Guide to W&M Computational Resources

Eric J. Walter

Executive Director of Research Computing

September 21st, 2023

RC/HPC Staff

Eric J. Walter
Executive Director of Research Computing
Joined in 3/13

Jay Kanukurthy
Applications Analyst
(software installs / debugging, batch system, user onboarding)
Joined in 1/16

Matt Kennedy
Network, Filesystems & Data Specialist
(network, filesystems, data science hardware, tape backup, RC storage)
Joined in 4/20

Daanish Fiaz
Systems Administrator
(hardware, user accounts, security patching, tape backup)
Joined in 9/20

Malcolm Slaughter
Research Computing Specialist
(RC field support, RC storage, Linux workstation/software, special projects)
Joined in 7/23

Research Computing / HPC Training

1. Thursday – September 14th 5-6pm Zoom – Introduction to Linux command line
2. **Thursday – September 21st 5-6pm Zoom – Introduction to HPC at W&M/VIMS**
3. Thursday – September 28th 5-6pm Zoom – Linux in the HPC environment

- https://www.wm.edu/it/rc - shortcut to RC/HPC web
- hpc-help@wm.edu - help desk email – **best way to reach us**
- hpc-announce@wm.edu - list for announcements from RC (auto-subscribed)
- Office hours: by appointment – email hpc-help@wm.edu

Using HPC / Web Docs

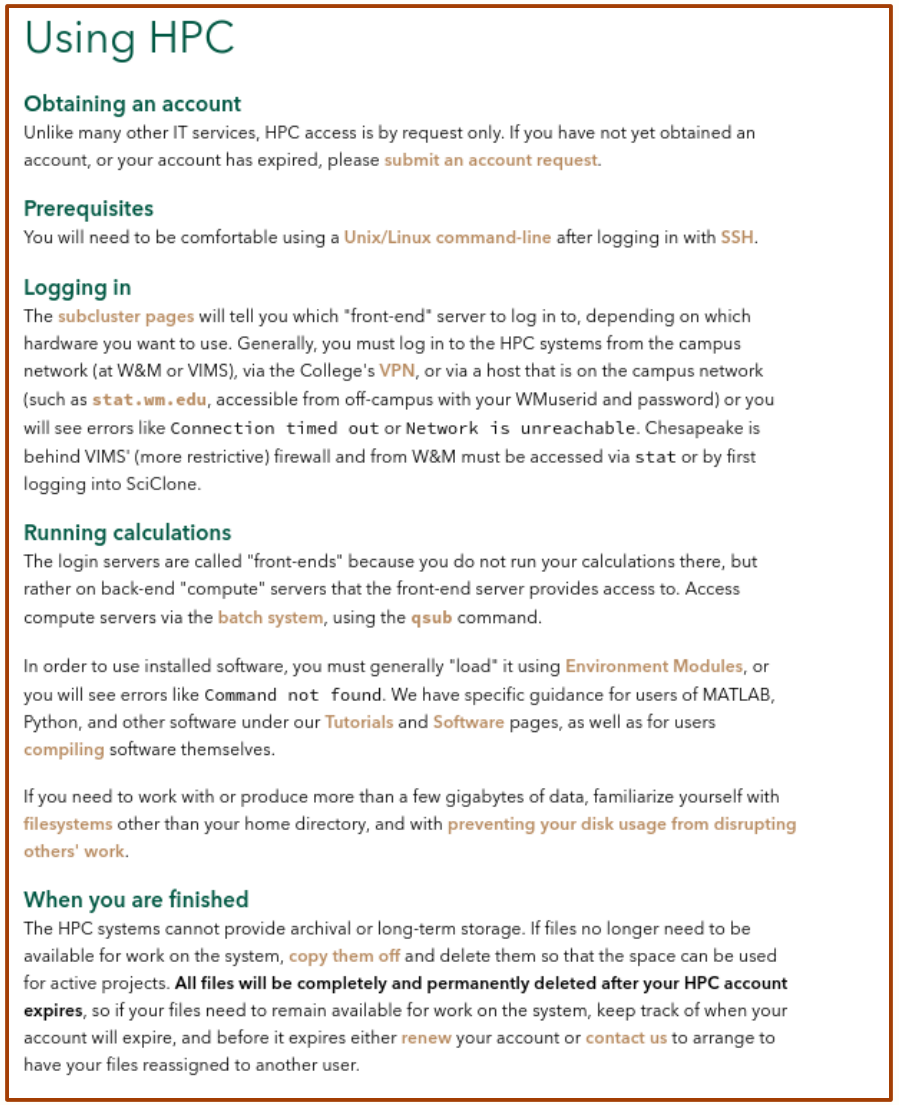

Using HPC

Obtaining an account

Unlike many other IT services, HPC access is by request only. If you have not yet obtained an account, or your account has expired, please submit an account request.

