OAK RIDGE National Laboratory

OLCF Overview for New Users

Bill Renaud
OLCF User Support

ORNL is managed by UT-Battelle, LLC for the US Department of Energy

U.S. DEPARTMENT OF ENERGY

General Information

- The goal of this presentation is to give a brief introduction to using OLCF systems & OLCF policies
- This is by no means an all-inclusive presentation
- Feel free to ask questions
- Much of this may be a reminder
  - Largely an update to previous presentations:
    https://vimeo.com/channels/olcftraining/254494153
    https://vimeo.com/channels/olcftraining/343636411

Overview

1. Staying Informed
2. OLCF Policies
3. Authentication with RSA Tokens
4. Projects & User Accounts
5. Data Management
6. Using HPSS
7. Finding and Building Software
8. Using Summit
9. Training Opportunities
10. Getting Help

OAK RIDGE National Laboratory

Staying Informed

Staying Informed

- OLCF provides multiple layers of user notifications about system status and downtimes
  - Email lists
  - Status indicators on https://www.olcf.ornl.gov (at the bottom)
  - Twitter (@OLCFStatus)
- For more information, see the OLCF website: https://www.olcf.ornl.gov/for-users/user-assistance/

Staying Informed – *Email Lists*

- *Announce* lists
  - All users are required to be members
  - System-specific (*-announce@email.ornl.gov) and center-wide (ccs-announce@email.ornl.gov) lists
  - Used for major announcements, weekly notice, etc.
- *Notice* lists
  - By default, "recent" users (active in last 2 weeks)
  - Permanent opt-in/opt-out possible (contact help@olcf.ornl.gov)
  - Used for minor updates, system status, etc.

Staying Informed – *Weekly Update*

- Sent before noon (Eastern) on Wednesday
- Announcements about outages, training, etc.
- ***ALL USERS* should receive this email**
- Also posted on OLCF home page

