IBM HPC Technology & Strategy

Hyperion HPC User Forum
Stuttgart, October 1st, 2018

Klaus Gottschalk
gottschalk@de.ibm.com

HPC Strategy

- Deliver End to End Solutions for traditional HPC, HPDA, and AI including SW, HW, Storage, and Services
- Use Emerging Cognitive and AI technologies to provide new Business and Technical Insights as well as enhance traditional HPC computation
- Advance computing frontiers with innovative and disruptive technology such as universal, fault-tolerant quantum computing
- Collaborate with technology leaders, open standards, and open source communities to provide the highest value platforms and ecosystem

HPC and HPDA towards Cognitive, AI and Deep Learning

HPC Simulation

Oil and Gas

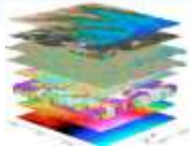

CAE

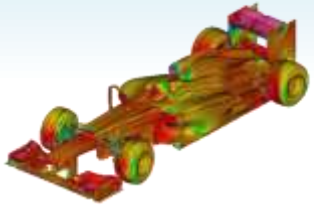

Material Science

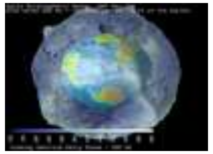

Life Sciences

Cognitive and Deep Learning

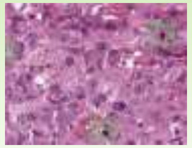

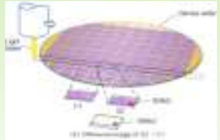

High Performance Data Analysis

Big Data

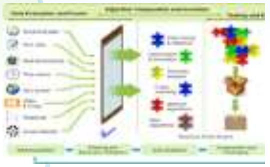

Business Intelligence

Financial Analytics

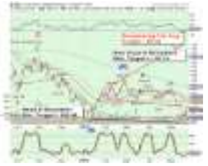

Social Analytics

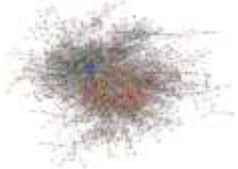

POWER Designed for Cognitive Discovery:

High-performance core, bandwidth, accelerator differentiation

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

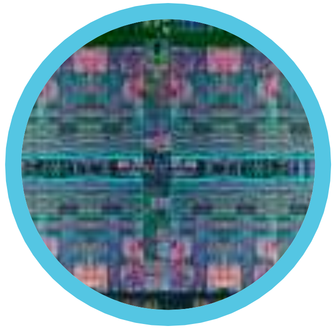

High Performance Cores

Strong cores
4/8 SMT

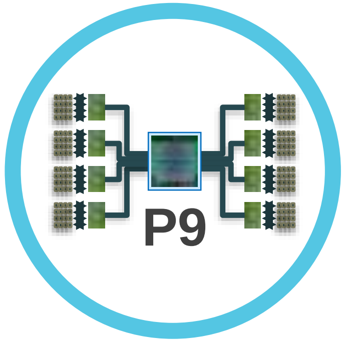

Fast & Large Memory System

1.3X Memory Bandwidth
vs. Intel

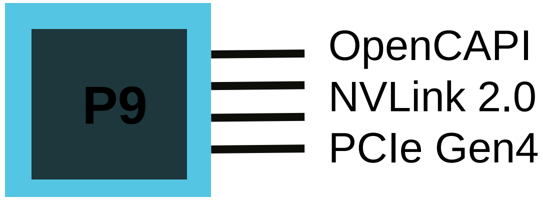

Faster PowerAccel Interconnect for Accelerators

OpenCAPI / NVLink 2.0
PCIe Gen 4

Proliferation of Acceleration with key partners: Nvidia, Xilinx, Mellanox, AMD, open to for many others

