

# Table of Contents

<b>Troubleshooting</b> .....	<b>1</b>
<u>Common Login Failures or Issues</u> .....	1
<u>Post Login Issues</u> .....	6
<u>Common Reasons for Being Unable to Submit Jobs</u> .....	8
<u>Common Reasons Why Jobs Won't Start</u> .....	10
<u>Using pdsh_gdb for Debugging Pleiades PBS Jobs</u> .....	14

# Troubleshooting

## Common Login Failures or Issues

### SSH Known-Hosts Error

SSH offers the ability to verify the identity of the remote host to which you are connecting. A successful host verification indicates that your SSH client has established a secure connection with the SSH server and no intermediate machines have access to that connection.

The identity of the remote host can be verified by checking the host public key of the remote host stored either in the system-wide `/etc/ssh/ssh_known_hosts` file (or `/etc/ssh/ssh_known_hosts` for some systems) or your personal `~/.ssh/known_hosts` file on your localhost.

SSH has three ways it can react to an unrecognized or changed SSH host key, based on the value of the `StrictHostKeyChecking` variable in either the system-wide `/etc/ssh/ssh_config` file (or `/etc/ssh/ssh_config` for some systems) or your personal `~/.ssh/config` file:

`StrictHostKeyChecking=no`

This is the most insecure setting as it will blindly connect to the server. It will add the server's key if it's not present locally, and if the key has changed it will add the key without asking.

`StrictHostKeyChecking=ask`

With this setting, if you have no host key for the server, it will show you the fingerprint and ask you to confirm. If you connect and the key does not match, it will prevent you from logging in, and will tell you where to find the conflicting key inside the `known_hosts` file.

`StrictHostKeyChecking=yes`

This is the most secure setting. If you have no host key for this server, then it will prevent you from logging in at all.

If `StrictHostKeyChecking` is set to `yes` and you get errors similar to the following the first time you log in to one of the SFEs, the simple solution is to add `-o "stricthostkeychecking=ask"` in your ssh command. `sfe1` is used in this example.

### Sample Error

```
your_localhost% ssh sfel.nas.nasa.gov
```

```
WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!  
IT IS POSSIBLE THAT SOMEONE IS DOING SOMETHING NASTY!  
Someone could be eavesdropping on you right now (man-in-the-middle attack)!  
It is also possible that the RSA host key has just been changed.  
The fingerprint for the RSA key sent by the remote host is  
11:9f:ae:09:56:2d:45:66:8e:9a:df:15:52:d6:88:5e.  
Please contact your system administrator.  
Add correct host key in /Users/userid/.ssh/known_hosts2 to get rid  
of this message.  
Offending key in /etc/ssh_known_hosts2:24  
RSA host key for sfel.nas.nasa.gov has changed and you have  
requested strict checking.  
No RSA host key is known for sfel.nas.nasa.gov and you have  
requested strict checking.  
Host key verification failed.
```

## Solution

```
your_localhost% ssh -o "stricthostkeychecking=ask" sfel.nas.nasa.gov
```

```
-----  
The authenticity of host 'sfel.nas.nasa.gov (198.9.4.3) '  
can't be established.  
RSA key fingerprint is  
11:9f:ae:09:56:2d:45:66:8e:9a:df:15:52:d6:88:5e.  
Are you sure you want to continue connecting (yes/no)? yes
```

```
-----  
Warning: Permanently added 'sfel.nas.nasa.gov,198.9.4.3' (RSA) to  
the list of known hosts.
```

```
-----  
Plugin authentication  
Enter PASSCODE: type your passcode
```

## Common SecurID Passcode Problems

If you cannot log in after you created your new PIN, it is possible that you did not successfully complete the "New-Pin Process". You might have used a special character in your PIN. Your PIN must consist of EXACTLY 8 alphanumeric characters, with at least one letter and at least one number, but no special characters.

```
your_localhost% ssh sfel  
user@sfel's password:  
Authenticated with partial success.  
Plugin authentication  
Enter PASSCODE:  
Plugin authentication  
You may create your own PIN or accept a server assigned PIN.  
Would you like to create your own new PIN (yes/no)? yes
```

```
Plugin authentication
Enter a new PIN of 8 alphanumeric characters:
Re-enter new PIN to confirm:
Enter PASSCODE:
Plugin authentication
Enter PASSCODE:
Permission denied, please try again.
user@sfel's password:
Permission denied, please try again.
user@sfel's password:
Permission denied ().
```

Each SecurID 6-digit tokencode can be used only once. If you try to use a tokencode that has just been used, you will be prompted again to enter a new passcode.

```
your_localhost% ssh sfel.nas.nasa.gov
PAM Authentication
Enter PASSCODE: enter PIN and a tokencode which was just used
PAM Authentication
Enter PASSCODE: ener PIN and a new token code
```

If you failed to provide a correct PIN + tokencode in a few consecutive attempts, your SecurID fob will be temporarily disabled. NAS Help Desk can help you unlock your fob (even for fobs provided by your local center). Call NAS Help Desk at (800) 331-USER (8737) or (650) 604-4444 for assistance.

## Common Passthrough Problems

If you set up SSH passthrough correctly, you should be prompted for your SecurID passcode only, and be transferred to the desired host. The following describes a few common problems and their solutions:

### Authentication Failure

If the authentication failed when you entered the passcode, it is possible that you have different usernames between your localhost and the NAS systems and you did not include your NAS username in your `.ssh/config` file.

