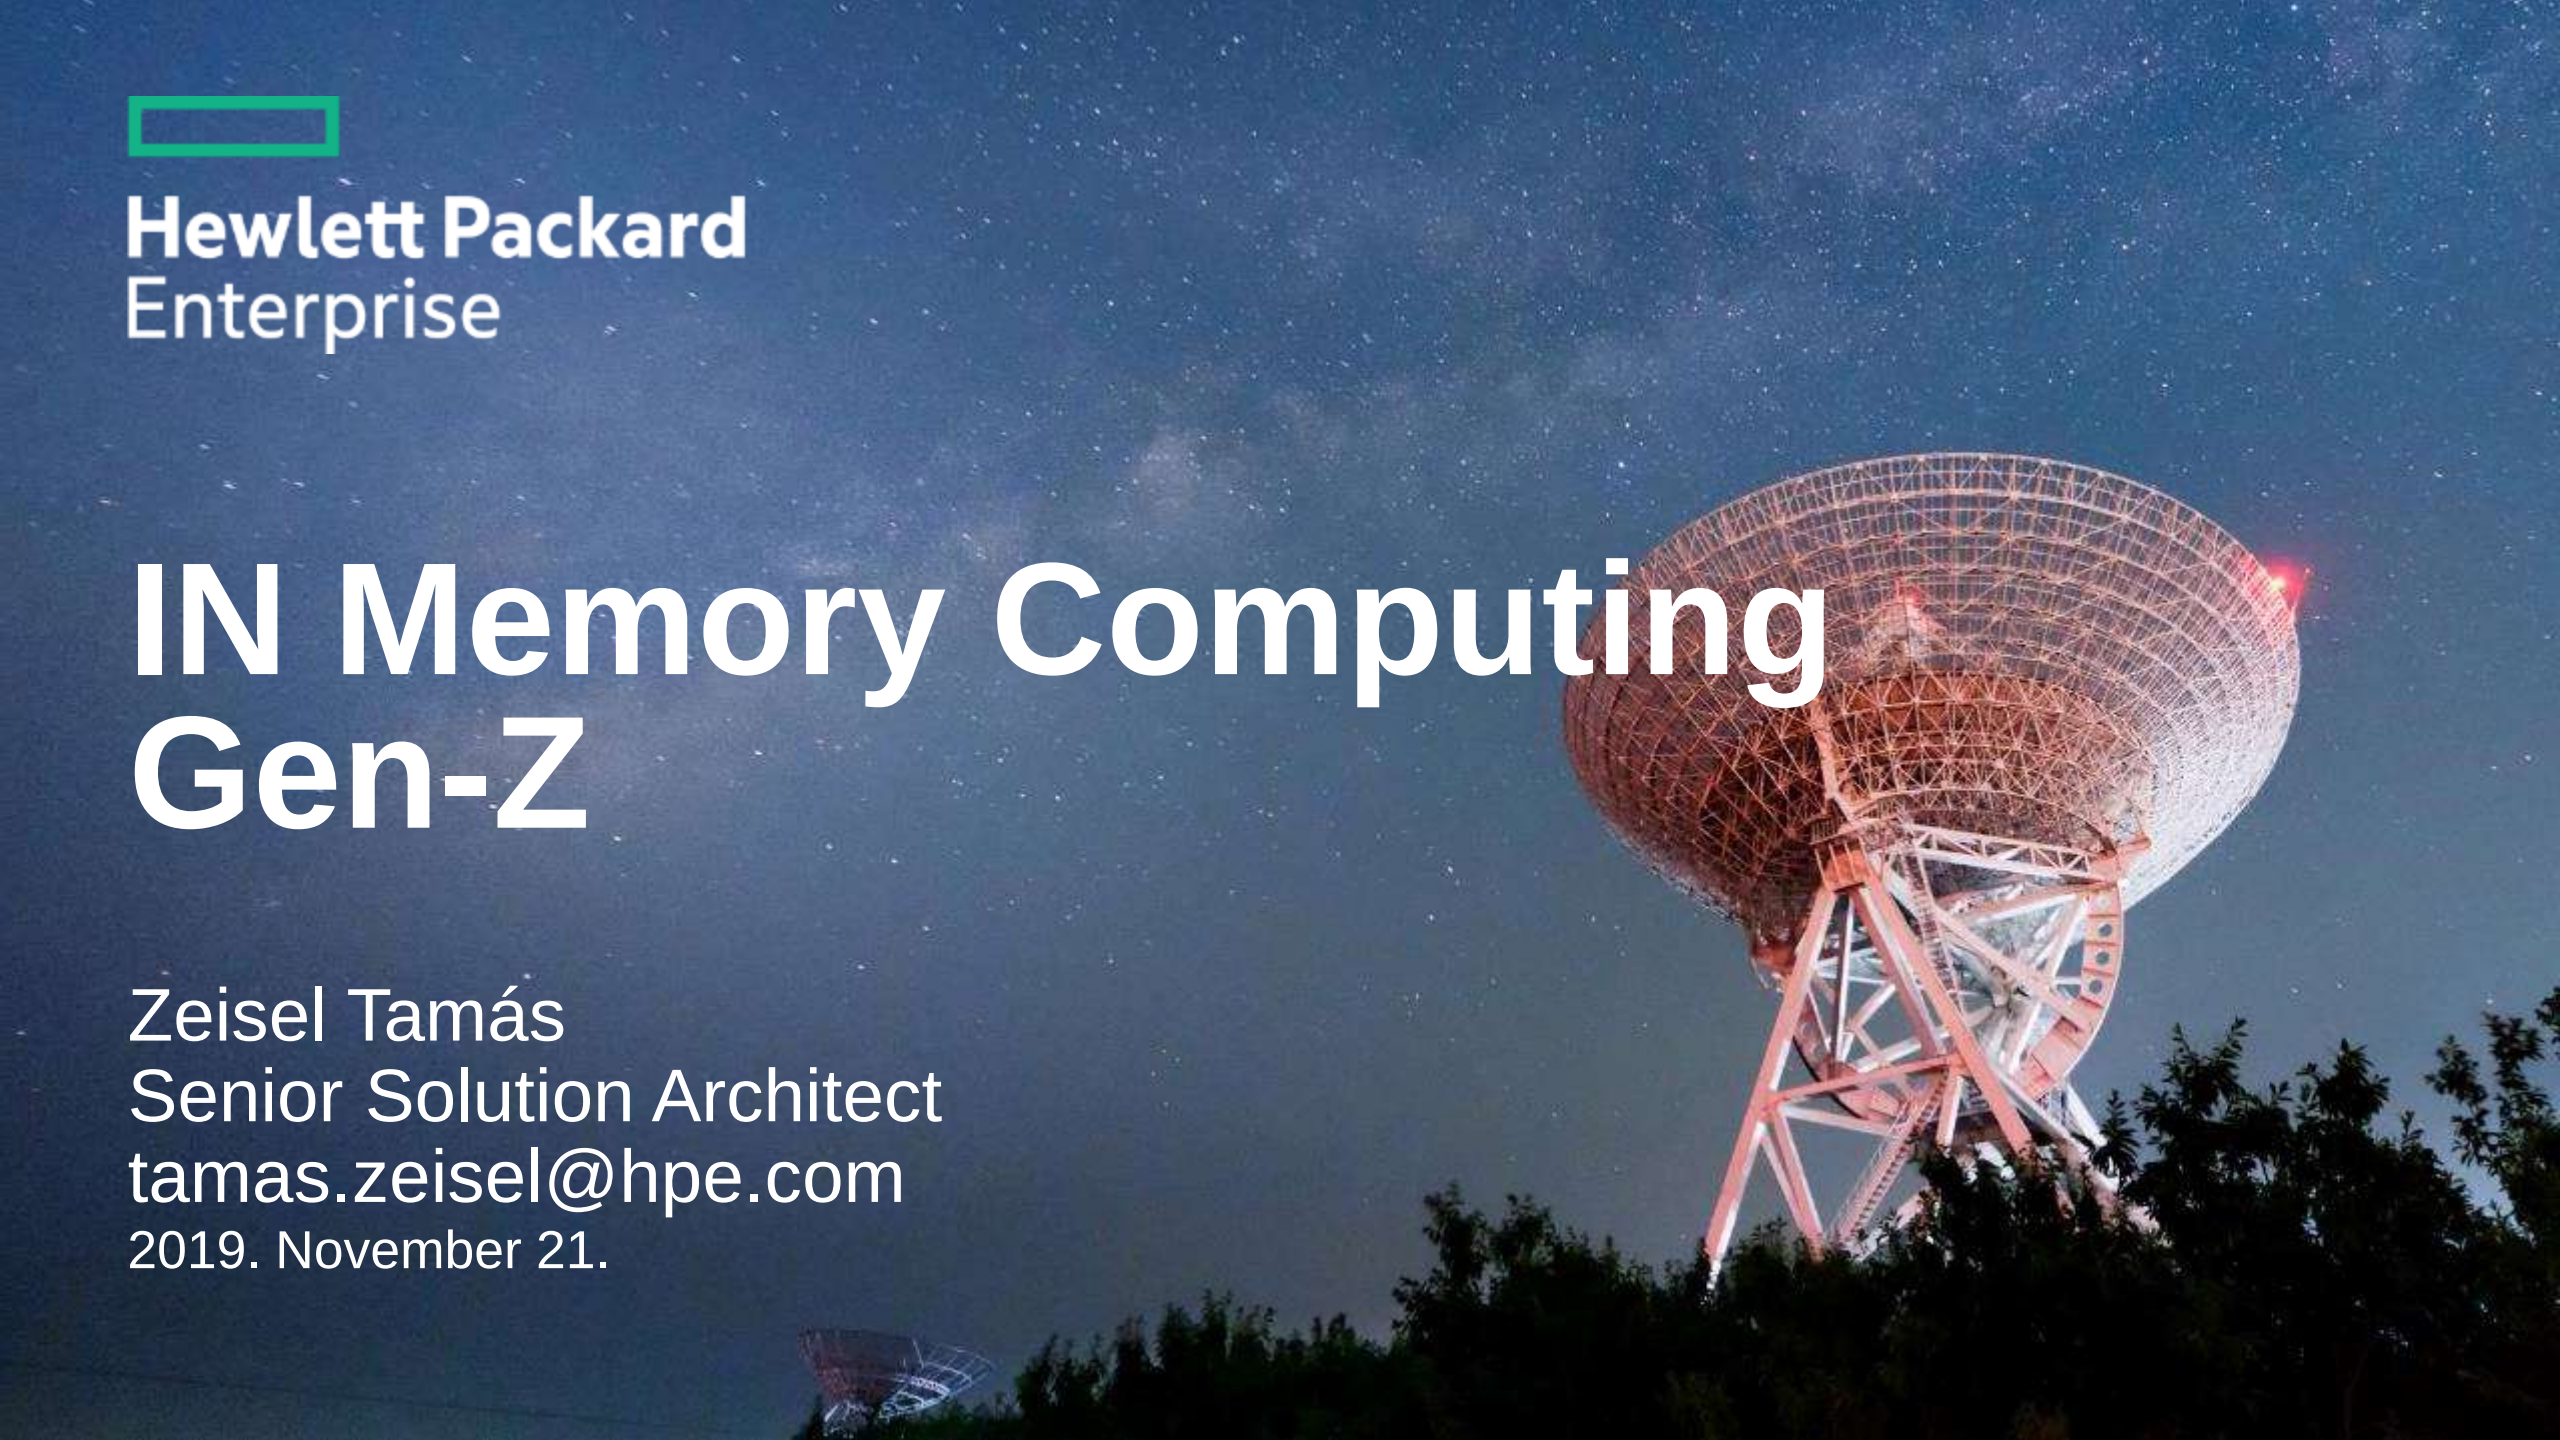

Hewlett Packard Enterprise

IN Memory Computing Gen-Z

Zeisel Tamás
Senior Solution Architect
tamas.zeisel@hpe.com
2019. November 21.

