Intro to HPC @ W&M and VIMS

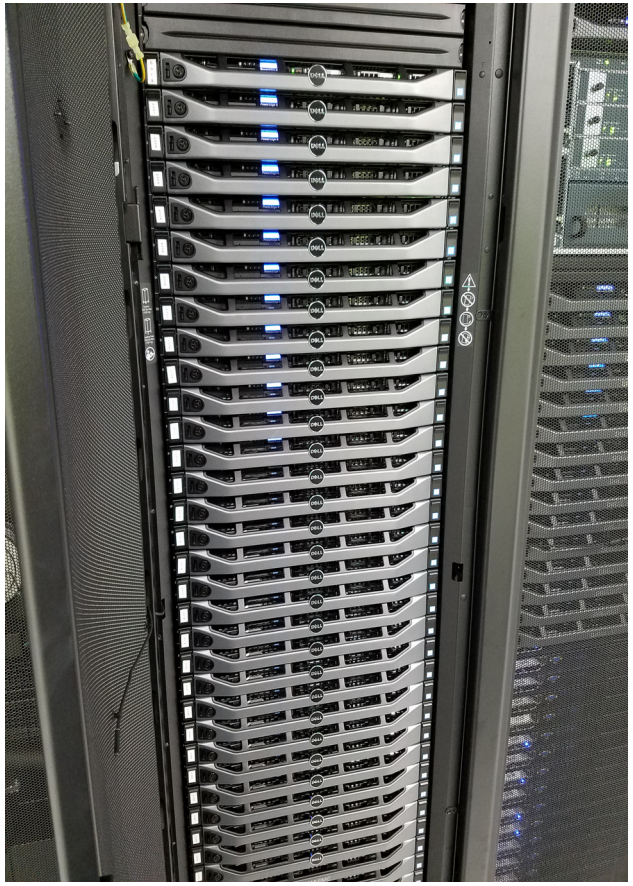

Feb 18th, 2019

Laura Hild
System Engineer

Eric J. Walter
Manager of HPC

Jay Kanukurthy
Applications Analyst

Using HPC

Departments & Offices > IT > Services > HPC > Using

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.

https://hpc.wm.edu → Using HPC

https://hpc.wm.edu **OR**
https://www.wm.edu/offices/it/services/hpc/atwm/index.php

Important topics:

- Getting an account
- Linux command line / text editors
- Logging into the clusters
- Selecting software
- How to use filesystems efficiently
- How to use to use the batch system
- Compiling / installing your own applications

