Power Systems

IBM

POWER chips

Lance Thompson

POWER Systems, IBM Systems

Recent and Future POWER Processor Roadmap

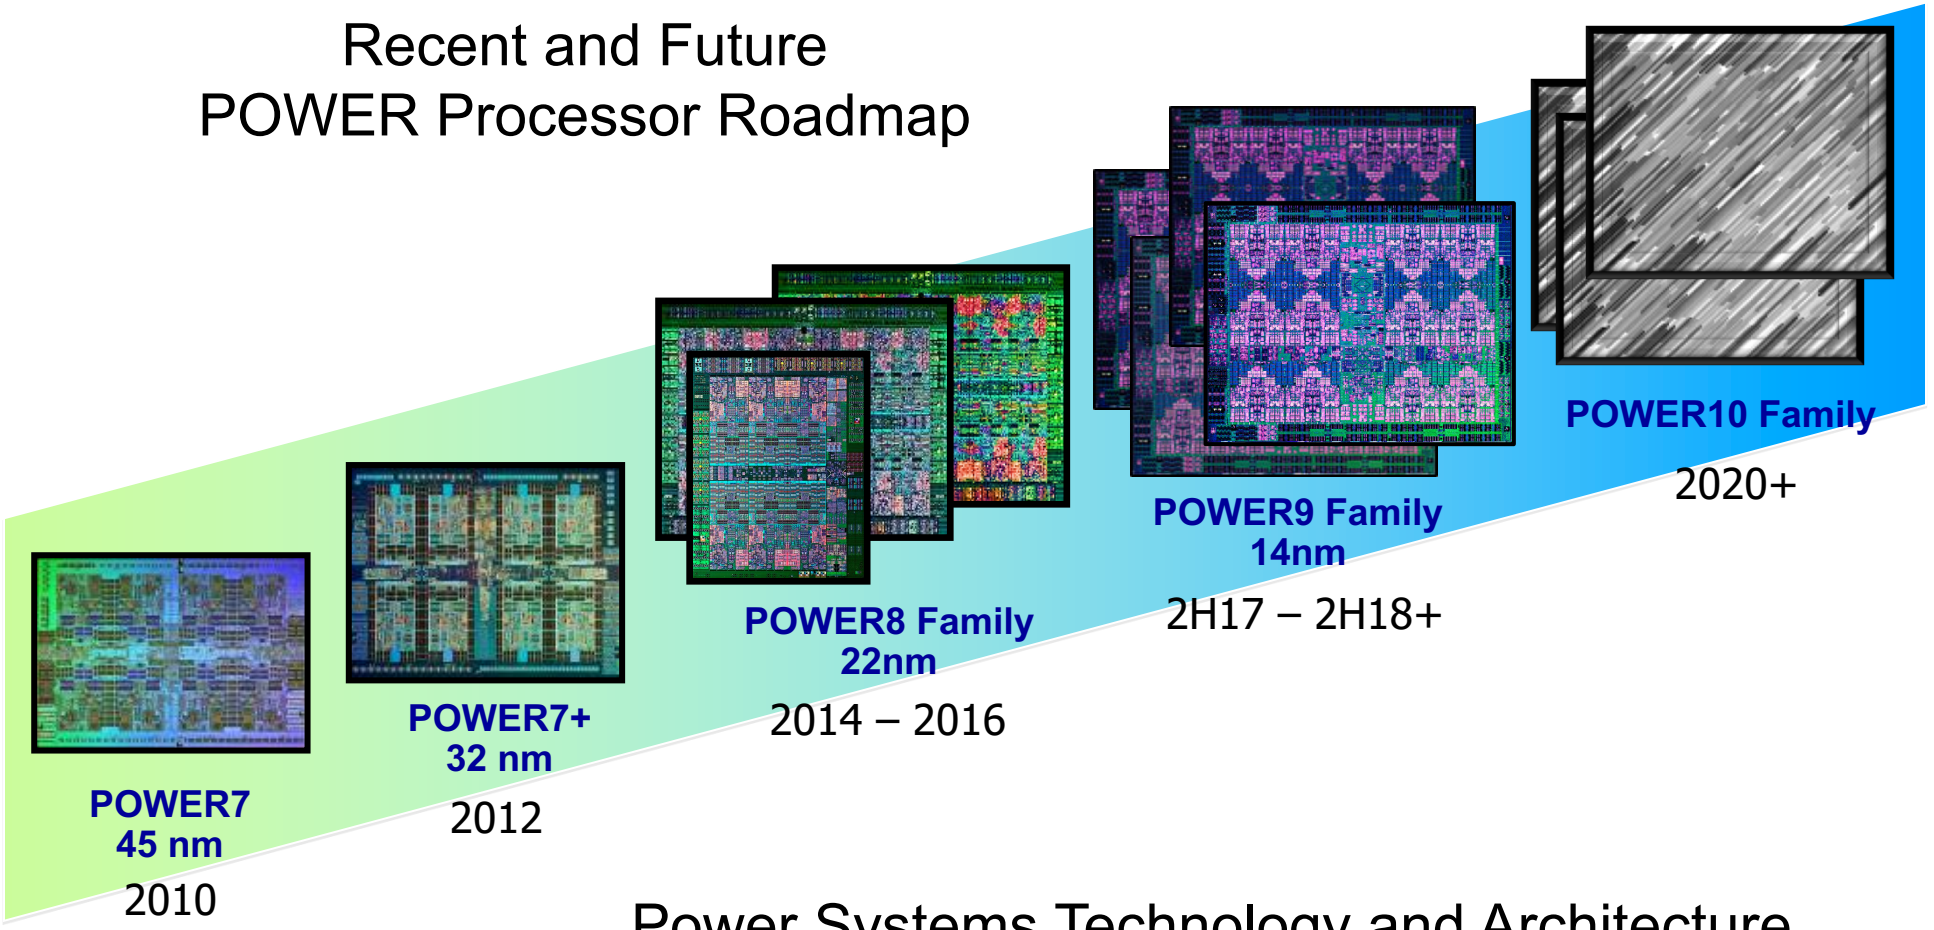

Power Systems Technology and Architecture
Leveraging the economics of the New Era

Power Systems

IBM

POWER8

Jeff Stuecheli

POWER Systems, IBM Systems

POWER8 Highlights

- 12 Cores per Socket/Chip
- Significantly Strengthened Cores
  - 8 threads per core (SMT8)
  - Wider fetch / dispatch/issue of instructions (8 fetch / dispatch, 10 issue)
  - Doubled highly utilized execution units
- Larger Caches
  - L1: 64K data “D” Cache, 32K instruction “I” Cache per core
  - L2: 512K private per core
  - L3: 8M per core (up to 96MB per chip)
  - L4: external to the chip
- 2 Integrated Memory Controllers w/ Improved Latency & BW
  - ~ 25% memory latency improvement via on-chip fastpath interconnect
  - 16MB mem L4 cache in each buffer chip
- Integrated SMP Interconnect w/ improved “Flatness”
  - 2-Hop fabric topology
- Integrated IO Subsystem
  - On Chip PCIe Controller
- Fine Grained Power Management
  - On Chip Power Management Controller & Power gating

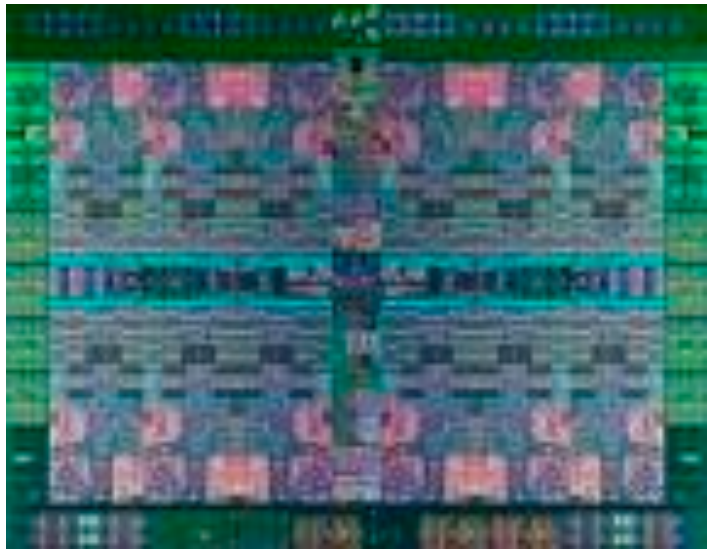

POWER8 Core

Wider Load/Store
- 32B → 64B L2 to L1 data bus
- 2x data cache to execution dataflow

Enhanced Prefetch
- Instruction speculation awareness
- Data prefetch depth awareness
- Adaptive bandwidth & Topology awareness

Larger Caching Structures vs. POWER7
- 2x L1 data cache (64 KB)
- 2x outstanding data cache misses
- 4x translation Cache

Execution Improvement vs. POWER7
- Enhanced threading: SMT4 → SMT8
- 8 dispatch / 10 issue / 16 execution pipes:
  - 2 FXU, 2 LSU, 2 LU, 4 FPU, 2 VMX, 1 Crypto, 1 DFU, 1 CR, 1 BR
- Larger Issue queues (4 x 16-entry)
- Larger global completion, Load/Store reorder
- Improved branch prediction / unaligned storage access

Core Performance vs P7
~1.5x Single Thread
~2x Max SMT

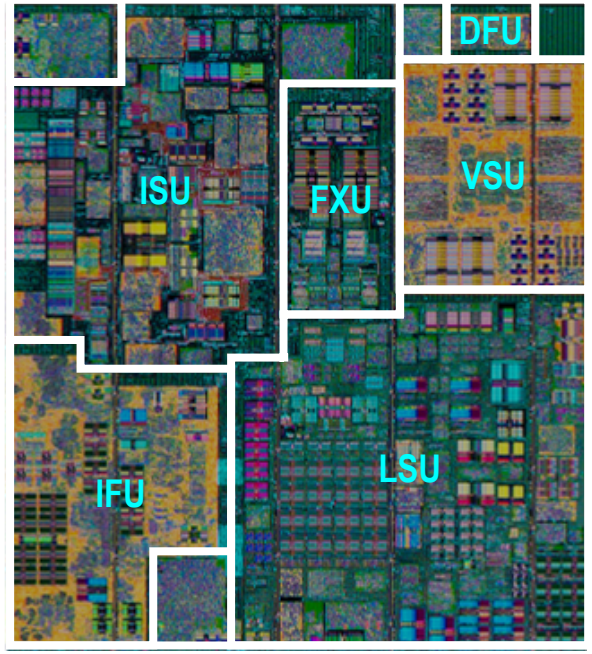

POWER8 Processor

Cores
- 12 cores (SMT8)
- 8 dispatch, 10 issue, 16 exec pipe
- 2X internal data flows/queues
- Enhanced prefetching

Caches
- 64K Data cache (L1)
- 512 KB SRAM L2 / core
- 96 MB eDRAM shared L3
- Up to 128 MB eDRAM L4 (off-chip)

Accelerators
- Crypto & memory expansion
- Transactional Memory
- Data Move / VM Mobility

Bus Interfaces
- Integrated PCIe Gen3
- SMP Interconnect
- *Nvidia NVLink 1.0*
- CAPI 1.0

Technology
- 22nm SOI, eDRAM, 15 ML 650mm2

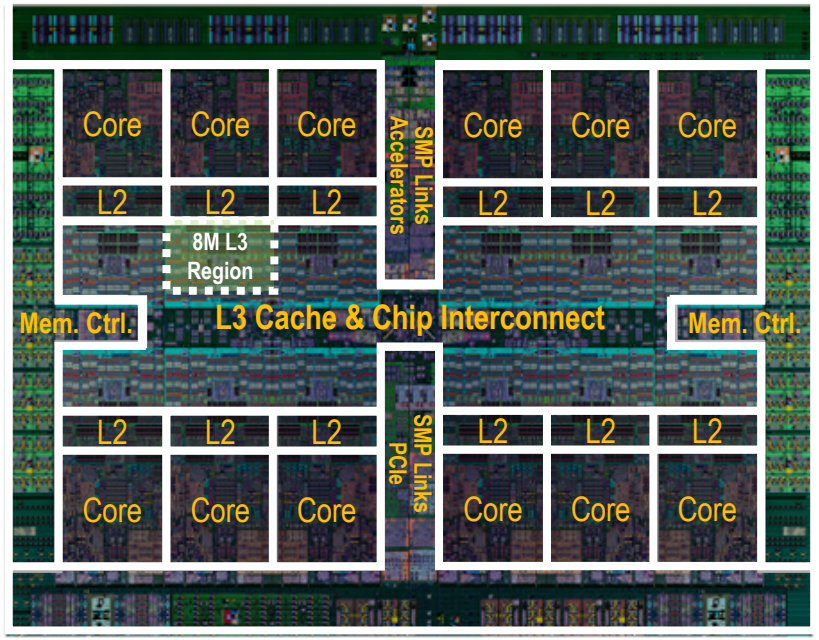

Energy Management
- On-chip Power Management Micro-controller
- Integrated Per-core VRM
- Critical Path Monitors