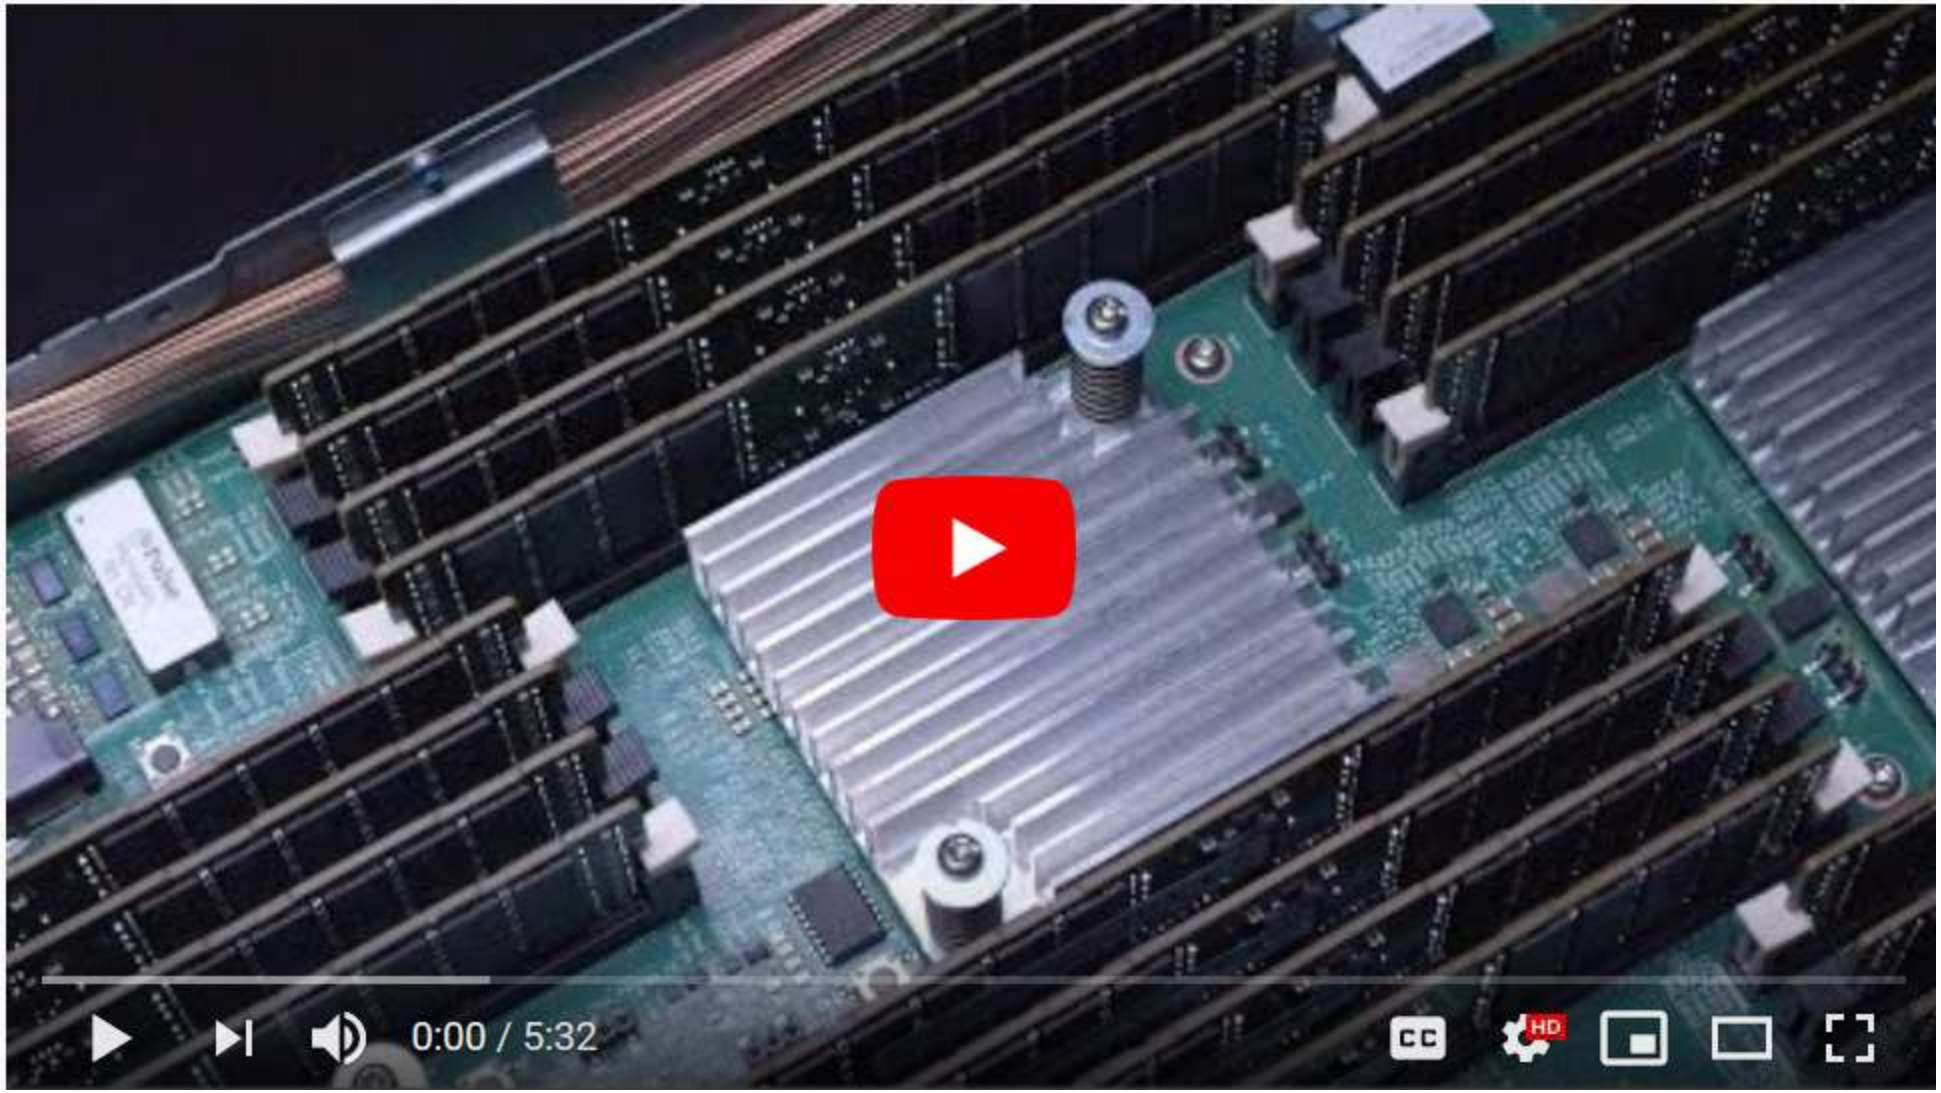

https://www.youtube.com/watch?v=2VG59FYkPdM#action=share

Traditional Computing Model

- Large data movement flows
  - Compute to Compute
  - Storage to Compute
- Data re-formatting
  - Byte-addressed, File, Block, Object
- Complex software
  - Large percentage of CPU moves data or waits
- Static workload placement

Memory-Driven, Data-Centric Computing

Powered by Gen-Z!

- Fabric-attached memory for active data
- Stop moving bulk data
  - Only access the data required
- Stop re-formatting data (block, file/object)
- Simplify the software (more CPU for workloads)
  - Create, Persist, Analyze all in one place, all in one format
- Enables dynamic workload placement
  - Move the workloads, not the data

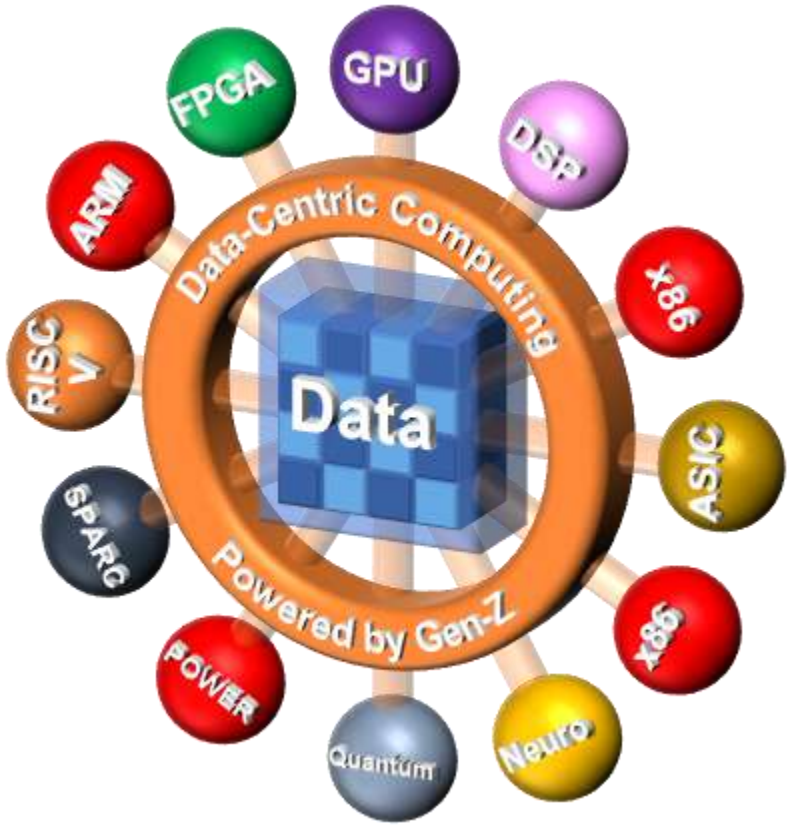

The Solution: Gen-Z!

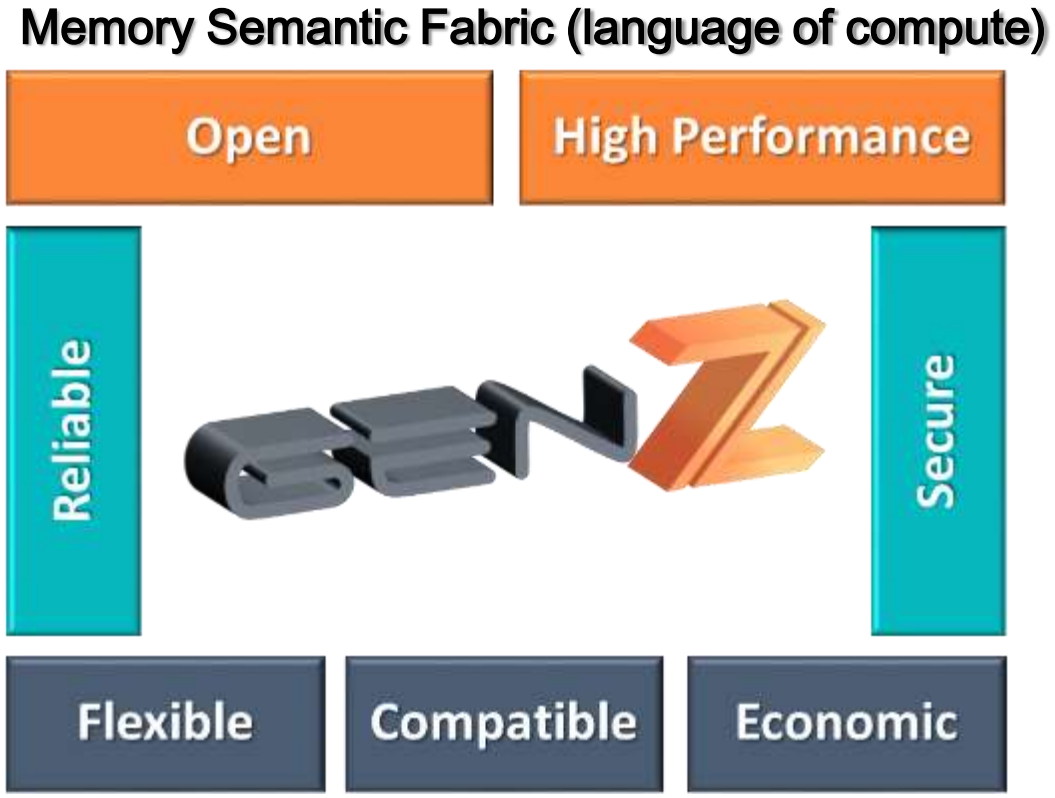

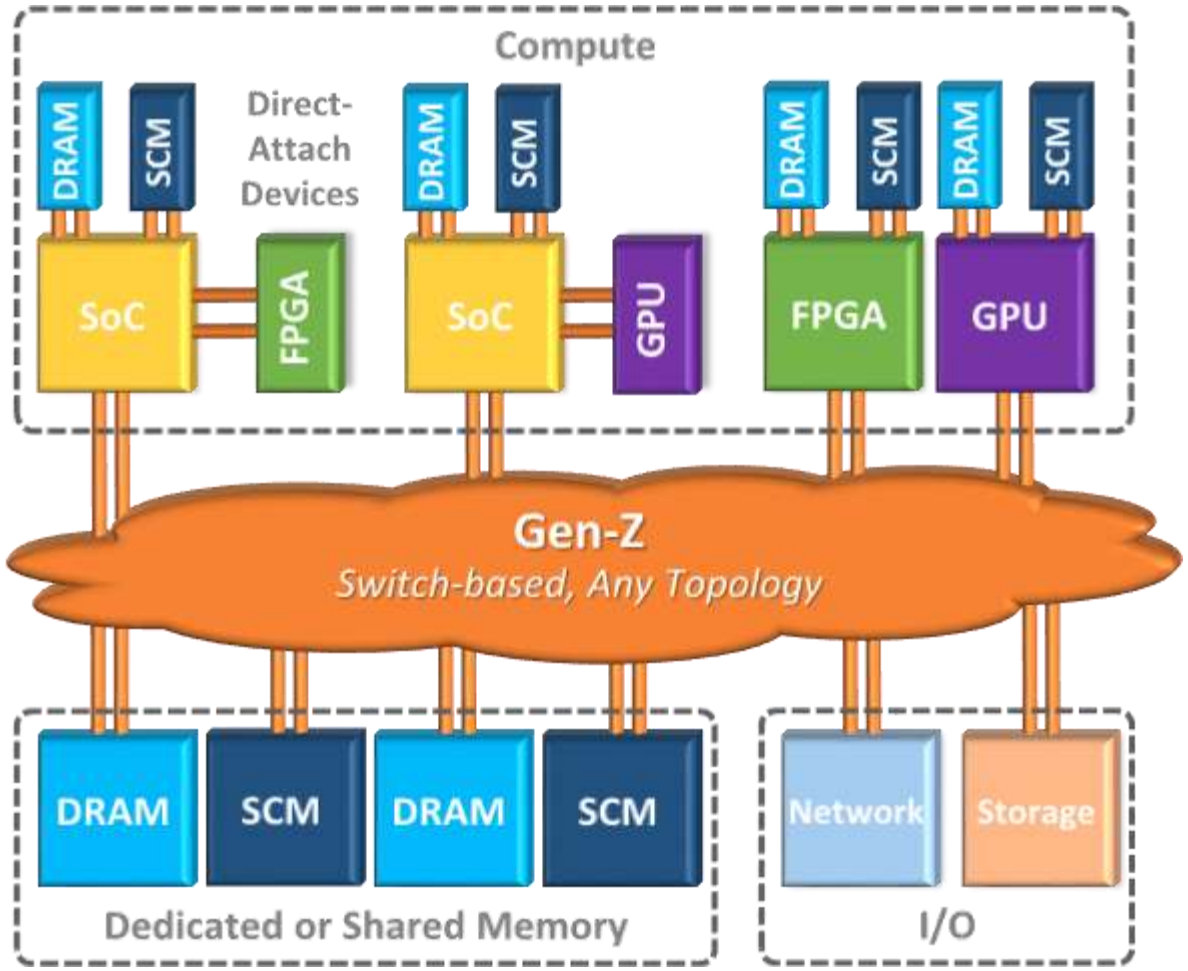

Hewlett Packard Enterprise

So What's an Example of How Gen-Z Helps Us?

