

Research Computing Support at the University of Utah: An Overview of CHPC

Anita Orendt

Assistant Director

Research Consulting & Faculty Engagement

anita.orendt@utah.edu

CHPC's Mission

In addition to deploying and operating high performance computational resources and providing advanced user support and training, CHPC serves as an expert team to broadly **support** the increasingly **diverse research computing needs** on campus.

These needs include support for big data, big data movement, data analytics, security, virtual machines, Windows science application servers, protected environments for data mining and analysis of protected health information, and advanced networking.

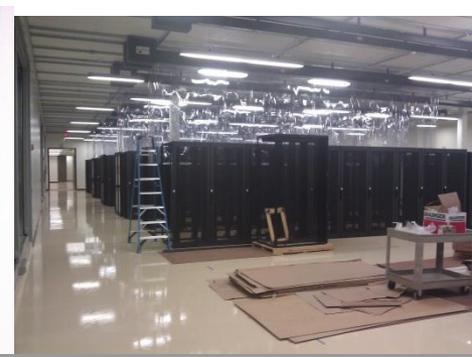
- CHPC can help if:
 - You need parallel processing
 - You need access to a single high-powered computer
 - You need the ability to run many individual jobs simultaneously
 - You have a large amount of data to store and/or process
 - You need an application you don't have on your computer
 - Your data is IRB-governed PHI
 - You have other computing needs your local resources cannot meet

Sources of Useful Information

- Getting Started Guide
 - <https://www.chpc.utah.edu/documentation/gettingstarted.php>
- CHPC policies
 - <https://www.chpc.utah.edu/documentation/policies/index.php>
- Cluster Usage Guides
 - <https://www.chpc.utah.edu/documentation/guides/index.php>
- Application Documentation
 - <https://www.chpc.utah.edu/documentation/software/index.php>
- Programming Guide
 - <https://www.chpc.utah.edu/documentation/ProgrammingGuide.php>
- How to Videos
 - <https://www.chpc.utah.edu/documentation/videos/index.php>