Memory
- Dual memory Controllers
- 230 GB/sec Sustained bandwidth

POWER8 Memory Organization (Max Config shown)

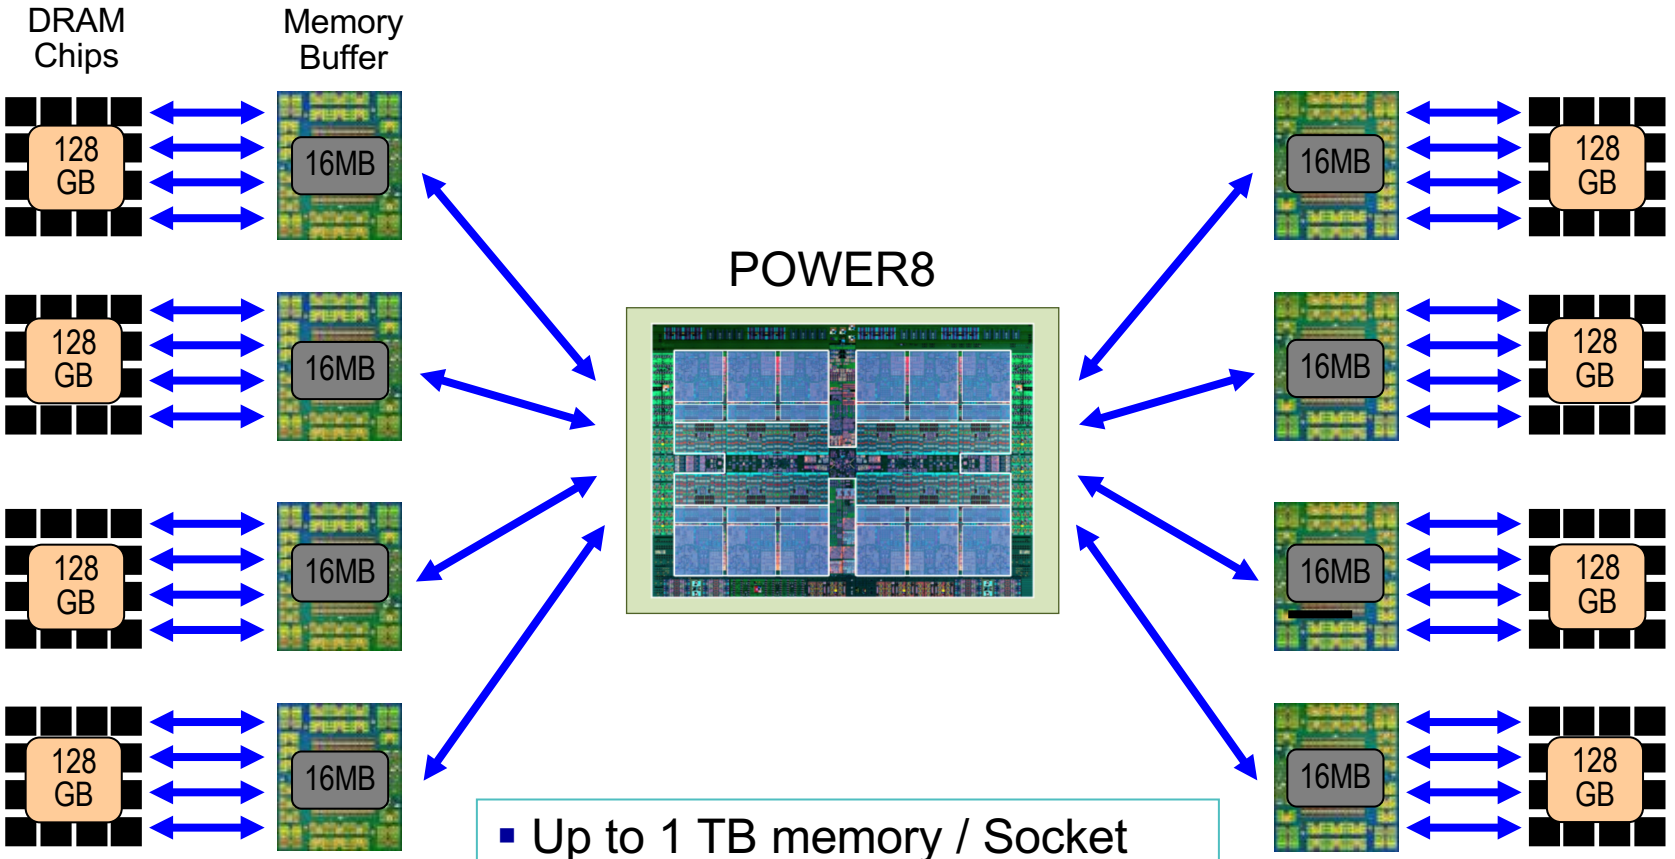

- Up to 1 TB memory / Socket
- Up to 128MB of L4 / Socket
- Up to 230 GB/s BW

POWER8 Integrated PCI Gen 3

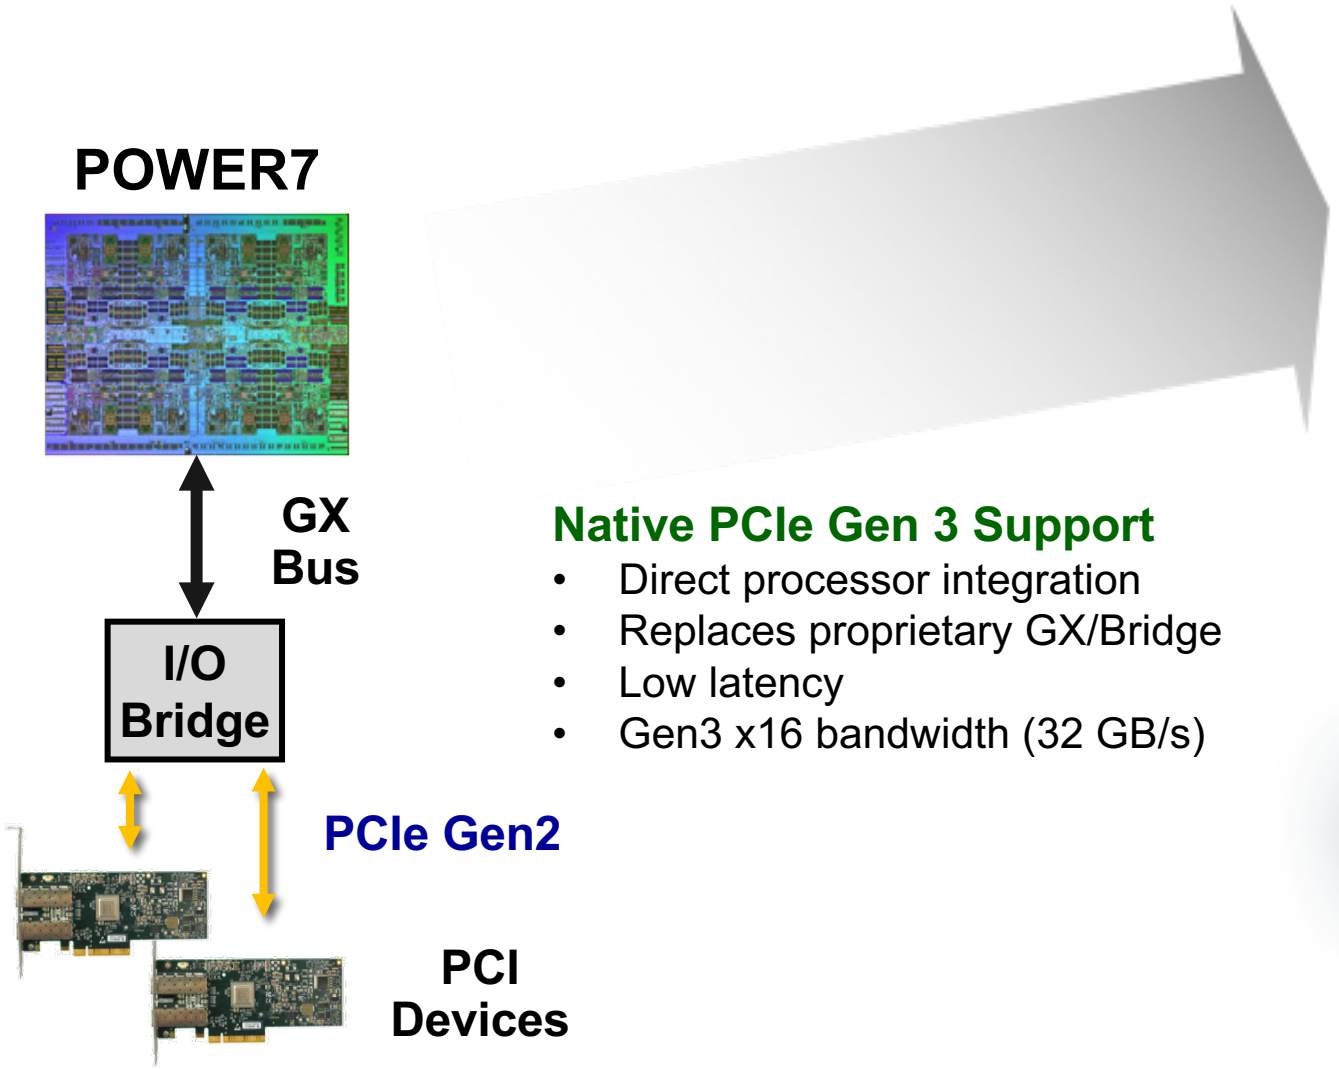

Native PCIe Gen 3 Support

- Direct processor integration
- Replaces proprietary GX/Bridge
- Low latency
- Gen3 x16 bandwidth (32 GB/s)

Power Systems

IBM

POWER9

Lance Thompson

POWER Systems, IBM Systems

POWER9 Family – Deep Workload Optimizations

Emerging Analytics, AI, Cognitive

- New core for stronger thread performance
- Delivers 2x compute resource per socket
- Built for acceleration – OpenPOWER solution enablement

SAP HANA, IBM Watson, TensorFlow, torch, DB2 BLU, Spark, Caffe

Technical / HPC

- Highest bandwidth GPU attach
- Advanced GPU/CPU interaction and memory sharing
- High bandwidth direct attach memory

OAK RIDGE National Laboratory, Lawrence Livermore National Laboratory

Cloud / HSDC

- Power / Packaging / Cost optimizations for a range of platforms
- Superior virtualization features: security, power management, QoS, interrupt
- State of the art IO technology for network and storage performance

Google, IBM SOFTLAYER, rackspace the open cloud company

Enterprise

- Large, flat, Scale-Up Systems
- Buffered memory for maximum capacity
- Leading RAS
- Improved caching

AIX, i for Business, Linux, PowerVM, DB2 10, SAP, ORACLE

New POWER9 Cores

Optimized for Stronger Thread Performance and Efficiency

- Increased execution bandwidth efficiency for a range of workloads including commercial, cognitive and analytics
- Sophisticated instruction scheduling and branch prediction for unoptimized applications and interpretive languages
- Adaptive features for improved efficiency and performance especially in lower memory bandwidth systems

Available with SMT8 or SMT4 Cores

8 or 4 threaded core built from modular execution slices

POWER9 SMT8 Core

- PowerVM Ecosystem Continuity
- Strongest Thread
- Optimized for Large Partitions

