

Status of Jobs (jobinfo)

(Last update: Jan 30, 2025.)

Status of submitted jobs can be checked with "jobinfo" command.

- Show Job Information ("jobinfo -c" and "jobinfo")
 - Job Status
 - Reason for Waiting
- Show Queue Usage and Availability Status (jobinfo -s)
- Show Working Directories of Jobs (jobinfo -w)
- Show Memory Usage of Jobs (jobinfo -m)
- Options of jobinfo

Show Job Status ("jobinfo -c" and "jobinfo")

When -c option is added, the latest status of jobs are shown. (You can add -l option as in the previous system case.)

```
$ jobinfo -c
-----
Queue Job ID Name      Status CPUs User/Grp   Elaps Node/(Reason)
-----
H      9999900 job0.csh     Run    16 zzz/---   24:06:10 ccc047
H      9999901 job1.csh     Run    16 zzz/---   24:03:50 ccc003
H      9999902 job2.sh      Run     6 zzz/---   0:00:36 ccc091
H      9999903 job3.sh      Run     6 zzz/---   0:00:36 ccc091
H      9999904 job4.sh      Run     6 zzz/---   0:00:36 ccc090
...
H      9999989 job89.sh    Run     1 zzz/---   0:00:11 ccg013
H      9999990 job90.sh    Run     1 zzz/---   0:00:12 ccg010
-----
```

If -c is not present, slightly old but detailed information (such as number of GPUs (1+1 represents 1 CPU core and 1 GPU) and group name) is also shown. Jobtype is also added to the queue name column ((c): core, (v): vnode, (l): largemem, (g): gpu).

```
$ jobinfo
-----
Queue Job ID Name      Status CPUs User/Grp   Elaps Node/(Reason)
-----
H(c)  9999900 job0.csh     Run    16 zzz/zz9   24:06:10 ccc047
H(c)  9999901 job1.csh     Run    16 zzz/zz9   24:03:50 ccc003
H(c)  9999902 job2.sh      Run     6 zzz/zz9   0:00:36 ccc091
H(c)  9999903 job3.sh      Run     6 zzz/zz9   0:00:36 ccc091
H(c)  9999904 job4.sh      Run     6 zzz/zz9   0:00:36 ccc090
...
H(g)  9999989 job89.sh    Run    1+1 zzz/zz9   0:00:11 ccg013
H(g)  9999990 job90.sh    Run    1+1 zzz/zz9   0:00:12 ccg010
-----
```

Job Status

In the job status column, one of the following values is shown.

- Run : job is running now.
- Queue : job is waiting now.
- Hold : job is held due to the reason like job dependency
- Array : not finished array job
- Exit : job process finished. Termination process is being performed by the system.

Reason for Waiting

When a job is waiting for running, the reason for waiting is shown at the far right column.

- (cpu) : not enough CPU cores
- (gpu) : not enough GPUs
- (cpu/gpu) : not enough resources (CPU and/or GPU etc.)
- (long) : too long; this job won't run until the end of next maintenance
 - If you submit long jobs before the scheduled maintenance and leave those jobs with this status, they will run after the maintenance. (Please be careful about walltime setting.)
- (group) : won't run due to the resource (usually CPU or GPU) limit for the group
- (user) : won't run due to the resource (usually CPU or GPU) limit for you
- (other) : other reasons. The status of jobs immediately after the submission is usually this value.
- (njob) : won't run due to the number of jobs limitation
- (never) : won't run. If you have question about this, please ask us with the ID of corresponding job.
- (error), (close) : there should be something wrong with the system. Please wait for a while. If this state lasts long, please ask us.

Show Queue Usage and Availability Status (jobinfo -s)

