

Research Computing and Data 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>
- Compiler Guide
 - <https://www.chpc.utah.edu/documentation/software/compilers.php>
- How to Videos
 - <https://www.chpc.utah.edu/documentation/videos/index.php>

Downtown Data Center

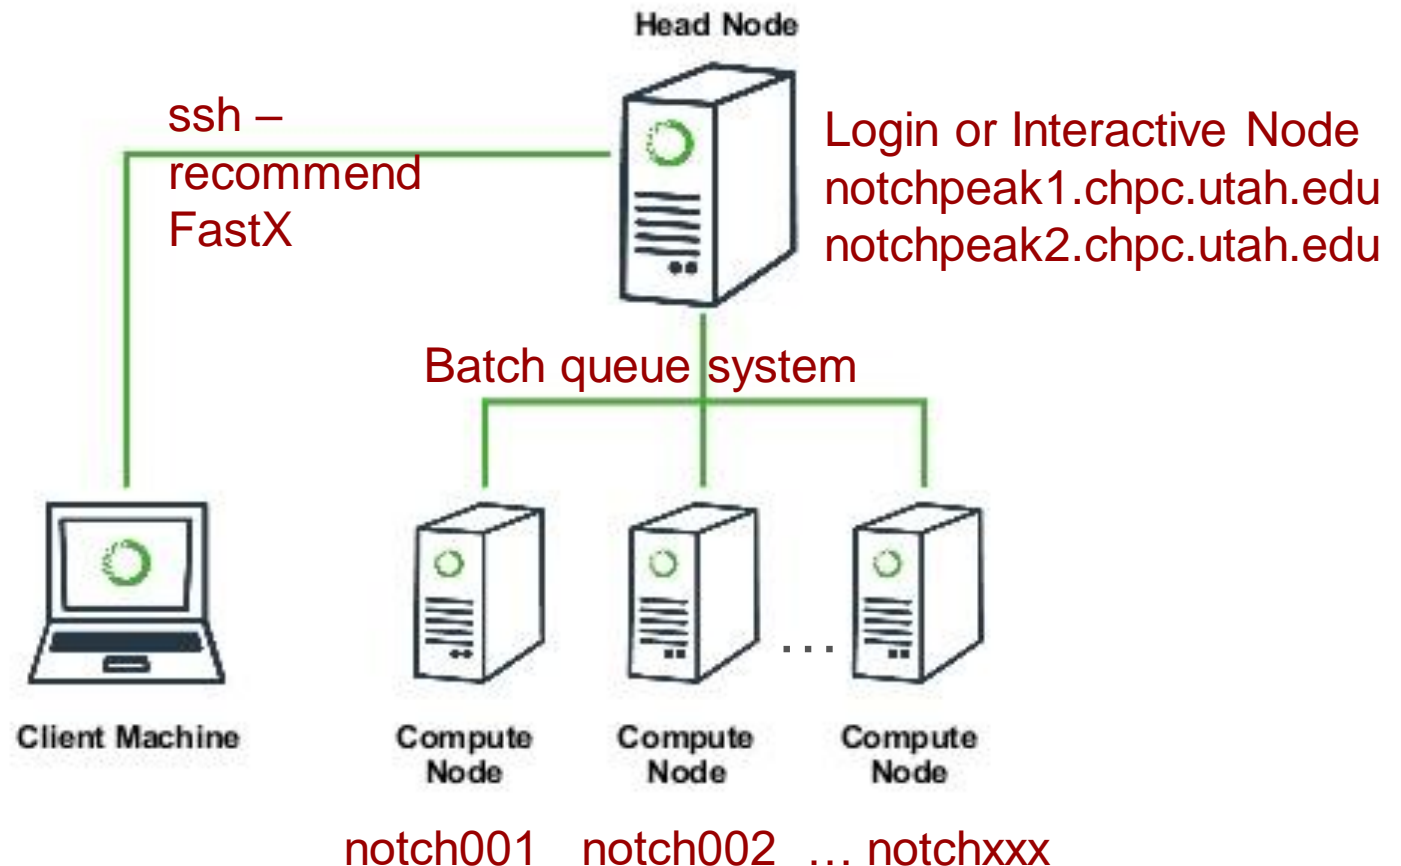
- Came online Spring 2012
- CHPC completed move to DDC Spring 2013
- Shared with enterprise (academic/hospital) groups
- 92 racks and 1.2MW of power with upgrade path to add capacity for research computing
- Metro optical ring connecting campus, data center, & internet2
- 24/7/365 facility



