

# Table of Contents

<b><u>Troubleshooting PBS Jobs</u></b> .....	<b>1</b>
<u>Common Reasons for Being Unable to Submit Jobs</u> .....	1
<u>Common Reasons Why Jobs Won't Start</u> .....	3
<u>Using pdsh_gdb for Debugging Pleiades PBS Jobs</u> .....	7

# Troubleshooting PBS Jobs

## Common Reasons for Being Unable to Submit Jobs

There are several common reasons why you might not be able to successfully submit a job to PBS:

### Resource Request Exceeds Resource Limits

If you get the following message after submitting a PBS job:

```
"qsub: Job exceeds queue resource limits"
```

Reduce your resource request to below the limit or use a different queue.

### AUID or GID not Authorized to Use a Specific Queue

If you get the following message after submitting a PBS job:

```
"qsub: Unauthorized Request"
```

It is possible that you tried submitting to a queue which is accessible only to certain groups or users. You can check the `qstat -fQ` output and see if there is an `acl_groups` or a `acl_users` list. If your group or username is not in the lists, you will have to submit to a different queue.

If your GID has no allocation left, you will also get this error. See the section 'Not Enough or No Allocation Left' below for more information.

### AUID Not Authorized to Use a Specific GID

If you get the following message after submitting a PBS job:

```
"qsub: Bad GID for job execution"
```

It is possible that your AUID has not been added to use allocations under a specific GID. Please consult with the principal investigator of that GID and ask him/her to submit a request to [support@nas.nasa.gov](mailto:support@nas.nasa.gov) to add your AUID under that GID.

## **Queue is Disabled**

If you get the following message after submitting a PBS job:

```
"qsub: Queue is not enabled"
```

You should submit to a different queue which is enabled.

## **Not Enough or No Allocation Left**

An automated script is used to check if a GID is over its allocation everyday. If it does, that GID is removed from PBS access control list and users of that GID will not be able to submit jobs.

Users can check the amount of allocations remaining using the `acct_ytd` command. In addition, if you see in your PBS output file some message regarding your GID allocation usage is near its limit or is already over, ask your PI to request for more allocation.

Once the request for more allocation is approved and added to your account, an automatic hourly script will add your GID back to the PBS access control list.

# Common Reasons Why Jobs Won't Start

If your job does not run after it has been successfully submitted, it might be due to one of the following reasons:

- The job is waiting for resources
- Your mission share has run out
- The system is going into dedicated time
- Scheduling is turned off
- The job has been placed on hold
- Your home filesystem or default /nobackup filesystem is down

You can often find out why a job does not start by running `tracejob jobid` on the PBS server `pbsp11`, which prints log messages for a PBS job, or `qstat -as jobid` on any Pleiades front-end systems, which displays all information about a job.

The following sections provide more details about each scenario, along with troubleshooting steps.

## The Job is Waiting for Resources

Your job might be waiting for resources for one of the following reasons:

- All resources are busy running jobs, and no other jobs can be started until resources become available again
- There is a higher-priority job that needs more resources than are currently available, so PBS is draining the system and not running any new jobs in order to accommodate the high-priority job
- Users have submitted too many jobs at once (for example, more than 100), so the PBS scheduler is busy sorting jobs and cannot start new jobs effectively
- If you requested a specific node or group of nodes to run your job, it might wait in the queue longer than if nodes were not specified

To view job status and events, run the `tracejob` utility on `pbsp11`. For example:

```
pbsp11 $ tracejob 234567
Job: 234567.pbsp11.nas.nasa.gov

06/15/2012 00:23:55 S Job Modified at request of Scheduler@pbsp11.nas.nasa.gov
06/15/2012 00:23:55 L No available resources on nodes
06/15/2012 00:38:47 S Job Modified at request of Scheduler@pbsp11.nas.nasa.gov
06/15/2012 00:50:30 S Job Modified at request of Scheduler@pbsp11.nas.nasa.gov
06/15/2012 01:51:21 S Job Modified at request of Scheduler@pbsp11.nas.nasa.gov
06/15/2012 01:55:38 S Job Modified at request of Scheduler@pbsp11.nas.nasa.gov
06/15/2012 02:16:44 S Job Modified at request of Scheduler@pbsp11.nas.nasa.gov
06/15/2012 07:39:48 L Considering job to run
06/15/2012 07:39:48 L Job is requesting an exclusive node and node is in use
```

**TIP:** If the scheduler has not yet reviewed the job, no information will be available and the `tracejob` utility will not provide any output.

If your job requests an exclusive node and that node is in use, you can wait for the requested node, request a different node, or submit the job again without requesting a specific node.

