



# Programming Environment

Cornell Center for Advanced Computing  
June 11, 2013

*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. Help



# 1. Accessing Stampede



## Before you Start

- Get an XSEDE Portal Account : <https://portal.xsede.org/>
- 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)
- Activate your account on TACC resources
  - Involves email handshake(s) with TACC user services
  - May take a few business days
  - Note that your TACC credentials (think ssh) may differ from XSEDE
  - TACC password resets can take 30+ minutes to propagate



## Logging into XSEDE Resources:

- Command Line (Unix/Linux) – ssh
- SSH / telnet client – e.g. Putty
- Single Sign On (SSO) from the XSEDE User Portal
- ... [and more](#)



## Login with SSH:

- [Putty](#) for Windows
- Built-in as “ssh” for Linux or Mac
- You will be connected to login#.stampede.tacc.utexas.edu
- **Do not** overwrite ~/.ssh/authorized\_keys

Using either Putty or ssh, login to stampede.tacc.utexas.edu:

All Programs | ClassFiles | putty  
use Host Name: stampede.tacc.utexas.edu

All Programs | Accessories | Command Prompt  
% ssh *username*@stampede.tacc.utexas.edu



## Login with SSO

- Go to the XSEDE User Portal: portal.xsede.org
- Log in
- Go to 'My XSEDE' tab
- Go to the 'Login to your XSEDE Accounts' link
- Use the appropriate 'login' link
- Note your username

Login using the XSEDE portal

XSEDE USER PORTAL  
Extreme Science and Engineering  
Discovery Environment

HOME MY XSEDE RESOURCES DOCUMENTATION ALLOCATIONS TRAINING USER FORUMS HELP STAFF

Allocations/Usage Accounts My Jobs Profile Tickets Registered DNs Change Portal Password Add User Community Accounts SSH Terminal

SEARCH:

RESOURCE NAME	LOGIN NAME	INSTITUTION	USERNAME	CONNECT
Blacklight	<a href="#">blacklight.psc.teragrid.org</a>	PSC	stanzio	<a href="#">Login</a>
Condor	<a href="#">tg-condor.purdue.teragrid.org</a>	Purdue	dstanzi	<a href="#">Login</a>
Dash	<a href="#">dash.sdsc.teragrid.org</a>	SDSC	dstanzi	<a href="#">Login</a>
Forge	<a href="#">login-forge.ncsa.xsede.org</a>	NCSA	dstanzi	<a href="#">Login</a>
Kraken	<a href="#">kraken-gsi.nics.utk.edu</a>	NICS		<a href="#">Login</a>
Lonestar	<a href="#">lonestar.tacc.teragrid.org</a>	TACC	dan	<a href="#">Login</a>
Longhorn	<a href="#">tg-login.longhorn.tacc.teragrid.org</a>	TACC	dan	<a href="#">Login</a>
Nautilus	<a href="#">login.nautilus.nics.xsede.org</a>	NICS		
Ranger	<a href="#">tg-login.ranger.tacc.teragrid.org</a>	TACC	dan	<a href="#">Login</a>
Spur	<a href="#">tg-login.spur.tacc.teragrid.org</a>	TACC	dan	<a href="#">Login</a>
Steele	<a href="#">tg-steele.purdue.teragrid.org</a>	Purdue	dstanzi	<a href="#">Login</a>
Trestles	<a href="#">trestles.sdsc.edu</a>	SDSC	dstanzi	<a href="#">Login</a>



## 2. Login Environment





## Account Info

Note your account number in the splash screen.

```
----- Project balances for user tg459571 -----
| Name          Avail SUs    Expires |
| TG-TRA120006   49998                |
-----
----- Disk quotas for user tg459571 -----
| Disk          Usage (GB)   Limit   %Used   File Usage   Limit   %Used |
| /home1        0.0          5.0     0.06    43          150000  0.03 |
| /work         0.0         400.0   0.00    3           30000000 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`

Get the lab files:

```
$ tar xvf ~tg459572/LABS/envi.tar
```

Change directory to the envi directory:

```
$ cd envi
```

List the lab files:

```
$ ls -la
```



## Experiment with Linux commands on Stampede

\$ pwd	(Print the current directory)
\$ ls -la	(List the content of the current directory)
\$ cd \$HOME	(Change the working directory to your home directory)
\$ cat .login	(Print the file <i>.login</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)
\$ rm -r testdir	(Delete the folders and all subfolders)
\$ man ls	(Show the manual page for ls, , 'q' to quit)
\$ env	(Show all environment/global variables)
\$ export newgreeting="Hello World"	(Set an environmental variable)
\$ echo \$newgreeting	( Print the variable <i>newgreeting</i> )



## Shells and Startup Scripts on Stampede

### Shells:

- bash is the default shell on Stampede
- TACC supports most major shells, e.g. csh, tcsh, zsh ...
- To change your default shell, submit a ticket (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
- Put your personal customization in `.login_user`


<http://www.tacc.utexas.edu/user-services/user-guides/stampede-user-guide#compenv-startup>



## 3. Stampede Overview

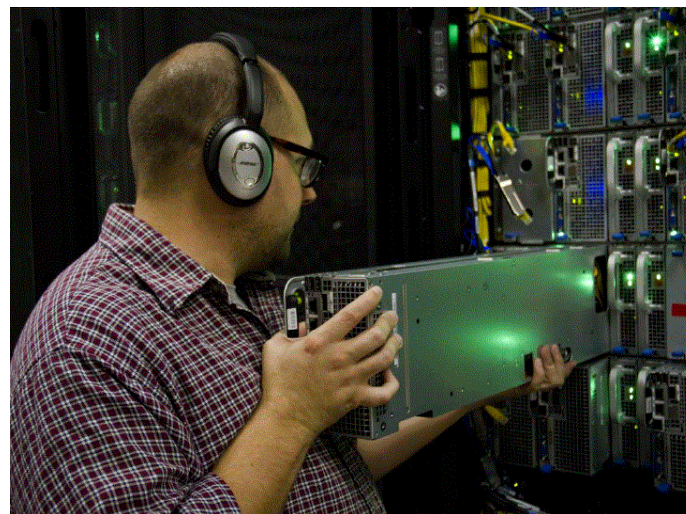
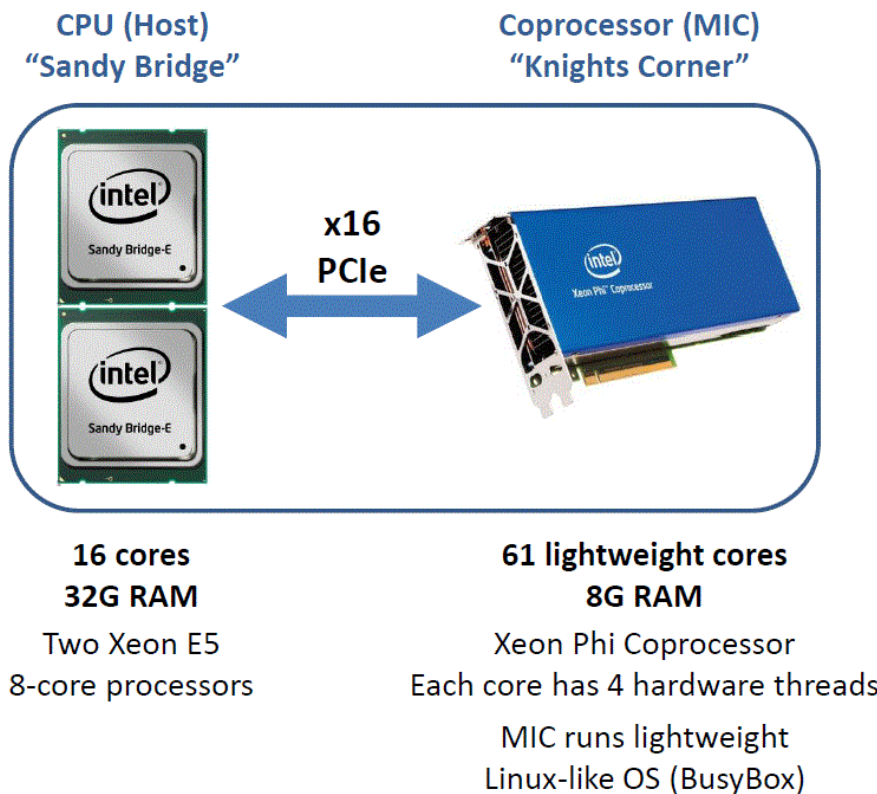


# The Generic Environment





## Typical Stampede Node ( = blade )

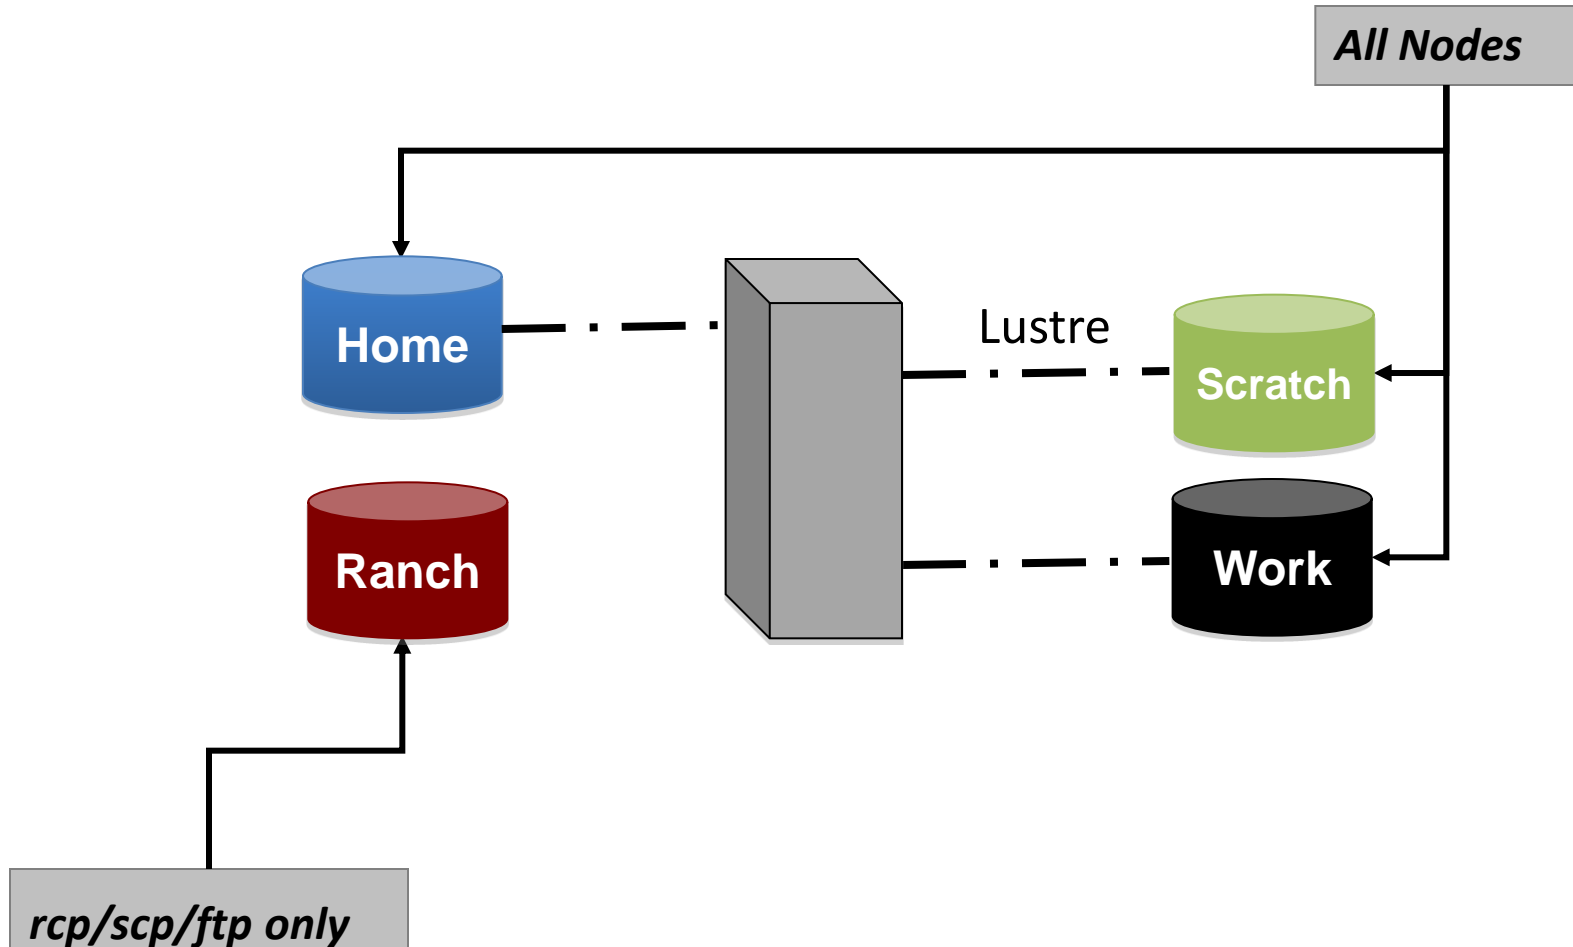




System	Stampede	Memory
Nodes	~6400 (in 160 racks) → 96,000+ total cores	
Typical Node	16 cores: 2 cpus/node x 8 cores/cpu	32GB RAM
	61 cores on MIC coprocessor,	8G RAM
Special Nodes	16 large memory nodes (32 Xeon cores)	1TB/node RAM
	128 GPU nodes (w/ NVIDIA Kepler 2 & MIC)	2GB/core
	Login nodes (don't have MIC)	
CPUs	Intel Sandy Bridge Intel Xeon Phi coprocessor	
Interconnect	56Gb FDR IB	
Disk	14PB Lustre (IB)	



# Available File Systems





# File System

Environment Variable	Purpose	User Access Limits	Lifetime
\$HOME	Source code	5 GB	Backups
\$WORK	Large file storage	400 GB	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
\$ARCHIVE	Archival tape	Essentially unlimited	Project



# File System

