

# Overview of CHPC

Martin Čuma, PhD

Center for High Performance Computing

[m.cuma@utah.edu](mailto:m.cuma@utah.edu)

Spring 2014

# Overview

- CHPC Services
- HPC Clusters
- Specialized computing resources
- Access and Security
- Batch (PBS and Moab)
- Allocations
- Getting Help

# CHPC Services

- HPC Clusters (focus of this talk)
- Advanced network services
- Specialized research services  
(statistics, database support,..)
- Campus infrastructure support

<http://www.chpc.utah.edu/docs/services.html>

# Telluride



# Telluride configuration

- Administrative nodes
- ~ 90 compute nodes, 2x4 core Intel Xeon E5430 2.8 GHz
- $\geq$  2 GB Memory per core
- InfiniBand interconnect



## Telluride access and use:

- ssh to the interactive login node  
`ssh -X telluride.chpc.utah.edu`
- submit job to run a program/calculation  
`qsub myjob.pbs`

# Cisco UCS standalone nodes



# UCS nodes - turretarch

- 8 2x6 core Intel Xeon X5650 2.67 GHz, 96 GB RAM
  - turretarch10-13, 18-21
- 8 2x10 core Intel Xeon E7-2850 2.0 GHz, 256 GB RAM
  - turretarch14-17,22-25
- Special access control
- No scheduling – run directly on the node you ssh to and be aware that it may be shared with others



# Ember



# Ember continued

- 382 dual-six core nodes (4584 cores) - Intel Xeon 2.8 GHz (Westmere)  
12 dual-six cores with each 2 M2090 GPU cards
- $\geq$  2 GB memory per core
- Interactive node:  
`ssh -X ember.chpc.utah.edu`