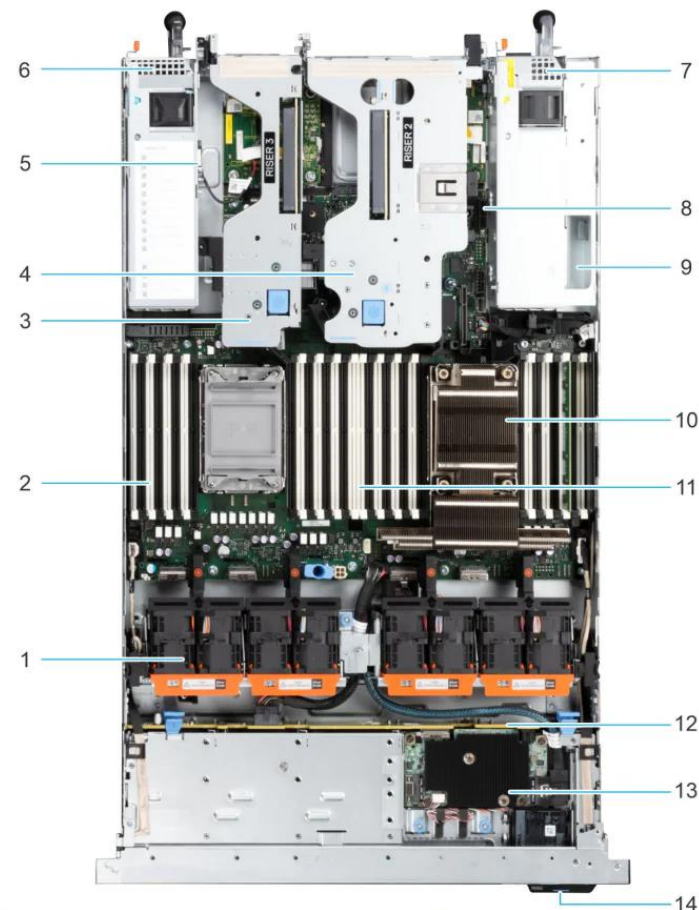
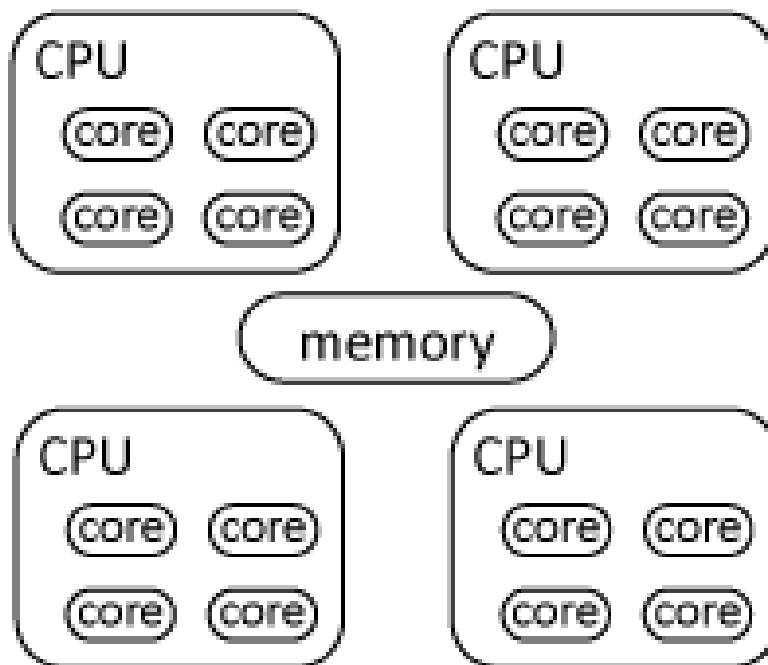
CHPC Resources & Services

