



Cornell University
Center for Advanced Computing

Programming Environment

Cornell Center for Advanced Computing
January 23, 2017

*Thanks to Dan Stanzione, Bill Barth, Lars Koesterke,
Kent Milfeld, Doug James, and Robert McLay
for their materials developed at TACC and XSEDE
that were incorporated into this talk.*



1. Accessing Stampede
2. Login Environment
3. Stampede Overview
4. Software
5. Compiling
6. Timing
7. Editing Files
8. Batch Job Submission: SLURM
9. Play Nice
10. Help



Cornell University
Center for Advanced Computing

1. Accessing Stampede



Before you can access Stampede

1. Get an XSEDE Portal Account : <https://portal.xsede.org/>
2. Get an Allocation (computing hours)
 - PI must request allocation through appropriate portal
 - PI may use portal to assign active users to an allocation
 - Note your allocation's "project name" (account code)
3. Activate your account on TACC resources (now automatic)
4. Decide how you're going to log on:
 - Use XSEDE Single-Sign-On
 - Set up Multi-Factor Authentication (MFA) at XSEDE
 - SSH directly to TACC
 - Log into TACC portal, and reset your password
 - Set up Multi-Factor Authentication (MFA) at TACC



2. Get an Allocation -- View My Allocations on XUP

- Go to the XSEDE User Portal (XUP): portal.xsede.org
- Log in
- Go to 'My XSEDE' tab
- Choose 'Accounts' on the nav bar
- See which resources you have access to
- Note: you may not have the same username on all resources

View accounts using the XUP

XSEDE | USER PORTAL
Extreme Science and Engineering Discovery Environment

Search XSEDE...

MY XSEDE RESOURCES DOCUMENTATION ALLOCATIONS TRAINING USER FORUMS HELP ECSS ABC

Summary Allocations/Usage Accounts Jobs Profile Publications Tickets Change Password Add User Community Accounts

XSEDE Single Sign on Login Hub
You can SSH into any XSEDE system with your PORTAL username and PORTAL password from the convenience of your desktop.

XSEDE recommends you use the XSEDE Single Sign on Login Hub to login to XSEDE resources with your local username and password. Use a client on your desktop to SSH to login.xsede.org with your portal username and password then easily gsi-ssh to any XSEDE system you have access to with no additional username or passwords. For more information please visit the [XSEDE Single Sign on Login Hub](#) document.

RESOURCE NAME	GSI-SSH LOGIN HOST	INSTITUTION	LOCAL USERNAME
Gordon ION	gordon.sdsc.xsede.org	SDSC	
Mason	mason.iu.xsede.org	IU	
Maverick	maverick.tacc.xsede.org	TACC	tg459571
SuperMIC	supermic.cct-lsu.xsede.org	LSU CCT	tg459571
OSG	submit-1.osg.xsede.org	OSG	shm7
Gordon	gordon.sdsc.xsede.org	SDSC	mehringe
XStream	xstream.stanford.xsede.org	Stanford U	xs-tg459571
Bridges	bridges.psc.xsede.org	PSC	mehringe



4. Decide How You're Going to Log On

SSH to the XUP Single Sign On (SSO) Login Hub

- Pros:
 - No need for a resource-specific username and password.
 - Easy way to switch between XSEDE resources

OR: SSH directly to TACC resources

- Pros:
 - Fewer clicks

****MFA is required to access TACC resources either way**



Either way you need SSH:

- SSH Secure Shell
 - SSH on the Command Line (Unix/Linux) or Terminal window (Mac)
 - SSH / telnet client for Windows – e.g. PuTTY
- **Do not** overwrite ~/.ssh/authorized_keys
- **Do** add stampede to the list of known hosts, if prompted

On linux: Open a command line window, then \$ ssh ...

On Mac: Open a terminal window, then \$ ssh ...


On Windows: Install and use a client like [PuTTY](#) (gui)





Single Sign On (SSO) Login Hub

Set up MFA <https://portal.xsede.org/mfa/>

- Install the Duo app on your smartphone or other device
Find  on iTunes or Google Play
Duo at Cornell: <https://it.cornell.edu/twostep>
- [Enroll your XSEDE Portal account in Duo](#)
- [Pair your Duo-enabled device with your XUP account](#)



Single Sign On (SSO) Login Hub

- **ssh** to **login.xsede.org** - single point-of-entry to XSEDE resources
- A 12 hour [proxy certificate](#) is automatically generated
- **gssssh** to any XSEDE compute resource
- Log in using your XUP username and password

Wait 2 slides if you've set up direct access

On linux or Mac:

```
localhost$ ssh XUPusername@login.xsede.org  
[shm7@sschub ~]$ gssssh stampede-knl
```

On Windows:

```
Start PuTTY, then enter Host Name login.xsede.org  
[shm7@sschub ~]$ gssssh stampede-knl
```

<https://portal.xsede.org/web/xup/single-sign-on-hub>



Connect Directly to TACC resources

Log into TACC portal, and reset your password (first time)

Set up MFA

<https://portal.tacc.utexas.edu/tutorials/multifactor-authentication>

- Manage Profile on the TACC User Portal
- Select Pairing Method
 - TACC Token app – download & pair TACC Token app
 - SMS (text) messaging - pair
 - TACC Hard token – request physical token (fob)



Connect Directly to TACC resources

- **ssh** to **login-knl1.stampede.tacc.utexas.edu**
- Log in using your TACC username and password

On linux or Mac:

```
localhost$ ssh TACCusername@login-knl1.stampede.tacc.utexas.edu
```

On Windows:

Start PuTTY, then enter Host Name

login-knl1.stampede.tacc.utexas.edu

(Stay logged in)

<https://portal.tacc.utexas.edu/user-guides/stampede#knl-system-access>



Cornell University
Center for Advanced Computing

2. Login Environment



Account Info

Note your account number in the splash screen.

```
----- Project balances for user tg459571 -----
| Name           Avail SUs   Expires   |
| TG-TRA140011   21016     2017-01-26 |
----- Disk quotas for user tg459571 -----
| Disk           Usage (GB)   Limit     %Used   File Usage   Limit     %Used   |
| /home1         0.0         5.0       0.07    231         150000    0.08    |
| /work          0.0        1024.0    0.00    13         30000000  0.00    |
| /scratch       0.0         0.0       0.00    2           0         0.00    |
-----
```