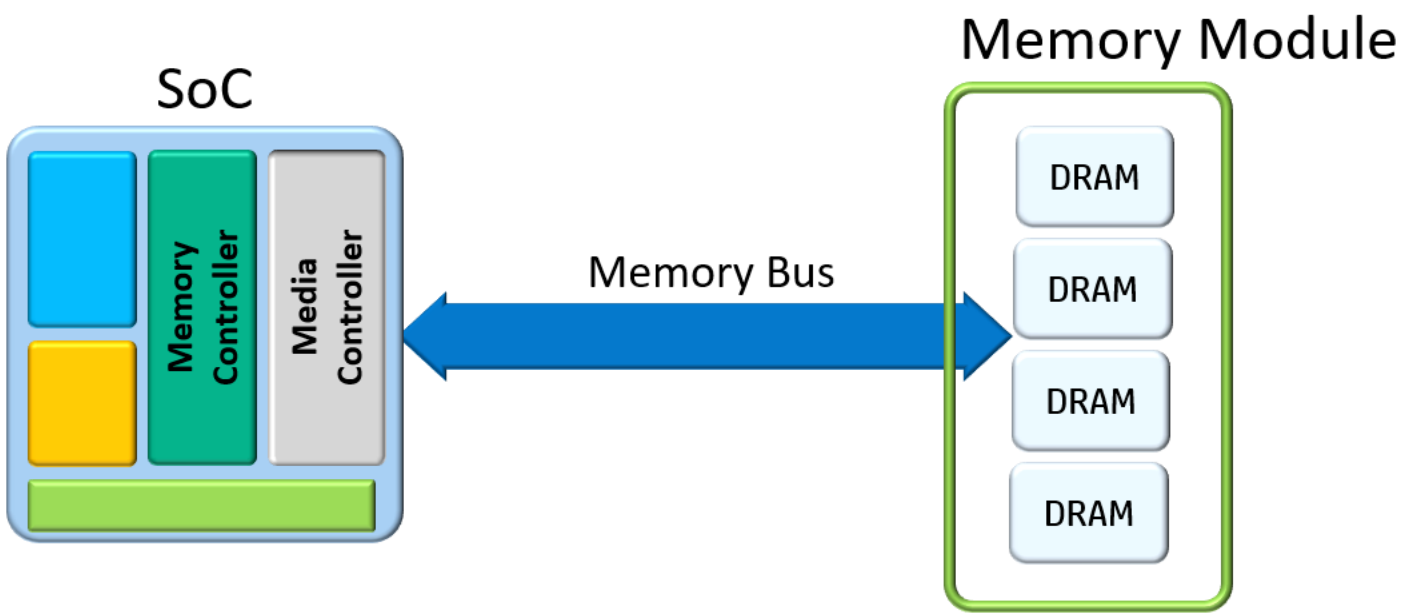

Traditional Memory Bus

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

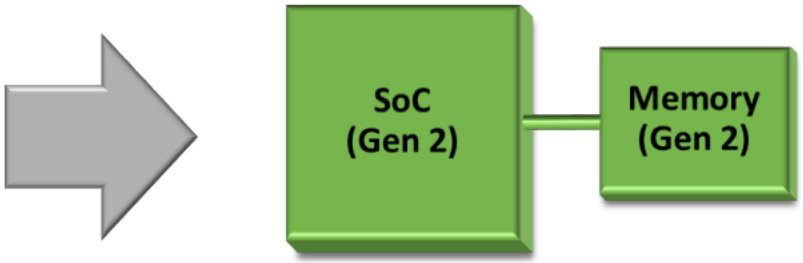

Traditional

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

Hewlett Packard Enterprise

So What's an Example of How Gen-Z Helps Us?

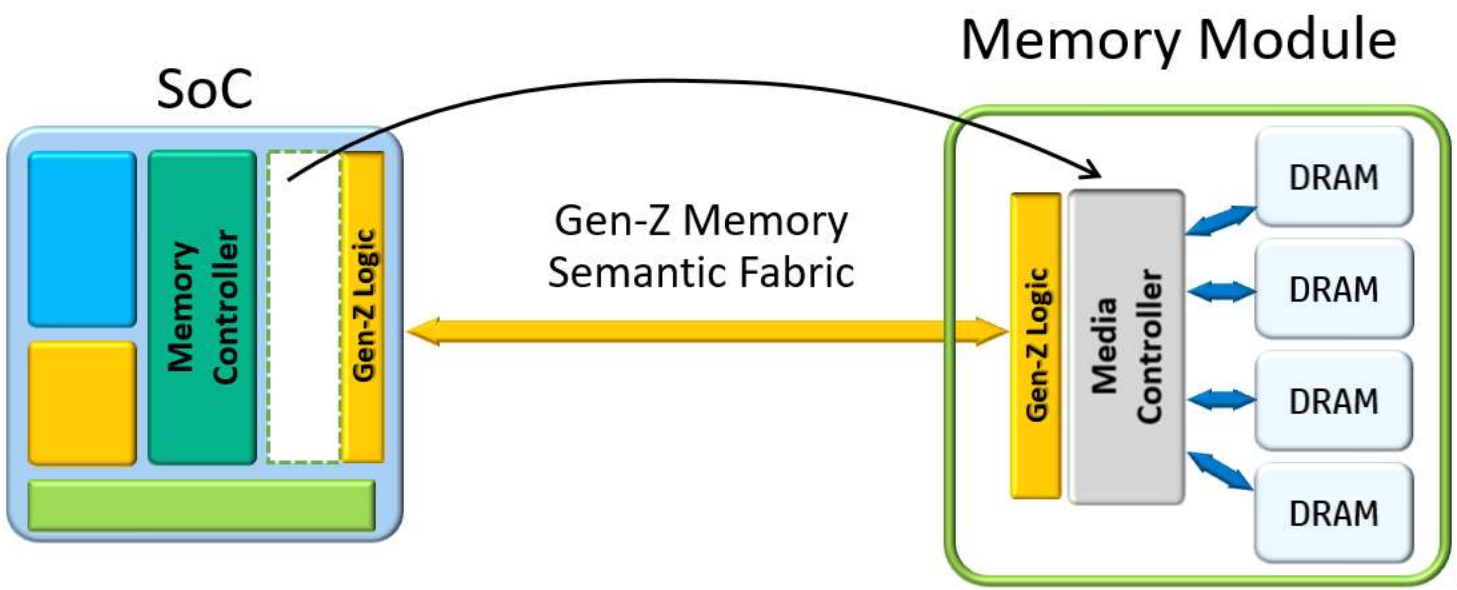

Gen-Z Memory Interconnect

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

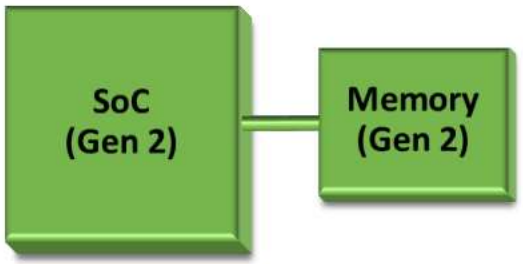

Traditional

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

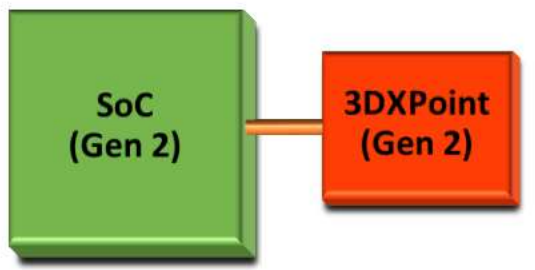

Gen-Z

- SoC & memory evolve freely
- Type & generation independent

Hewlett Packard Enterprise

Gen-Z is a high speed memory semantic fabric?

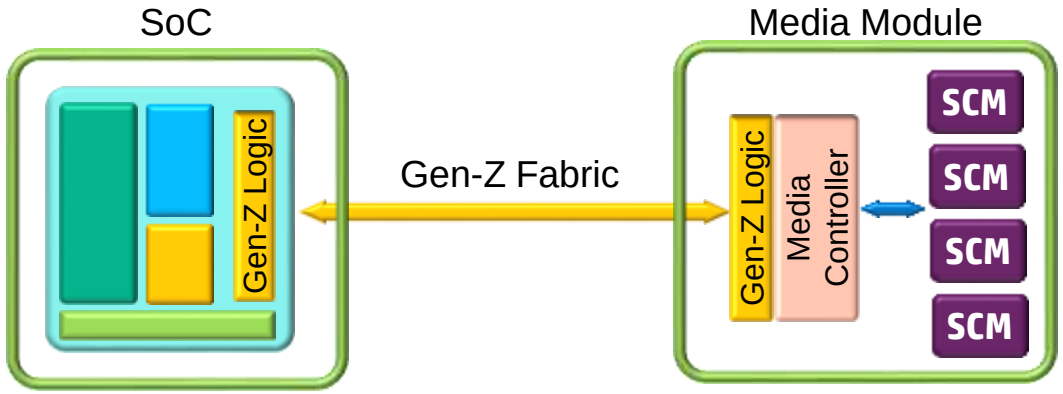

Storage Class Memory

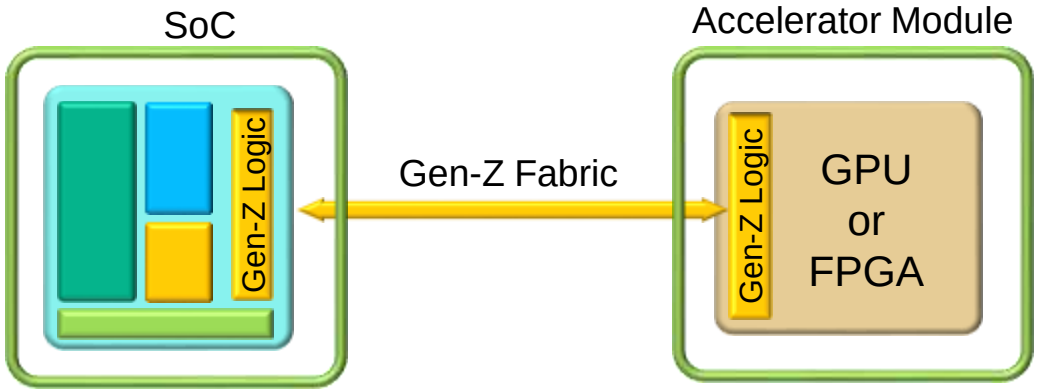

GPU or FPGA

- Supports DRAM, Flash, Memristor, PCRAM, MRAM, 3D-Xpoint…**Universal Interconnect**
- Decouples CPU/memory design
- Enables independent innovation