Prerequisites

You will need to be comfortable using a Unix/Linux command-line after logging in with SSH.

Logging in

The subcluster pages will tell you which "front-end" server to log in to, depending on which hardware you want to use. Generally, you must log in to the HPC systems from the campus network (at W&M or VIMS), via the College's VPN, or via a host that is on the campus network (such as `stat.wm.edu`, accessible from off-campus with your WMuserid and password) or you will see errors like `Connection timed out` or `Network is unreachable`. Chesapeake is behind VIMS' (more restrictive) firewall and from W&M must be accessed via `stat` or by first logging into SciClone.

Running calculations

The login servers are called "front-ends" because you do not run your calculations there, but rather on back-end "compute" servers that the front-end server provides access to. Access compute servers via the batch system, using the `qsub` command.

In order to use installed software, you must generally "load" it using Environment Modules, or you will see errors like `Command not found`. We have specific guidance for users of MATLAB, Python, and other software under our Tutorials and Software pages, as well as for users compiling software themselves.

If you need to work with or produce more than a few gigabytes of data, familiarize yourself with filesystems other than your home directory, and with preventing your disk usage from disrupting others' work.

When you are finished

The HPC systems cannot provide archival or long-term storage. If files no longer need to be available for work on the system, copy them off and delete them so that the space can be used for active projects. **All files will be completely and permanently deleted after your HPC account expires**, so if your files need to remain available for work on the system, keep track of when your account will expire, and before it expires either renew your account or contact us to arrange to have your files reassigned to another user.

What you need to get started

1. Getting an account
2. Linux command line / text editors
3. Logging into the clusters
4. Selecting/installing/requesting software
5. How to use file-systems effectively
6. How to use to use the batch system(s)
7. Compiling / installing your own applications
8. Saving your own files/projects

