Contents

Introduction

Below is a list of common problems users have experienced.

PBS Problems

Below is a list of common errors and how to fix them.

- You forgot to include the *#PBS -q*

```
user@hpc-login-prd-t1 ~]$ qsub syscalculation.pbs
qsub: No default queue specified MSG=cannot locate queue
```

- You specified an invalid queue.

```
user@hpc-login-prd-t1 ~]$ qsub syscalculation.pbs
qsub: Unknown queue
```

- You do not have access to the queue specified

```
user@hpc-login-prd-t1 ~]$ qsub syscalculation.pbs
qsub: Access to queue is denied
```

- PBS is unable to lookup your account

```
user@hpc-login-prd-t1 ~]$ qsub syscalculation.pbs
qsub: Bad UID for job execution
```

- PBS is unable to transfer information around cluster using ssh.
  To fix, from hpc-login-prd-t1 ssh to habeus then ssh back to hpc-login-prd-t1, accept ssh keys when asked, then logout of both hpc-login-prd-t1 and habeus.

```
Host key verification failed.
lost connection
>>> end error output
Output retained on that host
```

- When using modules in a PBS script you get *Unable to locate a modulefile*

```
user@hpc-login-prd-t1 ~]$ qsub syscalculation.pbs  OR
user@hpc-login-prd-t1 ~]$ module load matlab/r2017b
matlab(3):ERROR:105: Unable to locate a modulefile for 'matlab/r2017b'

Or

The shell being used in the PBS script, e.g. "/bin/sh", is not being
read as a login shell therefore it's not reading environmental
variables. To fix change "/bin/sh" to "/bin/sh -l".
```

HPC Job FAQ

Below is a list of common question about jobs.

- To find how much memory a job is using

```
[user1@hpc-login-prd-t1 ~]# qstat -a
Job id                    Name             User             Time Use S
Queue
------------------------- ---------------- ---------------- -------- -
-----
10720.hpc-admin           test4            user1            07:45:44 R
standard
```

from the output displayed find your job, then

```
[user1@hpc-login-prd-t1 ~]# jobmemusage 10720
The job 10720 has requested  10gb of memory
however
it is using  1301352kb of memory
and          2012660kb of virtual memory.
```

- To check if a job is running on a node

```
[user1@hpc-login-prd-t1 ~]# qstat -an
Job ID               Username Queue    Jobname          SessID NDS
TSK Memory Time  S Time
-------------------- -------- -------- ---------------- ------ -----
--- ------ ----- - -----
10720.hpc-admin      user1    standard test4             26241     1
--   10gb 72:00 R 02:12
   n007/3+n007/2+n007/1+n007/0
```

from the output displayed find your job, then

```
[user1@hpc-login-prd-t1 ~]# nodecheck n007
************************* hpc_cluster *************************
--------- n007---------
26397 user1  25   0  387m 311m  17m R 100.8  1.9 122:49.84 solver-
pvm.exe
26453 user1  21   0  350m 285m  15m R 100.0  1.8 115:43.65 solver-
pvm.exe
26452 user1  25   0  369m 305m  15m R 98.9  1.9  131:28.18 solver-
pvm.exe
26454 user1  25   0  354m 289m  15m R 98.9  1.8  128:50.67 solver-
pvm.exe
26285 user1  21   0 66832 4428 1100 S  0.0  0.0    0:00.15
10720.usqhpc00.
26331 user1  15   0  131m  49m 2252 S  0.0  0.3    0:02.53 perl
26389 user1  16   0  6188 1088  708 S  0.0  0.0    0:45.51 pvmd3
26398 user1  18   0 79448 7000 1544 S  0.0  0.0    0:00.06 perl
26399 user1  18   0 79448 7000 1544 S  0.0  0.0    0:00.06 perl
26400 user1  18   0 79448 7004 1544 S  0.0  0.0    0:00.06 perl
```

This is particular useful when you have requested a number of nodes and you need to check if node is being used. If you have no job running on a node nothing will be displayed, as below:

```
[user1@hpc-login-prd-t1 ~]# nodecheck n007
************************* hpc_cluster *************************
--------- n007---------
```

Copying data between clusters (deprecated)

Copy a file or directory from the old HPC to a particular directory on the new HPC;

An Ethernet cross-cable link has been established from the old to new HPC. Instruction to copy files from the old hpc is as below (first login to the 'usqhpcio' machine):

[userid@usqhpcio ~]$ scp -r $HOME/* userid@hpc-dmf-prd-t1:~

Note that the destination hostname "**hpc-dmf-prd-t1"** will take data to the new HPC's file server and respective home folder.
* Remix this command as required.

Copy a file or directory from the new HPC to a particular directory on the old HPC;

Yes. The other side's hostname is "hpc-dmf-prd-t1". Login to this machine ("ssh hpc-dmf-prd-t1" from hpc-login-prd-t1) and execute

[userid@hpc-dmf-prd-t1 ~]$ scp -r $HOME/* userid@usqhpcio:~

* Remix this command as required. You can replace ($HOME/*) with your specific file or directory and (~) means home directory at the destination. For more information "man scp"

Is there any ways to use WinSCP to directly take the data results from the new HPC to my computer?