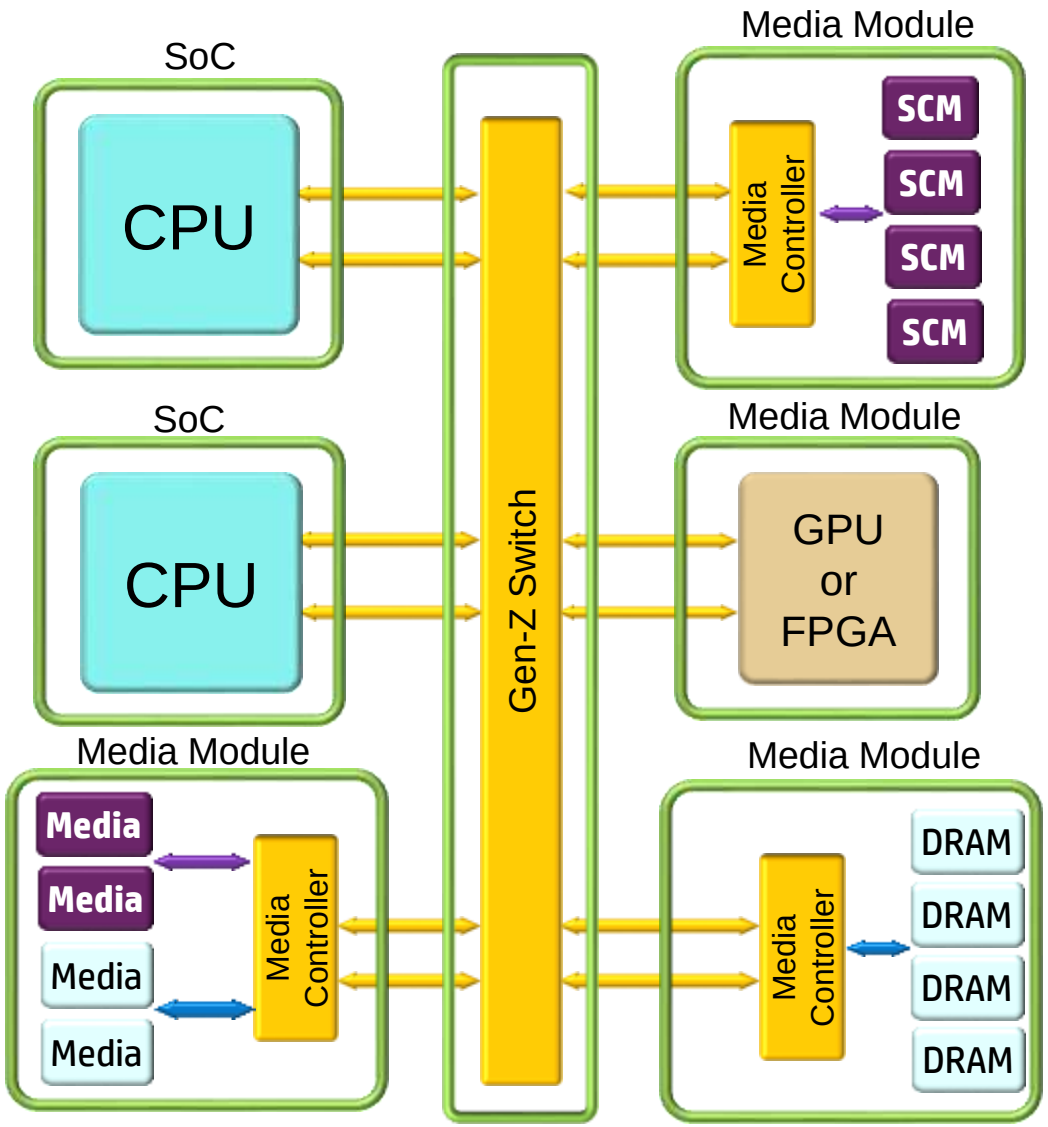

Multiple resources enabled by Universal Interconnect

Hewlett Packard Enterprise

Open Consortium with Broad Industry Support (65)

GEN Z Consortium Members

System OEM	CPU/Accel	Mem/Storage	Silicon	IP	Connect	Software
Cisco	AMD	Everspin	Broadcom	Avery	Aces	Redhat
Cray	Arm	Micron	IDT	Cadence	AMP	VMware
Dell EMC	IBM	Samsung	Marvell	Intelliprop	FIT	
H3C	Qualcomm	Seagate	Mellanox	Mentor	Genesis	**Govt/Univ**
Hitachi	Xilinx	SK Hynix	Microsemi	Mobiveil	Jess Link	ETRI
HP		Smart Modular	Sony Semi	PLDA	Lotes	Oak Ridge
HPE		Spintransfer		Synopsys	Luxshare	Simula
Huawei		Toshiba			Molex	UNH
Lenovo		WD			Samtec	Yonsei U
NetApp					Senko	ITT Madras
Nokia		**Technology SP**		**Eco/Test**	TE	
Yadro		Google		Allion Labs	3M	
		Microsoft		Keysight		
		Node Haven		Teledyne LeCroy		

Memory-Driven Computing: Transition to Photonics

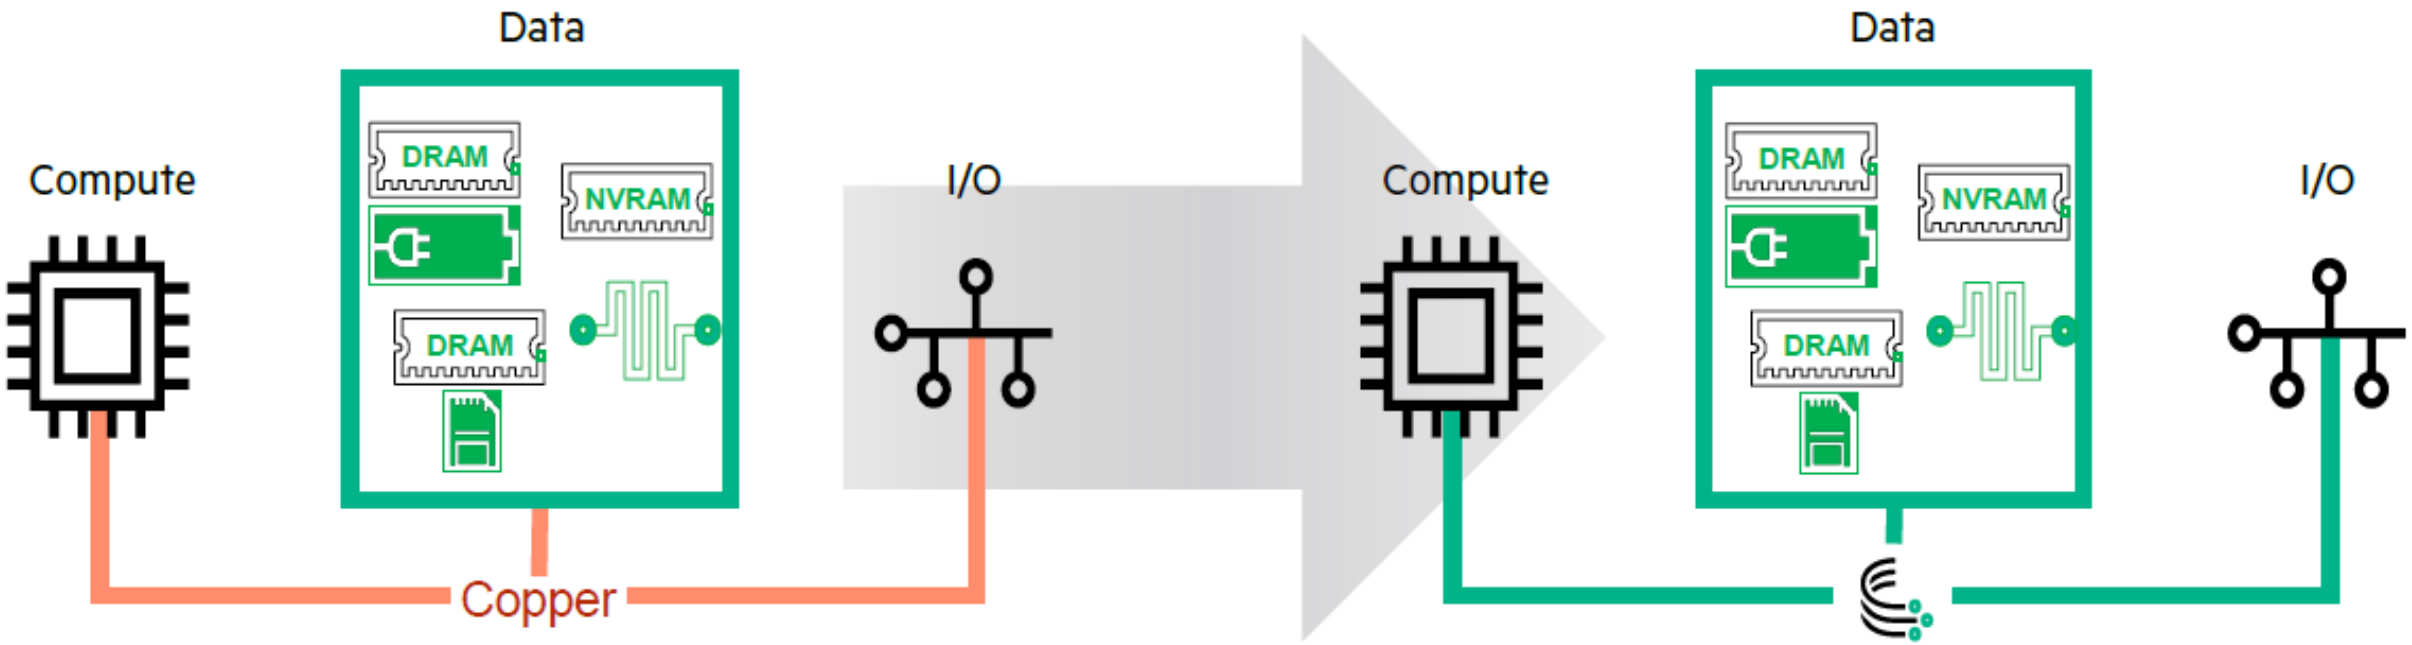

Hewlett Packard Enterprise

Levels of Optical Integration

Optical interfaces are migrating towards silicon die for advantages in:

- Cost
- Power
- Performance

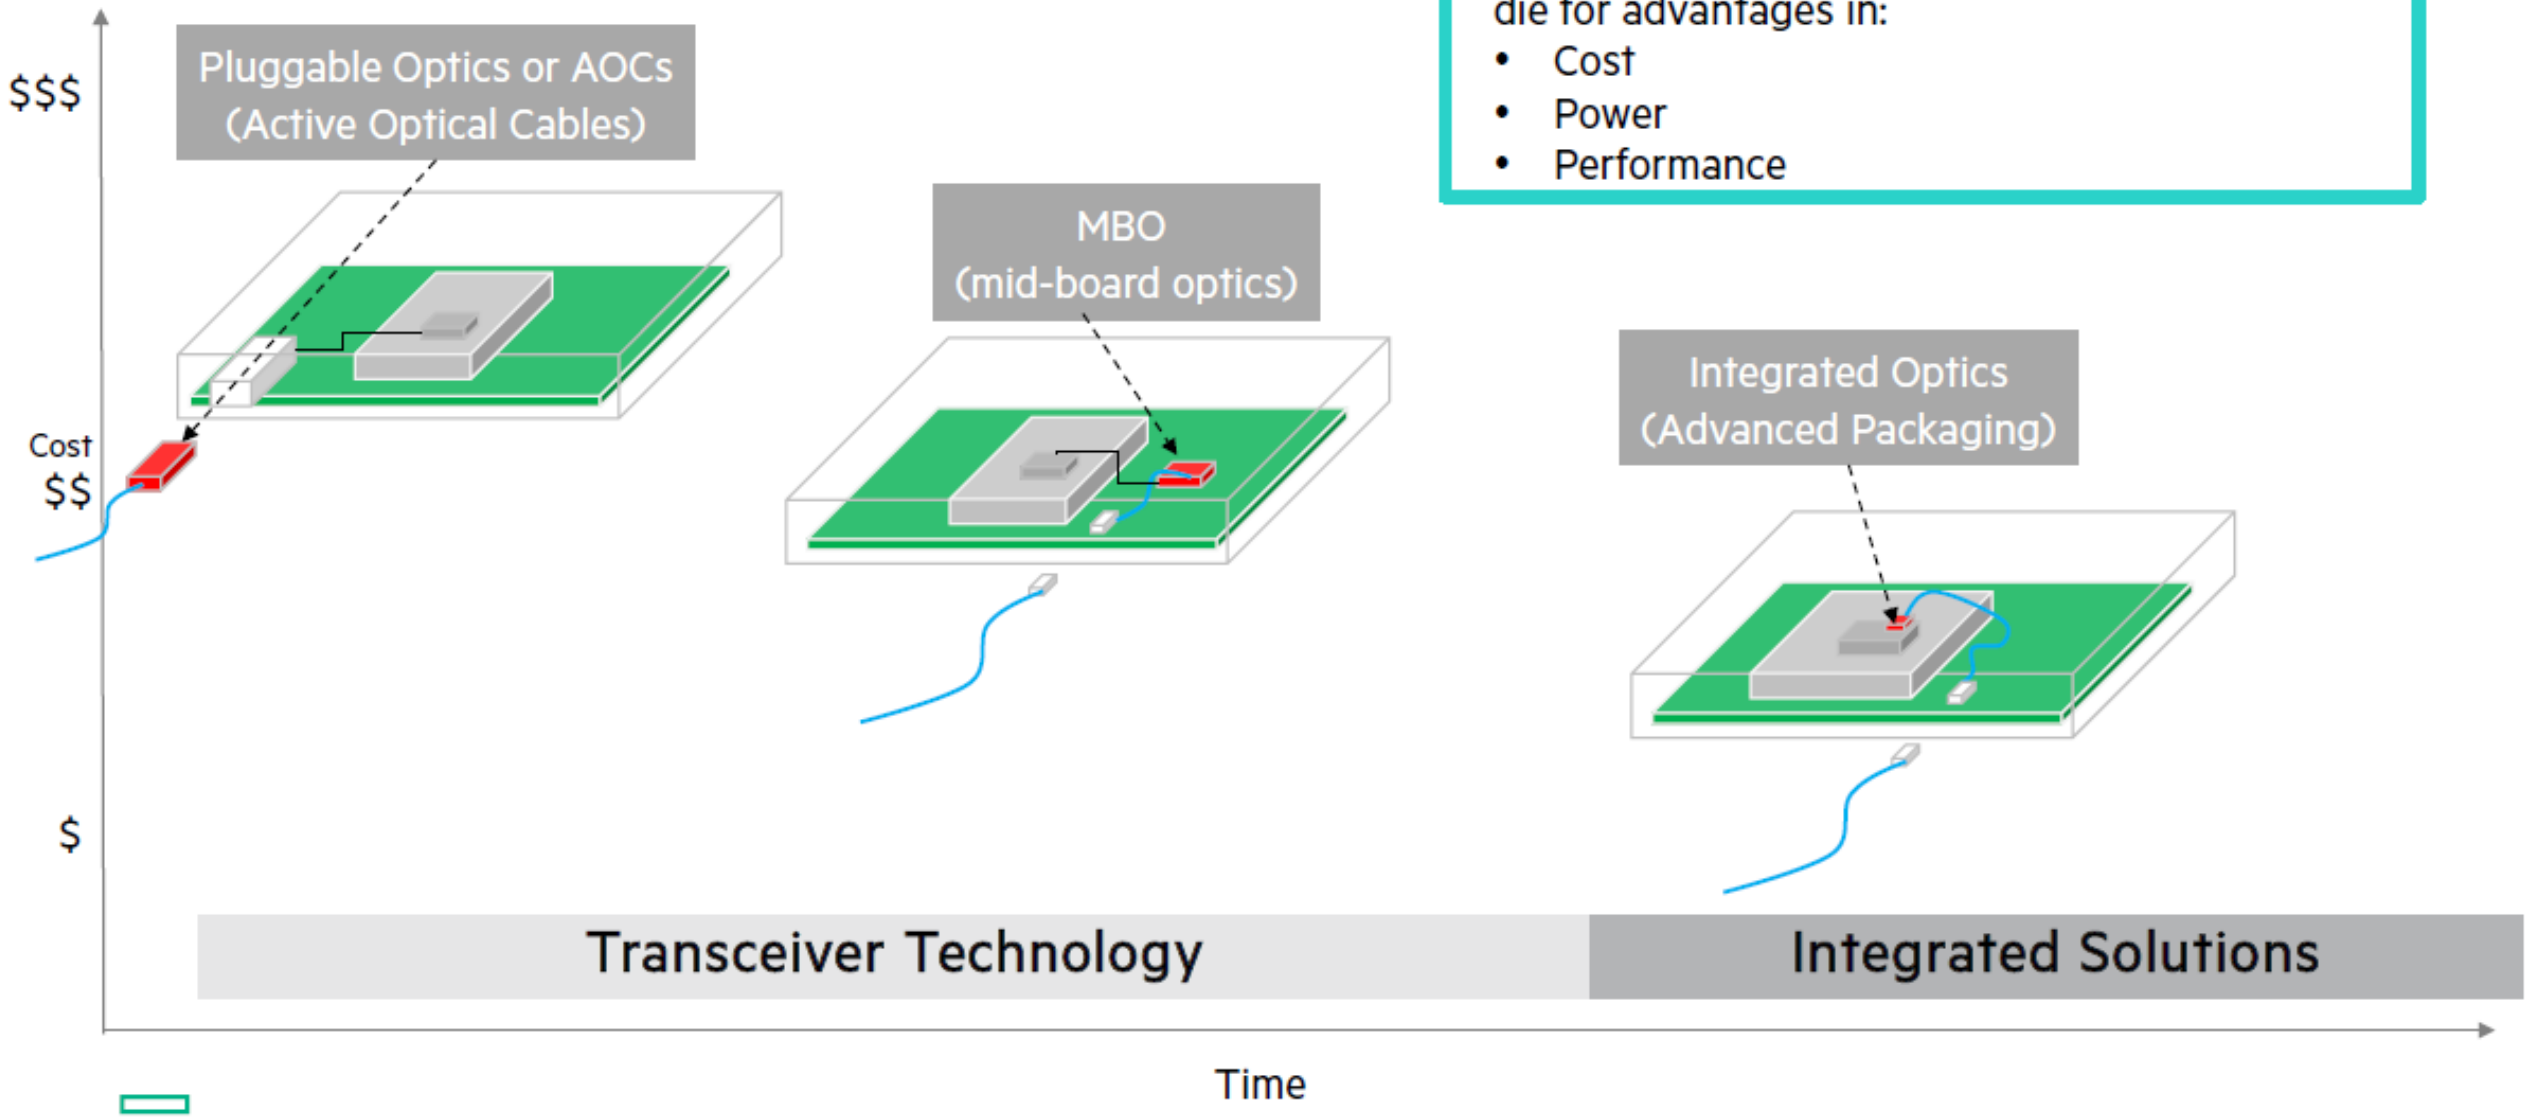

Hewlett Packard Enterprise

Advantage of HPE Integrated Optics

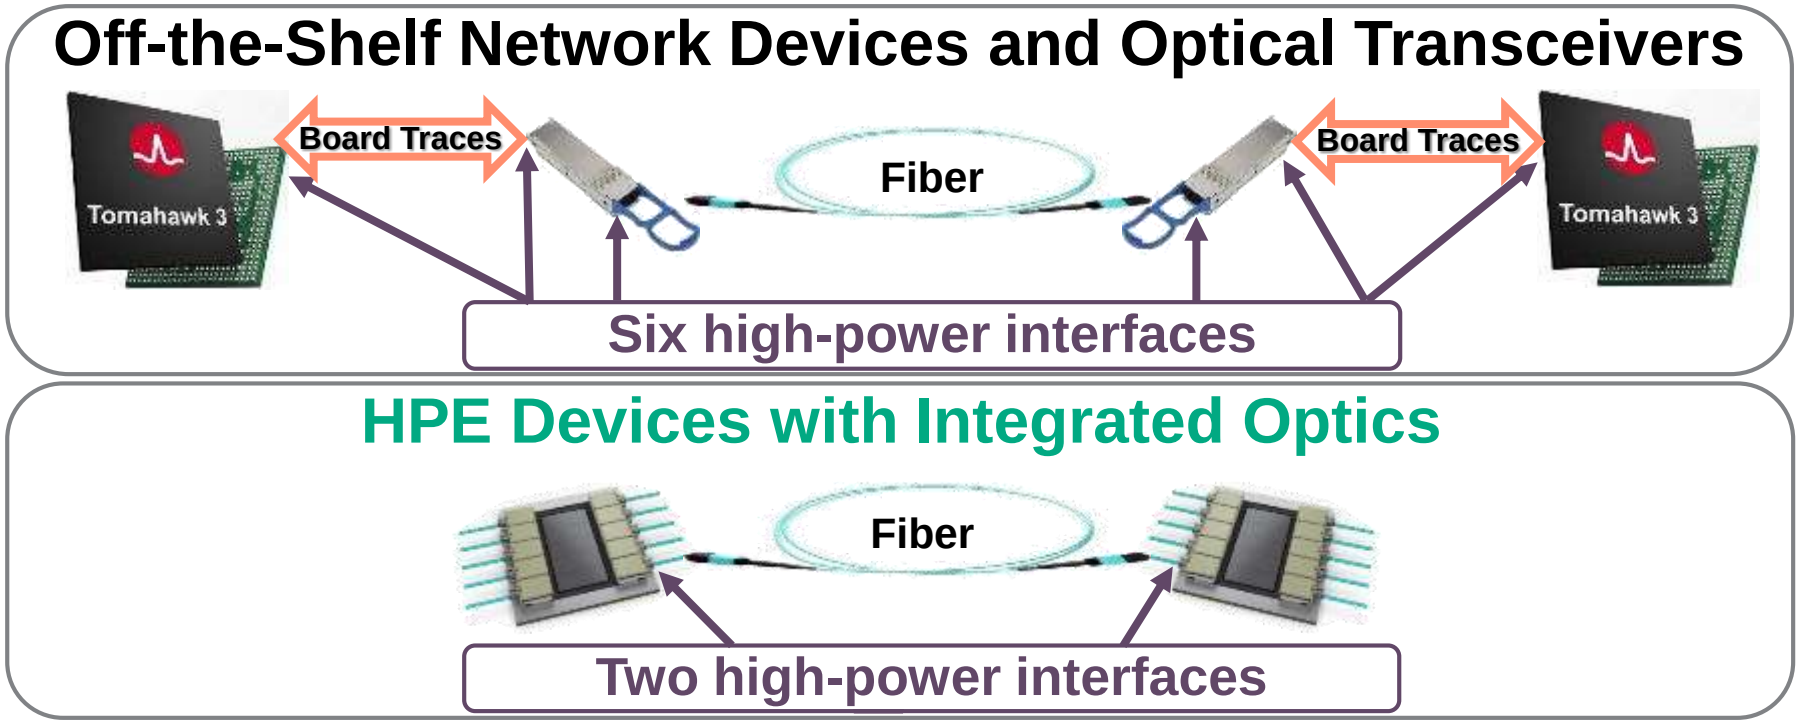

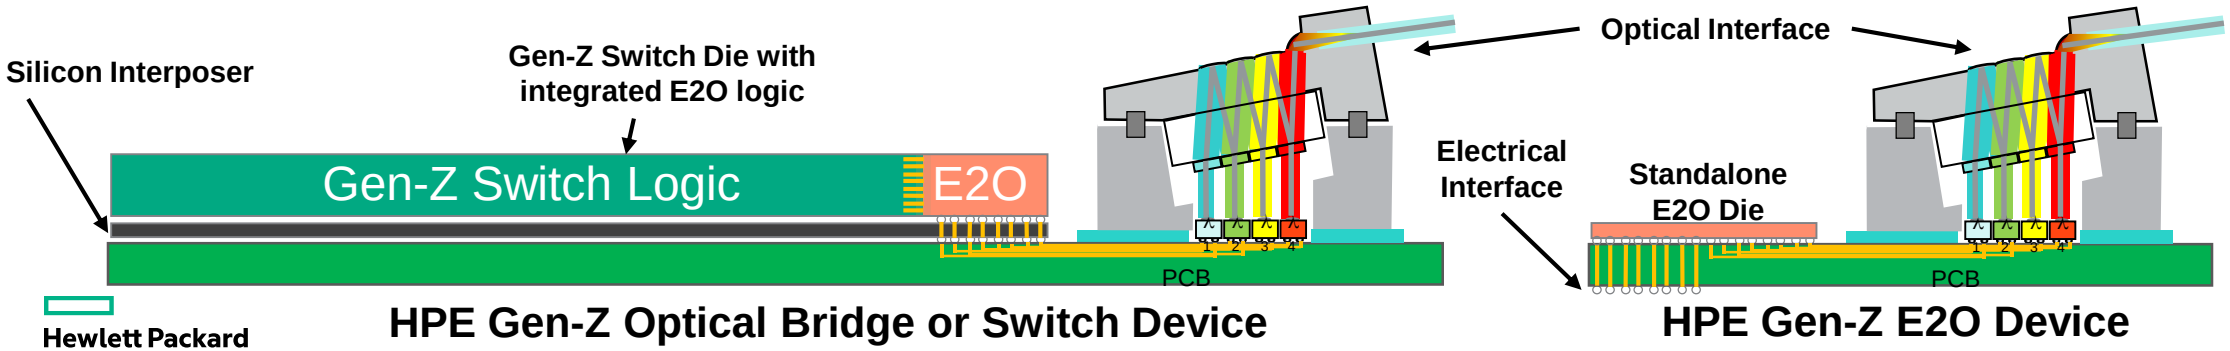

Hewlett Packard Enterprise

HPE's Integrated Optics: Increased Bandwidth per Fiber

CWDM-4
Coarse Wave Division Multiplexing

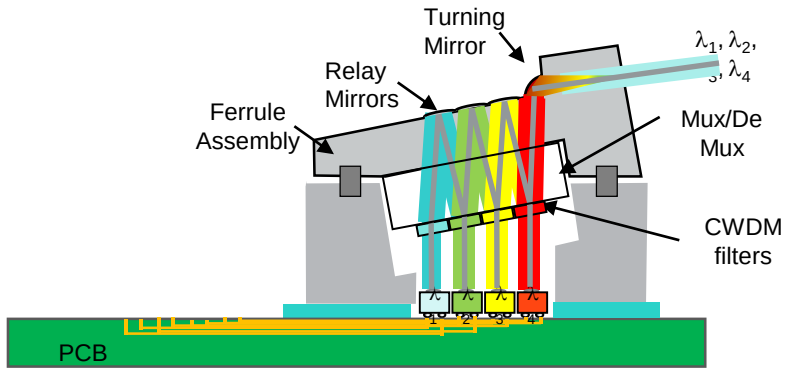

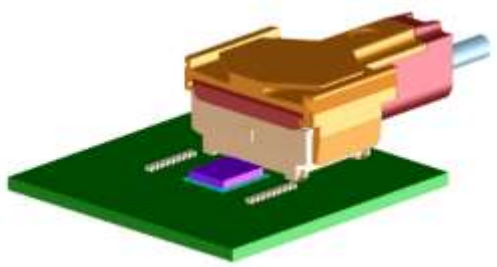

4x signals per fiber

PAM-4
Pulse Amplitude Modulation

