



HPC at IU – Hands On

Le Mai Weakley

languyen@iu.edu

Research Technologies

Indiana University



**RESEARCH
TECHNOLOGIES**

INDIANA UNIVERSITY

University Information Technology Services



**PERVASIVE TECHNOLOGY
INSTITUTE**

INDIANA UNIVERSITY

How do I get access?

- I am sure you all already did this...
- Create an account:
 - <https://access.iu.edu/Accounts/Create>
 - Request accounts on Karst and SDA for the demos
- Login domain names:
 - bigred2.uits.iu.edu
 - karst.uits.iu.edu
 - IU network ID and pass phrase
 - Read the message of the day (MOTD)

<https://kb.iu.edu/d/achr>



RESEARCH
TECHNOLOGIES

INDIANA UNIVERSITY
University Information Technology Services



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

Let's get our hands dirty

- We will do the following on Karst today:
 - Copy something from your desktop to Karst
 - Change your environment
 - Submit a job to the queue to run something
 - Get an interactive node and launch a GUI application
 - Access Karst via Karst Desktop
 - Explore Karst Desktop
- This covers the most common activities on HPC systems
- Break
- Cybergateway as a means to submit jobs



**RESEARCH
TECHNOLOGIES**

INDIANA UNIVERSITY
University Information Technology Services



**PERVASIVE TECHNOLOGY
INSTITUTE**

INDIANA UNIVERSITY

SSH and SCP


- Sounds great, where do I login? Not in a browser ☹
- SSH is used to login to a remote machine
- SCP is used to copy files to and from a remote machine
- Example: (in a terminal, different on Windows)

```
ssh username@karst.uits.iu.edu
```

```
scp local-file username@karst.uits.iu.edu:/path/to/remote/dir
```

<https://kb.iu.edu/d/aelc>

<https://kb.iu.edu/d/agye>



SSH
The Secure Shell



RESEARCH
TECHNOLOGIES

INDIANA UNIVERSITY
University Information Technology Services



PERVASIVE
INSTITUTIONS

INDIANA UNIVERSITY

You need clients, at least on Windows

- For SSH: Putty on Windows (<http://www.putty.org/>)
- For SCP: Winscp for windows (<http://winscp.net/eng/index.php>)
 - Filezilla is cross platform (<https://filezilla-project.org/>)
- On Mac and Linux: you can use the terminal for SSH
 - Use the terminal for SCP as well
 - Or use Filezilla for GUI interface



**RESEARCH
TECHNOLOGIES**

INDIANA UNIVERSITY
University Information Technology Services



**PERVASIVE TECHNOLOGY
INSTITUTE**

INDIANA UNIVERSITY

Environment – in your login session

- Home directory:
 - /N/u/<username>/Karst
- Scratch space for temporary files:
 - /N/dc2/scratch/<username>



<https://kb.iu.edu/d/avmj>



RESEARCH
TECHNOLOGIES

INDIANA UNIVERSITY
University Information Technology Services



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY



Task 1: Login and transfer files to Karst

- ssh username@karst.uits.iu.edu
- A. Copy something from your desktop/laptop to Karst
 - Use scp in the terminal or use an scp client
- B. Download something from the web onto Karst:
 - wget <http://mirrors.concertpass.com/gcc/releases/gcc-4.9.1/gcc-4.9.1.tar.gz>
 - wget lets you download from the web



RESEARCH
TECHNOLOGIES

INDIANA UNIVERSITY
University Information Technology Services



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

How to find and use software

- modulecommand for adding software to your environment
 - module avail
 - module list
 - module load <key>
 - module unload <key>
- Permanently configure your environment via the .modules file
- Modules arranged in groups, development, applications, life sciences, etc.



<https://kb.iu.edu/d/bcwy>



**RESEARCH
TECHNOLOGIES**

INDIANA UNIVERSITY
University Information Technology Services



**PERVASIVE TECHNOLOGY
INSTITUTE**

INDIANA UNIVERSITY



What does “module load” do?

- Makes the software ready/available for use
 - Makes the software findable
- Does not start or run the application
- There isn't anything like this on Windows

<https://kb.iu.edu/d/bcwy>



**RESEARCH
TECHNOLOGIES**

INDIANA UNIVERSITY
University Information Technology Services



**PERVASIVE TECHNOLOGY
INSTITUTE**

INDIANA UNIVERSITY



Task 2: Play with modules

- Check “module list”
- Check “echo \$LD_LIBRARY_PATH”
- Add intel to your programming environment
 - module load intel/17.0.1
 - Check “module list”
 - Check “echo \$LD_LIBRARY_PATH”
- Remove intel from your environment
 - module unload intel
 - Check “module list”
 - Check “echo \$LD_LIBRARY_PATH”



**RESEARCH
TECHNOLOGIES**

INDIANA UNIVERSITY
University Information Technology Services



**PERVASIVE TECHNOLOGY
INSTITUTE**

INDIANA UNIVERSITY

Contd...