Yes, you use "hpc-dmf-prd-t1" and port 22 and select sftp – ssh file transfer protocol.

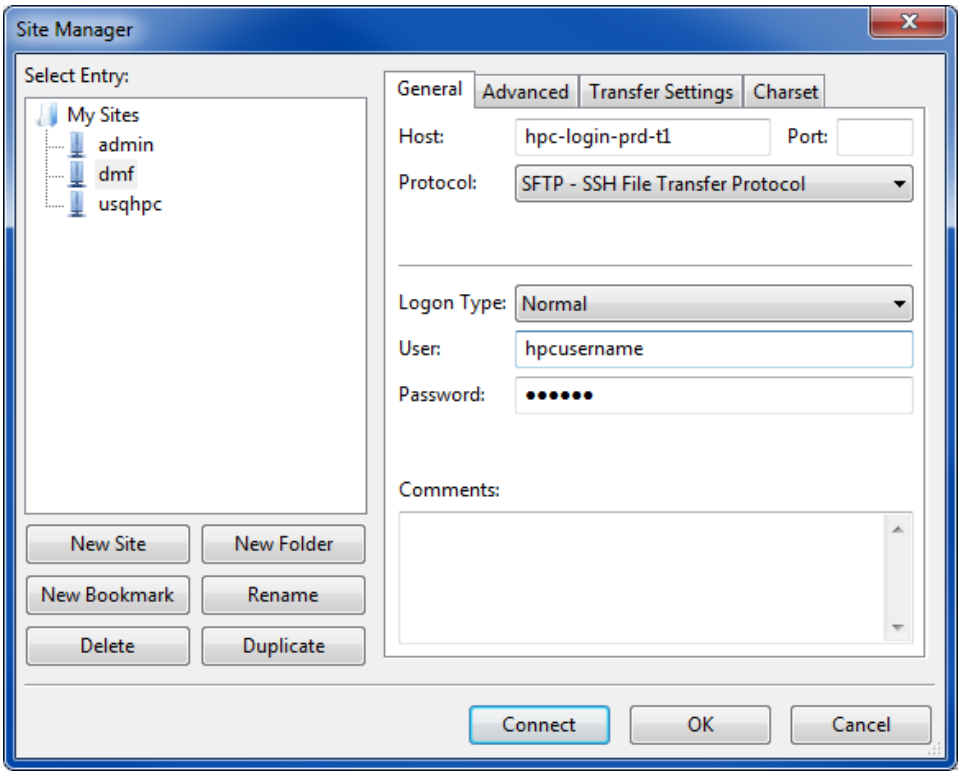

PBS and Updated Terminology

Question:

The term "nodes" is deprecated so it probably should not be used. You are using deprecated terminology and it is unclear to me how that is mapped by the software to the form it currently expects

The correct form is 'host' chunk and 'vnode'. It appears to me each host consists of 2 chunks (the two cpus) each chunk has 10 cores and each core has approximately 6GB (ignoring the hyperthreading) Whether the cpu should be defined as a chunk or a vnode I am unclear.

Vnode concept is a logical construct. It does NOT necessary map to one cpu. In your compute node, typically there are 2 sockets, so 20 cores in total per node, i.e. the Natural Node. This one Natural node can be divided into multiple Vnodes, which can be assigned any number of cores and mem, as long as they add up to the total amount on that Natural node. E.g. you can divide 1 node into 2 vnodes, each vnode with 10 cores and 64gb OR you can divide 1 node into 3 vnodes, like 8cores-50GB, 8cores-50GB, 4cores-28GB (total is 20 cores, 128GB on that node). So unless you have configured the system to do so, by default your system will have 1 vnode = 1 natural nodes. For more information see "3.1 Vnodes: Virtual Nodes" to setup PBS Pro 13 Administration Guide.

The *select* statement uses the Chunks concept - which can be within a vnode or spread across vnodes or natural nodes. Again in your case, 1 vnode is a real physical natural node (unless you setup differently)

e.g. -l select=2:ncpus=3:mem=4Gb+5:ncpus=7:mem=6Gb

This means, give me 2 chunks - with 3 cores and 4GB per chunk => 6 cores + 8GB PLUS

e.g. 5 chunks - with 7 cores , 6gb per chunk => 35 cores and 30GB

Each chunk sit inside one vnode. Multi chunks can share a vnode (other settings).
This concept of chunk-select statement enables a very flexible mix-match specification so if you have different nodes in your cluster, you can specify non-uniform resources on nodes.

Question:
The Intel MPI installation does not appear to include documentation. Need man pages installed or some indication of where the documentation can be found.

Intel MPI documentation is available at Intel or it's not in the man path. To add the man pages for Intel MPI to the man path, either:

module load mpi/impi-5.1.1.109

or

export MANPATH=${MANPATH}:/usr/local/opt/intel/impi/5.1.1.109/man