https://www.wm.edu/offices/it/services/researchcomputing/using/index.php

Cluster Nomenclature

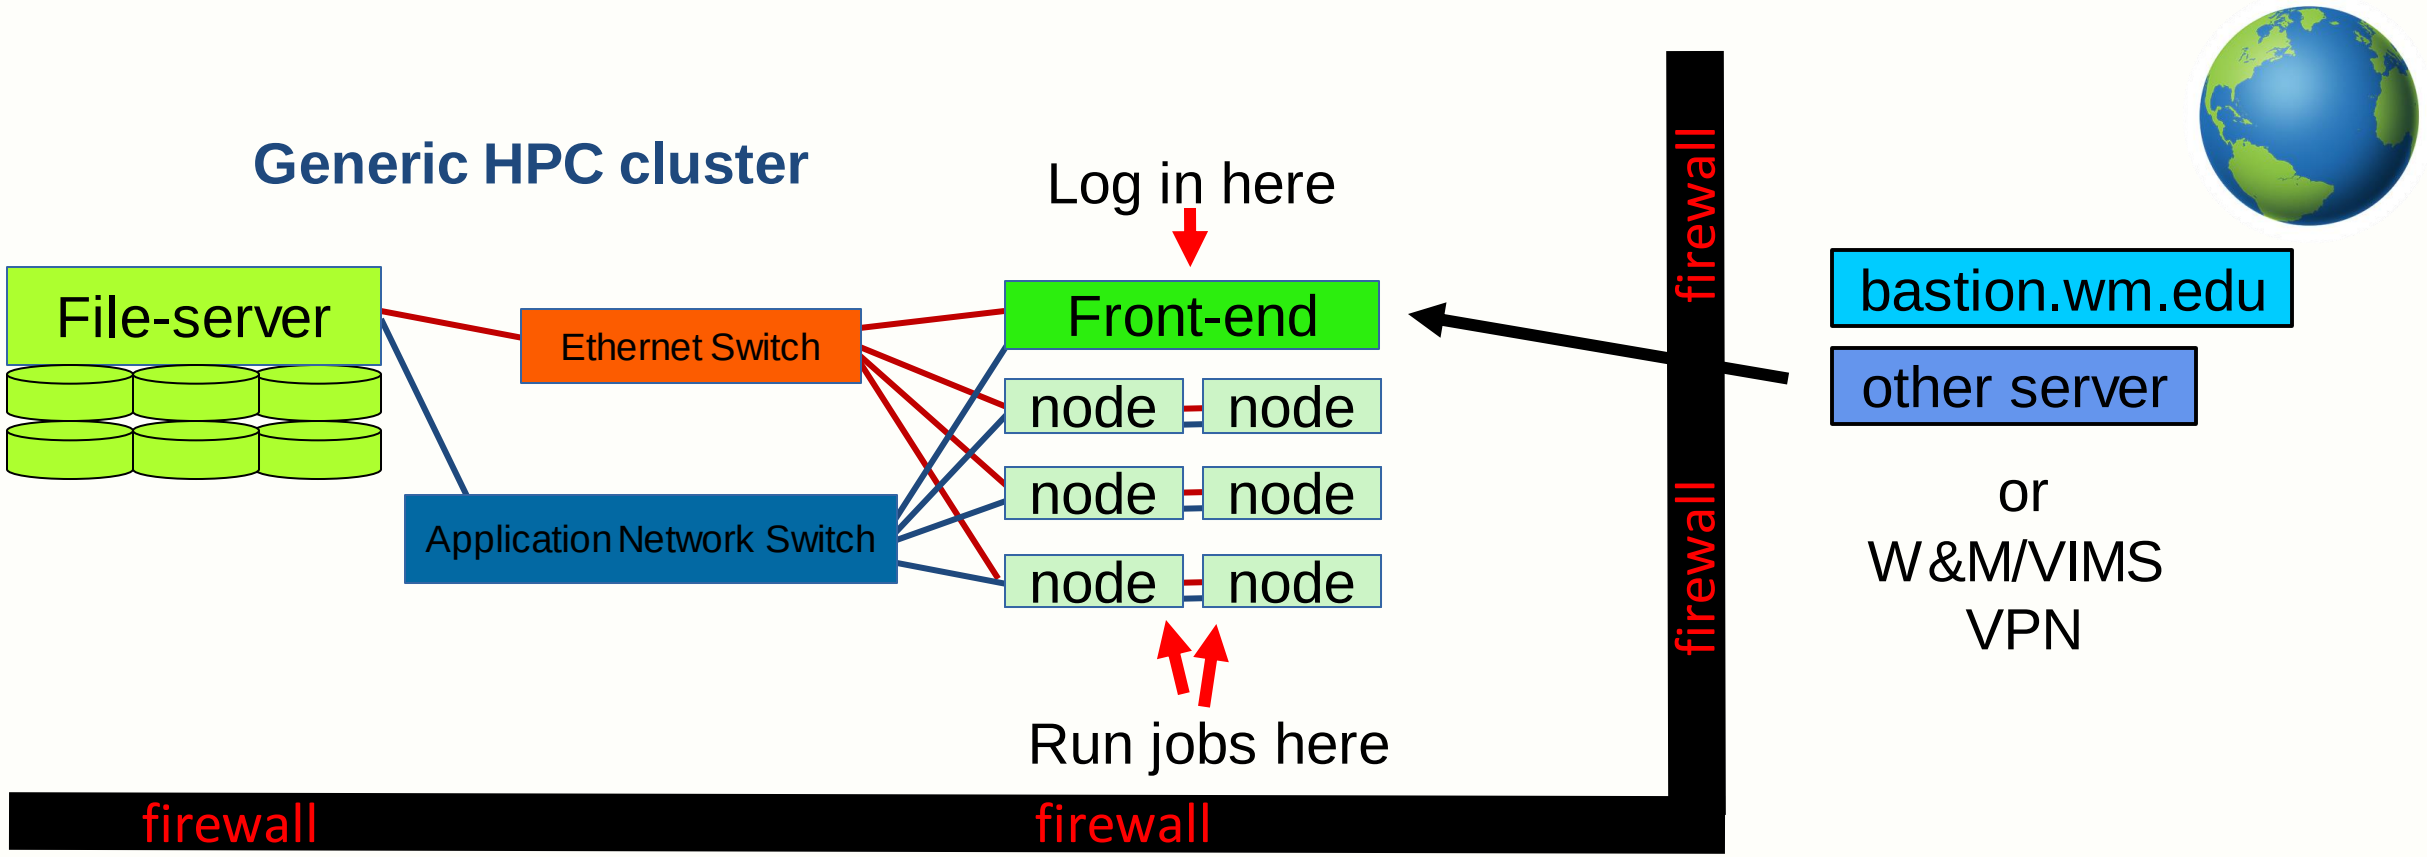

Must get through **W&M/VIMS firewall** to get to cluster resources
On campus you are already within firewall
Otherwise use **bastion.wm.edu**, another server, or VPN

Cluster Resources

	Cluster	Front-end	# Nodes	# Cores/Node	# Total Cores	Memory/Node (GB)	
HTC	vortex	vortex	31	12	372	32/128	ALL
HTC	hima	bora	7	32*	224	256	
HTC	gust	gust	2	128	256	512	
HPC	bora	bora	49	20	980	128	
HTC	meltemi	meltemi	88	64	5632	192	AUTH
HPC	cyclops	cyclops	30	24	720	32	
HPC	femto	femto	30	32	960	96	
HPC	james	james	26	20	520	64	VIMS
HTC	potomac	chesapeake	13	12	156	32	
HTC	pamunkey	james	2	64	128	256	

High Throughput Computing (HTC) – serial or small core count jobs
High Performance Computing (HPC) – parallel jobs with 10's-100's cores

VIMS has their own policies: email hpc@vims.edu if need more than 10 nodes for more than 48 hrs.

- W&M on .sciclone.wm.edu network (.e.g.: vortex.sciclone.wm.edu)
- VIMS on .hpc.vims.edu network (e.g.: james.hpc.vims.edu)

Main campus cluster policy:
ALL – resources available for all
AUTH – must get authorization
VIMS – must be a VIMS fac/staff/student

Connecting to W&M HPC

Must connect via Secure Shell Client (ssh)

- Linux / Mac can use built-in ***terminal / console***
- Windows – powershell/run or ***SSH Secure Shell Client / PuTTY***
- **ssh keys are allowed**

Am I on or off campus?

- On campus: can ssh right into HPC servers

```
15 [mycomputer] ssh vortex.sciclone.wm.edu
```

- Off campus, must jump through bastion host

```
15 [mycomputer] ssh -J bastion.wm.edu vortex.sciclone.wm.edu
```

Is my username the same on my local machine?

If it is different use: ssh <username>@<host>.<domain>

Do I need graphics?

If yes, must log in with –Y

Linux command-line / Text editors

Web documentation: https://www.wm.edu/offices/it/services/researchcomputing/using/prereqs/index.php

Intro to Linux slides / recording: https://www.wm.edu/offices/it/services/researchcomputing/using/tutorials/index.php

