

KALINGA CLUSTER BASIC MANUAL

1. Kalinga Cluster old master node IP Address: 10.10.0.136

2. Kalinga Cluster new master node IP Address: 10.10.0.135

Jobs can be submitted from both the master nodes.

3. List of installed software:

Command to check list of Installed software

```
[root@kalinga ~]# module avail
```

Command to load the particular software. (Quantum Espresso)

```
[root@kalinga ~]# module load codes/qe/6.5
```

Command to load the Intel compilers

```
[root@kalinga ~]# module load compilers/intel/parallel_studio_xe_2018.3.051
```

4. Queue Details:

To see the queue details:

sinfo

Queue Details:

Queue Name	Wall Time	Maximum number of nodes
1. Single	72 hours	1 node (40 core)
2. Small	48 hours	1 to 4 nodes (40 to 160 core)
3. Large	120 hours	5 to 16 nodes (200 to 640 cores)
4. Serial	336 hours	4 nodes (160 cores)
5. gpu1	28 days	1 node (32 cores) with 4 GPU
6. old_single	48 hours	1 node (32 core)
7. old_large	72 hours	4 to 32 nodes

NOTE: Single, Small, large, serial queue/partition is from new compute nodes.

In SLURM Scheduler, Queue is called as Partition.

5. PFS Quota:

User Quota (Soft Limit: 1.5TB, Hard Limit: 2TB)

6. Sample job scripts

Create your jobs script for jobs submission, Below is the sample script for some of the application. For other applications you just need to load the appropriate modules in the script. Also you need to specify the nodes, ppn and queue in the script.

Sample script is available in both the master nodes PATH: /opt/sample_script

```
=====
#!/bin/bash
#SBATCH --job-name=Ngen_test    # Job name
#SBATCH --partition=single    # Partition Name(queue)
#SBATCH --ntasks-per-node=40    # No. of CPU in one node
#SBATCH --nodes=1            # No. of Node for jobs
#SBATCH --output=relax_%j.log  # Standard output and error log
echo "Job started at date"
date
module load codes/qe/6.5
mpirun -np $SLURM_NPROCS pw.x -inp relax.in > relax.out
echo ----
echo "Job ended at date"
date
```

Sample script for serial jobs

Kindly use the below parameters for submitting jobs in serial queue

```
#!/bin/bash
#SBATCH --job-name=serial_job_test  # Job name
#SBATCH --partition=serial    # Partition Name (queue)
#SBATCH --qos=serial
#SBATCH --nodes=1            # No. of Node for job
#SBATCH --ntasks-per-node=1    # Run on a single CPU
#SBATCH --output=serial_test_%j.log  # Standard output and error log
```

For old_compute

```
#SBATCH --ntasks-per-node=32    # No. of CPU in one node
```

Additional Details:

```
#SBATCH --ntasks-per-core=20    # No. of CPU's in one socket (new compute nodes)
#SBATCH --ntasks-per-core=16    # No. of CPU's in one socket (old compute nodes)
```

7. To submit the job.

```
[root@kalinga ~]# sbatch <script file>
```

8. To check the details of the job

```
[root@kalinga ~]# scontrol show job <job-id>
```

9. To Cancel the running jobs

```
[root@kalinga ~]# scancel <job-d>
```

10. For checking the resources available

```
[root@kalinga ~]# sinfo
```

11. For Checking the status of the job

```
[root@kalinga ~]# squeue
```