Get the Lab Files

- TAR = Tape ARchive. Just concatenates files.
- `tar <switches> <files>`
 - z = compress or decompress
 - x = extract
 - c = create
 - v = verbose
 - t = list files
 - f = next argument is the file to read or write
- `~username` is the home directory of that user
- For example, to create a tar: `tar cvf myfiles.tar dir1 dir2 README`

```
$ tar xvf ~tg459572/LABS/cornellcac_labs.tar Get all lab files
```

```
$ cd envi Change directory to the files for this session
```

```
$ ls -R List the files in the current folder (envi)
```



Stampede's Operating System: Linux

\$ pwd	(Print the current directory)
\$ ls -la	(List the content of the current directory)
\$ cd \$HOME	(Change the working directory to home directory)
\$ cat .profile	(Print the file <i>.profile</i> to the screen)
\$ mkdir testdir	(Create the directory, <i>testdir</i>)
\$ touch test.txt	(touch renews a file's timestamp, but here is used to create an empty file)
\$ mv test.txt testdir	(Move text.txt into the directory testdir)
\$ ls -la testdir	(See the files in folder testdir)
\$ rm -r testdir	(Delete the directory and all subdirectories)
\$ man ls	(Show the manual page for ls, q to quit)
\$ env	(Show all environment/global variables)
\$ export newgreeting="Hello World"	(Set an environment variable)
\$ echo \$newgreeting	(Print the variable <i>newgreeting</i>)



Shells and Startup Scripts on Stampede

Shells:

- bash is the default shell on Stampede
- login1\$ **echo \$SHELL** *To determine your current login shell*
- login1\$ **cat /etc/shells** *To see list of available shells*
- To change your default shell, submit a ticket to TUP (chsh won't work)

Startup Scripts:

- When you log in, system-level startup files execute to allow administrators to enhance and customize the environment
- Enhance your shell environment, not your account
- Don't use "echo" in startup scripts, will break other tools
- TACC staff recommends that Bash shell users use ~/.profile rather than .bash_profile or .bash_login.

[Bash Users' Startup Files: Quick Start Guide](#)



Cornell University
Center for Advanced Computing

3. Stampede Overview



Transition: Stampede => Stampede 2

2011: NSF award to TACC to acquire and deploy Stampede

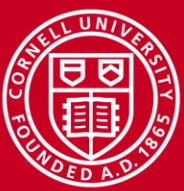
2012: Stampede deployed

2016: NSF award to TACC to acquire and deploy Stampede 2

2017: Stampede 2 fully deployed (Fall)

Over the next 9 months:

- Stampede will remain in production until Stampede 2 is in full production, ~Fall 2017
- Stampede will gradually have fewer nodes
- Stampede has 508 Knights Landing (KNL) nodes for testing; they will later become part of Stampede 2; will gradually increase
- File transfer: During transition Stampede's home and scratch file systems will be accessible from the Stampede 2 login nodes, and the work file system will be the same



Transition: Stampede => Stampede 2

Date	Event	Stampede (nodes)	Stampede-KNL (nodes)	Stampede 2 (nodes)
now		6,400	508	
Late Jan. 2017	Phase 1 start	3,800	508	
Summer 2017	Phase 1 production; Phase 2 start	1,800 (migrate files now)		4,204 (+ new filesystem)
October 2017	Phase 2 production			5,940

