

# Ranger Environment Documentation

Drew Dolgert

July 13, 2010

## Contents

|          |   |           |
|----------|---|-----------|
| <b>1</b> | <b>Introduction</b>                           | <b>2</b>  |
| 1.1      | Goals, Prerequisites, Resources               | 2         |
| 1.2      | Ranger Hardware                               | 2         |
| <b>2</b> | <b>Connect</b>                                | <b>3</b>  |
| 2.1      | Goals, Resources, Agenda                      | 3         |
| 2.2      | SSH from Linux or Mac                         | 4         |
| 2.3      | SSH from Windows                              | 4         |
| 2.3.1    | Installation of Secure Shell                  | 4         |
| 2.3.2    | Starting a Client Session                     | 4         |
| 2.4      | X-Windows from a Mac                          | 5         |
| 2.5      | X-Windows from Linux                          | 5         |
| 2.6      | X-Windows from Windows                        | 6         |
| 2.6.1    | Installation of X-Windows                     | 6         |
| 2.6.2    | Starting a Client Session                     | 7         |
| 2.7      | Exercise: Use SSH to Connect                  | 8         |
| 2.8      | Starting: Read Examples of Sessions           | 9         |
| 2.9      | Further Exercise: Using X-Windows to Connect  | 9         |
| 2.10     | Exercise: Connect with VNC                    | 10        |
| 2.11     | Further Exercise: Choose Your Shell           | 10        |
| 2.12     | Advanced: Make Login Faster                   | 11        |
| 2.12.1   | Making Shortcuts on Windows                   | 11        |
| 2.12.2   | Setup SSH Keys for No Password Login          | 11        |
| <b>3</b> | <b>Compile</b>                                | <b>11</b> |
| 3.1      | Goals, Resources, Agenda                      | 11        |
| 3.2      | About the Module Command                      | 12        |
| 3.3      | How to Read a Makefile                        | 12        |
| 3.4      | Modules List                                  | 13        |
| 3.5      | Discussion Exercise: Learn Modules            | 16        |
| 3.6      | Exercise: Use Modules to Load Boost Libraries | 17        |
| 3.7      | Exercise: Compile a Benchmark                 | 17        |
| <b>4</b> | <b>Run on Ranger</b>                          | <b>18</b> |
| 4.1      | Goals, Resources, Agenda                      | 18        |
| 4.2      | Text Editors on Ranger                        | 18        |
| 4.2.1    | VI Cheat Sheet                                | 18        |
| 4.3      | How to Construct a Batch File                 | 19        |
| 4.4      | How to Monitor a Job on Ranger                | 20        |

|          |   |           |
|----------|---|-----------|
| 4.5      | Exercise: Run in Batch . . . . .                              | 20        |
| 4.6      | Discussion: Why the Long Queue? . . . . .                     | 21        |
| 4.7      | Further Exercise: Using Scheduler-Defined Variables . . . . . | 21        |
| <b>5</b> | <b>Transfer and Archive Files</b>                             | <b>21</b> |
| 5.1      | Goals, Resources, Agenda . . . . .                            | 21        |
| 5.2      | Filesystems on Ranger . . . . .                               | 22        |
| 5.3      | Tape Archive . . . . .  | 22        |
| 5.4      | File Transfer with GridFTP . . . . .                          | 22        |
| 5.5      | Exercise: Retrieve Files with Secure Copy . . . . .           | 23        |
| 5.5.1    | From Linux or Mac . . . . .                                   | 23        |
| 5.5.2    | From Windows . . . . .  | 23        |
| 5.6      | Discussion: Transfer Techniques . . . . .                     | 24        |
| 5.7      | Exercise: File Transfer with BBCP . . . . .                   | 24        |

Contents:

# 1 Introduction

## 1.1 Goals, Prerequisites, Resources

**Goals:** This tutorial should prepare you to compile and run programs on the Ranger system at the Texas Advanced Computing Center.

**Prerequisites:** You must already have an account on Ranger. You have to understand how to work with files and directories. You need to be able to install software on your computer to allow you to connect to Ranger.

**Resources:** The Ranger User Guide has all details that distinguish this system from other Linux clusters. The beginning Unix tutorial will be tremendously helpful if you aren't well accustomed to Linux.

- Access to the training file `envi.tgz`, which contains exercises.
- Ranger User Guide - <http://services.tacc.utexas.edu/index.php/ranger-user-guide>
- Beginning Unix - <http://info.ee.surrey.ac.uk/Teaching/Unix>

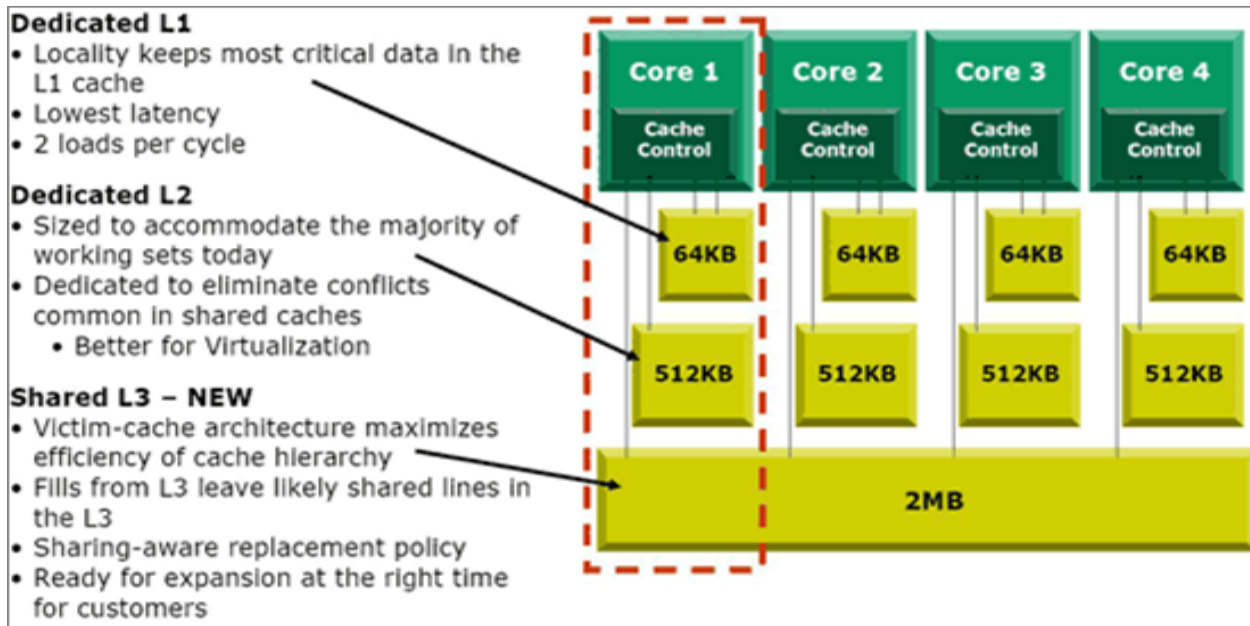
This document accompanies a talk introducing the working environment on the Ranger cluster at the Texas Advanced Computing Center. The exercises inside are enough for someone already familiar with the Linux command line to submit jobs to the Ranger cluster. There is also some reference information, mostly a convenient copy of the Ranger user guide.

