

# Getting Started with CISL Facilities and Support

CESM Tutorial  
July 30, 2012

Si Liu  
NCAR/CISL/OSD/USS  
Consulting Services Group

# CISL's Mission for User Support

CISL will provide a balanced set of services to enable researchers to utilize community resources securely, easily, and effectively.

CISL Strategic Plan

CISL also supports special colloquia, workshops and computational campaigns; giving users special privileges and access to facilities and services above normal service levels.



NCAR



# CISL Facilities Overview

Navigation and usage of the facilities require a basic familiarity with a number of the functional aspects of the facilities.

- Computing Systems
  - Bluefire (Mesa)
  - Lynx (Mesa)
  - Janus (CU)
  - **Yellowstone (NWSC)**
- Data Analysis and Visualization
  - Mirage and Storm
  - **Geyser and Caldera (NWSC)**
- Data Archival
  - HPSS
  - GLADE
- Allocations and Security
- User support
- Training

# Working with Bluefire



# Computing System - Bluefire

- **IBM clustered Symmetric Multi-Processing (SMP) system**
  - Operating System: AIX (IBM-proprietary UNIX)
  - Batch system: Load Sharing Facility (LSF)
  - File system: General Parallel File System (GPFS)
- **127 32-way 4.7 GHz nodes**
  - 4,064 POWER6 processors
  - SMT enabled (64 SMT threads per node)
  - 76.4 TFLOPS
- **117 compute nodes (70.4 TFLOPS peak)**
  - 3,744 POWER6 processors (32 per node)
  - 69 compute nodes have 64 GB memory
  - 48 compute nodes have 128 GB memory
- **10 other nodes**
  - 2 interactive sessions/login nodes (256 GB memory)
  - 2 debugging and share queue nodes (256 GB memory)
  - 4 GPFS/VSD nodes
  - 2 service nodes

# Compilers on Bluefire

- **Fortran 77 and Fortran 90/95 compilers:**
  - xlf, xlf\_r, xlf90, xlf90\_r
  - mpixlf, mpixlf\_r, mpixlf90, mpixlf90\_r
- **C and C++ compilers:**
  - xlc, xlc\_r, x1C, x1C\_r
  - mpcc, mpcc\_r, mpCC, mpCC\_r
- **The \_r versions are thread safe**
  - We recommend them over the non \_r versions.
- **Compile your source code**
  - xlc\_r hello\_world.c -o hello\_world\_c.exe
  - xlf90\_r hello\_world.f -o hello\_world\_f.exe
- **More information**

<http://www2.cisl.ucar.edu/docs/bluefire/compiling-and-optimization>

# Login to Bluefire with Yubikey

- **Security Shell(SSH)**
  - Cygwin, Putty, Terminal, etc.
- **Using your Yubikey token**
  - When you log in Bluefire,  
`ssh your_logon@bluefire.ucar.edu`  
your screen displays a response:  
`Token_Response:`
  - Enter your PIN number on the screen (**do not hit enter**), then touch the yubikey button. This will insert a new one-time password (OTP) and a return.
  - The yubikey is activated by the **warmth of your finger** not the pressure in pushing the button.
- **More information of Yubikey:**
  - <https://www2.cisl.ucar.edu/docs/enabling-your-yubikey-token>

# A Job Script on Bluefire

```
#!/bin/csh
# LSF batch script to run an MPI application

#BSUB -P 12345678 # project number (required)
#BSUB -W 1:00 # wall clock time (in minutes)
#BSUB -n 256 # number of MPI tasks
#BSUB -R "span[ptile=64]" # run 64 tasks per node
#BSUB -q workshop # queue setting: use 'workshop' queue this week
#BSUB -J myjob # job name
#BSUB -o myjob.%J.out # output filename
#BSUB -e myjob.%J.err # error filename

mpirun.lsf /usr/local/bin/launch ./MyProg.exe
```

