

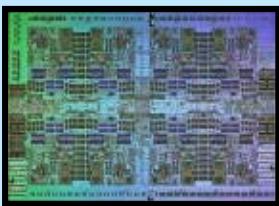
The IBM POWER9 Scale Up Processor

**Jeff Stuecheli
William Starke**

POWER Systems, IBM Systems



POWER Processor Technology Roadmap

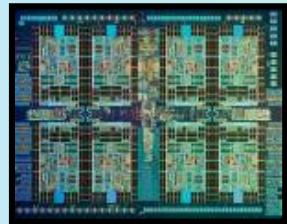


POWER7
45 nm

Enterprise

- 8 Cores
- SMT4
- eDRAM L3 Cache

1H10

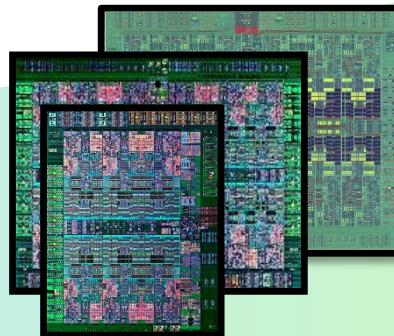


POWER7+
32 nm

Enterprise

- 2.5x Larger L3 cache
- On-die acceleration
- Zero-power core idle state

2H12

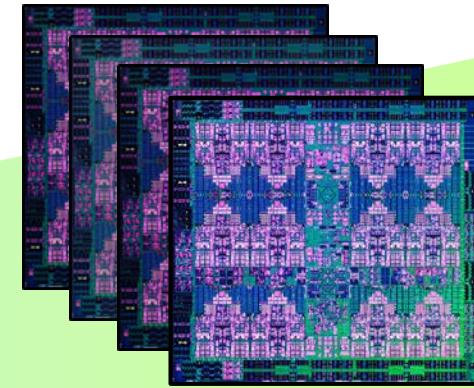


POWER8 Family
22nm

Enterprise & Big Data Optimized

- Up to 12 Cores
- SMT8
- CAPI Acceleration
- High Bandwidth GPU Attach

1H14 – 2H16



POWER9 Family
14nm

Built for the Cognitive Era

- Enhanced Core and Chip Architecture Optimized for Emerging Workloads
- Processor Family with Scale-Up and Scale-Out Optimized Silicon
- Premier Platform for Accelerated Computing

2H17 – 2H18+

New Core Microarchitecture

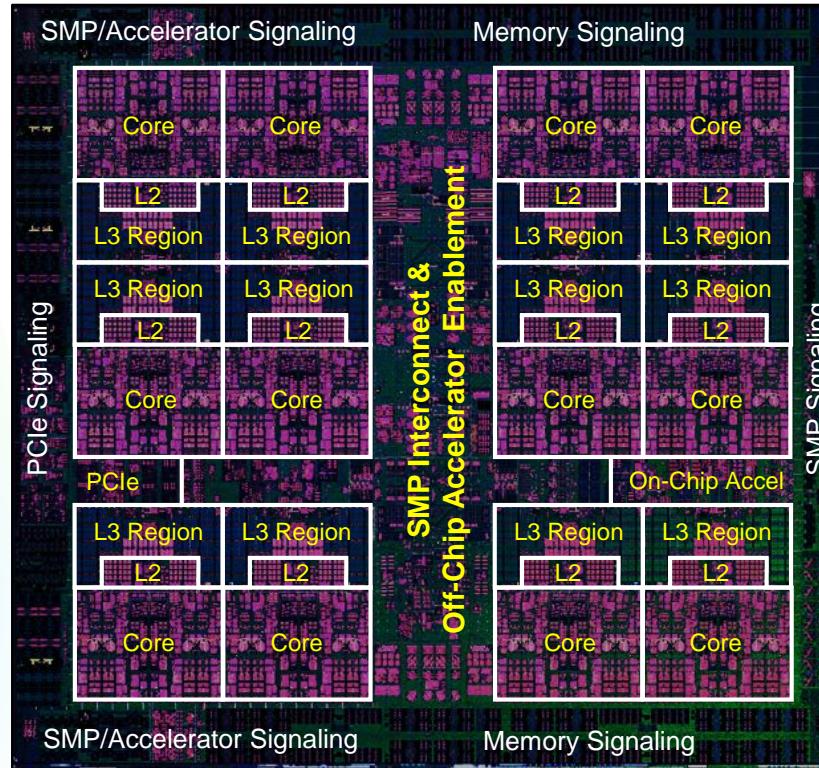
- Stronger thread performance
- Efficient agile pipeline
- POWER ISA v3.0

Enhanced Cache Hierarchy

- 120MB NUCA L3 architecture
- 12 x 20-way associative regions
- Advanced replacement policies
- Fed by 7 TB/s on-chip bandwidth

Cloud + Virtualization Innovation

- Quality of service assists
- New interrupt architecture
- Workload optimized frequency
- Hardware enforced trusted execution



14nm finFET Semiconductor Process

- Improved device performance and reduced energy
- 17 layer metal stack and eDRAM
- 8.0 billion transistors

Leadership Hardware Acceleration Platform

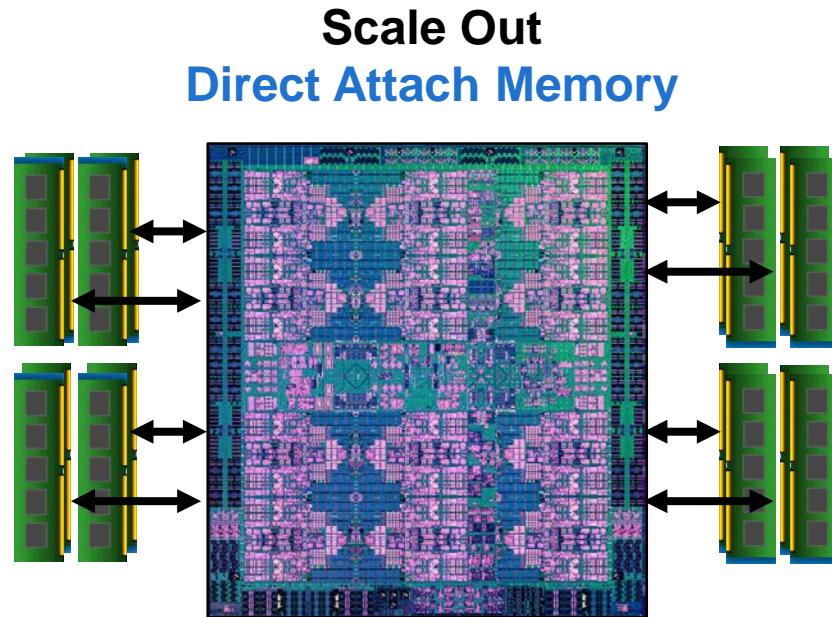
- Enhanced on-chip acceleration
- Nvidia NVLink 2.0: High bandwidth and advanced new features (25G)
- CAPI 2.0: Coherent accelerator and storage attach (PCIe G4)
- OpenCAPI: Improved latency and bandwidth, open interface (25G)

State of the Art I/O Subsystem

- PCIe Gen4 – 48 lanes

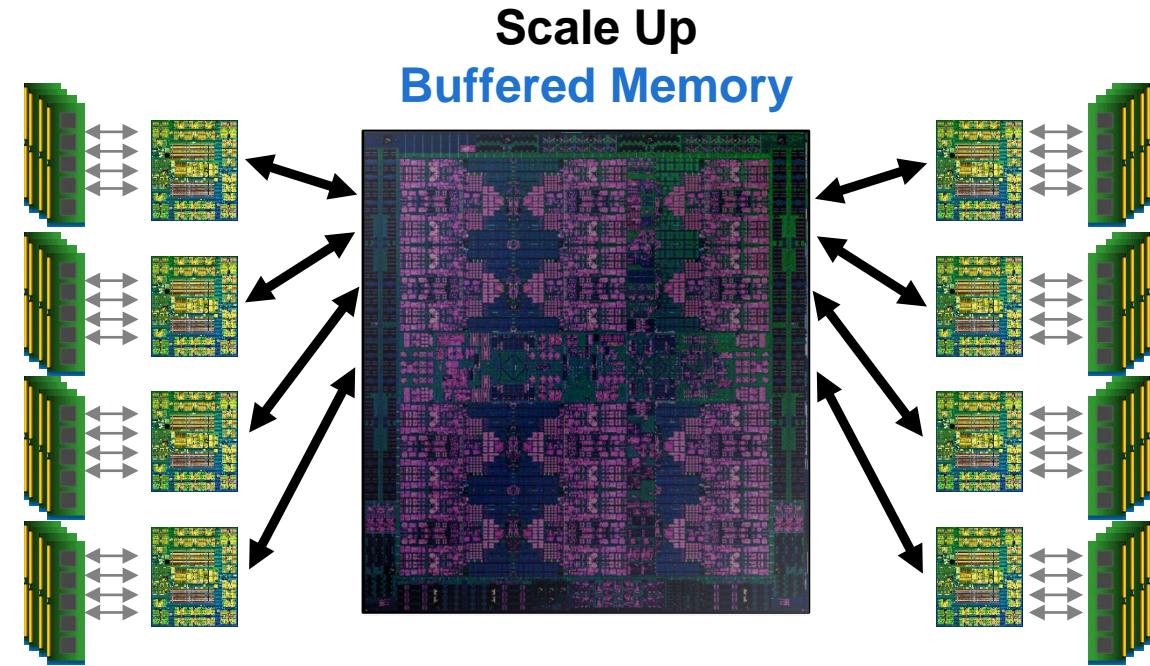
High Bandwidth Signaling Technology

- 16 Gb/s interface
 - Local SMP
- PowerAXON 25 GT/sec Link interface
 - Accelerator, remote SMP



