

You have been assigned to run a campus commodity cluster. There are 4 groups on campus, each with 2 users. All groups are already on the system submitting jobs.

Introduction

1. Get private ssh keys and ip address via instructors
2. Start ssh session using the key
 - a. If you need help importing the key, please tell us and we can accommodate
 - b. You will log in as centos
 - c. You can be root at this point
3. There are 7 nodes available for you to use for this practical
 - a. The head node is scheduler-*-0
 - i. This is the only publically accessible node
 - b. The compute nodes are compute-*-[0-1]
 - c. The storage nodes are storage-*-[0-3]
4. These nodes have the following specifications:
 - a. 2 CPUs
 - b. 1GB of memory
 - c. 20GB
 - d. Centos 7 latest
5. Today's portion will only work with the 3 scheduler oriented nodes
6. The home directories are NFS mounted from the scheduler nodes
7. Root has ssh keys only from the scheduler node to all other nodes
8. There are 8 users and 4 groups set up for the practical
 - a. There are ssh-keys between all nodes for these users
9. For the sake of time and ease, we are going to compress a week of scheduling time into 1 hour with 6 days in a week

Familiarize yourself with your cluster

1. Slurm is installed with the latest version
 - a. slurm 17.02.6
2. The cluster name is 'lcicluster'
3. There is currently 1 partition enabled named 'normal'
 - a. This has a 20 minute job
4. Your cluster is currently being used by "users" from a university
 - a. Jobs will be continually submitted by "users" throughout the session
 - b. We will tune the cluster to make it run more efficiently
5. Submit an interactive job
 - a.

```
$ salloc -N 1
$ srun hostname
```
5. There is currently a 20 minute limit on the single normal partition
6. MariaDB is enable with Slurm for accounting although nothing is enforced

Exercise 1: Enable Fairshare

Oh no! A professor Bob is complaining about the fairness of the cluster. They only run 4 core jobs a few times a week and other are running a ton of single core jobs. Enable fairshare to give all users an equal share.

1. The documentation for enabling fairshare is located here:
 - a. https://slurm.schedmd.com/priority_multifactor.html
2. The following are the parameters that should be modified in the `/etc/slurm/slurm.conf`
 - a. `PriorityType=priority/multifactor`
 - b. `PriorityDecayHalfLife=0-00:10:00`
 - c. `PriorityCalcPeriod=1:00`
 - d. We need to apply a weight to each factor in the multifactor equation
 - i. `PriorityWeightAge=1000`
 - ii. `PriorityWeightFairshare=10000`
 - iii. `PriorityWeightJobSize=1000`
 - iv. `PriorityWeightPartition=1000`
3. Reread the configuration
 - a. `scontrol reconfigure`
4. Verify job priority is being modified properly
 - a. `scontrol show config`
5. Let's check to see if anyone's priority has changed
6. We can now read the shares
 - a. `sshare`

Exercise 2: Modify Fairshare for Groups and enable accounting

The professors have talk with each other and they want to share the machine equally between groups.

1. We need to enable accounting for the groups
2. Add the cluster to the accounting
 - a. `sacctmgr add cluster lcicluster`
2. Next add the accounts (see table below)
 - a. `sacctmgr add Account <accountname>`
3. We know need to associate the users with the accounts
 - a. `sacctmgr add user <userid> Account=<accountname>`
4. We need to also make each account the user's default
 - a. `sacctmgr add user <userid> DefaultAccount=<accountname>`
 - b. Verify all associations
 - i. `sacctmgr show associations`
5. We need to configure fairshare to use the parent
 - a. Add the following to your `/etc/slurm/slurm.conf`

- i. `FairShare=parent`
- 6. Reread the configuration
 - a. `scontrol reconfigure`
- 7. Wait a bit and we will view the shares
 - a. `sshare`

Group	Users
bio	kara
	tim
physics	bob
	mary
math	rae
	david
astro	joe
	alice

Exercise 3: Issues with Priority

Professor Bob in the physics is back to having issues getting high priority jobs through. He has a conference coming up and needs the use of the entire cluster. Create a reservation that allows his group to use the entire cluster for the next 30 minutes. Then create a high partition with all the nodes that has a higher starting priority.