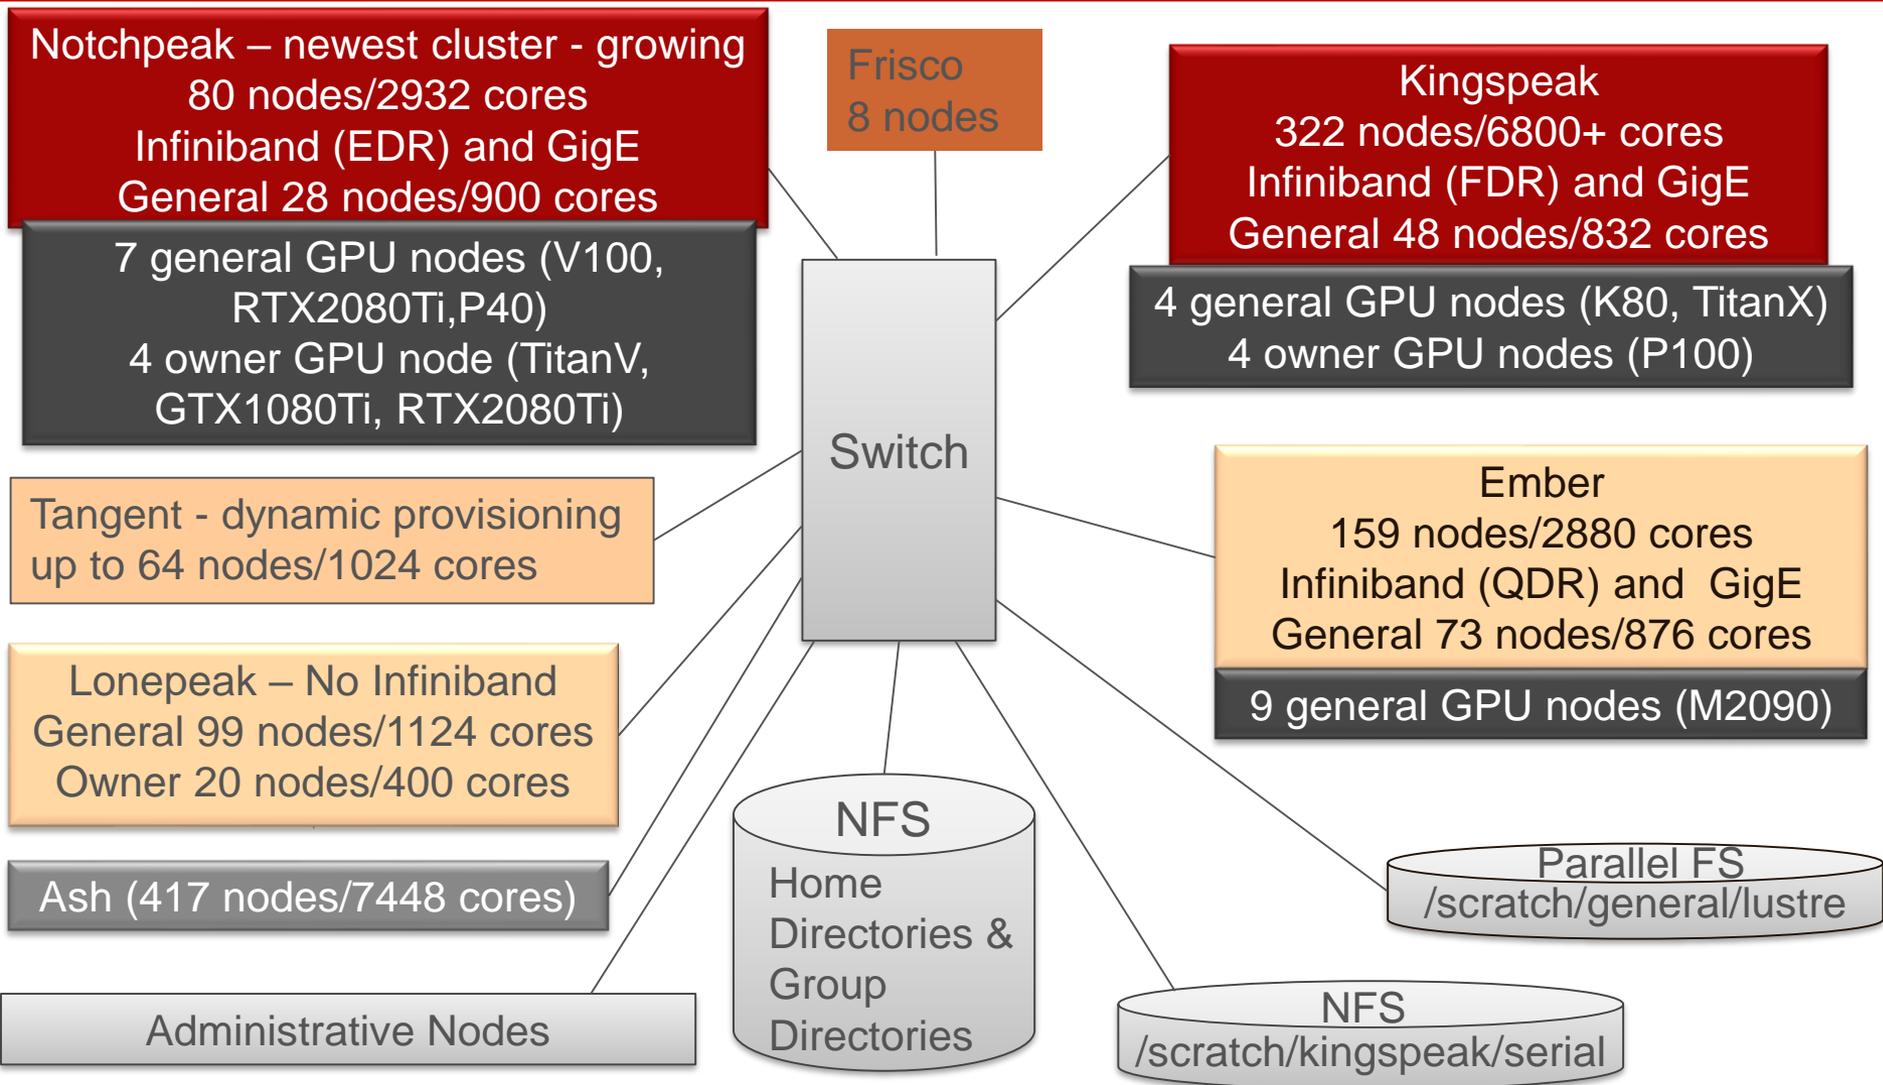
Downtown Data Center

- Came online Spring 2012
- CHPC completed move to DDC Spring 2013
- Shared with enterprise (academic/hospital) groups
- 92 racks and 1.2MW of power with upgrade path to add capacity for research computing
- Metro optical ring connecting campus, data center, & internet2
- 24/7/365 facility



