HPC Resources at W&M
...and *How To Use Them*

- What resources are available
- How to log into HPC machines
- Setting up your environment
- How to run a job on the cluster
- How to get more help

January 24, 2017

Laura Hild
System Engineer

Eric J. Walter
Manager HPC

Jay Kanukurthy
Applications Analyst

High Performance Computing

Refers to one or more of the following:

Advanced computer architecture
multi-processor / multi-core machines

High-speed networks
Gigabit Ethernet / Infiniband

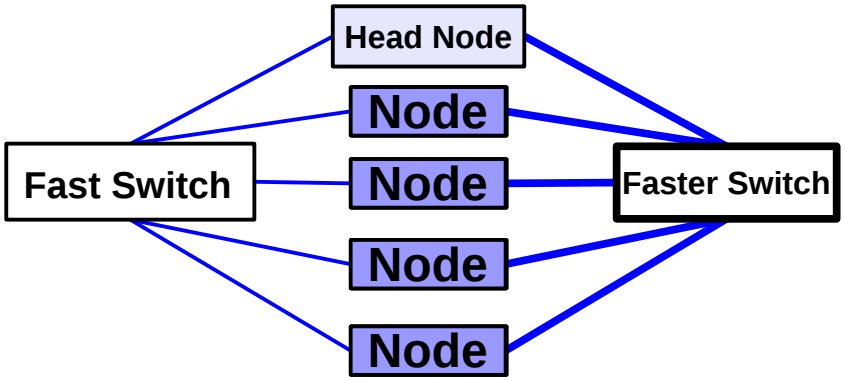

Parallel Algorithms
Use multiple cores and/or nodes

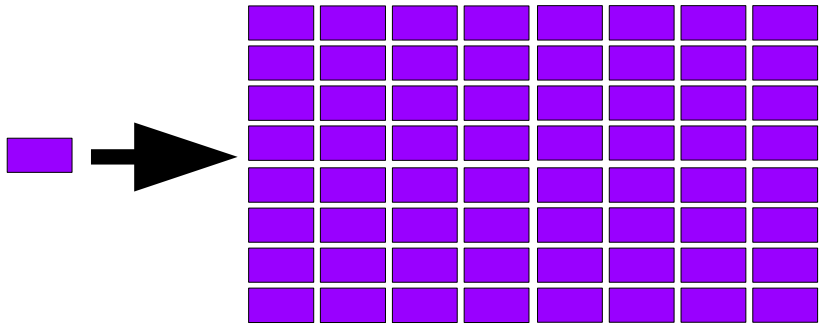

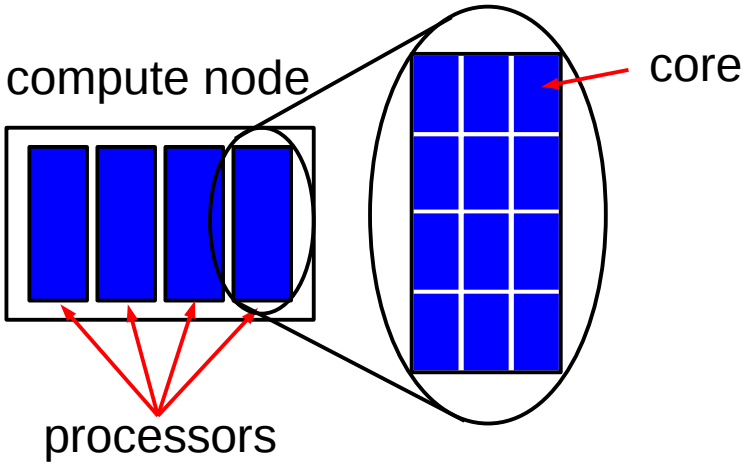

Numerical Algorithms
computing efficiently
computing in parallel

File Storage
managing disk space
efficient i/o

Differences between desktop/laptop vs. HPC

- One processor vs. multi-processor
- Run interactively vs. batch jobs
- Single job vs. multiple jobs
- Prepare on login/head node vs. direct
- Point and click vs. command line