IBM Power Systems: open to the core

OpenPOWER

>300 members

Mellanox Technologies
XILINX
NVIDIA

Open Source Workloads

mongoDB
Apache Spark
hadoop
EDB Postgres
cassandra

OpenCAPI

AMD
Hewlett Packard Enterprise
Google
DELL

Open Frameworks

torch
theano
Chainer
TensorFlow
Caffe

Expanding the ecosystem

AC922 POWER9 Accelerated Server

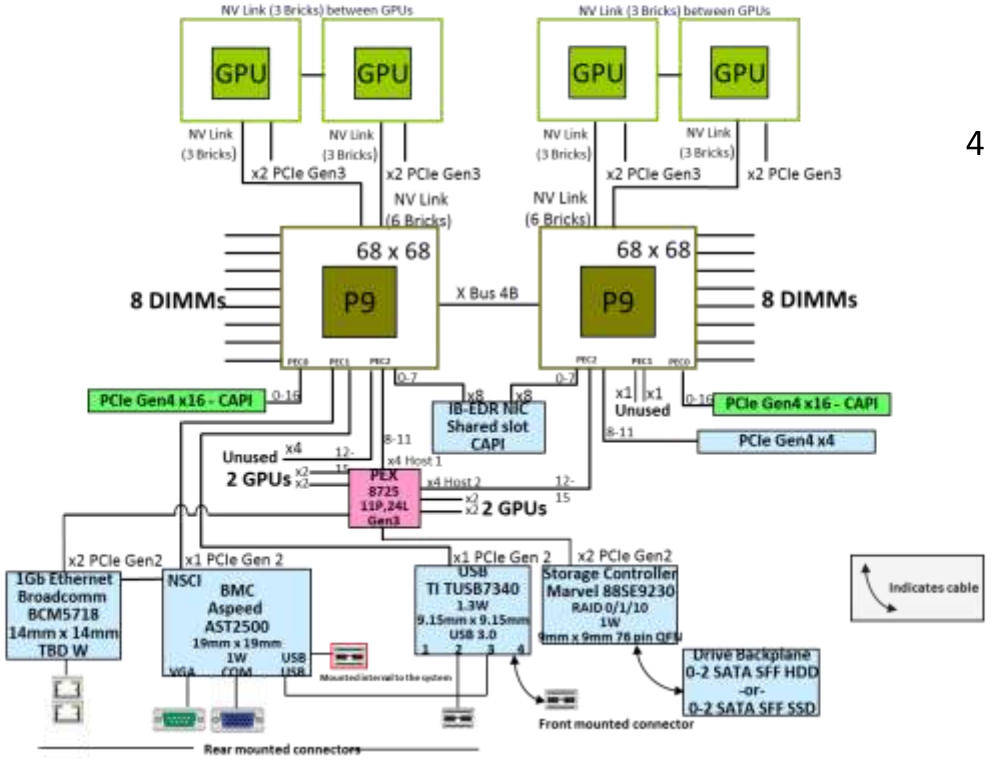

4 GPU

AC922 is the platform that delivers the commitments made in the CORAL contract

- 2 POWER9, 4 GPU for LLNL, water cooled
- 2 POWER9, 6 GPU for ORNL, water cooled

summit

New 200-Petaflops System is Worlds Most Powerful and Worlds Smartest Supercomputer for Science

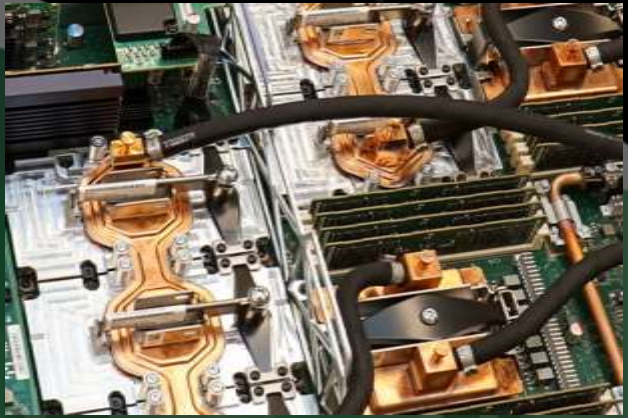

System Overview

System Performance

- Peak performance of 200 petaflops for modeling & simulation
- Peak of 3.3 ExaOps for data analytics and artificial intelligence

Each node has

- 2 IBM POWER9 processor
- 6 NVIDIA Tesla V100 GPUs
- 608 GB of fast memory
- 1.6 TB of NVMe memory

The system includes

- 4608 nodes
- Dual-rail Mellanox EDR InfiniBand network
- 250 PB IBM Spectrum Scale file system transferring data at 2.5 TB/s