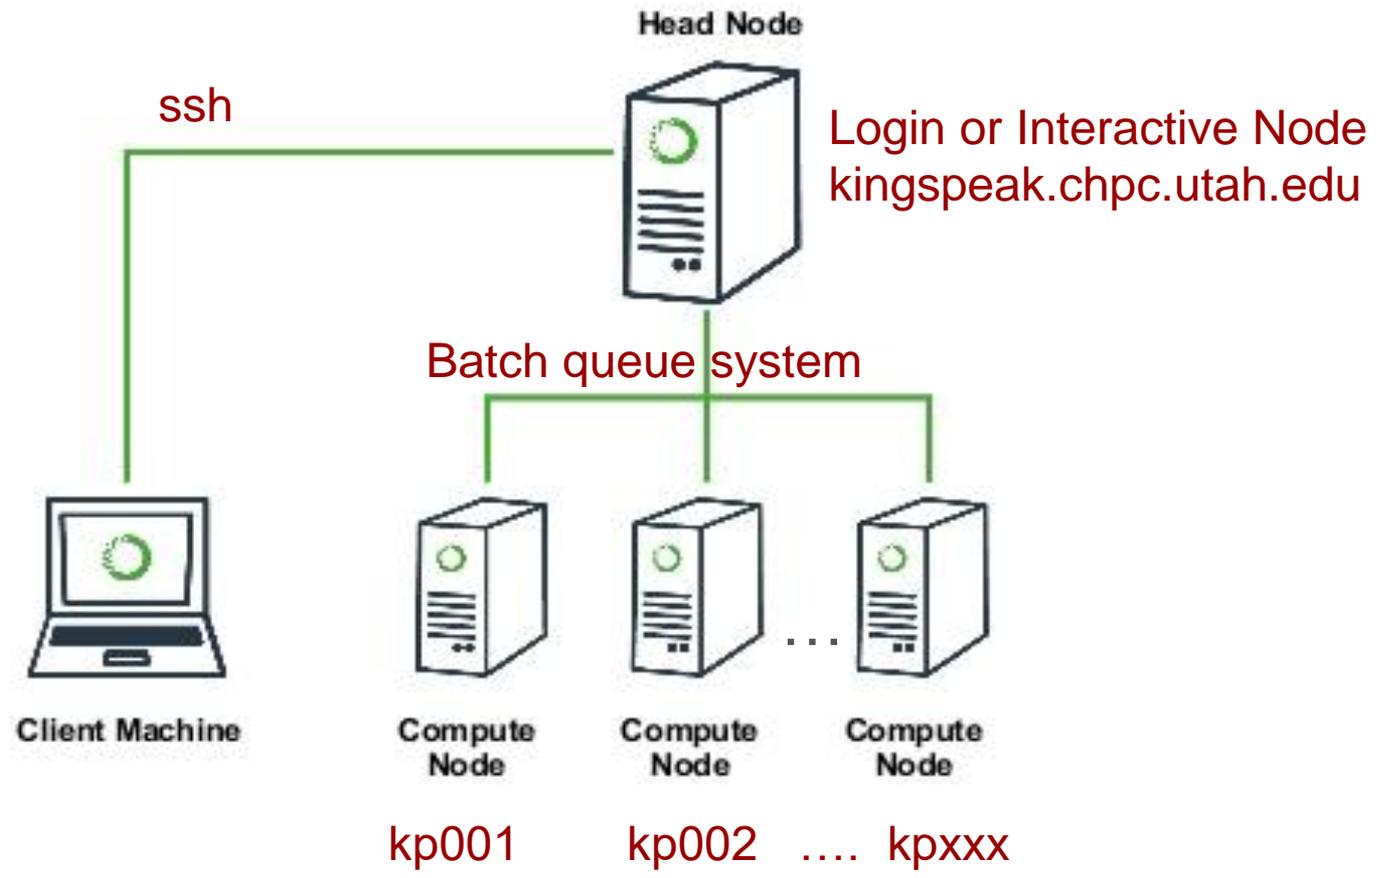
CHPC Resources & Services

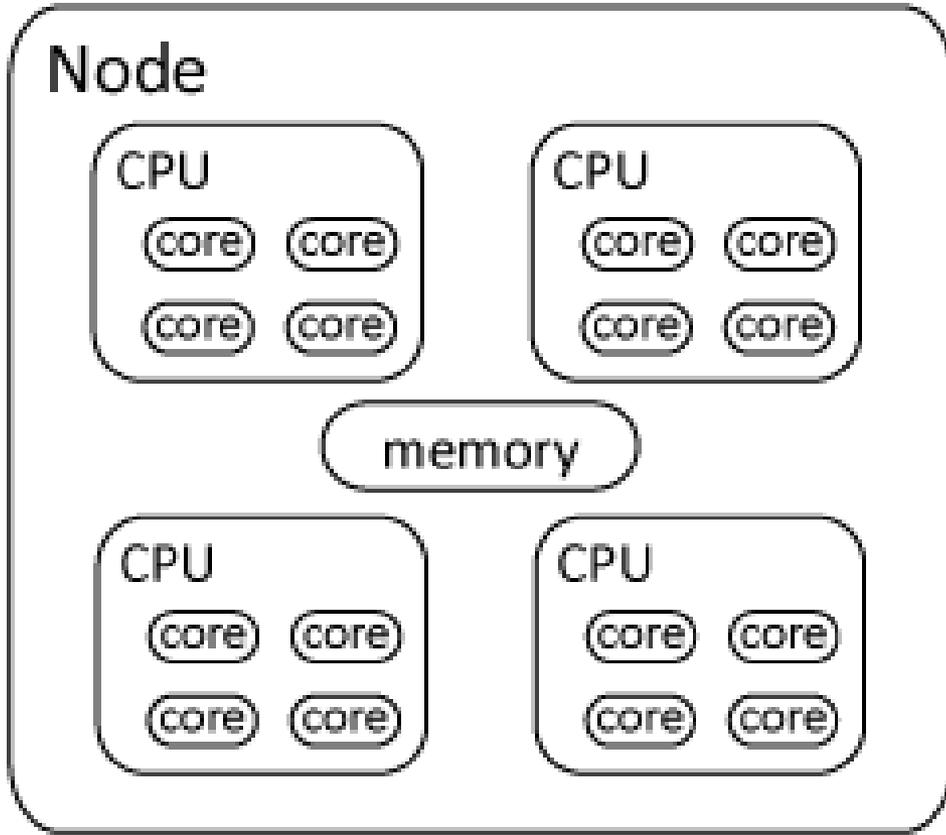
- **Computational Clusters** – Notchpeak, Kingspeak, Lonepeak, Ember, Ash, Tangent
- **Storage** – home, group, and scratch storage along with archive storage options
- **Windows Servers** – mainly statistics usage and windows only applications
- **Virtual Machines** – for needs not met with cluster and windows server
- **Protected Environment** –computational cluster Redwood, storage, VMs, and Windows Server
- **Networking Support** – support compute environment; work with researchers on data movement
- **User Support** – assistance with use of resources; installation of applications; training sessions





