Guide to W&M Computational Resources

Eric J. Walter

Director of Research Computing

May 7th, 2020

RC/HPC Staff

Eric J. Walter
Director

Jay Kanukurthy
Software

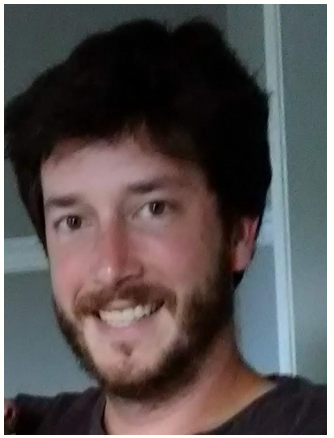

J. Matt Kennedy
Network/File-systems

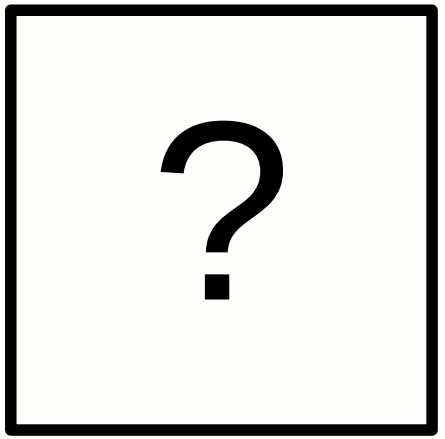

???
Systems Administrator

- 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

(subscribed when you open account)

