



Power Systems AC922 Overview

Chris Mann
IBM Distinguished Engineer
Chief System Architect, Power HPC Systems
December 11, 2017

IBM POWER HPC Platform Strategy

- High-performance computer and high-performance analytics drive common platform design
- Servers will be predominately 2-socket designs
- Developing deeper relationships with technology partners – ref OpenPOWER
- Majority of floating-point performance will come from GPUs
- Utilize Industry-standard compliant 19” racks and electronics enclosures
 - Air and water cooling options
- Platforms will be based on a common enclosure form factor
 - Enclosure provides working envelope that we will continue to enhance with the latest technology from IBM, NVIDIA, Mellanox and other OpenPOWER partners
 - Enclosure provides a platform with sufficient power, cooling capability to support these enhancements

IBM POWER GPU Intensive Roadmap



POWER S822LC

- 2 POWER8 Processors
 - 190 Turismo module
- 2 x16 Gen 3 FHFL PCIe slots
 - Supports 2 NVidia K80 GPU's
 - Supports 2 PCIe adapters
- 1 x8 Gen 3 HHHH PCIe, CAPI
- 1 x16 Gen 3 HHHH PCIe, CAPI
- 1 x8 Gen 3 PCIe
- 32 DDR3 IS DIMM's
 - 4, 8, 16, 32GB DIMMs
 - 32 – 1024GB Memory Capacity
- 2 SATA SFF HDD / SSD
- 2 1300W Power Supplies
 - 200VAC Input
- BMC support structure
 - IPMI, USB, EN, VGA
- Air cooled

2015



POWER S822LC for HPC

- 2 POWER8 w/ NVLink Processors
 - 190 module
- 1, 2, 4 NVidia "Pascal" GPU's
 - 300W, SXM2 Form Factor, NVLink 1.0
- 2 x16 Gen 3 HHHH PCIe, CAPI enabled
- 1 x8 Gen3 HHHH PCIe, CAPI enabled
- 32 DDR4 IS DIMM's
 - 4, 8, 16, 32GB DIMM's
- 2 SATA SFF HDD / SSD
- Pluggable NVMe storage adapter
 - 1.6, 3.2TB Capacity
- 2 1300W power supplies
 - 200VAC Input
- BMC Support Structure
 - IPMI, USB, EN, VGA
- Air and water cooled options

2016



POWER AC922

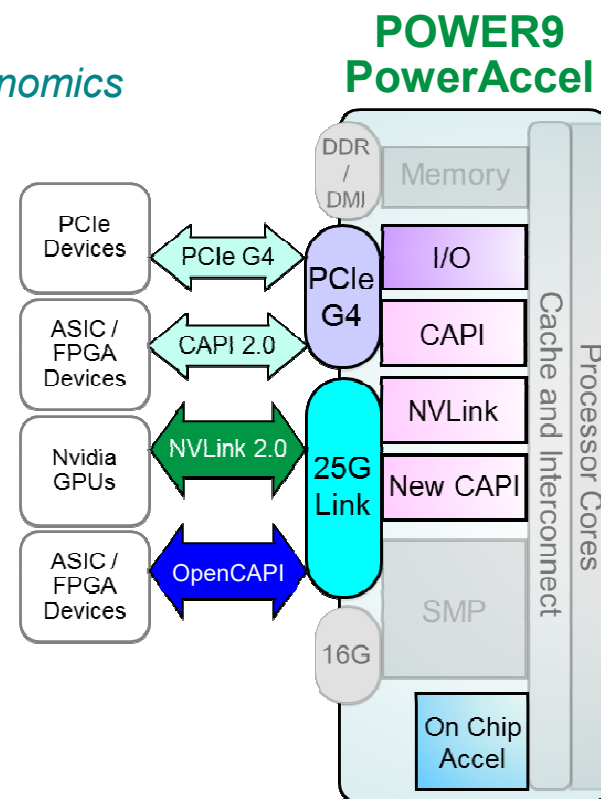
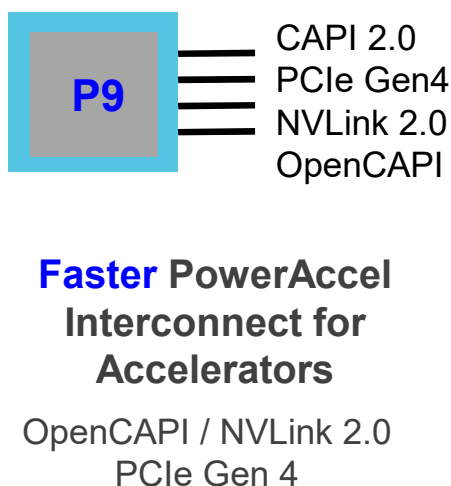

