Overview of CHPC

Anita Orendt
Assistant Director
Research Consulting & Faculty Engagement
anita.orendt@utah.edu

CHPC Mission

- The University of Utah's Center for High Performance Computing provides large-scale computer systems, storage, networking, and the expertise to optimize the use of these high-end computer technologies.
- CHPC supports faculty and research groups whose main focus requires computing, storage, and advanced networking as core instruments central to their research.

Downtown Data Center

- Came online spring 2012
- Shared with enterprise (academic/hospital) groups (wall between rooms)
- 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




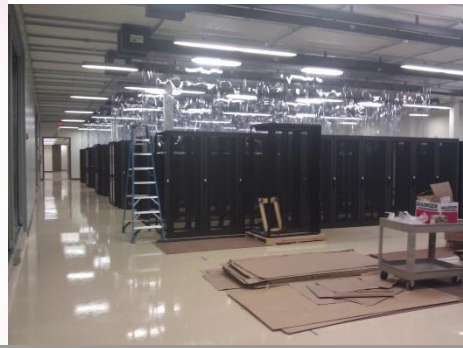

Overview

- HPC Clusters
- Storage
- Protected Environment
- Software
- Other Services
- Accounts & Allocation
- Getting Started
- Getting Help

CHPC Clusters

- Condominium Model
- General resources – No charge for usage
  - Allocation process for cycles on kingspeak and ember
  - Out of allocation – freecycle allowed (with preemption)
- Owner nodes
  - 20 core Intel IvyBridge (2.5GHz cpu speed), 64GB RAM, 1TB local hard drive, 5 year warranty, IB connectivity @ ~$5500/node
  - Others allowed on as owner-guest when not in use (preemption)
- Lonepeak – no allocation and no preemption on general nodes
- Owner cluster – ash

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
- Emerging “science DMZ” network

