OAK RIDGE National Laboratory

OLCF Best Practices and Overview for New Users

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 resources & covers systems, policies, and where to find help
  - This is by no means an all-inclusive presentation
  - Feel free to ask questions
- Much of this may be a reminder; it's largely an update to previous presentations

https://vimeo.com/channels/olcftraining/254494153
https://vimeo.com/channels/olcftraining/343636411
https://vimeo.com/channels/olcftraining/405885960
https://vimeo.com/channels/olcftraining/427792537

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* list
  - All users are required to be members
  - System-specific & center-wide announcements
  - olcf@olcf-communications.org
    (We use Mailchimp for this…It's an external address, but yes, it's us!)
  - 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)
  - 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*

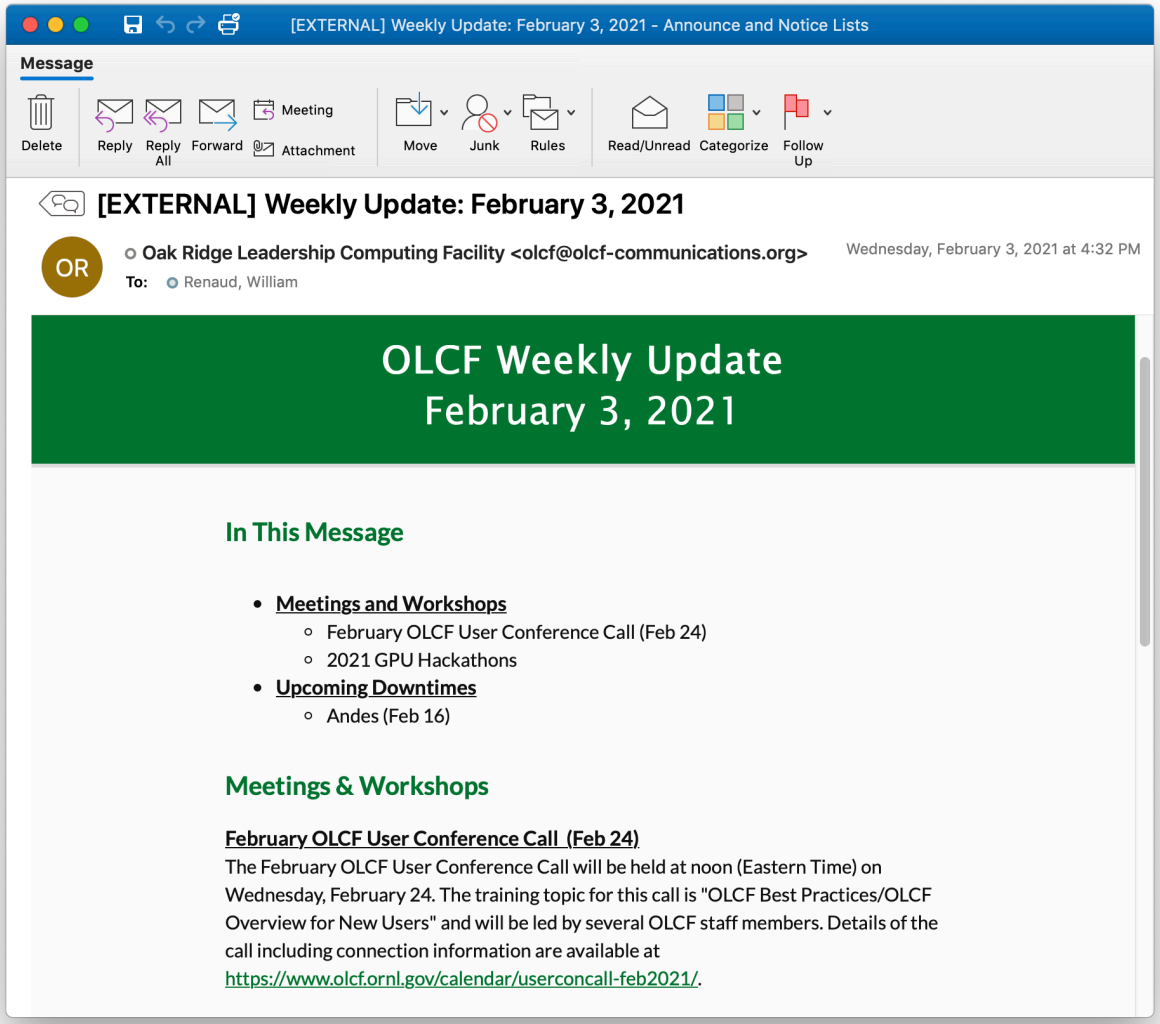

[EXTERNAL] Weekly Update: February 3, 2021 - Announce and Notice Lists

Message

Delete | Reply | Reply All | Forward | Meeting | Attachment | Move | Junk | Rules | Read/Unread | Categorize | Follow Up

[EXTERNAL] Weekly Update: February 3, 2021

OR — Oak Ridge Leadership Computing Facility <olcf@olcf-communications.org> — Wednesday, February 3, 2021 at 4:32 PM

To: Renaud, William

OLCF Weekly Update
February 3, 2021

In This Message

- Meetings and Workshops
  - February OLCF User Conference Call (Feb 24)
  - 2021 GPU Hackathons
- Upcoming Downtimes
  - Andes (Feb 16)

Meetings & Workshops

February OLCF User Conference Call (Feb 24)
The February OLCF User Conference Call will be held at noon (Eastern Time) on Wednesday, February 24. The training topic for this call is "OLCF Best Practices/OLCF Overview for New Users" and will be led by several OLCF staff members. Details of the call including connection information are available at https://www.olcf.ornl.gov/calendar/userconcall-feb2021/.