8 Direct DDR4 Ports

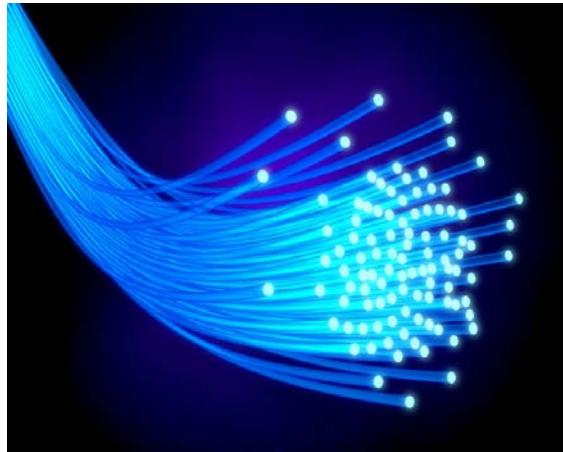
- Up to 150 GB/s of sustained bandwidth
- Low latency access
- Commodity packaging form factor
- Adaptive 64B / 128B reads



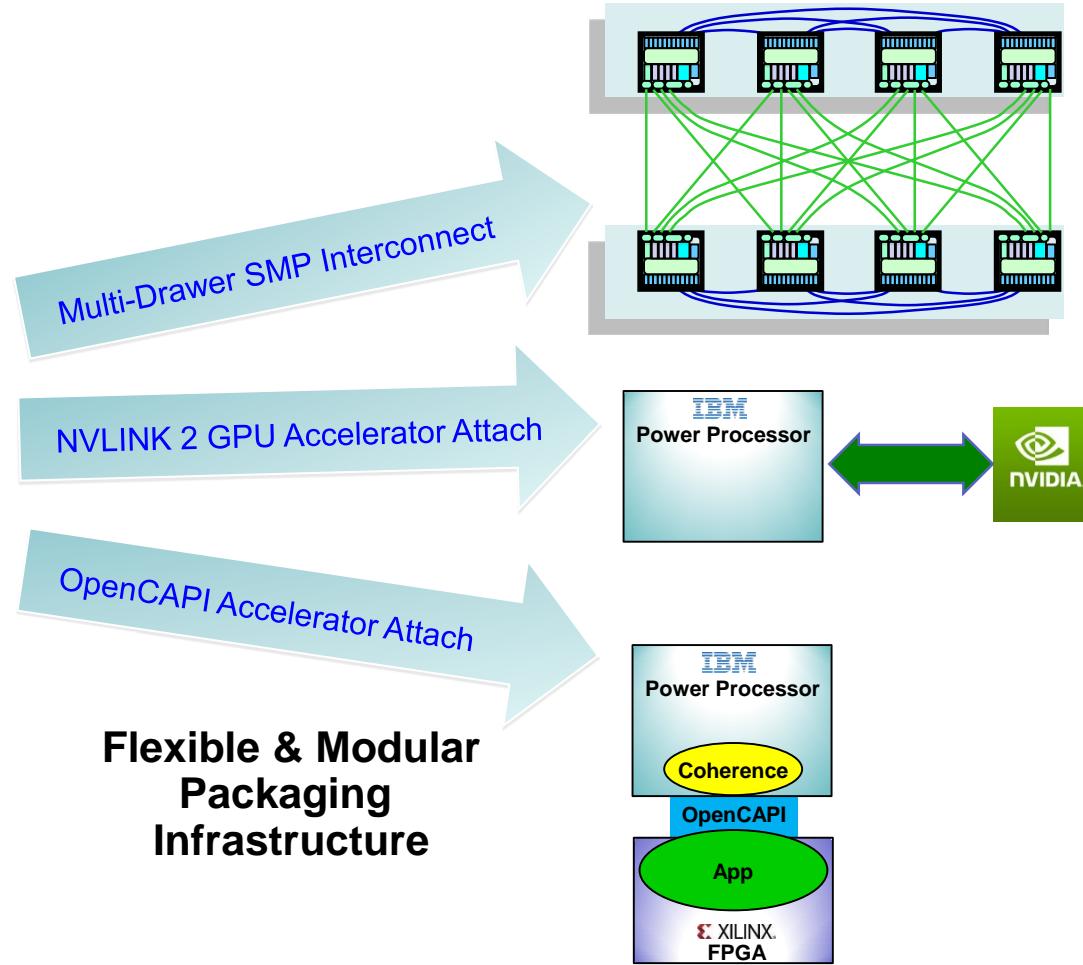
8 Buffered Channels

- Up to 230GB/s of sustained bandwidth
- Extreme capacity – up to 8TB / socket
- Superior RAS with chip kill and lane sparing
- Compatible with POWER8 system memory
- Agnostic interface for alternate memory innovations

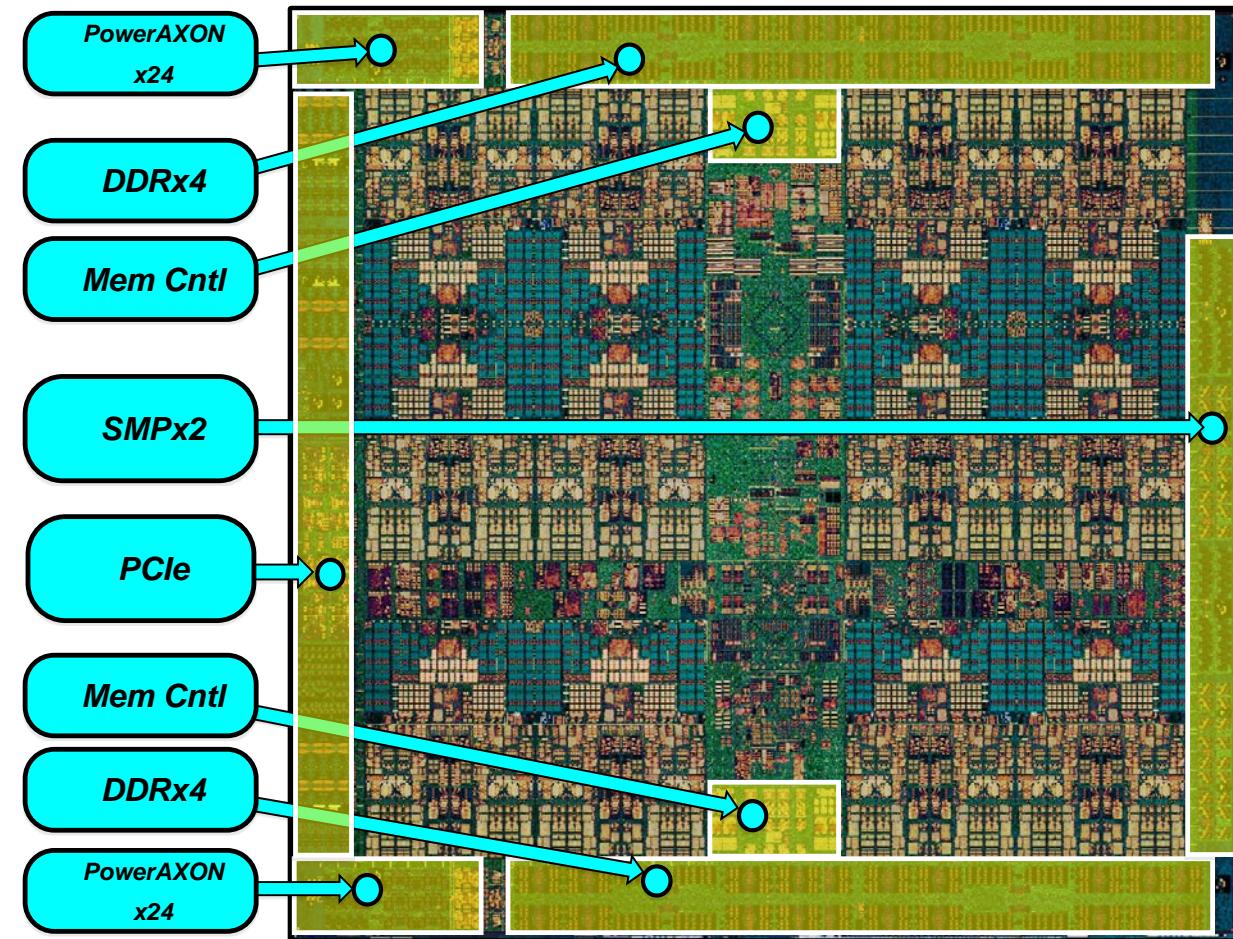
PowerAXON → High-speed 25 GT/s Signaling