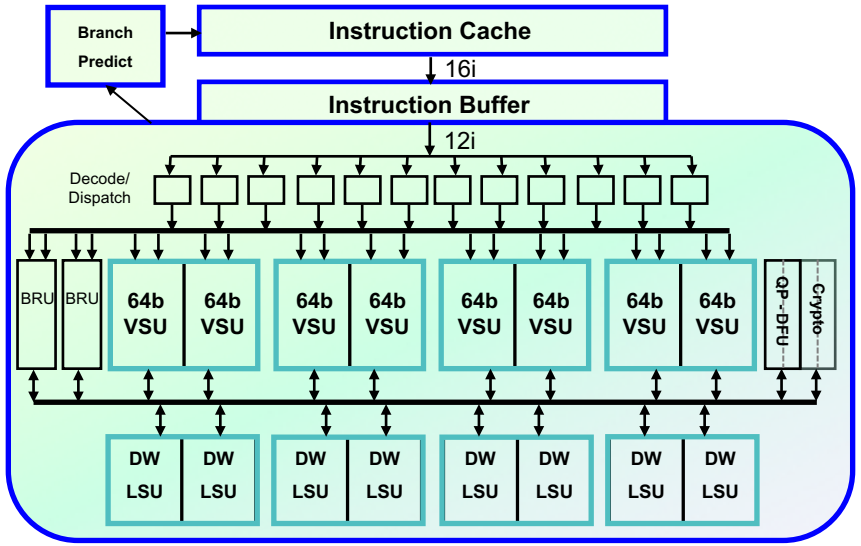

SMT8 Core

POWER9 SMT4 Core

- Linux Ecosystem Focus
- Core Count / Socket
- Virtualization Granularity

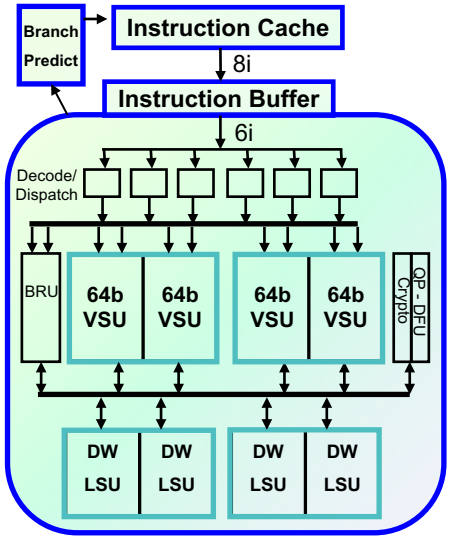

SMT4 Core

POWER9 Core Execution Slice Microarchitecture

Modular Execution Slices

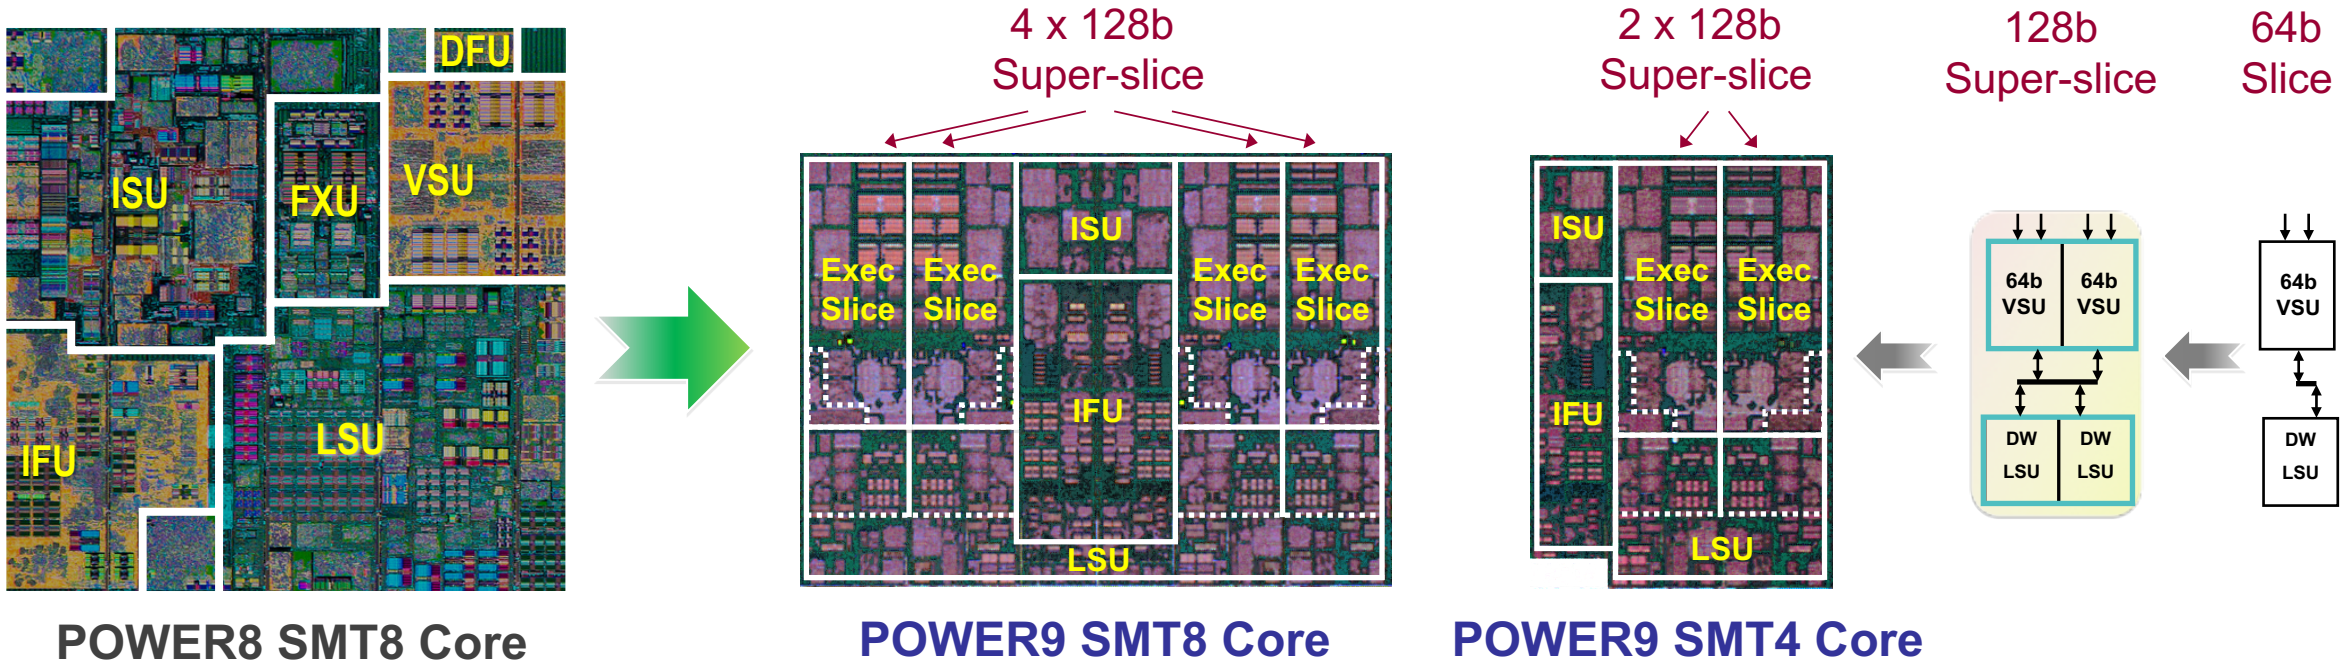

POWER8 SMT8 Core

POWER9 SMT8 Core

POWER9 SMT4 Core

Re-factored Core Provides Improved Efficiency & Workload Alignment

- Enhanced pipeline efficiency with modular execution and intelligent pipeline control
- Increased pipeline utilization with symmetric data-type engines: Fixed, Float, 128b, SIMD
- Shared compute resource optimizes data-type interchange

POWER9 Core Pipeline Efficiency

Shorter Pipelines with Reduced Disruption

Improved application performance for modern codes

- **Shorten fetch to compute by 5 cycles**
- **Advanced branch prediction**

Higher performance and pipeline utilization

- **Improved instruction management**
  - Removed instruction grouping and reduced cracking
  - Enhanced instruction fusion
  - Complete up to 128 (64 – SMT4 Core) instructions per cycle

Reduced latency and improved scalability

- **Local pipe control of load/store operations**
  - Improved hazard avoidance
  - Local recycles – reduced hazard disruption
  - Improved lock management

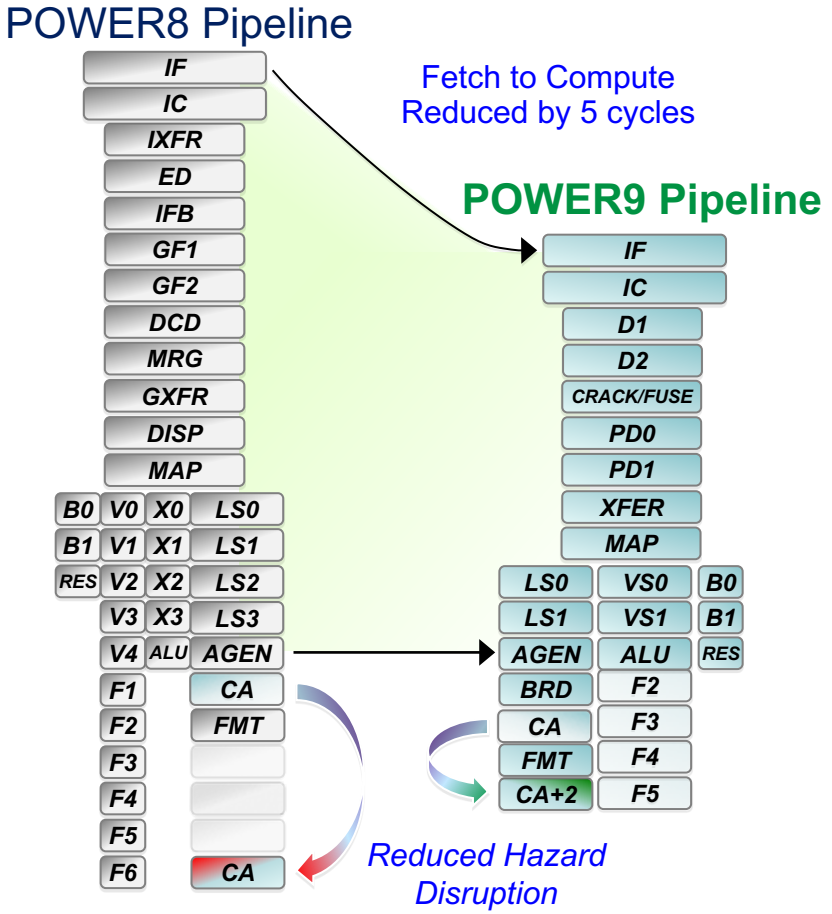

POWER9 – Core Compute

SMT4 Core Resources

Fetch / Branch
- 32kB, 8-way Instruction Cache
- 8 fetch, 6 decode
- 1x branch execution

Slices issue VSU and AGEN
- 4x scalar-64b / 2x vector-128b
- 4x load/store AGEN

Vector Scalar Unit (VSU) Pipes
- 4x ALU + Simple (64b)
- 4x FP + FX-MUL + Complex (64b)
- 2x Permute (128b)
- 2x Quad Fixed (128b)
- 2x Fixed Divide (64b)
- 1x Quad FP & Decimal FP
- 1x Cryptography

Load Store Unit (LSU) Slices
- 32kB, 8-way Data Cache
- Up to 4 DW load or store

Symmetric Engines Per Data-Type for Higher Performance on Diverse Workloads

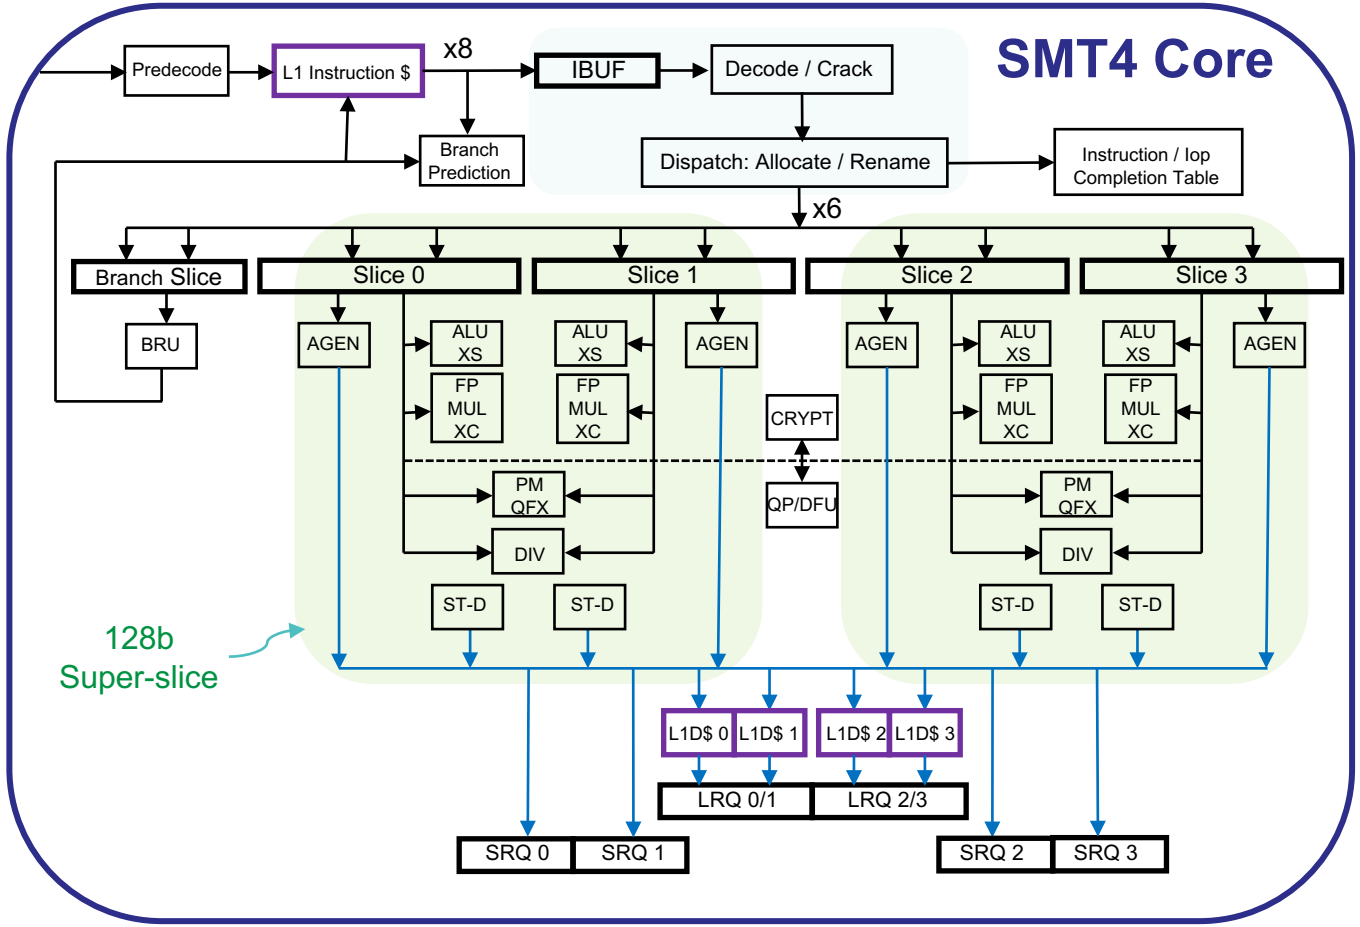

Efficient Cores Deliver 2x Compute Resource per Socket

POWER9 – Core Compute

POWER8 SMT8 Core Resources

Issue of VSU and AGEN
- 2x load AGEN / simple-ALU
- 2x load/store AGEN
- 2x scalar-64b / vector-128b
- 2x FXU