IBM Power9 Processor

- Up to 24 cores
  - Summit’s P9s have 22 cores for yield optimization on first processors
- PCI-Express 4.0
  - Twice as fast as PCIe 3.0
- NVLink 2.0
  - Coherent, high-bandwidth links to GPUs
- 14nm FinFET SOI technology
  - 8 billion transistors
- Cache
  - L1I: 32 KiB per core, 8-way set associative
  - L1D: 32KiB per core, 8-way
  - L2: 258 KiB per core
  - L3: 120 MiB eDRAM, 20-way

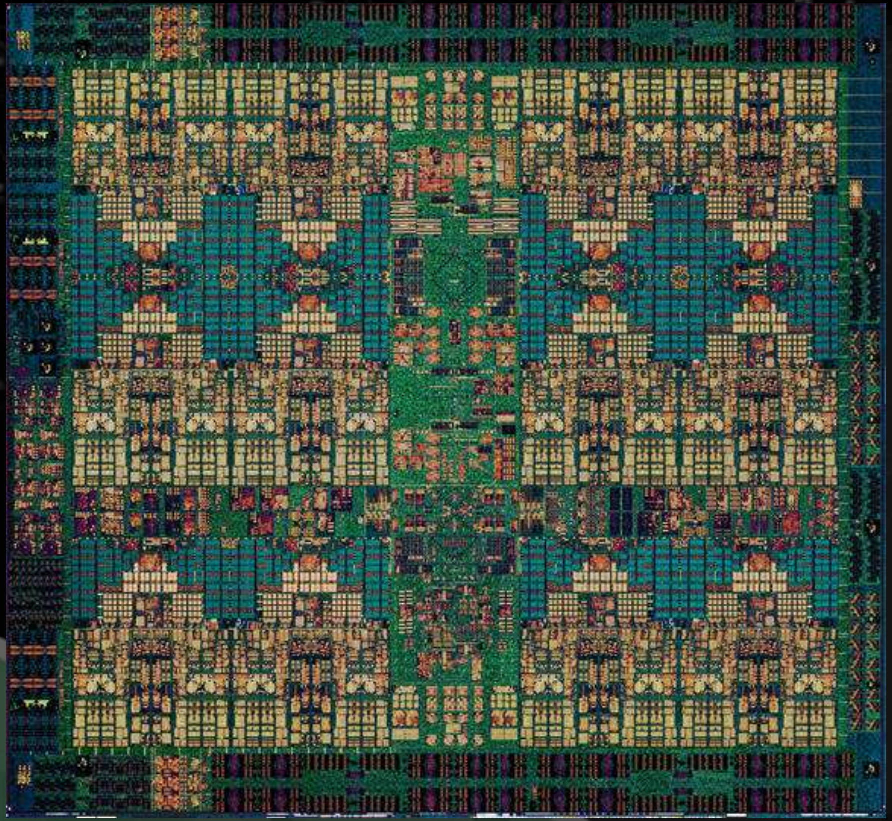

NVIDIA Volta Details

	Tesla V100 for NVLink	Tesla V100 for PCIe
PERFORMANCE with NVIDIA GPU Boost™	DOUBLE-PRECISION 7.8 TeraFLOPS	DOUBLE-PRECISION 7 TeraFLOPS
	SINGLE-PRECISION 15.7 TeraFLOPS	SINGLE-PRECISION 14 TeraFLOPS
	DEEP LEARNING 125 TeraFLOPS	DEEP LEARNING 112 TeraFLOPS
INTERCONNECT BANDWIDTH Bi-Directional	NVLINK 300 GB/s	PCIE 32 GB/s
MEMORY CoWoS Stacked HBM2	CAPACITY 16 GB HBM2	
	BANDWIDTH 900 GB/s	

Features of the Summit system

A **200-petaflop** machine, Summit can perform 200 quadrillion (peta-) floating point operations per second (flops). If every person on Earth completed one calculation per second, it would take **305 days** to do what Summit can do in **1 second**.

Occupying **5,600 sq. ft.** of floor space, Summit could fill **two tennis courts.**

Summit is connected by **185 miles** of fiber optic cables—or the distance from Knoxville to Nashville, Tennessee.