Learning to use the Linux command line for the first time is significantly more difficult than learning to submit simple jobs to Ranger. Finding a separate guide on C Shell will greatly increase your comfort.

## 1.2 Ranger Hardware

There are 3,936 nodes of 16 cores, which makes 62976 cores. Peak performance is 579 TFLOPS.

Each node on Ranger has four processors of four cores each. It's a 64-bit AMD Barcelona chipset at 2.3 GHz. With 9.2 GFlops/core, there are 36.8G Flops/chip, or 147.2 GFlops/node.



Check out the Ranger User Guide for more.

## 2 Connect

### 2.1 Goals, Resources, Agenda

Goals: Connect from the computer you are using now to Ranger in order to work there.

Resources: \* Ranger User Guide - <http://services.tacc.utexas.edu/index.php/ranger-user-guide>

\* Windows-only

- Putty Secure Shell Client - <http://www.chiark.greenend.org.uk/~sgtatham/putty/>
- Xming X-Windows Server - <http://www.straightrunning.com/XmingNotes/>
- (Optional) A VNC client, such as RealVNC or Tight VNC.
- Mac-only
  - Operating system installation discs, if you don't have secure shell installed.
  - (Optional) A VNC client, such as RealVNC or TightVNC.
- Linux-only
  - (Optional) A VNC client, such as RealVNC, TightVNC, vinagre, or vncviewer.
- How to set up no-password logins - <http://www.mtu.net/~engstrom/ssh-agent.php>

Agenda: In order to complete this section, you must do these sections:

1. *Exercise: Use SSH to Connect*
2. (For beginners) *Starting: Read Examples of Sessions*
3. *Further Exercise: Using X-Windows to Connect*
4. (Optional) *Exercise: Connect with VNC*

Everything before the exercises is a resource to do the exercises.

## 2.2 SSH from Linux or Mac

Secure Shell (SSH) is already installed on almost all Linux and Mac OS X systems.

Open a terminal window by right-clicking on the desktop. Type the part that follows the dollar sign:

```
$ ssh username@ranger.tacc.utexas.edu
```

## 2.3 SSH from Windows

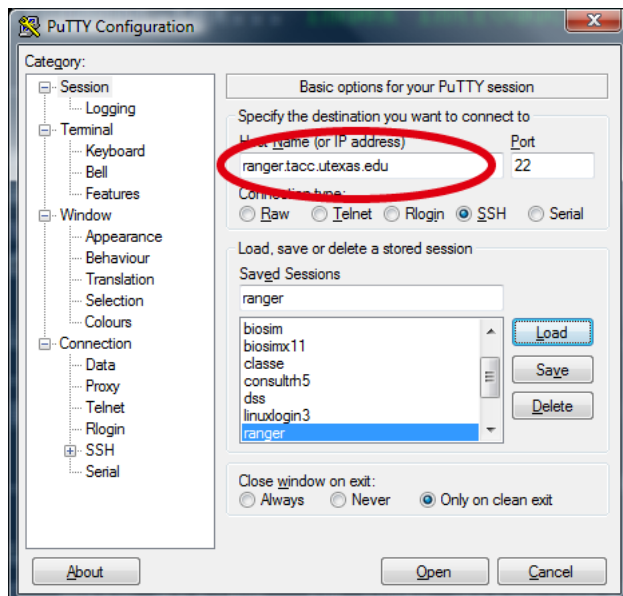
If you are working on a lab computer, skip the part about installation. Putty is installed on lab computers.

### 2.3.1 Installation of Secure Shell

Telnet is disabled for security reasons, but Secure Shell (ssh) clients work nicely as long as they support the SSH2 protocol. A popular client for Windows is the free Putty client, <http://www.chiark.greenend.org.uk/~sgtatham/putty/>. The simplest installation is to download the Windows installer, called putty-0.60-installer.exe, and run it to install putty into your Start menu.

### 2.3.2 Starting a Client Session

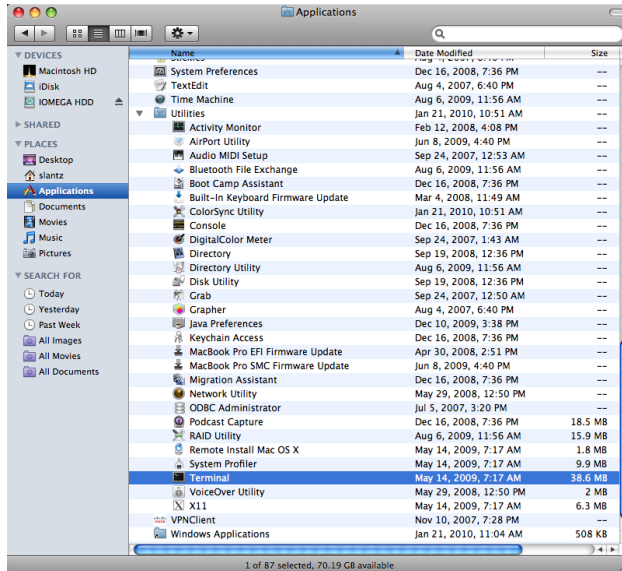
1. Start Putty from the Start Menu.
2. Type a host name, `ranger.tacc.utexas.edu`, where you see the red circle. and hit Enter.



The Ranger login node will respond with its status and the status of your account.

## 2.4 X-Windows from a Mac


If you are on a Mac, the big question is whether you have X11 installed. Navigate in the Finder to the Applications folder and Utilities sub-folder. Do you see an application here called X11, as shown in this picture?



If not, then you need your installation disc. X11 is in the Applications package, and the installer is in the Optional Installs folder.

If you do have X11 installed, then first open a terminal window as shown below. Then follow directions in [X-Windows from Linux](#).

Navigate in the Finder to the Applications folder and Utilities sub-folder. Double-click on the Terminal application to see a Bash command-line.