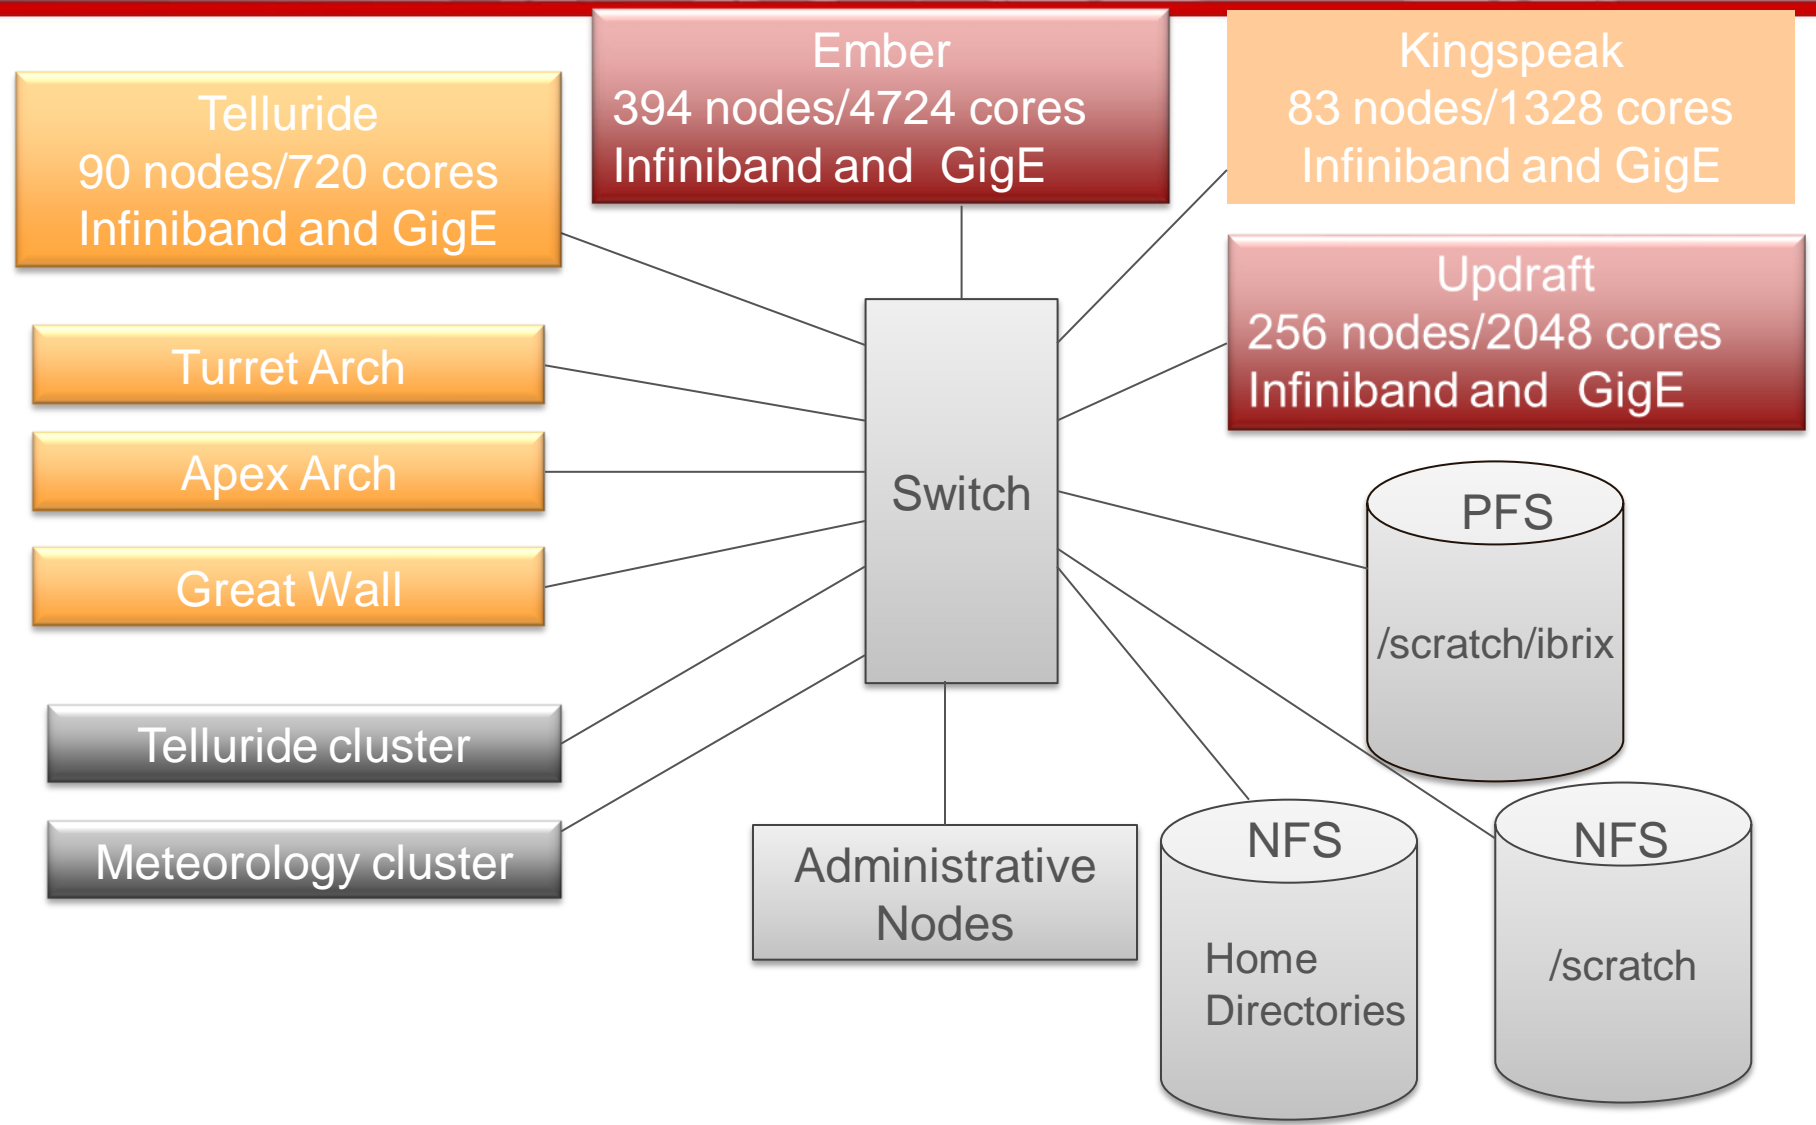
# Kingspeak

- 83 dual-eight core nodes (1328 cores) – Intel Xeon 2.6 GHz (Sandybridge)
- 4 GB memory per core
- Interactive node:  
`ssh -X kingspeak.chpc.utah.edu`

# Special purpose nodes

- apexarch: 10 nodes (2.8 GHz), memory  $\geq$  2GB per core
- turretarch: 4 cores, 16 GB (statistics: OS Linux)
- kachina: 1 node with 48 cores, 512 GB Memory, 1TB drive (statistics: OS Win.)
- swasey: 1 node - HIPAA (statistics: OS Win.)
- greatwall: 2 nodes (2.8 GHz/48 GB) & 2 nodes (2.2GHz/32 GB)