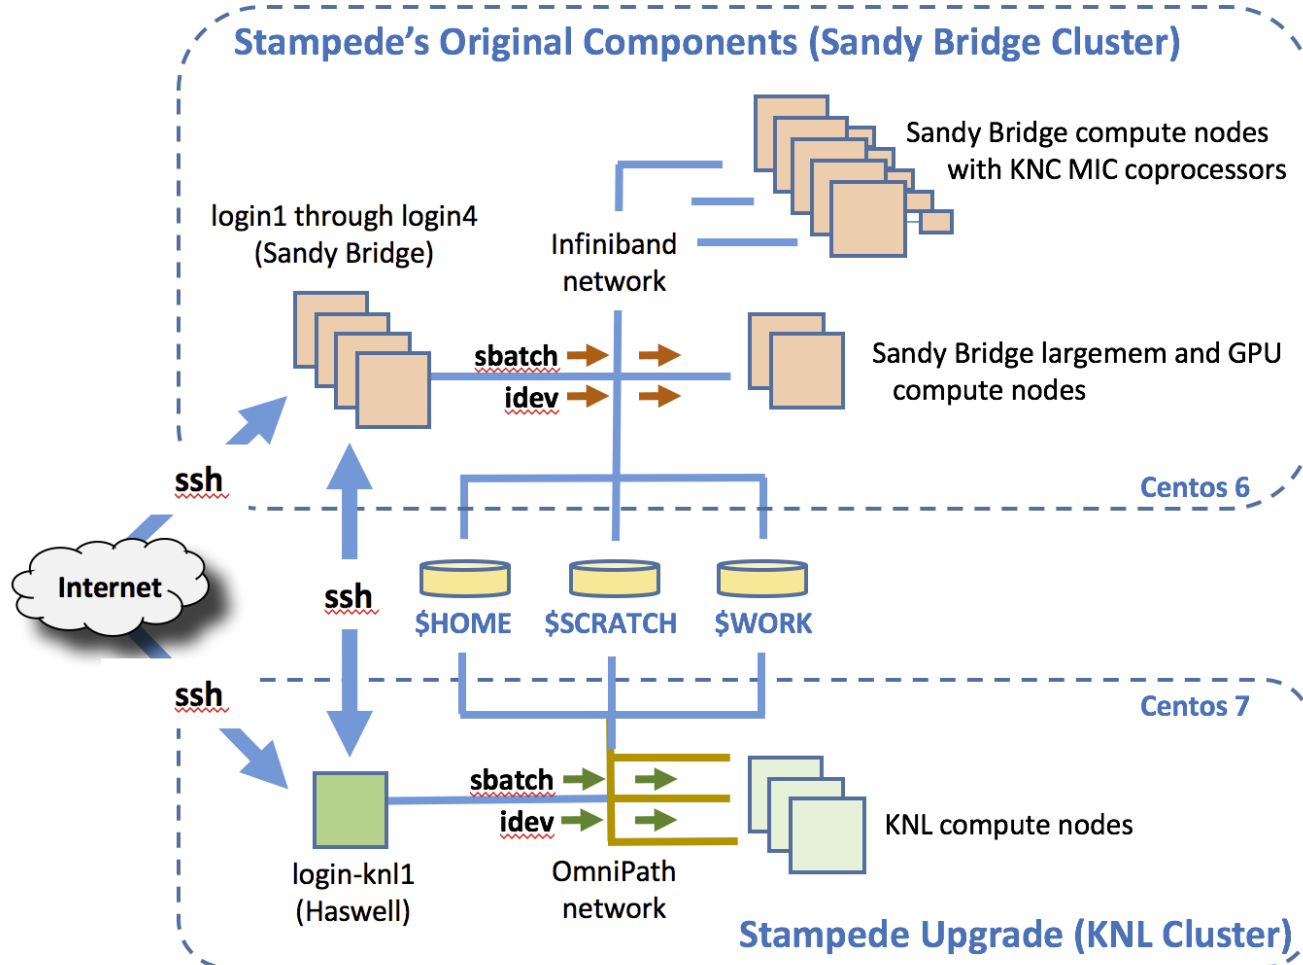
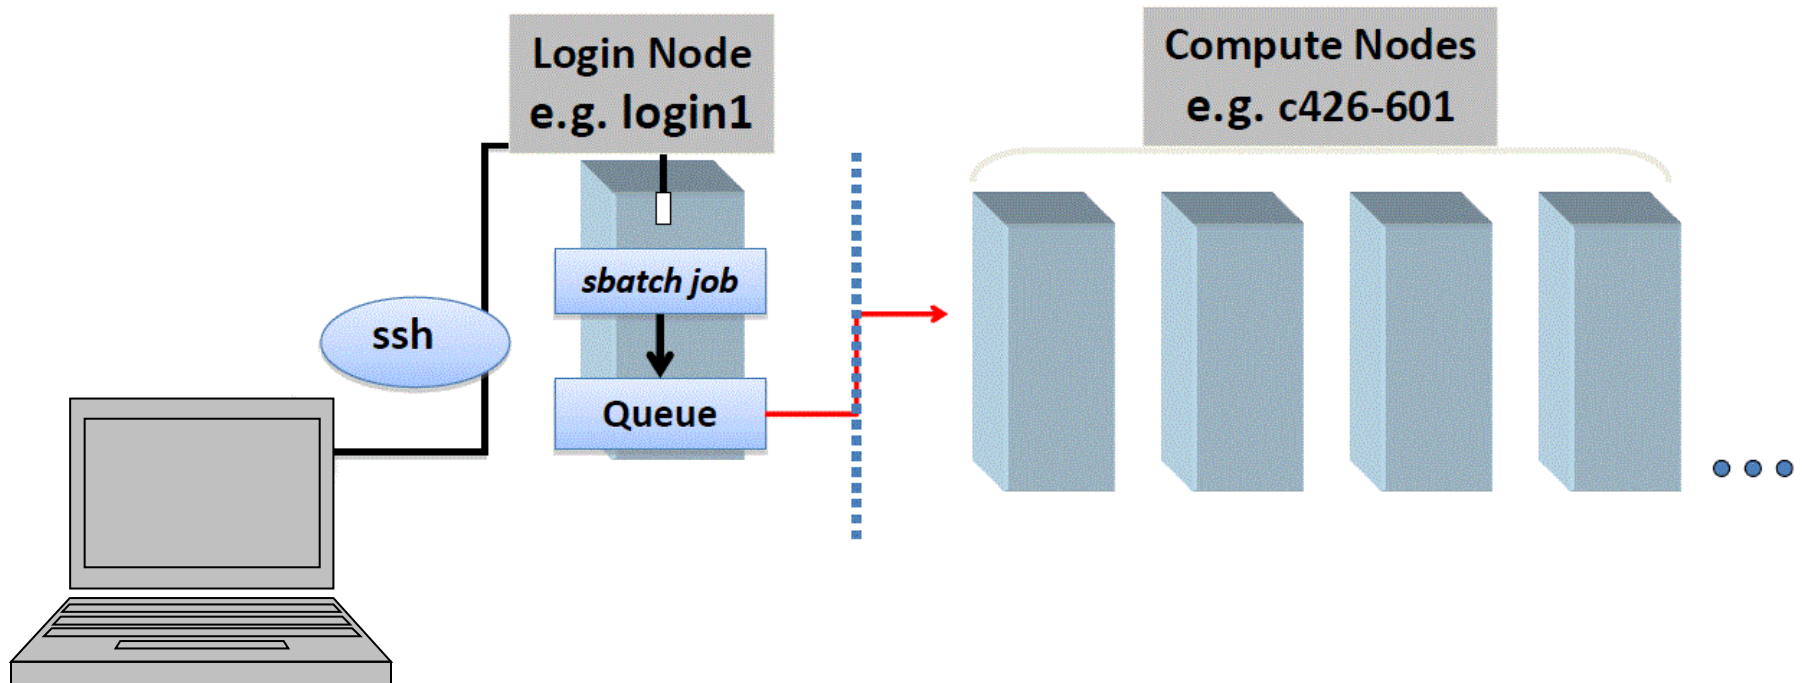