Vector Scalar Unit (VSU) Pipes
- 2x FP (64b/128b) + Complex (128b)
- 2x ALU (128b)
- 2x Permute (128b)
- 1x Decimal FP
- 1x Cryptography

Fixed Point (FXU) Pipes
- 2x ALU (64b)
- 2x FX-MUL + Fixed Divide (64b)

Load Store Unit (LSU) Slices
- 64kB, 8-way Data Cache
- Up to 4 DW load or 2 store
- 1x Store complete

4X Issue

2X VSU

4X DP

2X Ld Str

POWER9 SMT8 Core Resources

Issue of VSU and AGEN
- 8x scalar-64b / 4x vector-128b
- 8x load/store AGEN

Vector Scalar Unit (VSU) Pipes
- 8x ALU + Simple (64b slice)
- 8x FP + FX-MUL + Complex (64b slice)
- 4x Permute (128b)
- 4x Quad Fixed (128b)
- 4x Fixed Divide (64b)
- 2x Quad FP / Decimal FP
- 2x Cryptography

Load Store Unit (LSU) Slices
- 64kB, 8-way Data Cache
- Up to 8 DW load or store
- 2x Store complete

POWER9 – Dual Memory Subsystems

Scale Out
Direct Attach Memory

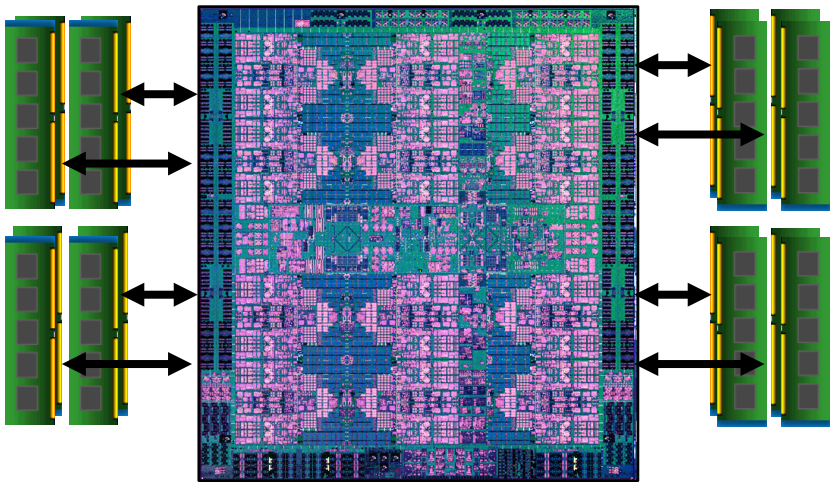

8 Direct DDR4 Ports

- Up to 120 GB/s of sustained bandwidth
- Low latency access
- Commodity packaging form factor
- Adaptive 64B / 128B reads

Scale Up
Buffered Memory

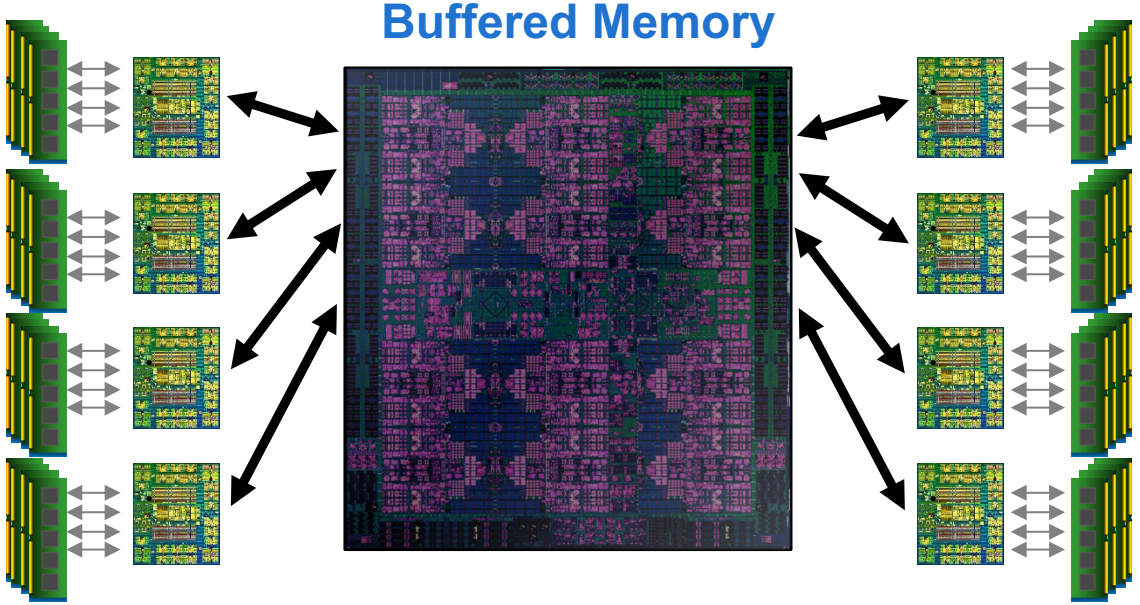

8 Buffered Channels

- 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

POWER9 Processor – Common Features

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

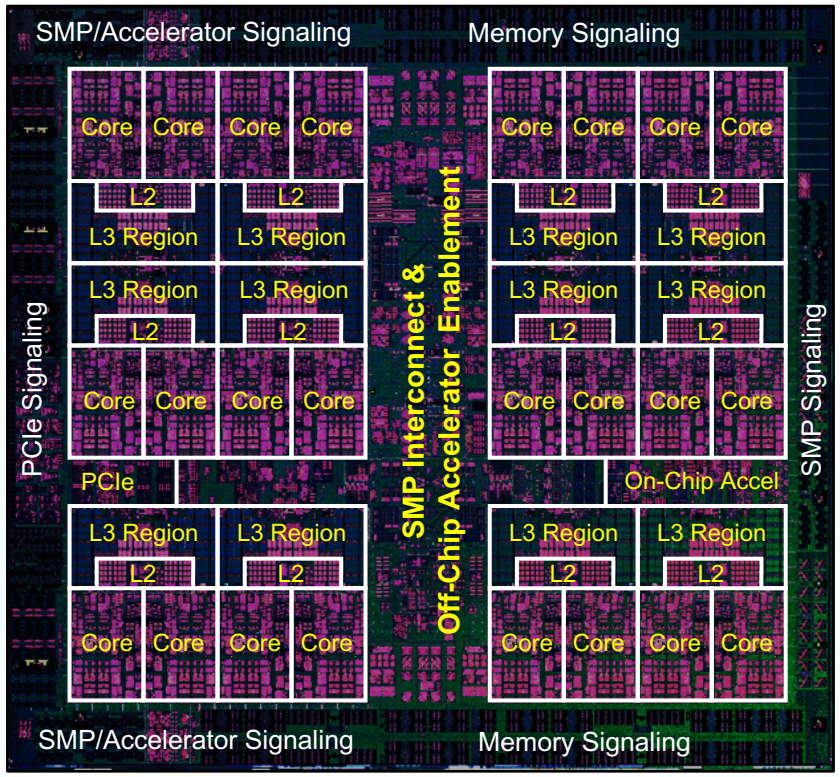

14nm finFET Semiconductor Process

- 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
- 25 Gb/s Common Link interface
  - – Accelerator, remote SMP