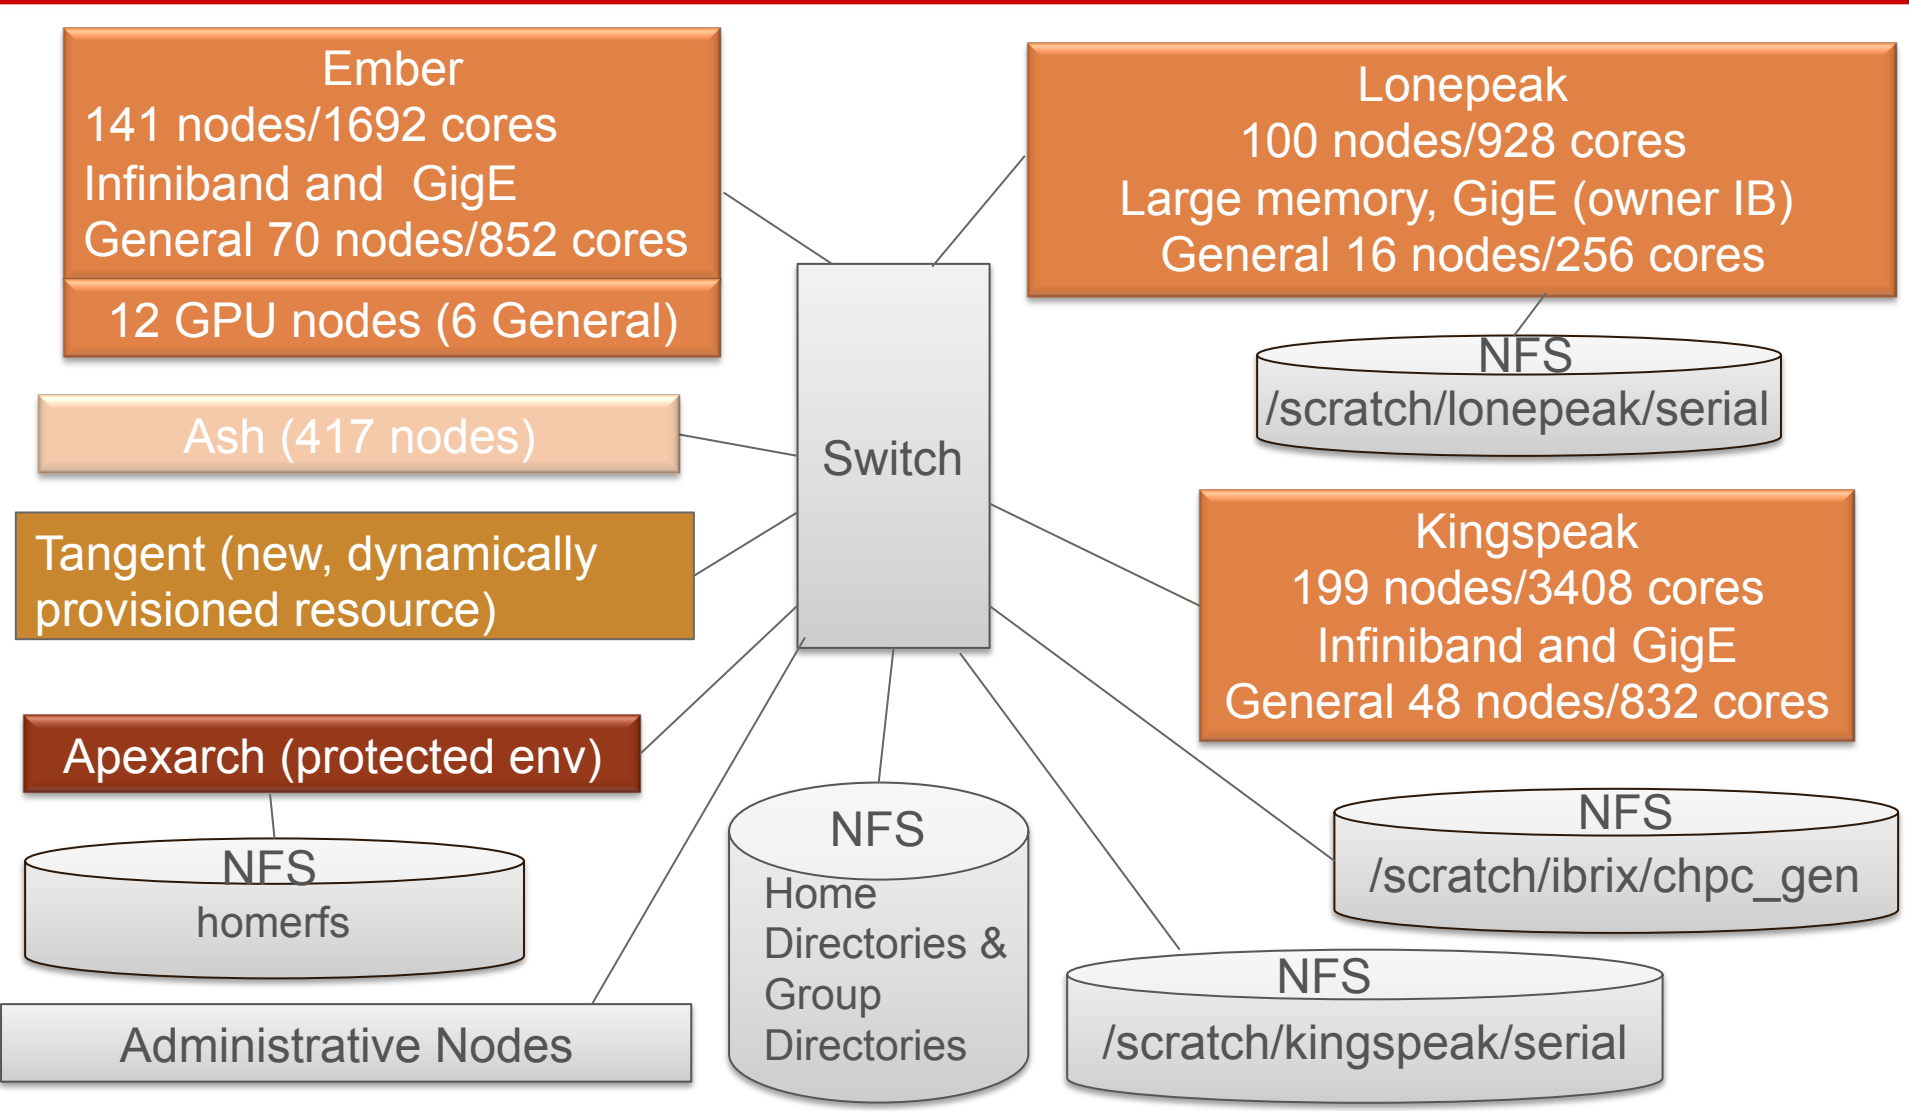
Ember
141 nodes/1692 cores
Infiniband and GigE
General 70 nodes/852 cores
12 GPU nodes (6 General)
Ash (417 nodes)
Tangent (new, dynamically provisioned resource)
Apexarch (protected env)
NFS
homerfs
Administrative Nodes
Switch
NFS
Home Directories & Group Directories
Lonepeak
100 nodes/928 cores
Large memory, GigE (owner IB)
General 16 nodes/256 cores
NFS
/scratch/lonepeak/serial
Kingspeak
199 nodes/3408 cores
Infiniband and GigE
General 48 nodes/832 cores
NFS
/scratch/ibrix/chpc_gen
NFS
/scratch/kingspeak/serial