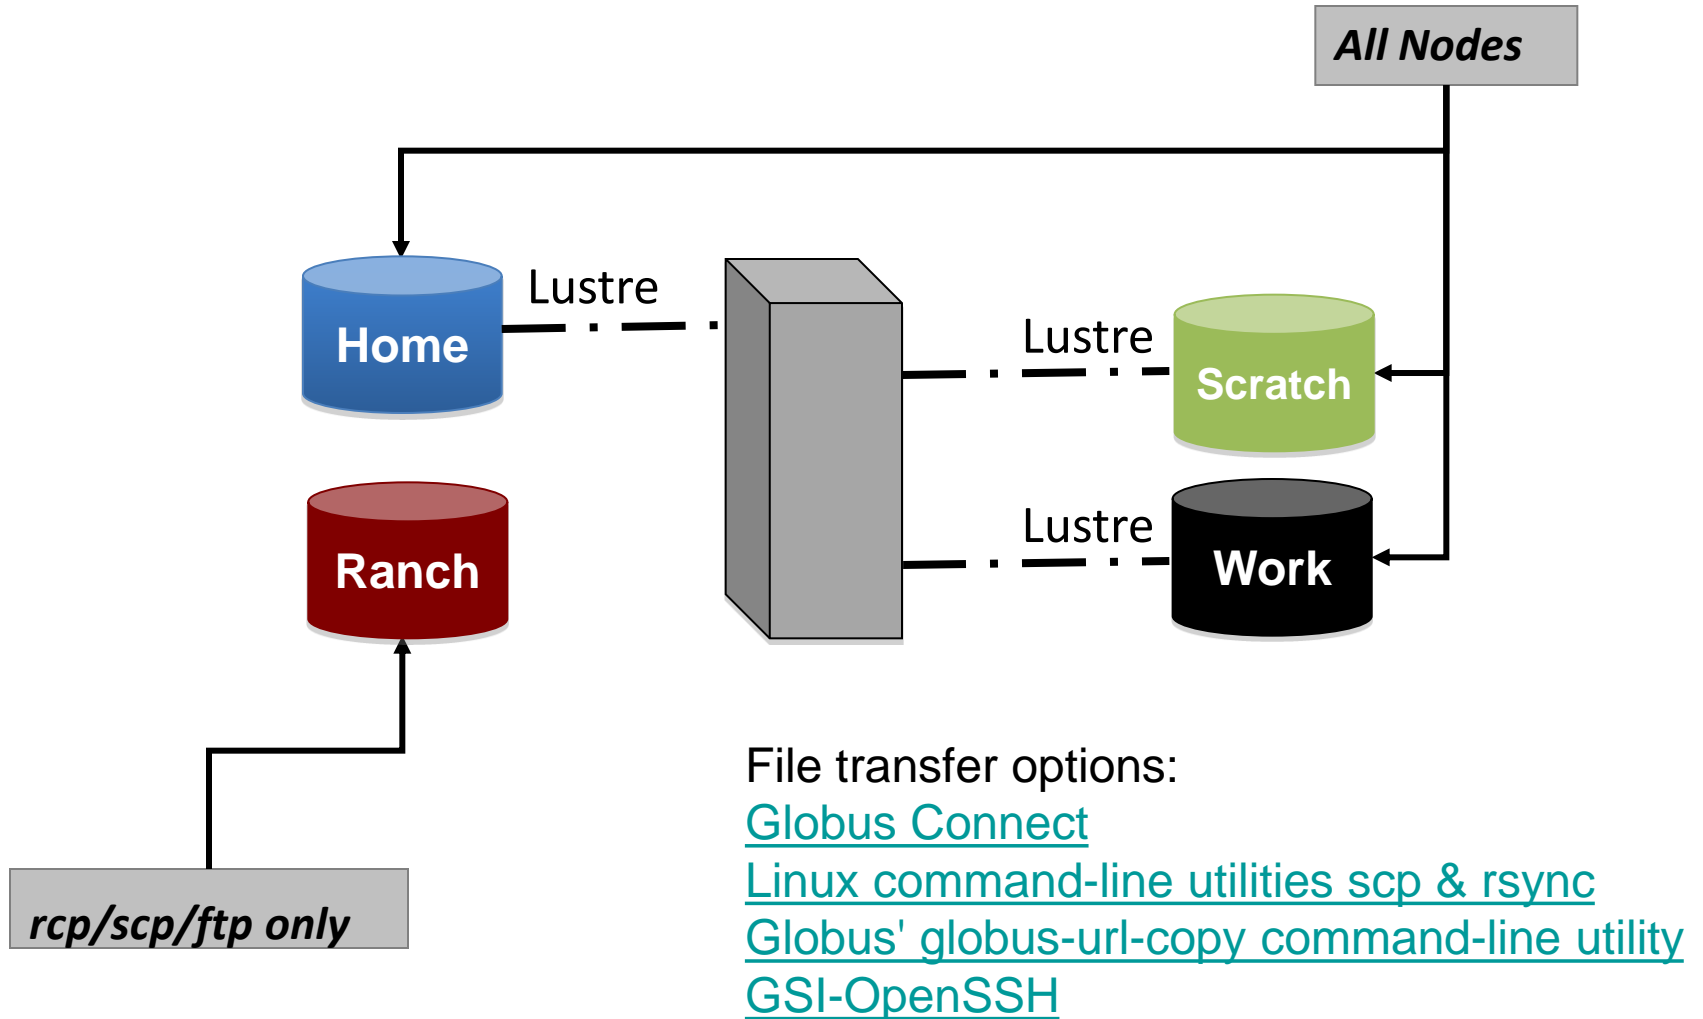
Image: <https://portal.tacc.utexas.edu/user-guides/stampede#intro>



The Generic Environment



Available File Systems





File System

Environment Variable	Purpose	User Access Limits	Lifetime
\$HOME	Source code	5 GB	Backups
\$WORK	Large file storage	1.0 TB	No backup
\$SCRATCH	Large files needed by compute jobs	~8.5PB total	Purged after 10 days
/tmp	Local disk on batch job node	~80 GB / node	Purged after job ends
\${ARCHIVER}:%ARCHIVE	Archival tape	Essentially unlimited	Project



Sharing Files with your Research Group

File sharing tutorial: <https://portal.tacc.utexas.edu/tutorials/sharing-project-files>

- All accounts have a default group when the account is created
- All usernames sharing an allocation should be in a common group
(If they are not, submit a ticket)