- **Computational Clusters** – Notchpeak, Kingspeak, Lonepeak, Ash
- **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; consultations

Cluster Architecture Diagram



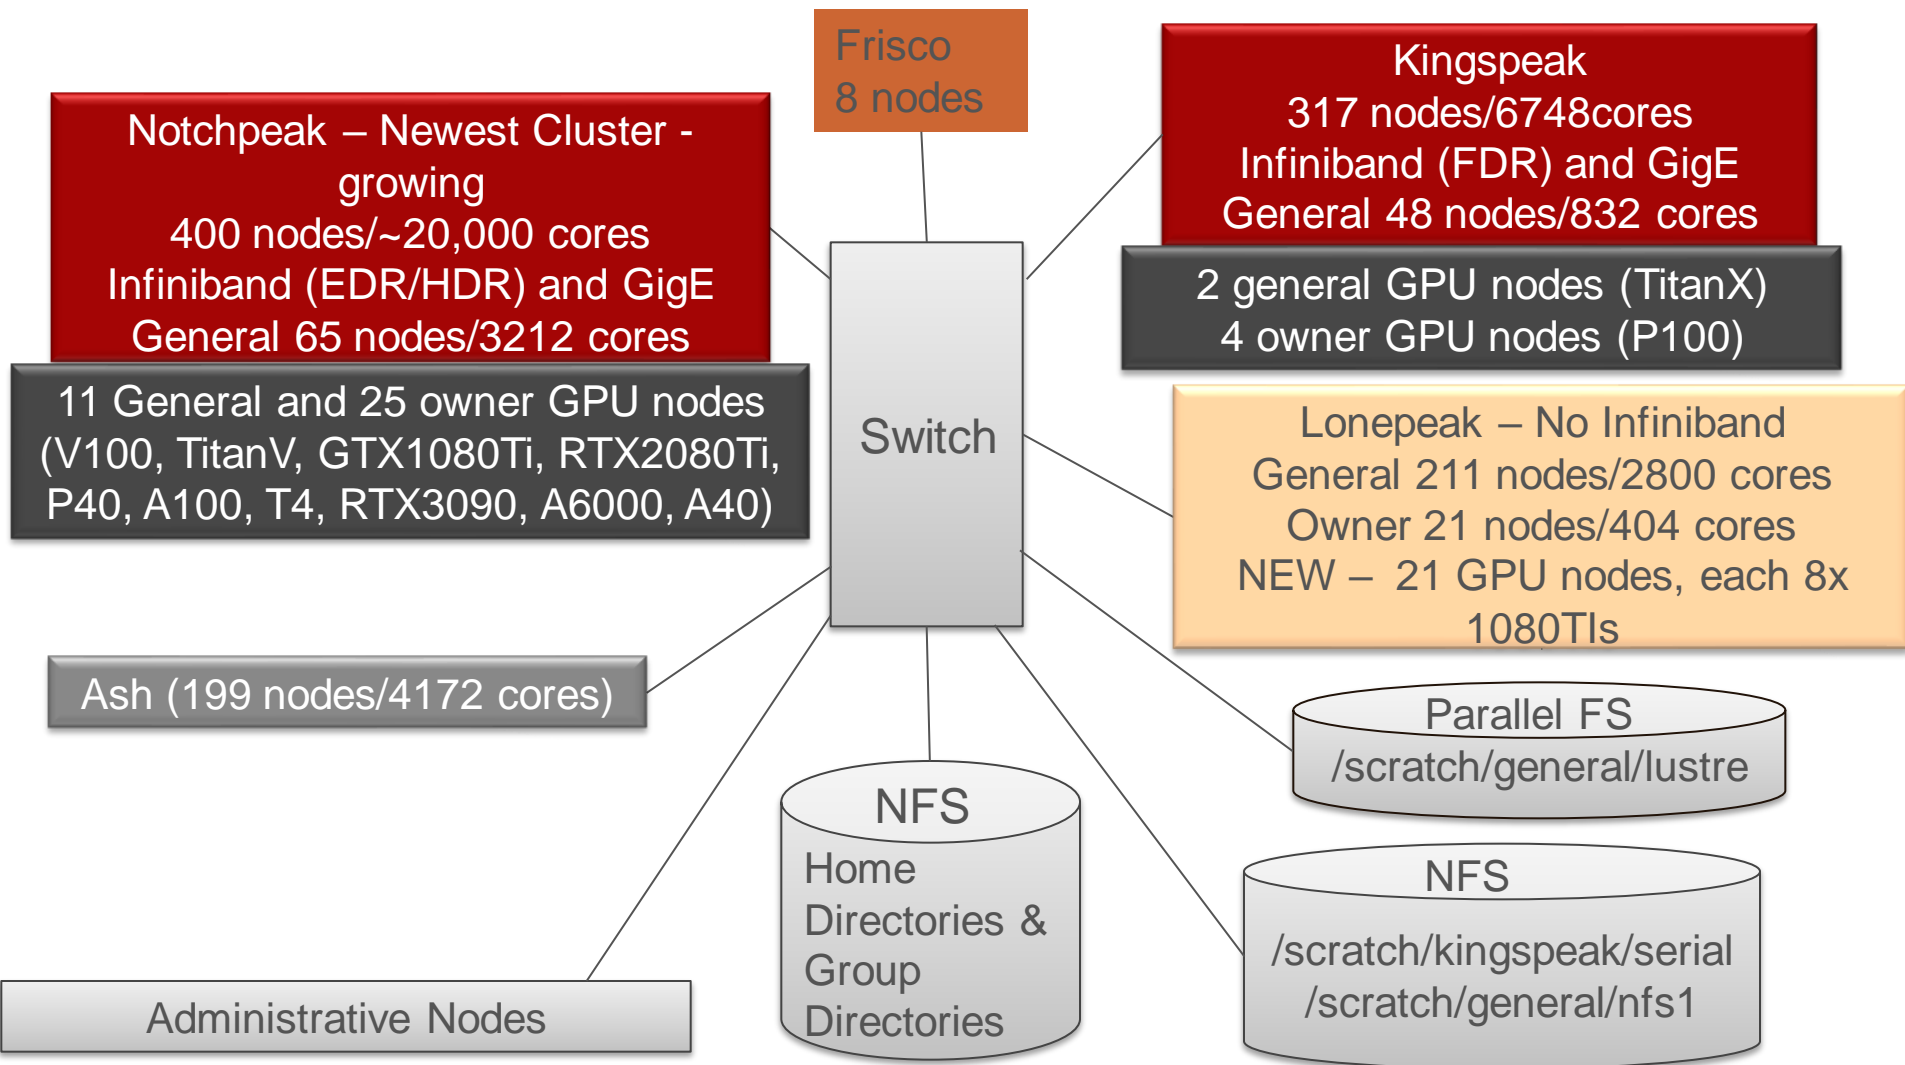
Node



1. Cooling fan
2. Memory DIMM socket for processor 2 (B2)
3. Riser 3
4. Riser 2
5. Intrusion switch
6. Power supply unit (PSU 2)
7. Power supply unit (PSU 1)
8. iSDM Internal USB card port
9. BOSS S2 module
10. Heat sink for processor 1
11. Memory DIMM socket for processor 1 (A6)
12. Drive backplane
13. fPERC
14. Service Tag