Common Linux Commands

cd – change directory
ls – list files
ls -l – list files (long)
cp – copy file
mv – move file
pwd – print working directory
mkdir – make directory
rmdir – remove directory
df – list disks
ps -fu <user> – list <user> process ids
kill <pid> – kill processes
cat <file> – dump file to screen
less <file> – page file on screen

Text Editors

nano : easiest / least powerful

https://www.howtogeek.com/42980/the-beginners-guide-to-nano-the-linux-command-line-text-editor/

vi/vim : advanced / powerful

emacs : advanced / powerful

vscode on local computer :
supports ssh for remote editing

- Linux session is called a "shell"
- '.' means current directory ; '..' parent directory
- '~' mean home directory
- <tab> for file-completion
- <up-arrow> mean "last command"
- bash/tcsh common shell flavor

All users get **tcsh** by default on W&M cluster / can request **bash**

Environment Modules / Software

Sets up environment for particular software packages

```
12 [vortex] module list
Currently Loaded Modulefiles:
  1) modules          2) maui/r156-GRes   3) torque/6.1.1.1   4) intel/2018       5) intel/2018-mpi
13 [vortex] module avail

------------------------------------------- /usr/local/Modules/versions --------------------------------------------
--
3.2.10 3.2.6

------------------------------------------ /usr/local/Modules/modulefiles ------------------------------------------
--
abcluster/2.0                           matlab/R2020a
ampl/20210531(default)                  maui/r156-GRes
anaconda2/2019.10                       mumps/5.4.0/gcc-4.8.5
anaconda3/2020.02                       mvapich2-ib/2.3.1/gcc-9.3.0
anaconda3/2021.05                       mvapich2-ib/2.3.1/intel-2018
anaconda3/2021.05n                      nciplot/4.0
boost/1.74.0/gcc-9.3.0                  nco/4.9.3-intel-2018_intel-mpi
.                                       .
.                                       .
.                                       .
```

