Gen-Z®: Built-In Security for the Data Centric World

Presented by: Michael Krause, HPE VP / Fellow, Gen-Z Lead Architect

FLASH MEMORY SUMMIT 8-7-19

Edge-to-Core Impacts to Data Generation & Access

Infrastructure Security Matters More Than Ever

August 21, 2018 Microsoft removes multiple websites allegedly created by *Fancy Bear* to influence US Elections

2016 Yahoo reveals millions of accounts hacked cut $350M from Verizon's Yahoo acquisition price

3Q2019 US Department of Defense to award $10B JEDI Cloud Contract

August 16, 2017 Maersk reported that the NotPetya cyberattack cost ~$300M in lost revenue; multiple executives terminated

December 19, 2013 Target retail store breach cost $252M & CEO terminated

Everything is an Attack Vector

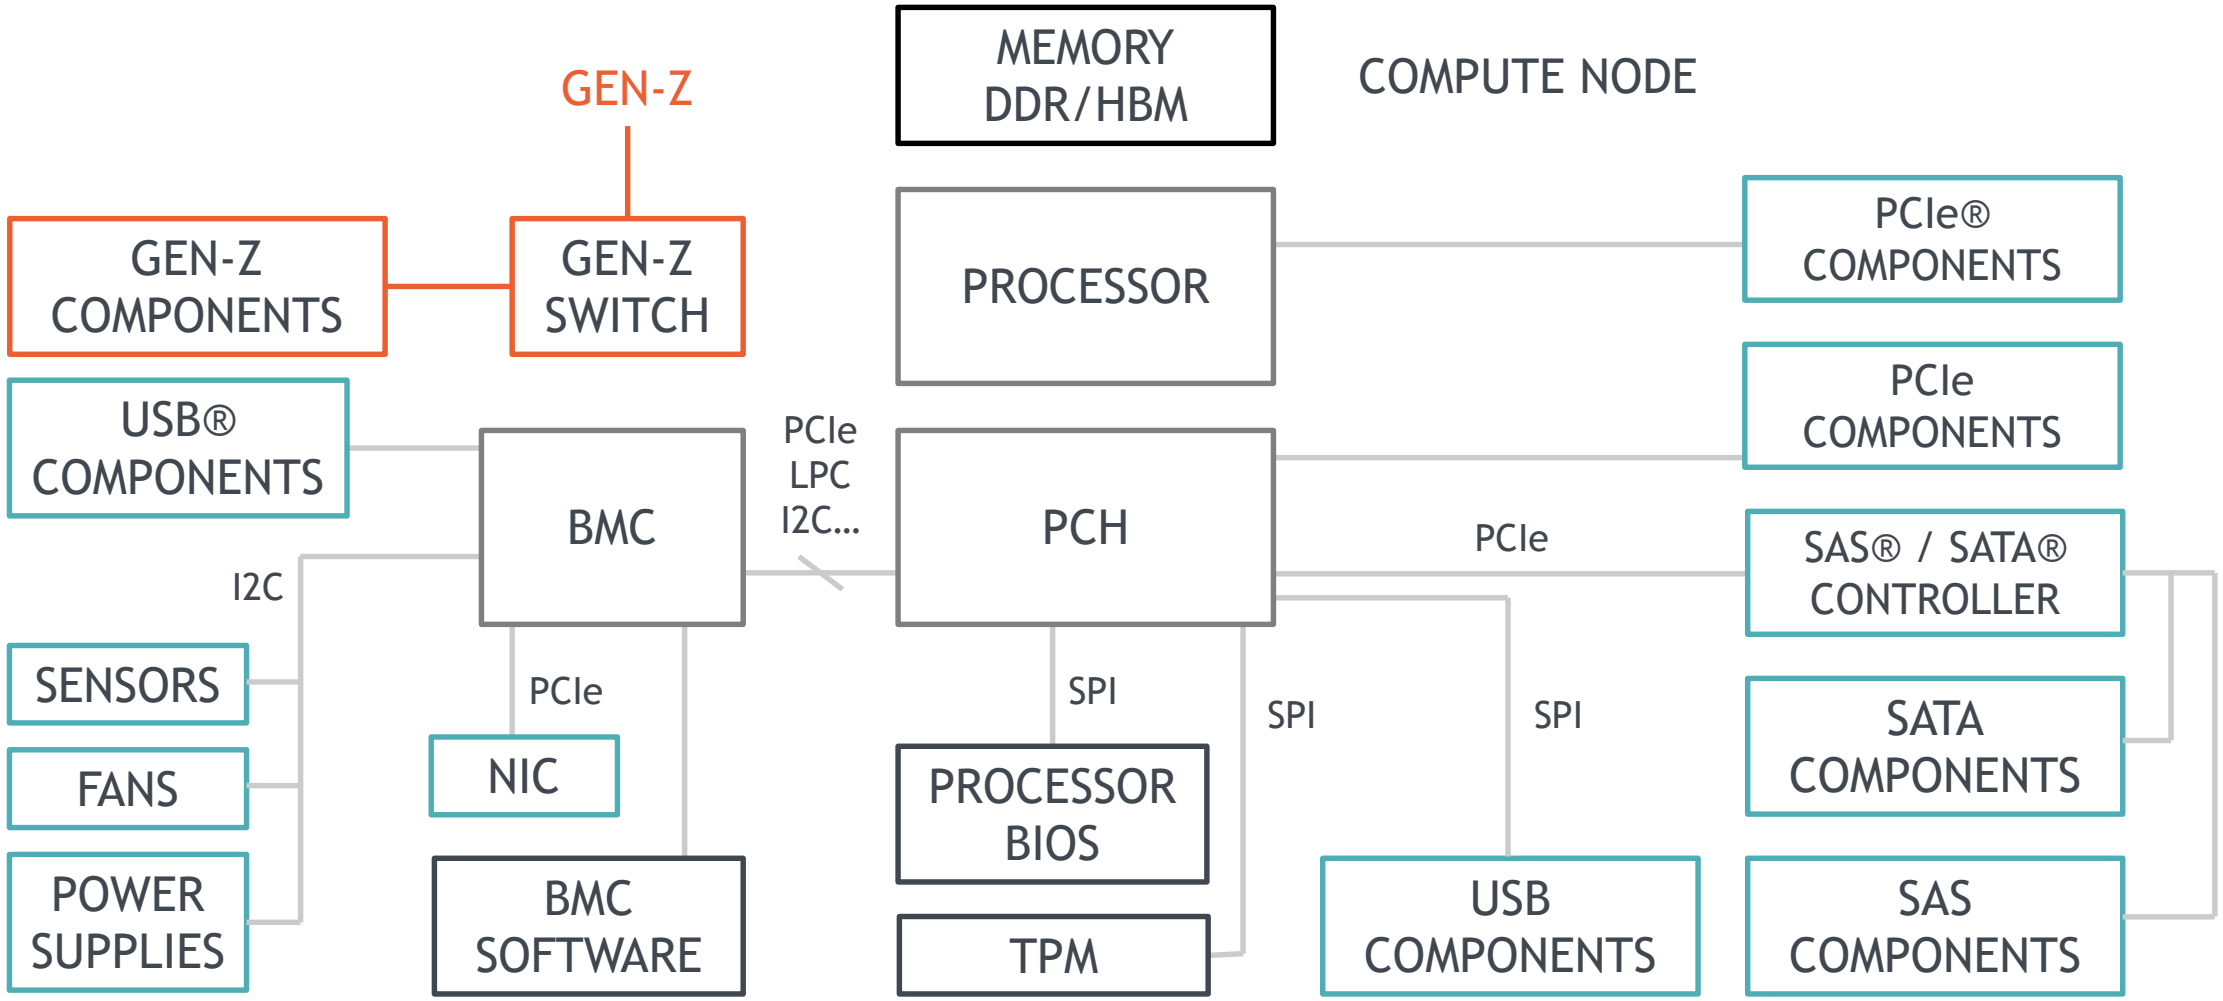

”Sea” of Memory/Storage Accessed by Optimized Compute

Industry embracing workload-optimized compute

- From edge-to-core with custom/semi-custom compute
  - General-purpose compute adding new ISA capabilities
- Workload drives performance, data movement, power, cost…
  - Edge-based compute & reduction critical to reducing backbone & data center communications load & improving responsiveness
  - Minimal data movement to optimize power, reduce latency, reduce infrastructure cost

Multi-petabyte data sets are common

- Data must be accessible by any compute module at any scale
  - Multiple compute modules need to access 100s to 1000s of memory & storage modules at any given time
    - NVMe over PCIe limited scalability / constrained sharing model
    - NVMe over Fabrics requires additional fabric (adds cost/complexity) & suffers similar native PCIe sharing constraints
    - NVMe over Gen-Z easily scales to 8192 NVMe SSDs per CPU or accelerator & simplifies multi-way sharing
    - Gen-Z Buffer Ops represent the next step for memory/storage data movement
- Data at rest must be secure/private
- Data movement must be secure/private