CENTER FOR HIGH PERFORMANCE COMPUTING





CHPC Clusters - Condominium Model

- Notchpeak (and Redwood in Protected Environment)
 - General resources – Allocation process for cycles
 - Out of allocation – freecycle allowed (with preemption)
 - Guest access to owner nodes
- Owner nodes – new purchased added to Notchpeak
 - Users from all groups allowed on as owner-guest when not in use (preemption)
 - Found on notchpeak, kingspeak, lonepeak (and redwood in PE)
 - 56 core Intel XeonSP, 3rd gen IceLake @ 2.0GHz cpu speed, 256GB RAM, 4 TB local SSD drive, 7 year warranty, HDR IB connectivity @ ~\$8,600/node; with 512 GB RAM ~\$10,000; with 1 TB ~\$13,400
 - 64 core AMD, 3rd gen Milan 7713P processor @2.0 GHz base cpu speed, 256 GB memory ~\$7,800; with 512 GB ~\$9,300; with 1 TB ~\$13,200
 - Getting pricing on Intel Sapphire Rapids and AMD Genoa based systems
- Kingspeak, Lonepeak
 - General nodes – no allocation, no preemptions
 - Guest access to owner nodes as mentioned above
- Ash – Owner cluster
 - All users have guest access as smithp-guest (preemption)

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
- Default 50 GB home directory
- HPC scratch space
- Some application and licensing costs
- High speed data transfer nodes (DTNs)
- 'Science DMZ' network
- CHPC Staff time

More on Access to Clusters

Allocations and node ownership status	What resource(s) are available
No general allocation, no owner nodes	<u>Unallocated general nodes</u> <u>Allocated general nodes in freecycle mode</u> - not recommended <u>Guest access on owner nodes</u>
General allocation, no owner nodes	<u>Unallocated general nodes</u> <u>Allocated general nodes</u> <u>Guest access on owner nodes</u>
Group owner nodes, no general allocation	<u>Unallocated general nodes</u> <u>Allocated general nodes in freecycle mode</u> - not recommended <u>Group owned nodes</u> <u>Guest access on owner nodes of other groups</u>
Group owner node, general allocation	<u>Unallocated general nodes</u> <u>Allocated general nodes</u> <u>Group owned nodes</u> <u>Guest access on owner nodes of other groups</u>

See <https://www.chpc.utah.edu/documentation/guides/index.php#parts>

Query your allocation

~]\$ myallocation

You have a **general** allocation on **kingspeak**. Account: **chpc**, Partition: **kingspeak**

You have a **general** allocation on **kingspeak**. Account: **chpc**, Partition: **kingspeak-shared**

You can use **preemptable** mode on **kingspeak**. Account: **owner-guest**, Partition: **kingspeak-guest**

You can use **preemptable GPU** mode on **kingspeak**. Account: **owner-gpu-guest**, Partition: **kingspeak-gpu-guest**

You have a **GPU** allocation on **kingspeak**. Account: **kingspeak-gpu**, Partition: **kingspeak-gpu**

You have a **general** allocation on **notchpeak**. Account: **chpc**, Partition: **notchpeak**

You have a **general** allocation on **notchpeak**. Account: **chpc**, Partition: **notchpeak-shared**

You can use **preemptable GPU** mode on **notchpeak**. Account: **owner-gpu-guest**, Partition: **notchpeak-gpu-guest**

You have a **general** allocation on **notchpeak**. Account: **notchpeak-shared-short**, Partition: **notchpeak-shared-short**

You can use **preemptable** mode on **notchpeak**. Account: **owner-guest**, Partition: **notchpeak-guest**

You have a **GPU** allocation on **notchpeak**. Account: **notchpeak-gpu**, Partition: **notchpeak-gpu**