- Terminal Output -- <http://go.iu.edu/IMI>
- Load the gnu openmpi module
 - module load openmpi/gnu/1.8.4
 - Check “module list”
 - Loads the libraries needed to run parallel applications
- Make this change permanent
 - Edit .modules in your \$HOME
 - Common editors: vi, nano, emacs
 - Output of “cat .modules” should look like this:
 - “module load openmpi/gnu/1.8.4”
 - “cat” prints contents of a file



**RESEARCH
TECHNOLOGIES**

INDIANA UNIVERSITY
University Information Technology Services



**PERVASIVE TECHNOLOGY
INSTITUTE**

INDIANA UNIVERSITY

Sample PBS script

- At the most basic level, on Karst:

```
#!/bin/bash
#PBS -l nodes=2:ppn=16,walltime=30:00
cd /path/to/working/dir
mpirun -np 32 -machinefile $PBS_NODEFILE ~/bin/binaryname
```

- mpirun – the command used to run applications in parallel
- np – number of processors; machinefile – list of processors
- Put these lines in a text file and run “qsub script.sh”

<https://kb.iu.edu/d/avmy>



**RESEARCH
TECHNOLOGIES**

INDIANA UNIVERSITY
University Information Technology Services



**PERVASIVE TECHNOLOGY
INSTITUTE**

INDIANA UNIVERSITY

Queue commands

- Some other useful commands to see what's going on with the queues:
 - `qstat -Q` #shows available queues
 - `qstat -u username` #shows the status of your queued jobs
 - `showq -i` #shows what's going to run next
- Use “`#PBS -q`” to specify a particular queue in the PBS script
 - Need to do this on Big Red II, to specify the cpu or gpu queue (default is cpu)
 - Not always required on Karst, the default queue routes the jobs
 - But needed if you want to use the debug or interactive queues

<https://kb.iu.edu/d/avmy>



**RESEARCH
TECHNOLOGIES**

INDIANA UNIVERSITY
University Information Technology Services



**PERVASIVE TECHNOLOGY
INSTITUTE**

INDIANA UNIVERSITY

On Karst (non-Cray machine)

- There is a login node (s)
- There are compute nodes
- You build, install and setup your job environment on the login node
- Run the job on the compute node
- Using mpirun if parallel or just the binary if it is serial
 - That is:
 - `mpirun -np 32 app_name`
 - Or just:
 - `./app_name`