# Filesystems

- 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
- Scratch space for HPC systems
  - Files older than 30 days eligible for removal

- On Ember & Kingspeak
  - /scratch/ibrix/chpc\_gen (56 TB)
  - /scratch/ibrix/icse (156 TB)
- On Telluride & Ember
  - /scratch/general (3.6 TB)
  - /scratch/serial (16 TB)
  - /scratch/uintah (7.7 TB)
- On Ember & Kingspeak
  - /scratch/kingspeak/serial (175 TB)

# Filesystems

- Access speed based on connectivity
- Local disk fastest – local to each node
  - /scratch/local
- NFS
  - /scratch/serial,/scratch/general,/scratch/uintah
  - home directories (don't use for large i/o!)  
/uufs/chpc.utah.edu/common/home/<uNID>
- PFS
  - /scratch/ibrix/{chpc\_gen,icse}



- All clusters have a local scratch: /scratch/local

Telluride: 66+ GB

Ember: 414+ GB

Kingspeak: 337+ GB

- To check the **current** status of the file system

<http://www.chpc.utah.edu/chpc/systems/graphs.php?g=scratch>

# Getting an Account

- Application Procedure  
[http://www.chpc.utah.edu/apps/profile/account\\_request.php](http://www.chpc.utah.edu/apps/profile/account_request.php)
- CHPC uses the campus uNID and password
- Passwords maintained by the campus information system

# Logins

- Access – interactive sites
  - ssh -X telluride.chpc.utah.edu
  - ssh -X turretarch10.chpc.utah.edu
  - ssh -X ember.chpc.utah.edu
  - ssh -X kingspeak.chpc.utah.edu
- No (direct) access to compute nodes

# 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
- Use a good password and protect it



# Security Policies

- 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](mailto:security@chpc.utah.edu))
- Please see <http://www.chpc.utah.edu/docs/policies/security.html> for more details

# Getting Started

- CHPC Getting started guide
  - [www.chpc.utah.edu/docs/manuals/getting\\_](http://www.chpc.utah.edu/docs/manuals/getting_started)  
[started](http://www.chpc.utah.edu/docs/manuals/getting_started)
- CHPC Environment scripts
  - [www.chpc.utah.edu/docs/manuals/getting\\_](http://www.chpc.utah.edu/docs/manuals/getting_started/code/chpc.tcshrc)  
[started/code/chpc.tcshrc](http://www.chpc.utah.edu/docs/manuals/getting_started/code/chpc.tcshrc)
  - [www.chpc.utah.edu/docs/manuals/getting\\_](http://www.chpc.utah.edu/docs/manuals/getting_started/code/chpc.bashrc)  
[started/code/chpc.bashrc](http://www.chpc.utah.edu/docs/manuals/getting_started/code/chpc.bashrc)

# Environment

- All new accounts provided with login or “dot” files.
- You should update these files periodically
- On the web at [http://www.chpc.utah.edu/docs/manuals/getting\\_started/code](http://www.chpc.utah.edu/docs/manuals/getting_started/code)
- `chpc.bashrc`, `chpc.tcshrc`
- e.g. “`wget http://www.chpc.utah.edu/docs/manuals/getting_started/code/chpc.tcshrc`”

# Environment – cont.

- Running the CHPC environment script will set your path correctly for that cluster using the \$UUFSCCELL environment variable.
  - If you are logged into one of the **telluride** nodes the path /uufs/\$UUFSCCELL/sys/bin is added to your PATH env variable
  - \$UUFSCCELL will be set to **telluride**
  - showq will show you the queues for telluride
  - Similarly for the other clusters

# Batch System

- Two components
  - Torque (OpenPBS )
    - Resource Manager
  - Moab (Maui) Scheduler
- Required for all significant work
  - Interactive nodes only used for short compiles, editing and very short test runs (no more that 15 minutes!)