More than **4,000 gallons** of water pump through Summit's cooling system every minute, carrying away about **13 megawatts** of heat.

For some AI applications, researchers can use less precise calculations than flops, potentially **quadrupling** Summit's performance to **exascale levels**, or more than a billion billion calculations per second.

Summit will replace Titan as the OLCF's leadership supercomputer

- Many fewer nodes
- Much more powerful nodes
- Much more memory per node and total system memory
- Faster interconnect
- Much higher bandwidth between CPUs and GPUs
- Much larger and faster file system

Feature	Titan	Summit
Application Performance	Baseline	5-10x Titan
Number of Nodes	18,688	4,608
Node performance	1.4 TF	42 TF
Memory per Node	32 GB DDR3 + 6 GB GDDR5	512 GB DDR4 + 96 GB HBM2
NV memory per Node	0	1600 GB
Total System Memory	710 TB	>10 PB DDR4 + HBM2 + Non-volatile
System Interconnect	Gemini (6.4 GB/s)	Dual Rail EDR-IB (25 GB/s)
Interconnect Topology	3D Torus	Non-blocking Fat Tree
Bi-Section Bandwidth	15.6 TB/s	115.2 TB/s
Processors	1 AMD Opteron™ 1 NVIDIA Kepler™	2 IBM POWER9™ 6 NVIDIA Volta™
File System	32 PB, 1 TB/s, Lustre®	250 PB, 2.5 TB/s, GPFS™
Power Consumption	9 MW	13 MW

Challenges for the world's smartest supercomputer for open science

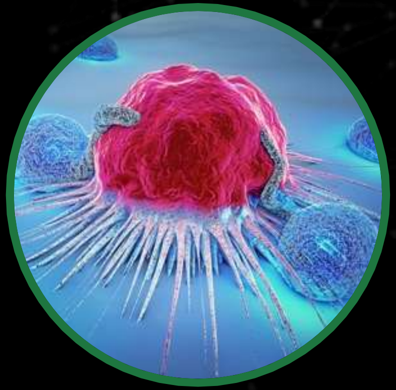

Combating Cancer

Through the development of scalable deep neural networks, scientists at the US Department of Energy and the National Cancer Institute are making strides in improving cancer diagnosis and treatment.

Predicting Fusion Energy

Predictive AI software is already helping scientists anticipate disruptions to the volatile plasmas inside experimental reactors. Summit's arrival allows researchers to take this work to the next level and further integrate AI with fusion technology.

Deciphering High-energy Physics Data

With AI supercomputing, physicists can lean on machines to identify important pieces of information—data that's too massive for any single human to handle and that could change our understanding of the universe.

Identifying Next-generation Materials

By training AI algorithms to predict material properties from experimental data, longstanding questions about material behavior at atomic scales could be answered for better batteries, more resilient building materials, and more efficient semiconductors.

ORNL: Delivering sustained leadership in HPC

Application firsts

- 1988: 1st gigaflops
  - Finite elements on Cray Y-MP
- 1998: 1st teraflops
  - Magnetic materials (LSMS) on Cray T3E
- 2008: 1st petaflops
  - Disorder in superconducting materials (DCA++) on Cray XT5 Jaguar
- 2018: 1st ExaOps
  - Genomics (CoMet) on IBM AC922 Summit

Simulation

QMCPACK

- Accurate quantum mechanics-based simulation of materials, including high-temperature superconductors
- First-of-a-kind studies of properties of complex oxide materials
- Runs on Summit with 25× speed-up on GPUs over CPUs only, and >97% parallel scalability
- 12× faster on Summit than on Titan

Data analytics

CoMet

- Coevolutionary relationships across a population of genomes at scale
- Enables discovery of epistatic relationships
- 1.88 ExaOps achieved on 4,000 Summit nodes in mixed precision
  - 25× faster than same method running on Titan
  - 10,000× faster than best comparable state of the art

Artificial intelligence

MENNDL

- Designs deep learning networks for image-based data sets, going beyond what human experts would design
- Enables a goal of nanotechnology: Atom-by-atom material fabrication
- Scaled to 3,000 available nodes of Summit, achieving 98.6 petaflops (FP32)

QMCPACK: Analyzing materials at the atomic level

Paul Kent
Oak Ridge National Laboratory

Developing the next-generation of materials, including compounds for energy storage, conversion, and production, depends on subatomic understanding of material behavior. QMCPACK, a quantum Monte Carlo application, simulates these interactions using first-principles calculations.

