



Summit System Overview



Tom Papatheodore
Oak Ridge Leadership Computing Facility
Oak Ridge National Laboratory


June 1, 2018

ORNL is managed by UT-Battelle
for the US Department of Energy



Coming in 2018: Summit will replace Titan as the OLCF's leadership supercomputer

Summit, slated to be more powerful than any other existing supercomputer, is the Department of Energy's Oak Ridge National Laboratory's newest supercomputer for open science.




Summit Overview



Compute Node

2 x POWER9
6 x NVIDIA GV100
NVMe-compatible PCIe 1600 GB SSD



25 GB/s EDR IB- (2 ports)
512 GB DRAM- (DDR4)
96 GB HBM- (3D Stacked)
Coherent Shared Memory



NVIDIA GV100

- 7 TF
- 16 GB @ 0.9 TB/s
- NVLink

Compute Rack

18 Compute Servers
Warm water (70°F direct-cooled components)
RDHX for air-cooled components



39.7 TB Memory/rack
55 KW max power/rack

Compute System


256 compute racks
4,608 compute nodes
Mellanox EDR IB fabric
200 PFLOPS
10.2 PB Total Memory
~13 MW



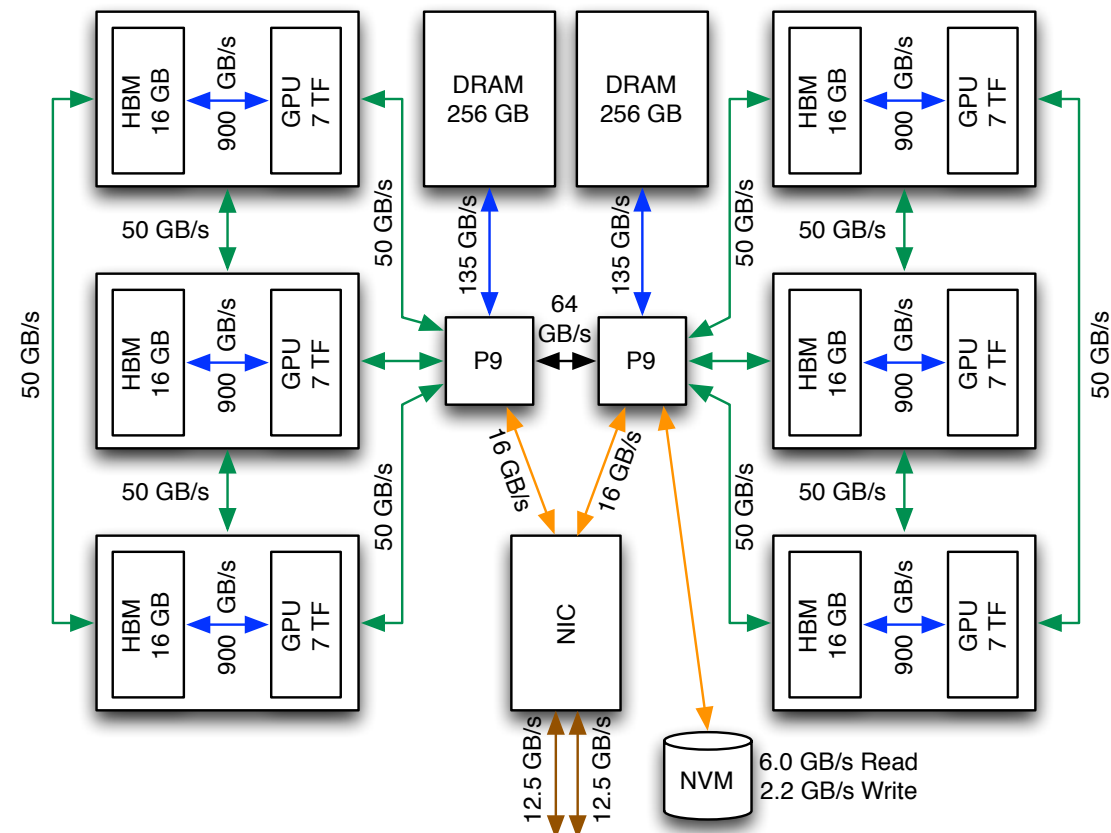

GPFS File System

250 PB storage

2.5 TB/s read, 2.5 TB/s write
(**2.5 TB/s sequential and 2.2 TB/s random I/O)



Summit Node Overview



TF
HBM
DRAM
NET
MMsg/s


42 TF (6x7 TF)
96 GB (6x16 GB)
512 GB (2x16x16 GB)
25 GB/s (2x12.5 GB/s)
83

↔ HBM/DRAM Bus (aggregate B/W)
 ↔ NVLink
 ↔ X-Bus (SMP)
 ↔ PCIe Gen4
 ↔ EDR IB

HBM & DRAM speeds are aggregate (Read+Write).
All other speeds (X-Bus, NVLink, PCIe, IB) are bi-directional.

IBM Power9 Processor

- Summit's P9s: 22 cores (4 hwthreads/core)
- 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: 32 KiB (per core, 8-way)
 - L2: 512 KiB (per pair of cores)
 - L3: 120 MiB eDRAM, 20-way (shared by all cores)



NVIDIA Volta Details

	Tesla V100 for NVLink	Tesla V100 for PCIe
PERFORMANCE with NVIDIA GPU Boost™	DOUBLE-PRECISION	DOUBLE-PRECISION
	7.8 TeraFLOPS	7 TeraFLOPS
	SINGLE-PRECISION	SINGLE-PRECISION
	15.7 TeraFLOPS	14 TeraFLOPS
	DEEP LEARNING	DEEP LEARNING
	125 TeraFLOPS	112 TeraFLOPS
TensorCores™ Mixed Precision (16b Matrix-Multiply-Add and 32b Accumulate)		
INTERCONNECT BANDWIDTH Bi-Directional	NVLINK	PCIe
	300 GB/s	32 GB/s
MEMORY CoWoS Stacked HBM2	CAPACITY	
	16 GB HBM2	
	BANDWIDTH	
	900 GB/s	



Note: The performance numbers are peak and not representative of Summit's Volta

Coming in 2018: 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	112 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

Summit is still under construction

- We expect to accept the machine in Summer of 2018, allow early users on this year, and allocate our first users through the INCITE program in January 2019.





Thank You.