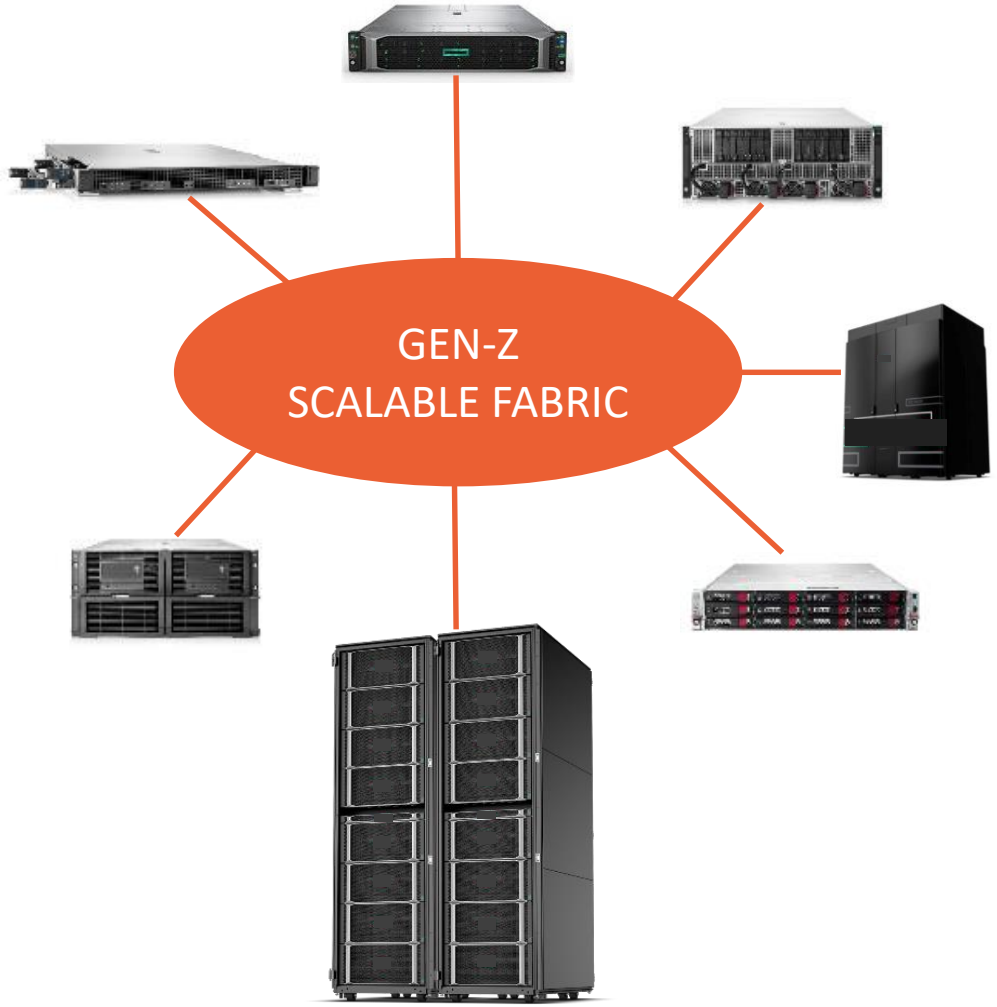

CAPEX/OPEX-Optimized Composable Infrastructure

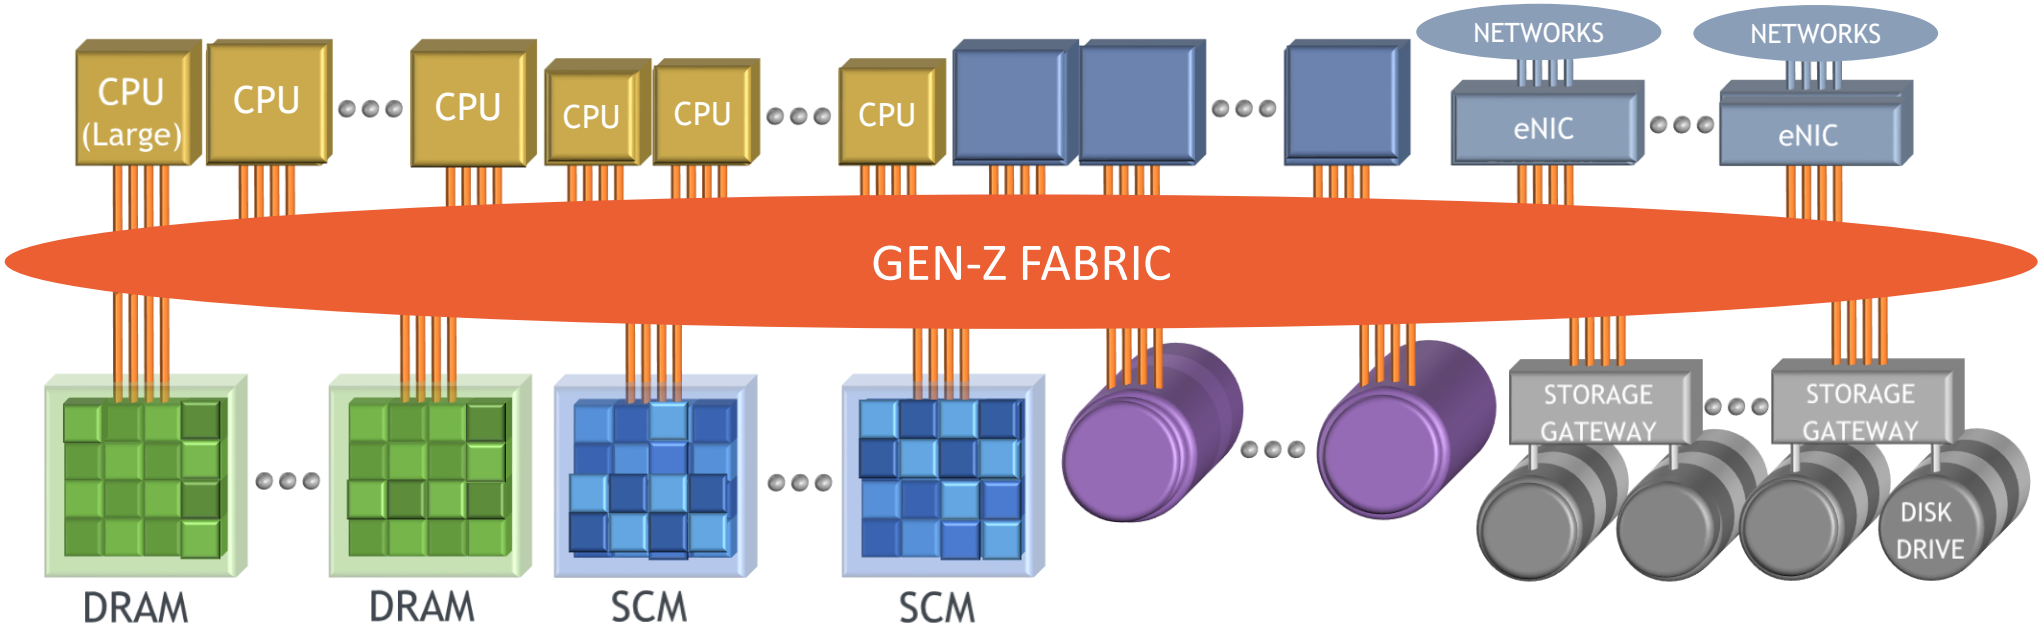

- Composability eliminates “hard choices”
- Logical systems composed of physical components
  - Or subparts or subregions of components (e.g. memory/storage/I/O)
- Logical systems match exact workload requirements
  - No stranded resources overprovisioned to workloads
- Facilitates data-centric computing via shared memory
  - Eliminates data movement: Do more with less, reduce/ optimize cost