Can change modules on demand: module load/unload
Also list what the module sets: module show
Can even write your own modules to make custom environments

https://www.wm.edu/offices/it/services/researchcomputing/using/modules/index.php

Startup Modules / Environment

In user home directories, there are startup files which control default modules (bash users use .bashrc.XXX)

Start-up file	Controls modules for
.cshrc	all front-ends
.cshrc.vortex	Vortex
.cshrc.el7-xeon	Bora & Hima
.cshrc.gust	Gust
.cshrc.el7-phi	Meltemi
.cshrc.femto	Femto
.cshrc.cyclops	Cyclops
.cshrc.el7-x86_64	James
.cshrc.potomac	Potomac
.cshrc.rhel7-opteron	Pamunkey

Be careful when modifying start-up files

User shell cannot have output when invoked (PBS/Torque)

$PLATFORM variable:

```
11 [vortex]  echo $PLATFORM
vortex
```

This means that startup is controlled by ***.cshrc.vortex*** for ***vortex cluster***

Paths, ENV variables, symbolic links

PATH – environmental variable ; list of directories which are searched for executable

```
22 [vortex] echo $PATH
/sciclone/home10/ewalter/bin:/usr/local/intel-2018/compilers_and_libraries_2018/linux/mpi/intel64/bin:/usr/local/inte
l-2018/compilers_and_libraries_2018.5.274/linux/bin/intel64:/usr/local/torque-6.1.1.1/bin:/usr/local/torque-6.1.1.1/s
bin:/usr/local/maui-r156-GRes/bin:/usr/local/Modules/3.2.10/bin:/usr/local/torque-6.1.1.1/bin:/usr/local/torque-6.1.1
.1/sbin:/usr/lib64/qt-3.3/bin:/usr/lib64/ccache:/usr/local/bin:/bin:/usr/bin:/usr/local/sbin:/usr/sbin:/sbin:/opt/ibu
tils/bin:/opt/puppetlabs/bin:/usr/local/bin:/usr/local/sbin:/usr/sbin:/usr/local/intel-2018/parallel_studio_xe_2018.4
.057/bin
```

Current directory (.) is usually not in path – (must type ./<executable>)

Symbolic link – allows an alias for another file/directory

```
23 [vortex] ls -l

lrwxrwxrwx 1 ewalter hpcf    24 Apr  9  2012 data10 -> /sciclone/data10/ewalter
lrwxrwxrwx 1 ewalter hpcf    18 Apr  9  2012 lscr -> /local/scr/ewalter
.
```

ln -s <PATH> <TARGET>

Files & I/O

Web Documentation

https://www.wm.edu/offices/it/services/hpc/using/files/index.php

- There are multiple files-systems available
- Some are for ongoing / project storage

  data, home
- Some are for running jobs (90 day purge)

  scrXX, pscr, /local/scr
- **Only** data/home/home00 backed up
- Use **local scratch** when possible (every node has some)
- Users are responsible for using disk space **responsibly**
- Misue can disturb other jobs / cause **administrative action**
- Don't use home/data10 for writing or large reads
- Use scratch space for jobs
- Lustre (pscr) best practices:

https://www.wm.edu/offices/it/services/researchcomputing/using/files/lustre/index.php

Name	Appropriate for	Backups	Purged	Per-formance
`/sciclone/home` `/ches/home00`	Source code, executables, configuration files, scripts, and small (<100MB total) data files. Unless we have directed you otherwise, **you should not have a job read or write any substantial amount of data to your home directory**, as doing so is extremely likely to impact others' interactive work.	Weeknightly, on-site only	After account expiration.	Low
`/sciclone/data10` `/ches/data10`	Input data files that are needed on an ongoing basis for active projects on the cluster and cannot be easily re-created or re-uploaded. **Please do not have jobs write a substantial amount to data filesystems.** Please use the scratch filesystems (below) for job output unless already given permission from HPC staff.	Weekly, on-site only		

