

Instructions for setting up MPI daemon:

These steps are needed to be done only once. However for whatever reason the cluster is rebooted then you may have to do this again. So keep these instructions handy.

To be able to run MPI programs you need to have MPI daemons (mpd) running on the cluster. To do this please follow the below directions:

There are two approaches to setting up mpd:

Approach #1 (recommended)

1. First, you will need to create a machinefile with the name of each of the nodes in the cluster and how many processors each has (separated by a colon). The order of the nodes doesn't matter. However for better load balancing I suggest you randomize the order. This is to ensure the individual MPI runs are likely to seek out unused nodes when multiple users are using the system.

Example Machine file (saved as machinefile.txt in your home directory):

```
glx3.eecs.wsu.edu:4  
glx1.eecs.wsu.edu:4  
glx4.eecs.wsu.edu:4  
glx5.eecs.wsu.edu:4  
glx2.eecs.wsu.edu:4
```

2. Next, you need to cleanup any running mpd processes on the 5 nodes. To do this ssh to individually to each of glx nodes and run the command *"mpdcleanup"* on each.
3. Next, you need to launch mpd. To do this, first ssh to any of the compute nodes (can be any of the nodes glx1-5, which one doesn't matter), and run the following command:

```
"mpdboot -n 5 -f <path>/machinefile.txt"
```

(where <path> is the folder where you saved the machinefile.txt)

The "-n 5" part of the argument specifies the number of compute nodes on which to run mpd on (note: there will be only one mpd process running on each compute node, regardless of the number of cores each has, by the end of this step).

After you press *enter* for the previous command, you will be prompted to enter your password to each node. If the terminal seems to stop after you enter a password, simply press "enter" again and you should be prompted for the next password. If this step finishes successfully then the console prompt will be returned. To check if mpds are running fine, use the command *"mpdtrace"* from any of the compute nodes and that should list all the 5 nodes with mpd on it.

Approach #2 (use if #1 fails)

If the first approach fails, you will need to manually create a head *mpd* and subsequently *ssh* into each of the nodes separately and create additional *mpds* that link to your head *mpd*. The steps are as follows:

- 1) First make sure you run *mpdcleanup* from all nodes.
- 2) Next, *ssh* into one of the nodes (doesn't matter which one) and type "*mpd &*" into the command line. You should get an output similar to "[1] 23654". From here, you will need to know two things to link to this head *mpd*: the name of the node and the port number. To find the node name and port number, simply type "*mpdtrace -l*" and you will get an output similar to *glx5.eecs.wsu.edu_40305* (10.99.0.147). *glx5* is the name of the node in this case and the port number is 40305. Now comes the tedious part, you will need to *ssh* into each of the other nodes in the cluster and follow these instructions for each one:
 - a. Once you *ssh* into a node, type "*mpd -h glx5.eecs.wsu.edu -p 40305 &*" (I am using the above example). You can test if the *mpd* is linked correctly with a "*mpdtrace*". If it is not set up correctly, type "*mpdcleanup*" to delete the *mpd*, and then you can retry.
 - b. Finally, when you are done setting up each *mpd*, *ssh* into the node with your *head mpd* and type "*mpdtrace*" to check if it lists every node linked to your head *mpd* node including the head node itself.