Staying Informed – *Weekly Update*

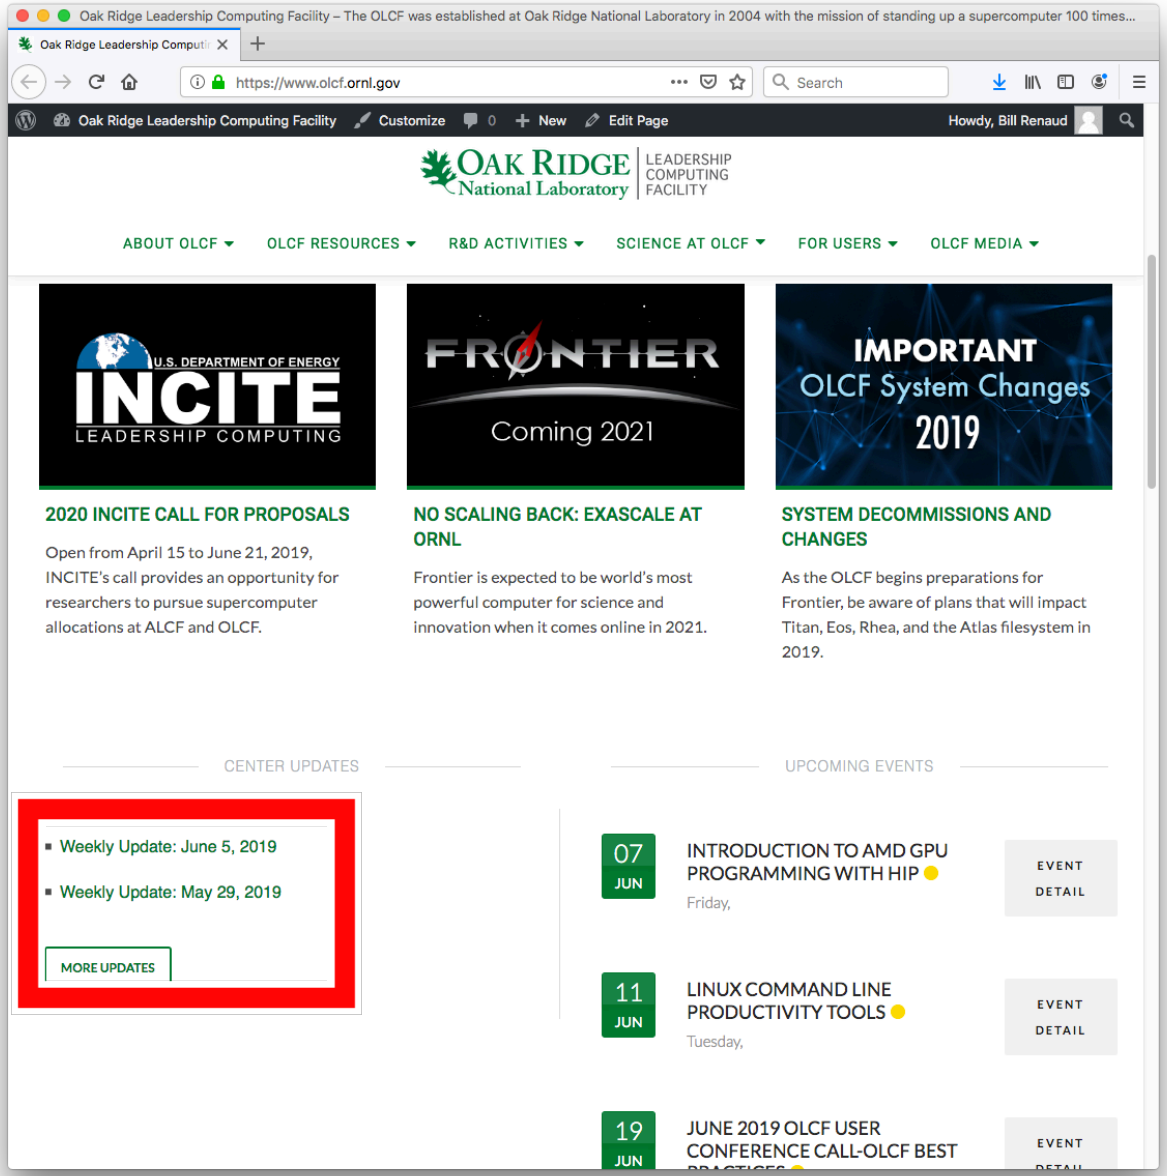

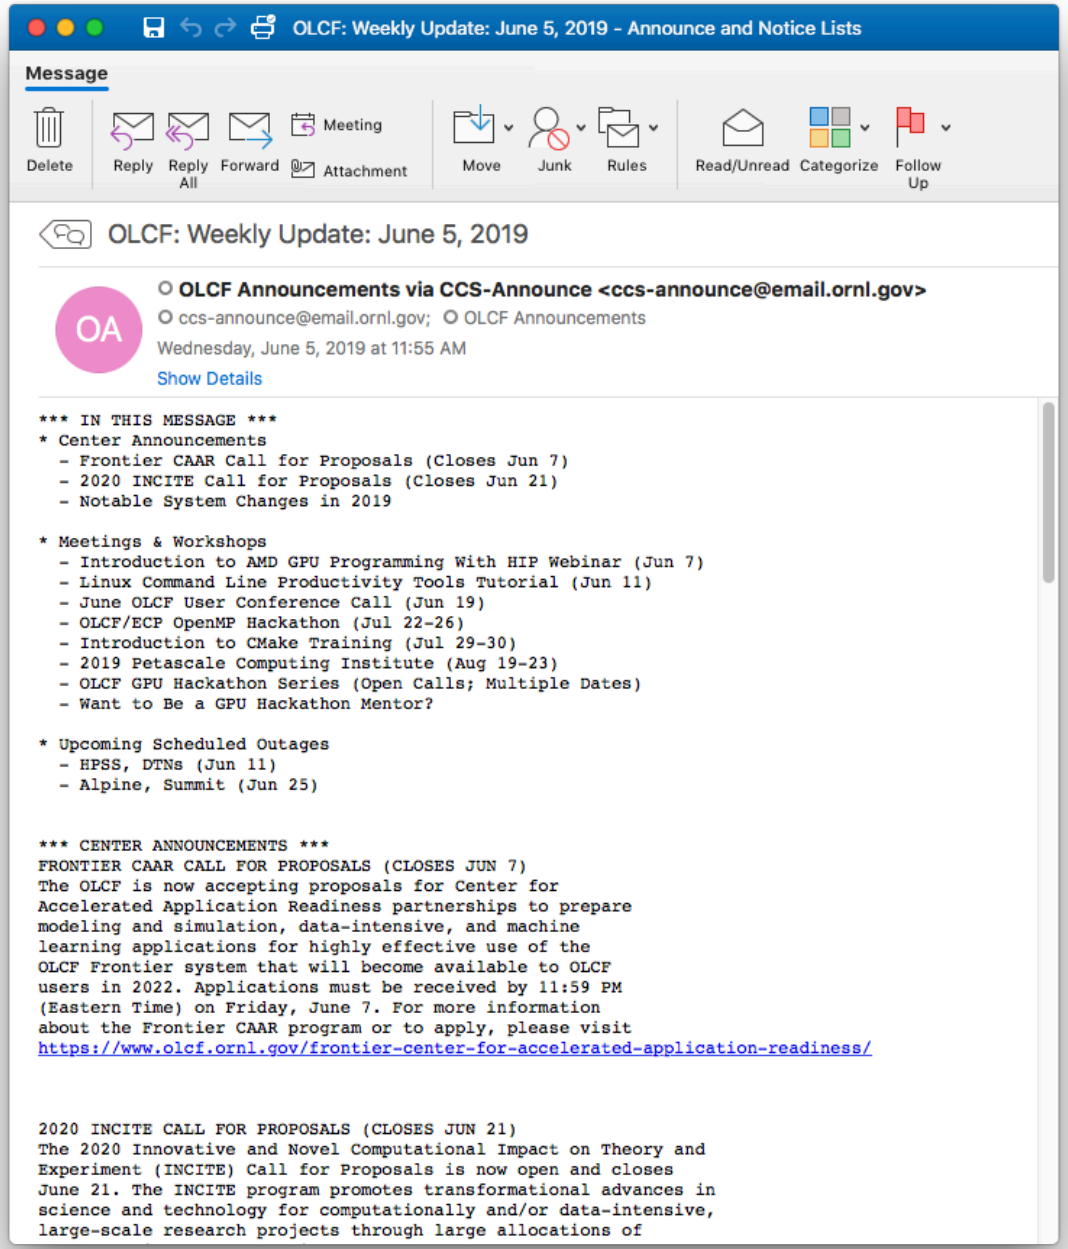

Staying Informed – *System Status*

- Logs from monitoring software are parsed to make educated guess on system status which is sent to multiple destinations
  - OLCF Website
  - Twitter (@OLCFStatus)
  - Appropriate *notice* list
- Fairly accurate, but still a fully automated process
  - Possibility of both false positives and false negatives
  - We do take some measures to mitigate this