Useful Information

- Getting Started Guide
  - https://www.chpc.utah.edu/docs/manuals/getting_started/
- CHPC policies
  - https://wiki.chpc.utah.edu/display/policy/CHPC+Policy+Manual
- Cluster Usage Guides
  - https://wiki.chpc.utah.edu/display/DOCS/Kingspeak+User+Guide
  - https://wiki.chpc.utah.edu/display/DOCS/Ember+User+Guide
  - https://wiki.chpc.utah.edu/display/DOCS/Lonepeak+User+Guide
- Application Documentation
  - https://wiki.chpc.utah.edu/display/DOCS/CHPC+Software+Documentation

Storage

- CHPC home directory space NFS mounted on all HPC platforms
  - /uufs/chpc.utah.edu/common/home/<uNID>
  - Home directories not backed up (some exceptions)
  - Periodic Archive Option available with purchase of tapes
- Groups can purchase home and/or group space
  - 1TB max home directory/group @$500/TB – 5 year warranty; tape backup (nightly incremental/weekly full)
  - Group space @$210/TB; can get quarterly archives with purchase of tapes

Scratch File Systems

- Ember, Kingspeak, and Ash

  /scratch/ibrix/chpc_gen (56 TB)

  /scratch/kingspeak/serial (175 TB)

- Lonepeak

  /scratch/lonepeak/serial (33 TB)

- NFS mounted on all interactive nodes
- Scrubbed of files older than 60 days
- All nodes also have a local scratch space

File Systems

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

Remember NFS mounted spaces are a shared resource!

To check the current status of the file systems (and clusters)

-- https://www.chpc.utah.edu/chpc/systems/graphs.php

Protected Environment

- Dedicated protected resources for handling of data/ projects with protected information
  - Currently HIPAA, looking at FISMA & FERPA
- Have both general, owner and VM resources
- HPC – apexarch with associated homerfs file system
- Significant area of growth for CHPC
- Described in recent paper (Bradford *et al*.) *www.ncbi.nlm.nih.gov/pubmed/23911553*