```
$ lfs quota -u <username> $HOME           see quota limits & usage
$ lfs quota -u <username> $WORK
$ lfs 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
```



## 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 and software stacks
- Works in your batch file, Makefile, and scripts, but not on MICs
- Affects \$PATH, \$MANPATH, \$LIBPATH
- Order matters! First choose compiler, then application software.

[Software](#) available on Stampede

[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 gcc	{add a module}
\$ module load intel	{try to load intel}
\$ module swap gcc intel	{swap two 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}



## 5. Compiling





## Compiling Serial Code

- The default compiler on Stampede is Intel C++ and Fortran
  - This is the only compiler that support the Phi coprocessors
- Compilers are available on login and compute nodes
  - But not on MIC coprocessors; compile from a Sandy Bridge host
- Use **man** or **-help** option, e.g. **man icc**.

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

- Use the **module** command to list modules & versions & to change the default compiler.
- Three versions of gcc suite are also available
- Other specialized compilers also supported, e.g. cuda support (nvcc):  
**module load cuda**



## Compiler Options

- Use compiler options to achieve optimal performance.
- Obtain best results by
  - 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.f**                Simulate changing a file

**\$ make**                    suba.f (and only suba.f) is recompiled



## 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)
- More detail can be found in the [Profiling and Debugging](#) Virtual Workshop module.