`/local/scr`
`/ches/scr10`
`/sciclone/scr10`
`/sciclone/scr20`
`/sciclone/scr-mlt`
	Scratch space: job outputs and working data that can be easily re-created or re-uploaded, or which will be copied elsewhere for longer-term storage.	**Never**	**Any files not accessed for 90 days**, and after account expiration.	Medium

`/sciclone/pscr`				High

Transferring Files

Web Documentation

https://www.wm.edu/offices/it/services/researchcomputing/using/files/xfers/index.php

Filesystem	Recommended node and Endpoint Display Name
/sciclone/* (New endpoint to be used for all sciclone transfers once Globus V4 is deprecated)	io.sciclone.wm.edu W&M RC/HPC/IO Collection
/sciclone/home	comet.sciclone.wm.edu Sciclone Cluster at W&M, home (Comet)
/sciclone/pscr	bora.sciclone.wm.edu Sciclone Cluster at W&M, pscr (Bora)
/sciclone/data10	tempest.sciclone.wm.edu Sciclone Cluster at W&M, data10 (Tempest)
/sciclone/scr10	polar.sciclone.wm.edu Sciclone Cluster at W&M, scr10 (Polar)
/ches/home00	james.hpc.vims.edu VIMS Cluster at W&M, home00 (James)
/ches/data10	choptank.hpc.vims.edu VIMS Cluster at W&M, data10 (Choptank)
/ches/scr10	rappahannock.hpc.vims.edu VIMS Cluster at W&M, scr10 (Rappahannock)

Globus - https://www.globus.org/
We have endpoints for all file-systems

Each file-system has a server that runs it
For direct access you are **STRONGLY** encouraged to use the recommended node

e.g. : Logged into **comet**; cd'd into data10 ; transfer off-site

Do this from **tempest** since files won't have to hop through **comet** to get off-site.