Staying Informed – *System Status*

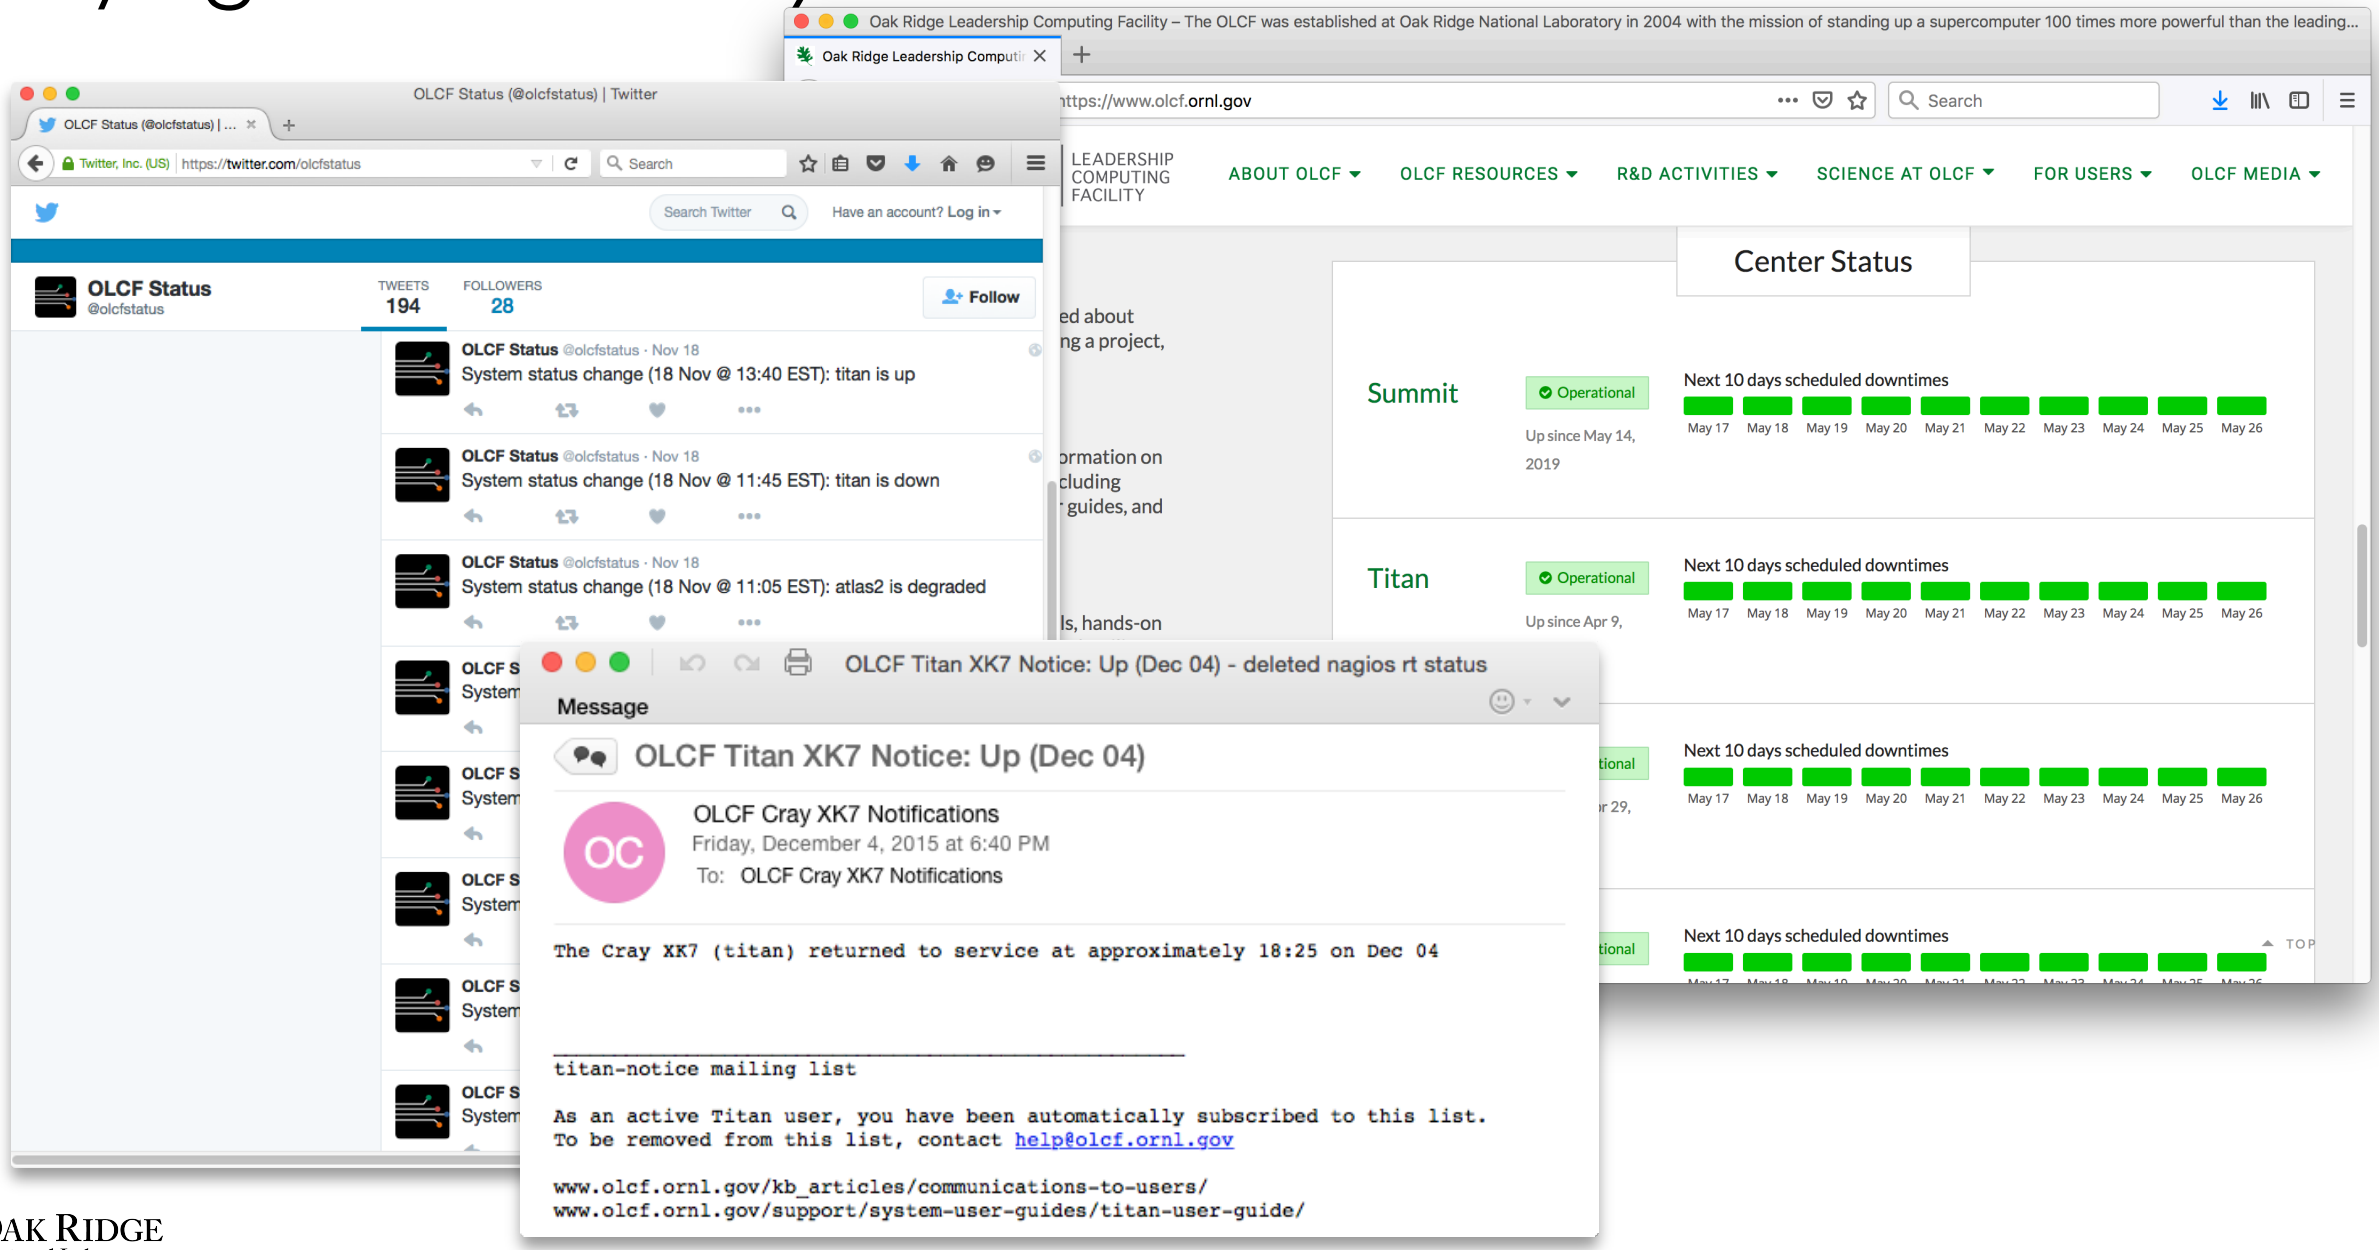

OAK RIDGE National Laboratory

OAK RIDGE National Laboratory

OLCF Policies

OLCF Policies

- Various OLCF policies are available at https://docs.olcf.ornl.gov/accounts/olcf_policy_guide.html
  - Computing Policy
  - Security Policy
  - Data Management Policy
  - Various Project Policies (Reporting, Allocation Utilization, User Agreement)
- The site also includes the acknowledgement statement for publications related to work done on OLCF resources.

OAK RIDGE National Laboratory

Authentication with RSA Tokens

Authenticating to OLCF Systems

- Interactive login access is via Secure Shell (SSH)
- Systems use “two-factor” authentication via user-selected PINs and RSA SecurID tokens
- Other authentication methods (password, public key, etc.) are not permitted

Common Login Issues

- SSH (typically) doesn't prompt for a username
  - By default, it uses your username on the client system
  - You must tell SSH if your username differs
    - Command line: `ssh olcfusername@home.ccs.ornl.gov`
    - Can set this in `~/.ssh/config` file
    - Various ways to do this in graphical SSH clients
  - Unless you request verbose output (`-v[v[v]]`), SSH won't tell you what username it's using

Common Login Issues

- SSH prompts for a password after PASSCODE fails
  - Typically happens after three PASSCODE failures
  - This is a fallback behavior of SSH (the PASSCODE didn't work, so it's trying something else)
  - We don't use passwords (so nothing you type in will permit login); kill the process (`Ctrl-C`) & try again.

```
Enter PASSCODE:
Enter PASSOCDE:
Enter PASSCODE:
user1@summit.olcf.ornl.gov's password:
```

Common Login Issues

- RSA token gets out of sync with the server
  - Sometimes you may be prompted for the 'next tokencode'
    - *Tokencode*: The 6 digit number on your RSA token
    - *PIN*: An alphanumeric string of 4-8 characters known only to you
    - *PASSCODE*: Your PIN followed by the current tokencode
  - When this happens, enter (only) the next tokencode your RSA token generates

```
Enter PASSCODE:
Wait for the tokencode to change, then enter the new tokencode :
```

Additional RSA/SSH Tips

- Once you've used a tokencode, you can't re-use it
- If your PASSCODE has failed twice, let the tokencode change before you try again (to avoid token being locked out)