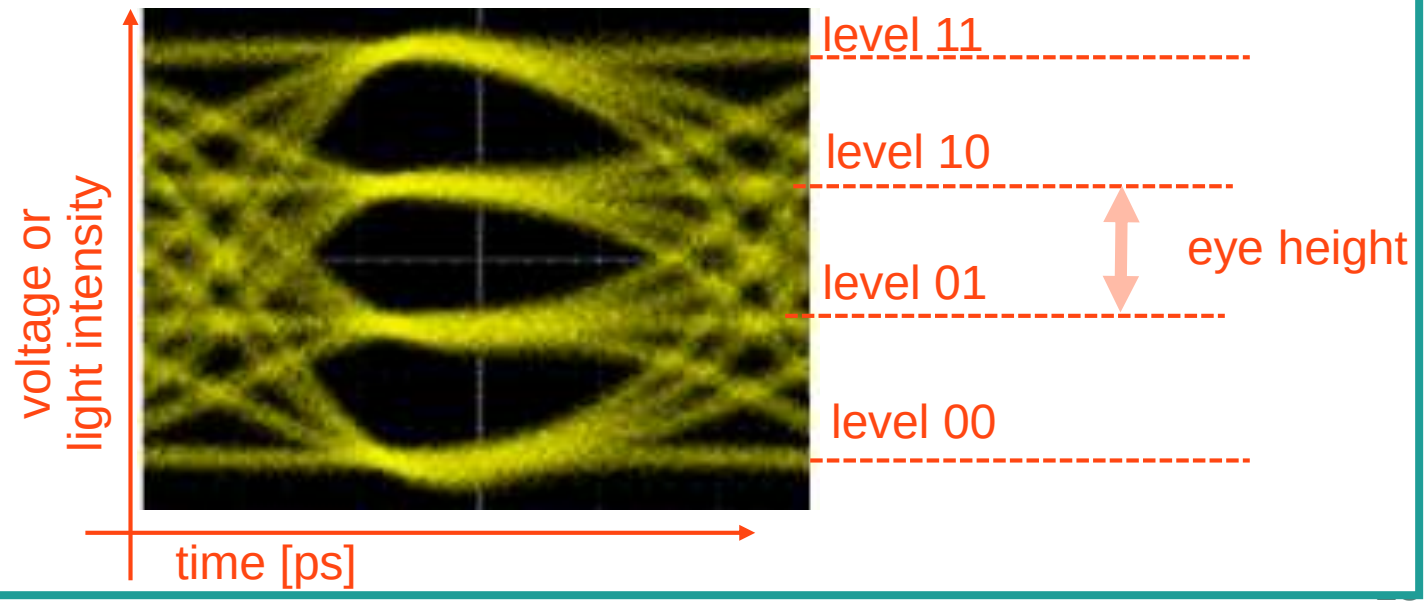

2x signals per fiber

Hewlett Packard Enterprise

HPE Integrated Optics

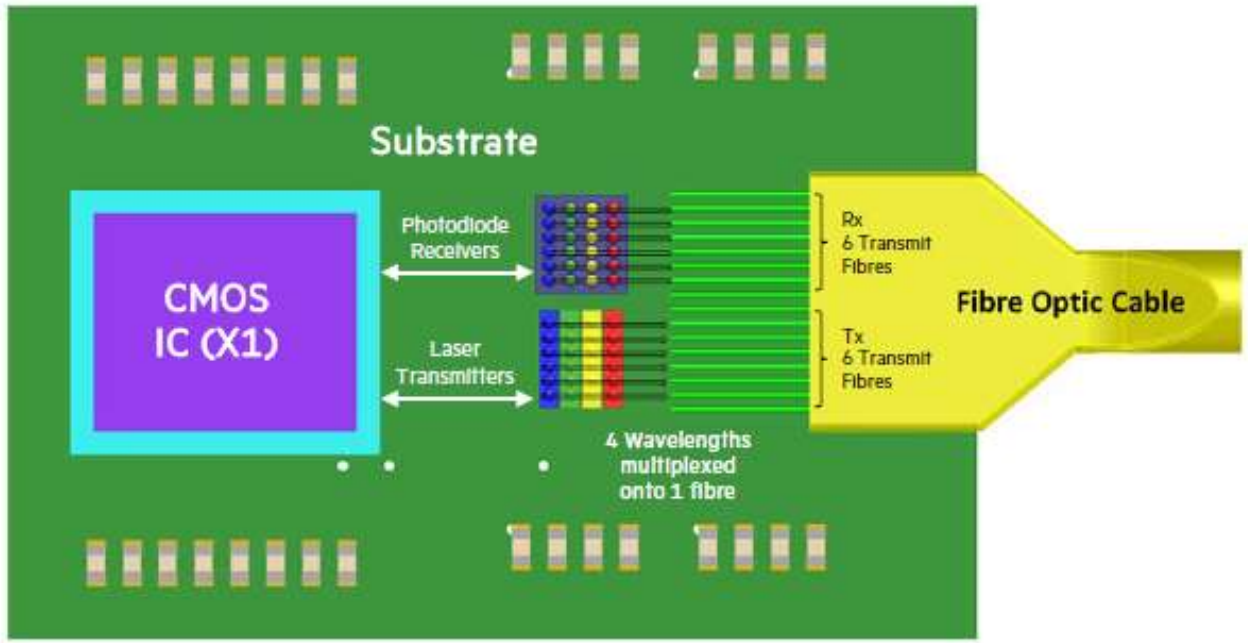

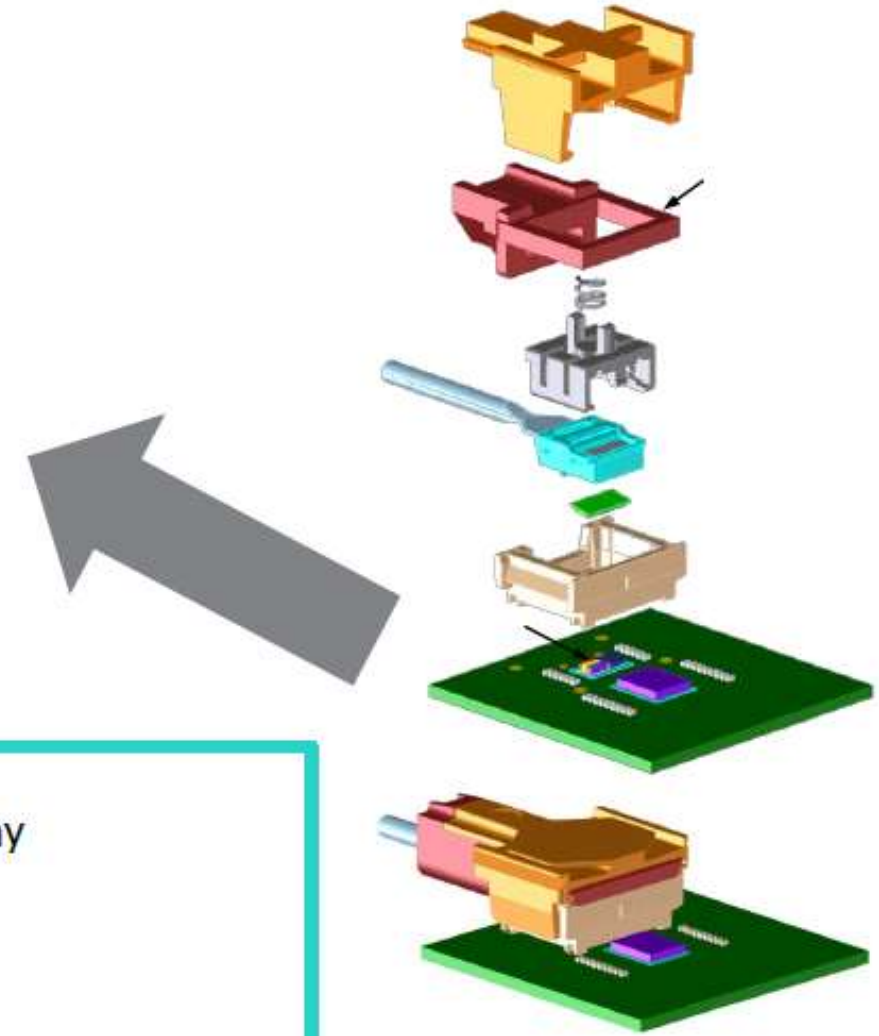

1. 4 λ Coarse Wavelength Division Multiplexing (CWDM)
   - Vertical-cavity surface-emitting laser (VCSEL) array
   - Photodiode (PD) array
2. CMOS Integrated Circuits
   - Optical (VCSEL) Direct Drive + PD Receiver
3. Micro-Optics socket attach and passive alignment
   - Integrated CWDM Mux/Demux filters and lenses