HPC ticket system
mail: *hpc-help@wm.edu*

Cluster Diagram

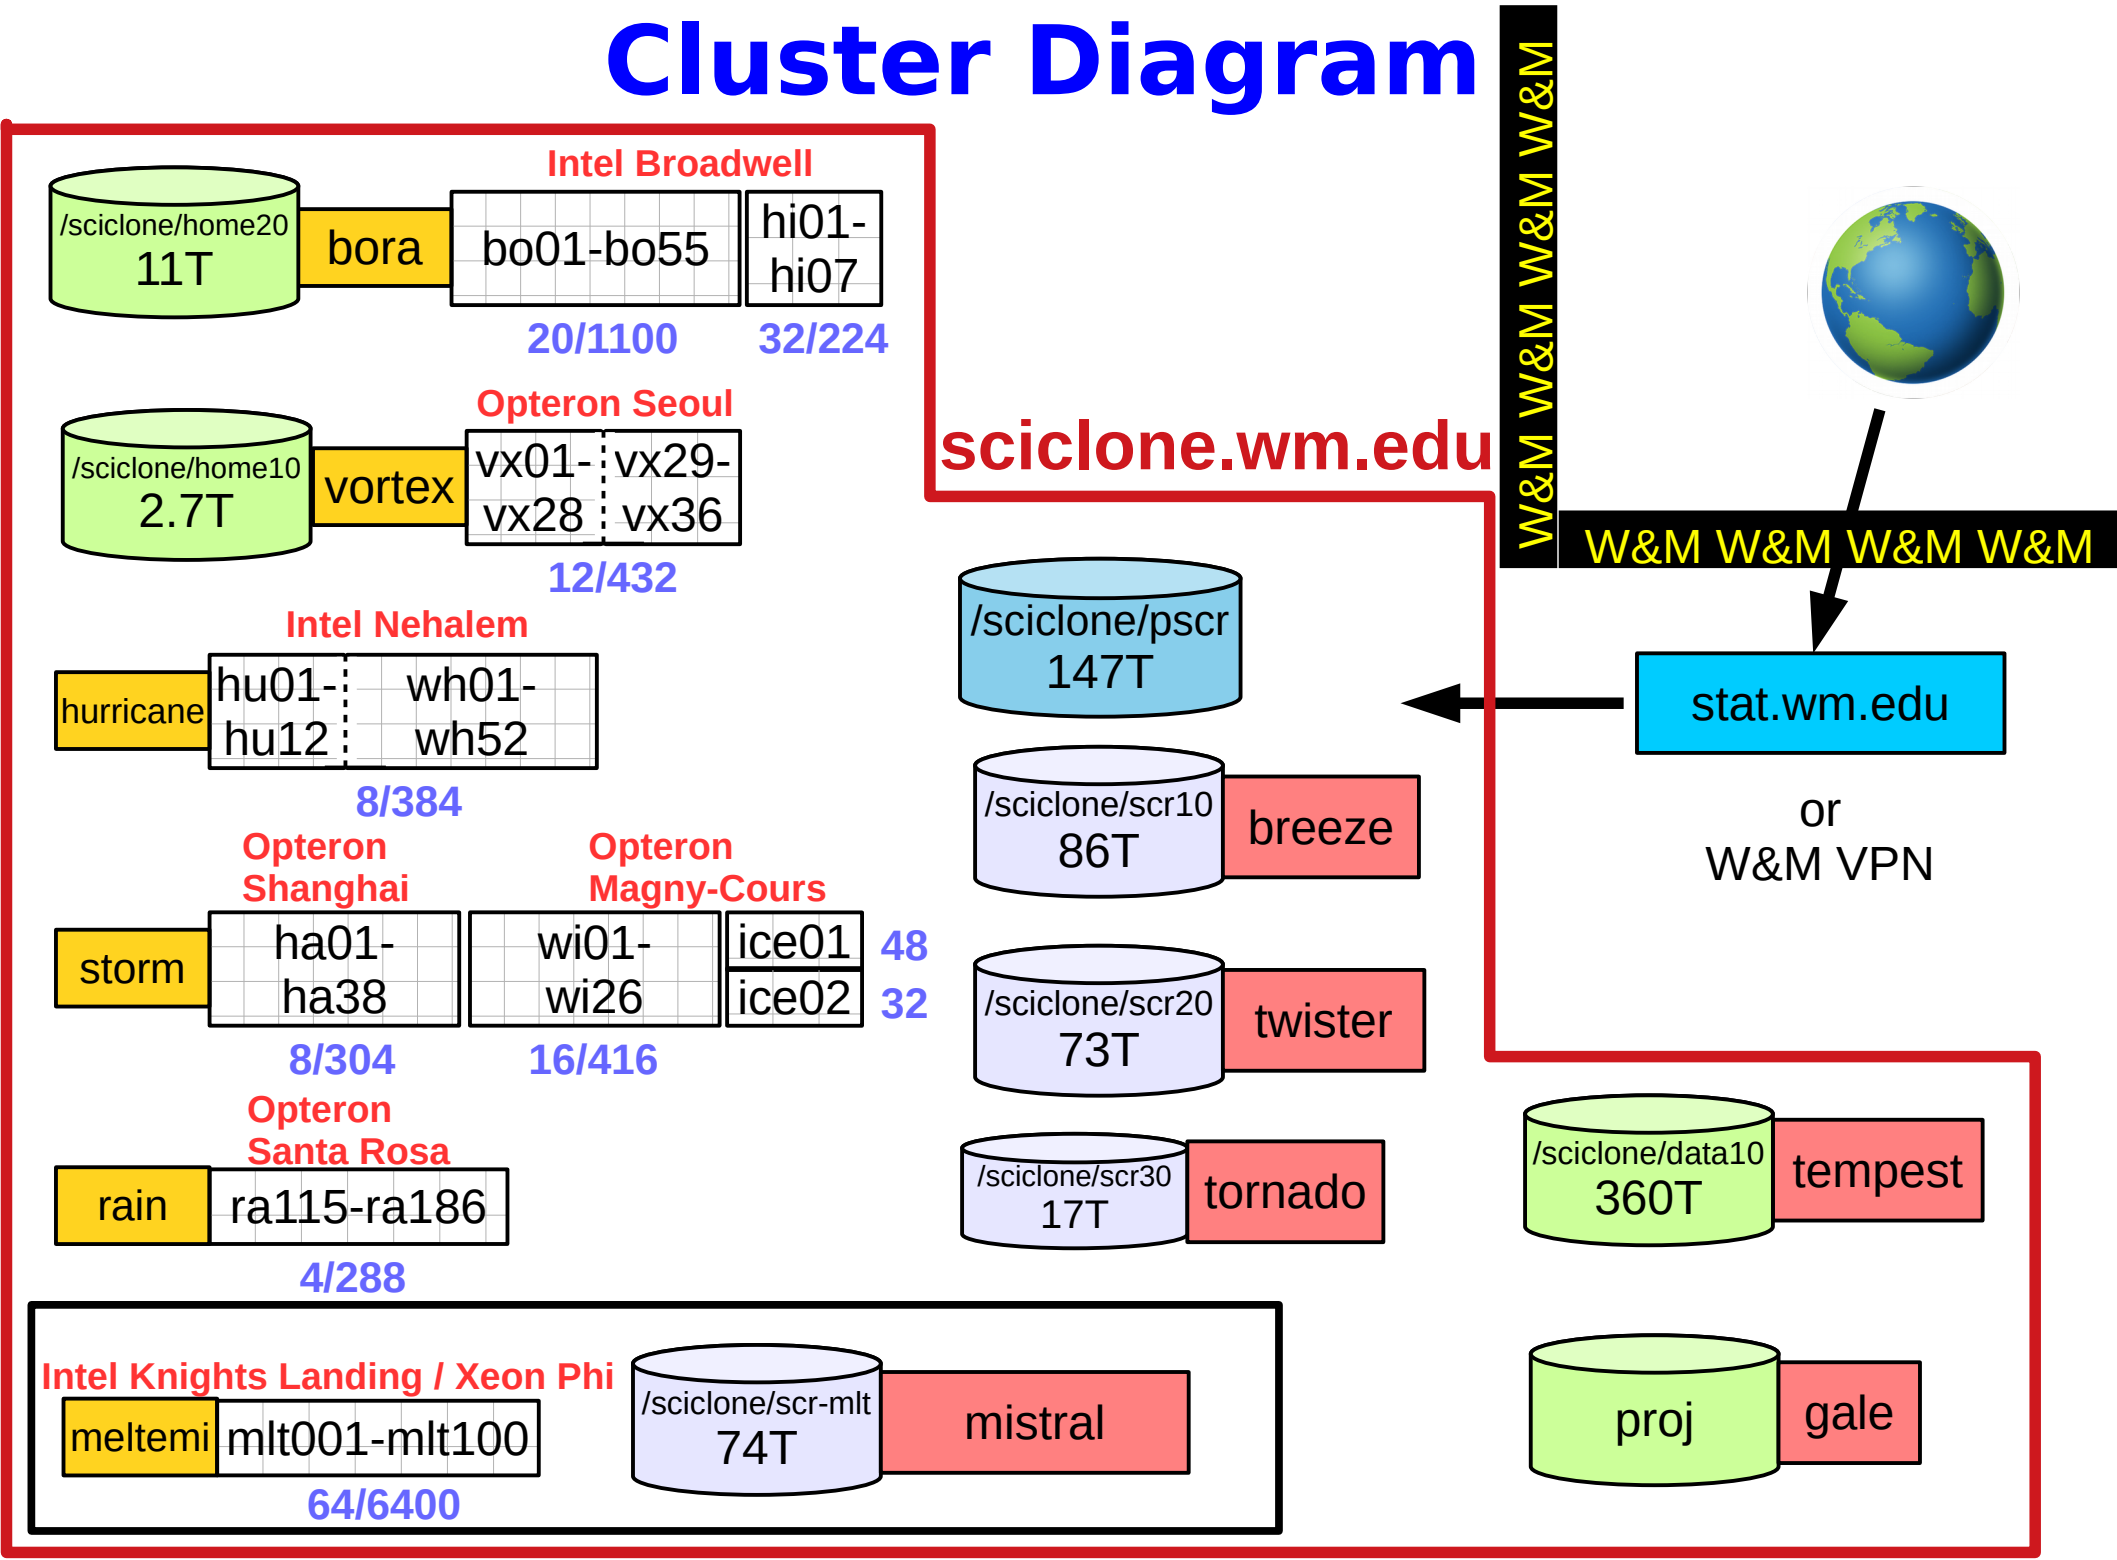

Cluster Diagram

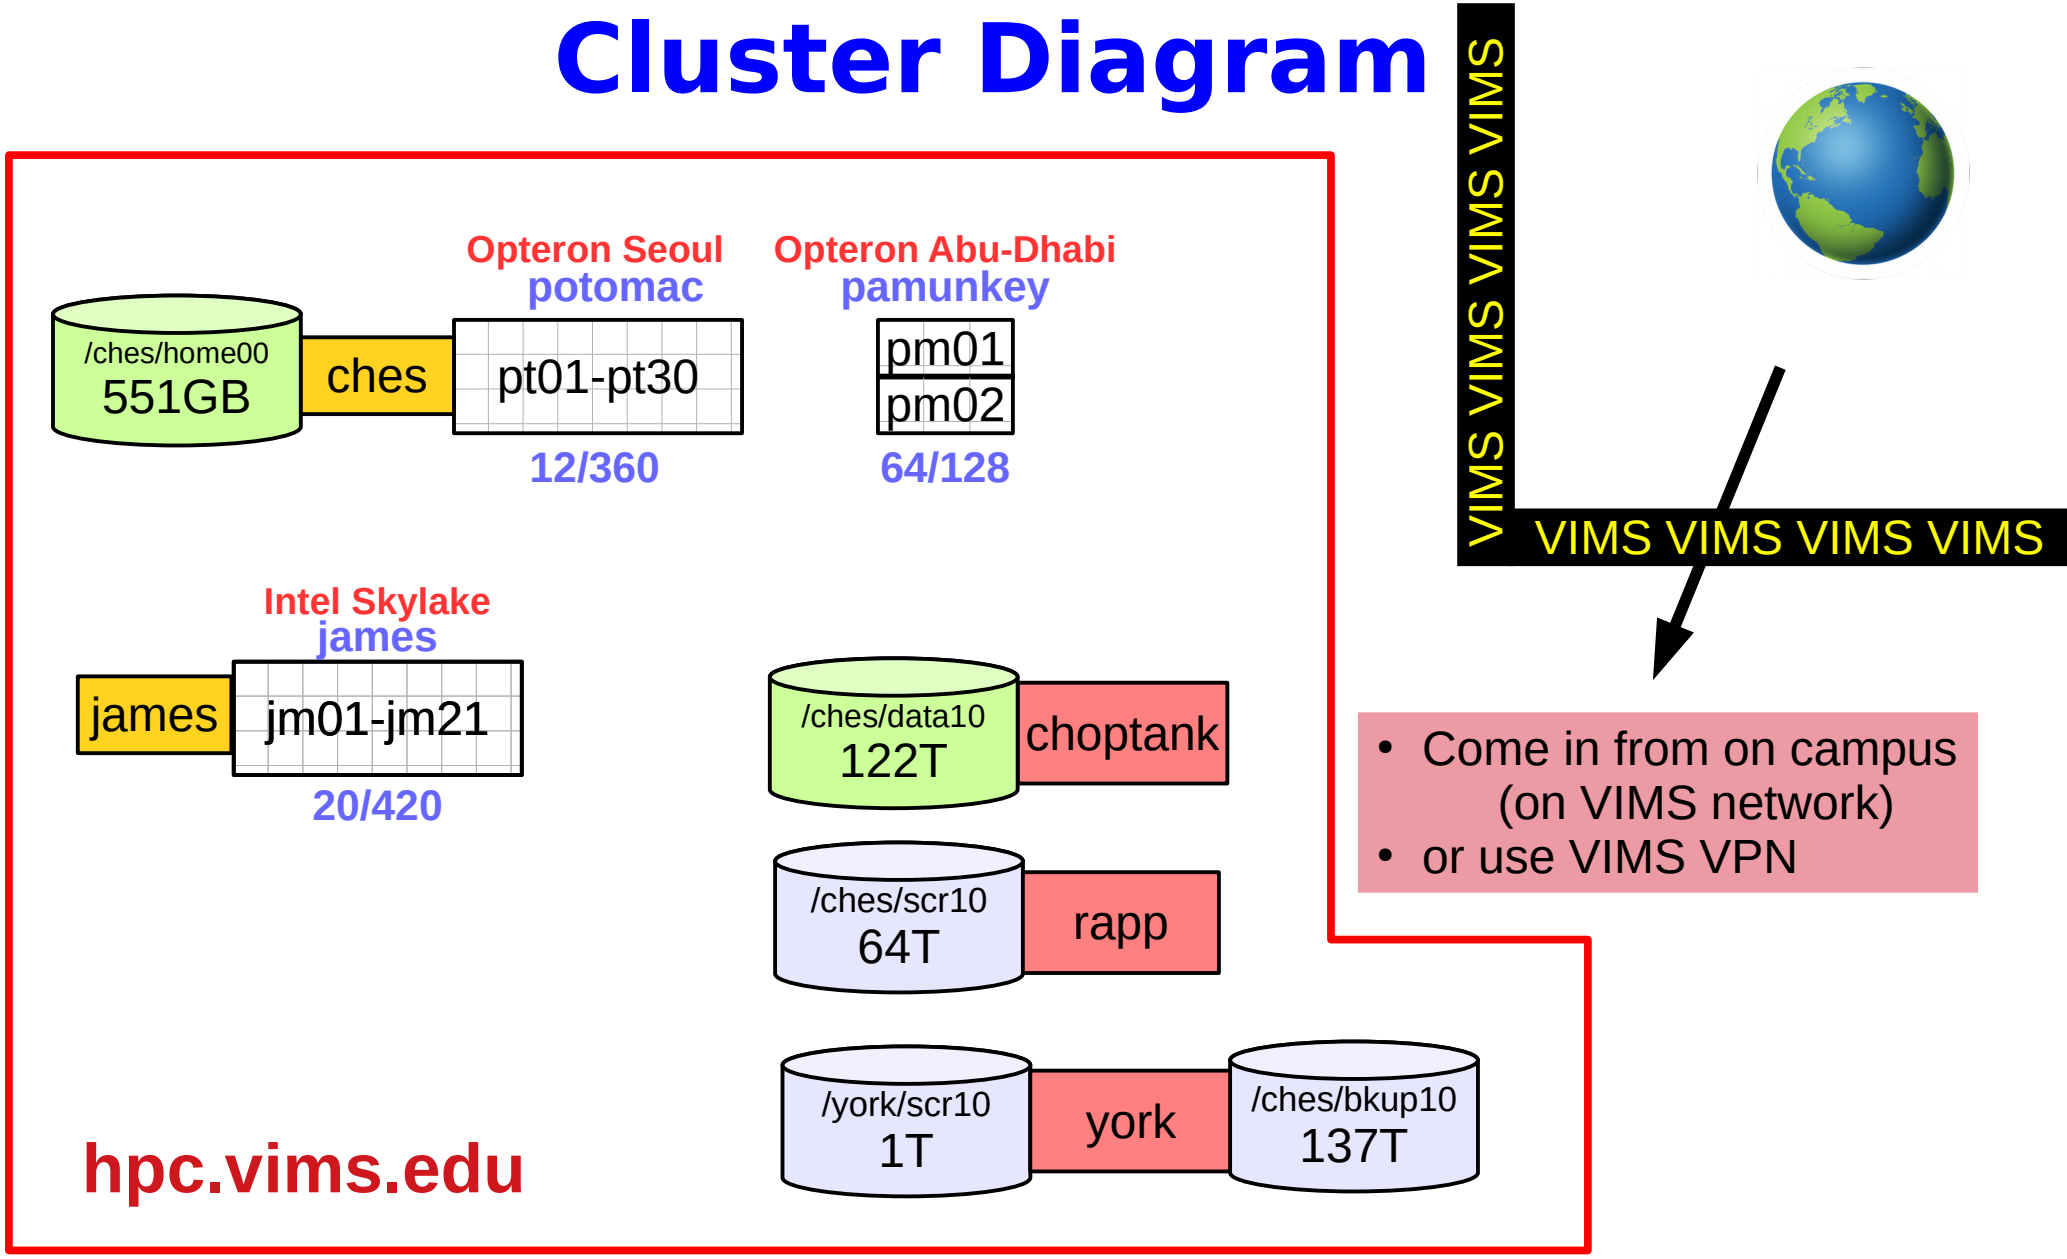

Ches is the front-end for the potomac and pamunkey cluster (both Opteron)
James is the front-end for the james cluster (Intel)

Logging In

Must use Secure Shell client (SSH)

- Linux / Mac built-in (terminal)
- Windows – SSH Secure Shell Client / PuTTY

```
[ejwalt@stat ~]$ ssh ewalter@ches.hpc.vims.edu
Password:
Last login: Fri Feb  8 17:38:29 2019 from vortex.sciclone.wm.edu
------------------------------------------------------------------------
                          Chesapeake Cluster
      College of William & Mary / Virginia Institute of Marine Science
```

Questions to ask yourself:

- **Am I on or off campus?**

If you are off-campus

  - use VIMS VPN to access VIMS network
  - log into *stat.wm.edu* first using your **W&M username and password**

-
- **Is my username the same as my current machine?**

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

- **Do I need graphics?**

If yes, then log in with -X

Start-up Files & Modules

Start-up files control global and cluster specific settings

.cshrc	personal environment settings for all subclusters
.cshrc.storm	for Hail, Ice, and Wind
.cshrc.rhel6-xeon	for Hurricane, Whirlwind, &c.
.cshrc.rhel6-opteron	for Vortex and Potomac
.cshrc.rhel7-opteron	for Pamunkey
.cshrc.el7-xeon	for Bora, Hima, &c.
.cshrc.el7-phi	for Meltemi
.cshrc.el7.x86_64	for James

Using start-up files is the recommended way to select software for particular cluster