HPC @ W&M

Recent events:

Moved to ISC-3 Q3/4 2016
Consolidated servers/nodes from Jones Hall, JLab, and (some from) physics
into Hot-aisle containment APC NetShelter 24 racks, 250 kW power

Total cores available for general use: **2152**

Coming on-line soon (by end Q1):

- Bora subcluster: 30 2xIntel E5 2640v4 (40 cores,128GB / node)
  1200 cores total

- Meltemi subcluster: 100 Intel Xeon Phi / Knights Landing (64 cores,32GB / node)
  6400 cores total – Shared with W&M Physics

- More disk space: ~200 TB parallel file system

New server room
ISC 1251

SciClone

Front-end/login

sub-cluster

Typhoon

Typhoon (ty01-ty72) 288 cores
4 cores/node Opteron Santa-Rosa (2.6 GHz)
60: 8 GB/node
12: 24 GB/node
~2007

Hurricane

Hurricane (hu01-hu12) 96 cores
8 cores/node Intel X5672 (3.2 GHz)
48 GB/node
~2011

Whirlwind (wh01-wh52) 448 cores
8 cores/node Intel X5672 (3.2 GHz)
44: 64 GB/node
8: 192 GB/node
~2011

Vortex

Vortex (vx01-vx36) 432 cores
12 cores/node Opteron Seoul (3.1 GHz)
28: 32 GB/node
8: 128 GB/node
~2015

Storm

Wind (wi01-wi28) 448 cores
12 cores/node Opteron Magny-Cours (2.4 GHz)
32 GB/node
~2012

ice (ice01,ice02) 48/32 cores
Opteron Magny-Cours (2.4 GHz)
96/64 GB/node
~2012

hail (ha01-ha38) 304 cores
8 cores/node Opteron Shanghai (2.4 GHz)
28: 32 GB/node
8: 128 GB/node
~2011

For beginners...

ssh to typhoon.sciclone.wm.edu
to run a job on typhoon

Necessary Topics for HPC

- How to log into HPC machines?
- Linux Shell / Text editors – basic Linux skills.
- What software do I want to run? (do I need to compile?)
- What sub-cluster will I use?
- What file-system should I use?
- Using the batch system.

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***

Logging into HPC machines

Must use Secure Shell client (SSH)

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

```
[ewalter@particle ~]$ ssh hurricane.sciclone.wm.edu
Password:
Last login: Tue Feb  2 13:57:59 2016 from particle.hpc.wm.edu
------------------------------------------------------------------------
           William and Mary Information Technology / SciClone Cluster
.
.
.
1 [hurricane]
```

Questions to ask yourself:

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

If you are off-campus

– 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

Linux Shell Usage

Main things to learn about linux/shell

- learn to log in and out of front-end servers
  ssh, X forwarding, alternate user name

- manipulating files/folders
  cd – change directory
  cp – copy file
  mv – rename file
  man – read manual page
  rm – remove file
  mkdir – make directory
  rmdir – remove directory

- configuring your environment
  env – print your environmentals variables
  editing your .cshrc.$PLATFORM file

http://www.wm.edu/offices/it/services/hpc/using/shell/index.php - page of linux tutorials

Linux Text Editors

Popular text editors: emacs or vi/vim

emacs: huge, bloated, not installed by default, but the champion!
vi/vim: tiny, always available, some users love it (?)

nano: editor with training wheels, very easy to use, not very powerful

Emacs Tutorials

http://www.gnu.org/software/emacs/tour/
http://www.jesshamrick.com/2012/09/10/absolute-beginners-guide-to-emacs/
http://www2.lib.uchicago.edu/keith/tcl-course/emacs-tutorial.html