POWER9 Processor Family

Four targeted implementations

Core Count / Size

SMP scalability / Memory subsystem

SMT4 Core

24 SMT4 Cores / Chip

Linux Ecosystem Optimized

SMT8 Core

12 SMT8 Cores / Chip

PowerVM Ecosystem Continuity

Scale-Out – 2 Socket Optimized

Robust 2 socket SMP system

Direct Memory Attach

- Up to 8 DDR4 ports
- Commodity packaging form factor

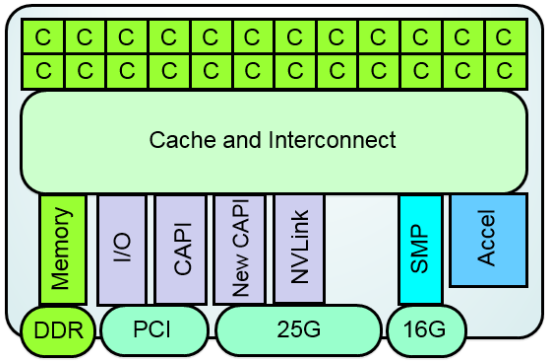

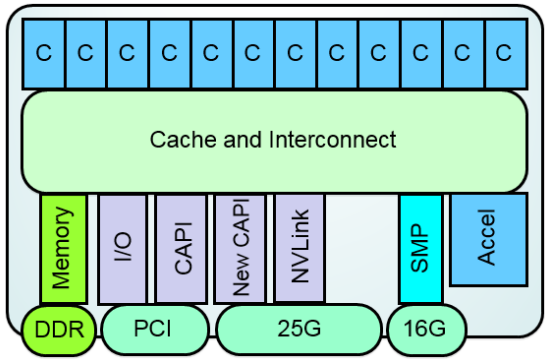

Scale-Up – Multi-Socket Optimized

Scalable System Topology / Capacity

- Large multi-socket

Buffered Memory Attach

- 8 Buffered channels

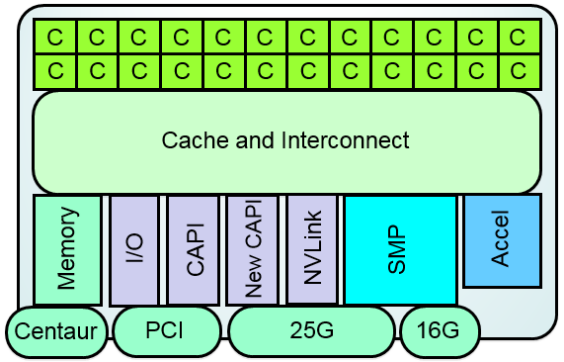

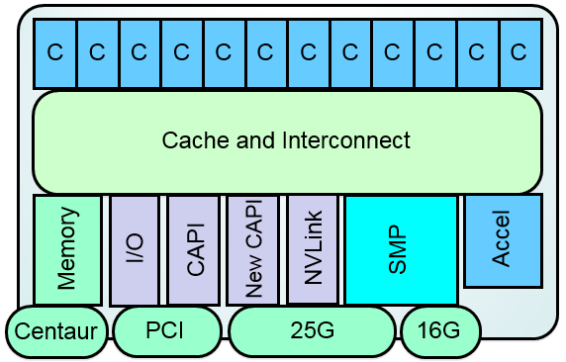

POWER9 – CPU Core Performance

Socket Performance

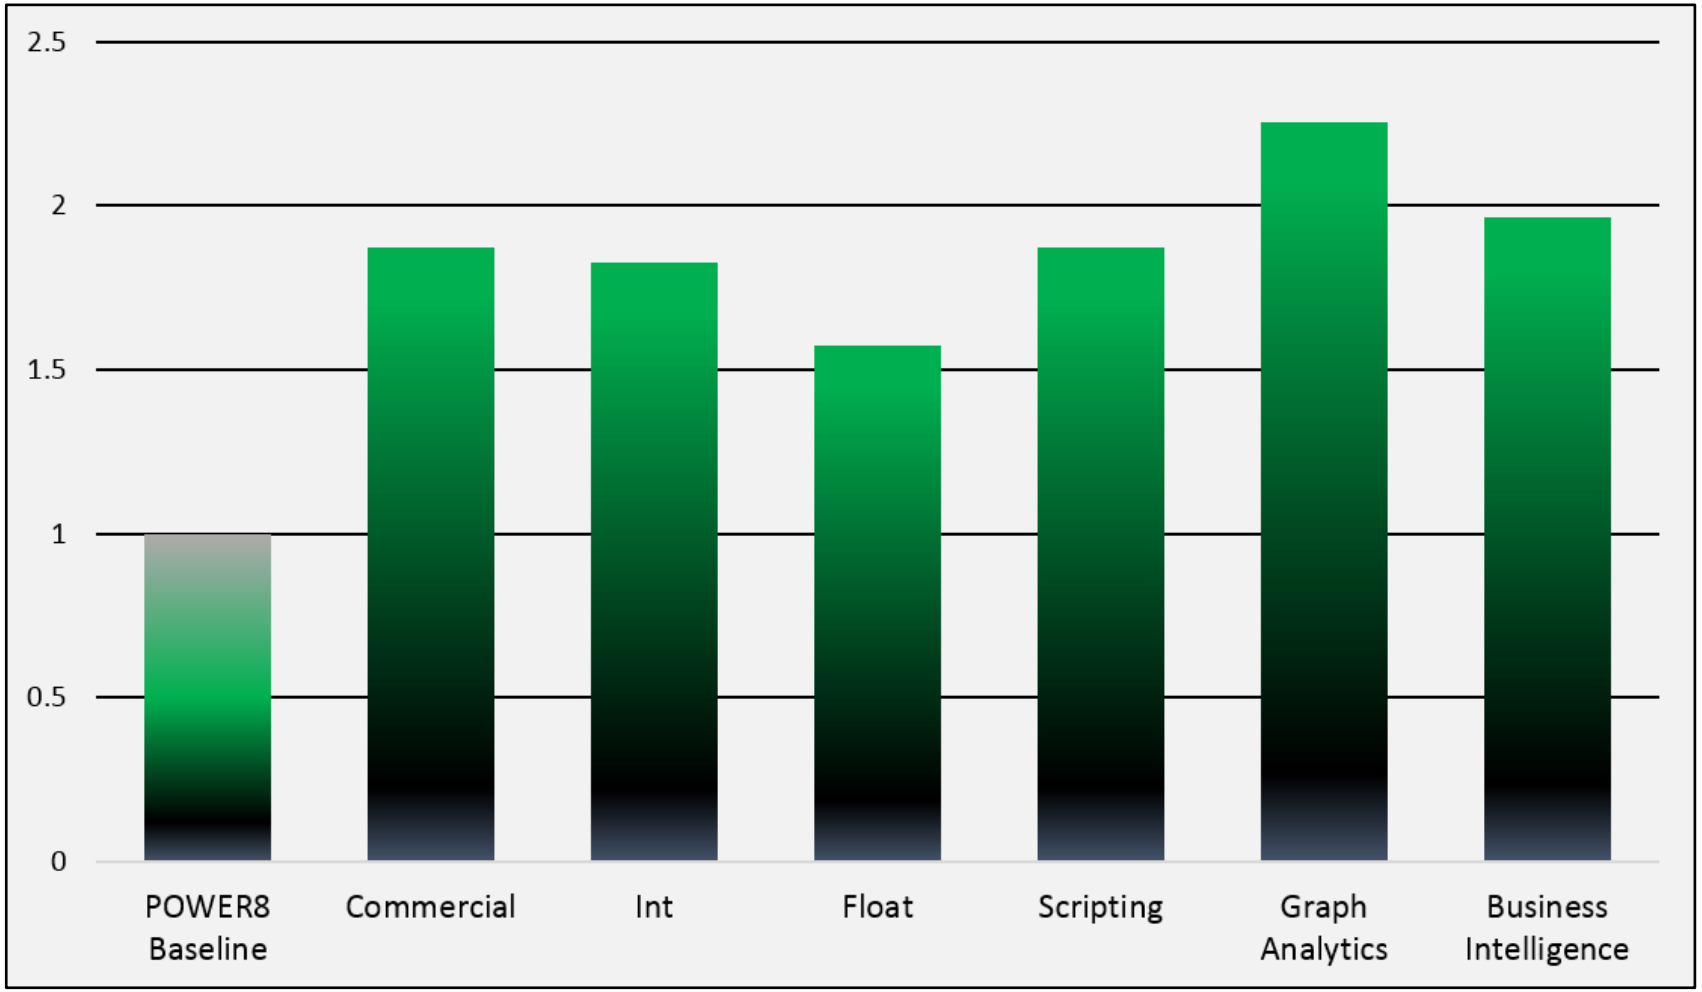

Scale-Out configuration @ constant frequency

POWER ISA v3.0

New Instruction Set Architecture Implemented on POWER9

Broader data type support

- 128-bit IEEE 754 Quad-Precision Float – Full width quad-precision for financial and security applications
- Expanded BCD and 128b Decimal Integer – For database and native analytics
- Half-Precision Float Conversion – Optimized for accelerator bandwidth and data exchange

Support Emerging Algorithms

- Enhanced Arithmetic and SIMD
- Random Number Generation Instruction

POWER

Accelerate Emerging Workloads

- Memory Atomics – For high scale data-centric applications
- Hardware Assisted Garbage Collection – Optimize response time of interpretive languages

Cloud Optimization

- Enhanced Translation Architecture – Optimized for Linux
- New Interrupt Architecture – Automated partition routing for extreme virtualization
- Enhanced Accelerator Virtualization
- Hardware Enforced Trusted Execution

Energy & Frequency Management

- POWER9 Workload Optimized Frequency – Manage energy between threads and cores with reduced wakeup latency

POWER9 – Data Capacity & Throughput

Big Caches for Massively Parallel Compute and Heterogeneous Interaction

L3 Cache: 120 MB Shared Capacity NUCA Cache

- 10 MB Capacity + 512k L2 per SMT8 Core
- Enhanced Replacement with Reuse & Data-Type Awareness
  12 x 20 way associativity

Extreme Switching Bandwidth for the Most Demanding Compute and Accelerated Workloads

