Indian Institute of Technology Ropar
HPC User Manual, v1.4

Accessing the cluster

1. Shell Access
Unix-like OS (Linux, Mac OS X, BSDs, etc.)

If you are using a Unix-like OS, you can use SSH to log in to your cluster account.

Syntax:
$ ssh username@10.1.1.52

Microsoft Windows
If you are using Microsoft Windows, you can use PuTTY to access your cluster account.

2. Transferring Files
Unix-like OS (Linux, Mac OS X, BSDs, etc.)

If you are using a UNIX-like OS, you can use SCP to transfer files between your local machine and the cluster.

Syntax (uploading files):

$ scp <local path>/<local file(s)> username@10.1.1.52:<remote path>

Syntax (downloading files):

$ scp username@10.1.1.52:<remote path>/<remote file(s)> <local path>

Microsoft Windows
If you are using Microsoft Windows, you can use PSCP that comes bundled with PuTTY. It follows

a syntax similar to SCP, as mentioned above.
Alternatively, you can use WinSCP for a GUI based option.

Available Software
LAMMPS
GAMESS
CPMD
Gaussian 09 (Department Licensed)
VASP (Group Licensed)

Instructions for running programs

> No program shall be run directly on any of the nodes.

> You need to write a jobscript for running any of your programs, be it parallel or serial
code.

> You need to specify the number of CPU cores and the required wall time (time that youwant the job to run for) in the jobscript.

> Note: If your job exceeds the walltime, it will be automatically killed by the server.

> Use the qsub command to submit the jobscript to the server.

Syntax:
$ qsub jobscript.sh

Here's a **sample jobscript** for your reference. You can save it with any name (with or withoutextension), that you want. Though, a good convention is to name it jobscript.sh:

```
#!/bin/bash
#PBS l nodes=1:ppn=1
#PBS -l walltime=02:00:00
#PBS -e "$PBS_JOBID".err
#PBS -o "$PBS_JOBID".out
echo "PBS job id is $PBS_JOBID"
echo "PBS nodefile is at $PBS_NODEFILE"
NPROCS=$(wc -l < "$PBS_NODEFILE")
echo "NPROCS is $NPROCS"
cat "$PBS_NODEFILE" > nodes
mpirun -machinefile "$PBS_NODEFILE" -np "$NPROCS"
/apps/gcc/vasp.4.6/vasp
```

Queues

As a fair usage policy, the following types of queues (see table below) have been implemented
as of now. Depending on the number of CPU cores specified in your jobscript, your job will be automatically assigned to one of the queues specified in the table given below.

> Jobs will be run on a First Come, First Served basis.
> If your jobscript doesn't satisfy any of the contraints specified in the following table, itwill be rejected by the server.

• If enough free resources aren't available in the cluster, your job will have to wait in the"Idle Jobs" queue.

• If you submit more jobs than are permissible as per the "Run/user" column below,
your job will be deferred to the "Blocked Jobs" queue.

No. of CPU cores	**Time (hours)**	**Max. no. of CPU Core**	**Min. no.of CPU Core**	**Run/user**	**Queue/user**	**Nodeset**	**Queue**
1	**240**	**1**	**1**	**20**	**20**	**serial**	**serial**
20	**120**	**20**	**20**	**4**	**4**	**ompq**	**ompq**

Node reservation:

5 nodes (100 cores) are reserved for the serial queue. This includes the 2 GPU nodes as

well.
Rest of nodes are reserved for the ompq queue.

Job management

Apart from the "qsub" command that you have already been introduced to, the following
commands will come in handy for you.

To see the status of all jobs, you can use either of these commands:
$ showq
or
$ qstat -a

To see detailed information regarding a job, you can use:
$ checkjob <jobid>

To cancel a submitted job, you can use:
$ canceljob <jobid>

To see the estimated start/complete time for a job, you can use:
$ showstart <jobid>

Cluster Information

Total .NO of nodes	**34**
No. of master nodes	**2**
No. of compute nodes	**30 (28 CPU, 2 GPU)**
No. of storage nodes	**2**

CPU node names	**c1 (a-d) to c7 (a-d)**
GPU node names	**g1 and g2**
CPUs/compute node	**2 x Intel Xeon E5-2670v2 (10 cores, 2.5 GHz)**
Total CPU cores (compute nodes)	**600**
GPUs/GPU node	**2 x NVIDIA Tesla K20Xm**
Memory/compute node	**96 GB (CPU nodes), 128 GB (GPU nodes)**
Storage	**28TB**

Software

Operating System: CentOS 6.5 (x86_64)
Resource Manager: TORQUE
Job Scheduler: Maui
MPI Library: Open MPI

Charging Policy

At present there are no charges for using the HPC. Later, if any charging policy is introduced,users will be informed.

Contact

System Administrator: Shubham Mishra <shubham.m@netwebindia.com>