<code>\$ groups <username></code>	Display groups that username belongs to
<code>\$ groups</code>	Display groups that you belong to
<code>\$ id -g -n</code>	Display your default group
<code>\$ id <username></code>	Display username, default group, all groups, for that user
<code>\$ touch test.txt</code>	Create a file
<code>\$ ls -la</code>	Display your files, including group information. Note that the file you just created has your default group ownership
<code>\$ chgrp -v G-<grp number> test.txt</code>	Change the group ownership of a file to a different group (verbose output)
<code>\$ chmod 644 test.txt</code>	Modify permissions so everyone in that group has access



Sharing Files with your Research Group

Want to share files with your colleagues, but you have different default groups? You can

- a) submit a ticket to get your default group changed, or
- b) create a common folder with the proper settings:

<code>\$ mkdir /scratch/<grp 1>/<username>/test</code>	Create a directory for everybody to share
<code>\$ chmod g+rwx /scratch/<grp 1>/<username>/test</code>	Set permissions allow group read/write/execute (also make sure the parent dir's permissions aren't too restrictive).
<code>\$ chgrp G-<grp 2> test</code>	Change its group via 'chgrp' to the common group
<code>\$ chmod g+s test</code>	Set the setgid bit, so that everything created underneath it will inherit its group.
<code>\$ umask 177</code>	Everyone in the group should use an appropriate umask such as 002 or 117, so that files they create are actually group readable and writable. (Put this into a login script!)



File System

```
$ ls quota -u <username> $HOME           see quota limits & usage
$ ls quota -u <username> $WORK
$ ls quota -u <username> $SCRATCH
$ cd                                     change directory to $HOME
$ pwd
$ cdw                                     change directory to $WORK
$ pwd
$ cds                                     change directory to $SCRATCH
$ pwd
$ du -sh                                  see how much space is available in the
                                           current user-owned directory
$ df -k .                                  see the amount of disk space used in a file
                                           system, "." meaning in the current directory
```



Cornell University
Center for Advanced Computing

4. Software



Software

Use the [module](#) utility on Stampede to provide a consistent, uniform method to access software

- Loads specific versions of libraries/executables
- Manages dependencies between multiple compilers & software stacks
- Works in your batch file, Makefile, and scripts
- Affects \$PATH, \$MANPATH, \$LIBPATH
- Order matters! First choose compiler, then application software.
- Warning: the module system treats the Sandy Bridge and KNL clusters as separate systems.

[Software](#) search available on XSEDE

[Lmod](#) is TACC's Module System



Setting your Default Software Environment

Set and save your personal default module environment:

```
$ module reset                # return to the default environment
$ module load ddt
$ module load fftw3
$ module save                  # will load at login or restore
```

Create a named collection of modules for reliability and repeatability:

```
$ module save chemtools
...
$ module restore chemtools
```



Module

This utility is used to set up your PATH and other environment variables:

\$ module help	{lists options}
\$ module avail	{lists available modules}
\$ module list	{lists loaded modules}
\$ module load boost	{add a module}
\$ module unload boost	{remove a module}
\$ module help <module_name>	{module-specific help}
\$ module spider	{lists all modules}
\$ module spider petsc	{list all versions of petsc}
\$ module load impi	{unload module A, load module B}
\$ module list	{lists loaded modules}
\$ module reset	{return to system defaults}



Cornell University
Center for Advanced Computing

5. Compiling



Compiling Serial Code

- The default compilers on Stampede are Intel C++ and Fortran (& support KNL)
- Don't try to compile once to run on both KNC and KNL (Network stack incompatible)
- Use **man** or **-help** option, e.g. **man icc**
- Compilers are available on login and compute nodes, but--
- The **login node is a Haswell, not KNL**, processor;
use the "-xMIC-AVX512" switch at both compile and link time for KNL:

```
knl-login1$ icc -xMIC-AVX512 -o mycode.exe mycode.c
```

```
knl-login1$ ifort -xMIC-AVX512 -o mycode.exe mycode.f90
```

Compiler	Language	File Extension	Example
icc	C	.c	icc <i>compiler_options</i> prog.c
icpc	C++	.C, .cc, .cpp, .cxx	icpc <i>compiler_options</i> prog.cpp
ifort	F77	.f, .for, .ftn	ifort <i>compiler_options</i> prog.f
ifort	F90	.f90, .fpp	ifort <i>compiler_options</i> prog.f90



Compiler Options

- Use compiler options to achieve optimal performance.
- To obtain best results:
 - Select the appropriate optimization level
 - Target the architecture of the computer (CPU, cache, memory system)
 - Allow for interprocedural analysis (inlining, etc.)
- No single answer for all cases; test different combinations.

Optimization Level Description

-O0	Fast compilation, full debugging support. Automatically enabled if using -g.
-O1 -O2	Low to moderate optimization, partial debugging support:
-O3	Aggressive optimization - compile time/space intensive and/or marginal effectiveness; may change code semantics and results (sometimes even breaks code!)

See the User Guide for [additional compiler options](#).



Makefiles

\$ cd \$HOME/envi/using_makefiles

\$ cat Makefile Read over the Makefile

\$ make Compile the program, generate a.out

\$ make Reports “up to date”, i.e. not recompiled

\$ touch suba.f90 Simulate changing a file

\$ make suba.f90 (and only suba.f90) is recompiled



Cornell University
Center for Advanced Computing

6. Timing



Timers

- Time your code to see how long your program runs and estimate if it's having gross difficulties. Gauge effectiveness of code and software changes.
- Wall-clock time in a dedicated environment is most accurate
- **/usr/bin/time -p** is preferred over the shell's time command (-p specifies traditional precision output in seconds)

```
$ cd $HOME/envi/intro
$ make
g++ hello.c -o hello
$ /usr/bin/time -p ./hello
Hello world!
real 0.01
user 0.00
sys 0.01
$
```

You can also [time specific sections](#) of your code by inserting timer calls before and after important sections.