To view the current node usage for each processor type, run the `node_states.sh` script. For example:

```
/u/scicon/tools/bin/node_stats.sh
```

```
Pleiades Nodes Total: 12457
Pleiades Nodes Used : 11675
Pleiades Nodes Free : 782

Harpertown   Total: 4083, Used: 3948, Free: 135
Nehalem      Total: 1279, Used: 1246, Free: 33
Westmere     Total: 4662, Used: 4436, Free: 226
SandyBridge  Total: 1723, Used: 1513, Free: 210
GPU nodes    Total: 62, Used: 2, Free: 60
Devel wes    Total: 504, Used: 418, Free: 86
Devel san    Total: 144, Used: 112, Free: 32
```

```
Currently queued jobs are requesting: 1734 Harpertown, 1502 Nehalem, 5219 Westmere,
```

**TIP:** Add `/u/scicon/tools/bin` to your path in `.cshrc` or other shell start-up scripts to avoid having to type in the complete path for this tool.

## Your Mission Share Has Run Out

If all of the cores within your mission directorate's share have been assigned, or if the new job would exceed your mission share, your job will not run. If resources appear to be available, they belong to other missions.

To view all information about your job, run `qstat -as jobid`. In the following sample output, a comment indicates that the job would exceed the mission share:

```
Qstat as 778574
JobID          User      Queue   Jobname  CPUs  wallt  Ss  wallt  Eff  wallt
-----
778574.pbsp13  msmith3  normal  my_GC    12    04:00  Q   07:06  --   04:00
Job would exceed mission CPU share
```

To view the share distribution among all mission directorates, run `qstat -W shares`. For example:

```
pfe20 $ qstat -W shares
Group  Share% Use%  Share Exempt  Use  Avail Borrowed Ratio Waiting
-----
```

Overall	100	0	167456	0	0	167456	0	0.00	0
ARMD	25	20	41864	0	34480	7384	0	0.82	11744
HEOMD	23	17	38514	0	29632	8882	0	0.77	8304
SMD	50	45	83728	2912	76728	7000	0	0.92	129376
NAS	2	0	3349	0	384	2965	0	0.11	32

If your job exceeds your mission share, you might be able to borrow nodes from other mission directorates. To borrow nodes, your job must not request a wall-clock time that is too long (more than 4 hours on Pleiades). See [Mission Shares Policy on Pleiades](#) for more information.

## The System is Going into Dedicated Time

When dedicated time (DED) is scheduled for hardware and/or software work, the PBS scheduler will not start a job if its projected end-time runs past the beginning of the DED time.

If you can reduce the requested wall-clock time so that your job will finish running prior to DED time, then your job can then be considered for running. To change the wall-clock time request for your job, follow the example below:

```
%qalter -l walltime=hh:mm:ss jobid
```

To find out whether the system is in dedicated time, run the command `schedule all`. For example:

```
pfe25$ schedule all
```

```
No scheduled downtime for the specified period.
```

## Scheduling is Turned Off

Sometimes job scheduling is turned off by NAS Control Room staff or a system administrator. This is usually done when there are system or PBS issues that need to be resolved before jobs can be scheduled to run. When this happens, you should see the following message near the beginning of the `qstat -au your_userid` output.

```
+++Scheduling turned off.
```

## Your Job Has Been Placed On Hold ("H" Mode)

A job can be placed on hold either by the job owner or by someone who has root privilege, such as a system administrator. If your job has been placed on hold by a system administrator, you should get an email explaining the reason for the hold.

## Your Home Filesystem or Default /nobackup Filesystem is Down

When a PBS job starts, the PBS prologue checks to determine whether your home filesystem and default /nobackup are available before executing the commands in your script. If your default /nobackup filesystem is down, PBS cannot run your job and will put the job back in the queue. If your PBS job does not need any file in that filesystem, you can tell PBS that your job will not use the default /nobackup to allow your job to run.

For example, if your default is /nobackupp1 (for Pleiades), you can add the following in your PBS script:

```
#PBS -l /nobackupp1=0
```

## Using `pdsh_gdb` for Debugging Pleiades PBS Jobs

A script called `pdsh_gdb`, created by NAS staff Steve Heistand, is available on Pleiades under `/u/scicon/tools/bin` for debugging PBS jobs *while the job is running*.

Launching this script from a Pleiades front-end node allows one to connect to each compute node of a PBS job and create a stack trace of each process. The aggregated stack trace from each process will be written to a user specified directory (by default, it is written to `~/tmp`).

Here is an example of how to use this script:

```
pfel% mkdir tmp
pfel% /u/scicon/tools/bin/pdsh_gdb -j jobid -d tmp -s -u nas_username
```

More usage information can be found by launching `pdsh_gdb` without any option:

```
pfel% /u/scicon/tools/bin/pdsh_gdb
```