Cluster Architecture Diagram



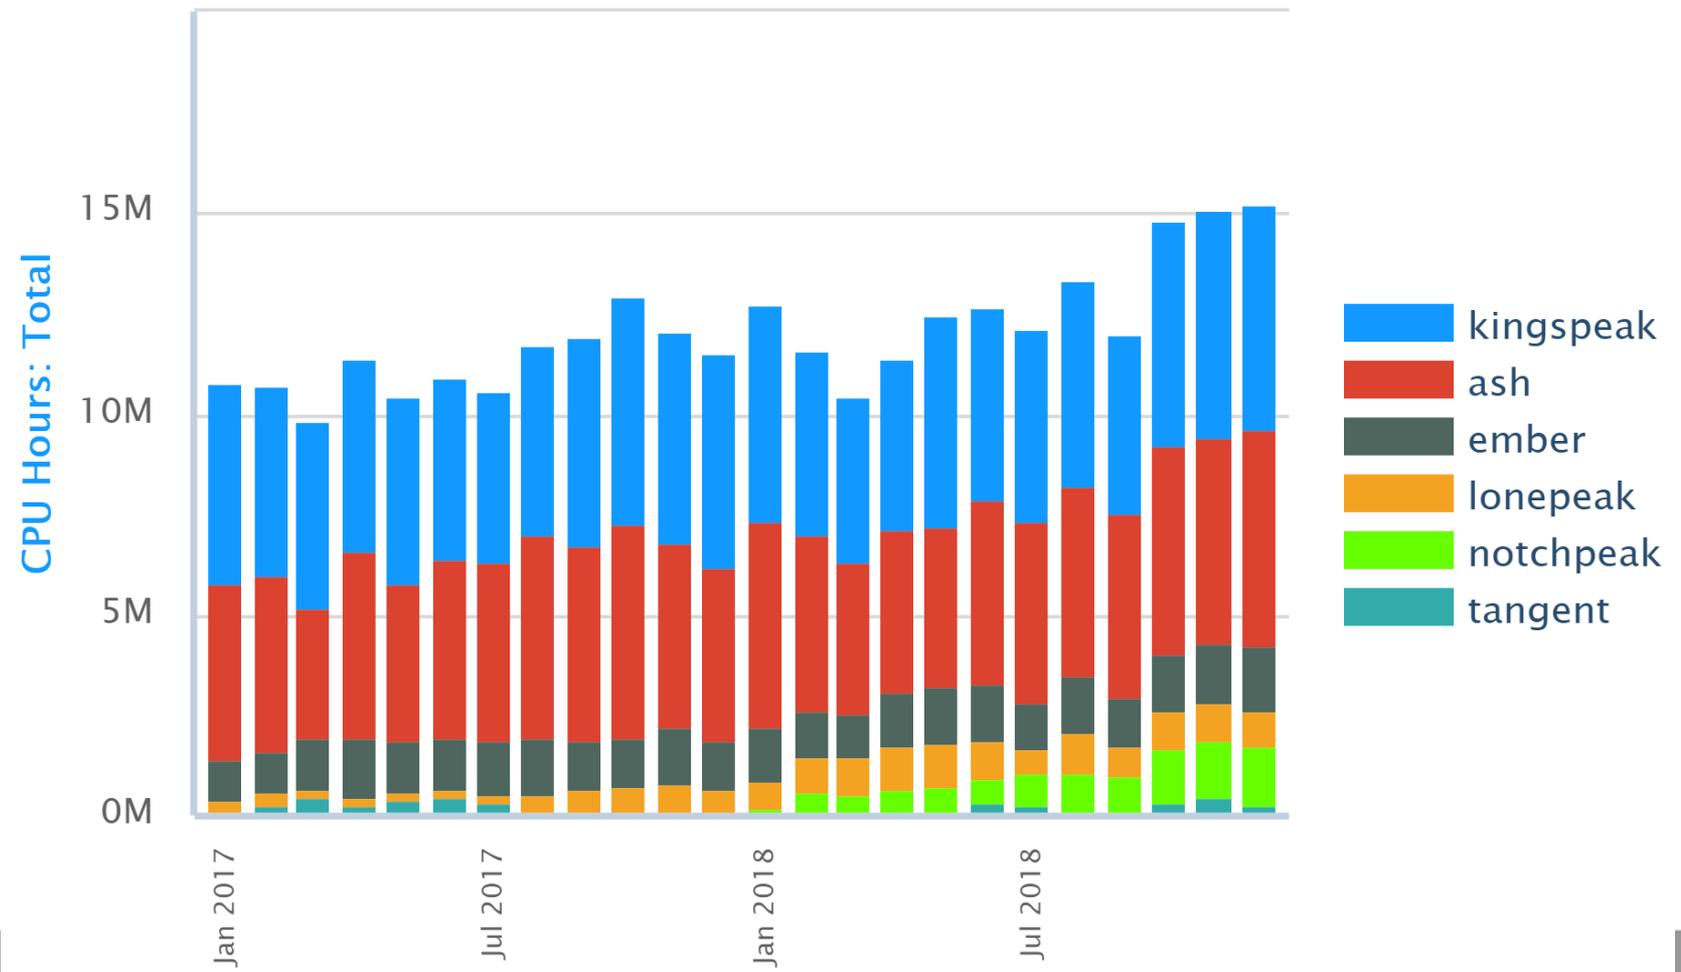


CHPC Clusters - Condominium Model

- Notchpeak, Kingspeak
 - General resources – Allocation process for cycles
 - Out of allocation – freecycle allowed (with preemption)
- Owner nodes – now being added to notchpeak
 - Users from all groups allowed on as owner-guest when not in use (preemption)
 - Found on notchpeak, kingspeak, ember, lonepeak
 - Latest nodes have 40 core Intel XeonSP (Cascade Lake) @ 2.1GHz cpu speed – Pricing to be announced
- Ember, Lonepeak
 - no allocation and no preemption
- Tangent
 - dynamically provisioned cluster resource; no allocation and no preemption
- Ash – Owner cluster
 - All users have guest access as smithp-guest (preemption)

Core Hour Usage

General Environment - Over 134M core hours provided in 2017; 154M in 2018



CHPC Provides Core Infrastructure

- Physical needs (racks, power, cooling)
- Core ethernet and IB fabric; IB to ethernet bridging
- Login/head/management nodes
- Scheduling, batch and allocation systems
- HPC scratch space
- Some application and licensing costs
- High speed data transfer nodes (DTNs)
- 'Science DMZ' network
- CHPC Staff

Recent and Upcoming Changes

- New NFS scratch file system – Coming soon
 - /scratch/general/nfs1
 - about 600 TB
- New Cascade Lake based nodes
- Additional notchpeak general nodes
 - Gpu just added
 - cpu coming soon
- New Windows server in general environment
- New backup options – moving away from tape
- Updated Arbiter monitoring interactive node usage

Storage Options

- *Home Directories* -- /uufs/chpc.utah.edu/common/home/<uNID>
 - Home directories 50 GB, not backed up
 - Groups can purchase 1 TB max home directory per group at \$1250/TB for hardware lifetime; this comes with backup (nightly incremental, weekly full, 2 week retention)
 - Compellent solution - two disk based copies mirrored for HA
- *Group Level File Systems*
 - Group space @\$150/TB
- *Scratch File Systems*
 - For use by all users; scrubbed of files older than 60 days
 - 700 TB Lustre Parallel file system (/scratch/general/lustre)
 - 175 TB NFS mounted file system (/scratch/kingspeak/serial)
- *Disk Based Archive Storage*
 - Archive currently at \$140/TB

Storage Reminder

****you should always have an additional copy or possibly multiple copies, on independent storage systems, for any crucial/critical data. While storage systems built with data resiliency mechanisms (such as RAID and erasure coding used at CHPC or other similar technologies) allow for multiple component failures, they do not offer any protection for large scale hardware failures, software failures leading to corruption, or for accidental deletion or overwriting of data. Please take the necessary steps to protect your data to the level you deem necessary.****

File Systems

- Access speed based on connectivity
- Local disk fastest – local to each node; varies in size
 - /scratch/local
- Network mounted scratch file systems
 - /scratch/general/lustre
 - /scratch/kingspeak/serial
 - home directories and group spaces (*don't use for large i/o!*)

