K-State Beocat Compute Cluster

2 Head Nodes

ICR-Clymene

ICR-Helios



Software and Hardware overview

- Rocky Linux with the Slurm batch scheduler
- Over 10,000 Intel/AMD cores on 310 compute nodes
- ♦ 53 TeraBytes of RAM memory
- Low-latency 30-100 Gbps InfiniBand/RoCE network
- 1/10/40 Gbps Ethernet to a 1 PetaByte file server
- 290 TB fast scratch space /fastscratch
- 170 32-bit NVIDIA GPUs and 4 64-bit NVIDIA P100s

310 Compute Nodes accessed through Slurm 4 Interactive nodes accessed only through OnDemand

31 Warlocks	39 Wizards	120 Moles	62 Dwarves	54 Heroes
32-128 core Epyc 128 GB - 1.5 TB 40 Gbps RoCE 40 Gbps Ethernet 16 GPUs	32-64 core Skylake 96 GB - 1.5 TB 100 Gbps OmniPath 10 Gbps Ethernet 103 32-bit GPUs 4 64-bit GPUs	20-core Broadwell 32 GB 32 Gbps QDR 1 Gbps Ethernet	32-core Broadwell 128-512 GB 56 Gbps InfiniBand 40 Gbps Ethernet 26 GPUs	24-core Broadwell 128-512 GB 40 Gbps RoCE 40 Gbps Ethernet

Beocat Group

HPC Director Prof Dan Andresen



Application Scientist Dr Dave Turner



Lead System Admin Nathan Wells

System Admin

CS System Admin Seth Galitzer





Accessing Beocat

Request a Beocat Account

- + https://beocat.ksu.edu
- + Click on the Get an Account button
- Log in with your KSU credentials
- You may need to use your DUO access
- Submit the form referencing your instructor as your advisor
- Seth will need to approve the account
- Nathan will create the account and email you

Macintosh/Linux users

- Open a terminal window
- ssh into Beocat ssh daveturner@beocat.ksu.edu
- DUO authentication is required

Windows users

- Download and install MobaXterm or PuTTY
- connect to Beocat
- DUO authentication is required

BeoShock Compute Cluster

2 Head Nodes



Headnode02



nodes

36 cores

186 GB RAM

Software and Hardware overview

- Rocky Linux with the Slurm batch scheduler
- 768 Intel cores on 21 compute nodes
- Over 7.2 TeraBytes of RAM memory
- Low-latency 10 Gbps Ethernet network
- 150 TeraByte file server
- 8 High-End NVIDIA GPUs for accelerating scientific apps

21 Compute Nodes accessed through Slurm. Also Available interactively through OnDemand graphical interface.



BeoShock Administrators

Director of HPC Terrance Figy



Data Center Directory Ryan Doll



ITS Director of Project Management Anita Barrett



KSU System Admin Nathan Wells



KSU Application Scientist Dr Dave Turner



KSU CS System Admin Seth Galitzer





Request a BeoShock Account

- You must have a WSU ID and VPN access
- https://www.wichita.edu/services/hpc/hpc-guides.php
- + Expand Obtain an Account and click on New User Request
- + Fill in the form
- + Allow 2-3 business days for your account to be created.

Macintosh/Linux users

- + If off campus log into the WSU VPN on your computer
- Open a terminal window (xterminal or xterm2 on Mac)
- + Use your WSU_ID to ssh into BeoShock ssh WSU_ID@hpc-login.wsu.edu
- Or use the graphical interface https://ondemand.hpc.wichita.edu/

Windows users

- Download and install MobaXterm or PuTTY
- + If off campus log into the WSU VPN on your computer
- Use your WSU_ID to ssh into BeoShock ssh WSU_ID@hpc-login.wsu.edu
- Or use the graphical interface https://ondemand.hpc.wichita.edu/

Graphical Interface using OnDemand

Request an Interactive Session

- + Beocat https://ondemand.beocat.ksu.edu
- + BeoShock https://ondemand.hpc.wichita.edu
- + Log in with your university credentials
- You may need to use your DUO authentication
- + Under Interactive Apps choose Beocat/Beoshock Desktop
- Specify number of hours, cores, and memory for your interactive job
- + You may need to wait a few minutes for Slurm to start your job
- This provides a more graphical interface to the HPC system