Continue with the next section, [X-Windows from Linux](#).

## 2.5 X-Windows from Linux

In order to use X-Windows applications on Ranger, you have to turn on *X11 Forwarding* when you make an SSH connection. This is as simple as adding `-X` to the SSH command:

```
$ ssh -X username@ranger.tacc.utexas.edu
```

Now that you are at a prompt at the login node, open an X-Windows application. The dollar sign is called the *prompt*. Type the part that follows it:

```
$ xclock
```

You should see it appear in the corner of your screen if X-Windows is working.

If ever you see a message that “trusted X11 forwarding has failed,” connect again with `-Y` instead of `-X`.

## 2.6 X-Windows from Windows

If you are working on a lab computer, skip the part about installation. Putty and Exceed (not Xming) are installed on lab computers.

### 2.6.1 Installation of X-Windows

You will need to install an X-Windows server on your Windows machine. Common choices include Xming (free), Cygwin/X (free), XFree86 (free), and Exceed.

Xming, available at <http://www.straightrunning.com/XmingNotes/>, is open source, with an available shareware version that has improved graphics. There are two pieces to download.

#### Links and release status

| Website Releases  | Version  | State/Notes     | Released   | MD5 signatures                 | Size MB |
|---|----------|-----------------|------------|--------------------------------|---------|
| <a href="#">Xming</a>   | 7.5.0.17 | Website Release | 7 Mar 2010 | <a href="#">MD5 signatures</a> | 3.77    |
| <a href="#">Xming-portablePuTTY</a>                               | 7.5.0.18 | Website Release | 7 Mar 2010 | <a href="#">MD5 signatures</a> | 1.62    |
| See <a href="#">Donations</a> for how to obtain a Donor Password. |          |                 |            |                                |         |

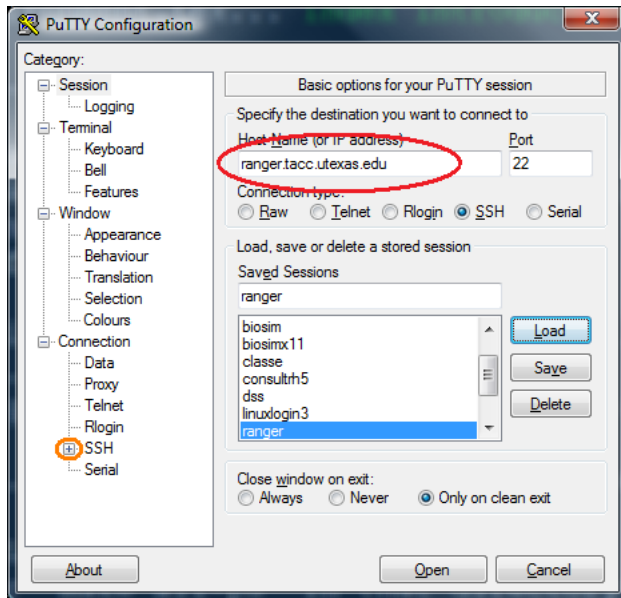
| Public Domain Releases                              | Version  | State/Notes   | Released    | MD5 signature  | Size MB      |
|---|----------|---------------|-------------|--|--------------|
| <a href="#">Xming-fonts</a>                         | 7.5.0.11 | Public Domain | 20 Oct 2009 | 01843baadd75105b962baafa2d08942e                                     | 30.6         |
| <a href="#">Xming</a><br><a href="#">Xming-mesa</a> | 6.9.0.31 | Public Domain | 4 May 2007  | 4cd12b9bec0ae19b95584650bbaf534a<br>e580debbf6110cfc4d8fcd20beb541c1 | 2.10<br>2.50 |

Figure 1: Download and install both Xming-mesa and Xming-fonts.

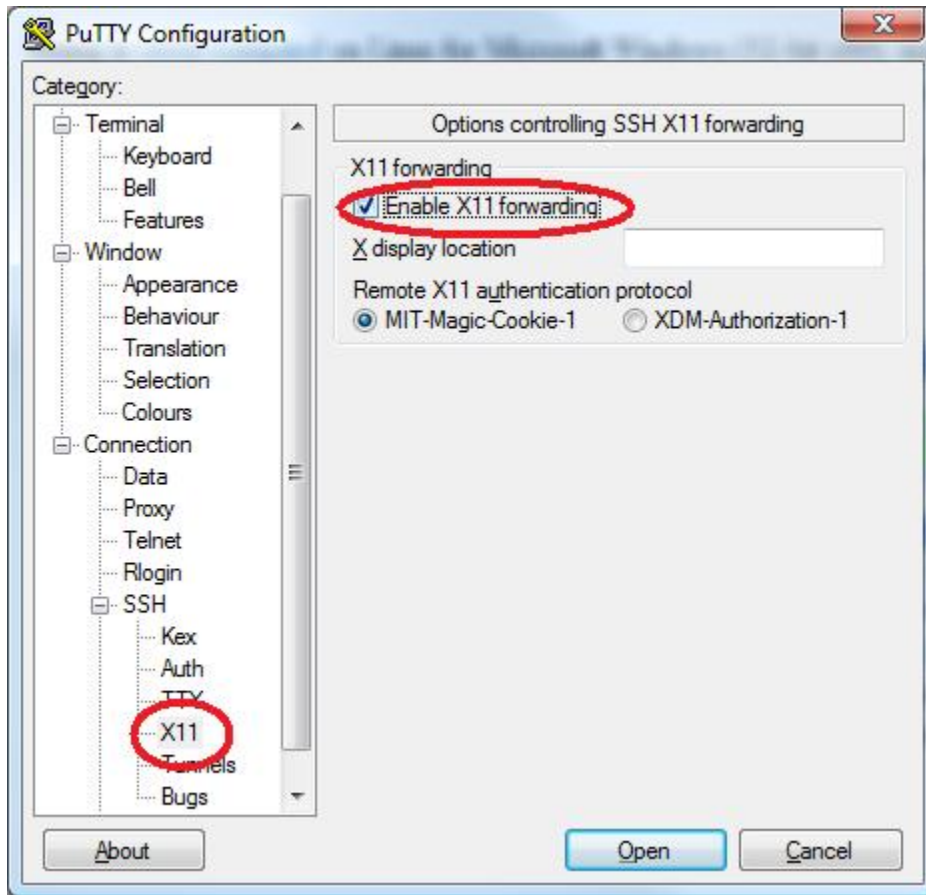
- Xming-mesa (public domain release). There are two links together, one for Xming, one for Xming-mesa. Either will work for this tutorial, but Xming-mesa has some newer features that might come in handy some time.
- Xming-fonts (public domain release) If you purchase the website release of Xming, remember to install the Xming-fonts, as well.