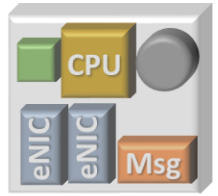

LOGICAL WEB SERVER

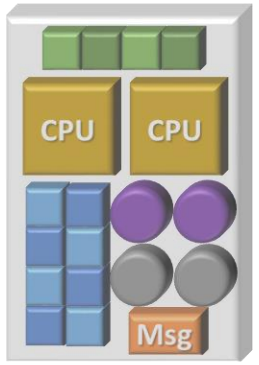

LOGICAL DATABASE SERVER

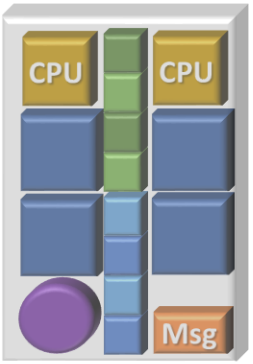

LOGICAL ANALYTICS SERVER

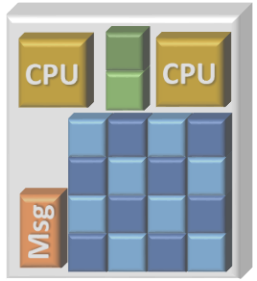

LOGICAL KEY-VALUE STORE

Gen-Z Basics

High Performance
- High Bandwidth, Low Latency, Scalable
- Eliminates protocol translation cost/complexity/latency
- Eliminates software complexity/overhead/latency

Reliable
- No stranded resources or single-point-of-failures
- Transparently bypass path and component failure
- Enables highly-resilient data (e.g. RAID / erasure codes)

Secure
- Provides strong hardware-enforced isolation and security

Flexible/Extensible
- Multiple topologies, component types, use cases, etc.

Compatible
- Supports PCIe PHY up to 32 GT/s, 802.3/OIF 25 GT/s to 112 GT/s PAM 4, unmodified OS support

Economic
- Lowers CAPEX/OPEX, unlocks/accelerates innovation

GEN-Z SPEAKS THE LANGUAGE OF COMPUTE

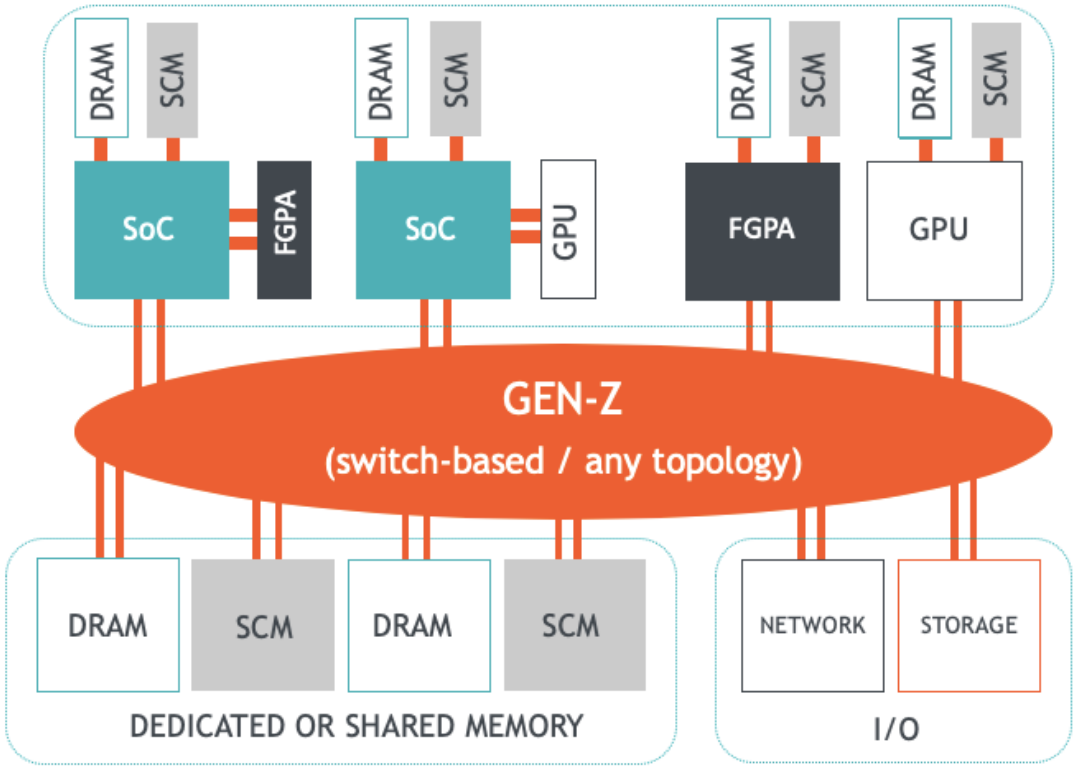

Gen-Z Hardware-Enforced Access Control

Supports multiple mechanisms to ensure authorized component and resource access

- Authorization is managed by software and enforced by hardware
- Authorization does not equate to security
  - Authorization mitigates the potential damage caused by erroneous or failing components
  - Authorization mitigates the potential damage caused by malicious actors

