

Queue Structure

On Gadi, users should submit jobs to a specific queue to run jobs on the corresponding type of node. For example, jobs need to run on GPUs have to be submitted to `gpuvolta` queue to get access to nodes with GPUs, while jobs requiring large amounts of memory may use the `hugemem` queue. If your job can run on the nodes in one of the `normal` queues, you should use those queues. The `normal` queues have more nodes available for your jobs, and this allows users and jobs that do require more specialised queues to get fair access to those queues.

Gadi queue structure also has two main levels of priority, express and normal, which is reflected in the queue names. Express queues (`express` and `expressbw`), are designed to support work needs rapid turnaround, but at a higher service unit charge.

Intel Xeon Cascade Lake

express

- Express priority queue for testing, debugging or other jobs need quick turnaround
- 2 x 24-core [Intel Xeon Platinum 8274 \(Cascade Lake\)](#) 3.2 GHz CPUs per node
- 192GB RAM per node
- 2 CPU sockets per node, each with 2 NUMA nodes
 - 12 CPU cores per NUMA node
 - 48 GB local RAM per NUMA node
- 400 GB local SSD disk per node
- Max request of 3200 CPU cores

normal

- Normal priority queue for standard computational intensive jobs
- 2 x 24-core [Intel Xeon Platinum 8274 \(Cascade Lake\)](#) 3.2 GHz CPUs per node
- 192GB RAM per node
- 2 CPU sockets per node, each with 2 NUMA nodes
 - 12 CPU cores per NUMA node
 - 48 GB local RAM per NUMA node
- 400 GB local SSD disk per node
- Max request of 20736 CPU cores, exceptions available on request

copyq

- Normal priority queue for data archive/transfer and other jobs that need network access, 6 nodes total
- 2 x 24-core [Intel Xeon Platinum 8268 \(Cascade Lake\)](#) 2.9 GHz CPUs per node
- 192GB RAM per node
- 2 CPU sockets per node, each with 2 NUMA nodes
 - 12 CPU cores per NUMA node
 - 48 GB local RAM per NUMA node
- 800 GB local SSD disk per node
- External network access (not available on any other nodes in any other queues)
- Access to the tape filesystem `massdata` (not available on any other nodes in any other queues, job needs to explicitly flag PBS with the directive `l storage=massdata/<project_code>` to ensure the access)
- Max request of 1 CPU core

hugemem

- Normal priority queue for jobs that use large amount of RAM, 50 nodes total
- 2 x 24-core [Intel Xeon Platinum 8268 \(Cascade Lake\)](#) 2.9 GHz CPUs per node
- 1.5 TB of [Intel Optane DC Persistent Memory](#) with 384 GB DRAM as a cache per node
- 2 CPU sockets per node, each with 2 NUMA nodes
 - 12 CPU cores per NUMA node
 - 384GB Optane DC Persistent Memory per NUMA node
 - 96GB RAM per NUMA node as cache
- 1.5 TB local SSD disk per node
- Max request of 192 CPU cores, exceptions available on request

gpuvolta

- Normal priority queue, nodes equipped with NVIDIA Volta GPUs, 160 nodes total
- 2 x 24-core [Intel Xeon Platinum 8268 \(Cascade Lake\)](#) 2.9 GHz CPUs per node
- 384 GB RAM per node
- 2 CPU sockets per node, each with 2 NUMA nodes
 - 12 CPU cores per NUMA node
 - 96 GB local RAM per NUMA node
- 4 x Nvidia Tesla Volta V100-SXM2-32GB per node
- 480 GB local SSD disk per node
- Max request of 960 CPU cores (80 GPUs)

Intel Xeon Broadwell (ex-Raijin)

expressbw

- Express priority queue for testing, debugging or other jobs need quick turnaround on the Broadwell nodes

- 2 x 14-core [Intel Xeon E5-2690v4 \(Broadwell\)](#) 2.6GHz CPUs per node
- 128 or 256GB RAM per node
- 2 CPU sockets per node, each with 1 NUMA node
 - 14 CPU cores per numa NODE
 - 64 or 128 GB local RAM per NUMA node
- 400GB local SSD disk
- Max request of 1848 CPU cores

normalbw

- Normal priority queue for standard computational intensive jobs on the Broadwell nodes
- 2 x 14-core [Intel Xeon E5-2690v4 \(Broadwell\)](#) 2.6GHz CPUs per node
- 128 or 256GB RAM per node
- 2 CPU sockets per node, each with 1 NUMA node
 - 14 CPU cores per numa NODE
 - 64 or 128 GB local RAM per NUMA node
- 400GB local SSD disk
- Max request of 10080 CPU cores, exceptions available on request

hugemembw

- Normal priority queue for jobs that use large amount of RAM on the Broadwell nodes, 10 nodes total
- 2 x 14-core [Intel Xeon E5-2690v4 \(Broadwell\)](#) 2.6GHz CPUs per node
- 1TB RAM per node
- 2 CPU sockets per node, each with 1 NUMA node
 - 14 CPU cores per numa NODE
 - 512 GB local RAM per NUMA node
- 400GB local SSD disk
- Minimum memory request is 7 cores and 256 GB memory
- Max request of 140 CPU cores

megamembw

- Normal priority queue for jobs that use large amount of RAM on the Broadwell nodes, 3 nodes total
- 4 x 8-core [Intel Xeon E7-4809v4 \(Broadwell\)](#) 2.1 GHz CPUs per node
- 3TB RAM per node
- 4 CPU sockets per node, each with 1 NUMA node
 - 8 CPU cores per numa NODE
 - 768 GB local RAM per NUMA node
- 800GB local SSD disk
- Minimum memory request is 32 cores and 1.5TB
- Min request of 32 CPU cores

Intel Xeon Skylake (ex-Raijin)

normalsl

- Normal priority queue for standard computational intensive jobs on the Skylake nodes, 192 nodes in total
- 2 x 16-core [Intel Xeon Gold 6130 \(Skylake\)](#) 2.1GHz CPUs per node
- 192GB RAM per node
- 2 CPU sockets per node, each with 1 NUMA node
 - 16 CPU cores per numa NODE
 - 96 GB local RAM per NUMA node
- 400 GB local SSD disk
- Max request of 640 CPU cores, exceptions available on request