## 2.6.2 Starting a Client Session

1. Before you connect to the login node, start the X-Windows server on your machine. In the Start menu, click the eXceed or Xming icon. The application will briefly appear, then disappear into the background, waiting for windows to show.
2. Start Putty.
3. Type a host name, `ranger.tacc.utexas.edu`, where you see the red circle.



4. Open the SSH menu, in orange on the left of the picture above. Choose X11 settings, shown on the left below, to find the button to “Enable X11 forwarding.”



1. Now use the menu on the left to navigate back to the Session screen. Enter “ranger” under Saved Sessions, and click the Save button. You can use this short cut the next time you connect.
2. Click Open, and it will connect to the head node.

Now that you are at a prompt at the login node, open an X-Windows application. The dollar sign is called the *prompt*. Type the part that follows it:

```
$ xclock
```

You should see it appear in the corner of your screen if X-Windows is working.

## 2.7 Exercise: Use SSH to Connect

*Goal: Connect to the login computer where we will work.*

The Ranger User Guide tells you what the login node names are. You have a username and password for your Ranger account. Use Secure Shell to connect to Ranger. If you don't know how to do that from your machine, look for a tutorial above for help.

- When you login, you will see a splash screen. Read it carefully, this once. What information here will you need later today? Save it somewhere.



## 2.8 Starting: Read Examples of Sessions

*Goal: Familiarize with general command line use.*

The Linux command line may look like an unfair game because there is no indication what to type. The great news is that there are just a few commands that account for almost everything you might want to do. In this exercise, **just read** transcripts to see how work is done.

Following are transcripts of Csh sessions at a command line. The hash, #, is a comment character. The part you would type, but which you are **not going to type right now**, is after the dollar sign. Any responses from Csh follow.

### Create a New Workspace

```
# Comments like this start with hash marks.
# First, where am I?
$ pwd
/home/ajd27
# Make a directory called intro.
$ mkdir intro
# Check that the directories were made
$ ls
bin course intro
# Move into intro.
$ cd intro
# Run a command from another user's directory
$ ~ajd27/bin/lu.S
NAS Parallel Benchmarks (NPB3.2-SER) - LU Benchmark
Size: 12x 12x 12
Iterations: 50
...
# Copy that command to your own bin directory
$ cp ~ajd27/bin/lu.S ~/bin
# See that it worked by listing the bin subdirectory of the parent (..) directory.
$ ls ../bin
farm lu.A lu.C      lu.omp.B lu.omp.S lu.S
htop lu.B lu.omp.A lu.omp.C lu.omp.W lu.W
# Save the command results to a file.
~/bin/lu.S > s0.txt
# Check that you got the file and see its size in bytes.
$ ls -la
total 3428
drwxr-xr-x  5 ajd27 Domain Users   4096 Apr 20 17:05 .
drwxrwxrwx 35 ajd27 Domain Users   4096 Apr 14 15:17 ..
-rw-r--r--  1 ajd27 Domain Users   1932 Apr 20 17:05 s0.txt
# Read the file carefully, using a text reader called less.
# In less: 'b' for back. 'q' to quit. '/' lets you enter a search string.
less s0.txt
```

## 2.9 Further Exercise: Using X-Windows to Connect

*Goal: Use the mouse and menus with remote applications.*

1. Log out of your current SSH session with “logout” or “Ctrl-D.”
2. Start X-Windows on your machine, if necessary. If unsure, look above for information about how X-Windows works on your computer.

3. Use SSH to connect to Ranger, this time enabling X11 Forwarding.
4. Test that your X-Windows connection works by starting “xclock”.

## 2.10 Exercise: Connect with VNC

VNC is a protocol that gives you a shows you a complete desktop, as though Ranger’s login node were your a local computer. The great advantage of VNC is that it can be fast enough to do remote visualization because it transfers only mouse information and bitmap images of the desktop.

Like SSH and X-Windows, a VNC session requires a client and a server. In this case, you have to start a server manually on the Ranger login node. That server will define a *screen* to which you can connect the VNC client, which is installed on your local computer. Each VNC server listens on a TCP/IP port designated by its screen number, so screen 1 listens at port 5901, screen 27 at port 5927.

There are many good VNC clients. On Windows and the Mac, try RealVNC or TightVNC. On Linux, try `vinagre` or `vncviewer`.

Login to Ranger, if you have not already done so. First set your vnc password, which you will need when your client connects to the server. You are free to choose a simple password. Then start the server:

```
$ vncpasswd
$ vncserver
New 'login3.ranger.tacc.utexas.edu:1 (train200)' desktop is
login3.ranger.tacc.utexas.edu:1
Starting applications specified in /share/home/0002/train200/.vnc/xstartup
Log file is /share/home/0002/train200/.vnc/login3.ranger.tacc.utexas.edu:1.log
```

The message above indicates that this VNC session is on screen 1 of login3. We can connect to that from our local VNC client using the address `login3.ranger.tacc.utexas.edu:1`.

When you are done with the session, be sure to kill the VNC server with:

```
% vncserver -kill :1
Killing Xvnc process ID 11406
```

Note that `:1` represents screen 1.

It is possible to tunnel VNC through Secure Shell for a safer connection. If you are using TightVNC over long distances, you can greatly speed the connection by increasing compression. Another trick is to offer your VNC server address and password to a colleague in order to share a display.

## 2.11 Further Exercise: Choose Your Shell

The default shell is `/bin/csh`. If you prefer another shell, now is a good time to change it. List available shells with:

```
$ chsh -l
```

and then change your shell, for instance to `tcsh`, with:

```
$ chsh /bin/tcsh
```

Every shell reads configuration files on login. The two main types of shells, the Csh family and Bash family, are configured on Ranger so that your personal configurations belong in `$HOME/.cshrc_user` or `$HOME/.login_user`, respectively.

## 2.12 Advanced: Make Login Faster

There are two things you can do to make it easier to login to ranger. One is to make a shortcut, or alias, to the login node. The other is to create SSH keys that will allow you to login to the node without typing a password.

### 2.12.1 Making Shortcuts on Windows

*Goal: Open SSH terminal window with a single click.*

If you are working on your laptop, make a shortcut or alias to login to `ranger.tacc.utexas.edu`.