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

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

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

Cluster Nomenclature

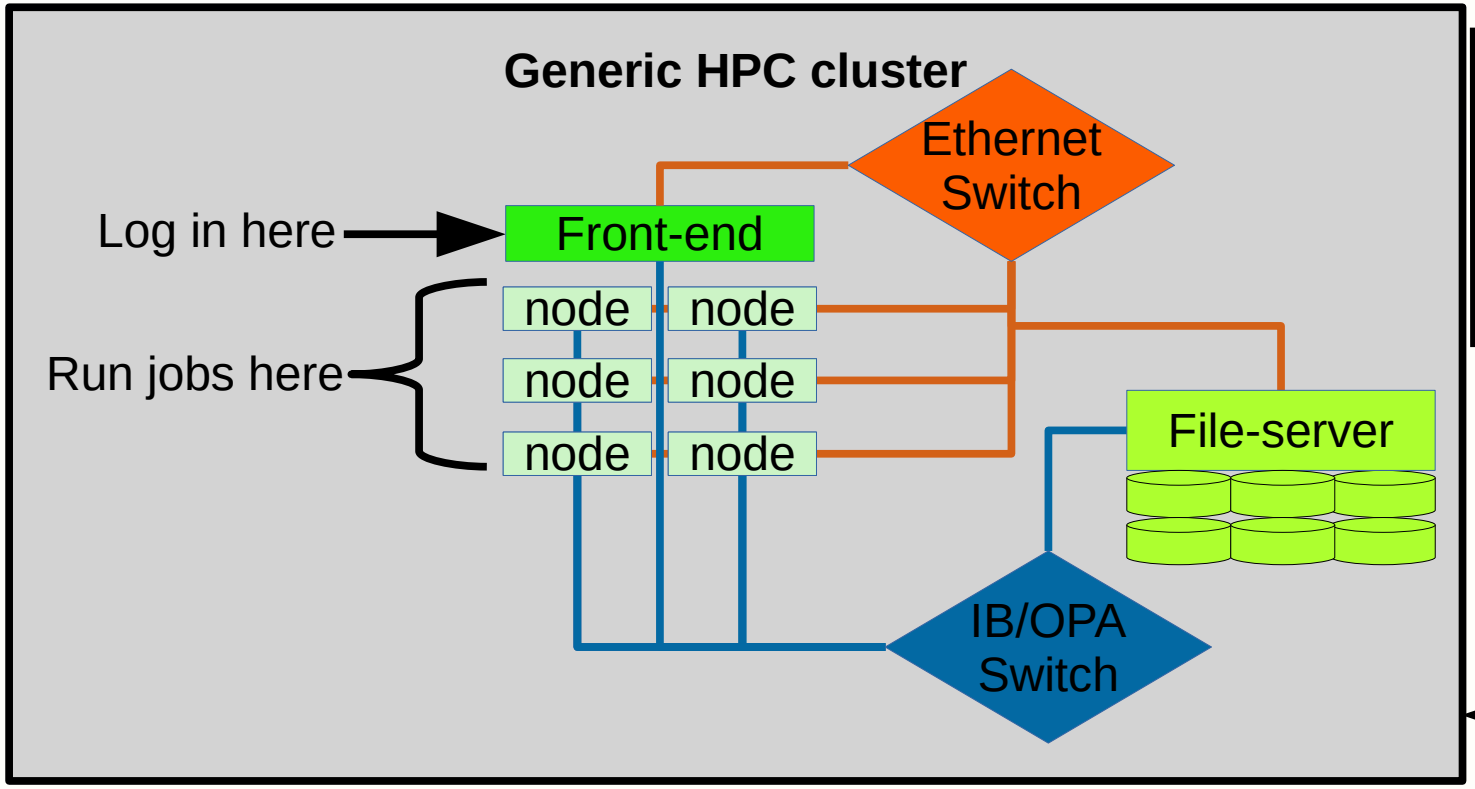

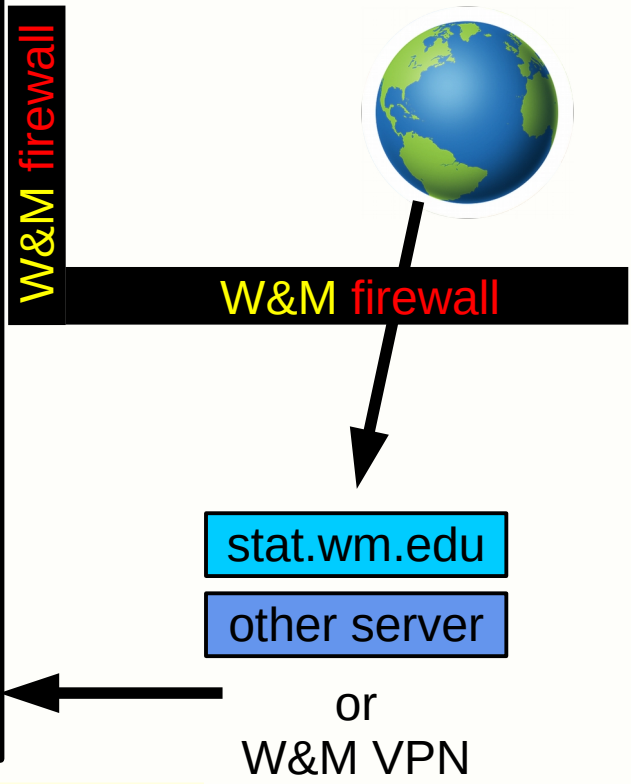

Must get through W&M firewall to get to cluster resources
On campus you are already within firewall
Otherwise use stat.wm.edu, another server, or W&M VPN

Cluster Resources

	cluster	front-end	# nodes	# cores/node	# total cores	memory(GB)	
HTC	vortex	vortex	36	12	432	32/128	OPEN
HTC	hurricane/whirlwind	hurricane	50	8	400	48/64/192	OPEN
HTC	wind/ice	storm	26	16/32/48	384/32/48	32/64/96	OPEN
HTC	hima	bora	7	64	448	256	OPEN
HPC	bora	bora	55	20	1100	128	OPEN
HPC	meltemi	meltemi	100	64	6400	192	AUTH
HPC	femto	femto	30	32	960	96	AUTH
HPC	cyclops	cyclops	30	24	720	32	AUTH

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

OPEN – resources available for everyone
AUTH – must get authorization before using

All on .sciclone.wm.edu network (.e.g.: vortex.sciclone.wm.edu)

Connecting to W&M HPC

Must connect via Secure Shell Client (ssh)

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

Am I on or off campus?