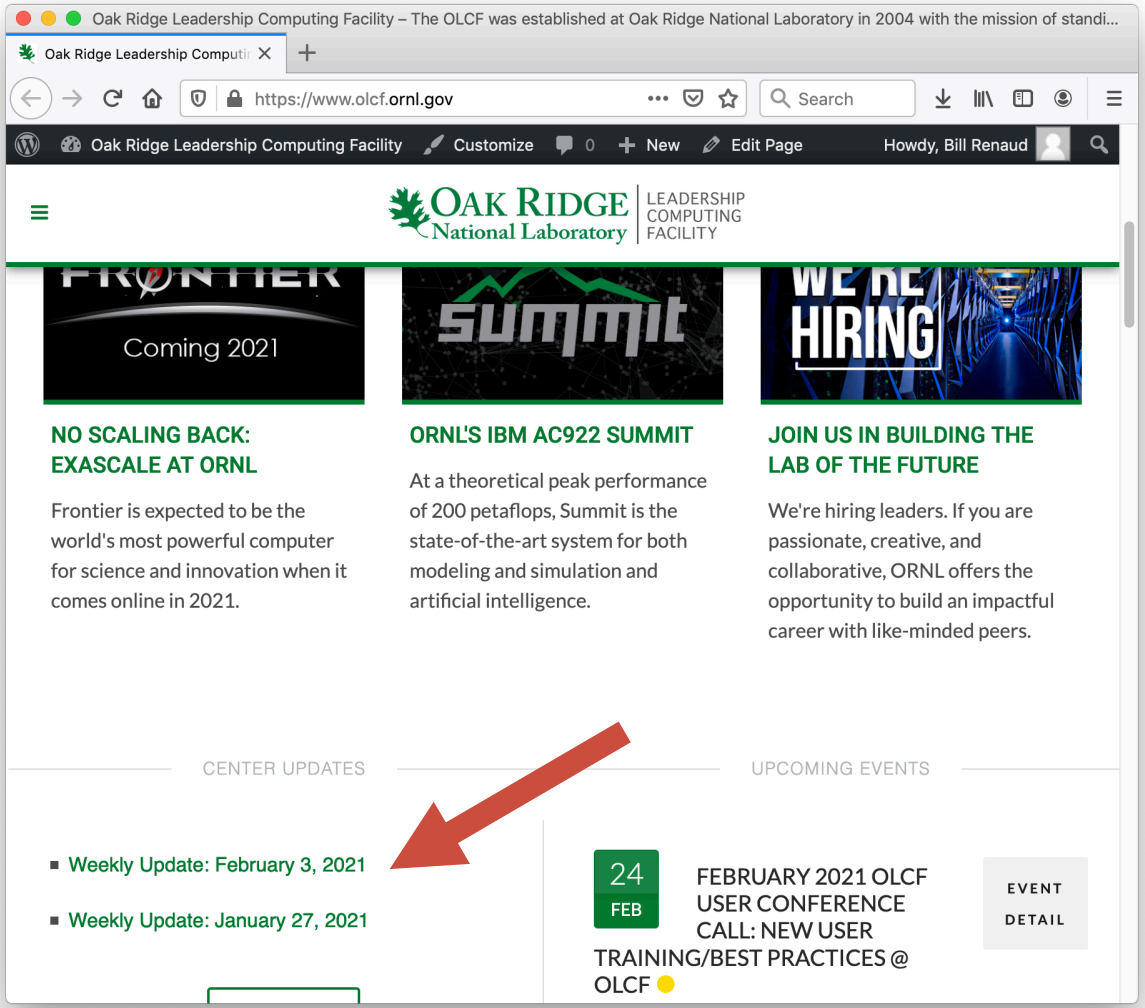

Oak Ridge Leadership Computing Facility – The OLCF was established at Oak Ridge National Laboratory in 2004 with the mission of standi...

Oak Ridge Leadership Computi

https://www.olcf.ornl.gov

Search

Oak Ridge Leadership Computing Facility | Customize | 0 | New | Edit Page | Howdy, Bill Renaud

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

NO SCALING BACK: EXASCALE AT ORNL

Frontier is expected to be the world's most powerful computer for science and innovation when it comes online in 2021.

ORNL'S IBM AC922 SUMMIT

At a theoretical peak performance of 200 petaflops, Summit is the state-of-the-art system for both modeling and simulation and artificial intelligence.

JOIN US IN BUILDING THE LAB OF THE FUTURE

We're hiring leaders. If you are passionate, creative, and collaborative, ORNL offers the opportunity to build an impactful career with like-minded peers.