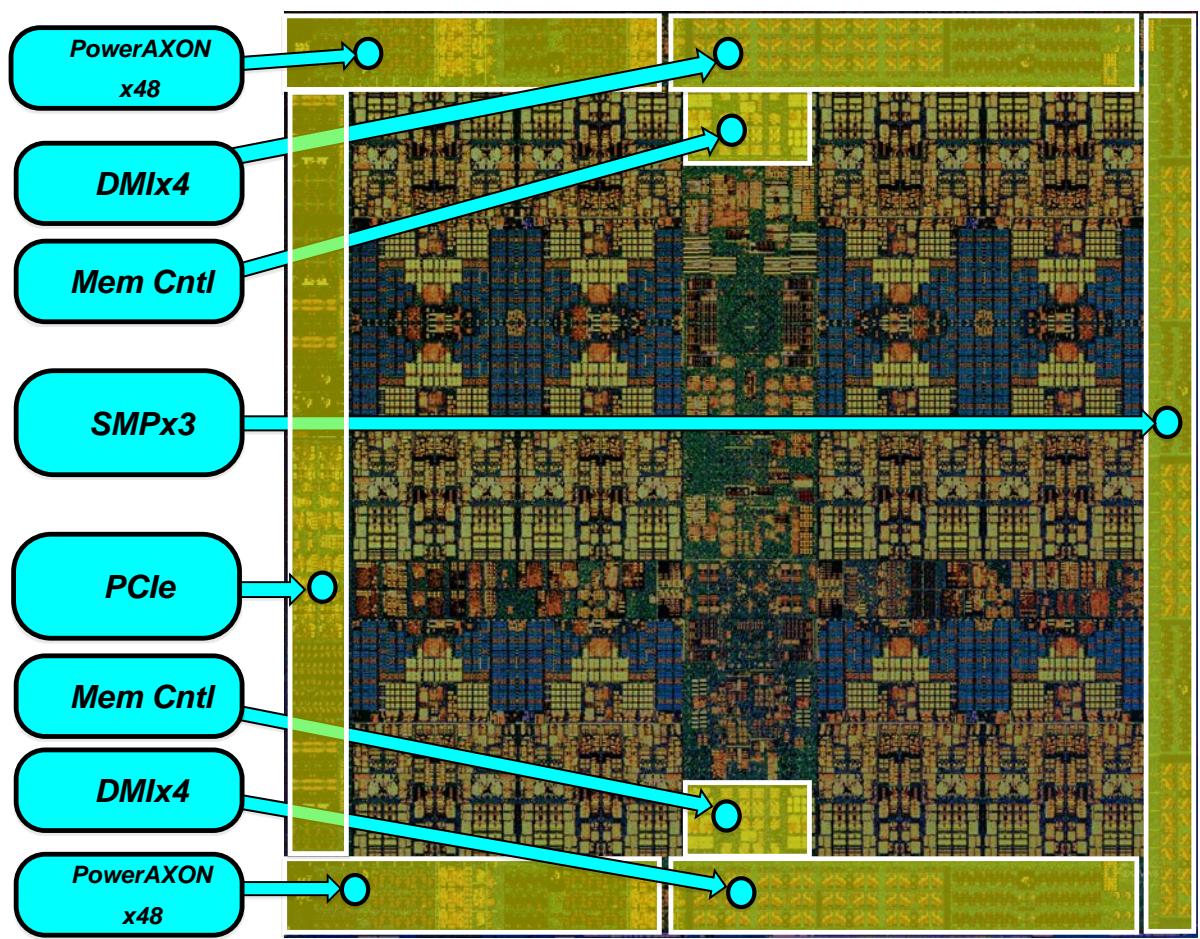
**Utilize Best-of-Breed
25 GT/s Optical-Style
Signaling Technology**