Supported techniques (access violations reported to management):

- Access Keys—these provide component group-level communication access control
- Access Request and Access Response controls—these provide fine-grain component-level access controls
- Region Keys (R-Keys)—provide page-level access control (required to share resource)
- R-Key Domains—provide Requester R-Key range filter access control
- Switch packet filtering—used to filter which packets can be relayed and where
- Component Destination Structure—used to configure authorized peer components
- Nonce validation to prevent rogue hardware replacing a component during deep low-power
  - Validated whenever a link transitions from low-power state to operational state

Gen-Z Component Authentication

DMTF®-based Component Authentication[1]

- Gen-Z members donated authentication security objects[2] to DMTF to enable multiple industry bodies to support a consistent & robust security ecosystem
- Authenticates hardware, configuration, & firmware/ software
- Component authentication performed at component manufacturing, component integration, initialization & power cycle events, runtime, component addition/ replacement/decommissioning

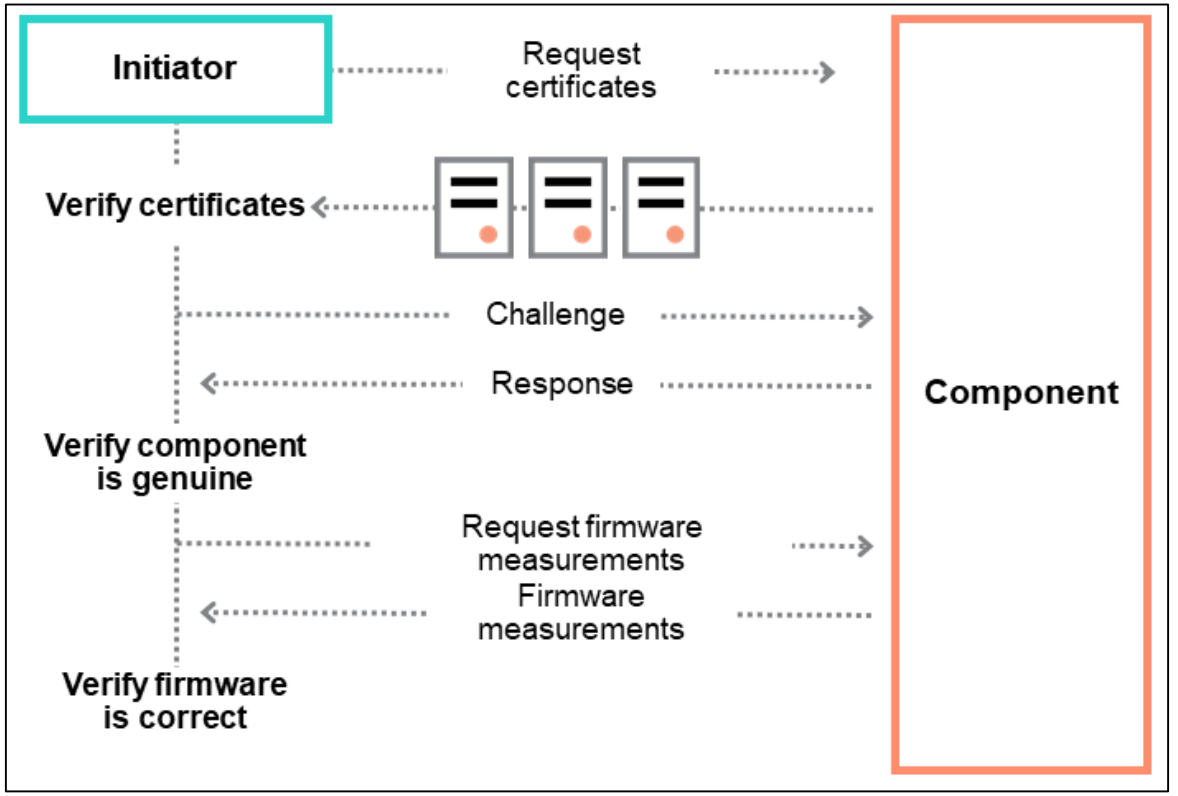

Data objects exchanged using MCTP over Gen-Z

- MCTP operates over multiple interconnects: I2C/I3C, PCIe, Gen-Z, …
- MCTP simplifies management => improved security infrastructure

1 High-level component authentication white paper applicable to multiple interconnects / technologies is available at: https://genzconsortium.org/wp-content/uploads/2019/03/Gen-Z-Component-Authentication-Secured-Infrastructure.pdf

2 See DMTF Security Protocol and Data Model (SPDM) specification

Gen-Z Packet Authentication and Encryption

Gen-Z data privacy + packet authentication = **maximum security protection**

Gen-Z uses authenticating encryption to provide:

- AES 256 encryption
- Packet tamper and anti-replay attack protection

Packet authenticating encryption may be selectively enabled

- Unique session per communicating peer
- Data plane, control plane, or both

Gen-Z encryption sessions will use DMTF SPDM session establishment (WIP)