$PLATFORM variable:

```
11 [chesapeake] echo $PLATFORM
rhel6-opteron
```

This means that startup is controlled by ***.cshrc.rhel6-opteron*** for ***vortex***

Modules and Start-up

```
23 [chesapeake] module avail

------------------------------------------ /usr/local/Modules/modulefiles --------------------
acml/5.3.1/gcc                  hdf5/1.8.13/intel-2018          netcdf/fortran-4.2/intel-2017
acml/5.3.1/open64               hdf5/1.8.13/pgi                 netcdf/fortran-4.2/intel-2018
acml/5.3.1/pgi                  intel/2016                      netcdf/fortran-4.2/pgi
acml-int64/5.3.1/gcc            intel/2017                      netcdf/fortran-4.4.4/gcc
acml-int64/5.3.1/open64         intel/2018                      netcdf/fortran-4.4.4/intel-2016
acml-int64/5.3.1/pgi            isa/abu-dhabi                   netcdf/fortran-4.4.4/intel-2017
acml-mp/5.3.1/gcc               isa/magny-cours                 netcdf/fortran-4.4.4/intel-2018
acml-mp/5.3.1/open64            isa/seoul                       netcdf/fortran-4.4.4/pgi
acml-mp/5.3.1/pgi               isa/valencia                    numpy/1.8.1
```

```
24 [chesapeake] more .cshrc.rhel6-opteron

module load isa/seoul
module load mvapich2-ib/2.2/intel-2017 intel/2017
```

```
25 [chesapeake] more .cshrc.rhel7-opteron

module load isa/abu-dhabi
module load bio-informatic
```

```
26 [chesapeake] more .cshrc.el7-x86_64

module load intel/2017 mvapich2-ib/2.3/intel-2017
module load netcdf/fortran-4.4.4/intel-2017 netcdf/4.4.1.1/intel-2017
```

Must start new session (log out and back in) to load new start-up modules

http://www.wm.edu/offices/it/services/hpc/using/modules/index.php - online module help

Files and I/O

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

Summary of W&M HPC public user filesystems

Name	Appropriate for	Backups	Purged	Per-formance

`/sciclone/home10`
`/sciclone/home20`
`/ches/home00`
	Source code, executables, configuration files, scripts, and small (<1GB total) data files. Unless you have been directed 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`	Data 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 batch jobs write a substantial amount to data filesystems. Please use the scratch filesystems for job output unless already given permission from HPC staff.**	Weekly, on-site only		
`/sciclone/scr30`	**Scr**atch 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.	

`/local/scr`
`/ches/scr10`
`/sciclone/scr10`
`/sciclone/scr20`
				Medium

`/sciclone/pscr`				High

- There are multiple files-systems available
- Some are for ongoing / project storage
  data, homeXX
- Some are for running jobs
  scrXX, pscr
- **Only** data/homeXX backed up
- Use local scratch when possible (every node has some)
- Users are responsible for using disk space responsibly!!

Misuse of file-systems can disturb other jobs and can result in a administrative action!

Don't use data10 for job writes or large job reads

Use scratch space for jobs (90 day purge)

Transferring Files

Filesystem	Recommended node
/sciclone/aiddata10	gale.sciclone.wm.edu.
/sciclone/baby10	
/sciclone/gluex10	
/sciclone/home10	vortex.sciclone.wm.edu.
/sciclone/home20	bora.sciclone.wm.edu.
/sciclone/pscr	
/sciclone/data10	tempest.sciclone.wm.edu.
/sciclone/scr10	breeze.sciclone.wm.edu.
/sciclone/scr20	twister.sciclone.wm.edu.
/sciclone/scr30	tornado.sciclone.wm.edu.
/sciclone/scr-mlt	mistral.sciclone.wm.edu.
/ches/home00	chesapeake.hpc.vims.edu.
/ches/data10	choptank.hpc.vims.edu.
/ches/scr10	rappahannock.hpc.vims.edu.
/york/scr10	york.hpc.vims.edu.

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

hpc.wm.edu → Using HPC → Files & Filesystems → Transferring Files

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 ches ; cd'd into data10 ; transfer off-site

Do this from choptank since files won't have to hop through ches to get off-site.