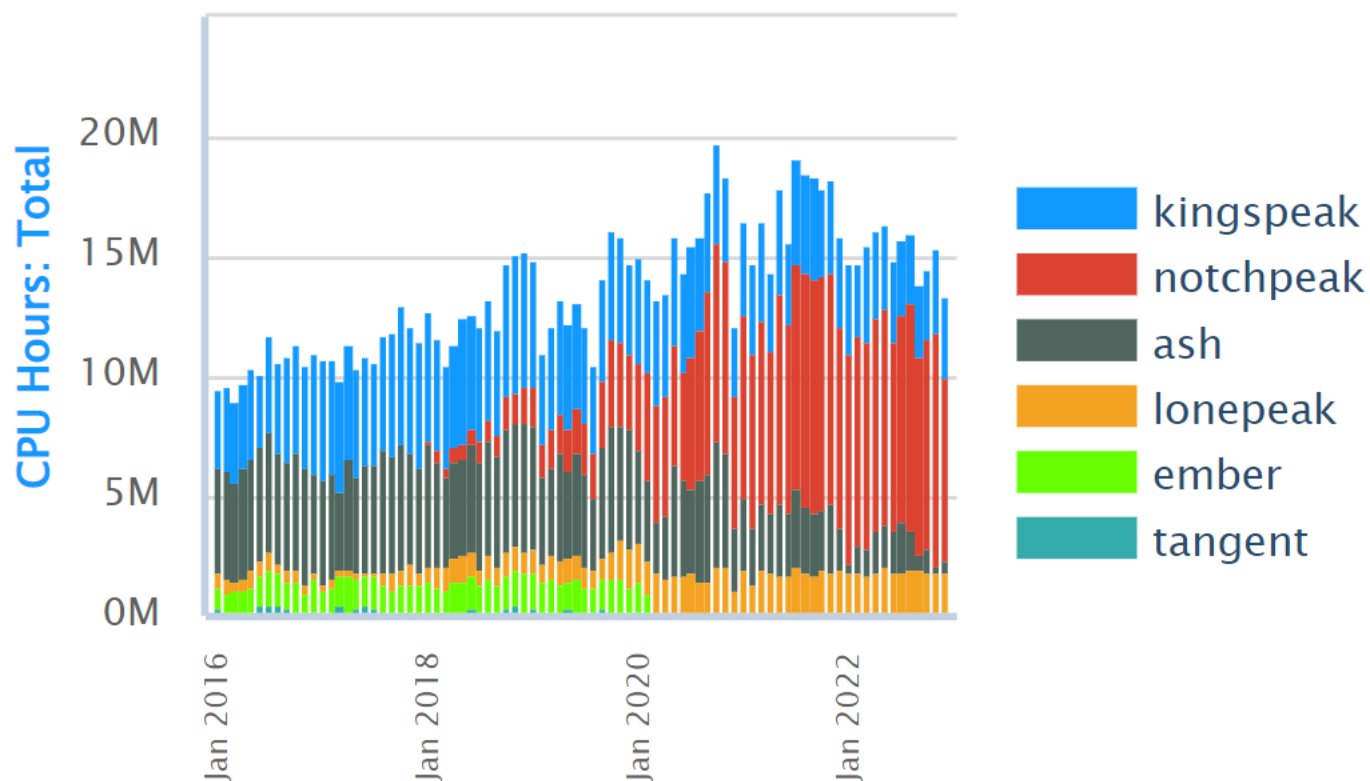
You have a **general** allocation on **lonepeak**. Account: **chpc**, Partition: **lonepeak**

You have a **general** allocation on **lonepeak**. Account: **chpc**, Partition: **lonepeak-shared**

You can use **preemptable** mode on **lonepeak**. Account: **owner-guest**, Partition: **lonepeak-guest**

You can use **preemptable** mode on **ash**. Account: **smithp-guest**, Partition: **ash-guest**

Core Hour Usage



YEAR	CORE HOURS
2016	125M
2017	135M
2018	161M
2019	178M
2020	202M
2021	221M
2022	199M

Since Jan 2016 – over 3,600 users from ~700 groups ran more than 35M jobs

Storage Options

- *Home Directories* -- /uufs/chpc.utah.edu/common/home/<uNID>
 - New solution in place – VAST
 - Home directories 50 GB, not backed up
 - Groups can larger home directory per group at \$900/TB for 5 year (warranty period); this comes with backup (nightly incremental, weekly full, 2 week retention)
- *Group Level File Systems*
 - NFS mounted group space @\$150/TB – no backup included in this price
- *Scratch File Systems*
 - For use by all users; no quotas; scrubbed of files older than 60 days
 - 600 TB NFS mounted file system (/scratch/general/nfs1)
 - 1 PB VAST scratch space (/scratch/general/vast)
- *Disk Based Archive Storage*
 - Archive at \$150/TB