For Windows and Putty, first create a saved session. When you first open Putty, it gives you a list of “saved sessions.” Enter a name in the box above that list, and click the save button. Then make a shortcut with the `-load` option, where *ranger* is the whatever you called your session:

```
C:\util\putty.exe -load ranger
```

### 2.12.2 Setup SSH Keys for No Password Login

*Goal: Make login not require typing a password.*

This can take a few minutes if you have not done it before.

Start a key agent on your local computer which then logs in for you. You have to type your pass phrase once when you start your local computer. There is an excellent tutorial already at <http://www.mtu.net/~engstrom/ssh-agent.php> which covers Linux, the Mac, and Windows.

## 3 Compile

### 3.1 Goals, Resources, Agenda

Goals: Compile a parallel application on Ranger. Know how to find libraries for that application and choose which library versions your application will use.

Resources:

- Ranger User Guide - <http://services.tacc.utexas.edu/index.php/ranger-user-guide>
- Linux man command - The help command on linux is called “man”, for manual. Find help on the command “module” with “man module”.
- Google - For most Linux commands, a Google search for “Linux shell” and the name of the command will be fruitful.

Agenda: In order to complete this section, you must do these sections:

1. *Exercise: Use Modules to Load Boost Libraries*
2. *Exercise: Compile a Benchmark*

### 3. (Optional) *Discussion Exercise: Learn Modules*

The other sections are resources to help with exercises.

## 3.2 About the Module Command

The `module` command changes variables in your shell in order to choose which applications you use. For instance, Ranger has about twenty different versions of MPI because there are MPI builds for two compilers, three versions of MPI, and for past and development builds.

To load or unload a module, the `module` command changes environment variables the shell uses to choose what it runs.

- **\$PATH** - list of directories where executables are found.
- **\$MANPATH** - list of locations to find help files.
- **\$LD\_LIBRARY\_PATH** - list of directories where applications can find shared libraries.

Each loaded module may define a new set of environment variables. For instance, loading the Intel Math Kernel Libraries (MKL) defines variables you can use in a makefile.

- **\$TACC\_MKL\_DIR** - The base directory of the MKL installation.
- **\$TACC\_MKL\_DOC** - Location of documentation.
- **\$TACC\_MKL\_LIB** - Library directory.
- **\$TACC\_MKL\_INC** - Include directory for headers.

The `module` command tells you about possible conflicts. For instance, you will have the Intel compilers loaded when you start. If you add the PGI compilers, you will see a message that Intel is already loaded. Heed that advice, and delete the Intel compilers when you add the PGI compiler.

The basic `module` commands are

- **module avail** - List available modules
- **module avail string** - Search for the module “string”.
- **module list** - Show currently-loaded modules.
- **module add** - Add a module to the environment.
- **module del** - Remove a module from the current environment.
- **module help string** - Get help on a module named “string”.
- **module purge; module load TACC** - Return to original state.

Some versions of `module` require you to load them in a particular order. The version installed (and developed) at TACC does not.

## 3.3 How to Read a Makefile

A Makefile is a script you can run with the “make” command. When you type “make”, the “make” program looks for a file called Makefile, or makefile, and searches it for a list of things to do. The list has targets, dependencies, and commands:

```
target: dependency, dependency
      command..
      command..
```

For instance, if you want to create a Makefile to compile a single file into an executable, the Makefile would contain:

```
xhpl: xhpl.c
      cc xhpl.c -o xhpl
```

The `cc` command reads the dependency, `xhpl.c`, to create the target, `xhpl`. You execute this with the command “`make xhpl`”.

Makefiles can be much more complicated than this. In these tutorials, you will see Make used to execute all kinds of commands besides compilation. In every case, Make is a shorthand to execute the same commands you could type at the command line.

### 3.4 Modules List

This is a partial list of modules. Some entries that appear more than once on Ranger are here only in one section.

- `/opt/apps/pgi7_2/mvapich1_1_0_1/modulefiles`
  - `amber` - Molecular dynamics of proteins and nucleic acids
  - `arpack` - Fortran subroutines to solve large scale eigenvalue problems.
  - `charm++` - Object-oriented language in C++ with adaptive MPI (NAMMD, cosmology, QCD).
  - `espresso` - electronic-structure calculations and materials modeling at the nanoscale.
  - `fftw2` - A fast, free C FFT library; includes real-complex, multidimensional, and parallel transforms.
  - `fftw3` - A fast, free C FFT library; includes real-complex, multidimensional, and parallel transforms.
  - `gamess` - ab initio quantum chemistry.
  - `gromacs` - molecular dynamics for biomolecular systems.
  - `gulp` - General utility lattice program for 3D periodic solids, gas phase clusters, and more.
  - `hypre` - library for solving large, sparse linear systems of equations
  - `ipm` - integrated performance monitoring.
  - `kojak` - performance-analysis tool for parallel applications supporting the programming models MPI, OpenMP, SHMEM, and combinations thereof
  - `mpiP` - lightweight profiling library for MPI applications.
  - `pdtoolkit` - Program database toolkit for analyzing source code.
  - `petsc` - data structures and routines for the scalable (parallel) solution of scientific applications modeled by partial differential equations.
  - `phdf5` - Parallel data model, library, and file format.
  - `plapack` - Parallel linear algebra matrix manipulations.

- pmetis - parallel graph partitioning.
  - pnetcdf - parallel interface, library, and format for data.
  - scalapack - linear algebra routines using MPI.
  - scalasca - open-source analysis of parallel applications.
  - slepc - Scalable library for eigenvalue problems.
  - sprng - scalable parallel pseudo random number generation.
  - tao - toolkit for advanced optimization.
  - trilinos - algorithms for large-scale, complex multi-physics engineering and scientific problems.
- /opt/apps/pgi7\_2/modulefiles
    - acml - AMD Core Math Library, lapack.
    - autodoc - docking tools for small molecule binding to known 3D structure.
    - glpk - GNU linear programming kit.
    - gotoblas - Hand-tuned Basic linear algebra subroutines.
    - hdf5 - data model, library, and file format.
    - hecurea - MPI-2 implementation for Infiniband
    - mvapich - MPI-1 implementation for InfiniBand
    - nco - Programs for manipulating and analyzing NetCDF files.
    - netcdf - interface, library, and format for data.
    - openmpi - MPI-2 implementation for Infiniband
  - /opt/apps/teragrid/modulefiles
    - ctssv4 - Coordinated TeraGrid software and services.
    - globus-4.0 - Toolkit for building grids.
    - teragrid - presumably TeraGrid tools.
    - apache-ant - A Java build tool.
    - condor-g - Grid version of Condor scheduler.
    - condor - Methods for high throughput computing.
    - gridshell - login shell for grid computing.
    - gsissh - SSH for TeraGrid.
    - gx-map - Utilities to maintain grid security directory.
    - koomie - Internal for testing software builds.
    - mycluster - Virtual clusters on demand.
    - pacman - Pacman lets you define, install, configure and setup software.
    - srb-client - Storage request broker client.
    - tg-policy - Displays TACC teragrid policy.
    - tginfo - TeraGrid Information Service Discovery Client

