

# HPC Knowledge Portal

## Slurm Training '15

### Slurm Workload Manager Hands-On 02

In this hands-on, you are going to simulate a very large environment thanks to tuned Slurm setup. You will be able to setup different QoS and scheduling strategies, and submit some jobs to this virtual cluster in order to see how it behaves.

Estimated time : 45 minutes

#### ToDo

- Login to SuSE Studio Slurm Appliance  
<https://susestudio.com/a/MnLYey/slurm>
- Click on Test Drive, accept the EULA
- Enable the Networking
- Connect via SSH
- Add PriorityFlags values i.e.  
`PriorityFlags=SMALL_RELATIVE_TO_TIME`
- Activate the Multi Factor Priority  
`PriorityType=priority/multifactor`
- Since we want to test the behavior in short period of time, the following values are very short (and unrealistic).
- We suggest to submit large jobs with short walltime using different user accounts in order to see how the competition evolves between these users.  
`PriorityDecayHalfLife=00:10:00`  
`PriorityFavorSmall=NO`  
`PriorityMaxAge=00:01:00`  
`PriorityUsageResetPeriod=HOURLY`  
`PriorityWeightAge=500`  
`PriorityWeightFairshare=10000`  
`PriorityWeightJobSize=500`  
`PriorityWeightPartition=1000`  
`PriorityWeightQOS=0`
- Create three partitions attending to relative priority (low, medium and high).
- Limit those partitions with reasonable walltime limits.
- Setup different priority for each partition.
- Setup low partition as default.
- Start the daemon with the following command line:  
`service slurm startclean`
- Change the previous values and restart the daemons in order to see the impact in the scheduler.
- In order to explore the impact of QoS, change the following parameter:  
`PriorityWeightQOS=100000`
- Setup a default QoS with a huge priority difference for the existing users and restart the daemons.  
`service slurm stop`  
`service slurm startclean`
- After that, you can simulate activity in order to explore the behavior.

#### Requirements

- Desktop or Laptop with SSH client.

#### Quick Reference

##### Slurm commands

- `sacct` Extract accounting information from cluster.
- `scancel` Job deletion.
- `sinfo` Shows status information about cluster.
- `squeue` Status listing of jobs and queues.
- `sview` GUI to view job, node and partition information.
- `smap` CLI to view job, node and partition information.
- `sbatch` Command line interface to submit jobs.

##### QoS Examples

Add a new qos

```
sacctmgr add qos NameOfQoS MaxCpusPerUser=100
```

Add a qos to a user account

```
sacctmgr modify user name=user01 set  
qos+=NameOfQoS defaultqos=NameOfQoS
```

Set QOS priority

```
sacctmgr modify qos NameOfQoS set priority=100
```

Set Max CPU minutes limit (60 minutes \* 24 hours)

```
sacctmgr modify qos NameOfQoS set GrpCPUMins=1440
```

Set Max CPUs per group

```
sacctmgr modify qos NameOfQoS set GrpCpus=1000
```

Set Max Jobs per group

```
sacctmgr modify qos NameOfQoS set GrpJobs=1000
```

Set Max Nodes per group

```
sacctmgr modify qos NameOfQoS set GrpNodes=1000
```

Set Max Submit Jobs per group

```
sacctmgr modify qos NameOfQoS set  
GrpSubmitJobs=1000
```

##### Simulate Usage Activity

- You can use submit script based on examples folder.
- Or use the following long command line:  

```
sbatch -p high -n 512 \  
--wrap="env; srun -n 1 sleep 120"
```
- You can populate the virtual cluster with jobs using job arrays : `--array=1-1000`
- You can submit as a different user with `su` command line:  

```
su - hpckp01 -c "submit command"
```

#### Reference

- [Slurm Training Slides](#)
- [Slurm Rosetta Stone](#)
- [Slurm Official Documentation](#)

# HPC Knowledge Portal

Slurm Training '15