- DMTF SPDM will be supported by multiple interconnects thus enhancing solution and infrastructure security
- See: https://www.dmtf.org/content/dmtf-releases-security-protocol-and-data-model-spdm-architecture-work-progress

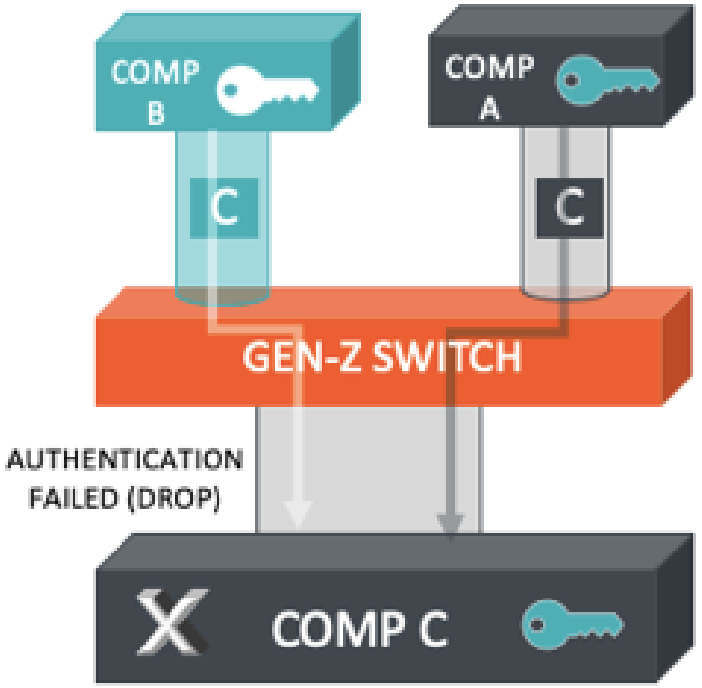

Violations reported to management

Secured Gen-Z Buffer Operations

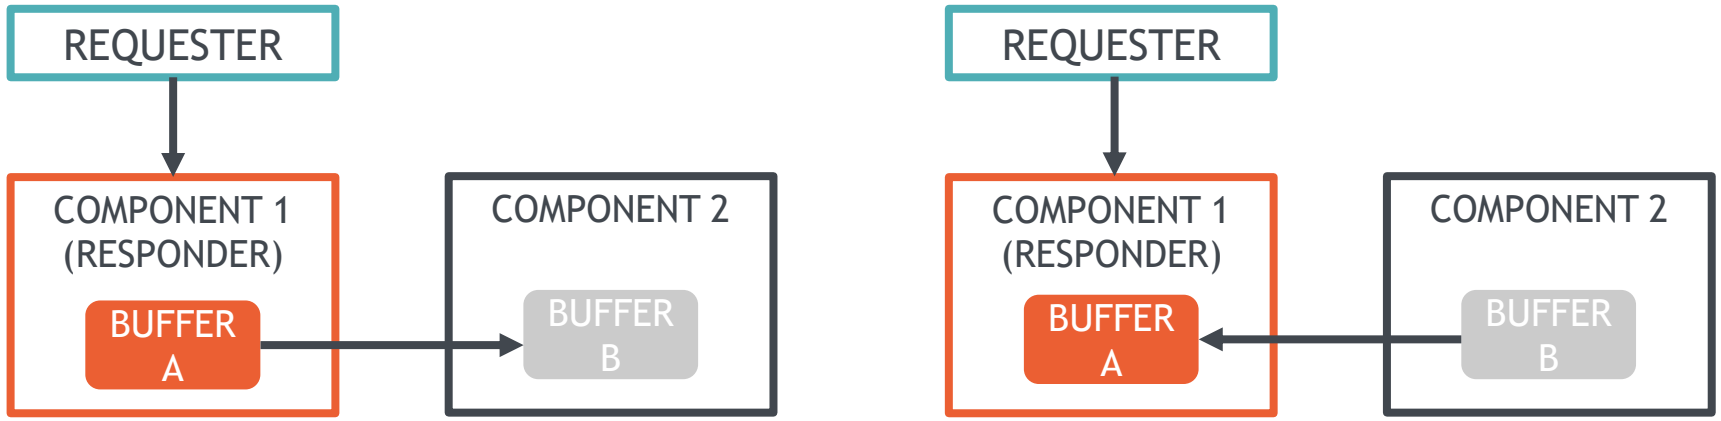

Gen-Z specifies 16 Buffer operations that simplify data movement, buffer allocation, security, etc.

- No work queues, completion queues, etc. to manage—single buffer request / response
- One-, Two-, and Three-party data movement

Buffer operations can improve performance in multiple ways

- Three-party data movement eliminates need for all data to flow through a SoC
- Signaled buffer operations eliminate multi-thread / multi-node coordination overhead and complexity
- T10 DIF / PI acceleration and Secured Hash / Encryption acceleration
- Simplifies software / management, reduces control packet overheads, multipath aggregation, etc.