OAK RIDGE National Laboratory

Projects & User Accounts

System & Group Access

- ***Projects***
  - Are granted system access
  - Are assigned Unix groups
- ***Users***
  - Are assigned to projects
  - "Inherit" the project's groups & system accesses
  - Are ***not*** directly added to systems/groups

Finding Your Project's ID and Allocation

- Use **`showproj`** to list your projects

```
$ showproj

user1 is a member of the following project(s) on summit:
  abc123
```

- Use **`showusage`** to show usage on your project(s)

```
$ showusage

summit usage for the project's current allocation period:
                                  Project Totals          user1
 Project      Allocation        Usage    Remaining        Usage
_____________________________|_____________________________|____________
 abc123          735000    |      11138       7338862 |          12
```

NOTE: *Both commands display a help message if invoked with* **`-h`**

Managing Your Allocation

- Projects are not disabled for going over allocation; they receive priority reduction
  - slight for 100-125% of allocation
  - larger for >125% of allocation (also incurs a limit of 1 running job)
- Since we don't disable projects, we don't give refunds (*per se*)
  - We can delay priority reduction if lots of jobs were affected by a particular system issue (same effect but easier to manage)

Project Closeout

- When your project ends, you'll no longer be able to access OLCF resources associated with that project
  - Even if you're continuing on other projects (and thus retain access to systems like Summit), you won't be able to access the storage areas for the project that ended
- Users will be given a month for data retrieval
  - You won't be able to access the main resources…you will need to use the Data Transfer Nodes

OAK RIDGE National Laboratory

Data Management

Storage Areas

- Divided by scope
  - User
  - Shared among project members
  - Shared among (potentially) all users
- Divided by use
  - Home
  - Scratch
  - Archive

Data Storage Locations

Storage Area	Technology	Purpose of Area
User home	NFS	Frequently accessed user data
Project home	NFS	Frequently accessed project data
User work	Spectrum Scale/GPFS	User's scratch files
Project work	Spectrum Scale/GPFS	Project's scratch files
Global work	Spectrum Scale/GPFS	Sharing data among all users/projects on the system
User archive	HPSS[1]	Long-term storage for user data
Project archive	HPSS[1]	Long-term storage for project data
Global archive	HPSS[1]	Long-term storage for system-wide shared data

[1]HPSS is not mounted as a filesystem; access is discussed in upcoming slides

Data Storage Locations

Storage Area	Location[1]	Default Permissions	Change Permissions?	Backed Up?
User home	`/ccs/home/$USER`	0750	Yes	Yes
Project home	`/ccs/proj/proj_id`	0770	No[3]	Yes
User work[2]	`/gpfs/alpine/proj_id/scratch/$USER`	0700	No[3]	**No**
Project work[2]	`/gpfs/alpine/proj_id/proj-shared`	2770	No[3]	**No**
Global work[2]	`/gpfs/alpine/proj_id/world-shared`	2775	No[3]	**No**
User archive	`/hpss/prod/proj_id/users/$USER`	0700	No[3]	**No**
Project archive	`/hpss/prod/proj_id/proj_shared`	2770	No[3]	**No**
Global archive	`/hpss/prod/proj_id/world_shared`	2775	No[3]	**No**

[1] These are recommended ways to reference directories, not necessarily absolute path names.

[2] Or, `$MEMBERWORK/proj_id`, `$PROJWORK/proj_id`, `$WORLDWORK/proj_id`

[3] Top-level directory permissions "enforce" proper scope of user/project/global directories.

Data Backups

- NFS directories are backed up to a limited extent
  - Accessed via the `.snapshot` subdirectory of an NFS directory
  - The hourly/daily/weekly subdirectories of `.snapshot` are copies of the current directory as of the snapshot time (directory names are the snapshot timestamp)
  - Note that `.snapshot` won't show up in "`ls -al`"
- GPFS ***is not*** backed up (*AND* is purged…see next slide)
- HPSS ***is not*** backed up

Purge Policy

- "Work" directories (i.e. GPFS) are scratch areas and are not intended for long-term storage
- To ensure available space, they are regularly purged of files not accessed/modified in a certain period of time
  - Current purge threshold is in our Data Management Policy
- You should archive data in these areas to a more permanent location as soon as possible

Monitoring Storage Usage

- For home directories, use the `quota` command

```
$ quota -Q
Disk quotas for user user1(uid 98765):
     Filesystem  blocks   quota   limit   grace   files   quota   limit   grace
nccs-svm1.lb.ccs.ornl.gov:/nccs/home1
                 116432  52428800 52428800            1007  4294967295 4294967295
```

