HPC basics and introduction to Xanadu cluster

Vijender Singh

Computational Biology Core

UCONN
UNIVERSITY OF CONNECTICUT

HPC architecture : Xanadu (Queue/Partition info)

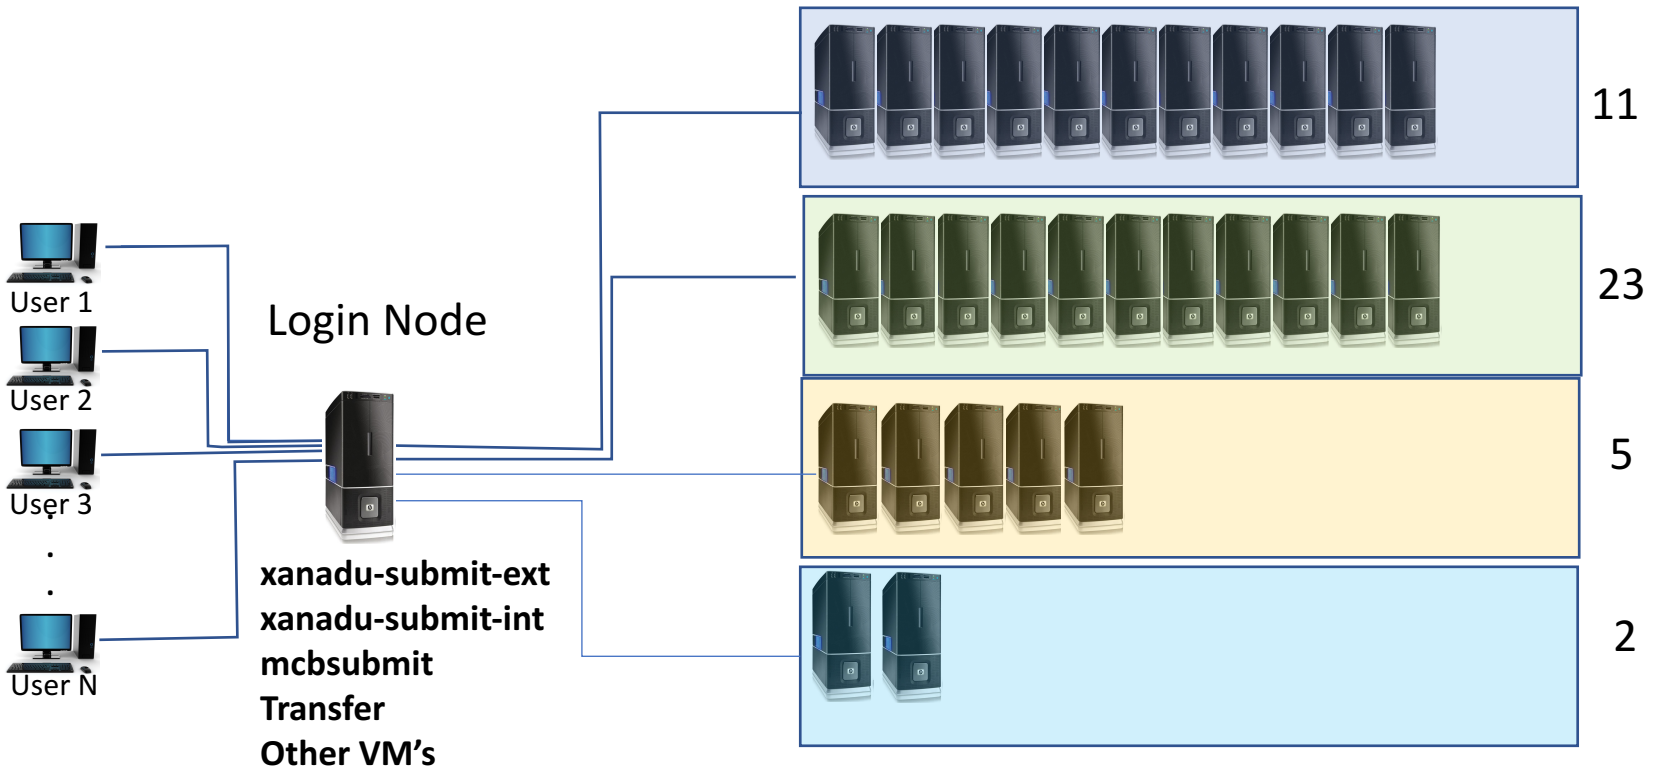

Queue/Partition	Nodes
xeon/general	xanadu03-13
amd/general	shangrila01-19 xanadu20-28
himem.	xanadu29-31
mcbstudent	xanadu01-02

cpu : 4
RAM: 8GB

cpu : 32 – 48
RAM: 128 – 512 GB

Cluster Related Commands:

hostname	Displays node name on which user is active.
sinfo	Summary of cluster nodes with their status
sinfo –N –l	Detailed information on compute nodes.
sinfo –N –l –p himem	Detailed information on compute nodes associated with a queue
srun –pty –p general –qos=general bash	interactive session on **general** partition with **one core** and **128MB RAM** (default)
srun –pty –p general –qos=general –mem=2G bash	interactive session on **general** partition with **one core** and **2GB RAM.**
srun –pty –p general –qos=general –mem=2G -c 4 bash	interactive session on **general** partition with **4 core** and **2GB RAM.** (Increase core number only if you are using multithread applications)

Cluster Related Commands:

Resource request header for script

```
#!/bin/bash
#SBATCH --job-name=myscript
#SBATCH -n 1
#SBATCH -N 1
#SBATCH -c 8
#SBATCH –mem=40G
#SBATCH –qos=general
#SBATCH --partition=general
#SBATCH --mail-type=END
#SBATCH --mail-user=first.last@uconn.edu
#SBATCH -o myscript-%j.out

hostname
```

sbatch myscript.sh	Submit your batch job on cluster.
squeue	list all the jobs on the cluster cluster.
squeue -u $USER **or** squeue –u userid	list all the jobs of a user.