- On campus: can ssh right into HPC servers

```
>> ssh vortex.sciclone.wm.edu
```

- Off campus: hop through stat.wm.edu or another on-campus server

```
>> ssh stat.wm.edu
```

then

```
>> ssh 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 -X

Linux command-line / Text editors

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

vi/vim : advanced / powerful

emacs : advanced / powerful

- Linux session is called a "shell"
- '.' means current directory ; '..' parent directory
- '~' mean home directory
- <tab> for file-completion / CNTRL-D
- <up-arrow> mean "last command"
- sh/bash/csh/tcsh common shell flavors
  All users get tcsh by default on W&M cluster

Environment Modules / Software

Sets up environment for particular software packages

```
[1 ewalter@vortex ~ ]$module list
Currently Loaded Modulefiles:
  1) modules          3) torque/6.1.1.1   5) intel/2018
  2) maui/r156-GRes   4) isa/seoul        6) intel/2018-mpi

[2 ewalter@vortex ~ ]$module avail
------------------------------------ /usr/local/Modules/modulefiles -------------------------------------
abcluster/2.0                   hdf/4.2.10/gcc                  netcdf/3.6.3/gcc-5.2.0
acml/5.3.1/gcc                  hdf/4.2.10/pgi                  netcdf/3.6.3/intel
acml/5.3.1/open64               hdf5/1.8.13/gcc                 netcdf/3.6.3/pgi
acml/5.3.1/pgi                  hdf5/1.8.13/intel               netcdf/3.6.3/pgi-11.10
acml-int64/5.3.1/gcc            hdf5/1.8.13/pgi                 netcdf/3.6.3/pgi-16.3
acml-int64/5.3.1/open64         hyperworks/18                   netcdf/4.3.2/gcc
acml-int64/5.3.1/pgi            hyperworks/19(default)          netcdf/4.3.2/intel
```

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

Web Documentation: **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

Startup file	Sub-cluster
.cshrc.rhel6-opteron	Vortex
.cshrc.rhel6-xeon	Hurricane & Whirlwind
.cshrc.storm	Ice and Wind
.cshrc.el7-xeon	Bora & Hima
.cshrc.el7-phi	Meltemi
.cshrc.femto	Femto
.cshrc.cyclops	Cyclops

Be careful when modifying start-up files

User shell can not have output when invoked

$PLATFORM variable:

```
[11 ewalter@vortex ~ ]  echo $PLATFORM
rhel6-opteron
```

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

Paths, ENV variables, symbolic links

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

```
[5 ewalter@vortex ~ ]$echo $PATH
/sciclone/home10/ewalter/bin:/usr/local/intel-2018/compilers_and_libraries_2018/linux/mpi/intel64/bin:/usr/local/intel-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/sbin:/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/ibutils/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

```
[33 ewalter@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, homeXX
- Some are for running jobs (90 day purge)
  scrXX, pscr, /local/scr
- **Only** data/homeXX 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

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

Transferring Files

Web Documentation

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

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.

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

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

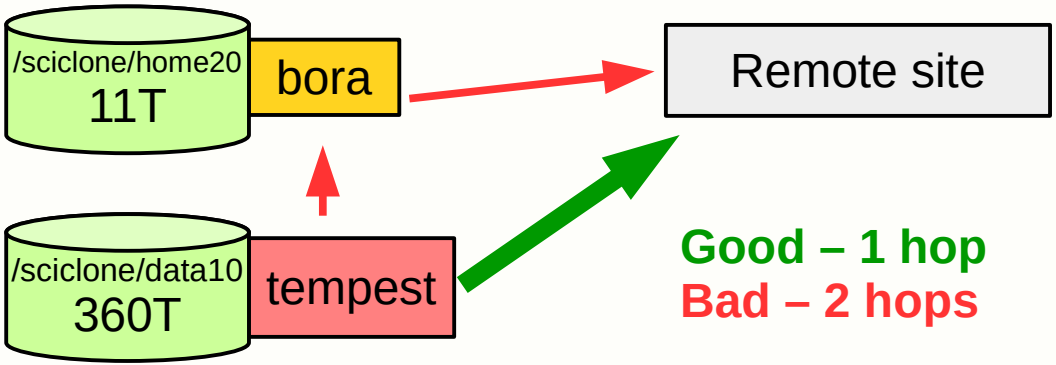

Globus - **https://www.globus.org/**
We have endpoints for most file-systems

Permissions / sharing files

see http://linuxcommand.org/lc3_lts0090.php for more information

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

[45 ewalter@vortex ~ ]$ ls -ld results                        (long list just <dir>)
drwx------ 2 ewalter hpcf 4096 May  6 12:27 results
[46 ewalter@vortex ~ ]$ ls -l results                         (long list <dir> contents)
total 28
-rw------- 1 ewalter hpcf 25905 May  6 12:27 out

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

[48 ewalter@vortex ~ ]$ chgrp vasp -R results/                (change group)

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

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

[51 ewalter@vortex ~ ]$ chmod g+rX -R results                 (change group permissions recursively)

[52 ewalter@vortex ~ ]$ ls -ld results
drwxr-x--- 2 ewalter vasp 4096 May  6 12:27 results

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

[54 ewalter@vortex ~ ]$ ls -ld /sciclone/home10/ewalter
drwxr-xr-x 467 ewalter hpcf 131072 May  6 12:27 /sciclone/home10/ewalter
```

```
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

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”
- 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 and Installing

Web docs: **https://www.wm.edu/offices/it/services/researchcomputing/using/compiling/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 OPEN 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 validity and of results.

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

Batch system & Jobs I

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

sub-cluster	node spec	max walltime (hrs)
vortex	vortex, c18a, c18b, compute	180
hurricane	hurricane, x5672, compute	180
whirlwind	whirlwind, x5672, compute	180
wind	wind,compute	180
ice	ice, compute	180
hima	hima	72
bora	bora	72

All OPEN nodes can submit jobs **to each other**

All “open” nodes use PBS/Torque to run jobs
Nodes are selected via the ***node spec***
qsub – submits the job to the batch system

node spec