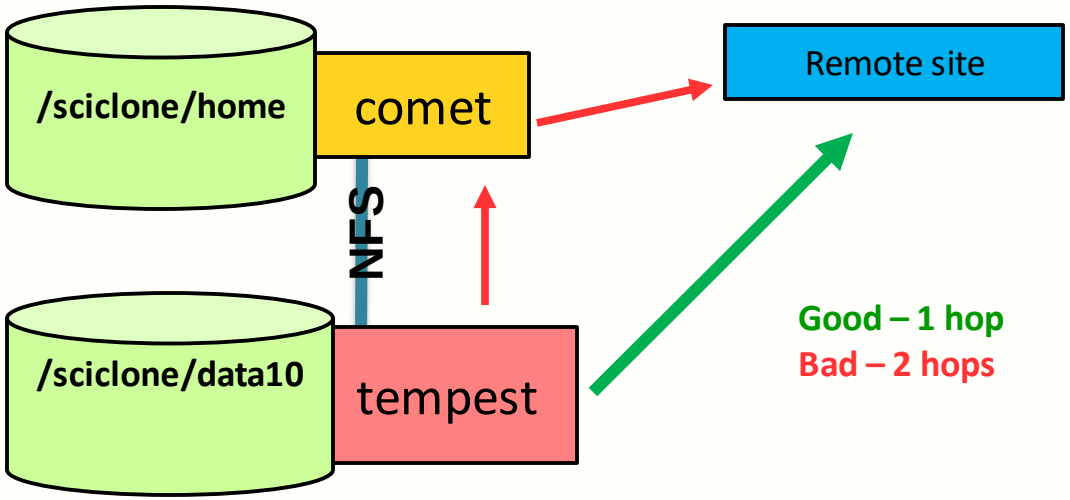

Permissions / sharing files

Want to allow users in the VASP group to read my results file (out)

```
44 [vortex] pwd                                                    (where am I?)
/sciclone/home10/ewalter

45 [vortex] ls -ld results                                  (long list just <dir>)
drwx------ 2 ewalter hpcf 4096 May  6 12:27 results

                                                  (long list <dir> contents)
46 [vortex] ls -l results
total 28
-rw------- 1 ewalter hpcf 25905 May  6 12:27 out

47 [vortex] groups ewalter                          (what group am I in?)
ewalter : hpcf wheel hpcstaff hpcadmin sysadmin www seadas
vasp wm wmall hugepage

                                                                (change group)
48 [vortex] chgrp vasp -R results/

49 [vortex] ls -ld results
drwx------ 2 ewalter vasp 4096 May  6 12:27 results

50 [vortex] ls -l results
total 28
-rw------- 1 ewalter vasp 25905 May  6 12:27 out

                                                                (change group
51 [vortex] chmod g+rX -R results                               permissions
                                                                recursively)
52 [vortex] ls -ld results
drwxr-x--- 2 ewalter vasp 4096 May  6 12:27 results

53 [vortex] ls -l results
total 28
-rw-r----- 1 ewalter vasp 25905 May  6 12:27 out
```

```
drwx------ 2 ewalter hpcf 4096 May  6 12:27 results
```

user group other — user — group

d-directory
r-read
w-write
x-execute/enter

pwd – print working directory
groups – print groups that user is in
chgrp – change **group** ownership
chmod – change **permissions**
chown – **change user ownership**

umask – controls default permissions
- change in .cshrc/.bashrc

Whole path needs to be accessible to share!

Software

There are many software packages available on the HPC systems!

Common packages are all available: Python, R, Gaussian16, Matlab, etc.

- Check the modules on a particular cluster with: “module avail”
- Install it yourself
- Email hpc-help@wm.edu

We encourage **users to install their own software** in their home directory if possible
We can help install, but we get **LOTS** of requests so try not to abuse
Packages used by multiple users can be considered for installing globally

Compilers and Installing

Web docs: https://www.wm.edu/offices/it/services/researchcomputing/using/compiling/index.php
https://www.wm.edu/offices/it/services/researchcomputing/using/software/index.php

- All clusters are equipped with GNU and Intel compilers – some have older PGI
- All popular flavors of MPI supported (Intel, Mvapich2, OpenMPI)
- Also (for PBS/Torque nodes) special wrapper mvp2run available for easier MPI use:

https://www.wm.edu/offices/it/services/researchcomputing/using/jobs/mvp2run/index.php

Compiler web page lists suggested compiler flags for each node type and best practices guides. It is **extremely** important to check the validity of results.

Don't assume if the job runs correctly, it has correct results!

Batch system & Jobs

Web documentation:https://www.wm.edu/offices/it/services/researchcomputing/using/jobs/index.php

Cluster	Front-end	Max Walltime (hrs)	Batch System	
vortex	vortex	180	PBS/Torque	ALL
hima	bora	72	PBS/Torque	
bora	bora	72	PBS/Torque	
gust	gust	72	SLURM	
meltemi	meltemi	72	PBS/Torque	AUTH
cyclops	cyclops	72	SLURM	
femto	femto	72	SLURM	
james	james	72	PBS/Torque	VIMS
potomac	chesapeake	180	PBS/Torque	
pamunkey	james	180	PBS/Torque	

Front-end – same architecture as **cluster** used for compiling, preparing jobs, analyzing data, etc.

Only newer resources use SLURM (gust)

PBS/Torque vs. **SLURM** ; both share same concepts
Commands are different

PBS/Torque clusters -
Vortex/hima/bora all share batch system
e. g. can submit/query jobs from
bora front-end to vortex/bora/hima cluster
vortex front-end to bora vortex/bora/hima cluster
Same with james/potomac/pamunkey
Meltemi uses a separate batch system from these two

DON'T SUBMIT JOBS longer than 72hrs to hima,bora, or james – they won't start running.

SLURM clusters -
Submit from cluster front-end only

Batch system: Interactive Jobs

Cluster	Front-end	Max Walltime (hrs)	node spec
vortex	vortex	180	vortex,c18a,c18b
hima	bora	72	Hima,hima:gpu,hima:nogpu
bora	bora	72	bora
gust	gust	72	N/A
james	james	72	james
potomac	chesapeake	180	potomac
pamunkey	james	180	pamunkey

special node subsets

c18a – 32GB nodes
c18b – 128GB nodes
hima:gpu – want node with gpu
hima:nogpu – want node without gpu

For PBS/Torque nodes are selected via the ***node spec***

qsub – submits the job to the batch system

```
1 [vortex] qsub -I -l walltime=30:00 -l nodes=1:vortex:ppn=1
qsub: waiting for job 10274099 to start
qsub: job 10274099 ready

2 [vx03] python prog.py
```

SLURM front-ends are tied to a single cluster --- no need for node-spec.