Software

- CHPC supports a wide array of software – partial list at https://wiki.chpc.utah.edu/display/DOCS/CHPC+Software+Documentation
- For ALL clusters (built with mpich2): /uufs/chpc.utah.edu/sys/pkg/***packagename***
- Specific build for EACH cluster (usually due to Infiniband/MPI needs: /uufs/***cluster*** /sys/pkg/***packagename***
- We work with research groups to install needed software, included software licensed only to a given group

Other Services

- Windows Statistics Servers
  - Two servers each with 48 core and 512 GB RAM, one for general (kachina) and the other for use in protected environment (swasey)
  - SAS 9.3 with text miner, SPSS, R, STATA, Mathematica, Matlab
- Virtual Machine (VM) Farms
  - For applications that do not fit in our normal HPC environment
  - Allows for maximum use of infrastructure and minimizes physical servers to be managed
  - Two farms – one for general needs and the second for needs in the protected environment
  - Setting up MOU/SLA with VM owners prior to provisioning of VM

Accounts

- CHPC uses campus uNID and password
- Account Application Procedure – Online process
  - Complete CHPC account form

    http://www.chpc.utah.edu/apps/profile/account_request.php
  - PIs must have account and also approve accounts for any members of their research group (can delegate)
  - 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

Security Policies

- No clear text passwords, use ssh and scp
- You may not share your account under any circumstances
- Don't leave your terminal unattended while logged into your account
- Do not introduce classified or sensitive work onto CHPC systems unless on Protected Environment
- Do not try to break passwords, tamper with files etc.
- Do not distribute or copy privileged data or software
- Report suspicions to CHPC (security@chpc.utah.edu)
- See http://www.chpc.utah.edu/docs/policies/security.html for more details

Allocation

- General Allocation Process Information
  - https://wiki.chpc.utah.edu/display/policy/1.4+Allocation+Policies
- 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 SU
    - 1 core-hour = 1SU on ember & 1 core-hour = 1.5 SU on kingspeak
- Quick allocation
  - https://www.chpc.utah.edu/apps/profile/allocation_quick_form.php
- Check usage -- http://www.chpc.utah.edu/docs/allocations/

Accessing Clusters

- Login or interactive nodes with each cluster
  - ssh –X kingspeak.chpc.utah.edu
  - ssh –X ember.chpc.utah.edu
  - ssh –X lonepeak.chpc.utah.edu
  - ssh –X ash-guest.chpc.utah.edu
- 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!***

SSH clients for windows

- PuTTY
  - http://www.chiark.greenend.org.uk/~sgtatham/putty/
- XShell4
  - http://www.netsarang.com/download/down_xsh.html

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 closing, allowing users to reattach to the session from the same or other systems
- Server on all interactive nodes as well as several dedicated fastx servers
- 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://wiki.chpc.utah.edu/display/DOCS/FastX
- Download client from CHPC site at https://www.chpc.utah.edu/apps/profile/software.php
- Do install
- Start program
- Set host to kingspeak1.chpc.utah.edu OR kingspeak2.chpc.utah.edu
- Have dedicated fastX servers – frisco1-frisco6
  - Frisco5,6 for vgl applications

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 or edit!***
- Shell set to tcsh – can change to bash at www.chpc.utah.edu (choose PROFILE, login, select edit profile)
- Will occasionally request that users update these files
- At end '#source .aliases' – use this for customization
- Login scripts (see with ls –a)
  - .tcshrc. .bashrc, .bash_profile

Batch System Information

- Used to access compute nodes
- http://docs.adaptivecomputing.com/
- Two components
  - Torque (OpenPBS) -- Resource Manager
  - Moab (Maui) – Scheduler – enforces policies and sets job priorities

Some Basic Batch commands

- qsub <script> -- *submit job*
- qstat -- *lists all jobs in queue*
- qdel $jobid -- *deletes job*
- pbsnodes –a -- *lists all nodes in cluster*
- showq -- *lists all jobs in queue {-r for running, -i idle, -b blocked}*
- showstart $jobid -- *shows estimated start time*
- checkjob $jobid -- *give more detailed info about job*

