

Supercomputers and you.



A guide to using the compute resources available
at Indiana University
Kenneth Steimel

Overview

1. Systems available at IU
 - 1.1 Karst
 - 1.2 Carbonate
 - 1.3 Big Red II
2. Accessing these Machines
 - 2.1 Getting Accounts
 - 2.2 Logging in
 - 2.3 Accessing file system
3. Adding software
4. Queues

Karst

High throughput computing

- karst.uits.iu.edu
- Best for single threaded, compute intensive applications that run for days (or less)
- Large span of pre-installed software
- 228 compute nodes
 - 32 GB of ram each
 - 2 octa core processors
- 16 dedicated data nodes
 - 64 GB of ram
 - Same processors

Carbonate

High memory computing

- carbonate.uits.iu.edu
- Newest system
- Best for high memory workloads
- 72 compute nodes
 - Each has 256 GB ram
 - Each has dual 12-core processors
- 8 large memory nodes
 - Each has 512 GB ram
 - Probably shouldn't queue up on these

Big Red II

- bigred2.uits.iu.ed
- Parallel jobs, gpu programming, distributed jobs
- Cray XE6/XK7 super computer
- Theoretical compute output of 1 petaflop
- 344 compute nodes (CPU only)
 - Two 16 core processors
 - 64 GB of ram
- 676 gpu accelerated nodes
 - One 16 core processor
 - 32 GB of ram
 - Tesla K20 gpu
- Very different architecture (MPP versus commodity cluster)
means that jobs are ran slightly differently

Getting set up with computing accounts

- Visit <https://one.iu.edu/task/iu/account-creation?searchTerms=create%20accounts>
- Sign in using CAS and then select the computing accounts you'd like to add. The prompt should be pretty straight forward.
- You'll receive an email in an hour or so saying your account has been created.

Logging in

- All of the computers have a terminal interface (shell)
- On windows, use PuTTY
(<https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html>)
- On Macs and Linux, use ssh¹²
- Use your iu network id as the username (e.g. ssh ksteimel@karst.uits.iu.edu)
- Karst has Karst desktop, this works great for people who are unfamiliar with the terminal
 - Visit <https://kb.iu.edu/d/bfwp> for more information.

Accessing file system

- You can connect to the shared storage between the computers using sftp
- For Windows, download winSCP, put in username, password and the hostname, the rest of the configs can be left alone
- For Mac and Linux, you can usually mount sftp servers as networked folders using your file manager
- Alternatively, you can scp files up there. Scp is solid but doesn't provide much advanced functionality

Adding software

- You can view the software that is available to add at the following locations:

<https://cybergateway.uits.iu.edu/iugateway/modulesInfo?machine=big>

<https://cybergateway.uits.iu.edu/iugateway/modulesInfo?machine=kar>

<https://cybergateway.uits.iu.edu/iugateway/modulesInfo?machine=carl>

(doesn't seem to be up to date)

- Alternatively, you can type 'module avail' to see what modules are available on the system you're running on.
- All of the computers include python, anaconda, nltk, nltk data etc.
- To add software, run 'module add <module name>'
- You may get an error saying Tcl command execution failed:
conflict <some module name>'
- You need to unload the module name it said and load the new

Software

- Removing software can be done by running 'module unload <module name>'
- Alternatively, you can swap two modules by running 'module swap <module to remove> < module to add>'
- There is a file in your home directory on each server that is called .modules
- module commands you put in this file will run every time you login
 - This is ideal for big red2 where you drop down into different nodes based upon the scheduler

Queues

- To view the queues that are available and see how long the line is, run 'qstat -q'
- Jobs that run on the login nodes are limited to 20 minutes of cpu time. After that you'll get booted
- Submitting to a queue may take a little bit but it will allow you to run longer jobs

Submitting to a Queue

- `qsub -l nodes=1:ppn=4,vmem=10gb,walltime=24:00:00 death_star.script`
 - One node, four processes, 10 gigs of ram, reserved for the next 24 hours running `death_star.script`
- `qsub -l -q interactive -l nodes=1:ppn=4,vmem=10gb,walltime=4:00:00`
 - One node, four processes, four hour reservation running interactively
- Specify the queue to run in with `-q <queue name>`
- Copy environment variables to run node using `-V` (you want to do this!)

A note about Big Red II

- On Karst and Carbonate, just don't run on login nodes. On Big Red II, also don't run on aprun nodes
- These are service nodes that drop running jobs onto a compute node.
- If you're running interactively, you want to preface your commands with aprun (and then usually some parameters like the number of processes)
- If you're not, you want to prepare a <https://kb.iu.edu/d/bdkt>. These scripts begin with a shebang (!) and have the aprun and torque directives necessary to run your job.
- You have to include module load statements in your TORQUE script or have them in your .modules file so that the libraries are accessible