# Getting Started on Proteus

# Anatomy of the Cluster

- Login Nodes
  - proteusi01
  - proteusa01
- proteusmaster
  - Job scheduler
  - Administration
- Infinband
- Compute Nodes
- Storage
  - NFS
  - Lustre (temporary)

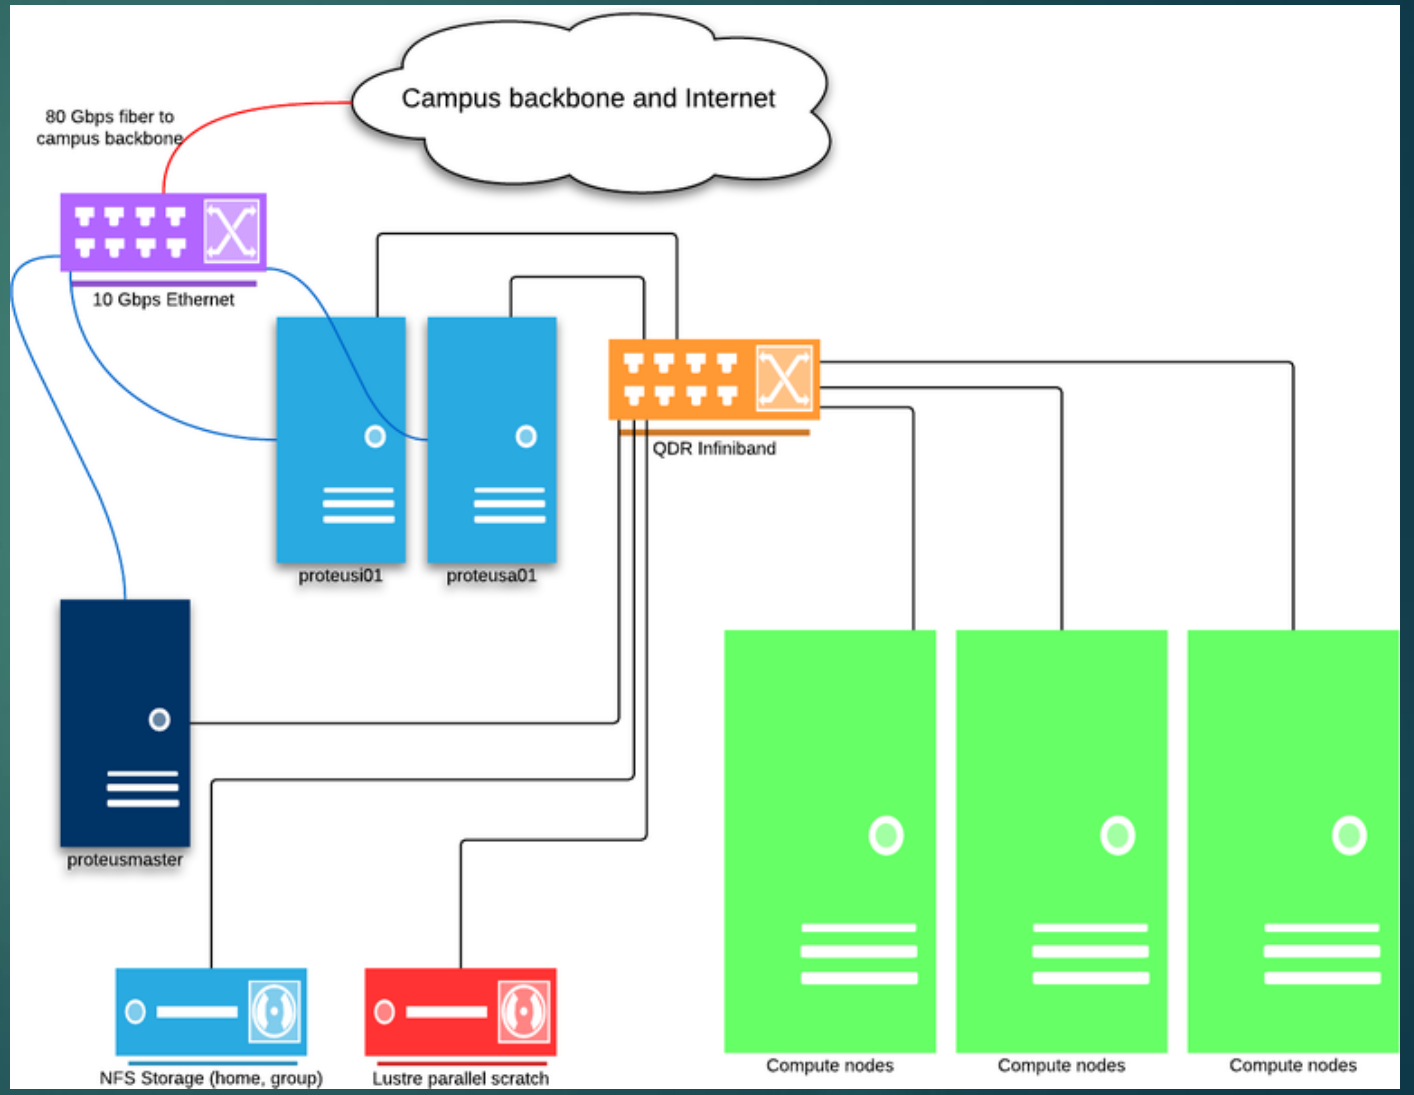

# Connecting to Proteus

- Two login nodes
  - proteusi01
  - proteusa01
- On Linux/MacOS
  - ssh proteusi01.urcf.drexel.edu
- On Windows
  - Putty
  - MobaXTerm
  - Terminus
  - Powershell
- If locked out of the system, use the Drexel VPN