**For more examples, see the /usr/local/examples directory.**





- # Submit, Delete, and Monitor Jobs on Bluefire
- **Job submission**
    - `bsub < script`
  - **Monitor jobs**
    - `bjobs`
      - `bjobs -u all`
      - `bjobs -q regular`
    - `bhist`
      - `bhist -n 0 jobid`
  - **Delete a job**
    - `bkill jobid`



# Module Utility on Bluefire

Modify environment to find alternative compilers or software

- To show all available module files  
`module av`
- To see which modules are in force  
`module list`
- To load a new module (e.g. ImageMagick)  
`module load ImageMagick-6.5.3-10`

<https://www2.cisl.ucar.edu/docs/bluefire/getting-started?tab=3>

# "Big 3": Get a Better Performance

- **Simultaneous Multi-Threading(SMT)**

- a second, on-board "virtual" processor
- 64 virtual cpus in each node

- **Multiple page size support**

- 64-KB page size when running the 64-bit kernel
- Large pages (16 MB) and "huge" pages (16 GB)

- **Processor binding**

<http://www2.cisl.ucar.edu/docs/bluefire/running-jobs?tab=3>

# Lynx



# Computing System – Lynx

## Single-cabinet Massively Parallel Processing Supercomputer

- Operating system: Cray Linux Environment
  - Compute Node Linux (CNL) – based on SuSE Linux SLES 10
- Batch System:
  - MOAB workload manager
  - Torque (aka OpenPBS) resource manager
  - Cray's ALPS (Application Level Placement Scheduler)
- File system: Luster file system

# Computing System – Lynx

- 76 compute nodes (**8.026 TFLOPS peak**)
  - 12 processors per node, 912 compute processors
  - Two hex-core AMD 2.2 GHz Opteron chips
  - Each processor has 1.3 GB of memory and totaling 1.216 TB of memory in the system.
  
- 10 I/O nodes
  - A single dual-core AMD 2.6 GHz Opteron chip and 8 GB of memory
  - 2 login nodes, 4 nodes reserved for system functions
  - 4 nodes are for external Lustre file system and GPFS file system testing

# Compilers on Lynx

- **Available compilers**
  - PGI Programming Environment
  - EKOPATH ("PathScale") Compiler Suite
  - Intel Compiler Suite
  - GNU Compiler Collection
- **Load the corresponding PrgEnv-xxx module to change compilers**
  - ftn: Fortran programs
  - cc: C programs
  - CC: C++ programs
- **More information**
  - <https://www2.cisl.ucar.edu/docs/lynx/compilers>

# A job script on Lynx

```
#!/bin/bash  
#PBS -q workshop  
#PBS -l mppwidth=60      ###Number of Processors  
#PBS -l walltime=01:30:00  
#PBS -N example  
#PBS -e testrun.$PBS_JOBID.err  
#PBS -o testrun.$PBS_JOBID.out  
  
cd $PBS_O_WORKDIR  
aprun -n 60 ./My_Prog.exe
```



# Submit, Delete, and Monitor Jobs on Lynx

- **Submit a job**
  - `qsub batch_script`
- **Check job status**
  - `qstat -a`
  - `showq`
- **Delete a job**
  - `qdel jobid`

# HPSS



# HPSS Introduction

- High Performance Storage System (12+ PB of data)
- Hierarchical Storage Interface (HSI) is the primary interface supporting for data transfer to/from HPSS along with metadata access and data management.
- HPSS Tape Archiver (HTAR) is used to package files on your file system to a single archive file and then send it to HPSS.
- HPSS files have NO expiration date. They remain in the archive until they are explicitly deleted. Once deleted, they cannot be recovered.