Profilers: gprof (GNU profiler)

- gprof reports a basic profile of time spent in each subroutine
- Find the most time-consuming routines, the hotspots
- As with all profiling tools, the code must be instrumented to collect the timing data and then executed to create a raw-data report file.
- Read the data file into an ASCII report or a graphic display.
- Instrument the code by recompiling using the -pg option (Intel)
- <https://portal.tacc.utexas.edu/user-guides/stampede#tools-profilers>
- Also available: [PerfExpert](#)

```
$ cd $HOME/envi/precision
$ ifort -pg precision.f90      instrument code with -pg
$ a.out                      produce gmon.out trace file
$ gprof                       reads gmon.out (default args: a.out gmon.out)
                              report sent to STDOUT
```



Cornell University
Center for Advanced Computing

7. Editing Files



vi (short for “visual”)

- “vi filename” will open it or create it if it doesn’t exist.
- Command mode: keystrokes are commands
- Input mode: keystrokes are text you are adding to the file
- Last line mode: start with : end with <return>
- Examples:
 - i Insert characters before current position (use ESC to exit)
 - dd Delete current line
 - R Overwrite existing text (until ESC)
 - u Undo last operation
 - :wq Writes a file to disk and exit editor
 - :q! Quit without saving

<http://www.tuxfiles.org/linuxhelp/vimcheat.html>



nano

- The commands for all operations are preceded by the Control key:

A screenshot of the nano editor's command list. The text is displayed in a monospaced font on a light background. A dark rectangular box highlights the text "[Read 8 lines]". Below this, two rows of commands are listed, each preceded by a control key symbol (^). The first row contains: ^G Get Help, ^O WriteOut, ^R Read File, ^Y Prev Page, ^K Cut Text, and ^C Cur Pos. The second row contains: ^X Exit, ^J Justify, ^W Where Is, ^V Next Page, ^U UnCut Text, and ^T To Spell.

- If you have modified the file and try to exit (^X) without writing those changes (^O) you will be warned.
- Makes text editing simple, but it has less powerful options than vi (search with regular expressions, etc..)



emacs

- emacs is actually a lisp interpreter with extensions to use it as a text editor
- Can perform the same operations as in vi
- Uses Control or ESC followed by keystroke combinations to execute commands
- “Hard to learn, easy to use”

<http://emacswiki.org/emacs/ReferenceCards>



Files edited in Windows? Warning:

- Beware line ending differences
- Do not use MS Word (e.g.) to create scripts or programs
- Do not copy/paste from PDF files
- Linux filenames are case-sensitive
- Blanks in filenames can be problematic
- Use Wordpad rather than Notepad if necessary
- [dos2unix](#) utility can convert text files with DOS or MAC line breaks to Unix line breaks



Cornell University
Center for Advanced Computing

8. Batch Job Submission: SLURM



Getting to the Compute Nodes

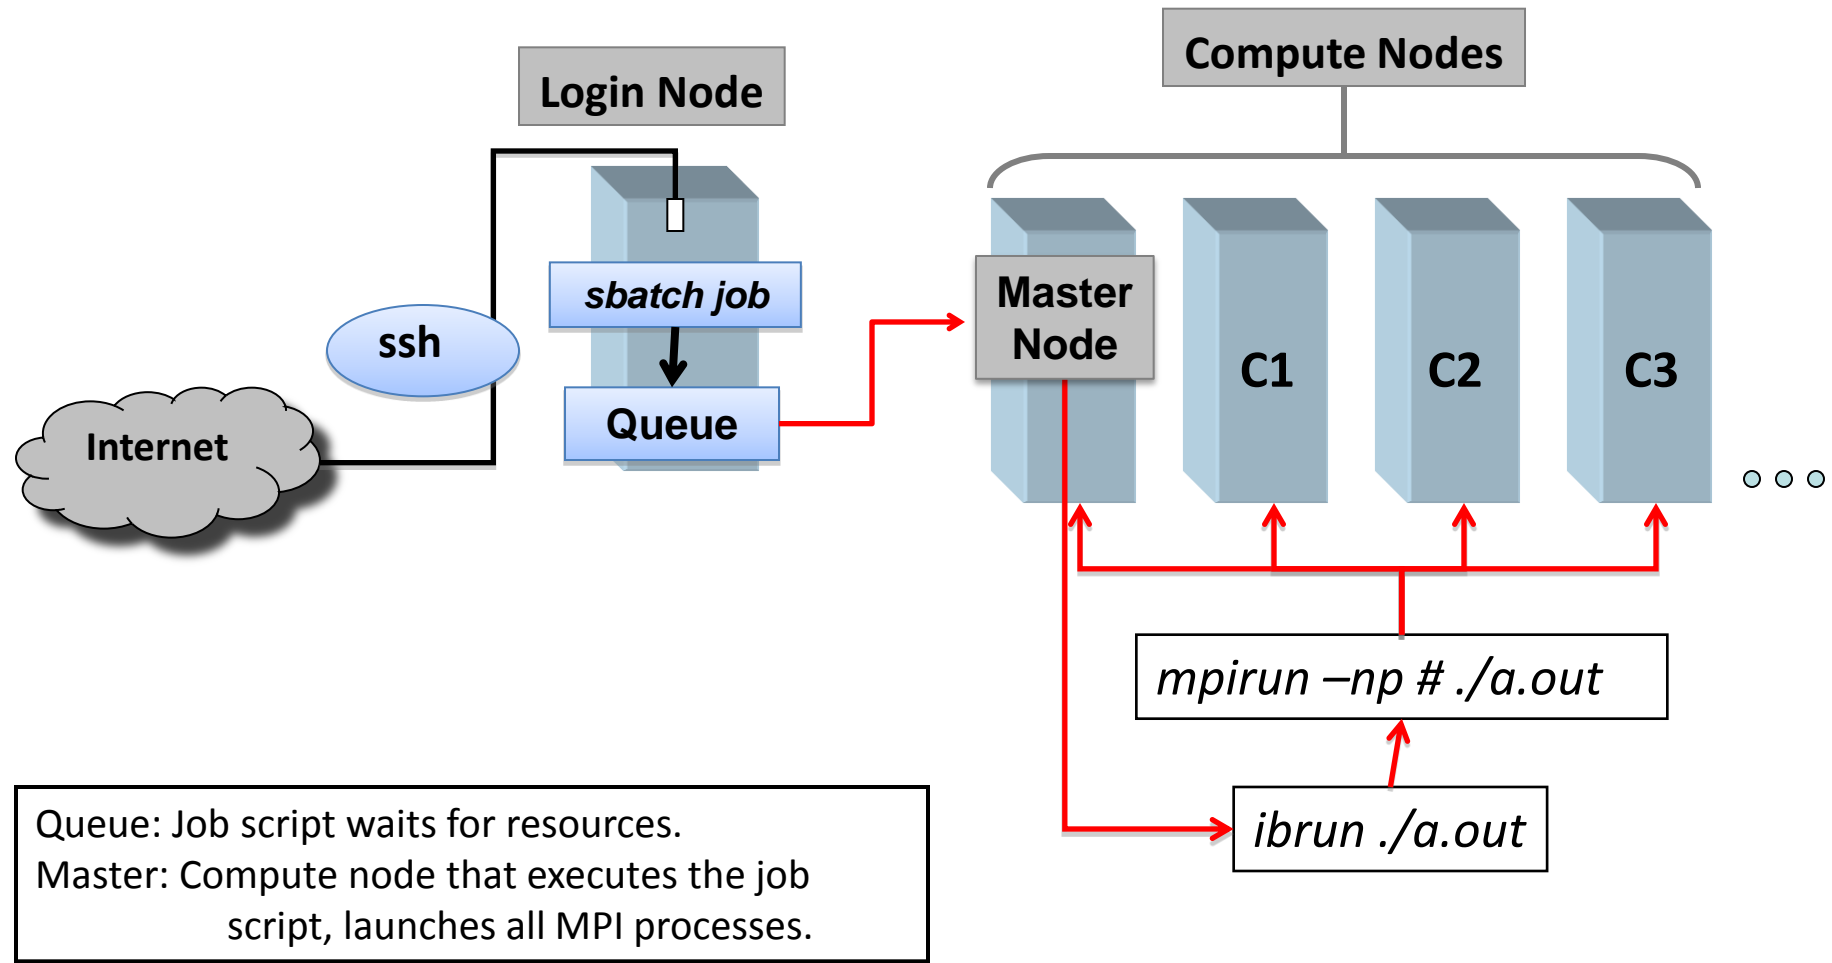
Four ways to get to the back end (compute nodes):