# Navigating a Linux/Unix Environment

- ls
- cd <*directory*>
- mkdir <*directory*>
- rm <*file*>
- less <*file*>
- Some editors:
  - nano
  - vi
  - emacs
  - Or use one on your computer and upload the file

# Modules

- Modules show software already on Proteus

- module list
- module avail
- module load <*module*>
- module unload <*module*>

# Submitting Jobs

- qsub myjob.sh
- Arguments passed through command line or in the script with #$

| Arguments | What they do |
|---|---|
| -cwd | Tells the Job Scheduler to start in the current directory |
| -j y | Merges output and error |
| -S /bin/bash | Tells the job scheduler the job is a bash script |
| -P myPrj | The project your job will run under |
| -l h_rt=02:00:00 | Maximum time for the job to run |
| -l m_mem_free=2G | Minimum memory **per slot** |
| -l h_vmem=3G | Hard cap on memory **per slot** |
| -pe shm 16 | Requests 16 total slots/cores |
| -q all.q | What queue the job will be scheduled in |

Writing Job Scripts

# Resource Requests

- Hardware
  - Intel nodes – 16 slots – 63G*
  - AMD nodes – 64 slots – 252G*
  - new.q nodes – 40 slots – 187G*
- Memory requests are **per slot**
- Maximum time request on all.q is 48 hours
- More is not necessarily better

Hardware

* Total installed, not total available

# Monitoring your jobs

- qstat
  - Shows status of your current jobs
- qstat –j xxxxxx
- qstat –f –u *
  - Gives info on the cluster

```
[cwf25@proteusi01 OpenMPI]$ qsub hello_mpi.sh
Your job 541499 ("hello_mpi.sh") has been submitted
[cwf25@proteusi01 OpenMPI]$ qstat
job-ID     prior   name       user         state submit/start at     queue                          jclass                         slots ja-task-ID
-----------------------------------------------------------------------------------------------------------------------------------------------
    541499 0.81000 hello_mpi. cwf25        r     07/28/2020 13:44:56 all.q@ic07n01.cm.cluster                                       16
[cwf25@proteusi01 OpenMPI]$
```

qstat

# Once Job is Done

- Output file in format *myjob.sh.oxxxxxx*
- If –cwd is provided, files in submit directory

- qacct –j xxxxxx

- qacct –j xxxxxx >& test.out &

```
[cwf25@proteusi01 OpenMPI]$ qacct -j 541499
==============================================================
qname        all.q
hostname     ic07n01.cm.cluster
group        urcfcoopGrp
owner        cwf25
project      urcfcoopPrj
department   defaultdepartment
jobname      hello_mpi.sh
jobnumber    541499
taskid       undefined
account      sge
priority     0
cwd          /mnt/HA/groups/urcfcoopGrp/opt/Examples/OpenMPI
submit_host  proteusi01.cm.cluster
submit_cmd   qsub hello_mpi.sh
qsub_time    07/28/2020 13:44:55.820
start_time   07/28/2020 13:45:00.453
end_time     07/28/2020 13:45:47.628
granted_pe   fixed16
slots        16
failed       0
deleted_by   NONE
exit_status  1
ru_wallclock 47.175
ru_utime     0.134
ru_stime     0.565
ru_maxrss    5556
ru_ixrss     0
ru_ismrss    0
ru_idrss     0
ru_isrss     0
ru_minflt    27735
ru_majflt    78
ru_nswap     0
ru_inblock   14624
ru_oublock   536
ru_msgsnd    0
ru_msgrcv    0
ru_nsignals  0
ru_nvcsw     2794
ru_nivcsw    577
wallclock    835.380
cpu          0.699
mem          0.000
io           0.000
iow          0.000
ioops        488
maxvmem      214.363M
maxrss       2.738M
maxpss       1.270M
arid         undefined
jc_name      NONE
```

# Interactive Jobs

- qlogin
- Same options as qsub – all command line
- If running a GUI application, launch a terminal remotely, and use qrsh

```
[cwf25@proteusi01 ~]$ qlogin -pe shm 16 -l h_rt=00:30:00 -l h_vmem=4G -l m_mem_free=2G
Your job 541500 ("QLOGIN") has been submitted
waiting for interactive job to be scheduled ...
Your interactive job 541500 has been successfully scheduled.
Establishing builtin session to host ic21n01.cm.cluster ...
(base) [cwf25@ic21n01 ~]$
```

# Array Jobs

- Launch multiple similar jobs that don't depend on each other
- If you intend to launch over 50 jobs you **must** use an array job

- -t *n[-m[:s]]* my_job.sh
  - n – start id
  - m – end id
  - s – step

- Keep track of task ids with $SGE_TASK_ID
- Output file format: myjob.sh.oXXXXXX.MMM

# Job Dependencies

- Some jobs depend on the output of other jobs, SGE allows for jobs to be held until another job completes
- qsub -hold_jid xxxxxx myjob.sh
- -terse will return the jobid

Dependencies

# Storage

- Files should be stored in your research group directory, not your home directory
- Home directory has 15GB soft limit and 17GB hard limit

- For large file sizes use Lustre
  - /scratch/
  - Lustre files not accessed within 45 days will be deleted automatically

Lustre

# Let's Test Some Things