"jobinfo -s" will show the number of running, waiting, holding jobs and the maximum number of jobs/CPU-cores/GPUs/etc. Also, queue usage and availability are shown, where "week jobs" and "long jobs" mean jobs within one week and more than one week, respectively.

```
$ jobinfo -s

User/Group Stat:
-----
queue: H          | user(***)      | group(***)
-----
NJob (Run/Queue/Hold/RunLim) | 1/ 0/ 0/- | 1/ 0/ 0/5000
CPUs (Run/Queue/Hold/RunLim) | 4/ 0/ 0/- | 4/ 0/ 0/6400
GPUs (Run/Queue/Hold/RunLim) | 0/ 0/ 0/- | 0/ 0/ 0/48
core (Run/Queue/Hold/RunLim) | 4/ 0/ 0/1200 | 4/ 0/ 0/1200
lmem (Run/Queue/Hold/RunLim) | 0/ 0/ 0/- | 0/ 0/ 0/896
sncore (Run/Queue/Hold/RunLim) | 0/ 0/ 0/- | 0/ 0/ 0/6400
-----

note: "core" limit is for per-core assignment jobs (jobtype=core/gpu*)
note: "lmem" limit is for jobtype=largemem
note: "sncore" limit is for 64 or 128 cores jobs

Queue Status (H):
-----
  job   | free | free | # jobs | requested
  type  | nodes | cores (gpus) | waiting | cores (gpus)
-----
week jobs
-----
1-4 vnodes | 705 | 90240 | 0 | 0
5+ vnodes  | 505 | 64640 | 0 | 0
largemem   | 0   | 0     | 0 | 0
core       | 179 | 23036 | 0 | 0
gpu        | 0   | 0(0)  | 0 | 0(0)
-----
long jobs
-----
1-4 vnodes | 325 | 41600 | 0 | 0
5+ vnodes  | 225 | 28800 | 0 | 0
largemem   | 0   | 0     | 0 | 0
core       | 50  | 6400  | 0 | 0
gpu        | 0   | 0(0)  | 0 | 0(0)
-----

Job Status at 2023-01-29 17:40:12
```

The first part shows the usage status of you and your group. Number of running jobs (NJob), number of using CPU cores (CPUs), number of using GPUs are shown. For each line, number of (running/queue(waiting)/holding/limit) for each resource is available.

In the example output above, this group can use 6,400 CPU cores and 48 GPUs.

The "core (Run/Queue/Hold/RunLim)" line is the usage status and limit for per core jobs (jobtype=core, gpu). In the example above, you can use up to 1,200 cores with jobtype=core (ncpus<64) or gpu (ngpus>0) jobs. The "lmem (Run/Queue/Hold/RunLim)" line is the usage status and limit for jobtype=largemem. In the example above, your group can use up to 896 cores (7 nodes) with jobtype=largemem. The "sncore (Run/Queue/Hold/RunLim)" line is the usage status and limit for single node jobs; total ncpus value is 64 or 128. The limit value matches with CPUs limit when there are enough free nodes. On contrary, when there are not enough free nodes, the limit value becomes small. Please be careful about this value if you are trying to run so many ncpus=64,128 jobs. Groups with fewer resources (CPUs limit is 768) are not influenced by "sncore" limit; the limit always matches with CPUs one.

The second part of the output shows the availability for each jobtypes.

Show Working Directories of Jobs (jobinfo -w)

Working directory of your jobs (\$PBS_O_WORKDIR) can be shown by adding "-w" option. This is useful when you cannot find the job you are looking for by the name or ID.

```
$ jobinfo -w
-----
Queue Job ID Name      Status Workdir
-----
H 9999920 H_12345.sh    Run  /home/users/zzz/gaussian/mol23
H 9999921 H_23456.sh    Run  /home/users/zzz/gaussian/mol74
...
```

This option is not compatible with "-c". If a job has just been submitted, you may have to wait a few minutes before it can be displayed.

Show Memory Usage of Jobs (jobinfo -m)

You can see memory usage of jobs by adding "-m" option. The "Used.Mem/MB" column shows the maximum memory usage of the job (in MB unit). The "Lim.Mem/MB" column shows the available memory amount for that job (memory used by system processes is allocated separately).

```
$ jobinfo -m
-----
Queue Job ID Name      Status CPUs User/Grp Used.Mem/MB Lim.Mem/MB
-----
H(v) 6245590 sample.sh    Run  128 ***/??? 66520 245760
H(v) 6245593 sample.sh    Run  128 ***/??? 78918 245760
```

This option is not compatible with "-c". If a job has just been submitted, you may have to wait a few minutes before it can be displayed. The memory usage information of finished jobs can be checked with ["joblog" command](#).

Options of jobinfo

Option	Meaning
--help	show help message
-c	show the latest information. Not compatible with -m, -w, or -s option.
-s	show queue usage and availability. Not compatible with other options.
-w	show working directory of jobs. Not compatible with "-c" etc.
-m	show memory usage status. Not compatible with "-c" etc.
-l	do nothing. This remains just to keep compatibility.
-L	show full list of computation nodes
-n	show status of nodes
-g	show also group members jobs
-g [group name]	if you belong to multiple groups, you can specify group name with this option.

Option	Meaning
--------	---------

-q [queue name]	specify queue. Usually you don't need to add this.
-a	show all the jobs. User names and job names are hidden.
[job id]	If used in combination with "-c", only specified jobs will be shown.