**** NFS mounted spaces – including file systems for group spaces – are a shared resource! ****

Protected Environment

- Refreshed in 2017 with award of NIH Shared instrumentation grant
 - HPC cluster – redwood
 - VM farm – prismatic
 - storage – mammoth, both home and project space; elm for archive object storage
 - windows compute – narwhal
- See <https://www.chpc.utah.edu/resources/ProtectedEnvironment.php>
- Dedicated protected resources for handling of data/projects with protected information
- Currently HIPAA
- Also for projects with other types of sensitive data/restrictions
- Significant area of growth for CHPC
- Preferred location for human genomic – meets NIH dbGaP requirements

Getting a CHPC Account

- CHPC uses campus uNID and password
- PIs must have account and will need to approve accounts for any members of their research group (can delegate)
- Account Application Procedure – Online process
 - Complete CHPC account form at https://www.chpc.utah.edu/role/user/account_request.php
 - For collaborators outside of University of Utah must complete affiliate form with HR to get uNID <https://www.hr.utah.edu/forms/affiliate.php> and then use account application

Accessing Clusters

- Login or interactive nodes with each cluster
 - ssh -Y **cluster**.chpc.utah.edu where **cluster** is notchpeak, kingspeak, ember, tangent, ash-guest, lonepeak (redwood in PE)
- Interactive nodes only used for short compiles, editing and very short test runs
- ***No more than 15 minutes and no jobs of any length that make heavy use of cpu or memory!***
- Have script which watches running processes and notifies users when in violation of the acceptable usage policy

Accessing Login nodes

- Use FastX from Mac, Windows, or Linux desktops -- preferred
 - <https://www.chpc.utah.edu/documentation/software/fastx2.php>
- Alternatively:
 - From windows need ssh client
 - PuTTY <http://www.chiark.greenend.org.uk/~sgtatham/putty/>
 - Xshell http://www.netsarang.com/products/xsh_overview.html
 - For X forwarding applications also need
 - Xming <http://www.straightrunning.com/XmingNotes/>
 - Look for “mesa” version
 - From mac/linux – use terminal ssh (with -Y for X forwarding)
- Access to protected environment needs Duo two factor authentication (and VPN if off campus)

FastX – Tool for Remote X

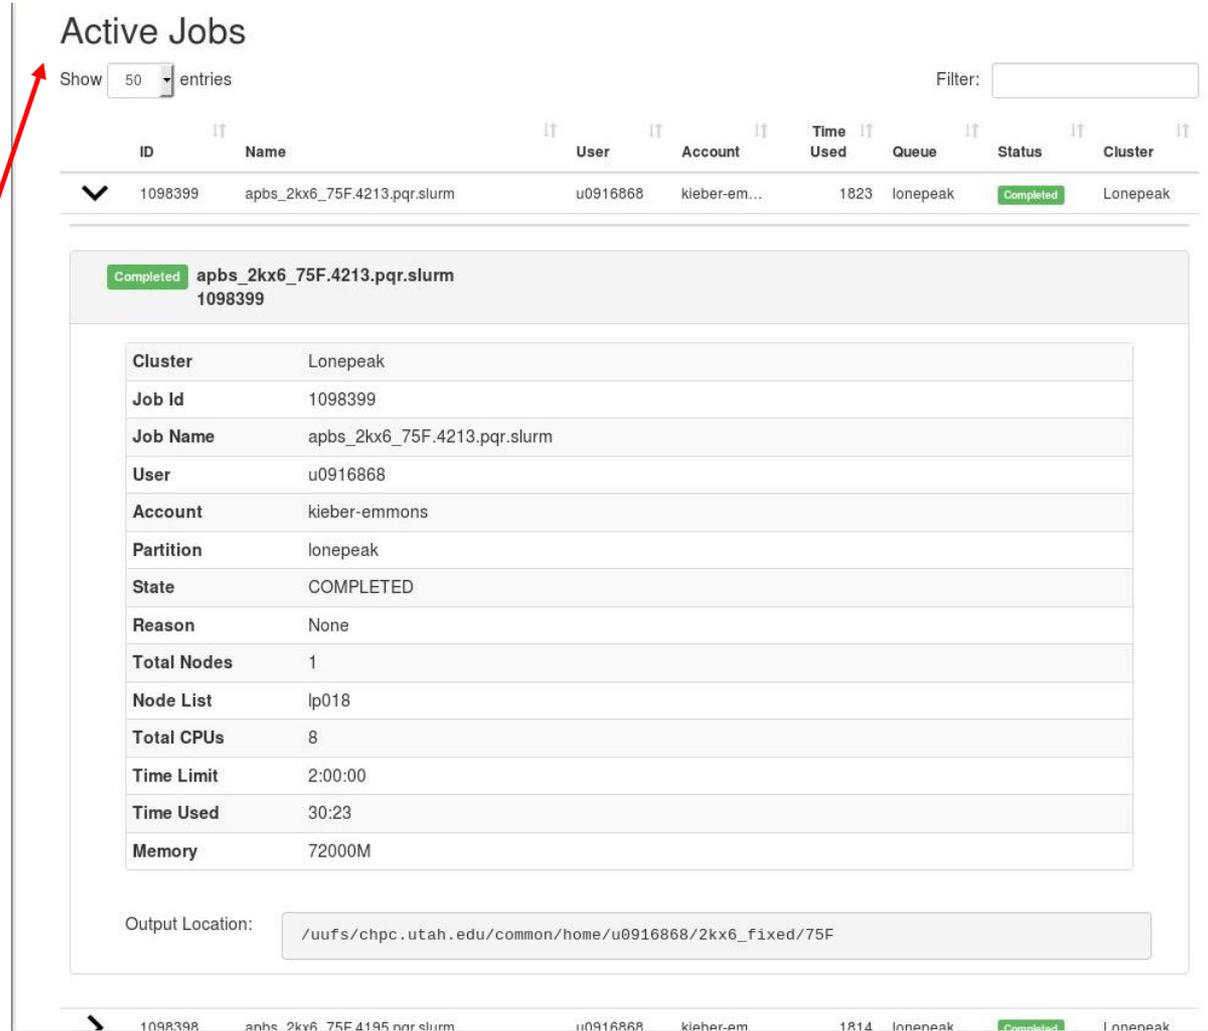
- <https://www.starnet.com/fastx>
- Used to interact with remote linux systems graphically in much more efficient and effective way than simple X forwarding
- Graphical sessions can be detached from without being closed, allowing users to reattach to the session from the same or other systems
- Server on all interactive nodes as well as the frisco nodes; some servers have graphics cards and support OpenGL
- Clients for windows, mac and linux; can be installed on both university and personal desktops.