1. Create the reservation for his group
 - a. `scontrol create reservation account=astro duration=60 nodecnt=2`
 - b. `scontrol create reservation starttime=now users=bob,mary duration=60 nodecnt=2 partition=normal`
 - c. Verify the reservation exists
 - i. `scontrol show res`
2. Create a new partition named high
 - a. Copy the partition line in the `/etc/slurm/slurm.conf` for normal and rename it high and change the default
 - i. `PartitionName=high Nodes=compute-0-[0-1] Default=NO MaxTime=20:00 State=UP`
 - b. Reread the configuration and verify the partition exists
 - i. `scontrol reconfigure`
 - ii. `sinfo`
 - c. Create a QOS for each partition and set a priority for high
 - i. `sacctmgr add qos <partitions>`
 - ii. `sacctmgr modify qos high set priority=1000`
 - iii. `sacctmgr show qos format=name,priority`

- d. Give QOS a fairshare weight
 - i. Add the following to the /etc/slurm/slurm.conf
 1. `PriorityWeightQOS=10000`
- e. Assign the QOS to each partition
 - i. Add the following to each Partition line in /etc/slurm/slurm.conf
 1. `QOS=<qosname>`
 - ii. `scontrol reconfigure`
- f. Check jobs and verify the qos is applied

Exercise 4: Limiting Groups with Accounting

The IT department has decided they want to sell the resources to each department to help fund the machine. Limit each project to 5 cpu hours per account. Modify partitions so that the high priority is double the cost of the normal partition

1. Set all groups to 5 cpus hours
 - a. `sacctmgr modify account <accountname> set GrpCPUMins=300`
2. Verify the accounts
 - a. `sacctmgr show Association`
3. We need to now make sure everything we are doing is enforced
 - a. Add the following to your /etc/slurm/slurm.conf
 - i. `AccountingStorageEnforce==associations,limits,qos`
 - ii. `scontrol reconfigure`
4. Change the change rate of the high queue
 - a. `sacctmgr modify qos high set usagefactor=2`

Exercise 5: Enable Preemption for a Low Partition

The users would like a low queue in which they can submit jobs that might be preempted but get changed less as well. Make sure these jobs only backfill and are half the cost of normal jobs.

1. Copy the partition line in the /etc/slurm/slurm.conf for normal and rename it low and change the qos
 - a. `PartitionName=low Nodes=compute-0-[0-1] Default=NO MaxTime=20:00 State=UP QOS=low`
2. Create the QOS for the low partition
 - a. `sacctmgr add qos low`
 - b. `sacctmgr modify qos low set priority=0`
 - c. `sacctmgr show qos format=name,priority`
3. Make sure the QOS has the Flag NoReserve to make sure it only backfills
 - a. `sacctmgr modify qos low set FLAGS=NoReserve`
4. Bump the normal partition's priority to 100 to be in between the low and high queue
 - a. `sacctmgr modify qos low set priority=100`
5. Change the change rate of the low queue

- a. `sacctmgr modify qos high set usagefactor=0.5`
6. `Scontrol reconfigure`

Exercise 6: Setup a Singularity for a User

Biology wants to use a Docker image for their genomics project. They need ubuntu. Follow the instructions to setup an Ubuntu image for them. Singularity has already been installed on the compute nodes and scheduler.

1. Create an image
 - a. `singularity create ubuntu.img`
2. Import your container
 - a. `singularity import ubuntu.img docker://ubuntu`
3. Get a shell inside of the container
 - a. `Singularity shell ubuntu.img`
 - b. `Cat /etc/debian_version`
 - c. `id`