Cluster Related Commands:

squeue | grep himem	List all the jobs running on himem node.
572983 general expre xyzav R 17:03:18 1 xanadu-27	
scontrol show jobid 572983	Details of a particular job running on cluster.
scancel 572983	Cancel a particular job running on cluster using jobid.
scancel --name expre	Cancel a particular job by name.
scancel –u $USER	Cancel all job of a user on cluster.
scontrol hold 572983	Pause or hold a job on cluster.
scontrol resume 572983	Resume a held job on cluster.
scontrol requeue 572983	Requeue (cancel and rerun) a particular job.

Himem nodes are meant to process jobs with large memory requirements **>250G**

Bad Examples of Resource usage

```
NodeList=xanadu-31
BatchHost=xanadu-31
NumNodes=1 NumCPUs=32 NumTasks=1 CPUs/Task=32 ReqB:S:C:T=0:0:*:*
TRES=cpu=32,mem=350G,node=1
```

```
NodeList=xanadu-31
BatchHost=xanadu-31
NumNodes=1 NumCPUs=12 NumTasks=1 CPUs/Task=12 ReqB:S:C:T=0:0:*:*
TRES=cpu=12,mem=125G,node=1
```

```
NodeList=xanadu-30
BatchHost=xanadu-30
NumNodes=1 NumCPUs=2 NumTasks=1 CPUs/Task=1 ReqB:S:C:T=0:0:*:*
TRES=cpu=2,mem=60G,node=1
```

Node configuration

Node	Cores	RAM
Himem	32	512G
Himem	32	512G
Himem	32	512G

Cluster Etiquette

Do not run code on the head node. Start a interactive session as soon as you can.
Do not ssh directly into a node.
Do not submit a large number of jobs without testing.
Do not Hog Resources.
Uses resources efficiently.
Himem nodes are meant to process jobs with large memory requirements >250G

Do monitor your jobs periodically
Monitor your disk usage:
Do not fill up the whole disk with unnecessary output files from your runs.