Scale Out
Direct Attach Memory
2 Socket SMP

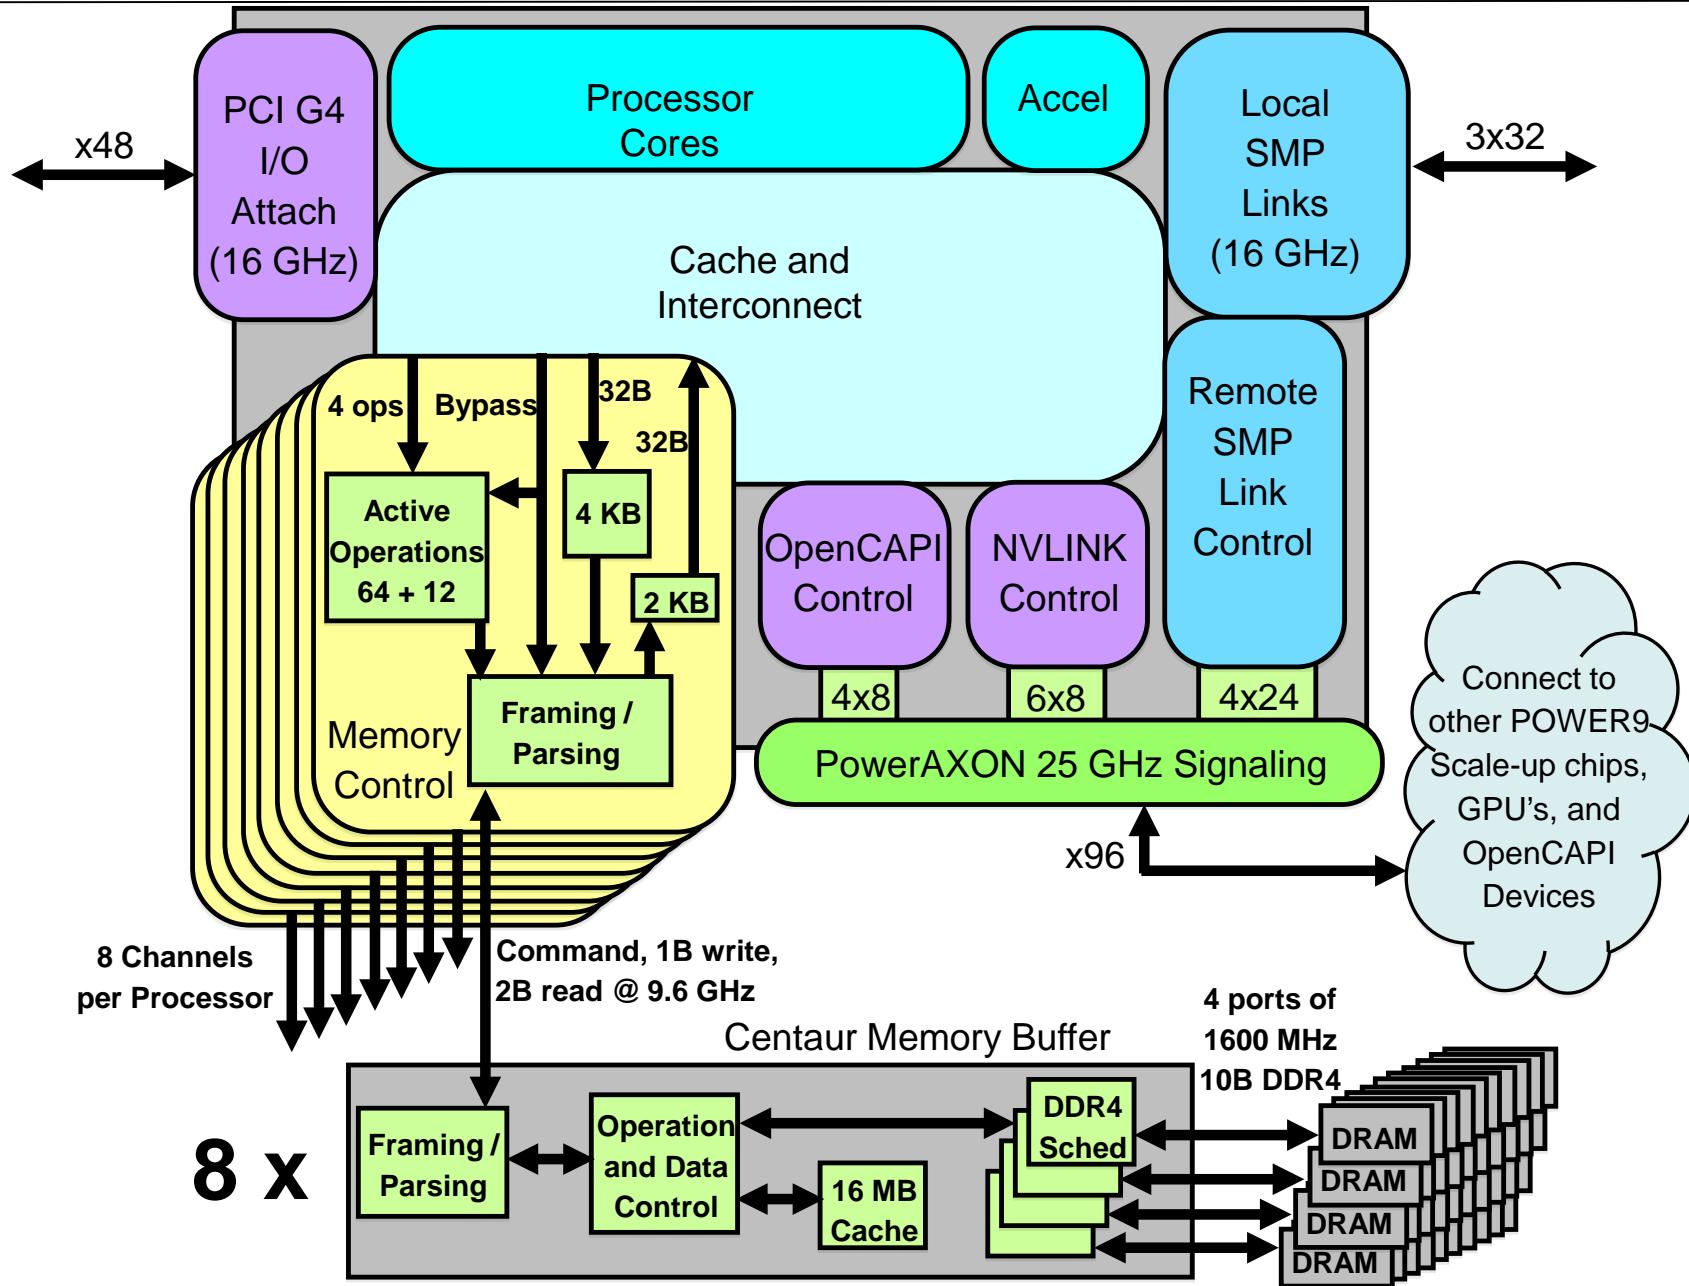


Scale Up
Buffered Memory
16 Socket SMP

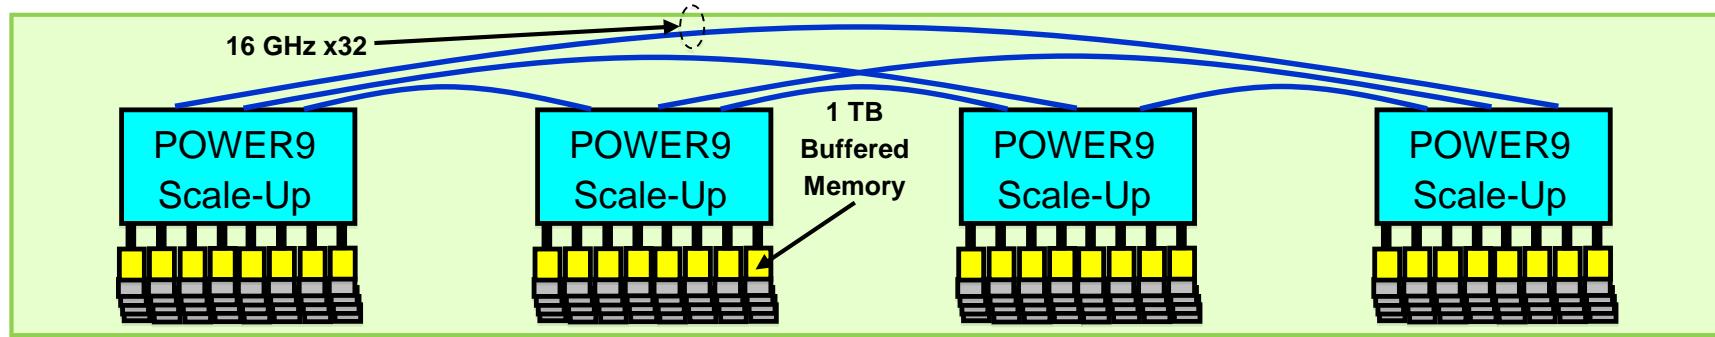




Scale Up Chipset Block Diagram