```
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
```

JobID

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

Batch system & Jobs II

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)*
cd $PBS_O_WORKDIR	*cd to where I submitted the job*
./a.out	*run the job*

```
35 [hurricane] qsub run
```

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

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

pbstop – graphical display of node usage

Batch system & Jobs III

MATLAB example

```
107 [hurricane] more run
#!/bin/tcsh
#PBS -N test
#PBS -l nodes=1:x5672: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

- Submit job with qsub
- Once job starts, check ‘OUT’ for output
- Check batch stdout/stderr for issues

```
108 [hurricane] head 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 3 of us

Grafana dashboard

Show real-time info on
- File-system I/O traffic
- Server loads & network traffic
- Lustre (pscr) traffic info

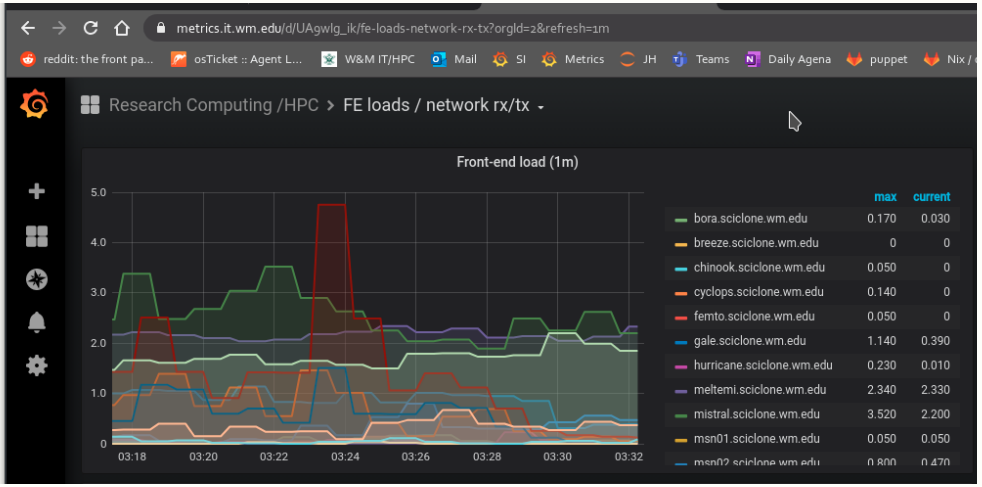

https://metrics.it.wm.edu/d/GIPjU-wmz/network-traffic-for-all-hpc-nfs-file-systems?orgId=3&refresh=1m
https://metrics.it.wm.edu/d/UA9wlg_ik/front-end-file-server-load-and-network-traffic?orgId=3&refresh=1m
https://metrics.it.wm.edu/d/kIGpSmwik/sciclone-pscr-lustre-stats?orgId=3&refresh=1m