Batch Policies

- https://wiki.chpc.utah.edu/pages/viewpage.action?pageId=281706582
- https://wiki.chpc.utah.edu/display/policy/2.1.3+Kingspeak+Job+Scheduling+Policy
- Walltime limit – 72 hours
- Accounts
  - **piname** for general allocation
  - **piname-{kp,em, lp}** for owner nodes on kingspeak, ember or lonepeak
  - **owner-guest** for guest jobs on owner nodes on kingspeak, ember or lonepeak
  - **smithp-guest** for guest jobs on ash

Sample Batch Script

```
#PBS -S /bin/tcsh
#PBS -l nodes=2:ppn=20,walltime=1:00:00
#PBS -A account
#PBS -N myjob

# Create scratch directory
mkdir -p /scratch/kingspeak/serial/$USER/$PBS_JOBID
# Change to working directory
cd /scratch/kingspeak/serial/$USER/$PBS_JOBID
# Copy data files to scratch directory
cp $HOME/work_dir/files /scratch/kingspeak/serial/$USER/$PBS_JOBID
#Set Environment
source /uufs/kingspeak.peaks/sys/pkg/mvapich2/1.9i/etc/mvapich2.sh
#Execute Job
mpiexec -np 40 -machinefile $PBS_NODEFILE ./executable
# Copy files back home and cleanup

cp * $HOME/work_dir && rm -rf /scratch/kingspeak/serial/$USER/$PBS_JOBID
```

Coming Soon

- Redesigned CHPC website
- Database of applications to make it easier to find what packages are installed
- Modules option for on clusters
- Exploring move from Moab/Torque to SLURM (already in place on tangent)

Getting Help

- CHPC website and wiki
  - www.chpc.utah.edu and wiki.chpc.utah.edu
    - Getting started guide, cluster usage guides, software manual pages, CHPC policies
- Jira Ticketing System
  - Email: issues@chpc.utah.edu
- Help Desk: 405 INSCC, 581-6440 (9-5 M-F)
- We use chpc-hpc-users@lists.utah.edu for sending messages to users; also have Twitter accounts for announcements -- @CHPCOutages & @CHPCUpdates

CHPC Spring Presentation Series

Date	Presentation Title
January 20	Protected Environment at CHPC**
January 27	Introduction to Linux, part 1*
February 3	Introduction to Linux, part 2*
February 6	SEDE HPC Monthly Workshop: OpenACC***
February 10	Introduction to Linux, part 3*
February 10	NLP and AI Services at CHPC**
February 17	Introduction to Parallel Computing
February 24	Hands-on Introduction to Python, Part 1*
March 3	Hands-on Introduction to Python, Part 2*
March 4-5	XSEDE HPC Monthly Workshop:MPI***
March 10	Hands-on Introduction to Numpy/Scipy*
March 24th	Intel Software Development Tools

All in INSCC Auditorium 1-2pm unless: *1-3pm hands on, **BMI Classroom (421 Wakara Way Room 1470) at 2-3pm, ***XSEDE workshops 9am-3pm

XSEDE Resource Support –xsede.org

Campus Champions: Julia Harrison & Anita Orendt

- Have CC allocation to get you started
- Can do startup allocation
- Yearly allocation for research
- Special allocations for teaching
- Gateways

XSEDE

Extreme Science and Engineering Discovery Environment

NSF

Advanced Cyberinfrastructure – Research and Education Facilitator (ACI-REF)