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/vast
 - /scratch/general/nfs1
 - home directories and group spaces (*don't use for large i/o!*)

Remember 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
- Pls must have account and will need to approve accounts for any members of their research group (can designate a 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
 - **cluster**1.chpc.utah.edu or **cluster**2.chpc.utah.edu where **cluster** is notchpeak, kingspeak, ash-guest, lonepeak (redwood in PE)
 - There are also eight frisco nodes (frisco1-8) that act as kingspeak interactive nodes (and two bristlecone nodes in the PE)
- Interactive nodes only used for short compiles, editing and very short test runs
- ***No more than 15 minutes (60 minutes on frisco/bristlecone nodes) and no jobs of any length that make heavy use of cpu or memory!***
- Have script (Arbiter) which watches running processes and notifies users when in violation of the acceptable usage policy

Accessing Clusters (2)

- Use FastX from Mac, Windows, or Linux desktops -- preferred
 - <https://www.chpc.utah.edu/documentation/software/fastx.php>
- Alternatively:
 - From windows need ssh client
 - MobaXterm <https://mobaxterm.mobatek.net/>
 - PuTTY <http://www.chiark.greenend.org.uk/~sgtatham/putty/>
 - Xshell <https://www.netsarang.com/en/xshell>
 - 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 then 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.

OpenOnDemand

- An alternate to open a shell is to use the Open OnDemand (OOD) portal to access CHPC resources
- Allows users to view, edit, upload and download files, create, edit, submit and monitor jobs, run GUI applications, and connect via SSH, all via a web browser and with a minimal knowledge of Linux and scheduler commands.
- Presentation on OOD – Thursday, February 9, 2023

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
- More about this and getting started on Linux in CHPC’s 3-part Introduction to Linux Series starting Thursday, June 8, 2023

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, June 1, 2023

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>
- Presentation on Slurm – Tuesday, June 6, 2023

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/r8
- Use ‘**module spider**’ to search for applications installed

Allocation

- General Allocation Process Information
 - <https://www.chpc.utah.edu/documentation/policies/1.4AllocationPolicies.php>
 - <https://chpc.utah.edu/userservices/allocations.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
 - Maximum request 300,000 core hours per quarter
 - Simplified small request form for requests up to 30,000 core hours/quarter
 - https://chpc.utah.edu/role/user/allocation_small_form.php
- Quick allocation
 - https://www.chpc.utah.edu/apps/profile/allocation_quick_form.php
 - Maximum quick allocation request – 1 quarter, up to 30,000 core hours
- Check usage -- <https://www.chpc.utah.edu/usage/cluster/current-project-general.php>

Windows Statistics Server

- Beehive – refreshed 2019
 - 48 physical cores, 512TB memory
- Presently has the following software installed
 - SAS 9.4 with text miner
 - 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 – hardware refreshed 2019; expanded 2021
- VM storage
- Have community mysql/mssql VMs, git repositories, web servers, etc
- All 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
1	4	2	50	\$425
2	8	2	100	\$615
4	16	4	200	\$990
8	32	8	400	\$1745
16	64	8	800	\$3250

Additional VM storage available, in 100GB increments, at a cost of \$1100/TB.

Presentation Schedule

- Have Fall, Spring and Summer schedules
- Current semester
 - <https://chpc.utah.edu/presentations/summer2023chpcpresentationschedule.php>
- Mix of presentations, hands on sessions (linux, python)
- If you have suggestions for other presentation contact us
- If you are interested in presentations for classes, research groups contact us

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
- Mailing Lists:
chpc-hpc-users@lists.utah.edu used to send messages to users