

Nyx Cluster Guide

The new CentOS7 Nyx cluster is now available for testing.

Nyx currently has 8 compute nodes na01-08 (288 CPUs) with Gclisi Isilon as the main file system. The login node is 'nyx'.

Each compute host has 932 GB of local /scratch.

Differences in LSF Configuration between Nyx and luna

1. **When specifying RAM for LSF jobs, specify GB of RAM per task (slot) on Nyx (unlike luna, where RAM is specified per job).**
2. Job may have -W (maximum execution Walltime). Please do not use -We on Nyx.
3. There is no /swap on compute nodes. Memory usage is enforced by cgroups so jobs never swap. A job will be terminated if memory usage exceeds its LSF specification.
4. There is no iounits resource on Juno.
5. To check jobs which are DONE or status EXIT, use "bhist -l JobID" or "bhist -n 0 -l JobID". bacct is also available. "bjobs -l JobID" only shows RUNNING and PEND jobs.

Queue

The Nyx cluster uses LSF (Load Sharing Facility) 10.1FP6 from IBM to schedule jobs. The default queue, 'general', includes all Nyx compute nodes.

Job Resource Control Enforcement in LSF with cgroups

LSF 10.1 makes use of Linux control groups (cgroups) to limit the CPU cores and memory that a job can use. The goal is to isolate the jobs from each other and prevent them from consuming all the resources on a machine. All LSF job processes are controlled by the Linux cgroup system. If a job's processes on a host use more memory than it requested, the job will be terminated by the Linux cgroup memory sub-system.

LSF Configuration Notes

Memory (-M or -R "rusage[mem=**]") is a consumable resource. specified as GB per slot/task (-n).

LSF will terminate any job which exceeds its requested memory (-M or -R "rusage[mem=**]").

Job Default Parameters

Queue name: general

Operating System: CentOS 7

Number of slots (-n): 1

Maximum Waltime (job runtime): N/A

Memory (RAM) : 2GB

Job Submission

To submit one slot and 6GB of RAM job to Nyx nodes :

```
bsub -n 1 -R "rusage[mem=6]" -J my_test -e %J.stdout -eo %J.stderr  
my_executable  
bsub -n 1 -R "rusage[mem=6]" -J my_test -o %J.stderrount my_executable
```

The `-o` flag without `-e` by itself will automatically combine the two logs into the single file. The `%J` variable is a special variable in the job submission which is replaced with your job ID.

LSF useful command: [LSF Job Submission](#)