

Juno Cluster Guide

The new Juno cluster is now available. Juno resources are specifically for investigators associated with the Center for Molecular Oncology. If you have an account on luna, you have access to junos.

Juno currently has 2556 CPUs. The login node is 'Juno', and runs CentOS 7 with GPFS as the main file system.

All nodes are running CentOS 7 with access to the new GPFS /juno storage.

All Juno nodes also have access to the Solisi Isilon file systems.

- [Configuration change log](#)
- [Differences in LSF Configuration between Juno and luna](#)
- [Queues](#)
- [Job Resource Control Enforcement in LSF with cgroups](#)
- [LSF Configuration Notes](#)
- [Job Default Parameters](#)
- [Short vs. Long Jobs and Node Availability](#)
- [Job Submission](#)

Configuration change log

New nodes jx01-34 added to the cluster.

New nodes jy01-03 added to the cluster. These nodes don't have NVMe /fscratch. To request node with /fscratch partition: `bsub -n 1 -R fscratch ...`

Slides from November 2019 User Group: [HPC-User-Group-2019-10.pdf](#)

Slides from March 21 2019 User Group [Juno_UG_032019_final.pdf](#)

Slides from Sep 27 2018 User Group meeting (updated Nov 28) [Juno_UG_092018_final.pdf](#)

December 17, 2018: The new queue "control" has been added. Please, check "Queues"

November 28, 2018: The default OS type is CentOS07. Please, check "Job Submission"

Differences in LSF Configuration between Juno and luna

Please see here: <http://mskcchpc.org/display/CLUS/Juno+vs.+Luna>

Queues

The Juno cluster uses LSF (Load Sharing Facility) 10.1FP8 from IBM to schedule jobs. The Juno cluster has two queues: general and control. The default queue, 'general', includes Juno compute nodes.

The control queue doesn't have wall-time limitation and has one node with 144 oversubscribed slots. The control queue should be used only for monitoring or control jobs (the jobs which doesn't use real CPU and memory resources).

To submit the job to the control queue:

```
bsub -n 1 -q control -M 1
```

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=**]").

All jobs should specify Walltime (-W), otherwise a default Walltime of 6 hours will be used.

LSF will terminate any job which exceeds its Walltime.

The maximum Walltime for general queue is 744 hours (31 days).

Job Default Parameters

Queue name: general

Operating System: CentOS 7

Number of slots (-n): 1

Maximum Waltime (job runtime): 6 hours

Memory (RAM) : 2GB

Short vs. Long Jobs and Node Availability

Juno has CMOPI and DEVEL SLAs. When CMOPI/DEVEL jobs are not filling their assigned nodes, 100% of those job slots are available to non-CMOPI jobs with a duration under 2 hours, 75% of slots are available to jobs under 4 hours, and 50% of slots are available to jobs under 31 days.

Nodes assigned to other SLAs are available to non-SLA jobs up to 6 hours.

Job Submission

By default jobs submitted on jun0 only run on CentOS 7 nodes (with GPFS).

To submit a job to CentOS 7 nodes use either of these formats:

```
bsub -n 1 -W 1:00 -R "rusage[mem=2]"
```

To submit a job to nodes with NVMe /fscratch:

```
bsub -n 1 -R "fscratch"  
bsub -n 1 -R "pic"
```