High-Throughput On-Chip Fabric

- Over 7 TB/s On-chip Switch
- Move Data in/out at 256 GB/s per SMT8 Core

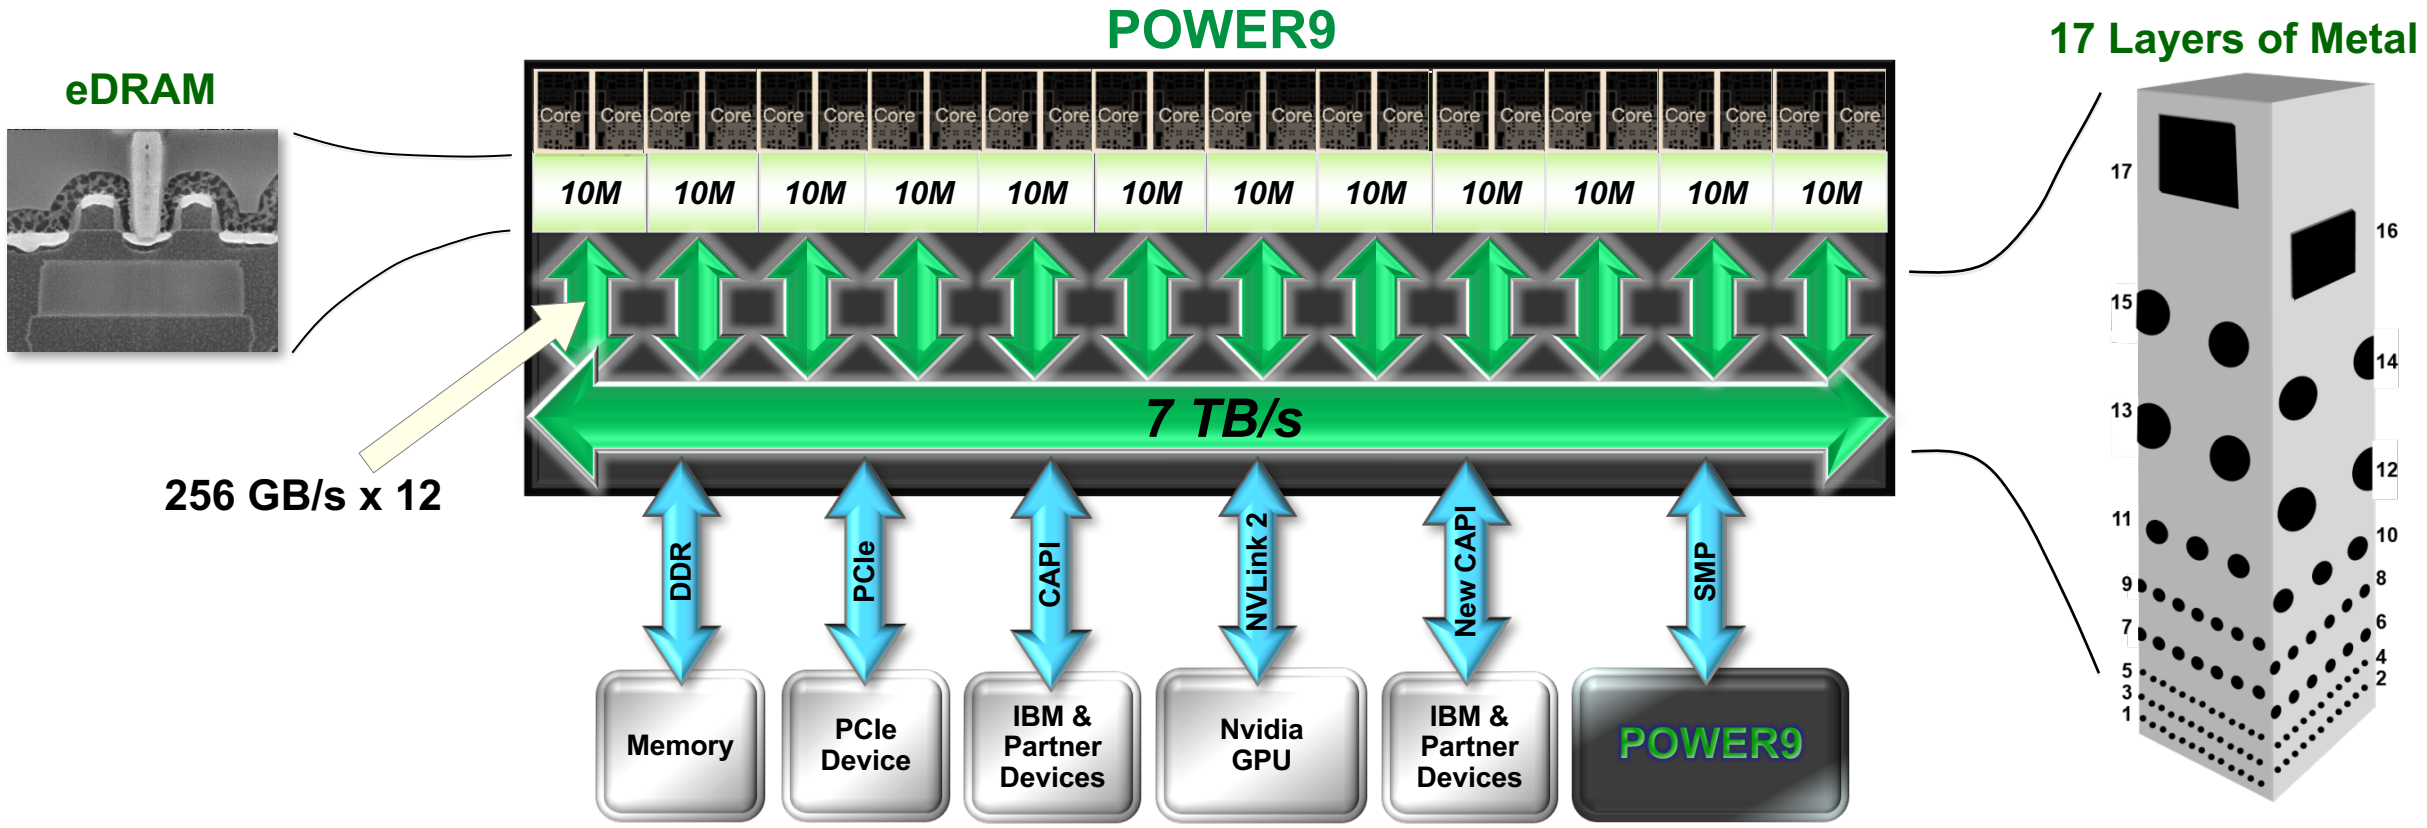

Modular Constructs → High-speed 25 Gb/s Signaling

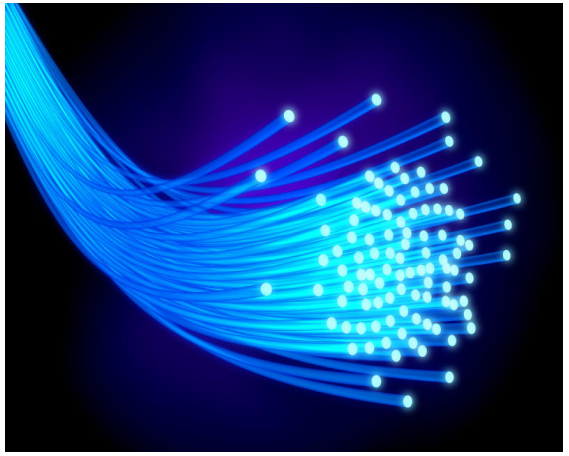

Utilize Best-of-Breed 25 Gb/s Optical-Style Signaling Technology

Multi-Drawer SMP Interconnect

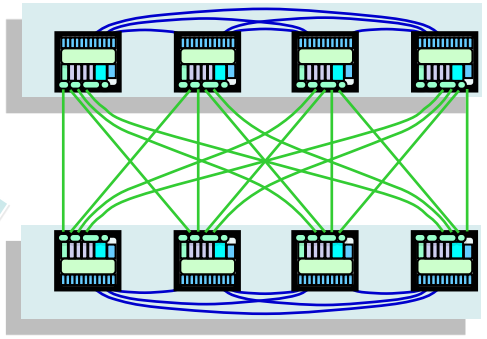

NVLINK 2 GPU Accelerator Attach

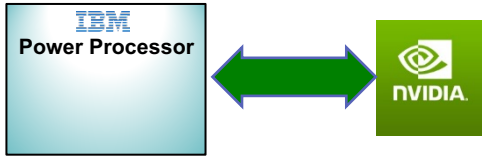

OpenCAPI Accelerator Attach

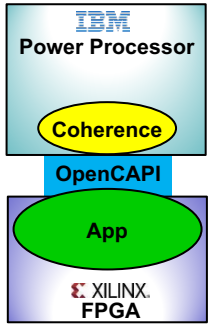

Flexible & Modular Packaging Infrastructure

POWER9 – Premier Acceleration Platform

- **Extreme Processor / Accelerator Bandwidth and Reduced Latency**
- **Coherent Memory and Virtual Addressing Capability for all Accelerators**
- **OpenPOWER Community Enablement – Robust Accelerated Compute Options**

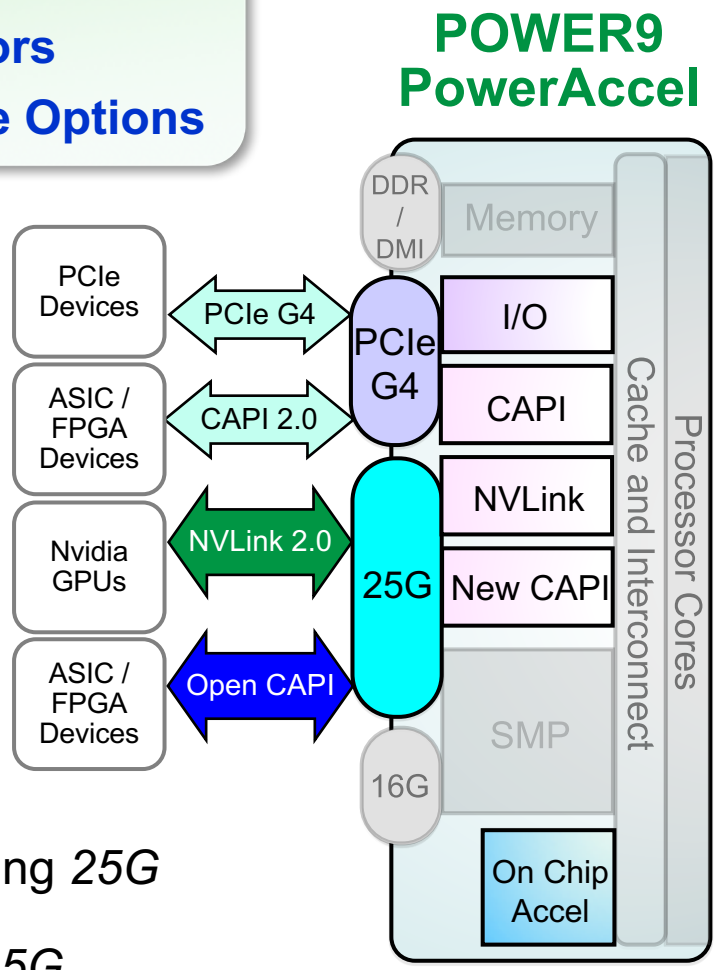

- **State of the Art I/O and Acceleration Attachment Signaling**
  - **PCIe Gen 4** x 48 lanes – 192 GB/s duplex bandwidth
  - **25Gb/s Common Link** x 48 lanes – 300 GB/s duplex bandwidth
- **Robust Accelerated Compute Options with OPEN standards**
  - **On-Chip Acceleration** – Gzip x1, 842 Compression x2, AES/SHA x2
  - **CAPI 2.0** – 4x bandwidth of POWER8 using *PCIe Gen 4*
  - **NVLink 2.0** – Next generation of GPU/CPU bandwidth and integration using *25G*
  - **Open CAPI 3.0** – High bandwidth, low latency and open interface using *25G*

POWER9 – Ideal for Acceleration

Extreme CPU/Accelerator Bandwidth

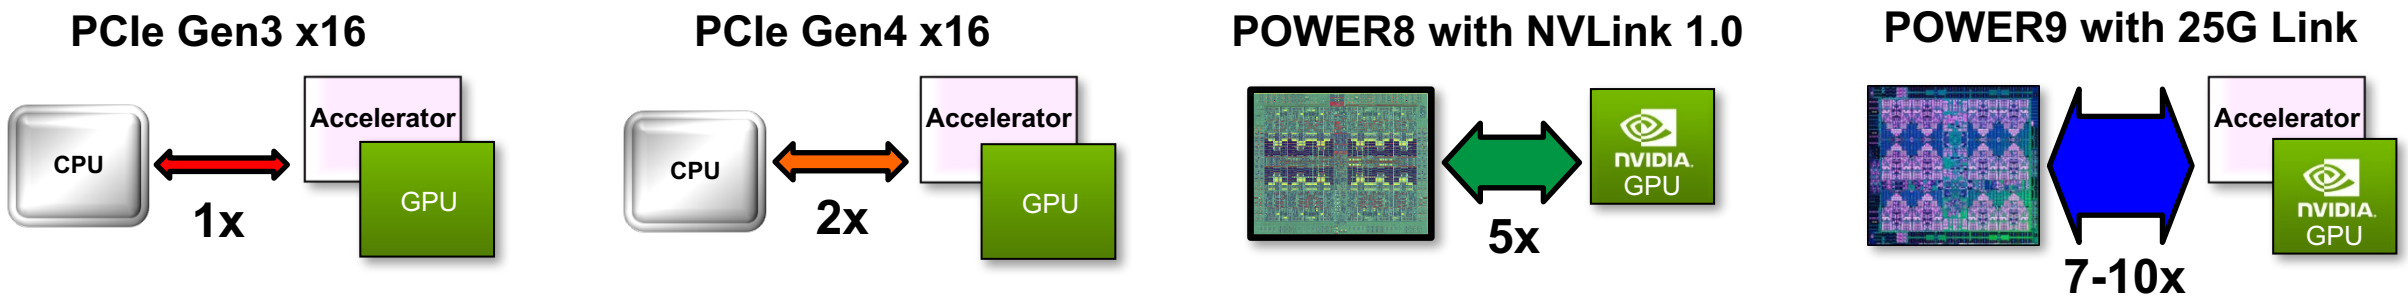

Increased Performance / Features / Acceleration Opportunity

Seamless CPU/Accelerator Interaction

- Coherent memory sharing
- Enhanced virtual address translation
- Data interaction with reduced SW & HW overhead

Broader Application of Heterogeneous Compute

- Designed for efficient programming models
- Accelerate complex analytic / cognitive applications

POWER9 – Ecosystem Enablement

OpenPOWER™ Foundation

- Accelerating Open Innovation
- Grown from 5 to over 200 members in less than 3 years

POWER9: Engineered for OpenPOWER Application

- Built for a Broad Range of Deployments and Platforms
- Open and Flexible Solutions
- Ideal for Developers

Power Systems

IBM

CAPI and OpenCAPI

Lance Thompson

POWER Systems, IBM Systems

What is CAPI/OpenCAPI

- Hardware, protocols and operating system enablement that allow attached devices to share virtual addresses with the applications requesting services from the devices
  - Accelerator
  - Memory
  - Network
  - Storage
  - Etc.

- Communication over open interfaces
  - CAPI via PCIe (pcisig.com)
  - OpenCAPI via OpenCAPI (opencapi.org)

A heterogeneous architecture with OpenCAPI

OpenCAPI

1. **Accelerators:** The performance, virtual addressing and coherence capabilities allow FPGA and ASIC accelerators to behave as if they were integrated into a custom microprocessor.

2. **Coherent Network Controller:** OpenCAPI provides the bandwidth that will be needed to support rapidly increasing network speeds Network controllers based on virtual addressing can eliminate software overhead without the programming complexity usually associated with user-level networking protocols.

3. **Advanced Memory:** OpenCAPI allows system designers to take full advantage of emerging memory technologies to change the economics of the datacenter.

4. **Coherent Storage Controller:** OpenCAPI allows storage controllers to bypass kernel software overhead, enabling extreme IOPS performance without wasting valuable CPU cycles.