- 2 POWER9 Processors
 - 190, 250W modules
- 4-6 NVidia "Volta" GPU's
 - 300W, SXM2 Form Factor, NVLink 2.0
- 6 GPU configuration, water cooled
- 4 GPU configuration, air or water cooled
- 2 Gen4 x16 HHHH PCIe, CAPI enabled
- 1 Gen4 x4 HHHH PCIe
- 1 Gen4 Shared x8 PCIe adapter
- 16 IS DIMM's
 - 8, 16, 32, 64, 128GB DIMMs
- 2 SATA SFF HDD / SSD
- 2 2200W power supplies
 - 200 VAC, 277VAC, 400VDC input
 - N+1 Redundant
- Second generation BMC Support Structure
- Pluggable NVMe storage adapter option

2017 - 2018



Introducing Power AC922 2U HPC Server – Premier Accelerator Platform Leveraging **POWER9** Innovation High-performance core, bandwidth, accelerator differentiation

For compute-intensive workloads, accelerators are key to Cognitive/AI economics



Power AC922 - POWER9 with increased GPU and IO bandwidth for differentiation

Realize unprecedented performance and application gains with POWER9 and NVLink 2.0

- 2 POWER9 CPUs and up to 4 “Volta” NVLink 2.0 GPUs in a versatile 2U Linux server
- PCIe Gen4 bus has double I/O Bandwidth vs. PCIe Gen3
- CPU (Turbo)/GPU (Boost) enabled for improved data center efficiency and performance to be maintained at high levels

High level System Overview

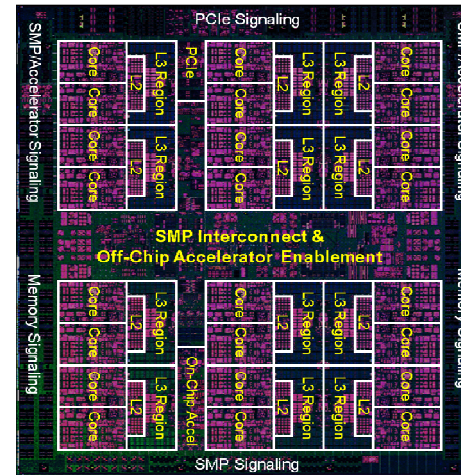
- 2-Socket, 2U Packaging
- 40 P9 Processor cores
- 4 NVIDIA Volta 2.0 GPUs
- 1 TB Memory (16x - 64GB DIMMs)
- 4 PCIe Gen4 Slots
- 2x SFF (HDD/SSD), SATA, Up to 7.7 TB storage
- Supports 1.6TB and 3.2TB NVMe Adapters
- Default 3 year 9x5 warranty, 100% CRU



POWER9 HPC Product Description


POWER9 Processor – Common Features

- **14HP Technology**
- **24 POWER9 Cores with new SMC uArchitecture**
 - 4 threads per core
 - 2 cores per chiplet
 - LINUX Radix Page Table support
- **Large, Low-latency Cache**
 - 512k private L2, 10MB NUCA L3 per chiplet
- **Direct Attach Memory Support**
 - 8 DDR4 channels
 - 2666 MHz DIMM support
 - 140 GB/s streaming bandwidth
- **SMP 2 Socket Support via 4B, 16 Gb/s X-Bus**
- **Leadership Hardware Acceleration Platform**
 - Enhanced on-chip acceleration
 - NVIDIA NVLink 2.0
 - CAPI 2.0, Coherent accelerator and storage attach via PCIe G4
 - OpenCAPI 3.0, Improved latency and bandwidth, open interface
- **Network Interconnect**
 - CAPI 2.0 attached Mellanox CX-5 support
 - Address translation support
 - Atomics
 - NVIDIA GPU direct



- **Gen4 PCIe, 48 Lanes @ 16Gb/s**
 - 6 PHBs
 - PCIe Bus 0 turbo-charged to support next generation HDR IB
- **Self Boot Capability**
- **Instant ON/OFF**
- **Cloud Management QoS**