Hewlett Packard Enterprise

Innovation: 3 types of HPE Gen-Z chiplets on a single reticle

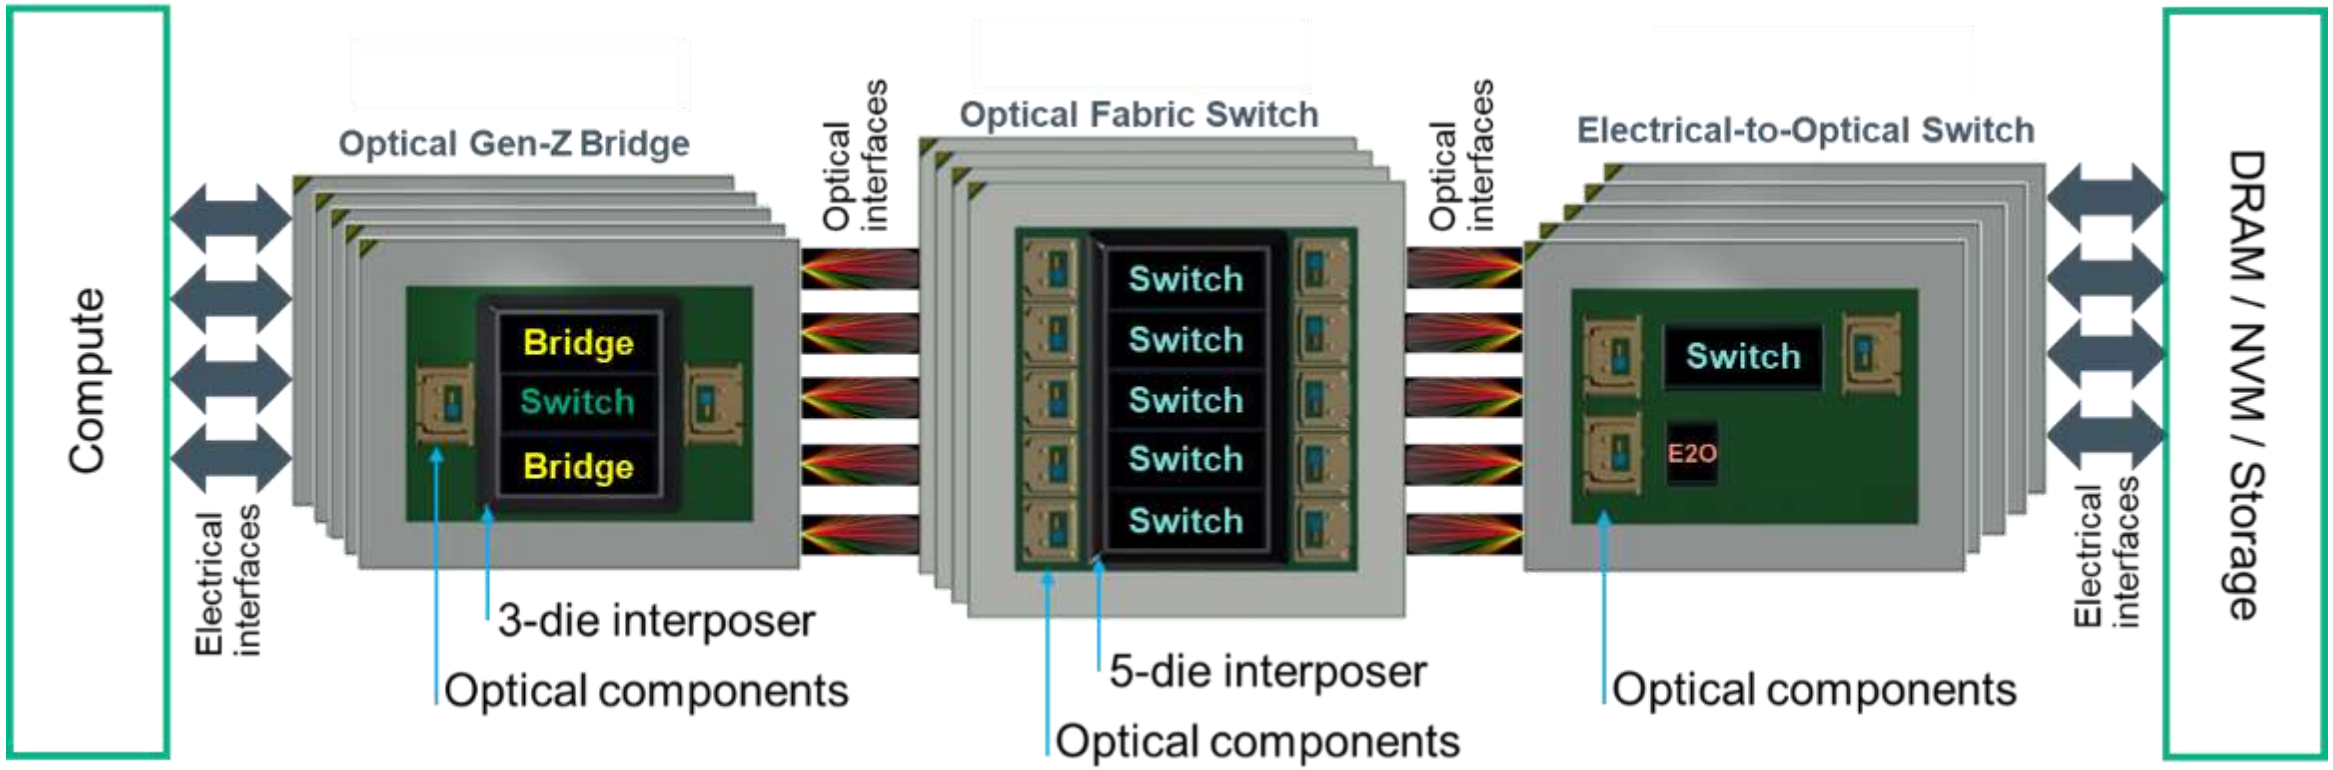

Gen-Z devices built as multi-chiplet modules (MCMs)

- Flexible system design
- Enables larger devices (e.g. switches) than single die
- Much lower development costs

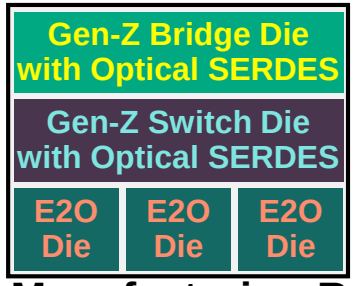

Chip Manufacturing Reticle

Hewlett Packard Enterprise

Gen-Z Optical Bridge

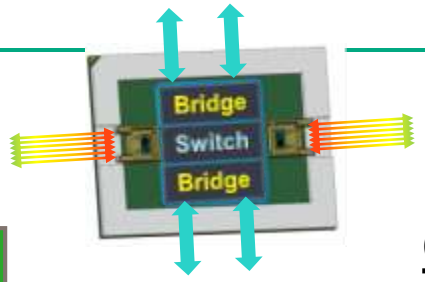

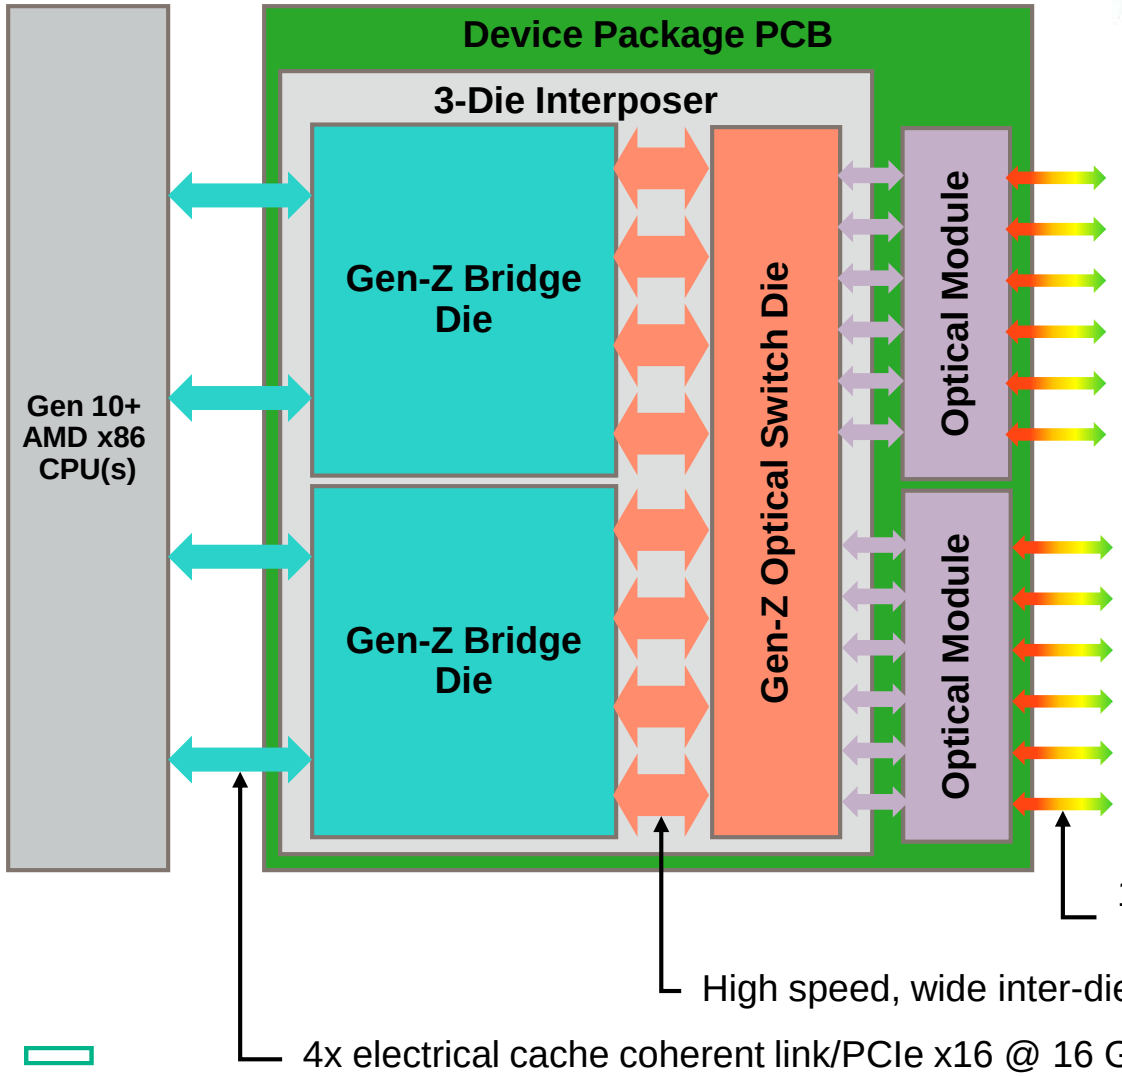

Gen-Z Bridge Operational Features

- Gen-Z **Core 64** OpClass Support
  - Read, Write, Write Partial, Write Poison, Persistent Flush, Interrupt, No-Op, Non-Idempotent Request Release
- Gen-Z **Context ID** OpClass Support
  - Reliable Write MSG, Non-Idempotent Request Release
- Gen-Z **Control** OpClass
  - Read, Write, Interrupt, No-Op, Unsolicited Event
- Gen-Z **Atomic** OpClass
  - Add, Swap, CAS, Logical AND, Logical OR, Logical XOR, Load Max, Load Min, Atomic Fetch
- Gen-Z Requester and Responder **ZMMUs**
- Gen-Z Transmit and Receive **Data Movers**

Hewlett Packard Enterprise

Building the PathForward System with Badger & Gen-Z

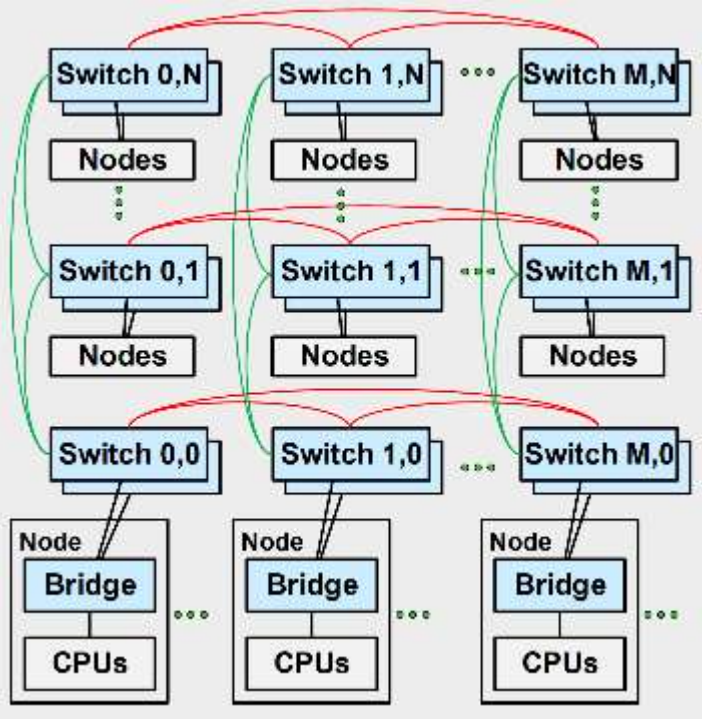

2D HyperX System Topology Optical Fabric

Gen-Z Switch Line Card

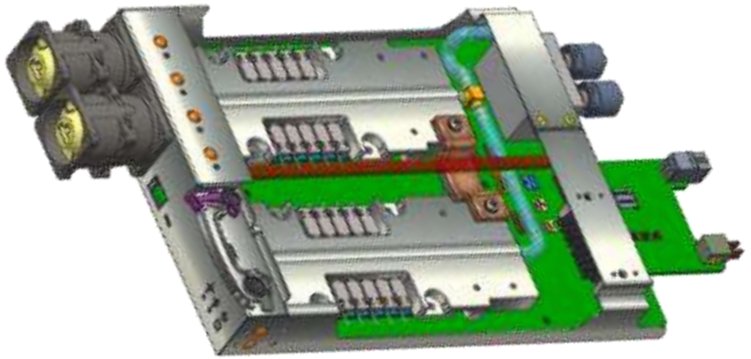

Gen-Z Connected x86 Compute Node

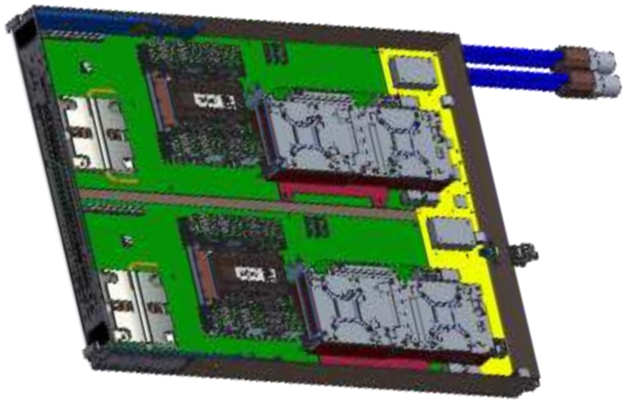

Gen-Z Connected I/O Node (Storage)

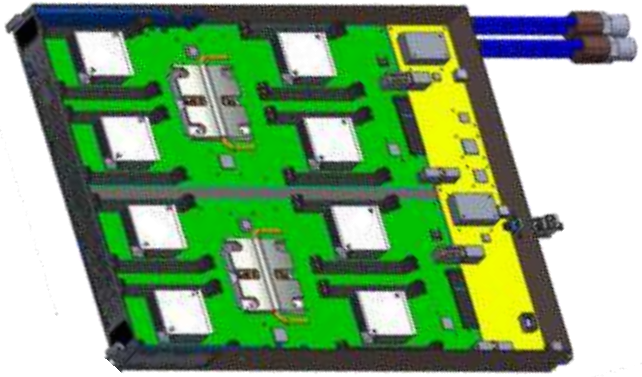

Hewlett Packard Enterprise

DL385 Gen10+ Gen-Z Development Kit Platform (DKP)

- Built with next generation Gen10+ AMD CPUs
  - New cache coherent link technology
    - Used to connect to the Gen-Z Optical Bridge
- Up to two Gen-Z Bridge Card assemblies
- Optional PCIe to Gen-Z management sidecar
  - To enable fabric and resource management development
- New front drive bay/midplane supports Gen-Z peripherals
  - ZFF DRAM Modules, PM Modules, NVMz Drives, etc.