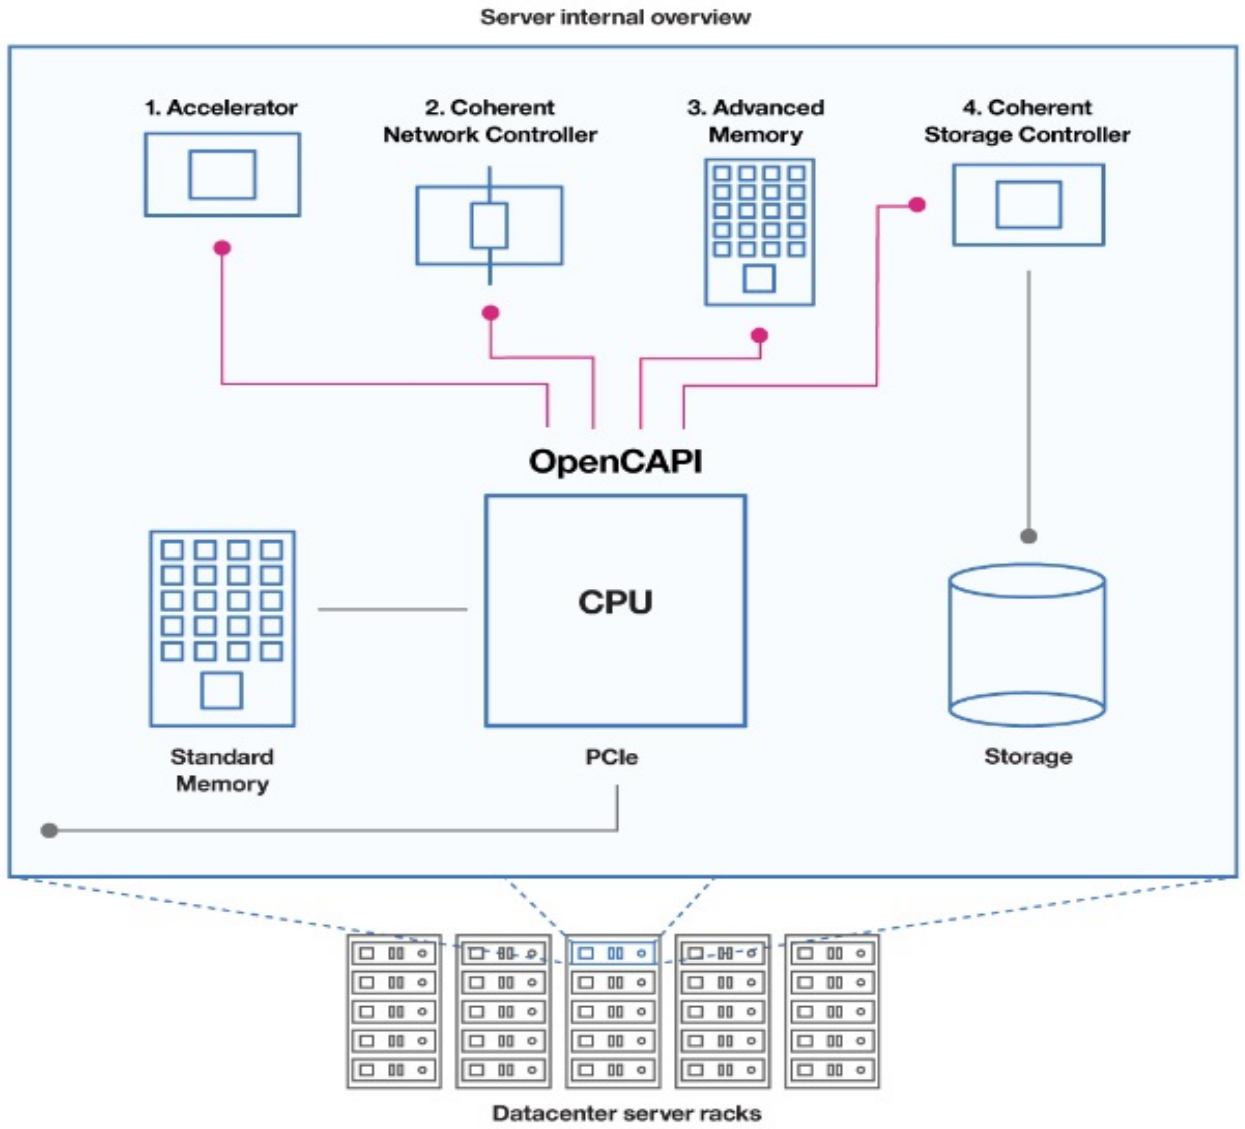

CAPI technology connections

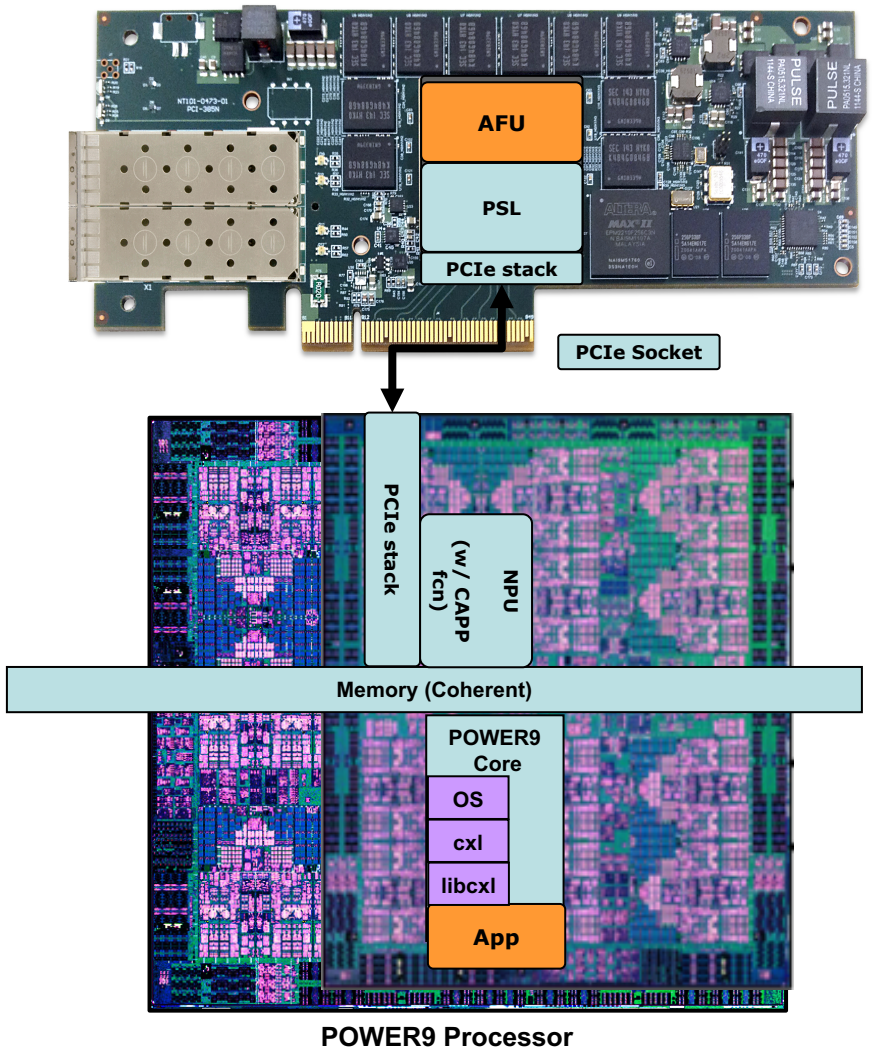

- Proprietary hardware for Memory Coherence
- CAPI compliant designs bridging to proprietary coherence protocols
- Operating system enablement
  - little endian linux
  - kernel driver (cxl)
  - user library (libcxl)
- Customer application and accelerator

- Open
  - PSL/AFU interface available from OpenPOWER.org
  - encrypted PSL also available
- shared virtual address space with host application
  - AFU utilizes the same virtual addresses as the Host Application

OpenCAPI technology connections

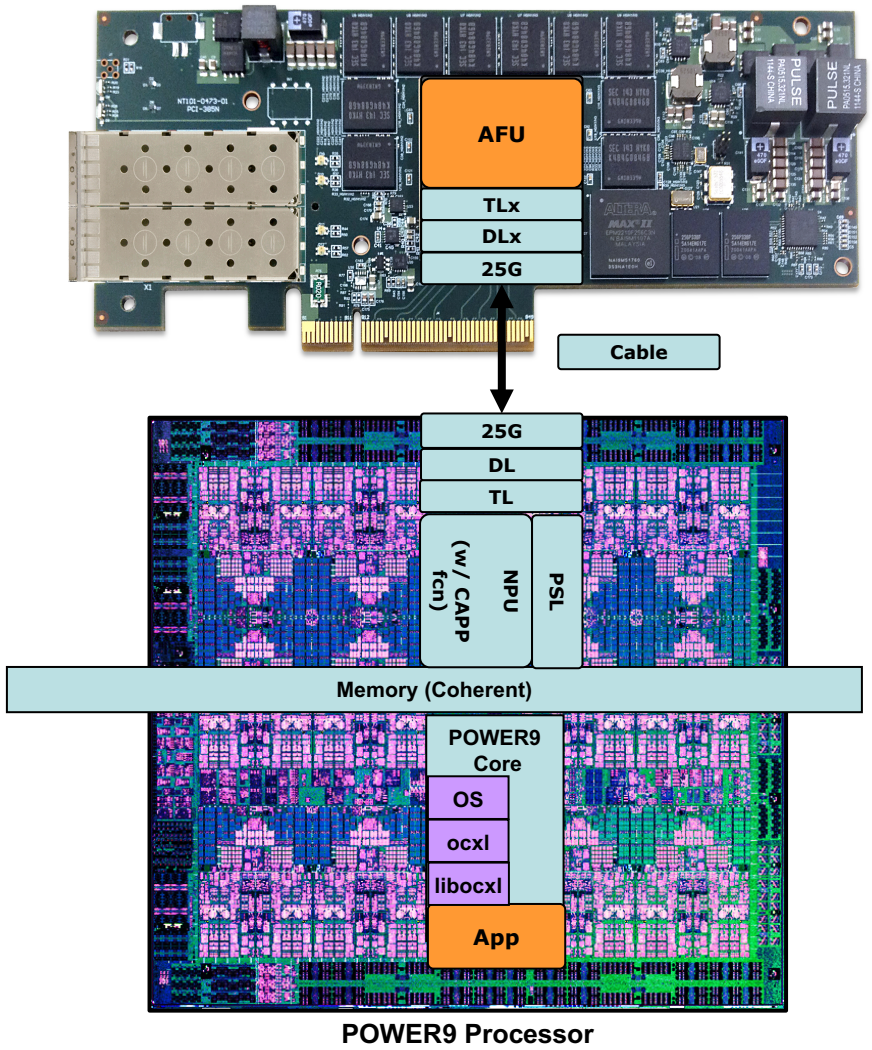

- Proprietary hardware for Memory Coherence
- OpenCAPI compliant designs and reference designs bridging to proprietary coherence protocols
- Operating system enablement
  - little endian linux
  - reference kernel driver (ocxl)
  - reference user library (libocxl)
- Customer application and accelerator

- Open
  - TLx/DLx/25G available through OpenCAPI.org
- low latency due to reduced TLx/DLx stack
- high bandwidth due to 25G phy
- shared virtual address space with host application
  - AFU utilizes the same virtual addresses as the Host Application

Reference Card Design

- Definition of FPGA reference card is being driven as part of the 25G workgroup within the OpenPower consortium
- Definition of the cable(s) are also driven as part of the 25G workgroup within the OpenPower consortium
- Currently IBM and Xilinx are driving the initial definition of a PCIE based form factor card
  - Representative Diagram is articulated below

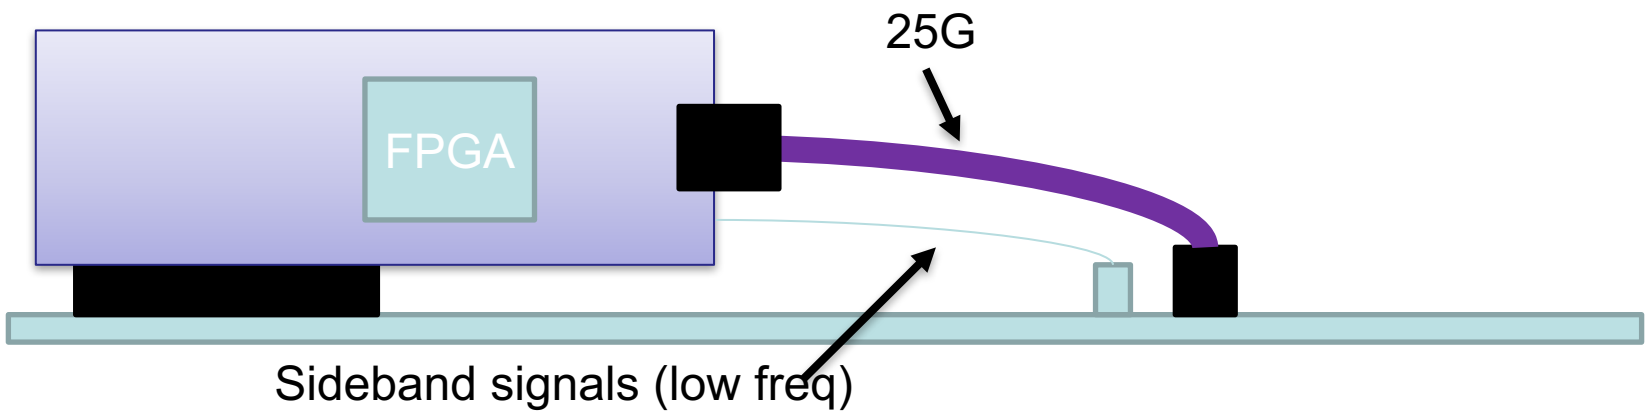

Development Environments

CAPI

- HDK/SDK
  - VHDL/Verilog FPGA design to PSL/AFU interface for accelerator development
  - low-ish level libcxl software interface for application
    - from github.com/ibm-capi/libcxl
  - Power Service Layer Simulation Engine (pslse) to co-verify application with accelerator
    - from github.com/ibm-capi/pslse
- CAPISNAP
  - layers on top the the HDK/SDK
  - permits C/C++ code for FPGA accelerator content
  - leverages Xilinx High Level Synthesis product
  - high level software interface for application
  - github.com/open-power/snap

OpenCAPI

- HDK/SDK
  - VHDL/Verilog FPGA design to TLx/AFU interface for accelerator development
    - interface definition from OpenCAPI.org
    - reference TLx/DLx designs from OpenCAPI.org
  - low-ish level libocxl software interface for application
  - OpenCAPI Simulation Engine (ocse) to co-verify application with accelerator
    - soon from OpenCAPI.org
- CAPISNAP
  - watch this space

Power Systems

IBM

Systems

Lance Thompson

POWER Systems, IBM Systems

POWER8 Scale-Out Family

S822
1 or 2 socket, 2U
4 (AIX) / 6 - 20 cores

S814
1 socket, 4U
4 - 8 cores

S824
2 socket, 4U
6 - 24 cores

S812L
1 socket, 2U, Linux
10 - 12 cores

S822L
2 socket, 2U, Linux
16 - 24 cores

S824L
2 socket, 4U, Linux
8 - 24 cores

OpenPOWER Linux Cluster (LC) Systems

S812LC

1 socket, 2U, Linux
8 or 10 cores
Up to 1 TB memory
Up to 112 TB Storage
12+2 Disk Bays
4 Available PCI Slots
KVM / Bare Metal

Hadoop / Spark

S822LC – GCA

2 socket, 2U, Linux
16 or 20 cores
Up to 1 TB memory
2 Disks
5 Available PCI slots
KVM / Bare Metal

Commercial

S822LC – GTA

2 socket, 2U, Linux
16 or 20 cores
Up to 1 TB memory
2 Disks
2 NVIDIA K80 GPUs
3 Available PCI Slots
Bare Metal

Technical

OpenPOWER LC Systems

S822LC for Big Data

Up to 20 Cores
Up to 512 GB memory
12 HDD/SSD / NVMe
5 PCI Slots
(4 CAPI, 2 K80 GPU)
4-Port 10 Gb NIC
KVM or Bare Metal

Big Data

S822LC for HPC

Up to 20 Cores
POWER8 with NVLink
Up to 1 TB memory
2 HDD/SDD
3 PCI Slots. 2 CAPI
Up to 4 P100 GPU
Bare Metal