Gen-Z Emergency/Planned Primary Media Backup

- Enables primary media to be automatically copied to a secondary media.
  - Works in a single enclosure or composable
  - Primary and secondary media can be dedicated or shared by multiple compute / systems
  - Primary and secondary media can be mechanically co-located or discrete
    - Shared amortizes backup media costs across multiple primary media modules
- Emergency backup can be initiated by a processor once it has flushed its caches or if the primary media controller detects h/w failure
- Planned backup can be initiated by software to create dynamic data checkpoints

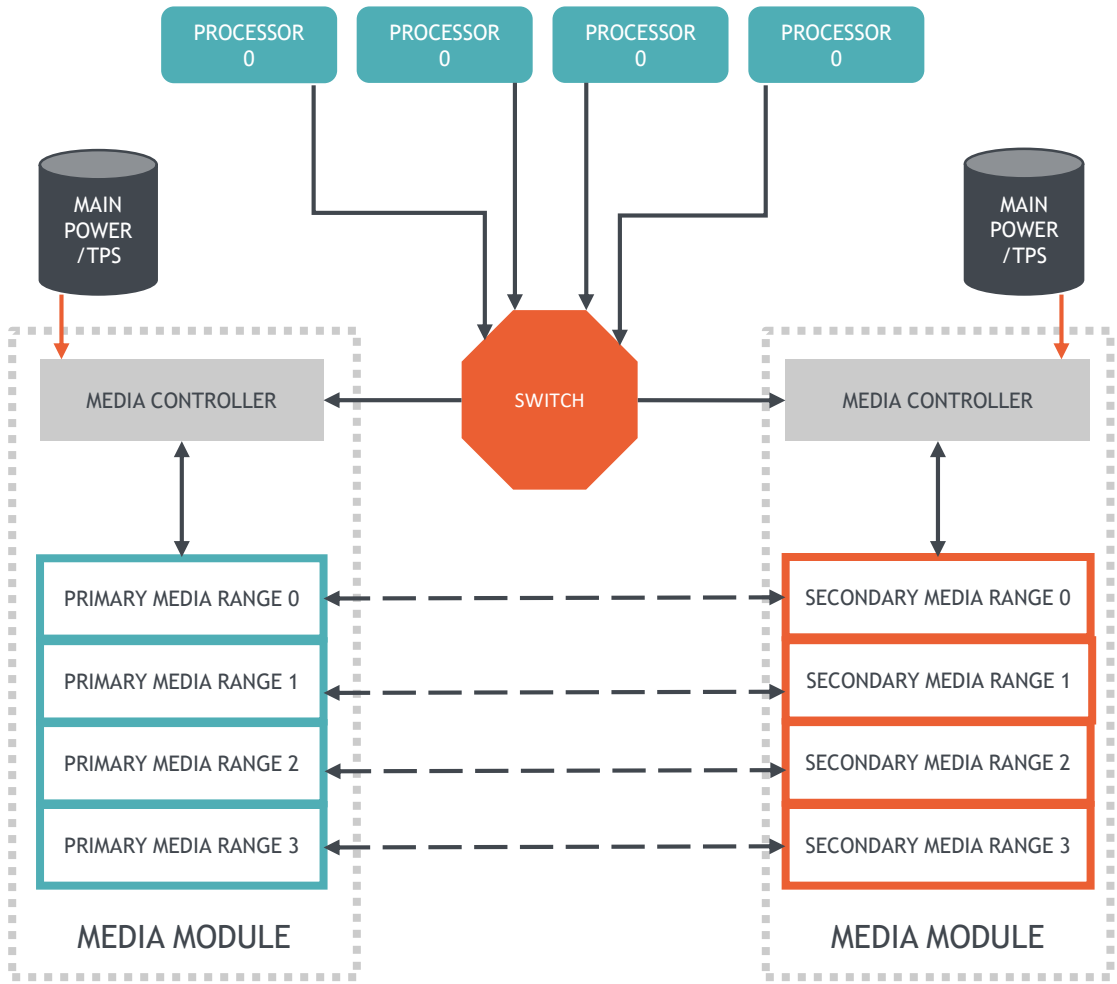

Secured Emergency/Planned Memory Backup

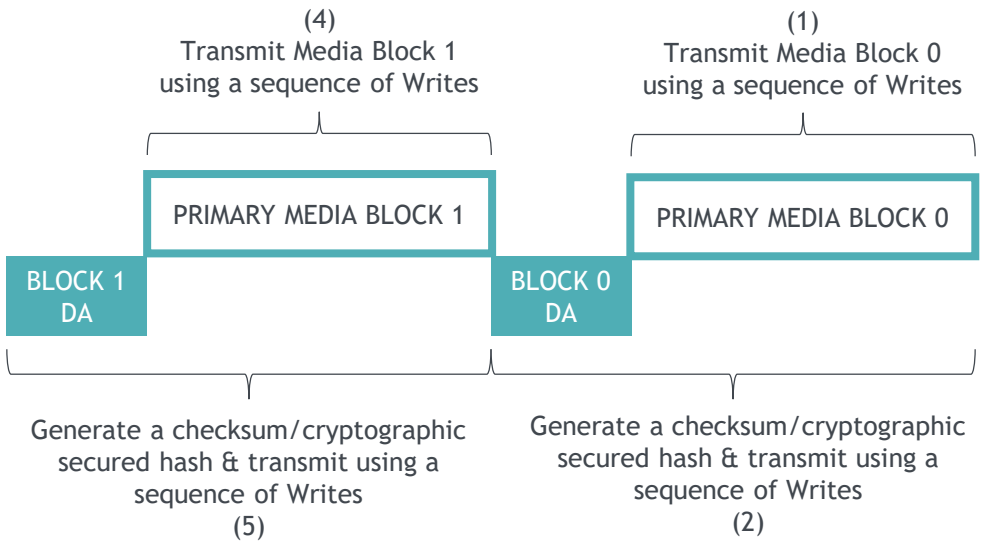

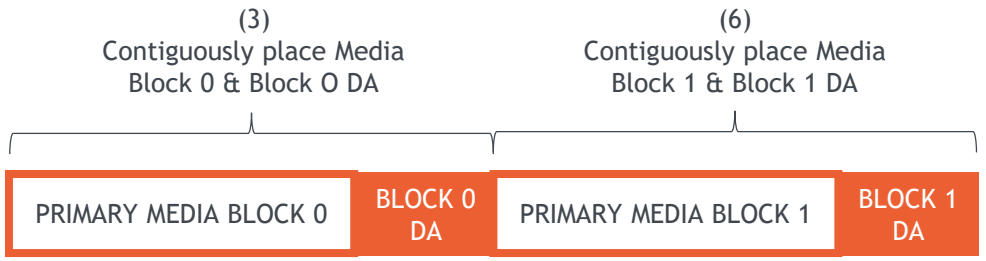

Direct-attached and shared / composable memory emergency and planned backup services

- Strong data integrity, cryptographically-secured hash, or encryption as moved to / from primary / secondary media