**RESEARCH
TECHNOLOGIES**

INDIANA UNIVERSITY
University Information Technology Services

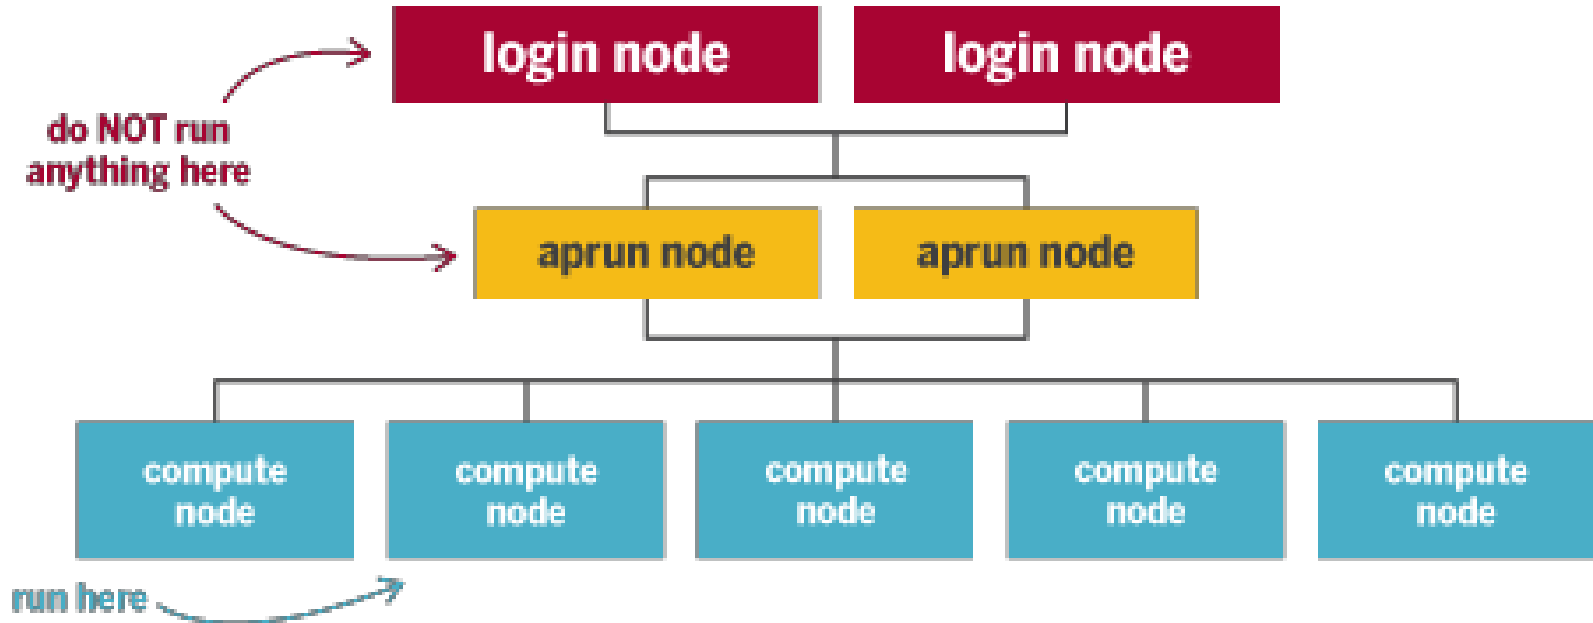


**PERVASIVE TECHNOLOGY
INSTITUTE**

INDIANA UNIVERSITY

On Big Red II and other Crays

- Login, aprun and compute nodes



<https://kb.iu.edu/d/bdkt>



RESEARCH
TECHNOLOGIES

INDIANA UNIVERSITY
University Information Technology Services



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY



CLM


CRAY LINUX ENVIRONMENT

ESM – *Extreme Scalability Mode*

- No compromise *scalability*
- Low-Noise Kernel for scalability
- Native Comm. & Optimized MPI
- Application-specific performance tuning and scaling

CCM – *Cluster Compatibility Mode*

- No compromise *compatibility*
- Fully standard x86/Linux
- Standardized Communication Layer
- Out-of-the-box ISV Installation
- ISV applications simply install and run



Launch everything with aprun or ccmrun, whether it is parallel or serial



RESEARCH
TECHNOLOGIES

INDIANA UNIVERSITY
University Information Technology Services



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

Contd.

- mpirun becomes aprun
- Flags and options are slightly different, but accomplish the same tasks
- Think of the aprun nodes as management nodes for the compute nodes



**RESEARCH
TECHNOLOGIES**

INDIANA UNIVERSITY
University Information Technology Services



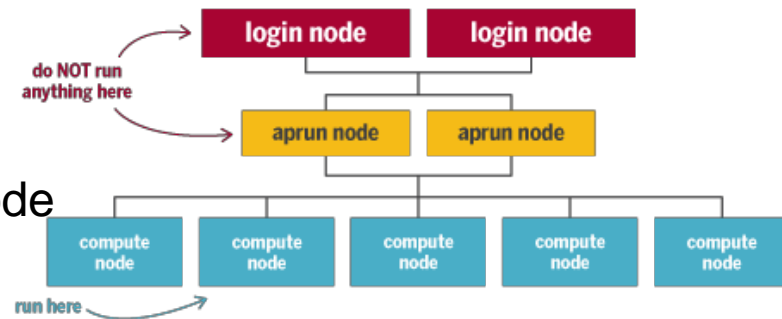
**PERVASIVE TECHNOLOGY
INSTITUTE**

INDIANA UNIVERSITY



Interactive jobs

- For graphical applications – you can launch GUIs, more later
- For testing
- `qsub -l`
 - Followed by options to specify walltime, queue, etc.
 - `qsub -I -l walltime=00:30:00 -q interactive`
- Straightforward on Karst
- Will land on an aprun node on Big Red II
 - `module load ccm`
 - Need to `ccmlogin` for the compute node



<https://kb.iu.edu/d/bdsi>



RESEARCH
TECHNOLOGIES

INDIANA UNIVERSITY
University Information Technology Services



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

Compilers

	Big Red II	Karst	Mason
Cray	module load PrgEnv-cray	---	---
c	cc	not available	not available
c++	CC	not available	not available
fortran	Ftn	not available	not available
GNU	module load PrgEnv-gnu	module load gcc	module load gcc
c	cc	gcc	gcc
c++	CC	g++	g++
fortran	ftn	gfortran	gfortran
Intel	module load PrgEnv-intel	module load intel	module load intel
c	cc	icc	icc
c++	CC	icpc	icpc
fortran	ftn	ifort	ifort
PGI	module load PrgEnv-pgi	module load pgi	module load pgi
c	cc	pgcc	pgcc
c++	CC	pgCC	pgCC
fortran	ftn	pgf77 or pgfortran (Fortran 90/95)	pgf77 or pgfortran (Fortran 90/95)

Java and MPI wrapper compilers (Open MPI and MPICH) are also available.