- SLURM batch job: **sbatch <batchfilename>**
- Interactive session: [srun](#) (SLURM) or [idev](#) (TACC)
- Run special app that connects to back end: e.g. **ddt**
- ssh to node on which you already have a job running

If you don't use sbatch, srun, or equivalent, you're running on the front end (login nodes) – don't do this!

- Don't launch exe (e.g. **./a.out**) on the command line
- One of the easiest ways to get your account suspended

Batch Submission Process





Stampede Batch Environment Queues

Sandy Bridge Queue	Max Runtime	Max Nodes/Procs	Max Jobs in Queue	Queue Multiplier	Purpose
normal	48 hrs	256 / 4K	50	1	normal production
development	2 hrs	16 / 256	1	1	development nodes
largemem	48 hrs	3 / 96	3	2	large memory 32 cores/node
serial	12 hrs	1 / 16	8	1	serial/shared_memory
large	24 hrs	1024 / 16K	50	1	large core counts (access by request ¹)

KNL Queue	Max Runtime	Max Nodes and Associated Cores per Job	Max Jobs in Queue	Charge per node hour	Configuration (Memory-Cluster)
development	2 hrs	4 nodes (272 cores)	1	16 SU	Cache-Quadrant
normal	48 hrs	80 nodes* (5440 cores)	10	16 SU	Cache-Quadrant
Flat-Quadrant	48 hrs	40 nodes* (2720 cores)	5	16 SU	Flat-Quadrant
Flat-All2All	12 hrs	2 nodes* (136 cores)	1	16 SU	Flat-All-to-All
Flat-SNC-4	12 hrs	2 nodes* (136 cores)	1	16 SU	Flat-SNC-4

<https://portal.tacc.utexas.edu/user-guides/stampede#knl-running>



Batch on Stampede: Select SLURM Commands

- **showq** - view summary of jobs in the batch system (not SLURM)
showq | more
showq -u <userid>
- **sacct** - report job or job step accounting information.
- **salloc** - allocate resources for a job in real time.
- **sbatch** - submit a job script for later execution.
sbatch <batchfilename>
- **sbcast** - transfer a file from local disk to local disk on the job nodes.
- **scancel** - cancel a pending or running job or job step.
scancel <jobid>
- **sinfo** - lists the availability and status of queues
sinfo -o "%20P %5a %.10l %16F"
- **squeue** - reports the state of jobs or job steps.
squeue | more
squeue -u <userid>
- **srun** - submit an interactive job (this example: 1-node 16 core)
srun --pty -n 16 -t 00:30:00 -p development -A <acct no> /bin/bash -l
- **ibrun** – run an MPI program (put this command in your batch script for MPI jobs)

Man pages exist for all SLURM daemons, commands, and API functions. The command option **--help** also provides a brief summary of options. Note that the command options are all case insensitive.



queue Options, Output, and Job State Codes

-i <interval>	Repeatedly report at intervals (in seconds).
-j <job_list>	Displays information for specified job(s)
-p <part_list>	Displays information for specified partitions (queues).
-t <state_list>	Shows jobs in the specified state(s)

JOBID	job id assigned to the job
USER	user that owns the job
STATE	current job status.

