

# Linux and the HPC Environment



**Eric J. Walter**

Executive Director of Research Computing  
September 28th, 2023

# 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**

2023 September						
SUNDAY	MONDAY	TUESDAY	WEDNESDAY	THURSDAY	FRIDAY	SATURDAY
					1	2
3	4	5	6	7	8	9
10	11	12	13	14 1.	15	16
17	18	19	20	21 2.	22	23
24	25	26	27	28 3.	29	30

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

# SSH/SCP/SFTP

*ssh* – is the standard app for connecting to a remote computer with Linux

*scp* – is the standard app for copying files from one site to another

*sftp* – more convenient for multiple hops

Both are easy within Linux. Mac also supports both via *terminal*. Windows users can use *powershell* to do this or use ssh client program (*putty*) and a separate scp client (*WinSCP*).

ssh to bora

- *ssh ejwalt@bora.sciclone.wm.edu*

ssh to bora through bastion host

- *ssh -J ejwalt@bastion.wm.edu ejwalt@bora.sciclone.wm.edu*

scp to bora

- *scp file1 ejwalt@bora.sciclone.wm.edu:*

scp from bora

- *scp ejwalt@bora.sciclone.wm.edu:file1 .*

sftp double hop

- *sftp -J ejwalt@bastion.wm.edu ejwalt@bora.sciclone.wm.edu*

See <https://www.wm.edu/offices/it/services/researchcomputing/using/connecting/>

- connecting/logging in to HPC

<https://www.wm.edu/offices/it/services/researchcomputing/using/xfers/>

- transferring files to/from HPC

# Globus for File Transfers

**Globus is another possibility:** <https://www.globus.org/data-transfer>

**Globus connect personal:** <https://www.globus.org/globus-connect-personal>

- Can use Globus to transfer files to/from and within clusters (Main and VIMS campus).
- Using Globus Connect Personal - attach your own computer to the Globus site and use the Globus GUI
- Available for Windows, Mac and Linux

HPC guide: <https://www.wm.edu/offices/it/services/researchcomputing/using/xfers/globus/personal/>

Recent talk on using Globus in HPC: <https://www.wm.edu/offices/it/services/researchcomputing/using/tutorials/>

# File/Folder Permissions

**umask** – controls the permissions set for new files. The default umask on all HPC systems is:

```
1 [bora] umask
0077
2 [bora] umask -S
u=rwx,g=,o=
```

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

*Annotations:*  
- Brackets under 'drwx' label 'user' (green)  
- Brackets under '---' label 'group' (green)  
- Brackets under '-----' label 'other' (green)  
- 'ewalter' is labeled 'user' (blue)  
- 'hpcf' is labeled 'group' (blue)

So, all new files get **u=rw**, and folders get **u=rwx** and nothing else set.

Other reasonable umasks are (set in .cshrc/.bashrc at end)

**027 (u=rwx, g=rx, o=)** - members of group can read and enter folders

**022 (u=rwx, g=rx, o=rx)** - everyone can read and enter folders

## Must have permissions correct the whole way through a path:

```
1 [bora] pwd
/sciclone/home/ewalter
2 [bora] ls -ld /sciclone
drwxr-xr-x. 20 root root 279 Apr 22 15:49 /sciclone
3 [bora] ls -ld /sciclone/home
drwxr-xr-x. 89 root root 4096 May 11 19:04 /sciclone/home
4 [bora] ls -ld /sciclone/home/ewalter
drwxr-x---. 565 ewalter hpcf 94208 Jun 21 18:43 /sciclone/home/ewalter
5 [bora] ls -ld /sciclone/home/ewalter/file1
-rw-rw-rw-. 1 ewalter hpcf 0 Jun 21 18:43 /sciclone/home/ewalter/file1
```

**"other" blocked from 'ewalter' folder**



# File/Folder Permissions – cont.

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

Annotations:   
 - **user** (green) under 'd'   
 - **group** (green) under 'rwx'   
 - **other** (green) under '-----'   
 - **user** (blue) under 'ewalter'   
 - **group** (blue) under 'hpcf'

Main reasons to change permissions:

- ***I want to give another user a file or folder(s).***

If it is a single file (tar file for multiple files or folders) and it is small enough (< ~1GB) ...