- For archive directories, use `showusage`

```
$ showusage -s hpss
HPSS Storage in GB:
                Project Totals      user1
Project         Storage             Storage
_______________|___________________|_____________
user1          |    550.65         |   550.65
abc123         |   2107.18         |   106.52
```

Data Considerations

- OLCF systems can generate large volumes of data very quickly
- Projects should develop a data strategy as soon as possible
  - Consider directory structures, permissions, group, etc.
  - Try to catch issues early (much easier to "fix" 100 files than 10,000)
  - The `chmod` and `umask` commands are vital
- Several things to consider
  - Where project members will store data
  - File ownership/permissions
  - Transferring data to other locations

Transferring Data

- Data transfer nodes (`dtn.ccs.ornl.gov`) are the preferred place for internal & external data transfers
- Several ways to transfer your data
  - External (to OLCF): `bbcp`, `scp`, `gridftp`, globus.org
  - Internal: `hsi/htar` along with those listed above
- Start early/transfer data as it's generated

OAK RIDGE National Laboratory

Using HPSS

The High Performance Storage System

- HPSS is the proper location for long-term storage
  - Not space-constrained like NFS "home" areas
  - Not subject to purge like GPFS
- Directories are project-centric with spaces for user, project, and globally-shared data
- HPSS is accessed via the `hsi` and `htar` commands
- HPSS is also accessible via Globus

HPSS Directory Structure Changes

- In January, moved from `/home/$USER` and `/proj/proj_id` to `/hpss/prod/...` directories
  - `/proj/proj_id` is now symlink to `/hpss/prod/proj_id/proj-shared`
  - `/home/$USER` is a linkfarm to a user's `/hpss/prod/proj_id` directories

  Existing users w/data in `/home/$USER` have not been migrated; they should move data into `/hpss/prod/...` and will be automatically migrated when all files are moved.
- Please begin using these directory structures for HPSS data
- (Existing users: Please migrate files in your HPSS home directory to an appropriate project area)

HPSS Best Practices

- Multiple simultaneous transfers won't necessarily yield benefits (especially on "get" operations)
- Avoid numerous consecutive **`hsi get`** calls (see next slide)
- File size best practices
  - With `htar`, no member file can be ≥ 64GiB (the archive itself can be)
  - For optimal transfer performance, use files ≥ 768GB
  - Minimum recommended file size is 512MB
    - Smaller files will be handled but I/O performance may be negatively affected
    - If you have many small files, try bundling with `htar` to achieve the 512MB threshold

HPSS Best Practices

- **Bad practice:** successive `hsi get` calls

```
$ hsi get /hpss/prod/abc123/users/someuser/file1

$ hsi get /hpss/prod/abc123/users/someuser/file2

$ hsi get /hpss/prod/abc123/users/someuser/file3
```

- **Good practice:** create a list file & call `hsi` once

```
$ cat getfiles.lst
get <<EOF
/hpss/prod/abc123/users/someuser/file1
/hpss/prod/abc123/users/someuser/file2
/hpss/prod/abc123/users/someuser/file3
EOF

$ hsi “in getfiles.lst”
```

The `hsi` transfer agent

- The `hsi` transfer agent offloads `hsi` transfers started on certain nodes to a dedicated set of transfer nodes
  - Reduces load on login nodes
  - Uses nodes optimized for transfer
- To use it, launch `hsi` from
  - Interactive DTNs (`dtn.ccs.ornl.gov`)
  - Batch/Scheduled DTNs (jobs submitted from `dtn.ccs.ornl.gov`)
- Running `hsi` from other nodes is discouraged

More HPSS Information

- Data transfer information is available in OLCF User Guides on the website (links at the end of this presentation)
- For even more HPSS information, see the presentation by OLCF's George Markomanolis
  - https://www.olcf.ornl.gov/wp-content/uploads/2018/12/storage_areas_summit_links.pdf
  - https://vimeo.com/306433952

OAK RIDGE National Laboratory

Finding & Building Software

Finding Software

- Basic commands are part of the default environment
- Other software via Lmod
  - Similar to Environment Modules (`module load`, etc.) but more powerful
  - https://lmod.readthedocs.io/
- Compilers, debugging/optimization tools, other libraries & some scientific apps available

Finding Software

- Special case: python
  - We provide python+anaconda distributions
  - Recommend using those with a conda or venv to build specific modules you need
- Other software
  - You're free to build what you need in your directories, subject to licensing and export control restrictions
  - If you think something might be of general interest, you can request we install it via http://www.olcf.ornl.gov/support/software/software-request or help@olcf.ornl.gov