PD	Pending
R	Running
S	Suspended
CA	Configuring
CG	Completing
CD	Completed
CF	Cancelled
F	Failed
TO	Timeout
PR	Preempted
NF	Node_fail



Batch Job Script Example: MPI

```
#!/bin/bash                                     # Don't miss this line!

#-----
# Generic SLURM script -- MPI
#-----

#SBATCH -J myjob                               # Job name
#SBATCH -o myjob.%j.out                       # stdout; %j expands to jobid
#SBATCH -e myjob.%j.err                       # stderr; skip to combine stdout and stderr
#SBATCH -p development                         # queue
#SBATCH -N 2                                  # Number of nodes, not cores (16 cores/node)
#SBATCH -n 32                                 # Total number of MPI tasks (if omitted, n=N)
#SBATCH -t 00:30:00                           # max time

#SBATCH --mail-user=myemail@myuniv.edu
#SBATCH --mail-type=ALL

#SBATCH -A TG-TRA120006                       # necessary if you have multiple project accounts

module load fftw3                               # You can also load modules before launching job
module list

ibrun ./main.exe                               # Use ibrun for MPI codes. Don't use mpirun or srun.
```



Batch Job Script Example: Serial

```
#!/bin/bash                                # Don't miss this line!

#-----
# Generic SLURM script
#-----

#SBATCH -J myjob                            # Job name
#SBATCH -o myjob.%j.out                     # stdout; %j expands to jobid
#SBATCH -e myjob.%j.err                     # stderr; skip to combine stdout and stderr
#SBATCH -p serial                            # queue
#SBATCH -N 1 -n 1                            # one node and one task
#SBATCH -t 00:30:00                         # max time

#SBATCH --mail-user=myemail@myuniv.edu
#SBATCH --mail-type=ALL

#SBATCH -A TG-01234                          # necessary if you have multiple project accounts

module load fftw3                            # You can also load modules before launching job
module list

./main.exe
```



Batch on Stampede: SLURM Commands

1. Use **sinfo -o "%20P %5a %.10I %16F"** to list queues, nodes, and system state
2. Issue **showq** to show all queued jobs
3. `$ idev -r -A TG-TRA140011`
4. Issue **cat** to take one last look at the batch script
`$ cd $HOME/envi/batch`
`$ cat job`

```
#!/bin/bash
#SBATCH -J myMPI                # Job name
#SBATCH -o myjob.%j.out         # stdout file (%j expands to jobId)
#SBATCH -p development          # Queue name
#SBATCH -N 2                    # Total number of nodes requested (16 cores/node)
#SBATCH -n 32                   # Total number of mpi tasks requested
#SBATCH -t 01:30:00             # Run time (hh:mm:ss) - 1.5 hours
#SBATCH -A TG-TRA140011        # Specify project/allocation number
ibrun ./a.out
```
5. Compile: **mpicc -xMIC-AVX512 -O3 mpihello.c -OR- mpif90 -xMIC-AVX512 -O3 mpihello.f90**
6. Issue **sbatch** to submit a batch script : **sbatch job**
7. Issue **squeue -u <your username>** to see the job status
8. Run **scancel <jobid>** to cancel the job, or **cat myjob.###.out** to view your output

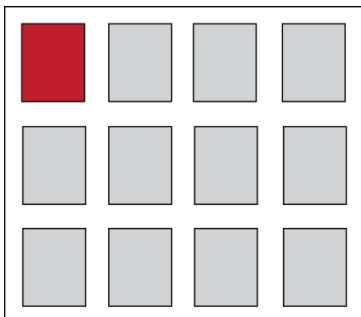


Resource Allocation on SLURM

- `-N` – Total nodes requested
- `-n` – Number of tasks per node
- Explicitly specify **both** "`-N`" (total nodes) and "`-n`" (tasks per node) in your Slurm job script; you'll get 272 tasks per node if only `n` is specified

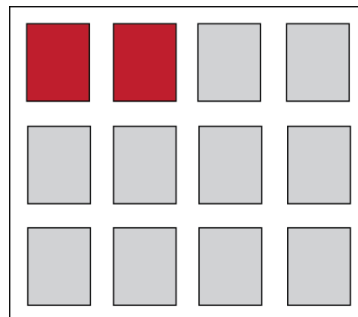
Serial Job

```
#SBATCH -N 1  
#SBATCH -n 1
```



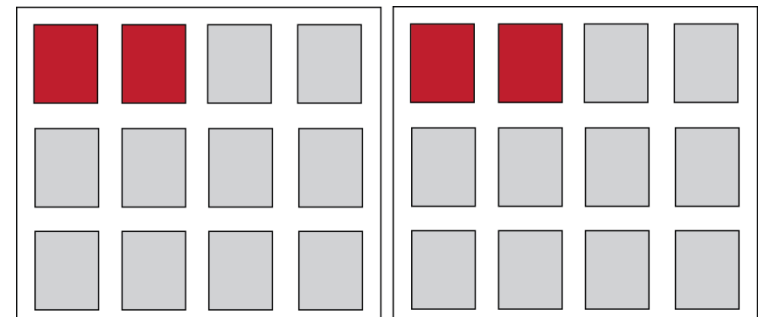
2 Tasks

```
#SBATCH -N 1  
#SBATCH -n 2
```



4 Tasks Parallel

```
#SBATCH -N 2  
#SBATCH -n 4
```





Cornell University
Center for Advanced Computing

9. Play Nice



- TACC's [Acceptable Use Policy](#)
- A large, shared, resource like Stampede works best when everyone practices [good citizenship](#)
 - [Respect the shared filesystems](#)
 - Run jobs should access files in \$WORK or \$SCRATCH, not in \$HOME
 - Don't exceed ~three concurrent file transfer sessions
 - Limit I/O intensive sessions (lots of reads and writes to disk)
 - [Don't run programs on the login nodes](#)
 - Login nodes are for compiling, file management, managing batch jobs, modest post-processing
 - Login nodes are not for running programs; submit a batch job instead



Cornell University
Center for Advanced Computing

10. Help



Questions?

Help Tickets

- CAC help@cac.cornell.edu Email to submit a ticket
- TACC <https://portal.tacc.utexas.edu/> Log in, Open a ticket
- XSEDE <https://portal.xsede.org/> Log in, Help, Help Desk

Training/Documentation

- <https://cvw.cac.cornell.edu/Topics/> Online training
- <https://portal.tacc.utexas.edu/user-guides/stampede> User Guide
- <https://portal.xsede.org/documentation-overview> Getting Started
- For shell commands, try **man <command>** or **man -k <command>**



Cornell University
Center for Advanced Computing

Appendix

Precision

The precision program computes and prints $\sin(\pi)$.

The π constant uses “E” (double precision) format in one case and “D” (single) in the other.

```
$ cd $HOME/envi/precision
$ cat precision.f90
$ module load intel
$ ifort precision.f90
$ ./a.out
```