Up to now, researchers have only been able to simulate tens of atoms because of QMCPACK's high computational cost. Summit, however, can support materials composed of hundreds of atoms, a jump that aids the search for a more practical superconductor—a material that can transmit electricity with no energy loss.

"Summit's large, on-node memory is very important for increasing the range of complexity in materials and physical phenomena. Additionally, the much more powerful nodes are really going to help us extend the range of our simulations."

Creating systems biology-wide understanding of humans as an organism

Daniel Jacobson
Oak Ridge National Laboratory

Applying data analytics and artificial intelligence to genetic and biomedical datasets offers the potential to accelerate understanding of bioenergy solutions as well as advance human health and disease outcomes.

Using a mix of AI techniques on Summit, researchers will be able to identify patterns in the function, cooperation, and evolution of human proteins and cellular systems. These patterns can collectively give rise to clinical phenotypes, observable traits of diseases such as Alzheimer's, heart disease, or drug addiction.

Through a collaboration between DOE and the Department of Veterans Affairs, researchers are combining clinical and genomic data with deep learning and Summit's advanced architecture to understand the genetic factors that contribute to conditions such as opioid addiction.

"The complexity of biological system is incredible. Summit is enabling a whole new range of science that was simply not possible before it arrived."

CoMet: First-ever exascale scientific application on Summit

Exploiting new technology

- Modified CCC algorithm uses NVIDIA Volta Tensor Cores and cuBLAS to compute counts of bit values
  - Speed: 4.5× previous highly optimized bitwise algorithm, 25× faster than Titan
  - >10,000× faster than competing codes

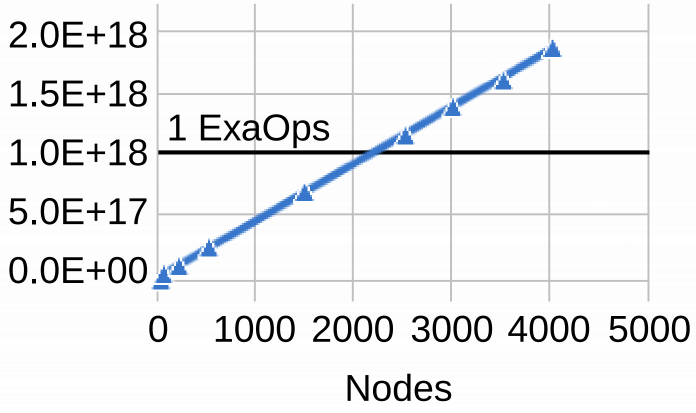

Scaling to Summit

- Near-perfect scaling to 4000 nodes: 1.88 ExaOps achieved
  - W. Joubert et al., *Parallel Comput.*, accepted

New discoveries enabled

- Coevolutionary relationships across a population of genomes at unprecedented scale
- Epistatic interactions for opioid addiction (Gordon Bell Prize submission)