HPC

S821LC

Up to 20 Cores
Up to 512 GB memory
4 HDD/SDD / NVMe
4 PCI Slots
(3 CAPI, 1 K80 GPU)
4 Port 10 Gb NIC
KVM or Bare Metal

Compute

IBM POWER9 Family

- More performance and scale via POWER9 processors
- More Memory capacity for in-memory DB
- Reduce latency and improve throughput with enhanced I/O support
  - PCIe Gen4
  - Integrated NVMe Flash (bootable)
- High-bandwidth (25Gb/s) links for GPU/OpenCAPI acceleration

2017

Robust 2-socket SMP
Direct memory attach

OpenPOWER

&
HSDC

AC922

Linux

2018

Scale Out 2-socket SMP
Direct memory attach

LC921
LC922

2-socket LC Infrastructure

H922
H924
S922
S944
L922

2 socket scale-out servers

Large-scale multi-socket SMP
Buffered memory attach

To Be Announced

4-socket Midrange

4- to16-socket Modular High-end

i for Business

Linux

AIX

AC922

An Acceleration Superhighway

Unleash accelerated computing potential in the post CPU-only era

Designed for the AI Era

Architected for the modern analytics and AI workloads that fuel insights

Delivering Enterprise-Class AI

Cutting-edge AI innovation data scientists desire, with dependability IT requires

AC922 4 GPU Air Cooled

- **Processor**
  - Two p9SO SCMs
    - 190W
    - Cores: 16c, 20c (SMT4)
- **Memory**
  - Maximum of 16 DDR4 IS RDIMM slots
    - Direct attach to P9 module
  - 8, 16, 32, 64, 128 GB DIMMs @ 2667 Mhz
- **Internal Storage**
  - RAID Controller: TBD
  - 0 to 2 SFF disks, default is 0 disks
  - PCIe attached non rotating storage, device and capacity is TBD
- **PCIe Gen4 Slots**
  - 2 PCIe x16 Gen4 LP slot
  - 1 PCIe x4 Gen4 LP Slot
  - 1, Dual port IB EDR NIC, shared by two P9 sockets.
- **GPU**
  - Up to 4 NVIDIA Volta GPU's (300W max)
  - 2 GPU's per P9 socket
    - SXM2 form factor, NVLink 2.0
- **Native I/O**
  - Host 2-port USB 3.0 (1F, 1R)
  - Management port
- **Enclosure Form Factor**
  - 2U 19" rack enclosure, 710mm (28") deep
    - Enclosure can be fixed rail or slide mounted
- **O/S:** Linux only
- **Sapphire, Opal**
- **RAS**
  - 5 year MTBF
  - Concurrent maintenance HDD (when installed)
  - N+1 redundant hot swap cooling
  - Concurrent maintenance of cooling
  - 2 1900W bulk power supplies
    - Non-redundant for HPC
    - 200VAC, 277VAC, 400VDC input voltage support
  - Support for 480V 3 phase rack input, requires 277VAC supplies
- **Energy Efficiency**
  - On-chip power management, power gating
  - Advanced thermal management
  - 80+ Platinum Power Supply Compliant
  - EPA Energy Star Compliant
- **Service Interface**
  - Industry BMC service controller w/ OpenPOWER firmware
  - Operator interface & FRU LEDs
- **Certifications**
  - FCC: Class A for Servers
  - Acoustics: Rack, Data Center Category 1A
  - Environment: ASHRAE A3 (5-40C, 8-85% RH, 3050m max)

AC922

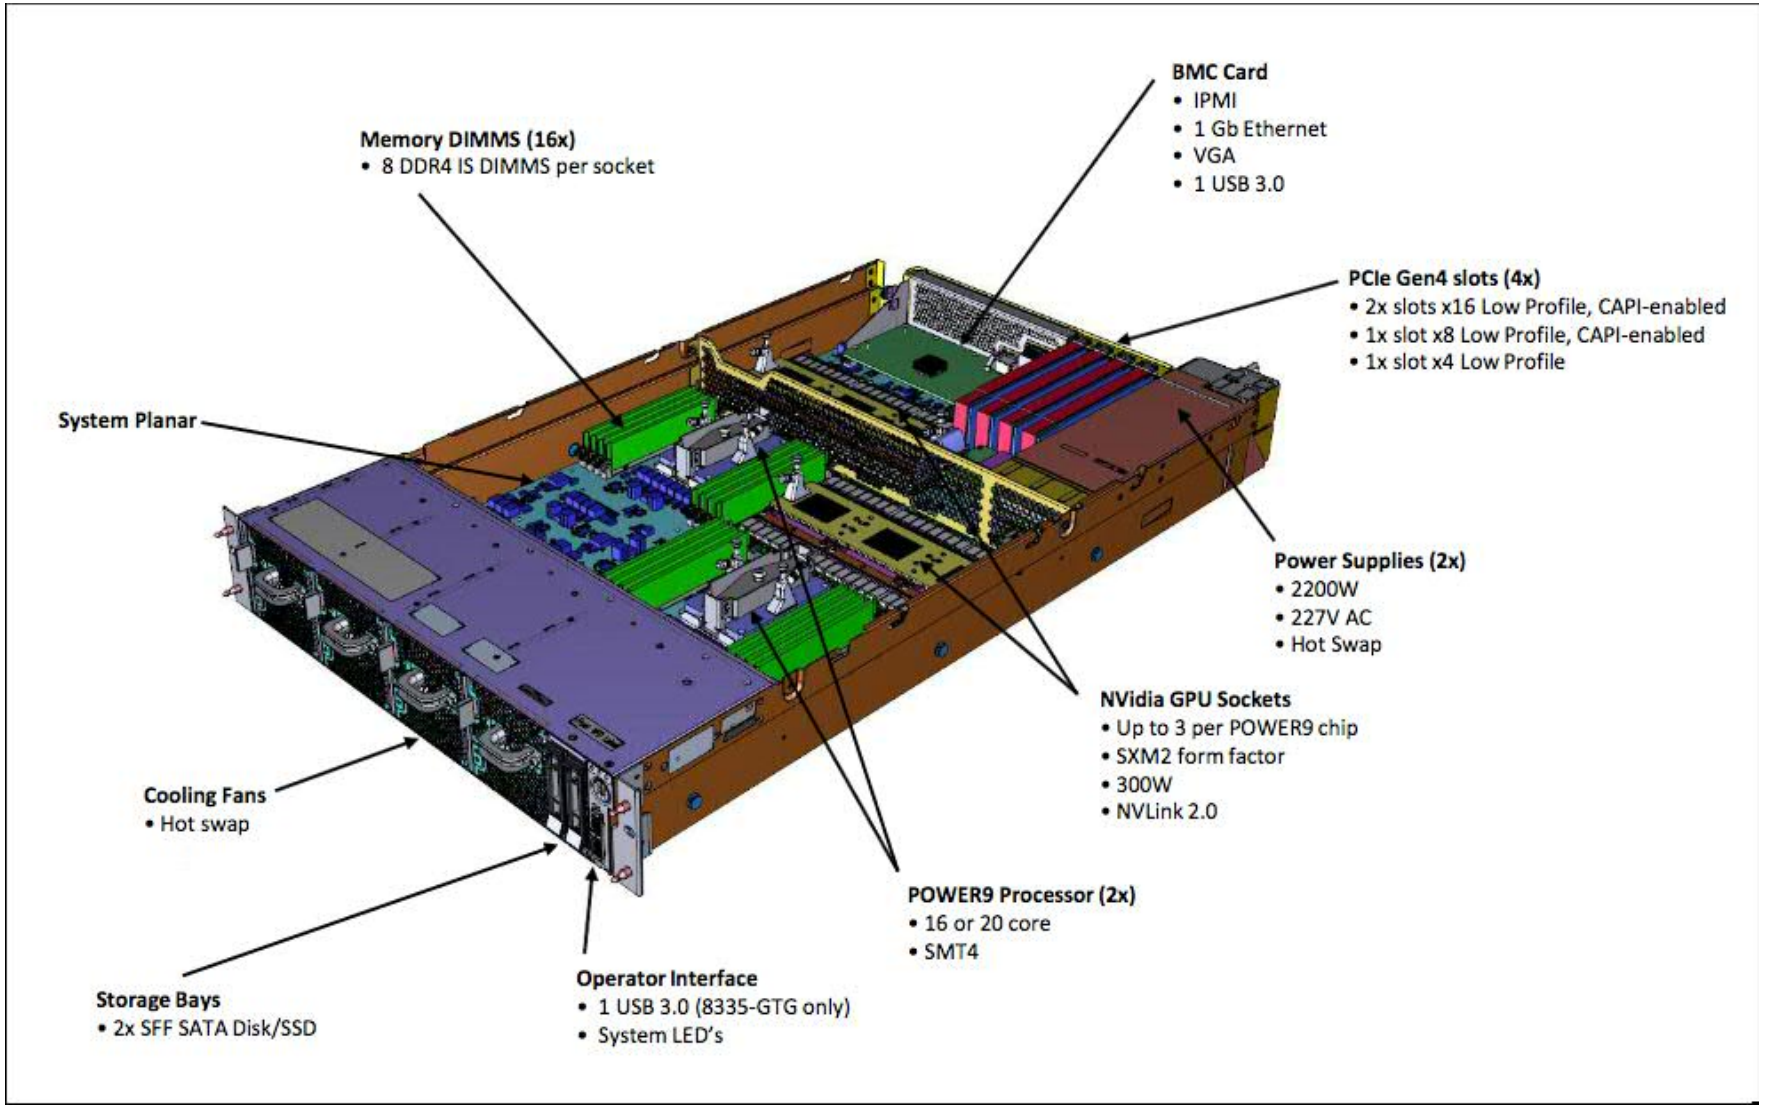

LC922 – 2 Socket 2U Server

- Form Factor
  - 19” Rack 2U Server
- P9 Processor
  - Up to 2 P9 chips
  - 190W
  - 2.91 GHz/2.7 GHz/2.6 GHz
  - 16/20/22 SMT4 Cores / socket
- Memory
  - Direct Attach Memory
  - 16 DDR4 IS DIMM Slots @ 2400 Mhz
  - 8, 16, 32, 64, 128 GB RDIMMs
  - 2 TB Max memory
  - 77 GB/s peak memory BW per socket
- 6 Integrated I/O Slots 2U
  - 2 PCIe G4 x16 FHFLDW Slots (Supports GPU cards)
  - 3 PCIe G4 x8 FHFL Slots
  - 1 PCIe G4 x8 LP Slot
- Internal Storage
  - Storage Controller – MicroSemi PM8069
  - 12 LFF Bays (4 Bays NVMe enabled)
  - 2 SATA DOM Flash modules
- Native I/O
  - 5 USB 3.0
  - 4x 10GbaseT (1 shared mgmt) + 1x 1Gb dedicated IPMI
  - Serial port, VGA port

- OS Support
  - Red Hat Enterprise Linux (RHEL) 7.5
  - Ubuntu Server 18.04 LTS
- RAS
  - Concurrent Maintenance disks
  - Redundant Hot plug Power
  - Redundant cooling (TBD)
  - Customer Setup and Install
  - Simplified Op Panel
- BMC Service Processor
- Certifications
  - FCC Class A
  - ASHRAE A2 Environment (10-35C)
  - Acoustics Datacenter 1A

LC922 – 2 Socket 1U Server

- **Form Factor**
  - 19” Rack 1U Server
- **P9 Processor**
  - Up to 2 P9 chips
  - 140W/160W
  - 2.2 GHz/2.13 GHz
  - 16/20 SMT4 Cores / socket
- **Memory**
  - Direct Attach Memory
  - 16 DDR4 IS DIMM Slots @ 2400 Mhz
  - 8, 16, 32, 64, 128 GB RDIMMs
  - 2 TB Max memory
  - 77 GB/s peak memory BW per socket
- **4 Integrated I/O Slots 1U**
  - 2 PCIe G4 x16 FHFL Slot (One supports GPU card)
  - 1 PCIe G4 x8 LP Slot (internal)
  - 1 PCIe G4 x8 LP Slot
- **Internal Storage**
  - Storage Controller – MicroSemi PM6089
  - 4 LFF Bays (4 Bays NVMe enabled)
  - 2 SATA DOM Flash modules
- **Native I/O**
  - 5 USB 3.0
  - 4x 10GbaseT (1 shared mgmt) + 1x 1Gb dedicated IPMI
  - Serial port, VGA port

- **OS Support**
  - Red Hat Enterprise Linux (RHEL) 7.5
  - Ubuntu Server 18.04 LTS
- **RAS**
  - Concurrent Maintenance disks
  - Redundant Hot plug Power
  - Redundant cooling (TBD)
  - Customer Setup and Install
  - Simplified Op Panel
- **BMC Service Processor**
- **Certifications**
  - FCC Class A
  - ASHRAE A2 Environment (10-35C)
  - Acoustics Datacenter 1A

Thank you

Lance Thompson

POWER Systems, IBM Systems

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquiries, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Revised September 26, 2006

Special notices (continued)