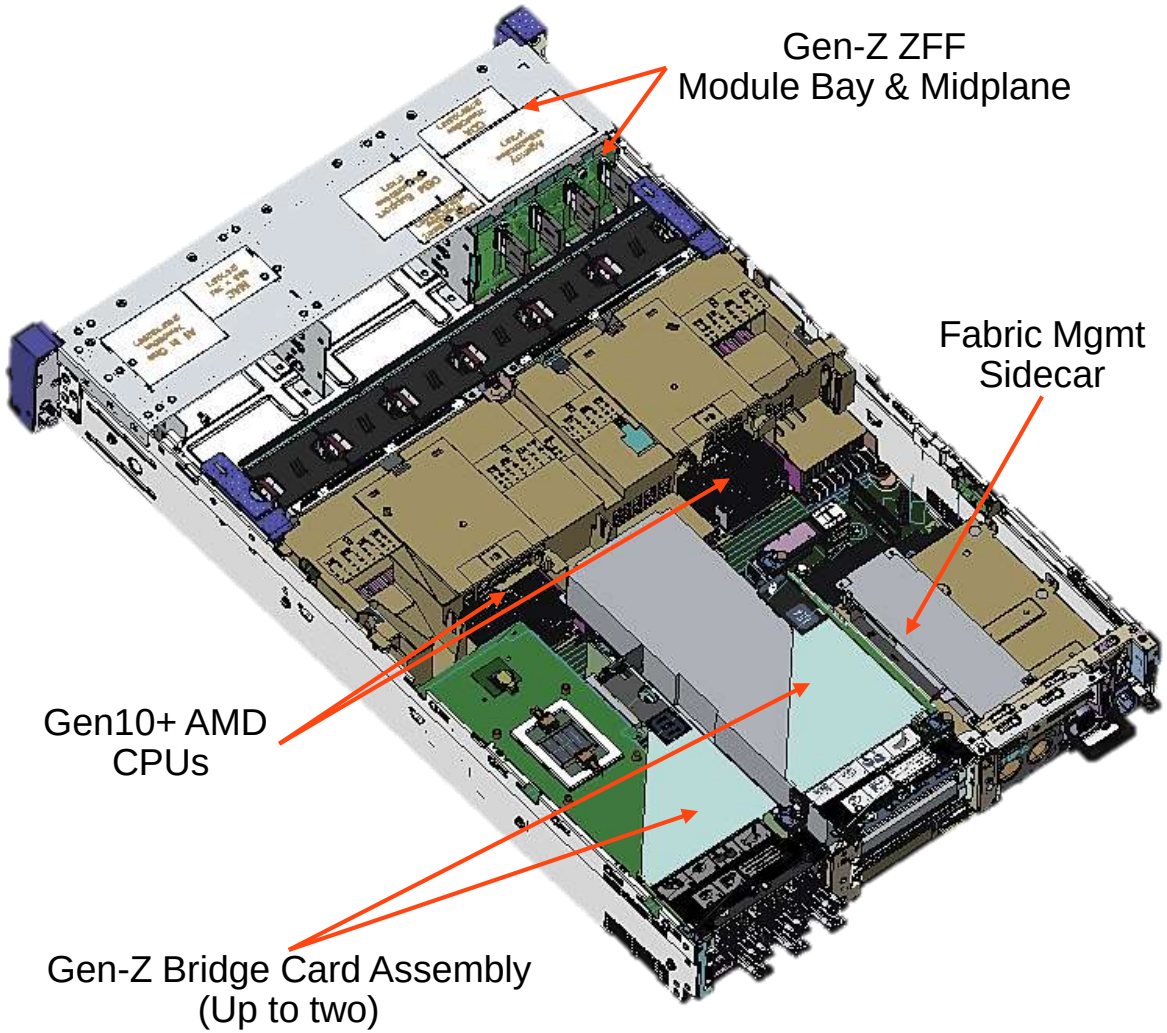

Hewlett Packard Enterprise

ZFF DRAM Memory Module

- 256GB Capacity
- Gen-Z Interfaces
  - SFF-TA-1002 4C card edge with four x4 Gen-Z links @ 25Gb/s
- Two links for host connection, 2 links for daisy chain

Mockup of DL385 bay with ZFF Modules

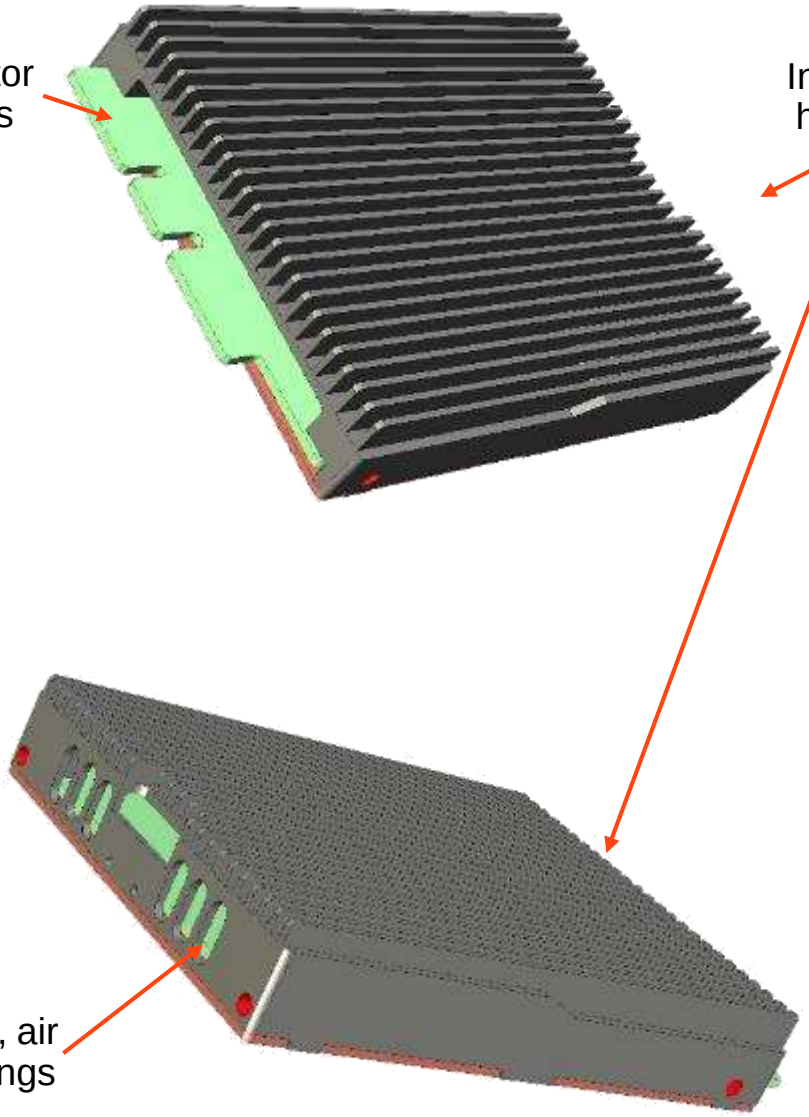

Hewlett Packard Enterprise

DL385 Gen-Z DKP Use Cases

Software Development

Firmware, OS, mgmt dev

In-memory application dev

High Performance Messaging

HPC MPI/OpenMPI app dev

HPC middleware/Libfabric dev

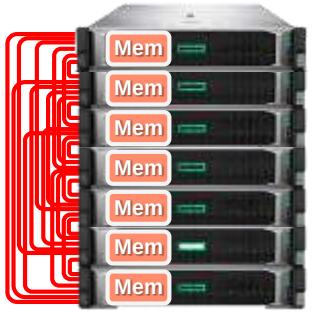

Composable Memory

Fabric management dev

Memory resource mgmt dev

Composable, in-memory app

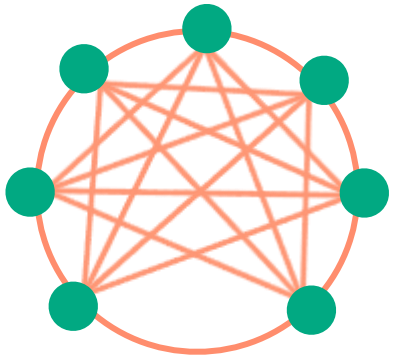

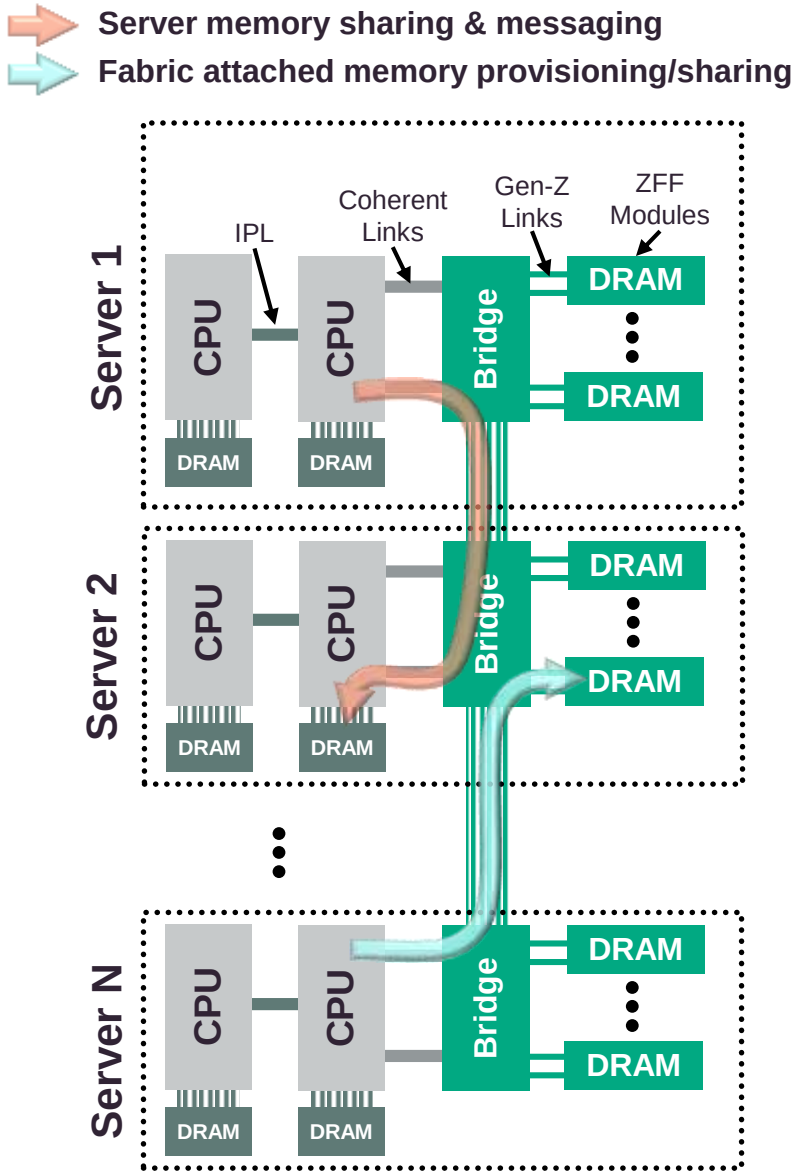

Hewlett Packard Enterprise

Hewlett Packard
Enterprise

Thank you

Adatközponti Rendszermérnök képzés

Időtartam

Indulás: 2020 március
Időtartam: 6 hónap – 376 kontaktóra
Időbeosztás:
– Két hétköznap: 17:00- 21:15
– Egy hétvégi nap: 9:00 – 18:00
Előképzettség: Nem szükséges, de fontos, hogy érdeklődj az informatika iránt és rendszerszemlélettel rendelkezz.

Miért minket válassz?

Jelenleg hazánkban egyedülállóan a Braining Hub képzési centrum kínál olyan intenzív programozó képzéseket, amelyek munka mellett végezhetőek és karrierváltásra is lehetőséget biztosítanak.

Az adatközponti technológiák világa izgalmas és kihívásokkal teli. Amennyiben szeretnél keresett rendszermérnökké válni, mindezt **olyan tematika és módszertan segítségével, amelyet a leendő munkáltatód igényei szerint alakítottak ki**, akkor itt az idő, hogy válts! Nálunk a legjobb helyen jársz!

Mik is azok az adatközponti megoldások?

Tipikusan minden olyan IT berendezés, amelyet a normál földi halandó soha nem lát (tehát nem a PC, nem a laptop vagy hasonló eszközök), bár azok szolgáltatásait minden nap igénybe veszi. Ezek géptermekben vannak, rendkívül nagy megbízhatóságúak, állandóan működnek, mindenki élvezi ha mennek, viszont bosszankodik, ha valamelyik szolgáltatás nem érhető el. Azaz igazán felemelő érzés ilyeneket tervezni, alkotni, mert sokaknak okoz örömöt a mindennapos használhatóság.

Kiket várunk? | **Jelentkezés** | **Helyszín** | **Konstrukciók**

Mindig is érdekelt az, hogyan működnek a nagy megbízhatóságú informatikai rendszerek? Kíváncsi vagy, milyen folyamatok zajlanak a háttérben?
Ha az informatikai érdeklődésed határtalan, precíz vagy, és jó logikai képességgel rendelkezel, valamint egy megbecsült, keresett informatikai szakemberré szeretnél válni, akkor Te vagy az, akit várunk!

Ha motivált és kitartó vagy, piacképes tudással leszel gazdagabb, és akár a Hewlett Packard Enterprise egyik partnercégénél is elkezdheted a karrieredet.

Hewlett Packard Enterprise

brainghub.hu/?s=adatközponti+rendszermérnők