POWER AC922 Design – 4 GPU



- PCIe slot (4x)**
- Gen4 PCIe
 - 2, x16 HHHHL Adapter
 - 1, Shared slot
 - 1 x8 HHHHL Adapter

- Power Supplies (2x)**
- 2200W
 - 200VAC, 277VAC, 400VDC input

- NVidia Volta GPU**
- 3 per socket
 - SXM2 form factor
 - 300W
 - NVLink 2.0
 - Air/Water Cooled

- Memory DIMM's (16x)**
- 8 DDR4 IS DIMMs per socket
 - 8, 16, 32,64, 128GB DIMMs

- BMC Card**
- IPMI
 - 1 Gb Ethernet
 - VGA
 - 1 USB 3.0

- Power 9 Processor (2x)**
- 18, 22C water cooled
 - 16, 20C air cooled

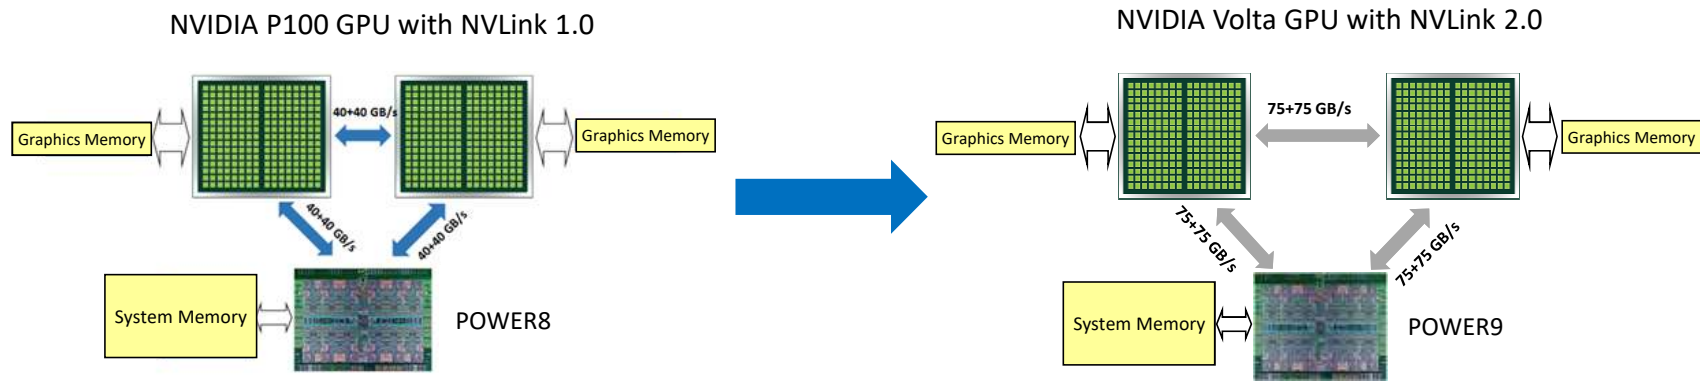
NVIDIA Volta Specifications

NVIDIA Volta GPU Features	
Peak double precision floating point performance	7.8 TFLOPS
Memory bandwidth	900 GB/sec
GPU Memory Size	16 GB
NVLink “Bricks” (8 lane interface)	6
NVLink Interconnect Bi-Directional	300GB/s
Maximum Power	300W