# Torque

- Build of OpenPBS
- Resource Manager
- Main commands:
  - `qsub <script>`
  - `qstat`
  - `qdel $jobid`
- Used with a scheduler

# Simple Batch Script

```
#PBS -S /bin/bash
#PBS -l nodes=2:ppn=8,walltime=1:00:00
#PBS -M username@utah.edu
#PBS -N myjob
# Create scratch directory
mkdir -p /scratch/serial/$USER/$PBS_JOBID
# Change to working directory
cd /scratch/serial/$USER/$PBS_JOBID
# Copy data files to scratch directory
cp $HOME/work_dir/files /scratch/serial/$USER/$PBS_JOBID
#Execute Job
/uufs/telluride.arches/sys/pkg/mvapich2/std/bin/mpirun -np 16
    -machinefile $PBS_NODEFILE ./hello
# Copy files back home and cleanup
cp * $HOME/work_dir && rm -rf /scratch/serial/$USER/$PBS_JOBID
```



# Moab scheduler

- Used on all HPC systems
- Enforces policies
- Sets priorities based on:
  - qos
  - out of allocation (“free cycle”)
- Optimizes throughput on the system

# Moab cont.

- Main commands:
  - showq
    - -i (idle jobs)
    - -r (running)
    - -b (blocked)
  - showstart <jobnumber>
  - checkjob <jobnumber>

# Batch Policies - Ember

- Qos and Time limits (CHPC nodes)
  - qos=general - 72 hours wallclock limit (**preemptor**)
  - qos=long - 14 days wallclock limit (**preemptor**) #maxproc:24
  - qos=freecycle - 72 hours wallclock limit (**preemptee**)
  - qos=general-gpu – 72 hours wallclock limit – need permission
- Dr. Philip Smith's 253 nodes: available for non-group members
  - qos=smithp-guest – 24 hours wallclock time (**preemptee**)
  - running on these nodes does **not** affect ember allocation
  - #PBS –A smithp-guest
- Use “pbsnodes –a” for node availability – also on web

# Bath Policies - Kingspeak

- Until Oct 1<sup>st</sup> -> freecycle (max. wallclock limit : 72h)
- After Oct 1<sup>st</sup> (using Ember allocation)
  - qos=general - 72 hours wallclock limit (preemptor)
  - qos=long – 14 days wallclock limit (preemptor) #maxproc:32
  - qos=freecycle – 72 hours wallclock limit (preemptee)

# Trouble Shooting

- Won't start:
  - showq -i
  - mdiag -p
  - mdiag [-v] -j <jobnumber>
  - qstat -f <jobnumber>
  - qstat -n (on running to get nodes)
  - showstart
  - **mshow -a -w [qos=general]**

# Software on the clusters

- For ALL clusters:

/uufs/chpc.utah.edu/sys/pkg/\$package/\$version

e.g: /uufs/chpc.utah.edu/sys/pkg/matlab/R2012a/bin/matlab

- Specific for EACH cluster:

/uufs/\$cluster.arches/sys/pkg/\$package/\$version

where cluster stands for kingspeak, ember, updraft, ...

e.g: /uufs/ember.arches/sys/pkg/openmpi/1.6.2i/bin/mpicc

# Allocations

- Requests for up to 4 quarters at a time due:
  - September 1<sup>st</sup> (for Oct-Dec period)
  - December 1<sup>st</sup> (for Jan-Mar period)
  - March 1<sup>st</sup> (for Apr-Jun period)
  - June 1<sup>st</sup> (for Jul-Sep period)
- Next Allocations Due **December 1<sup>th</sup>**
- Allocation should be completed online:  
[https://www.chpc.utah.edu/apps/profile/allocation\\_form.php](https://www.chpc.utah.edu/apps/profile/allocation_form.php)



# Allocations cont.

- Updraft (not on SDA)

1 SU = 1 Wallclock hour on 2GHz core

=> 1 hour on 1 Updraft node is  $(2.8/2)^* = 11.4$  SUs

- Ember & Kingspeak:

Ember: 1 SU/core => 1 h on 1 Ember node = 12 SUs

Kingspeak: 1.5 SU/core => 1h on 1 Kingspeak node = 24 SUs

- Check current balance/usage:

<http://www.chpc.utah.edu/docs/allocations/>

# Allocations cont.

## Note:

- Quick allocations (only once) may be provided for up to 5000 SU's for the current quarter only
- Unused allocations may not be carried forward to next quarter
- Allocation policies:  
<https://wiki.chpc.utah.edu/display/policy/1.4+Allocation+Policies>

# Getting Help

- <http://www.chpc.utah.edu>
- Email: [issues@chpc.utah.edu](mailto:issues@chpc.utah.edu)  
(Please do not email staff members directly –  
you may use cc)
- <http://jira.chpc.utah.edu>
- Help Desk: 405 INSCC, 581-6440  
(9-5 M-F)
- [chpc-hpc-users@lists.utah.edu](mailto:chpc-hpc-users@lists.utah.edu)