There are two ways to include your NAS username in the `.ssh/config` file:

- If you use the `.ssh/config` file for connecting to multiple computer sites, then you should add your NAS username to the ProxyCommand lines for corresponding NAS hosts, for example:

```
Host lou1
ProxyCommand ssh username@sfel.nas.nasa.gov /usr/local/bin/ssh-proxy lou1
```

In this case, you will still need to issue your **ssh** command as the following example in order to avoid being prompted for a password:

```
%ssh nas_username@lou1
```

- If you use the **.ssh/config** file only for connecting to NAS, then you can simply add the following at the beginning of your **.ssh/config** file:

```
User nas_username
```

You do not need to add your NAS username to the ProxyCommand line. In addition, you can simply use the following command and won't be prompted for a password:

```
%ssh lou1
```

## Prompted for Password

If you are prompted for password in addition to the passcode, there are multiple possible causes:

- You may have a different username and need to use:

```
your_localhost% %ssh nas_username@lou1
```

- Either your home directory or the **.ssh/authorized\_keys** file under your NAS account have write permission for group or others. You need to correct the permission so that they have write permission only to you.
- Either one or both of the following files is missing:
  - ◆ **.ssh2/authorization** of sfe[1,2]
  - ◆ **.ssh/authorized\_keys** of the NAS HECC host that you want to connect to

## Prompted for Passphrase

If you are prompted for passphrase in addition to the passcode, likely, you did not use the commands **ssh-agent** and **ssh-add ~/.ssh/id\_rsa** to forward your private key before you issue the **ssh** command.

## Incorrect Ciphers

A cipher is an algorithm for performing encryption or decryption. The SSH client and server must have a matching cipher in order to successfully verify the keys. If you get an error

similar to the following:

```
no matching cipher found: client blowfish-cbc  
server aes128-cbc,aes192-cbc,aes256-cbc,3des-cbc
```

Check your `./ssh/config` file or `/etc/ssh/ssh_config` file on your localhost and add the appropriate ciphers.

To modify `/etc/ssh_config`, system administrator's privilege may be needed.

## Password Expiration

Your NAS password is valid for 60 days. An email is sent to you by NAS reminding you to change your password. If it has expired, you can still log in. However, you will be prompted to change it right away.

## Account Expiration or Deactivation

Your NAS account is active if you have a valid project and a valid account request form on file at NAS. If your account has expired, it will be removed from NAS database and the `/etc/passwd` files. When this happens, no login will be allowed.

Please note that the account request form has to be filled out once every year. When the form has expired, you should receive an email from the NAS account administrator asking you to fill out a new one.

If you violate NAS security rules such as those listed in the [Acceptable Use Statement](#), your account can be deactivated.

# Post Login Issues

## Idle SSH Connections are Broken After 10-15 Minutes

If you find that idle SSH connections tend to be dropped after 10 to 15 minutes and you are using a router or other device that implements Network Address Translation (NAT), such as a DSL router, you may want to enable the `ServerAliveInterval` parameter in your `~/.ssh/config`.

The issue is that some devices that implement NAT will drop information from their state tables if a TCP connection has been idle for just a few minutes. In such cases, it may be necessary to ensure that the connection is never truly idle for an extended period.

The `ServerAliveInterval` parameter (supported in OpenSSH 3.8 and later) causes the SSH client to periodically send an encrypted query to determine if the server is still responsive. The packets will make the connection appear to be active and can prevent the NAT state table entry for the connection from timing out.

For example, adding the following to your `~/.ssh/config` will cause an encrypted query to be sent every five minutes:

```
ServerAliveInterval 5m
```

**TIP:** You may need to experiment with the length of the interval to determine an appropriate value to prevent your connection from being broken. If you use the NAS config template, `ServerAliveInterval` is set to 10 minutes.

## SSH Connection is Periodically Broken Even When Connection is Active

During an SSH session, the client and server will re-negotiate the keys used to encrypt the session every so often. Very old versions of OpenSSH did not implement this feature, which can cause a problem when connected to the SFEs. Your connection may appear to hang or may abort with an error message such as the following:

```
Dispatch protocol type 20
```

If you encounter this issue, you will need to upgrade to a newer version of OpenSSH. Note: Other implementations that were derived from the OpenSSH distribution, such as SUN Secure Shell, are likewise known to exhibit this problem.

## Slow Performance Transferring Data

While OpenSSH performs reasonably for local data transfers, the performance will tend to be reduced due to the latency of long-haul network connections.

Depending on the severity of the impact you may have several options to improve this situation.

## **Upgrade to OpenSSH 4.7 or Later**

If you are using a version of OpenSSH that is older than 4.7, you may see an improvement in file transfer performance by upgrading to OpenSSH 4.7 or later. This is due to the use of a larger channel buffer introduced in that version.

## **Use an HPN-enabled Version of OpenSSH**

The High-Performance Networking (HPN) patch set maintained by the Pittsburgh Supercomputing Center (PSC) allows the channel buffer used by OpenSSH to grow as needed. This may be useful in cases where OpenSSH 4.7 does not yield satisfactory performance.

In cases where you are seeing very poor performance (under 1.2 MB/s) and the data you are transferring will compress well, you may wish to enable compression.

You can do this by adding `-C` to your `scp` or `sftp` command-line.

Read [Tips for File Transfers](#) for more recommendations.

# 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
```