# Hierarchical Storage Interface (HSI)

- **POSIX like interface**
- **Different ways to invoke HSI**
  - Command line invocation
    - **hsi** cmd
    - **hsi** cget hpssfile (from your default dir on HPSS)
    - **hsi** cput myfile (to your default dir on HPSS)
  - Open an HSI session
    - hsi to get in and establish session; end, exit, quit to get out
    - restricted shell-like environment
  - hsi "in cmdfile"
    - File of commands scripted in "cmdfile"
- **Navigating HPSS while in HSI session**
  - pwd , cd, ls, cdl
  - <http://www2.cisl.ucar.edu/sites/default/files/HSI-command.pdf>

# Data Transfer

- **Writing data – cput command**

- `[HSI]/home/user1 > cput file.01`
- `[HSI]/home/user1 > cput file.01 : new.hpss.file`

- **Reading data – cget command**

- `[HSI]/home/user1-> cget file.01`
- `[HSI]/home/user1-> cget file.01 : hpss.file`

# GLADE centralized file service

- **The Globally Accessible Data Environment**
- **High performance shared file system technology**
- **Shared work spaces across CISL's HPC resources**
- **Multiple different spaces**
  - /glade/home/username 10 TB
  - /glade/users/username 80 TB
  - /glade/scratch/username 204 TB
  - /glade/proj\* 555 TB

# GLADE centralized file service

A centralized file service (Bluefire, Lynx, Mirage)

- **User home directory :**
  - /glade/home/username
  - 10 GB quota per user
  - Backup
- **Scratch (temporary computational space):**
  - /glade/scratch/username
  - 2 TB quota per user
  - NO Backup

# Data Analysis and Visualization

- **Data Analysis and Visualization**
  - High-end servers available 7 x 24 for interactive data analysis, data-post processing and visualization
- **Data Sharing**
  - Shared data access within the lab
  - Access to the NCAR Archival Systems and NCAR Data Sets
- **Remote Visualization**
  - Access to visual computing platforms from the convenience of your office using tcp/ip based remote image delivery service
- **Visualization Consulting**
  - Consult with CISL staff on your visualization problems





# Working on Mirage/Storm

- **Log on to 'mirage'**
  - `ssh -X -l username mirage[0-2].ucar.edu`
  - One-time password using CryptoCard or Yubikey
  - Use 'free' or 'top' to see if there is currently enough resources
- **Development environments**
  - Intel C, C++, F77, F90
  - GNU C, C++, Fortran, Tools
- **Software tools**
  - VAPOR, Paraview
  - NCL, NCO, NCARG, IDL
  - Matlab, R
  - ImageMagick

# Yellowstone Environment

- **Petascale computing resource**  
NCAR-Wyoming Supercomputing Center in Cheyenne, Wyoming
- **Production computing operations is expected to begin in summer/fall 2012.**
  - Computing resources: 30 times the workload throughput of NCAR's current Bluefire supercomputer
  - New centralized file system and data storage system (GLADE): 15 times the sustained I/O bandwidth and 12 times the capacity of CISL's current GLADE system
  - Combined data analysis and visualization (DAV) systems (Geyser and Caldera): 20 times increase in CISL's dedicated DAV resources
  - High Performance Storage System (HPSS): expanded more

# Future Reference

- Bluefire User Guide  
<http://www2.cisl.ucar.edu/docs/bluefire-user-guide>
- Lynx User Guide  
<http://www2.cisl.ucar.edu/docs/lynx-user-guide>
- Mirage/Storm User Guide  
<http://www2.cisl.ucar.edu/docs/mirage-storm>
- Glade File System  
<http://www2.cisl.ucar.edu/resources/glade>
- HPSS User Guide  
<http://www2.cisl.ucar.edu/docs/hpss>

# User Support

- **CISL Homepage:**
  - <http://www2.cisl.ucar.edu/>
- **CISL Consulting Services**
  - NCAR Mesa Lab Area 55, Floor 1B
- **CISL HELP**
  - Call (303)497-2400
  - Email to [cislhelp@ucar.edu](mailto:cislhelp@ucar.edu)
  - Submit an extraview ticket