<https://www.nvidia.com/en-us/data-center/tesla-v100/>

NVLink Evolution in POWER HPC




2016

2017-2018

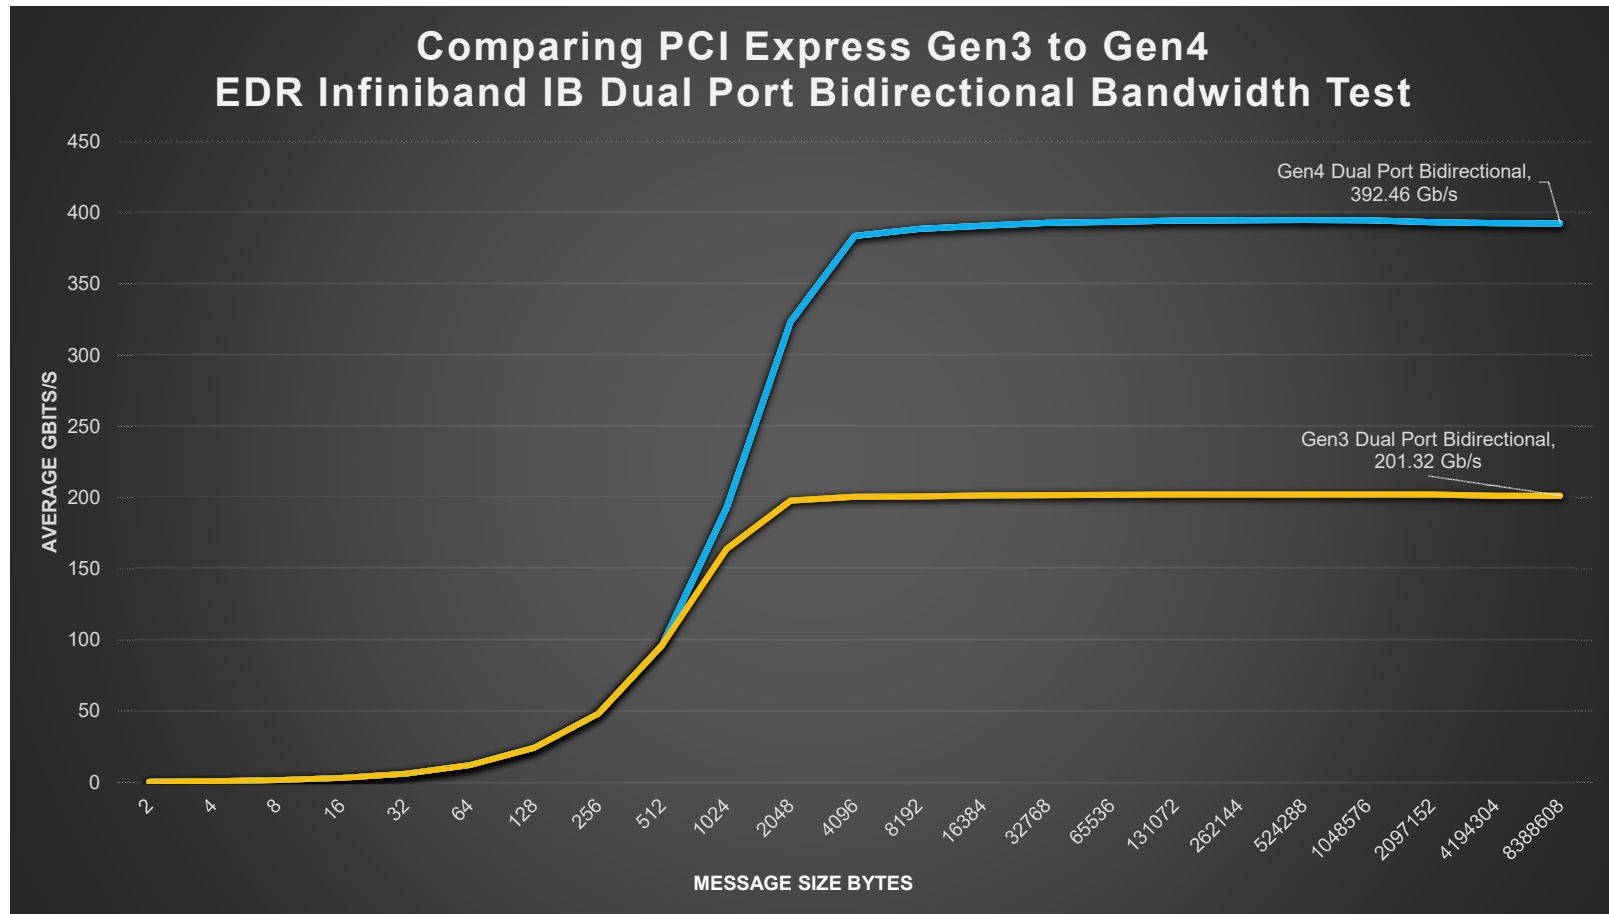
GPU Attach Bandwidth Comparison, PCIe Gen3 verses NVLink