FastX

- For FastX – see “To Use” section of documentation at <https://www.chpc.utah.edu/documentation/software/fastx2.php>
- Download client following directions on page
- Do install
- Start program
- Set host to kingspeak1.chpc.utah.edu OR kingspeak2.chpc.utah.edu OR other interactive node OR one of the frisco nodes (frisco1-frisco8.chpc.utah.edu)

Open OnDemand

- CHPC web portal
 - View/edit files
 - See queues and jobs
 - Create SLURM scripts from templates and submit jobs
 - Run interactive jobs and applications



Active Jobs

Show entries Filter:

ID	Name	User	Account	Time Used	Queue	Status	Cluster
1098399	apbs_2kx6_75F.4213.pqr.slurm	u0916868	kieber-em...	1823	lonepeak	Completed	Lonepeak

Completed apbs_2kx6_75F.4213.pqr.slurm
1098399

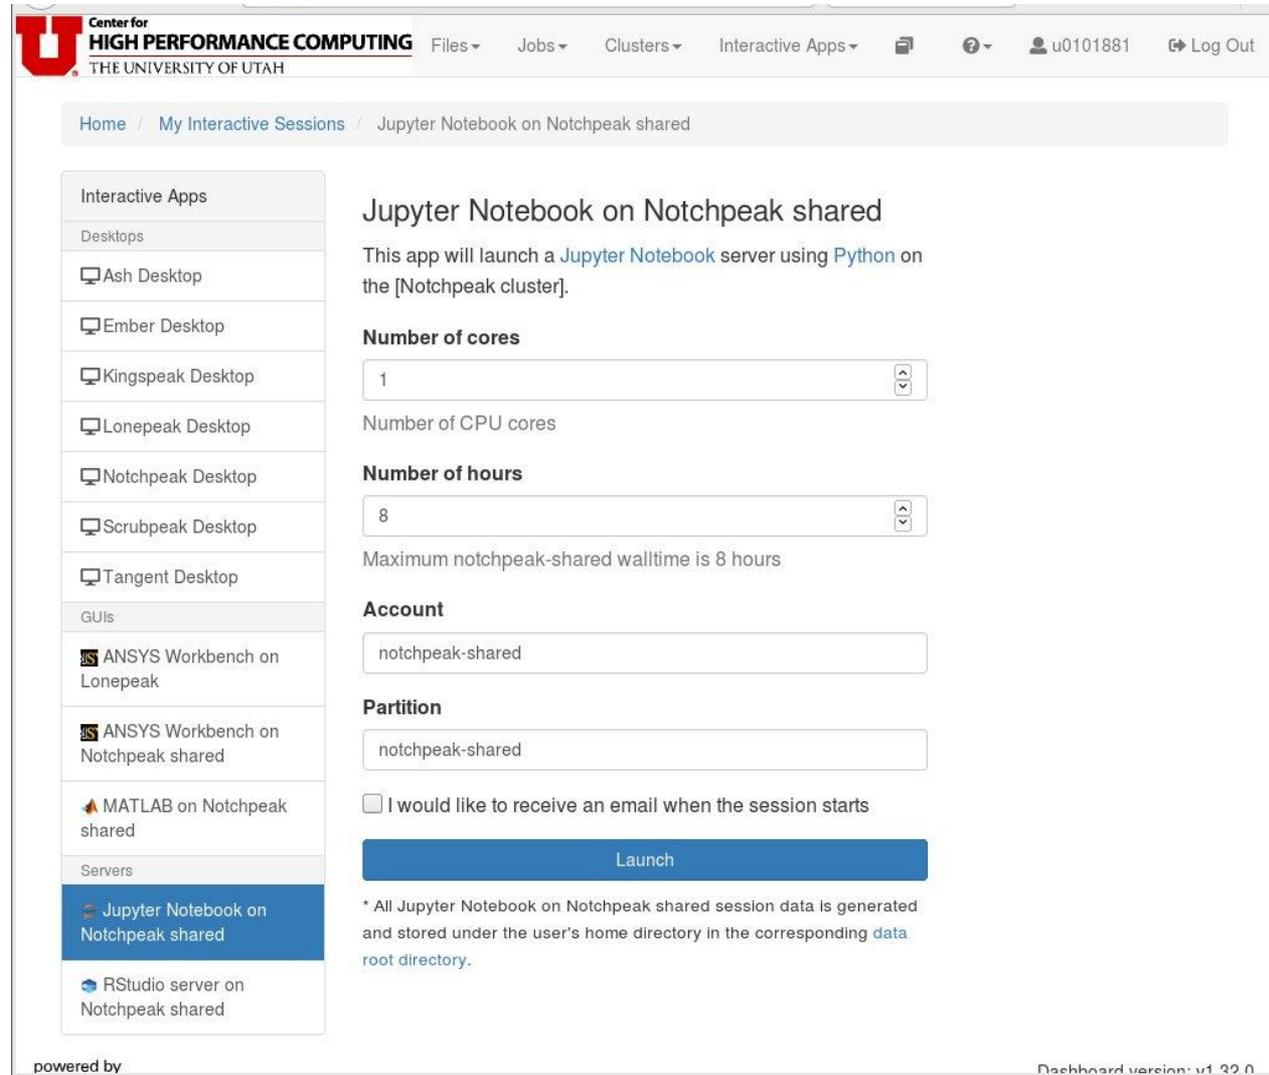
Cluster	Lonepeak
Job Id	1098399
Job Name	apbs_2kx6_75F.4213.pqr.slurm
User	u0916868
Account	kieber-emmons
Partition	lonepeak
State	COMPLETED
Reason	None
Total Nodes	1
Node List	lp018
Total CPUs	8
Time Limit	2:00:00
Time Used	30:23
Memory	72000M

