

Dogwood Overview and Technical Specifications

This new cluster is designed to run **large**, multi-node **parallel** jobs. Lenovo is the vendor integrating the system which consists of nine racks of compute, management and storage nodes connected with an InfiniBand switching fabric. The compute nodes are Intel Xeon nodes (details below) and the interconnect is Mellanox EDR InfiniBand, which is a low latency, high bandwidth fabric which designed for sending MPI messages between nodes. The fabric is connected in a fat tree topology. Subject to acceptance testing, we are planning to expand the cluster by another 45 nodes to just over 10,000 cores.

Cluster

- 183 nodes, 8052 total cores (planned: 228 nodes, 10,032 cores)
- 309.2 TF Rpeak
- GPFS file system (IBM Spectrum Scale)

Interconnect

- 9 leaf switches, 6 core switches
- each EDR link is 100 Gb/s bidirectional
- 2:1 blocking ratio Mellanox EDR IB
- Non-blocking within a rack (24 nodes, 1056 cores)

Storage

- 368 TB (raw) scratch space, 265 TB usable
- Non-blocking IB to the storage

Nodes

- Intel Xeon E5-2699Av4
- Broadwell-EP, 14 nm architecture
- 2.4 GHz clock speed
- dual socket, 22 cores per socket, 44 cores per node
- 2x9.6 GT/s QPI
- 16 x 32 GB DDR4, 2400Mhz Memory BW
- 55 MB L3 Cache
- 512 GB Memory/Node
- 145 W TDP

UNC Solution in the nutshell

- 183x compute nodes — x3550 M5
- 2x management nodes
- 2x login nodes with 6TB raw
- 4x login nodes
- 1x DSS G204 — 368TB raw / 265TB usable — 96x 3.84TB SSD SAS drives
- 8052 cores (compute only)
- 309.2 Tflops/s peak (compute only)
- 2:1 blocking ratio Mellanox EDR
 - Non blocking for the storage
 - 9x L1 Mellanox SB7890 leaf switches
 - 6x L2 Mellanox SB7890 core switches

Lenovo™