For more information, see "What compilers are available on the IU research systems?" kb.iu.edu/d/abby

<https://kb.iu.edu/d/abby>



**RESEARCH
TECHNOLOGIES**

INDIANA UNIVERSITY
University Information Technology Services



**PERVASIVE TECHNOLOGY
INSTITUTE**

INDIANA UNIVERSITY

Task 3: Submit a job to the queue

- The materials are available here: go.iu.edu/g2J
- Download an example program and job submission script:
 - `cd /N/dc2/scratch/username`
 - `wget http://rt.uits.iu.edu/ci/training/files/sc4everyone/fall-2015/job-submission.tar`
 - `tar xvf job-submission.tar`
- Change directory to job-submission
 - `cd job-submission`
- Compile the `mpi_hello.c` program
 - `mpicc mpi_hello.c`
 - Will create a binary “a.out”
 - Run “ls” to verify
 - “mpicc” is the MPI wrapper used to compile MPI programs



**RESEARCH
TECHNOLOGIES**

INDIANA UNIVERSITY
University Information Technology Services



**PERVASIVE TECHNOLOGY
INSTITUTE**

INDIANA UNIVERSITY

Cont....

- Edit the pbs.sh file to set the working directory and binary name
 - It will not work out of the box
 - You can check the current working directory by running “pwd”
 - Add “cd /N/dc2/scratch/username” to the script
 - Change “binary_name to a.out”
- Submit the job to the queue
 - “qsub pbs.sh”
- Change the PBS script to now submit the job to run across 2 nodes
 - #PBS -l nodes=2:ppn=16
 - ”qsub pbs.sh”
- Check output in the working directory or in your HOME directory if you did not specify this in your PBS script



RESEARCH
TECHNOLOGIES

INDIANA UNIVERSITY
University Information Technology Services



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

Launching GUI apps

- For X forwarding in SSH to work you must be running an X server program
- Most Linux systems have X server installed by default
- On Windows, can use Xming: <http://sourceforge.net/projects/xming/>
- On Mac, can use Xquartz: <http://xquartz.macosforge.org/landing/>
- Start these apps first before you launch SSH app
- Enable X11 forwarding by checking a box in your SSH app
- Use Karst Desktop – more in a bit

<https://kb.iu.edu/d/bdnt>



**RESEARCH
TECHNOLOGIES**

INDIANA UNIVERSITY
University Information Technology Services



**PERVASIVE TECHNOLOGY
INSTITUTE**

INDIANA UNIVERSITY

Interactive Job with X forwarding

1. On Karst, login with X forwarding enabled
 1. Launch xterm
 1. `ssh -X username@karst.uits.iu.edu`
2. On Karst, start an interactive job with X forwarding enabled
 1. Launch xterm from the compute node
 1. `qsub -l -X -q interactive -l nodes=1:ppn=1`

Find step by step instructions on the KB pages: <https://kb.iu.edu/d/bdnt>



**RESEARCH
TECHNOLOGIES**

INDIANA UNIVERSITY
University Information Technology Services



**PERVASIVE TECHNOLOGY
INSTITUTE**

INDIANA UNIVERSITY



Task 4: Karst Desktop, first time set up

- Download the suitable client for your OS from <https://www.cendio.com/thinlinc/download>
- Server: desktop.karst.uits.iu.edu
- Username: your IU username
- Password: your IU CAS password
- You will need an account on Karst (<https://one.iu.edu/> > Create Additional accounts > Karst)
- **WARNING:**
 - **Be sure to go through the options on the client app**
 - **Disable full screen mode under the “Screen” tab!**
 - **Mac users may not be able to go back to their Mac desktop without logging out or disconnecting**



RESEARCH
TECHNOLOGIES

INDIANA UNIVERSITY
University Information Technology Services



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

Task 5: Karst Desktop, demo

- KB Document
 - <https://kb.iu.edu/d/bfwp>
- Setting up Thinlinc Client
 - Connection settings
 - Screen Size setting
 - Authentication Settings
 - Icons
 - Exporting Local Drives
 - Reconnect Policy
 - Logging out vs Disconnecting
 - Submitting jobs



RESEARCH
TECHNOLOGIES

INDIANA UNIVERSITY
University Information Technology Services



PERVASIVE TECHNOLOGY
INSTITUTE

INDIANA UNIVERSITY

Anything else?

- Send us an e-mail/open a ticket!
- Gateways demo coming right up...



**RESEARCH
TECHNOLOGIES**

INDIANA UNIVERSITY
University Information Technology Services



**PERVASIVE TECHNOLOGY
INSTITUTE**

INDIANA UNIVERSITY