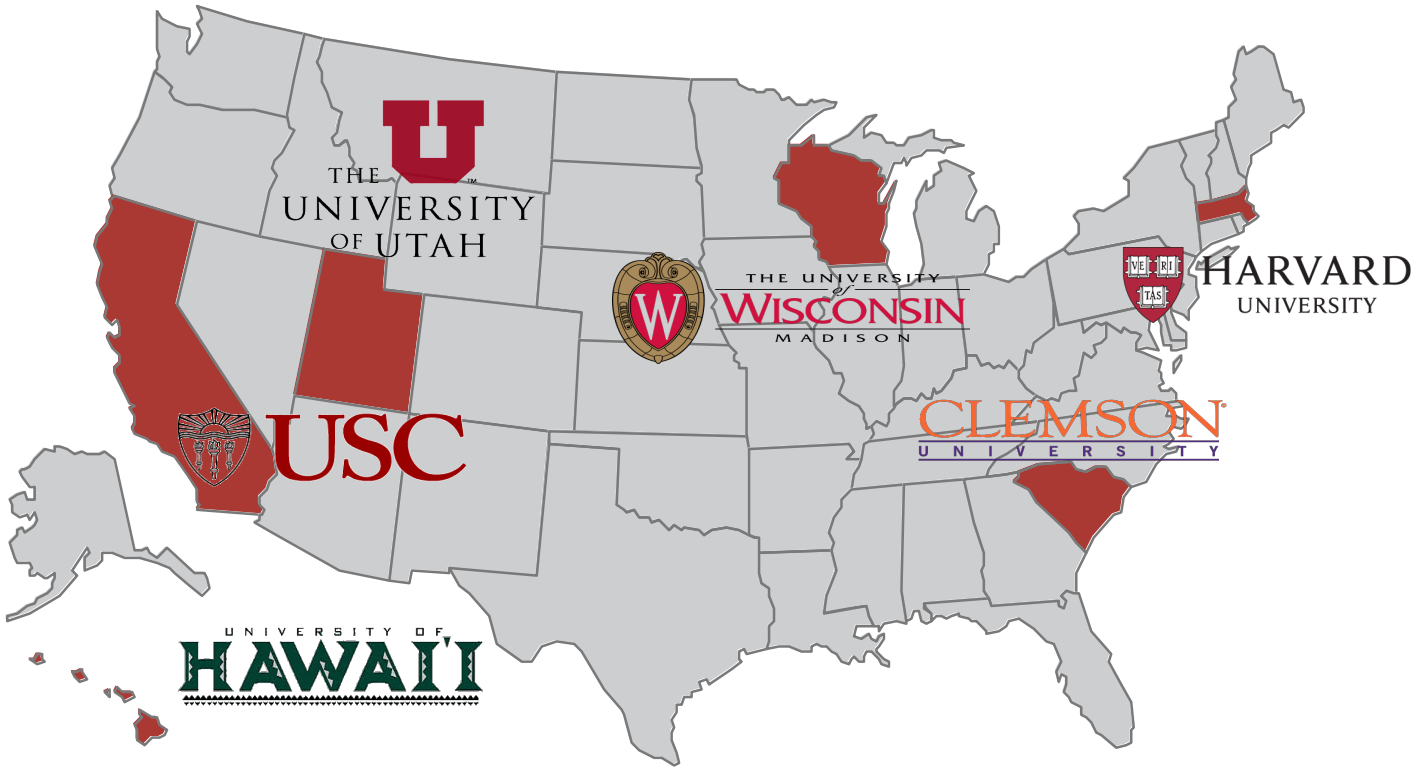

Advancing scientific discovery through a national network of Advanced Cyberinfrastructure (ACI) Research and Education Facilitators (ACI-REFs).

ACIREF Project Goals

- Enable scientific and research outcomes that may not be possible without the ACI-REF community
- Create a community that fosters mentorship and sharing of expertise at individual (ACI-REF to researcher) and campus (institution to institution) levels to help others advance their capabilities and outcomes
- Develop a framework for sharing best practices and lessons learned for facilitation of interactions with researchers & share with other campuses
- Increase the breadth and depth of interactions with researchers on campuses in various disciplines
- Enable researchers to be more competitive for research funding & foster new collaborations among researchers, campuses, and other organizations
- CIO-level commitment to ensure long-term sustainability
- Develop a plan for enabling other institutions to join the consortium