

ROGER Technical Summary

System Overview

ROGER is a Dell cluster consisting of a total of 108 Intel Xeon E5-2660v3 processors with a total of 13.3 TB of system memory available for computation. Each Xeon E5 v3 processor is capable of performing up to 416 Gflops¹, yielding a peak performance of 44.9 Tflops. In addition to the CPUs are 12 Nvidia Tesla K40M graphics units available for tasks that require GPU's. Each GPU is capable of performing at up to 1.68 Tflops², bringing the total cluster-wide performance to ~65 Tflops (theoretical). The cluster is connected by a high-speed network with 40Gb/s switches in the core and 10Gb/s uplinks to each node.

¹ 2.6GHz * 10 cores * 108 processors * 16 double-precision floating point operations per second with AVX2. ² see [table](#), max double precision.

Compute Nodes

ROGER is comprised of three distinct node types: traditional compute nodes, high memory nodes, and graphics nodes.

Batch Compute Nodes

The compute nodes are Dell Power Edge R730 servers with two Intel Xeon E5-2660 v3 chips. Each chip has 10 cores, each running at 2.6GHz, and a 25MB cache. Each server has 256GB of physical RAM and 500GB of local storage for swap and scratch space. There are 24 batch compute nodes. Their node names are in the format `cg-cmpXX` (where `XX` is the node number).

GPU Nodes

The GPU compute nodes are identical to the traditional compute nodes with the addition of an Nvidia Tesla K40M graphics processing unit. There are 12 GPU compute nodes. Their node names are in the format `cg-gpuXX` (where `XX` is the node number).

High Memory Nodes

The High Memory compute nodes are identical to the traditional compute nodes, except for being equipped with 800GB of local storage using SSDs. They are called "high memory" because before the upgrade of April 2016, they were the only nodes with 256GB of RAM. There are 16 high memory nodes. These nodes were designed to run Hadoop workloads, and can sustain an 11TB hadoop filesystem with their SSDs. However, the production Hadoop system is currently using the GPFS shared filesystem for storage, which still has excellent performance and allows much greater size: currently 175 Tb is allocated. Their node names are in the format `cg-hmXX` (where `XX` is the node number).

Compute Hardware

Node Type	Node Qty	CPU	CPU Qty	Total Cores	RAM	Connectivity	Storage	GPU
Compute ¹	24	Intel Xeon E5-2660 v3 2.6GHz 25M Cache 10 Cores	2	20	256 GB	(1x) 10Gb/s	500GB 7.2K rpm Nearline SAS 6Gb/s	n/a
GPU	12	(same)	2	20	256 GB	(same)	500GB 7.2K rpm Nearline SAS 6Gb/s	Nvidia Tesla K40M
Hadoop/ High Mem ²	16	(same)	2	20	256 GB	(same)	800GB SSD SAS 6Gb/s	n/a
GridFTP	2	(same)	2	20	256 GB	(1x) 40Gb/s	800GB SSD SAS 6Gb/s	n/a

¹Seven of these 24 are currently assigned to OpenStack.

²These were originally the only nodes with 256GB of RAM. They are currently allocated to Hadoop (11/16) and OpenStack (5/16).

The installation of seven more nodes is pending as of May 2016.

Service Nodes

In addition to the computing resources, the cluster also has 10 GPFS filesystem servers, 2 high memory service nodes and an administration node. Of the GPFS server nodes, 8 are identical to the high memory nodes except for local storage, where the single SSD is replaced by 3 600GB 15K rpm SAS drives for OS and filesystem use. The 2 remaining GPFS servers have an additional 6 SSD drives that connect via an internal 12Gb/s SAS connection, allowing extremely fast filesystem metadata access.

The additional 2 high memory service nodes have a 40Gb/s network connection and provide GridFTP services for data transfer into and out of the cluster. These nodes will also run various VM's, such as the user login node(s) and other cluster services.

Filesystem

All cluster nodes have access to the cluster-wide filesystems. The filesystems are built using the General Parallel File System (GPFS) software from IBM and backed by NetApp E2700 storage units. Each E2700 has 180 SATA drives. Total usable disk space is 4.5PB.

Storage Hardware

Node	Node Qty	CPU	CPU Qty	Total Cores	RAM	Disk	Disk Qty
GPFS server	6	Intel Xeon E5-2660 v3 2.6GHz 25M Cache 10 Cores	2	20	256 GB	600GB 15K RPM SAS 6Gbps 2.5in	3
GPFS server + meta-data	2	(same)	2	20	256 GB	(same as above) + 800 GB SSD SAS 12Gb/s	(3) + 6
Netapp E2700	1	n/a	n/a	n/a	n/a	4TB SATA	180
Netapp E2700	5	n/a	n/a	n/a	n/a	6TB SATA	180

OS Management Hardware

Node	Node Qty	CPU	CPU Qty	Total Cores	RAM	Disk	Disk Qty	Disk	Disk Qty
Admin Node	1	Intel Xeon E5-2407 v2 2.40GHz, 10M Cache 4 Cores	1	4	12 GB	1 TB	2 (Raid Mirror)	300 GB	2 (Raid Mirror)