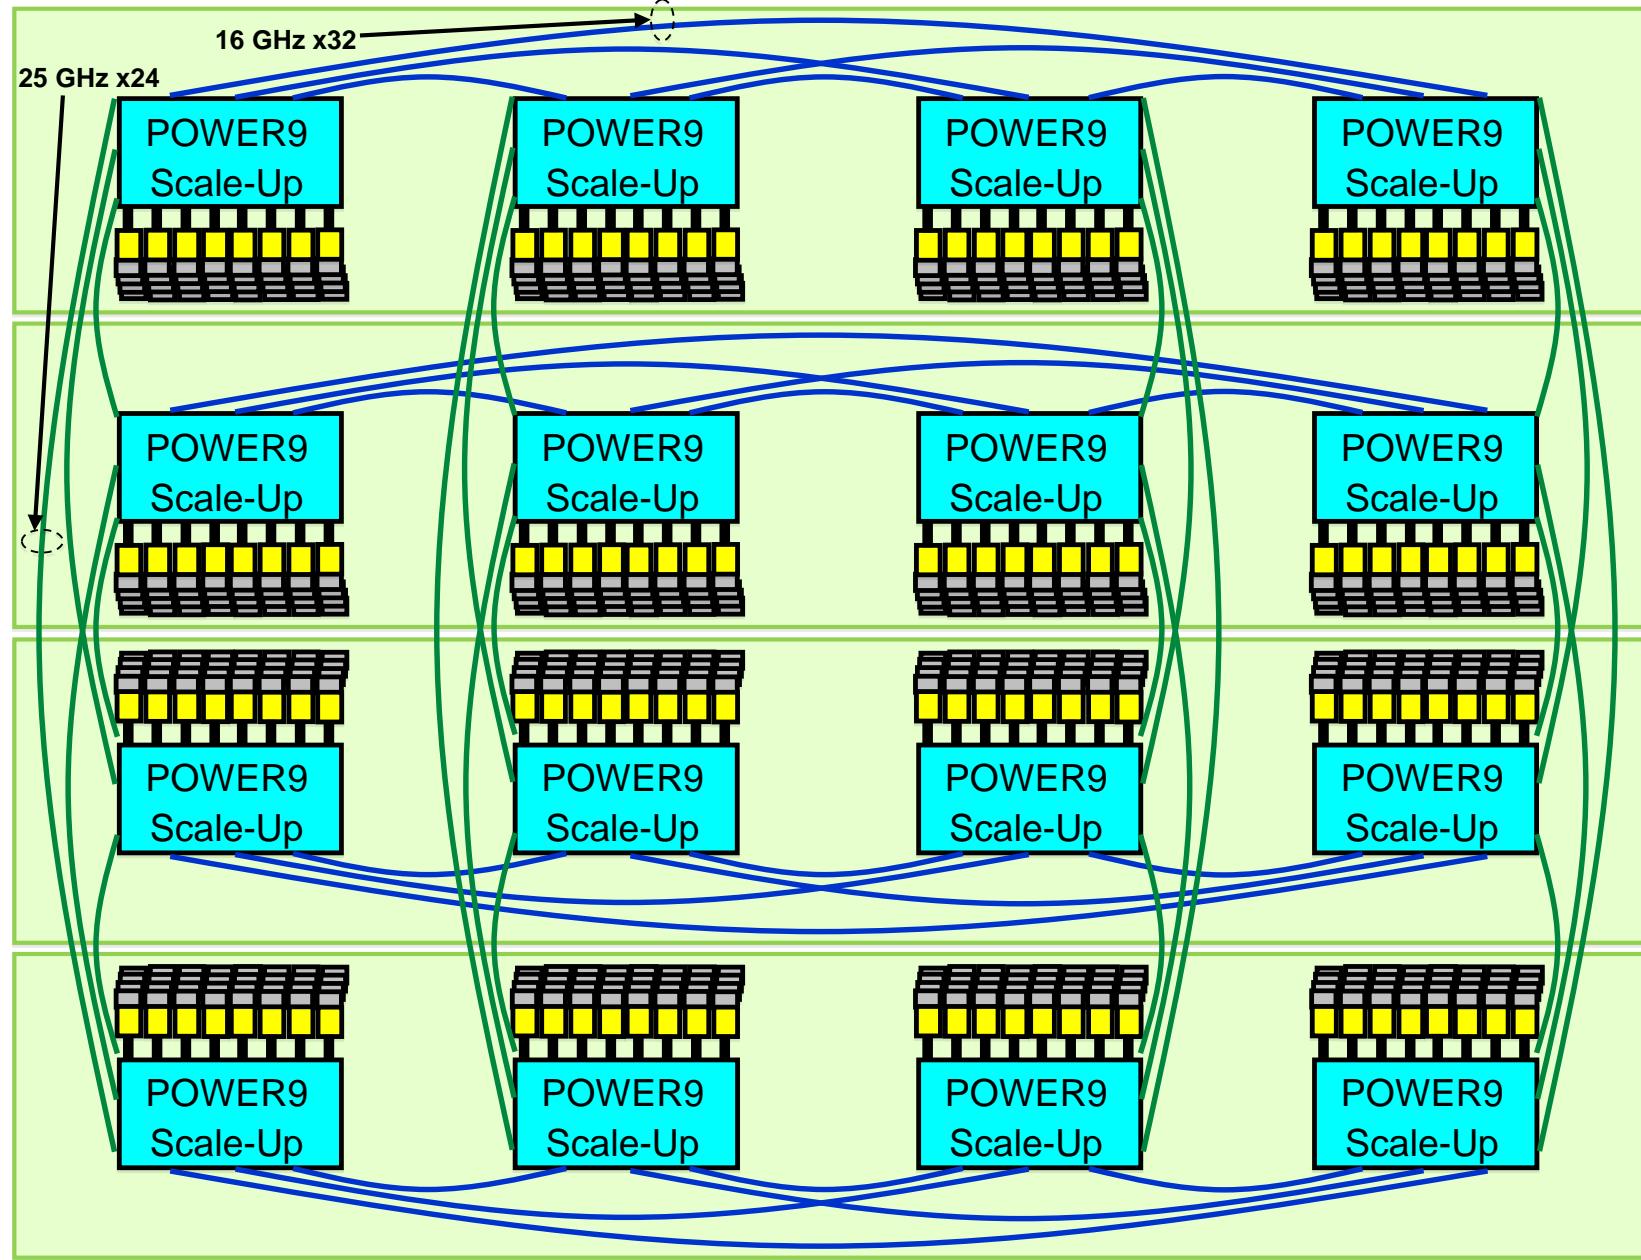


Four Socket Server Topology

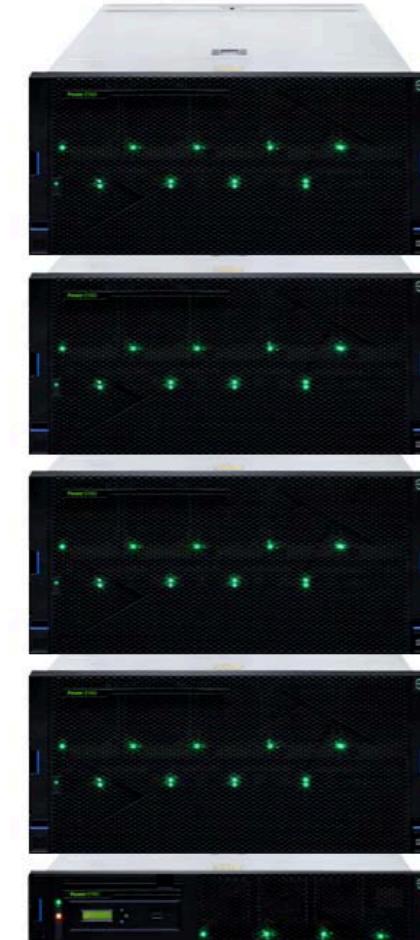


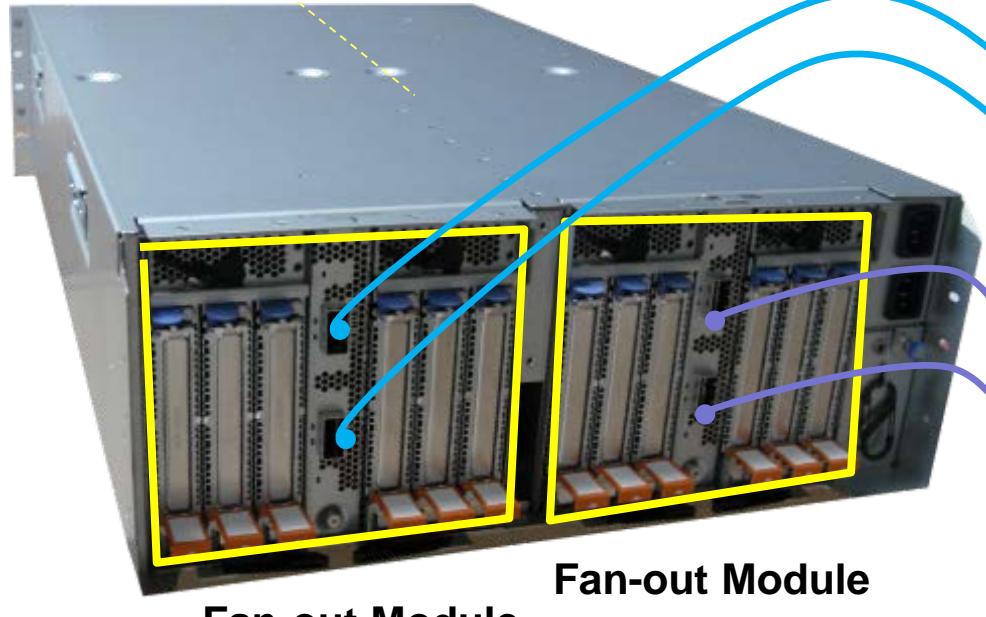
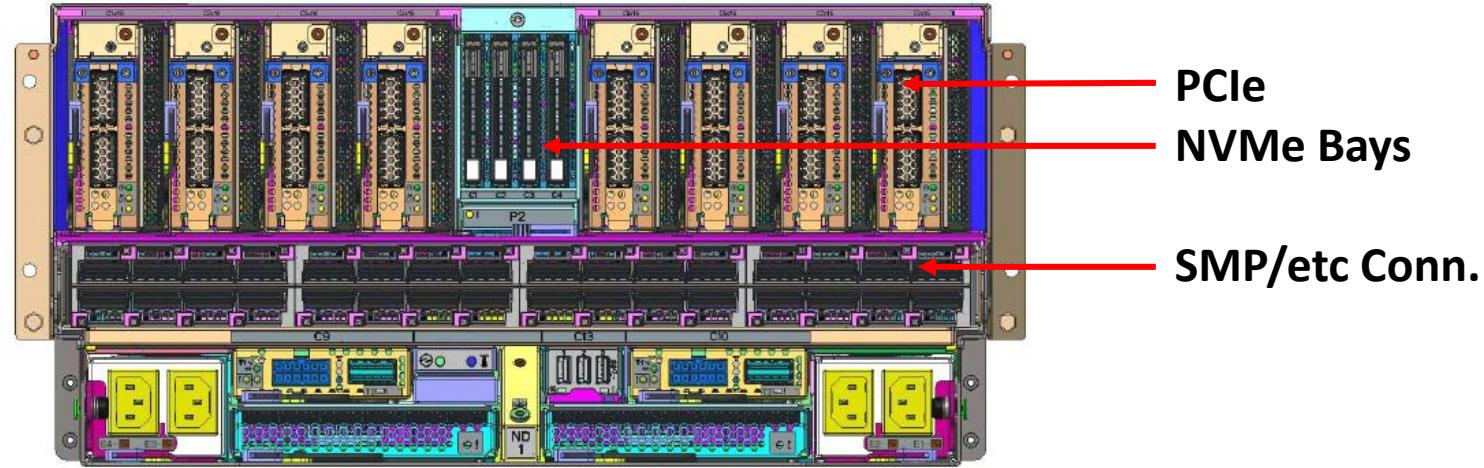