Output Location:

1098398	apbs_2kx6_75F.4195.pqr.slurm	u0916868	kieber-em...	1814	lonepeak	Completed	Lonepeak
---------	------------------------------	----------	--------------	------	----------	-----------	----------

OOD interactive job

- Launch Jupyter notebook on notchpeak-shared partition
- First select the resources needed and hit “Launch”



The screenshot shows the CHPC interactive job submission interface. The header includes the CHPC logo, navigation menus (Files, Jobs, Clusters, Interactive Apps), and user information (u0101881, Log Out). The main content area is titled "Jupyter Notebook on Notchpeak shared" and contains the following fields:

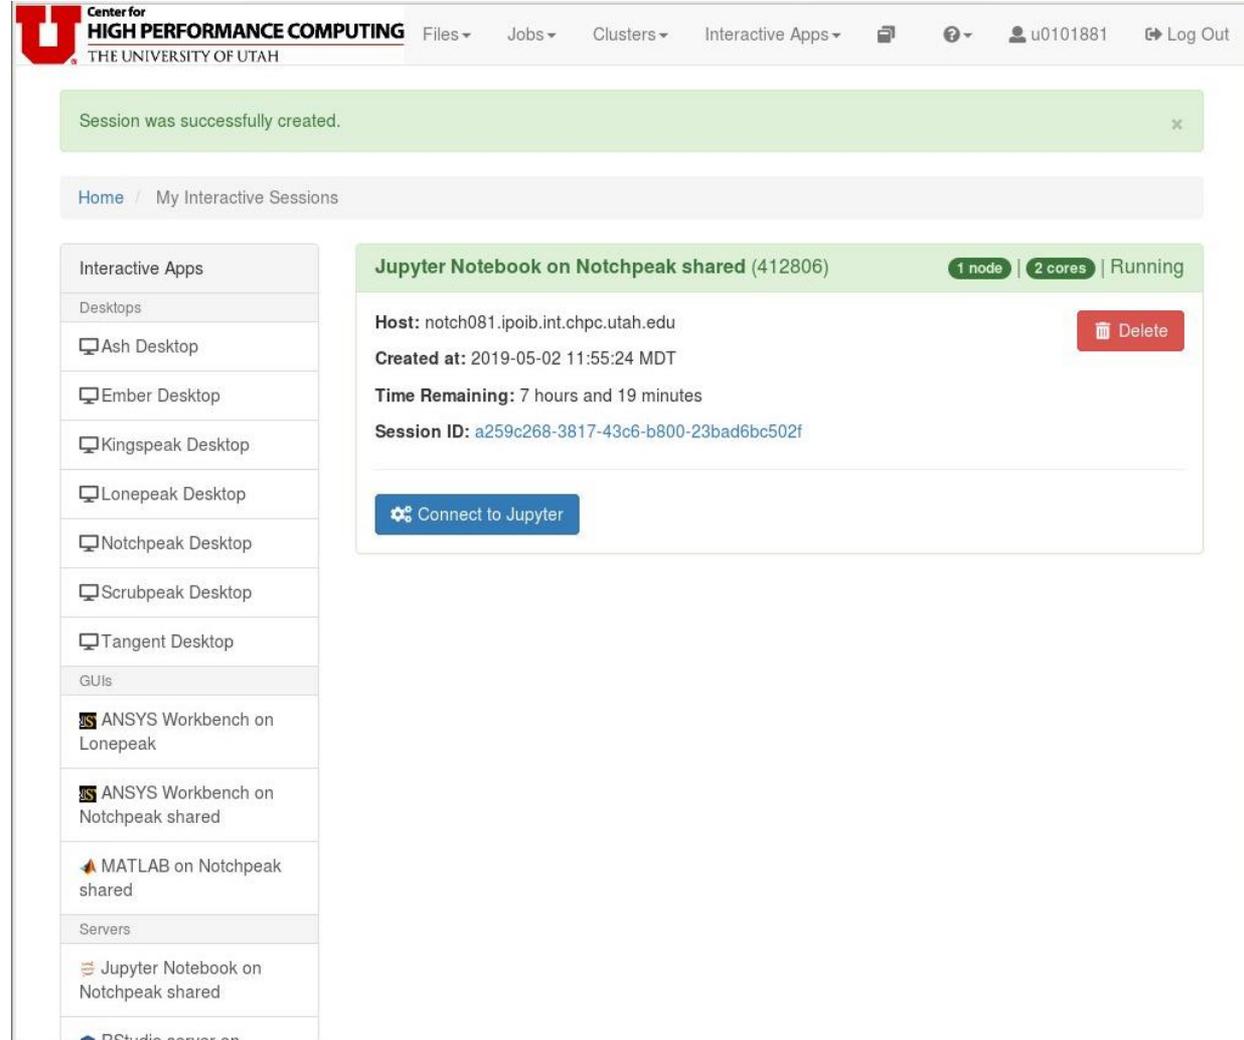
- Number of cores:** A dropdown menu set to 1.
- Number of hours:** A dropdown menu set to 8.
- Account:** A text input field containing "notchpeak-shared".
- Partition:** A text input field containing "notchpeak-shared".
- Launch:** A blue button labeled "Launch".

Below the "Launch" button, there is a note: "* All Jupyter Notebook on Notchpeak shared session data is generated and stored under the user's home directory in the corresponding data root directory."

The left sidebar lists various interactive apps and servers, with "Jupyter Notebook on Notchpeak shared" selected and highlighted in blue.

