

Systems staff optimizes evolving OSC computing and storage services



*A new system in 2016
“will exceed the peak
performance of all
the center’s existing
systems combined.”*

Doug Johnson

Chief Systems Architect
& Manager, HPC Systems
djohnson@osc.edu | (614) 292-6286

HPC Systems Services

With the April dedication of OSC’s newest cluster, the Ohio Supercomputer Center currently is offering researchers three mid-sized high performance computing (HPC) systems: the HP/Intel Xeon Phi Ruby Cluster, the HP/Intel Xeon Oakley Cluster and the IBM/AMD Opteron Glenn Cluster. OSC also is providing researchers with a storage environment with several petabytes of total capacity across a variety of file systems. With all of that already sitting on the floor at OSC’s nearby data center, many upgrades and installations await OSC clients in 2016.

The Ruby Cluster

“The seven racks of the new Ruby Cluster house 240 nodes and provide a total peak performance of more than 140 teraflops. We look at Ruby as a transitional system that features newer hardware than Oakley and additional capacity. Since the Glenn Cluster will have to be turned off to make room for the 2016 system, the value of the Ruby Cluster only increases. It will provide additional computational capacity while we are physically removing the Glenn Cluster and before the new system is available. The software environment on Ruby will also be a good springboard to the next system.”

Looking Ahead

In 2016, the remaining racks of the Glenn Cluster are expected to make way for a new system that will exceed the peak performance of all the center’s existing systems combined. Due to the significant increase in performance relative to the current systems, OSC will ensure that other facets of the infrastructure can keep pace with the new system.

We’ll be making enhancements to our external network connections. The core router that OSC uses to connect to the Internet through OARnet will be upgraded to support a 40-gigabit-per-second connection to OARnet. We will also upgrade our peer connection with The Ohio State University to the same speed. Both upgrades are designed for an eventual move to 100-Gbps connections to match what’s already available on the OARnet backbone.

The coming year will see upgrades of all the storage at OSC. This includes upgrades and replacements for users’ home directories, project storage and global scratch-file systems. We will increase our total file system capacity to over five petabytes, with aggregate throughput performance of approximately 200 gigabytes per second. Other improvements have been made to our storage environment to respond to the needs of our user community. These included expansions to our backup systems to not only accommodate the additional storage, but to also provide user accessible archive to tape for long-term retention of data. ”



In April 2015, Chancellor John Carey, right, and Statewide Users Group Chair Thomas Beck unveiled the center's newest supercomputing system—the 4,800-core HP/Intel Xeon Phi Ruby Cluster.

High performance computing systems



	Ruby (2014)	Oakley (2012)	Glenn (Phase III, 2009)
Theoretical Peak Performance	96 TF +28.6 TF (GPU) +20 TF (Xeon Phi) ~144 TF	88.6 TF +65.5 TF (GPU) ~154 TF	34 TF +6 TF (GPU) ~40 TF
# of Nodes / Sockets / Cores	240 / 480 / 4800	692 / 1384 / 8304	426 / 856 / 3408
Cores per Node	20 cores/node	12 cores/node	8 cores/node
Local Disks Space per Node	~800 GB in /tmp	~800 GB in /tmp	~400 GB in /tmp
Compute CPU Specs	Intel Xeon E5-2670 v2 CPUs • 2.5 GHZ • 10 cores per processor	Intel Xeon x5650 CPUs • 2.67 GHZ • 6 cores per processor	AMD Opteron 2380 CPUs • 2.5 GHZ • 4 cores per processor
Compute Server Specs	200 HP SL230 40 HP SL250 (for NVIDIA GPU/Intel Xeon Phi)	HP SL390 G7	IBM x3455
# / Kind of GPU/Accelerators	20 NVIDIA Tesla K40 • 1.43 TF Peak double-precision • 2880 CUDA cores • 12 GB memory 20 Xeon Phi 5110p • 1.011 TF Peak • 60 cores • 1.053 GHZ • 8 GB memory	128 NVIDIA M2070 • 515 GF Peak Double Precision • 6 GB memory • 448 CUDA cores	18 NVIDIA Quadro Plex 2200 S4 • Each with Quadro FX 5800 GPUs • 240 CUDA Cores/GPU • 4 GB memory/GPU
# of GPU / Accelerator Nodes	40 total (20 of each type)	64 Nodes (2 GPUs/node)	36 Nodes (2 GPUs/node)
Total Memory	~16 TB	~33 TB	~10 TB
Memory per Node / per Core	64 GB / 3.2 GB	48 GB / 4 GB	24 GB / 3 GB
Interconnect	FDR/EN IB (56 Gbps)	QDR (40 Gbps)	DDR IB (20 Gbs)