Sixteen Socket Server Topology

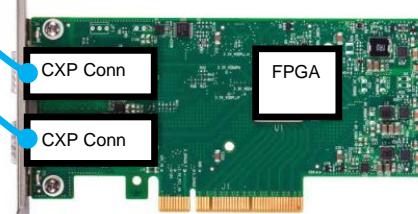


- ✓ Modular, Scalable POWER9 server
 - 1 to 4 x 5U CEC drawers + 2U Control Unit
- ✓ POWER9 Enterprise processor
- ✓ Up to 192 cores in a single system
- ✓ Up to 64 TB DDR4 memory
- ✓ Up to 32 PCIe Gen4 slots
- ✓ PowerAXON 25Gb/s ports
 - *Used for SMP cabling between nodes - 4x bandwidth improvement*
 - *Enabled for OpenCAPI accelerators*

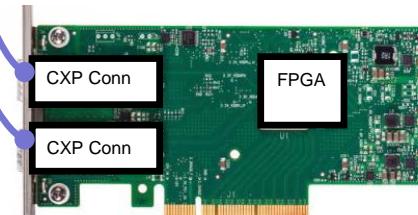




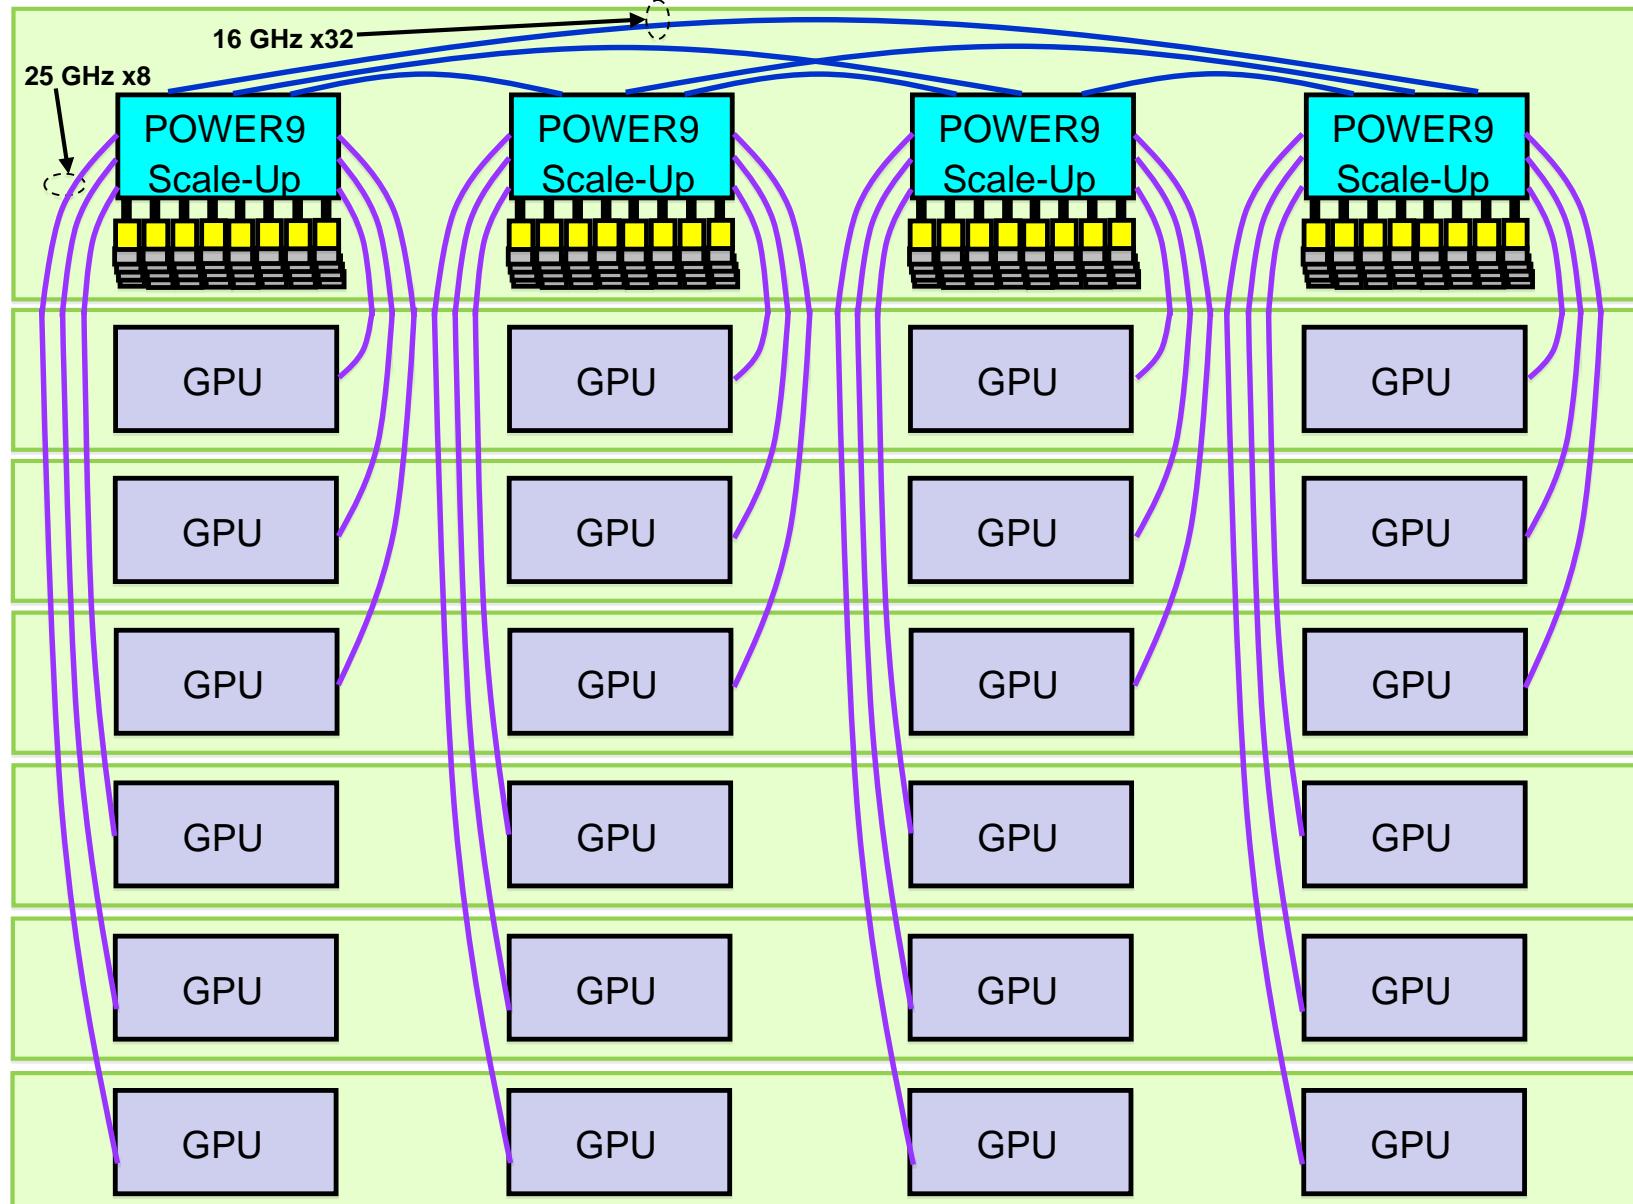
Cable Card



Cable Card

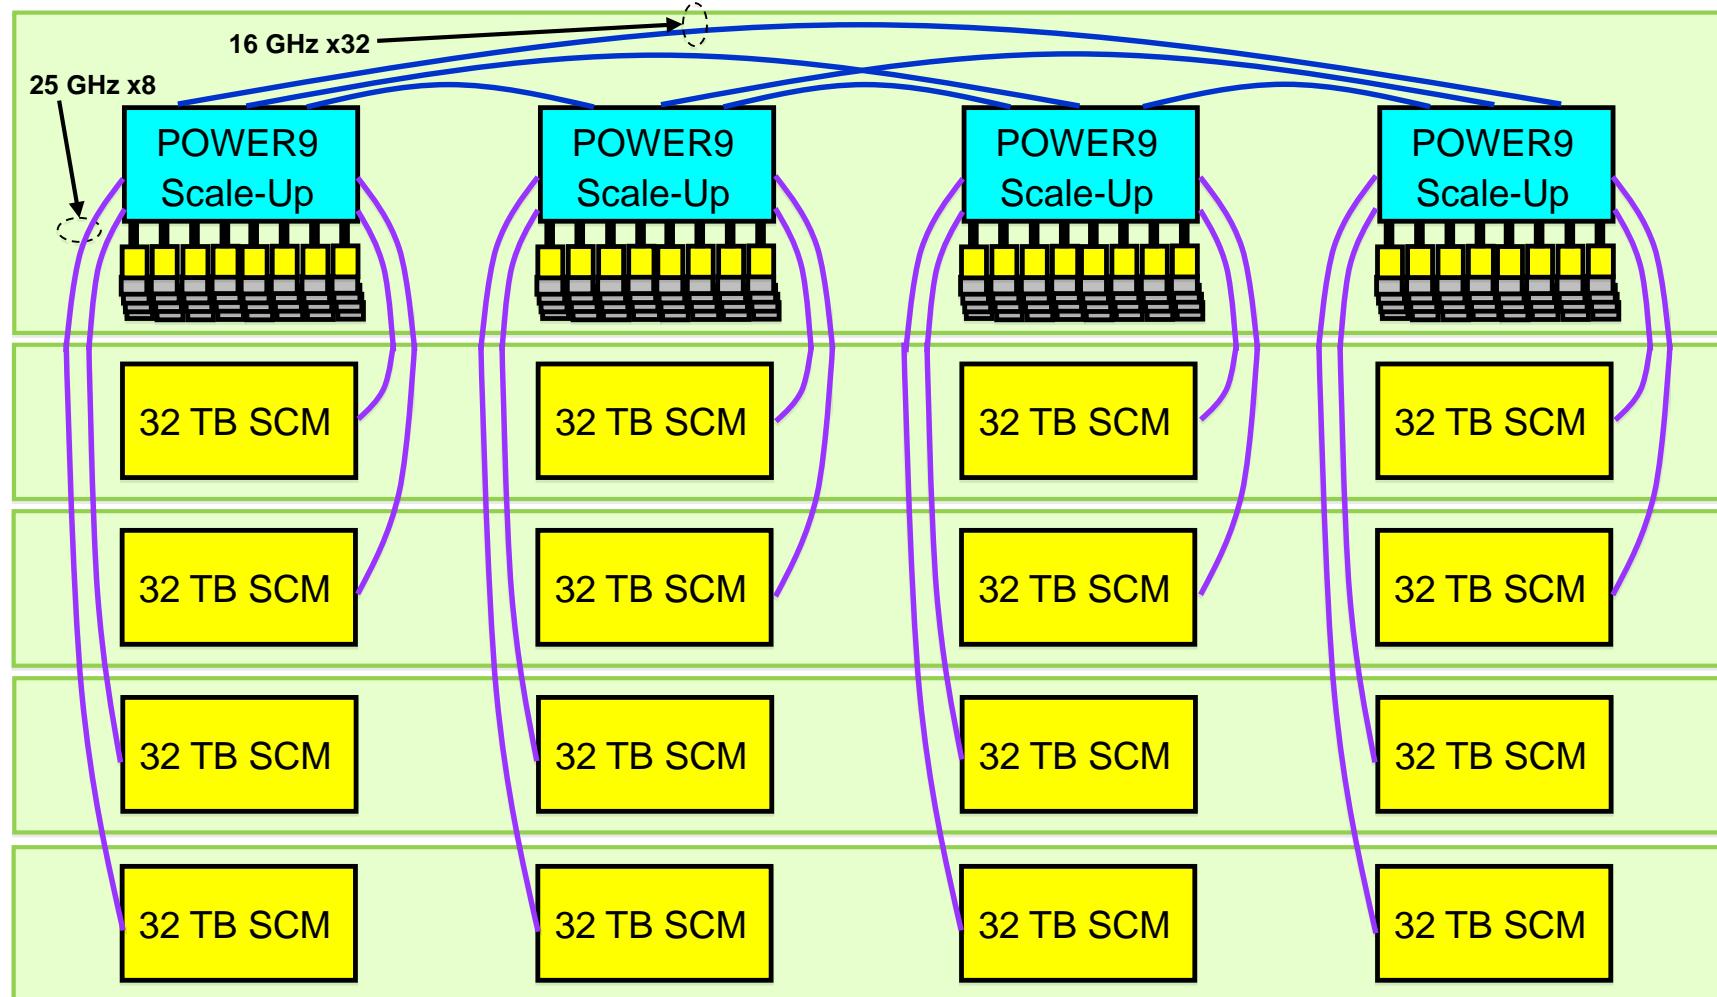


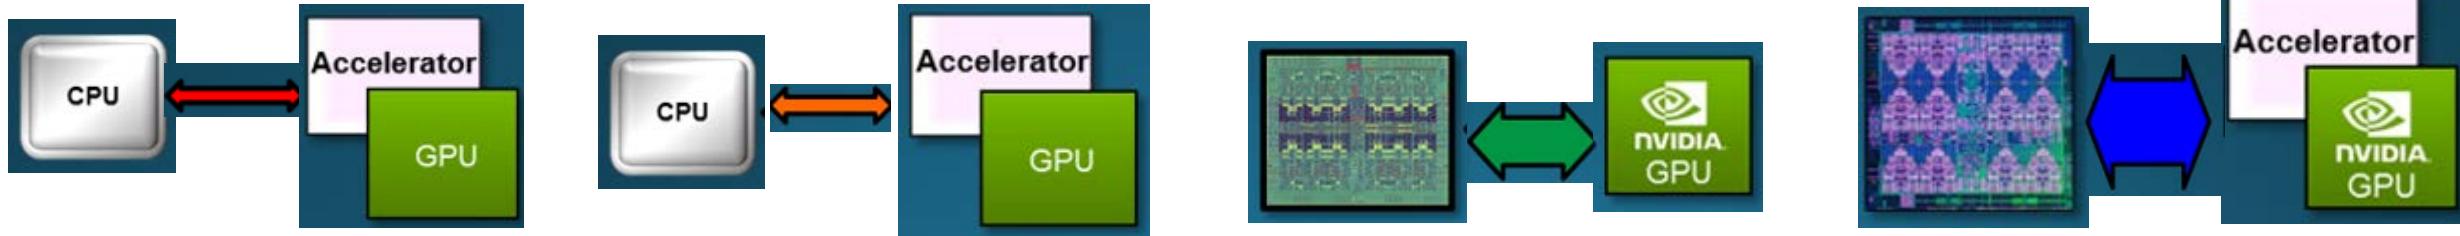
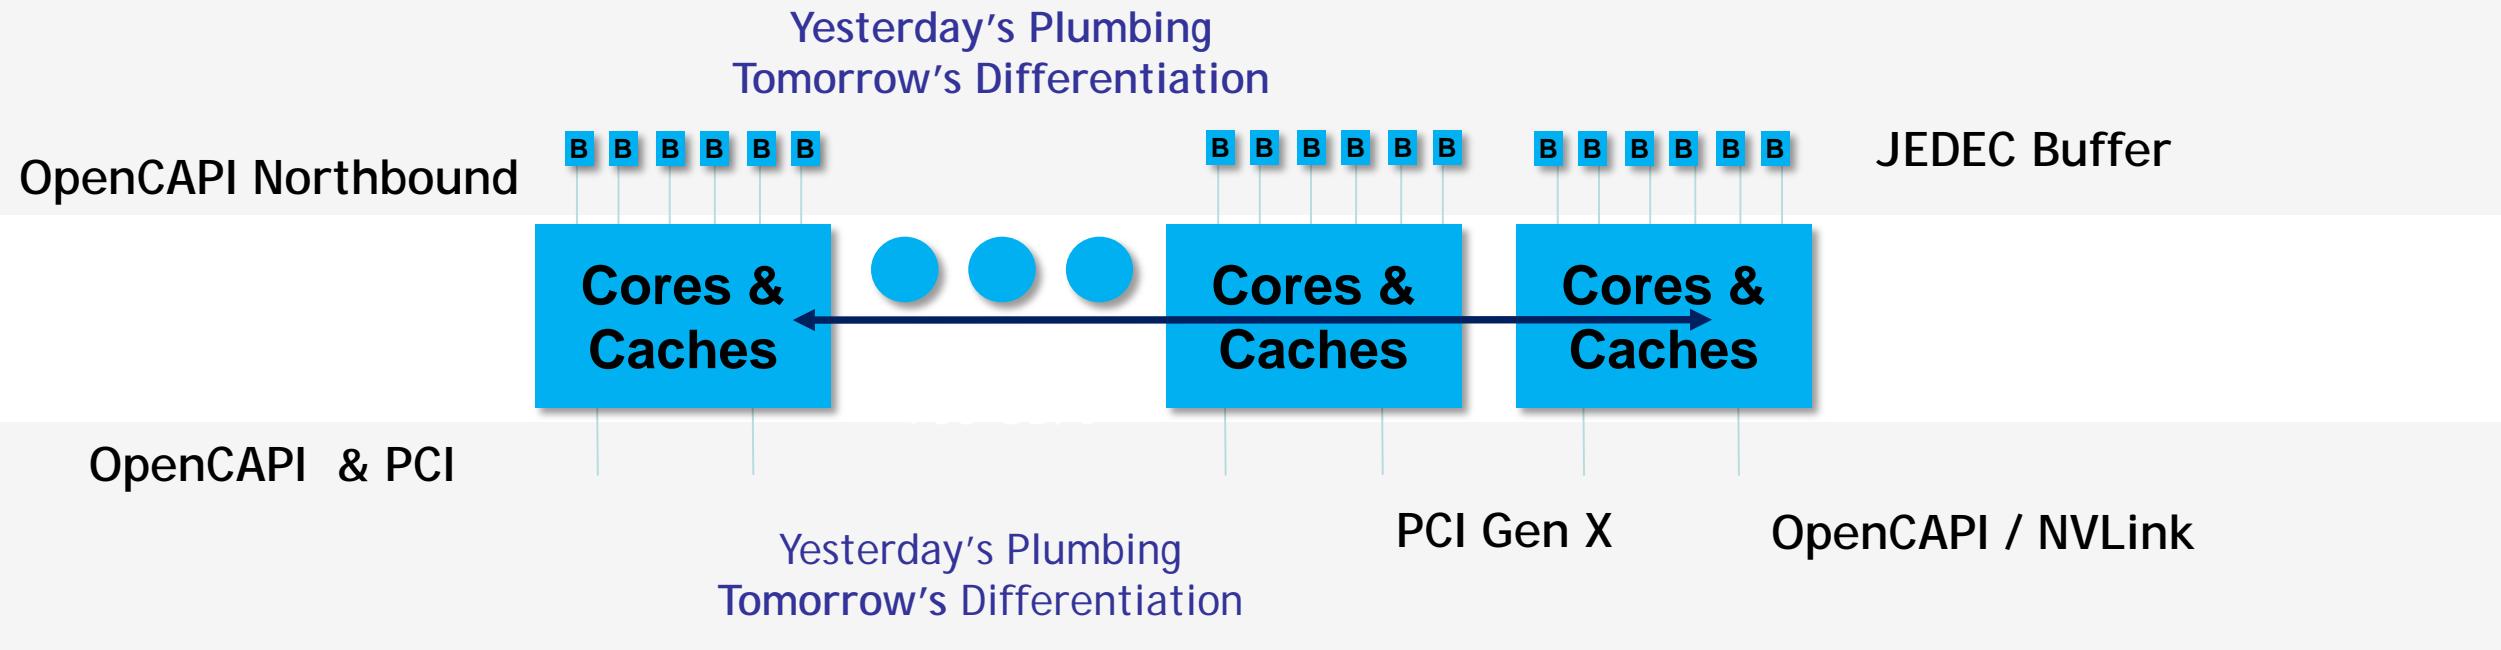
GPU on PowerAXON



Demonstration of POWER9 capability, does not imply specific system plans