Vim Tutorials

http://www.vim.org/
http://vim.wikia.com/wiki/Tutorial
http://linuxconfig.org/vim-tutorial

Nano homepage

http://www.nano-editor.org/

Using Software

- All HPC machines use *Environment Modules* to select software

- **Environmental variables** – determines a user's environment ;
  what commands/applications/libraries are available

$PATH – determines what can be executed
$LD_LIBRARY_PATH – determines what libraries are available at run time

- ***Environment Modules*** make these easy to change on the command line
- Before running applications on HPC machines, modules should be selected
- Default software modules can selected (best) by editing appropriate .cshrc.$PLATFORM
- Modules can be changed on demand (not as good but can be necessary)

Learning how to use Modules is important for using HPC

Using Software Modules

```
15 [hurricane] module avail

---------------------------------- /usr/local/Modules/modulefiles -----------------
acml/5.1.0/gcc                      mpc/0.8.2(default)
acml/5.1.0/pgi                      mpfr/2.4.2
acml-int64/5.1.0/gcc                mpfr/2.4.2a(default)
acml-int64/5.1.0/pgi                mpi4py/1.3.1/gcc
acml-mp/5.1.0/gcc                   mummer/3.23
acml-mp/5.1.0/pgi                   mvapich2-ib/1.2x1/pgi
acml-mp-int64/5.1.0/gcc             mvapich2-ib/1.9/gcc
acml-mp-int64/5.1.0/pgi             mvapich2-ib/1.9/gcc-4.8.4
admb/11.2b/gcc                      mvapich2-ib/1.9/pgi
allpathslg/47017/gcc                mvapich2-ib/1.9a2/gcc
         .                                  .
         .                                  .
         .                                  .
         .                                  .
```

```
16 [hurricane] module list
Currently Loaded Modulefiles:
1) modules            2) maui/3.2.6p21      3) torque/2.3.7       4) isa/nehalem        5) gcc/4.7.0
6) python/2.7.5/gcc
```

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

Can use module load and unload commands for current shell. **Best to use startup**

Configuring your Environment

$PLATFORM variable:

```
11 [hurricane] echo $PLATFORM
rhel6-xeon
```

This means that startup is controlled by ***.cshrc.rhel6-xeon*** for ***hurricane***

subcluster	front-end	$PLATFORM	isa	set?
typhoon	typhoon	sles10-opteron	amd64b	---
hurricane	hurricane	rhel6-xeon	nehalem	---
whirlwind	hurricane	rhel6-xeon	nehalem	---
vortex	vortex	rhel6-opteron	seoul	---
wind	storm	rhel6-storm	magny-cours	set_wind
ice	storm	rhel6-storm	magny-cours	set_ice
hail	storm	rhel6-storm	shanghai	set_hail

The .cshrc.$PLATFORM controls what modules are loaded for a batch job

HPC Filesystems / Backup

3 types of filesystems:

home – backed up nightly ; small used for input files and code, etc.
data – backed up weekly ; large files ; medium term storage
scratch – NOT backed up ; large output files short term storage

```
146 [hurricane] df -h
Filesystem            Size  Used Avail Use% Mounted on
.
.
.
/dev/mapper/VolGroup30-LogVol31
                      917G  482G  390G  56% /sciclone/home00
tn00:/usr/local       134G  113G   15G  89% /usr/local
tn00:/export           46G   15G   30G  33% /import
mh00:/var/spool/mail  7.9G  4.3G  3.3G  57% /var/spool/mail
gfs00:/sciclone/home04
                      591G  429G  157G  74% /sciclone/home04
ty00:/sciclone/scr02  273G   81M  273G   1% /sciclone/scr02
tn00:/sciclone/scr10  7.9G  2.3G  5.3G  31% /sciclone/scr10
tn00:/sciclone/scr30   17T   13T  4.7T  72% /sciclone/scr30
gfs00:/sciclone/data10
                       16T   15T  1.9T  89% /sciclone/data10
tw00-i8:/sciclone/data20
                       73T   57T   16T  79% /sciclone/data20
/dev/md1              8.1T  5.0T  2.8T  65% /sciclone/scr20
vx00:/sciclone/home10
                      2.7T  202M  2.6T   1% /sciclone/home10
vx00:/sciclone/scr00  318G  4.8G  297G   2% /sciclone/scr00
```

Use local scratch if you can! Will give the best performance...