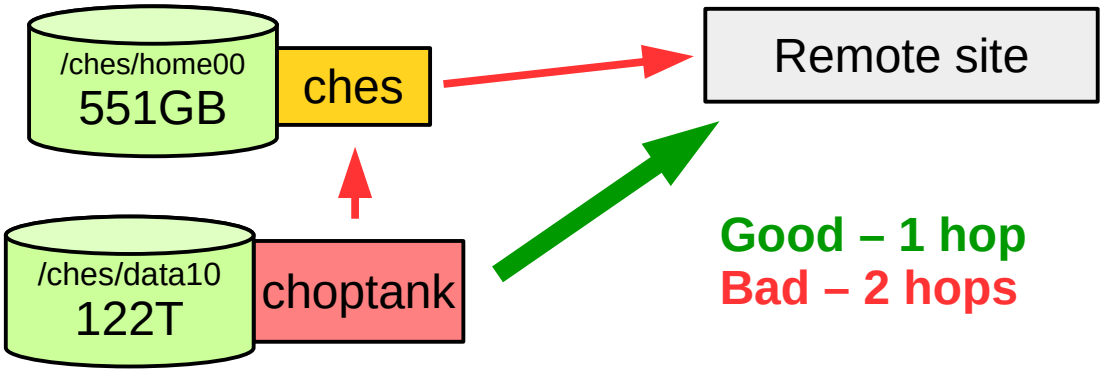

Globus - **https://www.globus.org/**
We have direct endpoints for most file-systems

Sharing Files & Folders : Permissions I

see
https://www.nersc.gov/users/storage-and-file-systems/unix-file-permissions
for more information

```
33 [hurricane] pwd                                    print working directory
/sciclone/home04/ewalter
```

```
34 [hurricane] ls -ld results                         list directory
drwx------ 2 ewalter hpcf 512 Jun 16 14:37 results
```

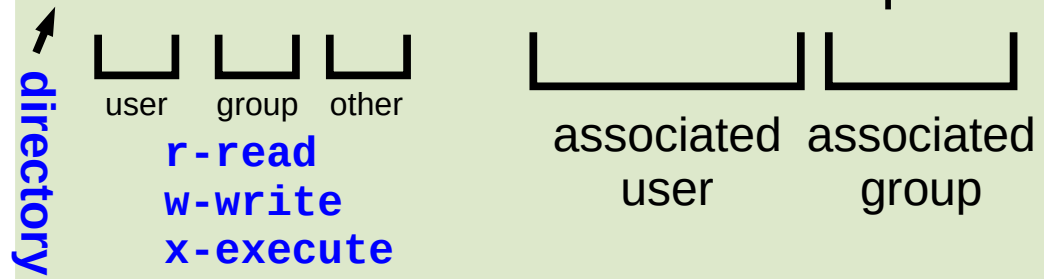

****directories need to be 'x' to be entered/passed through**

```
35 [hurricane] ls -l results                          list files in directory
total 3
-rw------- 1 ewalter hpcf 194 Jun 16 14:37 ww.dat
-rw------- 1 ewalter hpcf 194 Jun 16 14:37 yy.dat
-rw------- 1 ewalter hpcf 194 Jun 16 14:37 zz.dat
```

Common Permissions Tasks I

Change the permissions of a directory **chmod:**

```
[hurricane] ls -ld VASP
drwx------ 4 ewalter hpcf 512 Apr  4  2014 VASP
25 [hurricane] chmod go+rX VASP
26 [hurricane] ls -ld VASP
drwxr-xr-x 4 ewalter hpcf 512 Apr  4  2014 VASP
27 [hurricane] chmod o-rX VASP
28 [hurricane] ls -ld VASP
drwxr-x--- 4 ewalter hpcf 512 Apr  4  2014 VASP
29 [hurricane] chmod g-rX VASP
30 [hurricane] ls -ld VASP
```

Change the permissions of a directory and everything under it **chmod -R**:

```
32 [hurricane] ls -ld VASP
drwx------ 4 ewalter hpcf 512 Apr  4  2014 VASP
33 [hurricane] ls -l VASP
total 52457
-rw------- 1 ewalter hpcf 22932904 Apr  4  2014 potpaw_LDA.52.tar.gz
-rw------- 1 ewalter hpcf 25958479 Apr  4  2014 potpaw_PBE.52.tar.gz
drwx------ 2 ewalter hpcf    15360 Aug  6 00:01 vasp.5.3

34 [hurricane] chmod -R  g+rX VASP

35 [hurricane] ls -ld VASP
drwxr-x--- 4 ewalter hpcf 512 Apr  4  2014 VASP
36 [hurricane] ls -l VASP
total 52457
-rw-r----- 1 ewalter hpcf 22932904 Apr  4  2014 potpaw_LDA.52.tar.gz
-rw-r----- 1 ewalter hpcf 25958479 Apr  4  2014 potpaw_PBE.52.tar.gz
drwxr-x--- 2 ewalter hpcf    15360 Aug  6 00:01 vasp.5.3
```