- `tgproxy` - Helps you work with TeraGrid credentials.
- `tgresid` - TeraGrid resource identification
- `tgusage` - query TeraGrid allocations.
- `uberftp` - Grid-enabled FTP client.
- `/opt/apps/modulefiles`
  - `R` - functional language for statistical computing.
  - `autotools` - configuration of software.
  - `beta` - unknown
  - `binutils-amd` - Tools to work with binaries.
  - `cmake` - cross-platform, open-source build system.
  - `ddt` - large-scale parallel debugger
  - `gcc` - GNU compiler collection.
  - `git` - version control system
  - `gmake` - GNU make program.
  - `gmp` - Gnu Multiple Precision Library
  - `gnuplot` - 2D and 3D plotting.
  - `gsl` - GNU scientific library.
  - `gzip` - Open-source compression.
  - `intel` - Intel compilers
  - `irods` - Integrated Rule-oriented Data System.
  - `launcher` - TACC Parametric Job Launcher
  - `lua` - Programming language.
  - `mkl` - Intel Math Kernel Library.
  - `mpfr` - C library for multiple-precision floating point.
  - `mysql` - Open-source database.
  - `ncl_ncarg` - NCL/NCAR graphics library
  - `numpy` - Tools for numerical computing on Python.
  - `papi` - Interface for to use hardware performance counters.
  - `perfexpert` - Performance evaluation tool developed at TACC.
  - `pgi` - Portland Group compilers for C and Fortran.
  - `plplot` - Library to create scientific plots.
  - `postgres` - Open Source database.
  - `python` - Scripting language.
  - `star-ccm` - CAD preparation, meshing, and model setup.
  - `subversion` - version control system.
  - `sun` - Loads the Sun compiler environment.

- tar - Archive creation and manipulation.
- vis - Makes all of the visualization modules visible.
- zlib - compression library.
- /opt/modulefiles
  - java - programming language.
- /opt/apps/vis/modulefiles
  - VTSSV3 - TeraGrid visualization tools?
  - amira - Visualization stressing life sciences.
  - blender - Open Source 3D creation suite.
  - chromium - scalable rendering for clusters of workstations.
  - cuda-sdk - Software development kit for GPU programming.
  - ensight - The “extreme” visualization post-processor.
  - glew - OpenGL extension wrangler library.
  - glui - OpenGL User Interface library.
  - glut - OpenGL utility kit.
  - gmt - 60 unix tools to manipulate (x,y) and (x,y,z) data sets.
  - idl - Visualization programming language.
  - imagemagick - Convert, edit, compose images.
  - mesa - OpenGL offscreen rendering.
  - mplayer - Movie player.
  - netpbm - command-line utilities for image manipulation.
  - paraview - Open source scientific visualization.
  - qt - Application windowing toolkit.
  - sdl - Simple DirectMedia Layer for graphics and audio.
  - silo - file format used by VisIt.
  - vapor - Visualization package for ocean, atmosphere, solar.
  - visit - Open source, 3D visualization tool.
  - vtk - Visualization library, open-source, scriptable from C++, perl, python, java, tcl.

### 3.5 Discussion Exercise: Learn Modules

*Goal: Find out what is installed.*

Scan the list of modules in the next section for ones you recognize and think about the following questions, then discuss with the person beside you.

1. How many do you recognize as free software and which cost money?
2. Which scientific disciplines are well-represented?
3. How do you think applications and libraries are chosen?



If you need for work an application not already installed, you will be able to install it in your home directory.

### 3.6 Exercise: Use Modules to Load Boost Libraries

*Goal: Understand how to load the module you want.*

For this exercise, imagine you want to compile a program that uses the Boost libraries for C++. You were told it is on the system, now let's find and load it.

1. Type "module" at the command prompt, without any arguments, to see its help. Use some of these commands in the next steps. You can play with "load", "del", "list", and "avail" without doing irreparable damage. You can return to your original state with the command "module purge; module load TACC".
2. Can you load boost yet?
3. There are several compilers on TACC, and the module command helps you choose which one to use at any time. Which one are you using now? What others are there to use if you wanted them? There is a list above that describes what the various modules are.
4. The list of available modules separates sub-lists into sections. Take a look at the section titles to see how the available modules are organized.
5. What could you do to look for available modules that are not currently listed? Find and load the boost module.

### 3.7 Exercise: Compile a Benchmark

*Goal: Use modules to choose tools.*

You have to compile the HP Linpack benchmark suite, a suite sometimes used to test supercomputers. This code requires three kinds of libraries.

- C Compiler - Three compilers are installed on this system, two of which are listed by the module command.
- BLAS Library - Basic Linear Algebra Subroutines, which our program will use to invert large matrices.
- MPI - The message passing interface, which our program will use to run on many cores at the same time.

You can choose one of two combinations of libraries for your code.

1. The PGI version uses PGI compilers, the ACML BLAS library, and Mvapih MPI.
2. The Intel version uses Intel compilers, the MKL BLAS library, and Mvapih MPI.

You choose one of these combinations by using the module command to ensure the modules you want loaded are there and the other ones are not loaded.

For this particular library, you the Makefile offers these commands.

- "make arch=pgi\_acml" - This compiles using PGI compilers, ACML, and Mvapih.
- "make clean arch=pgi\_acml" - If you tried to compile and failed, this command deletes all intermediate files of the pgi\_acml build in order to give you a fresh start.
- "make arch=icc\_mkl" - This compiles using the Intel compilers, MKL, and Mvapih.

- “make clean arch=icc\_mkl” - If you tried to compile and failed, this command deletes all intermediate files of the icc\_mkl build in order to give you a fresh start.

Here are the steps to do.

1. Copy code from the train400 directory with “cd && tar zvxvf ~train400/envi.tgz”.
2. Choose either the PGI or Intel options, and load the correct modules for that option.
3. Use Make to compile the program. If the compilation ends with the word “Error”, then it failed to compile. The most popular reason for failure is that the correct modules aren’t loaded. If so, clean the directory, correct loaded modules, and try again.
4. Look in ~/hpl-2.0/bin/icc\_mkl or ~/hpl-2.0/bin/pgi\_acml for the program, xhpl, which you just built.