```
$ 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

** There is no output for this simple example **
```



## 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

- All operations commands are preceded by the Control key:
  - ^G Get Help
  - ^O WriteOut
  - ^X Exit
  - ....
- 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 series of multiple keystroke combinations to execute commands
- “Hard to learn, easy to use”

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



## Use Your Computer's Editor

Copying the file to your computer might be quicker than learning a new editor. Use a simple file transfer client:

Start menu

- All Programs

  - Class Files

    - SSH Secure Shell

      - Secure File Transfer Client ← Right click, "Pin to Start Menu"

Start Secure File Transfer Client

Use Quick Connect, specify hostname `lonestar.tacc.utexas.edu`

In the left pane, navigate to the desktop.

Drag files between panes to copy.

\*\* Beware line ending differences!



## 8. Batch Job Submission: SLURM



## Getting to the Compute Nodes

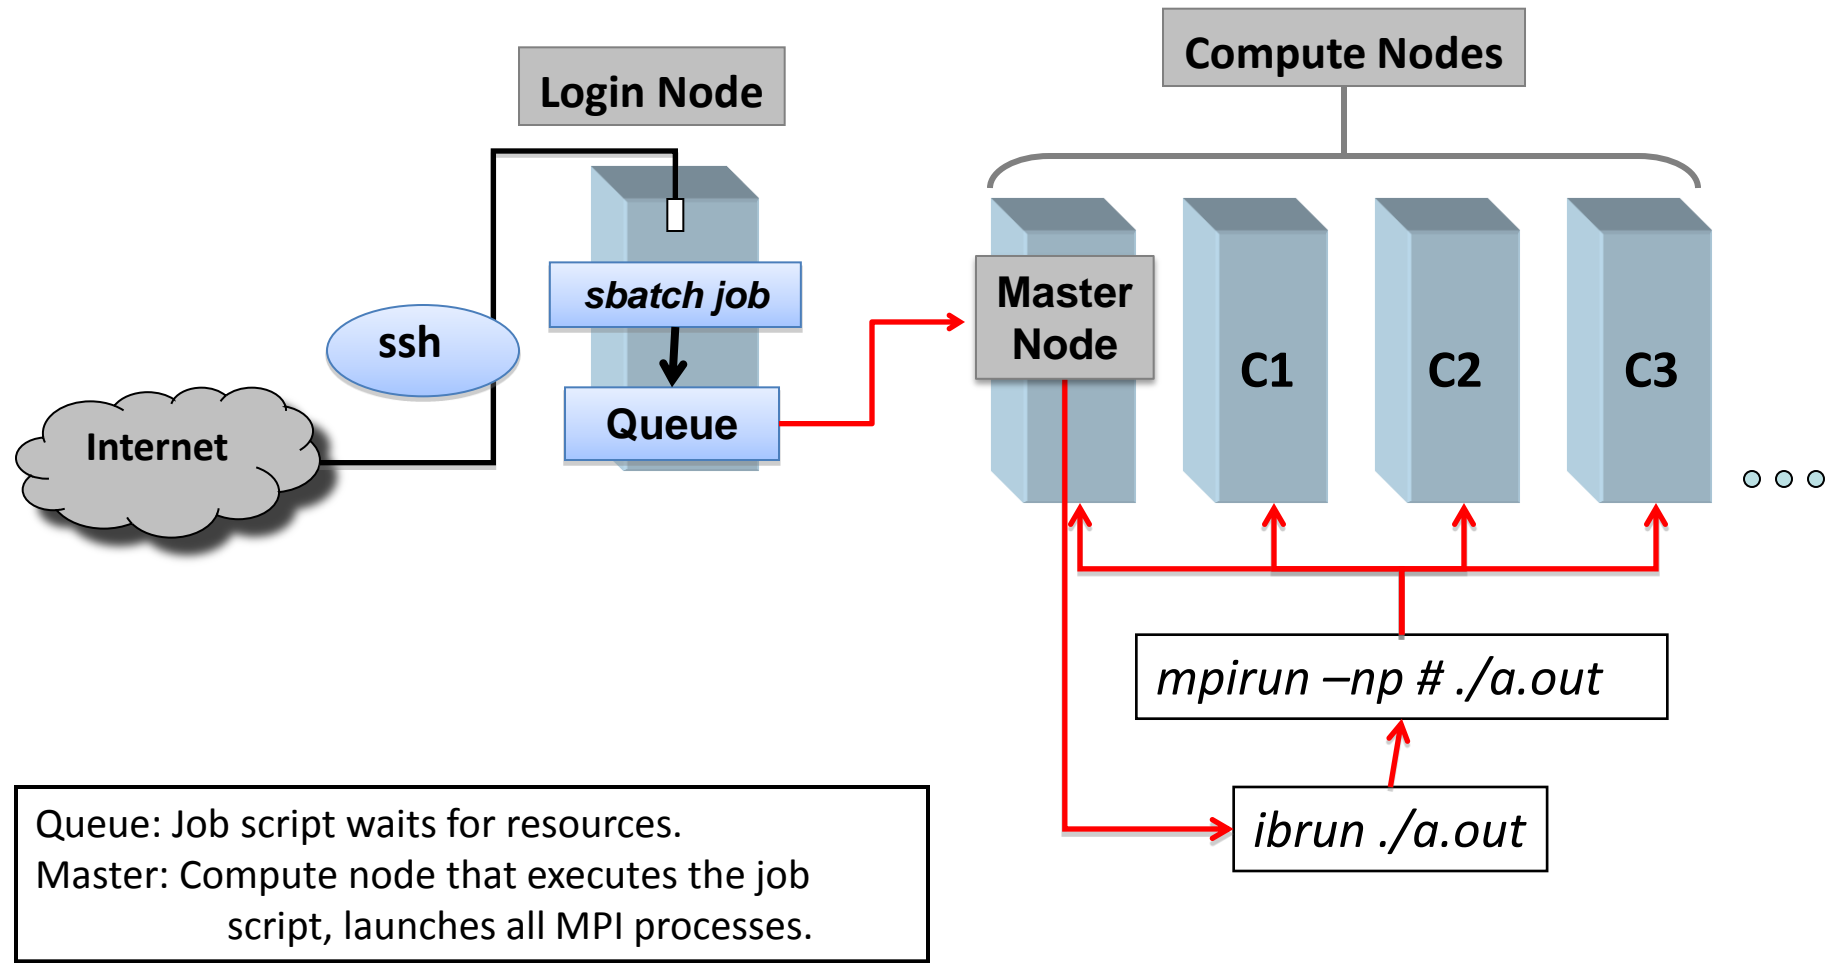
Four ways to get to the back end (compute nodes):

- SLURM batch job: **sbatch <batchfilename>**
- SLURM interactive session: **srun <flags>**
- Run special app that connects to back end: e.g. **ddt**
- ssh to node on which you already have a job running
  - once on compute node, **ssh mic0** gets you to its mic

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

Queue Name	Max Runtime	Max Nodes/Procs	SU Charge Rate	Purpose
normal	24 hrs	256 / 4K	1	normal production
development	4 hrs	16 / 256	1	development nodes
largemem	24 hrs	4 / 128	2	large memory 32 cores/node
serial	12 hrs	1 / 16	1	serial/shared_memory
large	24 hrs	1024 / 16K	1	large core counts **
request	24 hrs	--	1	special requests
normal-mic	24 hrs	256 / 4k	1	early production mic nodes
gpu	24 hrs	32 / 512	1	GPU nodes
gpudev	4 hrs	4 / 64	1	GPU development nodes
vis	8 hrs	32 / 512	1	GPU nodes + VNC service

<http://www.tacc.utexas.edu/user-services/user-guides/stampede-user-guide#running-slurm-queue>