... Just change the permissions on the file to be "world readable" and copy to /tmp (be sure it won't fill up /tmp)

```
cp <file> /tmp  
chmod ugo+r /tmp/<file>
```

- ***I want another user to be able to enter a directory of mine and read/copy any file.***

If you are members of the same group (**groups**), you could make the whole path readable and folders executable by group:

```
chgrp hpcstaff <folder>  
chmod g+rX -R <folder>
```

**Don't forget to check the folders above for correct permissions**  
**Research teams can request secondary groups to be created for them to share**

# Process Control

Control key



Sometimes, you will run a command and it takes too long / you want to kill it for some reason: **^c**  
Sometimes (interactive job session) you want to run a calculation in the "background": **&, jobs, kill**

```
1 [bora] find . -name hello &
[1] 14449

2 [bora] jobs
[1]+  Running                  find . -name hello &

3 [bora] kill %1

4 [bora]
[1]+  Terminated              find . -name hello
```

**&** - puts the calculation in the background

**jobs** – lists current running "jobs"

**kill** – kills the job (%1)

**^c** – cancel

**^z** – background

*Sometimes, you must use "kill -9" (always try without -9 first!)*

# Process Control – cont. 2

You can also move a bg job to the fg and vice/versa:

```
1 [bora] find . -name hello
```

```
^z
```

```
[1]+  Stopped                  find . -name hello
```

```
2 [bora] bg
```

```
[1]+  find . -name hello &
```

```
3 [bora] jobs
```

```
[1]+  Running                  find . -name hello &
```

```
4 [bora] fg
```

```
find . -name hello
```

```
^c
```



# Process Control – cont. 3

"job" control with %n *only works for the shell that launched the process*. Any other shell need to use `ps`.

`ps` – lists the current processes

`ps -ef` (all processes)

`ps -fu ewalter` (processes for ewalter)

Can kill using process id:

```
1 [bora] ps -fu ewalter
```

UID	PID	PPID	C	STIME	TTY	TIME	CMD
ewalter	20926	35637	23	19:46	pts/0	00:00:00	find . -name hello
ewalter	21014	35637	0	19:46	pts/0	00:00:00	ps -fu ewalter
ewalter	35636	35633	0	19:10	?	00:00:00	sshd: ewalter@pts/0
ewalter	35637	35636	0	19:10	pts/0	00:00:00	-bash

```
2 [bora] kill 20926
```

`top` – shows you the current running processes on a computer interactively

# Shell Scripting

- Shell scripting is essential to utilize Linux environment efficiently
- tcsh/csh and bash/sh are two different shell flavors
- All HPC users default to tcsh
- # - comment character
- #! - NOT comment if on first line
- Linux is case sensitive

```
#!/bin/tcsh
# comment :   #! Is a script shell interpreter ("do this with tcsh syntax")
foreach i (1 2 3 4 5)
    echo $i
end
```

```
[68 ewalter@bora ~ ]$ chmod u+x test.csh
```

**Change to executable**

```
[69 ewalter@bora ~ ]$ ./test.csh
```

**. Is not in your PATH, must add it explicitly**

```
1
2
3
4
5
```

# Shell Scripting – cont.

Say you have a file of all the inputs you want to run:

```
1 [bora] cat joblist
1.45
1.44
1.40
1.33
1.10
```

Say you run the job like this:

```
1 [bora] ./a.out -i 1.45

2 [bora] cat runjobs
#!/bin/tcsh
foreach i (`cat joblist`) ← (backtick) means "the result of this command"
    echo $i
    ./a.out -i $i
end

3 [bora] chmod u+x runjobs
4 [bora] ./runjobs
```