Interactive Jupyter Notebook Session

- + Beocat https://ondemand.beocat.ksu.edu
- + BeoShock https://ondemand.hpc.wichita.edu
- + Log in with your university credentials
- You may need to use your DUO authentication
- + Under Interactive Apps choose Jupyter
- + Specify number of hours, cores, and memory for your interactive job
- + You may need to wait a few minutes for Slurm to start your job
- + This will provide a typical Jupyter Notebook environment for Python

Interactive sessions are good for developing code but eventually it's best to run jobs from the command line interface.

Using kstat to see what's running

Slurm tools

squeue to see what's running **sacct** to see what has finished

kstat perl script

kstat --help kstat kstat -h warlock kstat -q kstat -c kstat -c kstat -u daveturner kstat -d 7 kstat -me

Yellow background is a warning Red background flashing is bad

Submitting a job with sbatch

Copy the sample sbatch script to your directory

Beocat cp /homes/daveturner/test/sb.hello . BeoShock cp /homes/c297w489/test/sb.hello .

Submit the sb.hello batch script

sbatch sb.hello

View with kstat --me

What node is it running on?

Look at the output file

slurm-#.out where # is the job ID number In the current working directory unless told otherwise

Kill the job off with scancel

scancel # where # is the job ID number

#!/bin/bash -I
#SBATCHjob-name=Hello
#SBATCHtime=0-1:00:00
#SBATCHnodes=1
#SBATCHntasks-per-node=1
#SBATCHmem=4G
host=`hostname`
echo "Hello from \$host"
sleep 300

Software Modules

- ✦ module avail
- module spider python/
- module avail I& grep -i python/
- ✦ module reset
- module load Python
- module load Python/3.7.4-GCCcore-8.3.0 loads a specific version
- Use the same toolchain for all packages
 - foss or iomkl and same version number
- ✦ If you don't find a software package you need
 - Install it yourself on your home directory
 - Talk to us to see if there is an EasyBuild script

clear modules before starting will load the default version

Software Installation

You are ultimately responsible for installing your own applications

- \bigstar Install your application on your home directory
- \bigstar We are always willing to help if you need it
- \bigstar We can provide advice on optimizing the code
- \bigstar Dave Turner is available when more serious optimization is needed

General instructions for installing a software package

- \bigstar Download the software package
- ★ Decompress everything
- ★ Read the documentations (**README** or **INSTALL** files, ./configure --help)
- ★ ./configure --prefix=/homes/daveturner/bio/abyss32
- ★ make
- \star make install

Many packages are more difficult to install

- More configure options may be needed
- Beocat has some optimized libraries installed that may be needed
- You may need to install other packages first
- You may need to manually edit the makefile
- Often things just don't work like they should

HPC Documentations

Beocathttps://beocat.ksu.edu/BeoShockhttps://www.wichita.edu/services/hpc/hpc-guides.phphttps://docs.hpc.wichita.edu/

- priority access through a project
- installed software modules and installing your own
- Linux basics
- ✦ Slurm basics
- Slurm advanced
 - File systems
 - Array jobs
- Getting help beocat@cs.ksu.edu
 - *kstat -d 7* to check for error messages
 - Give us lots of information
 - Job ID
 - directory and command you used
 - full description of the problem
 - Wichita include your wsuid
 - ssh into the node you are running on and use htop
- ✦ Zoom help sessions Wednesdays 1:30-2:30

Additional Information

Google search for 'linux tutorial'

Google search for 'vi tutorial'

Oklahoma University - Supercomputing in Plain English http://www.oscer.ou.edu/education.php

HPC User software carpentry module by Dr Dave Turner https://drdaveturner.github.io/HPC-User/index.html

message passing with MPI: https://computing.llnl.gov/tutorials/mpi/

Overview of **OpenMP**: https://computing.llnl.gov/tutorials/openMP/