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

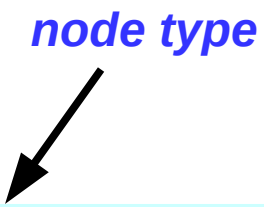


```
27 [vortex] qsub -I -l walltime=30:00 -l nodes=1:vortex:ppn=12
qsub: waiting for job 1552781.vortex.sciclone.wm.edu to start
qsub: job 1552781.vortex.sciclone.wm.edu ready

1 [vx01] python prog.py
```

Interactive job puts you on a node ready to work

all job ids are the form of ###.vortex.sciclone.wm.edu even on other subclusters...

There are many *node types*. The default node type is simply the sub-cluster name
vortex

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

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

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 [hurricane] cat run
#!/bin/tcsh
#PBS -N test
#PBS -l nodes=1:x5672:ppn=8
#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 [hurricane] qsub run
```

```
148 [vortex] 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 [hurricane] more run
#!/bin/tcsh
#PBS -N test
#PBS -l nodes=1:c9:ppn=1
#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 [hurricane] 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
.
.
.
```

http://www.wm.edu/offices/it/services/hpc/using/tutorials/index.php

Using the Batch System IV

Gaussian example: Can get test input files (.com files) and supplied answers from

/usr/local/gaussian/g09/tests/com and **/usr/local/gaussian/g09/tests/amd64**

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

cd $PBS_O_WORKDIR

setenv GAUSS_SCRDIR
/sciclone/scr10/ewalter
g09 < test0000.com > test.out0000
```

set $GAUSS_SCRDIR to something
defaults to current directory

g09 is the name of the program
test0000.com is the input file (from the com folder)
test.out0000 is the output file

Can we do something better than a serial job?

Gaussian Link 0 commands

from: https://www.msi.umn.edu/sites/default/files/IntroToGaussian09.pdf

%mem=n **sets the amount of dynamic memory (**n**), default is 32MB. Units allowed, kb, mb, gb, kw, mw, or gw.**

%nproc=n **sets the number of processors,** n**, to use**

%chk=file **location and name of checkpoint** file

%rwf=file **location and name of rwf** file

```
17 [typhoon] head input
%nproc=1
%mem=6gb
%rwf=/sciclone/scr10/ewalter/testrwf
%chk=/sciclone/data10/ewalter/testchk

#p pbepbe/6-311G sparse test scf
.
.
```

Gaussian Shared Memory Parallel

Gaussian09 has the ability to use multiple cores within a node (larger calculations)

Example: test0445.com in Gaussian directory; DFT optimization:

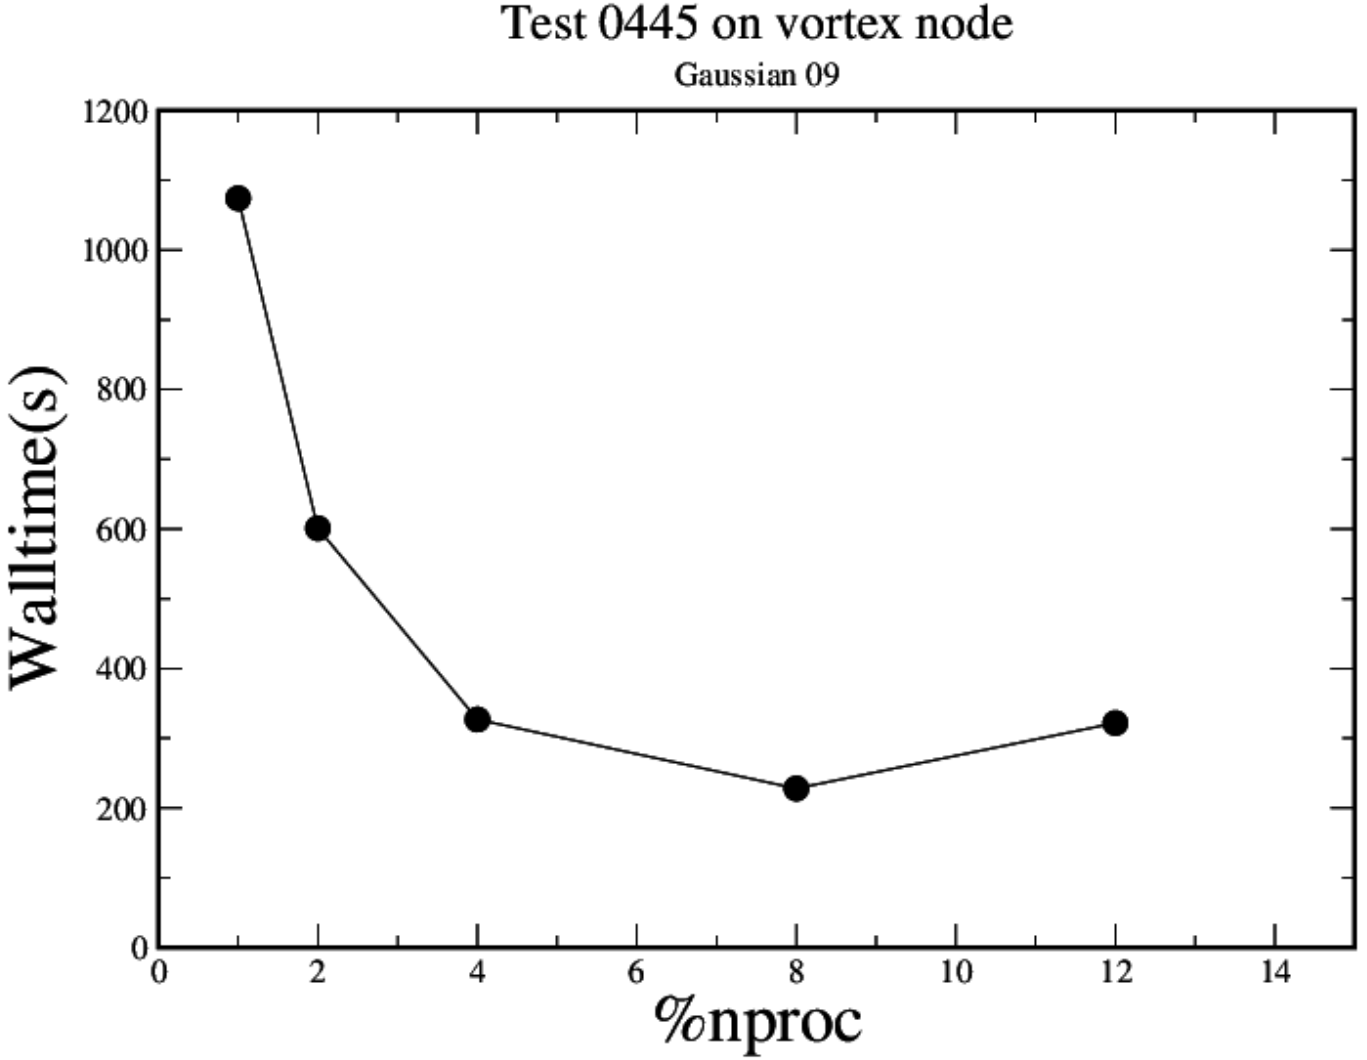

Always worth checking that parallelism is not slowing down run!

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!