Use ***salloc*** (***srun*** on on femto,cyclops)

All slurm clusters have 72hr max walltime

```
1 [gust] salloc -N1 -n1
salloc: Granted job allocation 388
salloc: Waiting for resource configuration
salloc: Nodes gt01 are ready for job

2 [gt01]$ python prog.py
```

Interactive **job puts you on a node ready to work**

Batch Jobs PBS/Torque

You can also submit a ***batch*** job which does not run interactively
First you must write a ***batch script***:

```
34 [vortex] cat run
#!/bin/tcsh
#PBS -N test
#PBS -l nodes=1:vortex:ppn=1
#PBS -l walltime=0:10:00

cd $PBS_O_WORKDIR

python prog.py >& prog.out
```

#!/bin/tcsh *interpret the following in **tcsh** syntax*
-N *name of the job*
-l *job specifications (walltime ; nodespec)*

PBS/Torque start session on node in user home directory!
cd $PBS_O_WORKDIR *cd to where I submitted the job*

```
35 [vortex] qsub run
```

most widely used batch commands
qsub <script> – submit job / interactive job
qdel <jobid> – delete job
qstat – list jobs
qsu – list my jobs

Generates *test.o<jobid>* and *test.e<jobid>*
- batch stdout & stderr output files

Batch Jobs SLURM

Just like any batch system, you must write a ***batch script***:

```
22 [gust] cat run

#!/bin/tcsh
#SBATCH -J test
#SBATCH -N 1 -n 1
#SBATCH -t 0:10:00

python prog.py >& prog.out
```

#!/bin/tcsh	*interpret the following in **tcsh** syntax*
-J	*name of the job*
-N	# nodes
-n	# cores
-t	walltime hh:mm:ss

SLURM starts session on node in ***submission directory***

```
24 [gust] sbatch run
```

Generates *slurm-<jobid>.out*
- batch stdout & stderr output file

most widely used batch commands
sbatch <script> – submit job script
scancel <jobid> – delete job
squeue – list jobs
squeue -u <userid> – list my jobs

Batch system & Jobs Example

PBS/Torque MATLAB example

```
107 [vortex] cat run
#!/bin/tcsh
#PBS -N test
#PBS -l nodes=1:vortex:ppn=12
#PBS -l walltime=12:00:00
#PBS -q matlab

cd $PBS_O_WORKDIR
module load matlab

matlab  -nodisplay -r "readMatrix" >& OUT
```

must add -q matlab for matlab jobs
(PBS/Torque only)

load matlab module (if needed)
PBS/Torque starts node session with default environment.
If modules are not loaded in module startup file (.cshrc.vortex, e.g.), must load in batch script

redirect stdout and stderr

file for stdout and stderr

- Submit job with **qsub**
- Once job starts, check '**OUT**' for output
- Check batch **stdout/stderr** files for issues - named *test.o<jobid>* and *test.e<jobid>*

108 [vortex] cat readMatrix.m

```
tic
%parpool(8)
syms a b c d;
meshpoints = meshgenerator();
eigfile = fopen('eigfile.txt', 'wt');
count = 1;
.
```

Batch system & Jobs Example

SLURM / MATLAB example

```
107 [gust] cat run
#!/bin/tcsh
#SBATCH -J test
#SBATCH –N 1 –n 1
#SBATCH -t 0:10:00

module load matlab

matlab  -nodisplay -r "readMatrix" >& OUT
```

load matlab module (if needed)

SLURM starts node session with current environment and in directory where you typed *sbatch <script>*.
If modules aren't loaded when script is submitted or in module startup file, they must be loaded.

redirect stdout and stderr

file for stdout and stderr

- Submit job with **sbatch**
- Once job starts, check '**OUT**' for output
- Check batch **stdout/stderr** file for issues - named *slurm-<jobid>.out*

```
108 [vortex] cat readMatrix.m
tic
%parpool(8)
syms a b c d;
meshpoints = meshgenerator();
eigfile = fopen('eigfile.txt', 'wt');
count = 1;
.
```

Getting more help

HPC webpage:	https://www.wm.edu/it/rc
HPC ticket system:	mail: *hpc-help@wm.edu*

Using the ticket system is useful since it is monitored by 5 of us

Thank you!