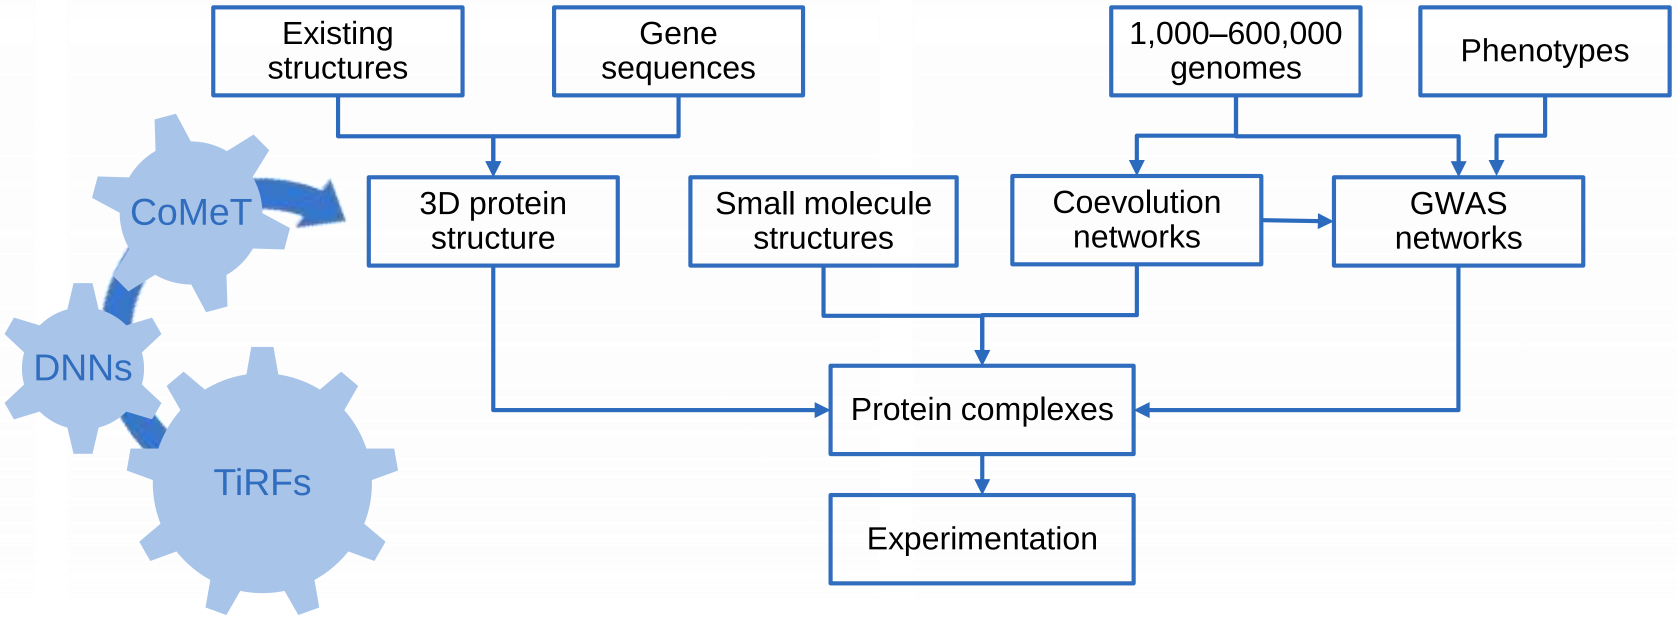

Advancing healthcare with big data

Georgia Tourassi
Oak Ridge National Laboratory

One of the keys to combating cancer is developing tools that can automatically extract, analyze, and sort existing health data to reveal previously hidden relationships between disease factors such as genes, biological markers, and environment. Paired with unstructured data such as text-based reports and medical images, machine learning algorithms scaled on Summit will help supply medical researchers with a comprehensive view of the US cancer population at a level of detail typically obtained only for clinical trial patients.

This cancer surveillance project is part of the CANcer Distributed Learning Environment (CANDLE), a joint initiative between DOE and the National Cancer Institute.

"Essentially, we are training computers to read documents and abstract information using large volumes of data. Summit enables us to explore much more complex models in a time efficient way so we can identify the ones that are most effective."

What makes Summit the most powerful and smartest supercomputer for science?

GPU Brawn: Summit links more than 27,000 deep-learning optimized NVIDIA GPUs with the potential to deliver exascale-level performance (a billion-billion calculations per second) for AI applications.

High-speed Data Movement: High speed Mellanox interconnect and NVLink high-bandwidth technology built into all of Summit's processors supply the next-generation information superhighways.