## 4 Run on Ranger

### 4.1 Goals, Resources, Agenda

Goals: Construct a job to run, run it, see it running, and check that it ran.

Resources:

- Queue information is in the Ranger User Guide - <http://services.tacc.utexas.edu/index.php/ranger-user-guide>
- *Text Editors on Ranger*
- *How to Construct a Batch File*
- *How to Monitor a Job on Ranger*

Agenda: You must do the following to complete this section:

1. *Exercise: Run in Batch*
2. (Optional) *Further Exercise: Using Scheduler-Defined Variables*

### 4.2 Text Editors on Ranger

There are several text editors available on Ranger.

- nano - Small and quick. Use Ctrl-o to save, Ctrl-x to quit.
- vi - Quick, powerful. You can get by in this with a cheat sheet below.
- emacs - Less quick, also powerful.

#### 4.2.1 VI Cheat Sheet

Vi has two modes, *insertion mode* for editing and *command mode* for saving. The editor starts in command mode.

In command mode,

- :w writes a file to disk
- :q quits

- :q! quits even if you didn't save changes
- i puts you in insertion mode.

In insertion mode,

- Cursors and typing work as expected.
- Escape key returns to command mode.

Ctrl-c will not quit this program.

For more, ask the Google about [Vi cheat sheets](#).

### 4.3 How to Construct a Batch File

Ranger uses the Sun Grid Engine scheduler. Submit jobs with qsub:

```
$ qsub batch.sh
```

You can put arguments to qsub on the command line, but they usually go in the batch files as comments at the start of the file where the scheduler searches for all lines that begin with #\$.

| Argument | Description  |
|----------|--|
| #\$ -N   | The job name, available to the script as \$JOB_NAME        |
| #\$ -cwd | Start the job in the same current working directory.       |
| #\$ -o   | The name of the output file.                               |
| #\$ -j y | Join the stdout and stderr streams in the output file.     |
| #\$ -A   | Your account name, listed at login.                        |
| #\$ -q   | Which queue to use, found in the Ranger User Guide.        |
| #\$ -pe  | The wayness of the job, as in 16way, and core count.       |
| #\$ -V   | Set job variables as though it were a child of this shell. |
| #\$ -l   | Specify resources, such as h_rt=00:10:00 for run time.     |

For example, a job to run on 64 cores, which is 4 nodes, for 10 minutes might look like:

```
#$ -N ht3d-hyb
#$ -cwd
#$ -o $JOB_NAME.o$JOB_ID
#$ -j y
#$ -A 20100519HPC
#$ -q development
#$ -pe 16way 64
#$ -V
#$ -l h_rt=00:10:00
set -x
ibrun $HOME/bin/xhpl
```

The `ibrun` command at the bottom starts a job under MPI using settings from the `-pe` line. The second argument on that line is the number of cores for your job. Given sixteen cores per node on all Ranger nodes, the number of nodes is the number of cores divided by sixteen.

*Wayness* tells `ibrun` how many processes to start on each node you have requested. You might request `1way` if you have a distributed, multi-threaded code. You might try `15way` or `14way` if you want to leave a free core to handle intensive I/O on a node.

The scheduler stores all stdout and stderr from your batch script when it executes on the head node of the job. Those files are stored, by default, in the directory where you executed `qsub`, unless you specify otherwise with the `-o` option.

## 4.4 How to Monitor a Job on Ranger

When you submit a job, the scheduler assigns it a *state* and changes that state over time. The four basic states are

- **Unscheduled** - Likely a sign something is wrong.
- **DepWait** - This job is set to run after another finishes.
- **Waiting** - It's in the queue and waiting for resources to become available.
- **Running** - As far as the scheduler is concerned, everything is working.

The main command for monitoring your own jobs is `qstat`.

| <b>qstat argument</b>               | <b>description</b>             |
|-------------------------------------|--------------------------------|
| <code>qstat</code>                  | Show status of only your jobs. |
| <code>qstat -j &lt;jobid&gt;</code> | Show details about one job.    |
| <code>qstat -ext</code>             | Extended job information.      |
| <code>qstat -r</code>               | Resource requirements.         |

The `showq` command will show jobs from all users.

| <b>showq argument</b>              | <b>description</b>             |
|------------------------------------|--------------------------------|
| <code>showq</code>                 | Show summary of jobs in queue. |
| <code>showq -sfb +priority</code>  | Sort by increasing priority.   |
| <code>showq -sfa -queuetime</code> | Sort by decreasing queue time. |

The possible sorting arguments are `priority`, `starttime`, `queuetime`, `wclimit`, `jobname`, and `remaining`.

To cancel a failing job, use `qdel`.

| <b>qdel argument</b>               | <b>description</b>               |
|------------------------------------|----------------------------------|
| <code>qdel &lt;jobid&gt;</code>    | Cancel a waiting or pending job. |
| <code>qdel -f &lt;jobid&gt;</code> | Cancel a running job.            |

## 4.5 Exercise: Run in Batch

*Goal: Monitor a job and evaluate its success.*

You compiled code for HPL-2.0 earlier. It's time to run that code.

1. Go to the same directory as the HPL-2.0 binary with “`cd ~/hpl-2.0/bin/icc_mkl`” or “`cd ~/hpl-2.0/bin/pgi_acml`”, depending on which you built.
2. There is a template for a batch file in both of those directories. Compare the template batch file, `ranger.sh`, with that of Section [How to Construct a Batch File](#). You want your job to use 16 cores on 1 node. Edit the file so that it has the right account and choose a queue. Where can you find a list of available queues to choose one?
3. Before you submit your job, look at the section [How to Monitor a Job on Ranger](#) to see what commands you will use to see whether your job is running.
4. Submit `ranger.sh` using the `qsub` command. Remember its job id.
5. Watch the job with `qstat`.
6. Can you tell when it will run? Mostly, not, so can you tell how many jobs are waiting?
7. When the job finishes, do you see its output file in the same directory from which you submitted the file?

## 4.6 Discussion: Why the Long Queue?

You can get a list of queues with `qconf -sql` and look at details of a particular queue with `qconf -sq queueName`.

1. Look at the `hostlist` and `user_lists` of development, normal, and large queues. Why make different host lists or user lists? What behavior are they trying to encourage?
2. Do any queues have higher or lower priority according to `qconf`? The `qstat -g c` command shows how many slots each queue is permitted to use. What incentives are there to prefer certain queues, and what are those incentives?
3. What are the consequences of a 24 hour max queue time?