**./runjobs will run all parameters in joblist file**

# Shell Scripting – cont. 2

## Main tcsh constructs:

### foreach loop

```
#!/bin/tcsh
foreach i (`cat joblist`)
    echo $i
    ./a.out -i $i
end
```

### If-then-else

```
#!/bin/tcsh
set scen = b
if ($scen == a) then
    echo $scen
else
    echo "wrong value"
endif
```

### while loop

```
#!/bin/tcsh
set i = 1
while ($i < 5)
    echo "i is $i"
    @ i++
end
```

### switch/case statement

```
switch ($myvar)
case 'foo':
    ./foo.csh
    breaksw
case 'bar':
    ./bar.csh
    breaksw
default:
    echo $usage
    breaksw
endsw
```

**Even though you have a tcsh environment, you can still use bash shell scripts  
Bash is considered more powerful for scripting / sometimes easier**

# PBS/Torque batch jobs

## Serial Gaussian16 job:

```
#!/bin/tcsh
#PBS -N Gtest
#PBS -l nodes=1:vortex:ppn=1
#PBS -l walltime=0:60:00
#PBS -j oe

cd $PBS_O_WORKDIR
setenv GAUSS_SCRDIR /sciclone/scr10/ewalter/GAUSS

module load gaussian/g16

g16 test0032.com
```

## Shared Memory Parallel Gaussian16 job:

```
#!/bin/tcsh
#PBS -N Gtest
#PBS -l nodes=1:vortex:ppn=12
#PBS -l walltime=0:60:00
#PBS -j oe

cd $PBS_O_WORKDIR
setenv GAUSS_SCRDIR /sciclone/scr10/ewalter/GAUSS

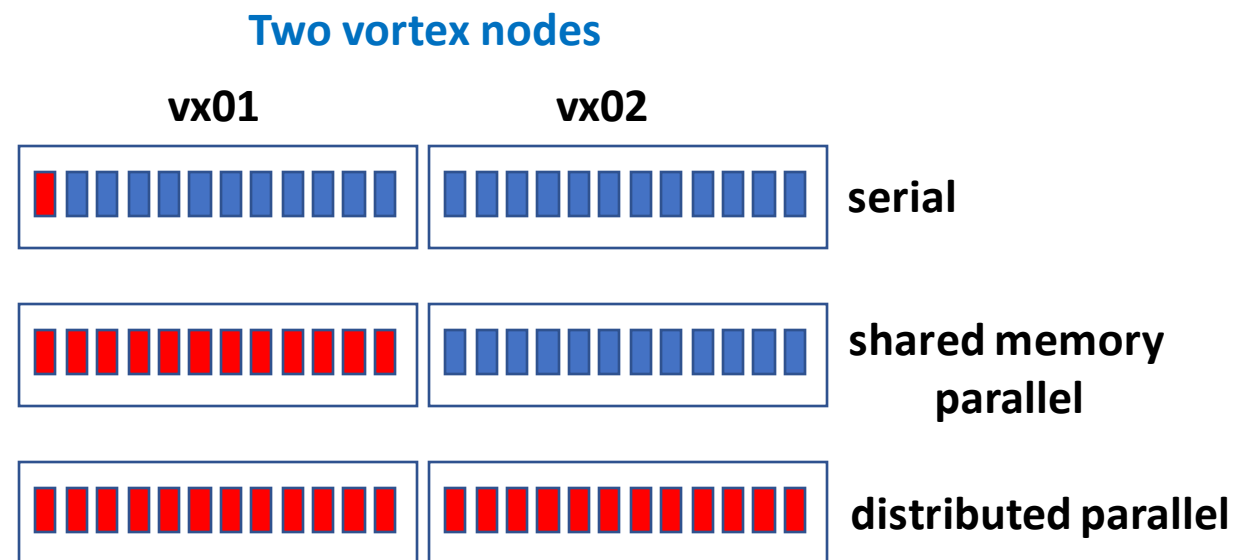
module load gaussian/g16

g16 -p=12 test0032.com
```

Once you have identified/prepared/installed software:

How can it run, serial only?