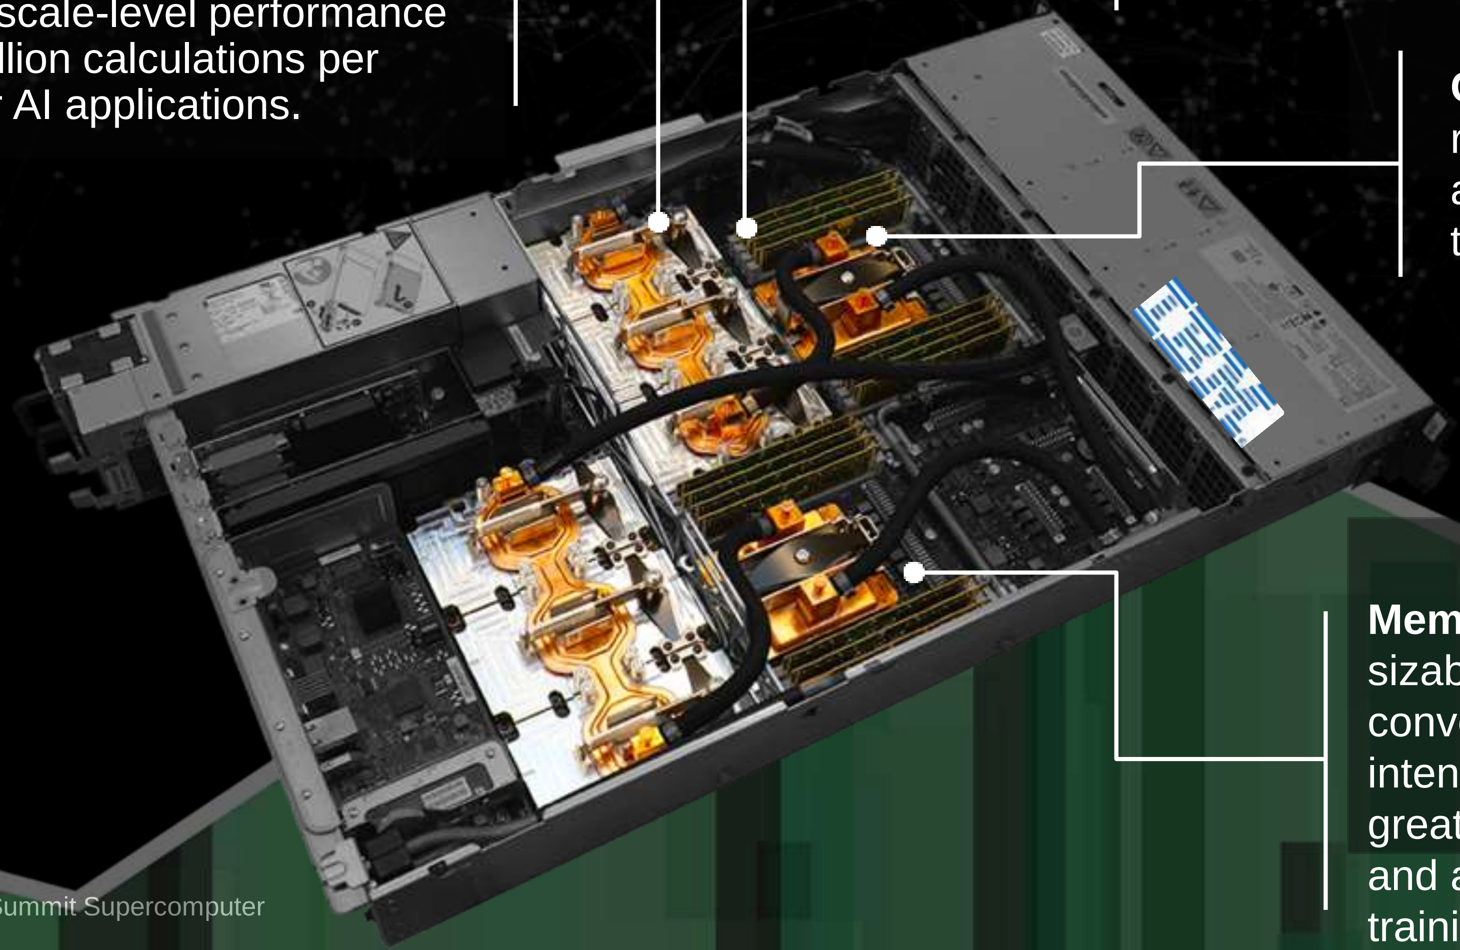

CPU Muscle: IBM Power9 processors to rapidly execute serial code, run storage and I/O services, and manage data so the compute is done in the right place.

Memory Where it Matters: Summit's sizable memory gives researchers a convenient launching point for data-intensive tasks, an asset that allows for greatly improved application performance and algorithmic accuracy as well as AI training.