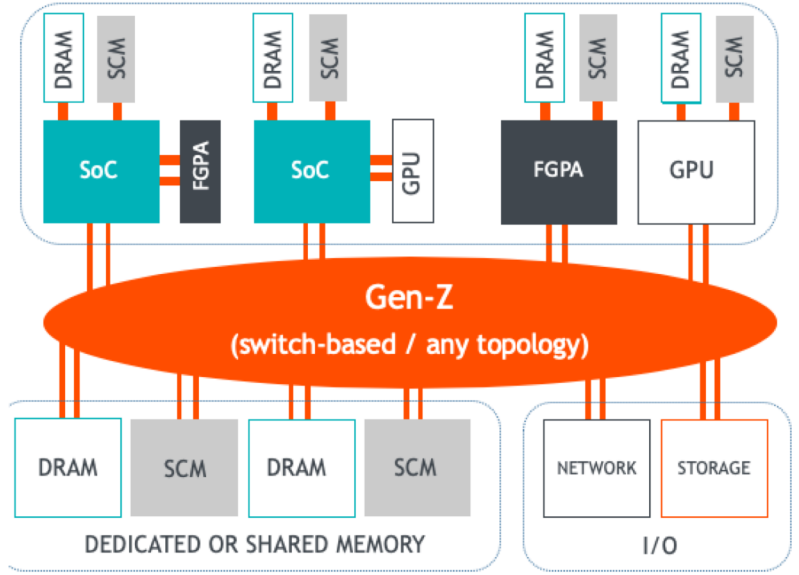

Memory: Freedom of Placement

GEN-Z ENABLED SYSTEM / CHASSIS

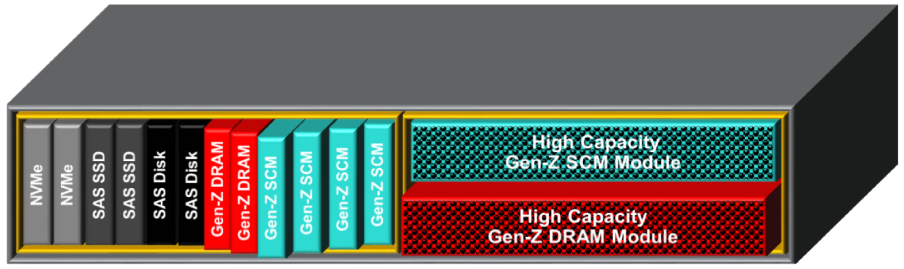

DRAM/SCM can be placed wherever storage is placed - utilize existing or new form factors

GEN-Z ENABLED MOTHERBOARDS

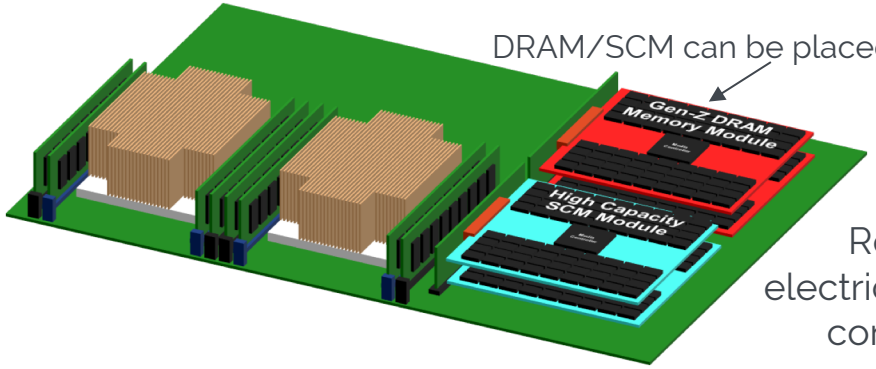

Relieves electrical/thermal constraints

GEN-Z ENABLED RACK SCALE

DRAM/SCM can be placed anywhere in a rack

Utilize compute, storage, or new form factors

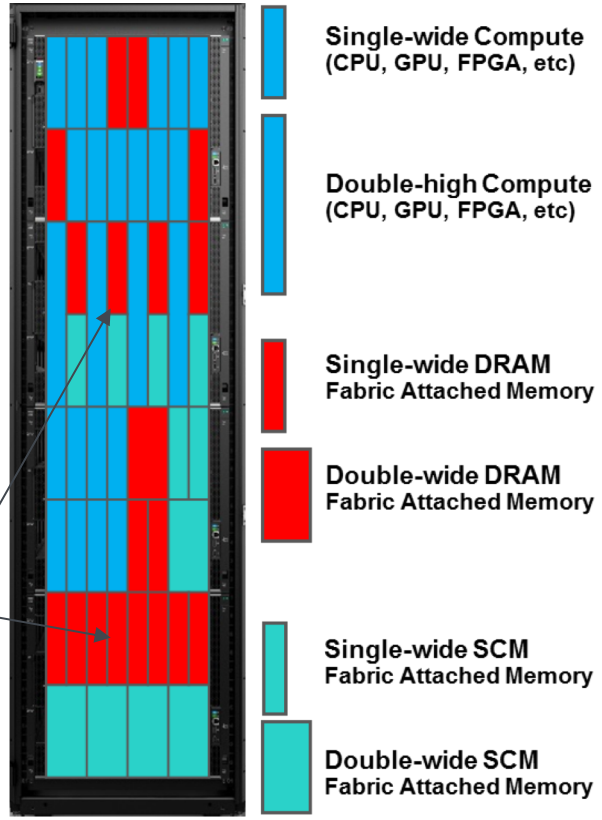

Open: Broad Industry Support

GEN Z CONSORTIUM MEMBERS

- 3M
- Aces Electronics
- Allion Labs
- AMD
- Amphenol Corporation
- ARM
- Avery Design Systems
- BizLink Technology, Inc.
- Broadcom Limited
- Cadence Design Systems, Inc.
- Cisco
- Cray
- Dell EMC
- Electronics and Telecommunications Research Institute
- Foxconn Interconnect Technology, Ltd. (FIT)
- Genesis Connected
- Google
- Hitachi
- HP
- HPE
- Huawei
- IBM
- IDT
- IIT Madras
- IntelliProp
- Jess-Link
- Jinwen University of Science and Technology
- Keysight Technologies
- Kingsignal Technology
- Lenovo
- Liqid
- Lotes
- Luxshare-ICT
- Marvell
- Mellanox
- Mentor Graphics Corporation
- Micron
- Microchip
- Microsoft
- Mobiveil
- Molex
- NetApp
- New H3C Technologies
- Node Haven
- Nokia Solutions
- Oak Ridge National Laboratory
- PLDA
- Qualcomm
- Red Hat
- Samsung
- Samtec
- Seagate
- Senko Advanced Components
- Simula Research Laboratory
- SK hynix
- SMART Modular Technologies
- Sony Semiconductor
- Synopsys
- TE Connectivity Corporation
- Teledyne LeCroy
- Toshiba Memory Corp.
- Triad National Security, LLC
- Univ. of New Hampshire InterOp. Lab
- VMWare
- Western Digital
- Xilinx
- YADRO
- Yonsei University
- Zhejiang Zhaolong Interconnect Technology

Flash Memory Summit

COMMUNICATION AT THE SPEED OF MEMORY

Universal Interconnect

Memory-Semantics

Transforms

www.genzconsortium.org

Resources

Gen-Z demo at our Booth!

View Gen-Z educational materials, membership details and links to related information at **www.GenZConsortium.org**

@GenZConsortium

Interested in MEMBERSHIP? Have QUESTIONS? EMAIL: admin@genzconsortium.org

THANK YOU