

Resources Request Examples

Category: PBS on Pleiades

Since Pleiades consists of four different processor types (Harpertown, Nehalem-EP, Westmere, and Sandy Bridge), you will benefit from keeping the following in mind when requesting PBS resources for your job:

Charging on the usage of the four Pleiades processor types is based on a common Standard Billing Unit which is on a per-node basis. The SBU rate for each of the Pleiades processor types is:

Processor Type	SBU Rate (per node)
Sandy Bridge	1.82 (16 cores in a node)
Westmere	1.0 (12 cores in a node)
Nehalem-EP	0.80 (8 cores in a node)
Harpertown	0.45 (8 cores in a node)

The actual amount of memory per node through PBS is slightly less than 7.6 GB per node for Harpertown, 22.5 GB per node for Nehalem-EP and Westmere, and about 31 GB per node for Sandy Bridge.

For the **normal**, **long**, and **debug** queue, use the **model=[har, neh, wes, san]** attribute to request the processor type(s) for your job. If the processor type is not specified in your PBS resource list, the job is routed to use the default processor type, Harpertown, for most queues. For the **devel** queue, see [Pleiades devel queue](#) for more information. For the **gpu** queue, see [GPU Basics](#).

Example 1

Here are some examples of requesting certain processor models for a 128-process job:

```
#PBS -l select=16:ncpus=8:model=har
#to run all 8 cores on each of 16 Harpertown nodes
```

```
#PBS -l select=32:ncpus=4:model=har
#to run on only 4 cores on each of 32 Harpertown nodes
#(note: will be charged for 32 nodes = 256 cores)
```

```
#PBS -l select=16:ncpus=8:model=neh
#to run all 8 cores on each of 16 Nehalem-EP nodes
```

```
#PBS -l select=11:ncpus=12:model=wes
#to run all 12 cores on each of 11 Westmere nodes
```

```
#(4 cores in 11th node will go unused)

#PBS -l select=8:ncpus=16:model=san
#to run all 16 cores on each of 8 Sandy Bridge nodes
```

Note that you can specify both the queue type (`-q normal, debug, long, low`) and the processor type (`-l model=har, neh, wes, san`). For example:

```
#PBS -q normal
#PBS -l select=16:ncpus=8:model=neh
```

If your application can run on any of the four processor types, you may want to submit your job to a processor type that has more nodes unoccupied by other running jobs. Doing this can possibly reduce the wait time of your job. The script `node_stats.sh` provides information about the total, used, and free nodes for each processor type. For example:

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

```
Pleiades Nodes Total: 11240
Pleiades Nodes Used : 10485
Pleiades Nodes Free : 755
```

```
Harpertown   Total: 4063, Used: 4015, Free: 48
Nehalem      Total: 1253, Used: 847, Free: 406
Westmere     Total: 3957, Used: 3916, Free: 41
SandyBridge  Total: 1288, Used: 1128, Free: 160
GPU nodes    Total: 56, Used: 2, Free: 54
Devel queue  Total: 623, Used: 577, Free: 46
```

```
Currently queued jobs are requesting: 8259 Harpertown, 2142 Nehalem, 9633 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.

For each job, you can also identify which processor models are used by looking at the "Model" field of the output of the command:

```
% qstat -a -W o=+model
```

Example 2

The Harpertown nodes in rack 32 have 16 GB memory/node instead of 8 GB per node.

This example shows a request of 2 nodes with bigmem in rack 32.

```
#PBS -l select=2:ncpus=8:bigmem=true:model=har
```

Example 3

For a multi-node PBS job, the NCPUs used in each node can be different. This is useful for jobs that need more memory for some processes, but less for other processes. Resource requests can be done in "chunks" for a job with varying NCPUs per node.

This example shows a request of two resource chunks. In the first chunk, 1 CPU in 1 node, and in the second chunk, 8 CPUs in each of three other nodes are requested:

```
#PBS -l select=1:ncpus=1+3:ncpus=8
```

Example 4

A PBS job can run across different processor types. This can be useful in two scenarios:

- When you cannot find enough free nodes within one model for your job
- When some of your processes need more memory while others need much less

This can be accomplished by specifying the resources in chunks, with one chunk requesting one processor type and another chunk requesting a different processor type.

Here is an example of how to request 1 Westmere node (which provides 24 GB per node) and 3 Harpertown nodes (which provides 8 GB per node), under either of the following circumstances:

```
#PBS -lplace=scatter:excl:group=model
```

```
#PBS -lselect=1:ncpus=12:mpiprocs=12:model=wes+3:ncpus=8:mpiprocs=8:model=har
```

Article ID: 188

Last updated: 26 Nov, 2012

Computing at NAS -> Running Jobs with PBS -> PBS on Pleiades -> Resources Request Examples

<http://www.nas.nasa.gov/hecc/support/kb/entry/188/?ajax=1>