## 4.7 Further Exercise: Using Scheduler-Defined Variables

The scheduler defines variables for the job to use when it runs on the node. The simplest are `$JOB_ID`, which is the number the scheduler assigns to the job and `$JOB_NAME`, which is the name you give with the `-N` switch.

You can see what variables are defined by submitting a batch script that has the typical header information but whose body contains the following:

```
env|sort>variables.txt
```

This will show many variables which begin with `$PBS`. Those that start with `$PBS_O` refer to the state of the shell when you submitted the job. So try submitting a job and take a look.

# 5 Transfer and Archive Files

Now you have run your job and begun to generate lots of data. Where do you put it? How do you get it home?

## 5.1 Goals, Resources, Agenda

Goals: After this section, you will know how to transfer files to and from Ranger. You will be able to choose where to put your files for medium and long term batch calculations.

Resources:

- Tape archive and BBCP information are in the Ranch User Guide <http://services.tacc.utexas.edu/index.php/ranch-user-guide>.
- The GridFTP tools are described at <https://www.teragrid.org/web/user-support/gridftp>, and Ranger's servers are described in the Ranger User Guide.
- The Windows Secure Copy client is at <http://tartarus.org/~simon/putty-snapshots/x86/pscp.exe>.
- *Filesystems on Ranger*

Agenda: To complete this section

1. *Exercise: Retrieve Files with Secure Copy*
2. *Discussion: Transfer Techniques*

3. (Optional, not during lab) *Exercise: File Transfer with BBCP*

## 5.2 Filesystems on Ranger

Ranger has three main filesystems, available from both login nodes and compute nodes.

| Filesystem | Total Size | Per User Quota | Shortcut | Retention Policy     |
|------------|------------|----------------|----------|----------------------|
| \$HOME     | ~100 TB    | 6 GB           | cd       | Nightly backup       |
| \$WORK     | ~200 TB    | 350 GB         | cdw      | No backup            |
| \$SCRATCH  | ~800 TB    | 400 TB         | cds      | Purged every 10 days |

These filesystems reside on 73 Sun x4500 I/O servers and 4 metadata servers, served by Lustre. They are highly parallel and very fast.

You can query quota information about individual directories with the `lfs` command. For instance, to see your quota on `$WORK`:

```
$ lfs quota -u $USER $WORK
```

Find your usage either from the login splash screen or with the `disk usage` command. This finds the number of megabytes under `$HOME`:

```
$ du -sm $HOME
1316    /share/home/00692/train00
```

## 5.3 Tape Archive

There is a 1.7 PB tape archiver, called Ranch, attached to Ranger. It is available for you to use as needed. To access it, use `secure copy` or the `bbcpc` command to copy files to your directory under the machine. The syntax for copying to your archive directory on Ranch is:

```
$ scp filename ${ARCHIVER}:${ARCHIVE}
```

For more on this, look at the [Ranch User Guide](#).

## 5.4 File Transfer with GridFTP

You could do this exercise on your own computer later. It requires an installation of GridFTP on the local machine.

GridFTP is a protocol designed to transfer files using grid credentials. With GridFTP, you can

- Login once with your TeraGrid credentials and transfer files many times without passwords.
- From one machine, execute a command that transfers files between two other machines.
- Encrypt files as they are sent for extra security.
- Transfer files quickly using many other features.

There are several applications that can do GridFTP, all described at <https://www.teragrid.org/web/user-support/gridftp>. You would have to install one on your local computer in order to exchange files with Ranger. Ranger's gridftp address is `gridftp.ranger.tacc.utexas.edu`.

## 5.5 Exercise: Retrieve Files with Secure Copy

This is a first experience with Secure Copy to show how easy it is to transfer files back and forth. An alternative, also easy, method is to use Secure FTP, or SFTP. On Windows, try WinSCP. On the Mac, try Fetch.

Find the Linux and Mac or Windows sections, as appropriate.

### 5.5.1 From Linux or Mac

1. Open a terminal window with a right-click on the desktop. You should now have two windows open, one that is on the local machine and one that is logged into Ranger.
2. On Ranger, create a 100 MB file with `cd && dd if=/dev/zero of=big bs=100M count=1`.
3. Do an `ls` to find out how large this file is.
4. On your local machine, copy that file from Ranger with `scp`. Put your username where the example shows `train2xx`. The single period at the end represents the current directory on the local drive as the place to store the retrieved file.

```
[local]$ scp username@ranger.tacc.utexas.edu:big .
```

About how long did that take? Now copy the file back and use the `time` command to time it.

```
[local]$ time scp big username@ranger.tacc.utexas.edu:big
```

That should transfer files from the local machine to the remote one.

### 5.5.2 From Windows

On Windows, you need to install a program called Pscp in order to transfer files with Secure Copy. This program is part of the suite with Putty, the SSH client.

1. You already have one Putty window open and logged into Ranger. Now open a local command line by going to the Start Menu, looking under All Programs, Accessories, and then open "Command Prompt". In what directory is this command prompt?
2. Download `pscp` into the same directory as the local command prompt. Do this by right-clicking on this link, <http://tartarus.org/~simon/putty-snapshots/x86/pscp.exe>, and choosing "save as..".
3. On Ranger, create a 100 MB file with `cd && dd if=/dev/zero of=big bs=100M count=1`.
4. Do an `ls` to find out how large this file is.
5. Run `pscp` in the local command prompt.

```
[local]$ pscp username@ranger.tacc.utexas.edu:big .
```

About how long did that take? Now copy the file back.

```
[local]$ pscp big username@ranger.tacc.utexas.edu:big
```

That should transfer files from the local machine to the remote one.

## 5.6 Discussion: Transfer Techniques

- Do you know of faster ways to transfer data?
- How long would it take to transfer a terabyte (a thousand gigabytes) with the fastest method?
- How could you avoid transfer of data?

## 5.7 Exercise: File Transfer with BBCP

You could do this exercise on your own computer later. It requires an installation of BBCP on the local machine.

BBCP is a program that functions much like secure copy, or SCP, except that it has several enhancements to transfer files much more quickly. In order to use this to transfer files to or from Ranger, you would have to install it on your local machine, and it seems to install best on Linux or other Unixes. Darwin builds for Mac are possible, too.

Download the source code from <http://www.slac.stanford.edu/~abh/bbcp/bbcp.tar.Z>. To copy a file to ranger, use it like scp:

```
[local]$ bbcp filename login3.ranger.tacc.utexas.edu:
```

More information is in the Ranger user guide, under [High speed transfers](#).