- Only primary media comprehends data protection applied—does not extend trust to secondary backup media

Summary

Gen-Z architecture designed from the start to support a fully-secured infrastructure

- Multiple hardware-enforced isolation mechanisms
- Component authentication and session management built on top of DMTF SPDM to simplify security infrastructure and management
- Strong privacy and packet tampering and anti-replay protection

Upcoming Gen-Z technology deliverables:

- New connector test fixture specifications
- New module-level optical connectivity through connector
- Gen-Z PHY 1.1 specification that supports:
  - PCIe PHY up to 32 GT/s
  - 802.3/OIF 50G PAM 4 (53.125 effective) with 2 ns FEC
- Compliance testing

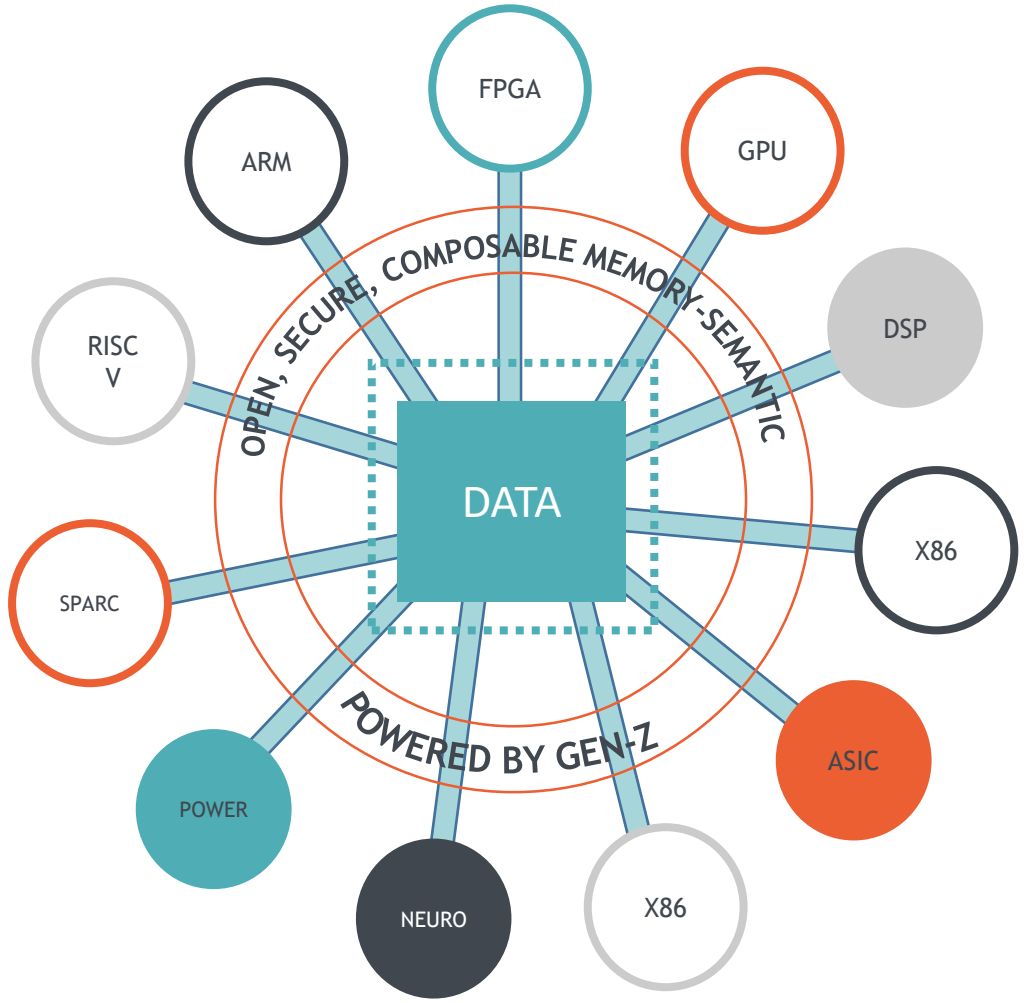

Resources

UPCOMING EVENTS

Gen-Z demos at Super Computing 2019 (November)

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