OOD interactive job

- Job gets submitted
- Once the resources are allocated and Jupyter started, “Connect to Jupyter” button appears
- Clicking on this button will start a new Jupyter notebook in a new browser tab



The screenshot displays the web interface of the Center for High Performance Computing. At the top, there is a navigation bar with the center's name and logo, and a user profile section showing 'u0101881' and a 'Log Out' button. A green notification banner at the top left states 'Session was successfully created.' Below this, the breadcrumb 'Home / My Interactive Sessions' is visible. On the left side, there is a sidebar menu with categories: 'Interactive Apps', 'Desktops', 'GUIs', and 'Servers'. Under 'Desktops', 'Notchpeak Desktop' is selected. Under 'GUIs', 'ANSYS Workbench on Notchpeak shared' is selected. The main content area shows a session titled 'Jupyter Notebook on Notchpeak shared (412806)' with a status of 'Running' (1 node, 2 cores). It includes details such as the host 'notch081.ipob.int.chpc.utah.edu', creation time '2019-05-02 11:55:24 MDT', time remaining '7 hours and 19 minutes', and session ID 'a259c268-3817-43c6-b800-23bad6bc502f'. A 'Delete' button is present next to the session title. A prominent blue button labeled 'Connect to Jupyter' is located at the bottom of the session details.

Login scripts

- CHPC provides login scripts (“dot” files) when creating account for both tcsh and bash shells
- These files set the environment so that applications are found, batch commands work – ***Do not remove***
- Choose shell at account creation – can change at www.chpc.utah.edu (sign in, select edit profile)
- Four files: .bashrc, .tcshrc, .custom.sh, .custom.csh
 - **The first two should not be edited**
 - **The second two is where to add custom module loads**
- Will automatically execute an .aliases file if it exists

CHPC Uses Modules for Setting Environment

- CHPC provides login scripts (“dot” files) when creating account for both tcsh and bash shells
- These files set the environment so that applications are found, batch commands work – ***Do not remove or edit!***
- <https://www.chpc.utah.edu/documentation/software/modules.php> for information
- Presentation on Modules – Thursday, May 16

Batch System Information

- Used to access compute nodes which must be used for any extensive use
- Use SLURM – Simple Linux Utility for Resource Management
- <https://www.chpc.utah.edu/documentation/software/slurm.php> for information
- Presentation on Slurm – Thursday, May 30

Software on Clusters

- Have a variety of compilers, mpi packages, math libraries and applications installed
- Some licensing restrictions may apply
- If you need a package we do not currently have installed – ask us!
- Currently we place installations at:
 - /uufs/chpc.utah.edu/sys/installdir
- Have a searchable application database
 - <https://www.chpc.utah.edu/software/chpc/>

Allocation

- General Allocation Process Information
 - <https://www.chpc.utah.edu/documentation/policies/1.4AllocationPolicies.php>
- Regular allocation form
 - https://www.chpc.utah.edu/apps/profile/allocation_form.php
 - Requests due Sept 1, Dec 1, Mar 1, and Jun 1
 - Allocation in core hours
- Quick allocation
 - https://www.chpc.utah.edu/apps/profile/allocation_quick_form.php
- Check usage -- <https://www.chpc.utah.edu/usage/cluster/current-project-general.php>
- Simplified quick allocation requests & general allocation requests for up to 20,000 core-hours per quarter

Windows Statistics Server

- Kachina – each 48 core, 512 TB memory
- Presently has the following software installed
 - SAS 9.4 with text miner
 - SPSS
 - R
 - STATA
 - Mathematica
 - Matlab
- *If you need other software, please contact us to discuss*

Virtual Machine Farm

- For needs and applications that do not fit in compute cluster or Windows server
- Multiple VM servers with failover
 - Working on migration to new hardware
- VM storage
- Have community mysql/mssql VMs, git repositories, web servers, etc
- New – user VMs (not use of community ones) will have a cost, both for the VM and for any customization needed.

Blocks	RAM (GB)	Cores	Storage (GB)	Price – General	Price - PE
1	4	2	50	\$475	\$630
2	8	2	100	\$705	\$1025
4	16	4	200	\$1175	\$1810
8	32	8	400	\$2115	\$3380
16	64	8	800	\$3995	\$6525

Summer 2019 Presentation Series

In INSCC Auditorium -- can join remotely via zoom

DATE	TIME	PRESENTATION TITLE	PRESENTER
Tuesday, May 14	1-2pm	<u>Introduction to HPC & CHPC</u> *	Anita Orendt
Thursday, May 16	1-2pm	<u>Module Basics</u> *	Anita Orendt
Thursday, May 30	1-2pm	<u>Slurm Basics</u> *	Anita Orendt
June 3-6	9am-3pm	<u>XSEDE Summer Boot Camp</u> ***	Brett Milash
Tuesday, June 11	1-3pm	<u>Hands-on Introduction to Linux, part 1</u> **	Anita Orendt, Brett Milash
Thursday, June 13	1-3pm	<u>Hands-on Introduction to Linux, part 2</u> **	Brett Milash, Anita Orendt
Tuesday, June 18	1-3pm	<u>Hands-on Introduction to Linux, part 3</u> **	Brett Milash, Robben Migacz
Thursday, June 20	1-3pm	<u>Introduction to Python, Part 1</u> **	Brett Milash , Robben Mlgacz
Tuesday, June 25	1-3pm	<u>Introduction to Python, Part 2</u> **	Brett Milash, Wim Cardoen
Thursday, June 27	1-3pm	<u>Introduction to Numpy Part 1 (Python, Part 3)</u> **	Wim Cardoen, Brett Milash
Tuesday July 2	1-3pm	<u>Introduction to Numpy Part 2 (Python, Part 4)</u> **	Wim Cardoen, Brett Milash
July 15-19	Start 10am End between 3-5pm	Build a Cluster Workshop Will be at Datacenter; watch for announcement for registration in June	Brian Haymore
August 6-8	9am-3pm	<u>XSEDE HPC Monthly Workshop: Big Data</u> ***	Martin Cuma

<https://www.chpc.utah.edu/presentations/Summer2019CHPCPresentationSchedule.php>

If you would like training for yourself or your group, CHPC staff would be happy to accommodate your request. Please contact anita.orendt@utah.edu or helpdesk@chpc.utah.edu

Getting Help

- CHPC website
 - www.chpc.utah.edu
 - Getting started guide, cluster usage guides, software manual pages, CHPC policies
- Ticketing System
 - Email: helpdesk@chpc.utah.edu
- Help Desk: 405 INSCC, 581-6440 (9-5 M-F)
- Mailing Lists:
chpc-hpc-users@lists.utah.edu used to send messages to users