Building Software

- Considering building/installing software in an NFS area
  - Avoids the scratch area purge
  - Also can be "friendlier" to the system, since compiling involves lots of metadata which can impact/be impacted by large parallel filesystems
- Might be possible to build in a ramdisk (such as `/tmp`) with an NFS directory as the installation prefix
  - Be careful with `/tmp` due to permissions

OAK RIDGE National Laboratory

Using Summit

Summit Considerations

- Running Jobs
  - Summit uses LSF for the batch queue system
  - Summit uses `jsrun` as its parallel launcher
    - Different approach to thinking about resources
    - OLCF provides tools to help (be sure to join the March user call for more info!)
- Processor Architecture
  - POWER9™
  - Not compatible with x86_64 binaries, so you may have issues with precompiled software (*LOTS* of stuff targets x86_64)

Summit Training Resources

- https://github.com/olcf-tutorials
- https://www.olcf.ornl.gov/for-users/training/training-archive/, in particular:
  - https://www.olcf.ornl.gov/calendar/introduction-to-nvidia-profilers-on-summit/
  - https://www.olcf.ornl.gov/calendar/summit-training-workshop/
  - https://www.olcf.ornl.gov/calendar/introduction-to-summit-webinar/
  - https://www.olcf.ornl.gov/calendar/jsrun-tutorial/

OAK RIDGE National Laboratory

Training Opportunities

Training Opportunities

- We host numerous training events through the year
  - Monthly User Conference Call
  - Software-specific courses
  - GPU Hackathons https://gpuhackathons.org/
- Most are in-person and are webcast with BlueJeans or Webex

Training Opportunities

- What training is coming up?
  - Watch for announcements in the Weekly Update
  - Check the OLCF Training Calendar https://www.olcf.ornl.gov/for-users/training/training-calendar/
- Videos of past training events are on the OLCF Vimeo site https://vimeo.com/channels/olcftraining

OAK RIDGE National Laboratory

Getting Help

Where do I find documentation?

- OLCF Website/System User Guides
  https://docs.olcf.ornl.gov/
- NVIDIA hosted documentation
  http://docs.nvidia.com

Working With User Support

- Email is often the best option to contact us
  - Especially for sending long/complicated error messages
  - Send as many error messages as possible
    (or place them in a file & direct us to the file)
- "Send" us codes by creating a `.tar` file & directing us to it
- Report new issues in new tickets (not a reply to an old ticket)
  - Helps us in classifying/searching through old tickets
  - Gives it greater visibility

Requesting Policy Exemptions

- The Resource Utilization Council accepts requests for temporary exemption from some policies (various job limits, purge exemptions, quota increases, etc.)
- Request can be made from our Documents & Forms page (in the third section of the page) https://docs.olcf.ornl.gov/accounts/documents_and_forms.html
  - Reviewed by RUC; make requests well in advance to allow for review
  - If requesting job priority, make sure you submit the job…they often run more quickly than you expect

OAK RIDGE National Laboratory

Finally

Finally…

- We're here to help you
- Questions/comments/etc. can be sent to the OLCF User Assistance Center
  - Staffed 9AM – 5PM US Eastern Time (exclusive of ORNL holidays)
  - help@olcf.ornl.gov
  - (865) 241-6536

List of Links

- General
  - https://www.olcf.ornl.gov
  - https://docs.olcf.ornl.gov/accounts/olcf_policy_guide.html
  - http://www.olcf.ornl.gov/support/software/software-request
  - https://docs.olcf.ornl.gov/accounts/documents_and_forms.html
- Documentation
  - https://docs.olcf.ornl.gov
  - http://docs.nvidia.com
  - https://lmod.readthedocs.io

List of Links

- Training
  - https://vimeo.com/channels/olcftraining
  - https://github.com/olcf-tutorials
  - https://www.olcf.ornl.gov/for-users/training/training-archive/
  - https://www.olcf.ornl.gov/calendar/introduction-to-nvidia-profilers-on-summit/
  - https://www.olcf.ornl.gov/calendar/summit-training-workshop/
  - https://www.olcf.ornl.gov/calendar/introduction-to-summit-webinar/
  - https://www.olcf.ornl.gov/for-users/training/training-calendar/
  - https://www.olcf.ornl.gov/calendar/jsrun-tutorial/

List of Links

- Training (Continued)
  - https://www.olcf.ornl.gov/wp-content/uploads/2018/12/storage_areas_summit_links.pdf
  - https://vimeo.com/306433952 (Summit storage/data transfers)
  - https://gpuhackathons.org/