Shared memory parallel, distributed parallel?



Every code/language/framework is different.  
How parallelism is achieved can vary!

# SLURM batch jobs

## Serial Gaussian16 job:

```
#!/bin/tcsh
#SBATCH -J Gtest
#SBATCH -N 1 -n 1
#SBATCH -t 0:60:00

setenv GAUSS_SCRDIR /sciclone/scr10/ewalter/GAUSS

module load gaussian/g16

g16 test0032.com
```

## Shared Memory Parallel Gaussian16 job:

```
#!/bin/tcsh
#SBATCH -J Gtest
#SBATCH -N 1 -n 12
#SBATCH -t 0:60:00

setenv GAUSS_SCRDIR /sciclone/scr10/ewalter/GAUSS

module load gaussian/g16

g16 -p=12 test0032.com
```

Once you have identified/prepared/installed software:

**How can it run, serial only?**

**Shared memory parallel, distributed parallel?**

### Two vortex nodes

vx01

vx02



serial



shared memory  
parallel



distributed parallel

**Every code/language/framework is different.**

**How parallelism is achieved can vary!**

# PBS/Torque batch jobs – cont.

## Distributed Parallel job:

```
#!/bin/tcsh
#PBS -N cpitest
#PBS -l nodes=3:vortex:ppn=12
#PBS -l walltime=1:00:00
#PBS -j oe

cd $PBS_O_WORKDIR

mvp2run -v ./a.out
```

How to run a job over **multiple processors** varies

Most parallel HPC codes use some flavor of **MPI**

The PBS/Torque clusters (only) have **mvp2run**

**mvp2run** will take your inputs and convert to the flavor of MPI you are using.

Can still run with mpirun, mpiexec, etc. But you may need to construct your own nodefile for some MPI (mvapich2).

# SLURM batch jobs – cont.

## Distributed Parallel job:

```
#!/bin/tcsh
#SBATCH -J cpitest
#SBATCH -N 3 -n 12
#SBATCH -t 1:00:00
```

```
srun ./a.out
```

How to run a job over **multiple processors** varies

Most parallel HPC codes use some flavor of **MPI**

ALL parallel executables should be run using **srun** in SLURM



# Help with Common Software

- <https://www.wm.edu/offices/it/services/researchcomputing/using/software/>  
basic info about using: *Matlab, Stata, Gaussian, and Python*
- Try an **interactive job** via batch system to work out job scripts, and workflow
- One thing to remember – your home directory is shared amongst all clusters / front-ends.  
Some software likes to install some files in your home directory.  
**Be careful to not overwrite software installed for one cluster by another.**
- One common example is **R** local packages. Likes to install local packages in home directory.  
If these packages are compiled on one cluster, may not work on another cluster.
- All Python on cluster should use Anaconda environments, not Python virtual environments.  
See software guide for more info.

# Getting more help

HPC webpage:

<https://www.wm.edu/it/rc>

HPC ticket system

mail: [hpc-help@wm.edu](mailto:hpc-help@wm.edu)

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

## Thank you!

# PBS/Torque jobs

## Interactive batch system job

```
27 [vortex] qsub -I -l walltime=30:00 -l nodes=1:vortex:ppn=12
qsub: waiting for job 1552781 to start
qsub: job 1552781 ready

1 [vx01] python prog.py
```

**Vortex** has 2 nodes/24 cores reserved for walltimes of **1hr max 8am-6pm M-F**  
– can be used for interactive tests

**Typical workflow is to test code with interactive job and run a batch job**

Other clusters may be hard to schedule an interactive job.

**pbstop** – bora/james will show open slots

**showstart** – will show when job should start if it were next

# SLURM jobs

## Interactive batch system job

```
1 [gust] salloc -N1 -n1 -t 0:30:00
salloc: Granted job allocation 388
salloc: Waiting for resource
configuration
salloc: Nodes gt01 are ready for job

2 [gt01] python prog.py
```

**shownodes** – will show available slots

**queue --start** – will show when job should start