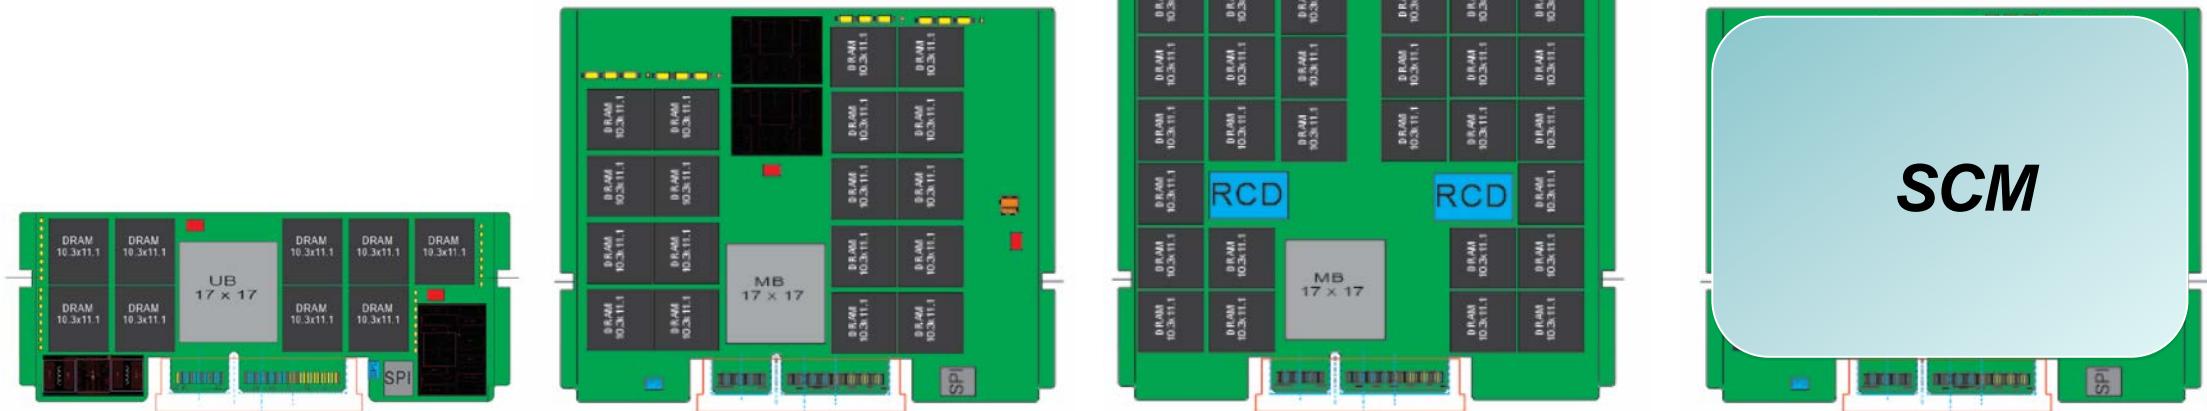
Storage Class Memory on PowerAXON





System bottleneck

- Signaling → AXON @25.6GHz vs DDR4 @ 3200 MHz
 - 4 times the bandwidth per IO
- Idle latency over traditional DDR
 - POWER8/9 Centaur design ~10 ns
 - OpenCAPI target of ~5 ns
- Centaur → One proprietary design
- OpenCAPI → Open



Todays talk

POWER7 Architecture	POWER8 Architecture	POWER9 Architecture	POWER10
2010 POWER7 8 cores 45nm New Micro-Architecture New Process Technology	2012 POWER7+ 8 cores 32nm Enhanced Micro-Architecture New Process Technology	2014 POWER8 12 cores 22nm New Micro-Architecture New Process Technology	2016 POWER8 w/ NVLink 12 cores 22nm Enhanced Micro-Architecture With NVLink
2017 P9 SO 12/24 cores 14nm New Micro-Architecture Direct attach memory New Process Technology	2018 P9 SU 12/24 cores 14nm Enhanced Micro-Architecture Buffered Memory	2019 P9 w/ Adv. I/O 12/24 cores 14nm Enhanced Micro-Architecture New Memory Subsystem	2020+ P10 TBA cores New Micro-Architecture New Technology
Sustained Memory Bandwidth	65 GB/s	65 GB/s	210 GB/s
Standard I/O Interconnect	PCIe Gen2	PCIe Gen2	PCIe Gen3
Advanced I/O Signaling	N/A	N/A	N/A
Advanced I/O Architecture	N/A	N/A	CAPI 1.0 , NVLink 1.0
			CAPI 2.0, OpenCAPI3.0, NVLink2.0
			CAPI 2.0, OpenCAPI3.0, NVLink2.0
			CAPI 2.0, OpenCAPI4.0, NVLink3.0
			TBA

Statement of Direction, Subject to Change

Thanks
Questions?