All values are MEASURED Single Direction Bandwidths




InfiniBand EDR 100Gb/s – PCIe Gen 4 verses PCIe Gen 3

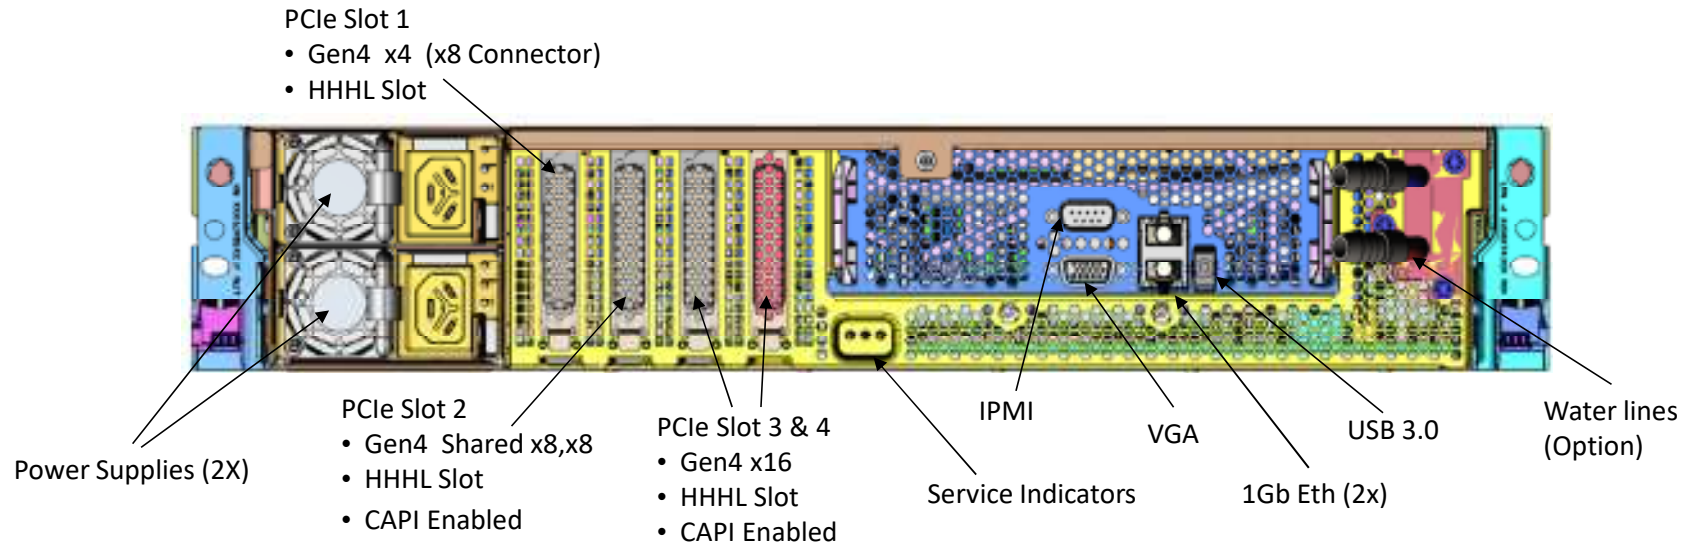


Front and Rear Views

Front




Rear



Rear View – AC Input

Unique AC input connection

- Requires unique AC cord to PDU
- Enables high voltage DC input to power supply
- Enables 480V Input voltage to the rack (future direction)
 - 277VAC distributed to power supplies



POWER AC922 Memory

- 16 direct attach industry standard DDR4 DIMMs are supported in Witherspoon
 - 8 DIMMs per P9
 - Maximum memory capacity increased to 2TB
 - Maximum memory bandwidth of 170GB/s per socket, 340GB/s for the enclosure
 - Measured Stream is 140GB/s per socket
- All memory DIMMs are installed in order to maximize memory bandwidth

DIMM Size	Maximum System Memory	Notes
16GB	256GB	DOE CORAL contract configuration LLNL
32GB	512GB	DOE CORAL contract configuration ORNL
64GB	1024GB	

CORAL

*IBM, Mellanox,
and NVIDIA
awarded \$325M
U.S. Department
of Energy's
CORAL
Supercomputers*


CORAL: Leadership Class Supercomputers

5x – 10x HIGHER APP PERF THAN CURRENT SYSTEMS



June 2018 System Acceptance

POWER AC922 Design – 6 GPU



CORAL Installation at LLNL



CORAL Installation at ORNL



Thank you!



ibm.com/systems/hpc