## 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** - reports the state of partitions and nodes managed by SLURM.  
sinfo -o "%20P %5a" *ignore queue limits reported*
- **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 20130418HPC /bin/bash -l
- **ibrun** - submit a batch job (not SLURM)

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.





## squeue 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 -- 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 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"** to list queues, nodes, and system state
2. Issue **showq** to show all queued jobs
3. Issue **srun** to run simple commands (e.g. an interactive shell) (ctrl-D to exit)  
\$ srun --pty -A TG-TRA120006 -p serial -t 00:10:00 -n 1 -N 1 /bin/bash -l
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  
ibrun ./a.out
5. Compile: **mpicc -O3 mpihello.c -OR- mpif90 -O3 mpihello.f90**
6. Issue **sbatch** to submit a batch script  
\$ sbatch job  
sbatch: Submitted batch job 469
7. Issue **squeue -u <your username>** to see the job status
8. Run **scancel <jobid>** to cancel the job, or look at your output




## Resource Allocation on SLURM

- **-N** – Number of node requested
- **-n** – Number of tasks to run


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
```





## 9. Help



## Questions?

- CAC [help@cac.cornell.edu](mailto:help@cac.cornell.edu)
- portal.xsede.org -> **Help** (in the navigation bar)
- portal.xsede.org -> My XSEDE -> **Tickets**
- portal.xsede.org -> Documentation -> **Knowledge Base**
- User Guide(s), Usage Policies, etc. and associated links:  
<http://www.tacc.utexas.edu/user-services>
- Try `man <command>` or `man -k <command>` or `<command> -h`  
or `<command> -help`



# Appendix



# Precision

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

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

```
$ cd $HOME/envi/precision
$ cat precision.f
$ module load intel
$ ifort -FR precision.f
(or)
$ ifort precision.f90
$ ./a.out
```

( The ifc compiler regards “.f” files as F77 fixed format programs. The -FR option specifies that the file is free format.)