Use this command if you want to allow group access to your home, scrXX, and dataXX directories

Common Permissions Tasks II

How do I change the initial permissions that files and folders are given when created:

You need to edit your .cshrc file in your home directory and add: **umask**

`umask 077` files get "`-rw-------`" folders get "`drwx------`"
`umask 027` files get "`-rw-r-----`" folders get "`drwxr-x---`"
`umask 022` files get "`-rw-r--r--`" folders get "`drwxr-xr-x`"

What groups are I in?: **groups**

```
52 [hurricane] groups ewalter
ewalter : hpcf wmall hpcstaff www seadas vasp sysadmin wm hpcadmin wheel hpsmh
```

My primary (default) group is `hpcf` and the rest are secondary

Change a group associated with a file or directory: **chgrp**

```
54 [hurricane] ls -ld project
drwx------ 2 ewalter hpcf 512 Aug 18 20:47 project
55 [hurricane] chgrp hpcstaff project
56 [hurricane] ls -ld project
drwx------ 2 ewalter hpcstaff 512 Aug 18 20:47 project
```

Software and Compilers

There are many software packages available on the HPC systems!
Ways to find out whether a package is available

- Check the modules on a particular cluster with: “module avail”
- Look at software web page (**http://hpc.wm.edu/software/**)
- Install it yourself
- Email hpc-help@wm.edu

We encourage users to install their own software in their home directory if possible
We will also do it for you or at least help, but we get LOTS of request so try not to abuse

- Compilers: **Intel**, PGI, GNU
- MPI libraries: Intel, mvapich2,openmpi
- mvp2run – wrapper for all three with extra functionality (node load checking)

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

Using the Batch System

HPC uses Torque (PBS) to schedule and run jobs
Nodes are selected via the ***node type***
qsub – submits the job to the batch system

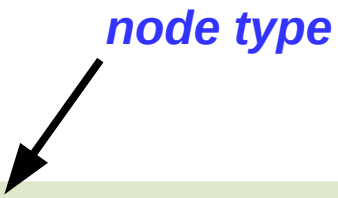


```
11 [chesapeake] qsub -I -l walltime=10:00 -l nodes=1:potomac:ppn=12
qsub: waiting for job 120423.cp00.hpc.vims.edu to start
qsub: job 120423.cp00.hpc.vims.edu ready

1 [pt01]
```

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

There are many *node types*. The default node type is simply the sub-cluster name
potomac, pamunkey, james, vortex, bora, hima, ...

It is also possible to select certain subsets within a cluster or a collection of sub-clusters
x5672 – any hurricane or whirlwind node
c18b – only large memory vortex nodes

See online documentation or send email to hpc-help@wm.edu for more information

https://www.wm.edu/offices/it/services/hpc/hw/nodes/index.php
https://www.wm.edu/offices/it/services/hpc/using/jobs/index.php

DO NOT RUN JOBS ON FRONT-END/LOGIN MACHINES!

Using the Batch System II

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

```
34 [chesapeake] cat run
#!/bin/tcsh
#PBS -N test
#PBS -l nodes=1:potomac:ppn=12
#PBS -l walltime=0:10:00
#PBS -j oe

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)*
-j	*combine stderr and stdout*
cd $PBS_O_WORKDIR	*cd to where I submitted the job*
./a.out	*run the job*

```
35 [chesapeake] qsub run
```

```
148 [chesapeake] more test.o2785870
Warning: no access to tty (Bad file descriptor).
Thus no job control in this shell.
tput: No value for $TERM and no -T specified
```

most widely used batch commands
qsub – submit job
qdel – delete job
qstat – list jobs
qsu – list my jobs

Using the Batch System III

MATLAB example

```
107 [chesapeake] more run
#!/bin/tcsh
#PBS -N test
#PBS -l nodes=1:potomac:ppn=8
#PBS -l walltime=12:00:00
#PBS -j oe
#PBS -q matlab

cd $PBS_O_WORKDIR
module load matlab

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

must add -q matlab for matlab jobs

load matlab module (if needed)

redirect stdout and stderr

file for stdout and stderr

```
108 [chesapeake] head readMatrix.m
tic
%parpool(8)
syms a b c d;
meshpoints = meshgenerator();
eigfile = fopen('eigfile.txt', 'wt');
count = 1;
count2 = 1;
%set(0, 'CurrentFigure', 1);
%plot3(0,0,0,'.');
%grid on
.
.
.
```

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

Where to get help?

HPC webpage:	**http://www.wm.edu/offices/it/services/hpc/atwm/index.php**
HPC ticket system	mail: ***hpc-help@wm.edu***

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

WE'RE HERE TO HELP!

Q&A