CENTER UPDATES

- Weekly Update: February 3, 2021
- Weekly Update: January 27, 2021

UPCOMING EVENTS

24 FEB — FEBRUARY 2021 OLCF USER CONFERENCE CALL: NEW USER TRAINING/BEST PRACTICES @ OLCF — EVENT DETAIL

OAK RIDGE National Laboratory

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*

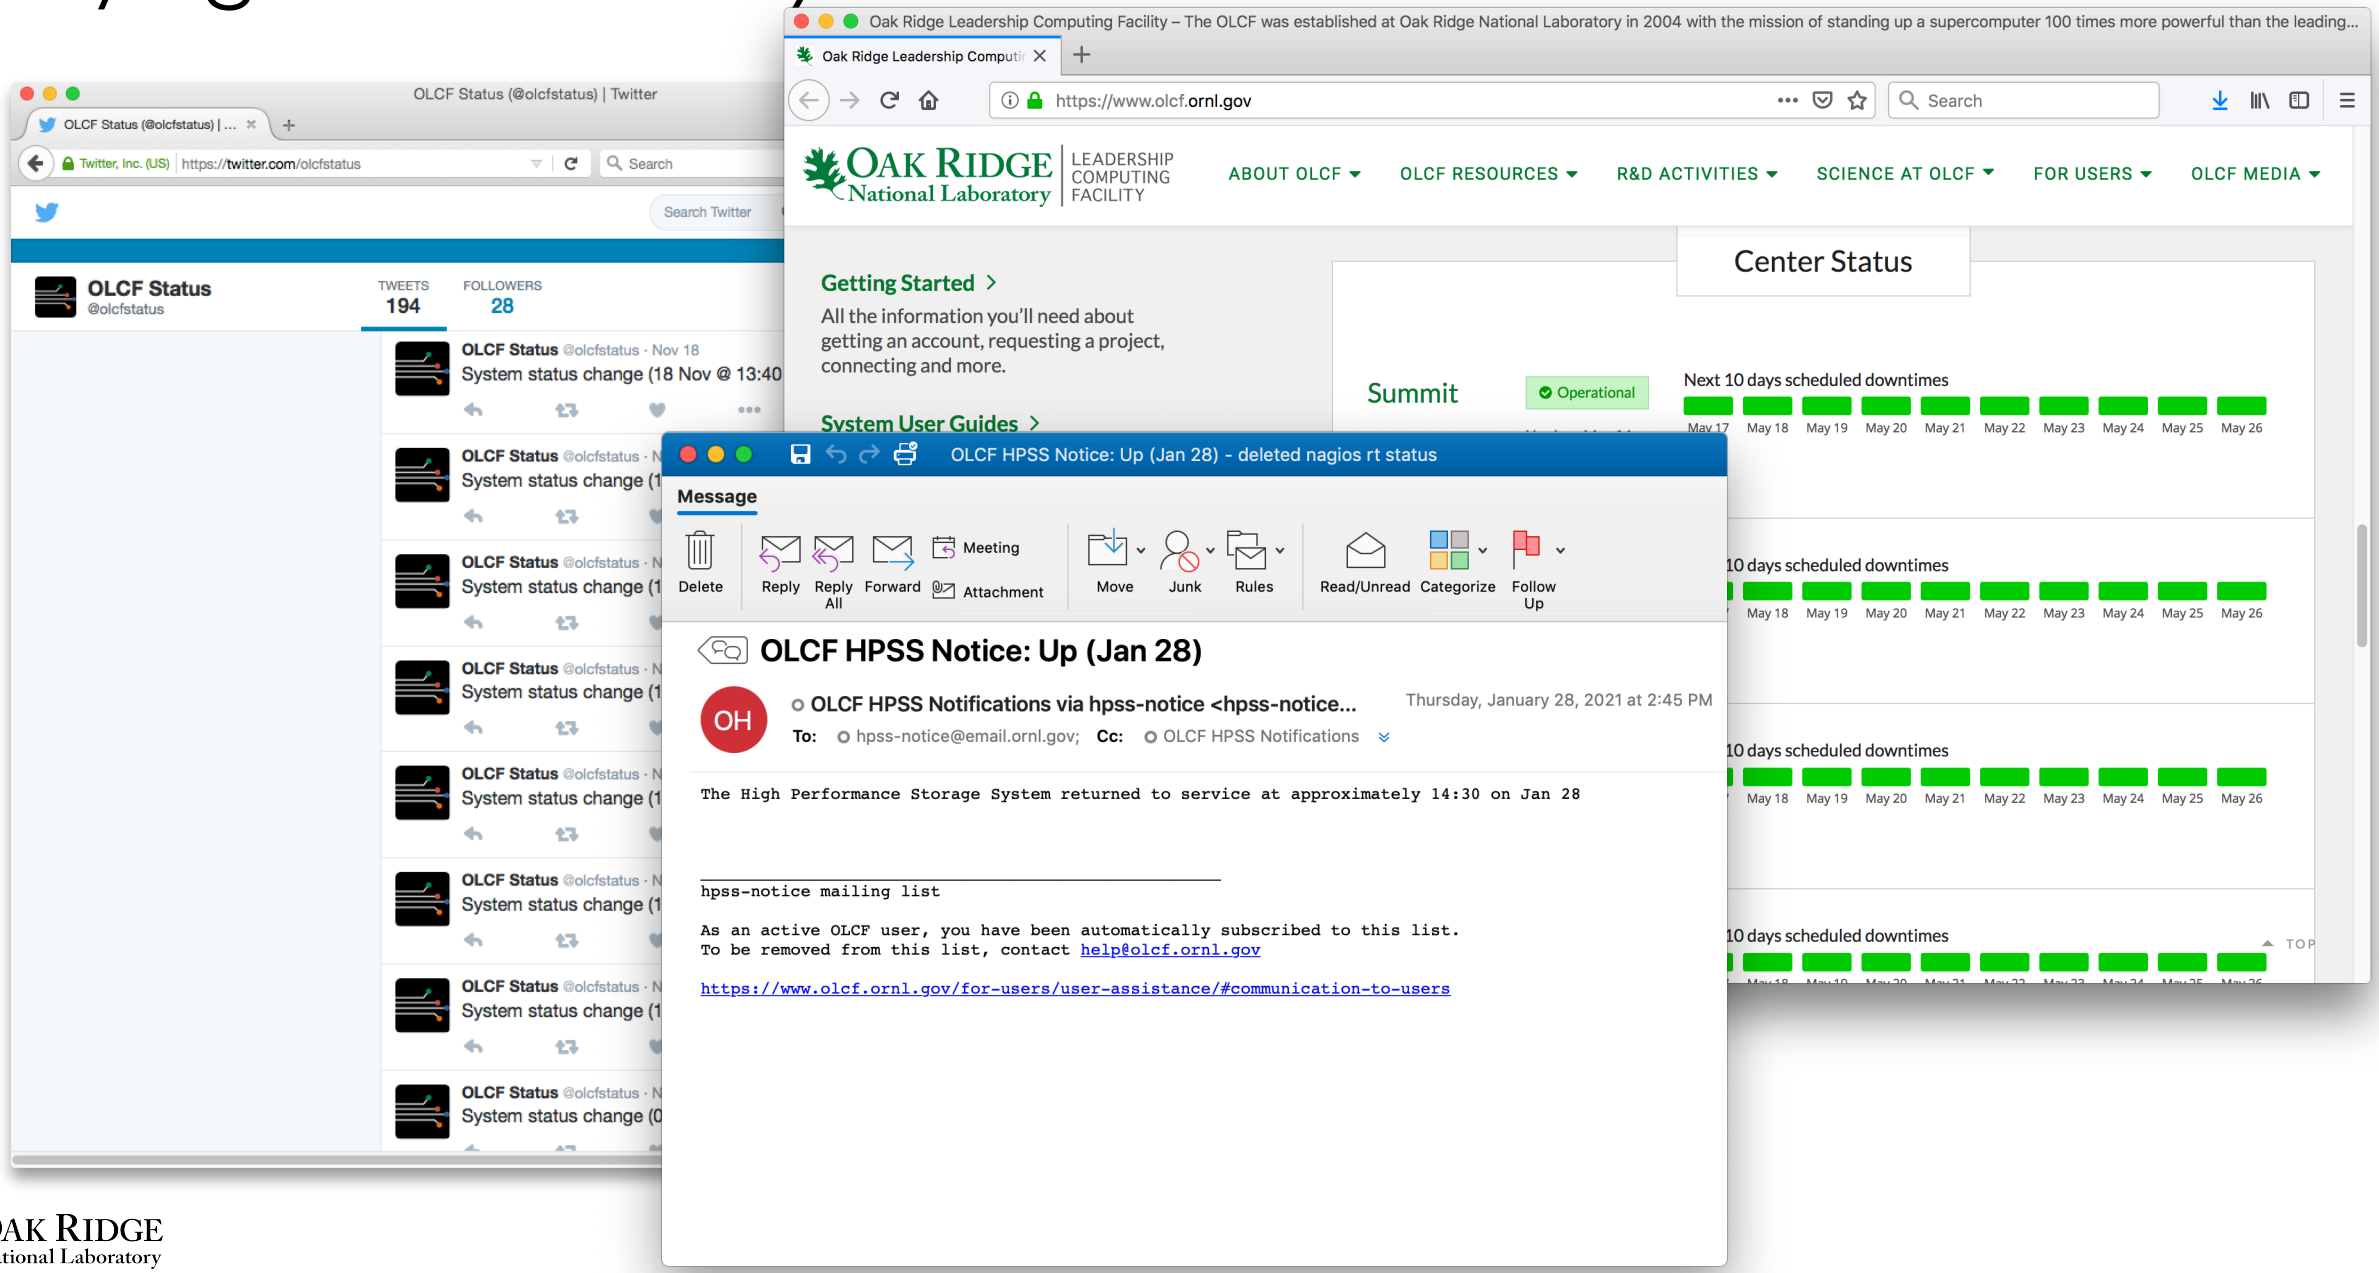

OAK RIDGE National Laboratory

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

So, to get access to a new group or system, you need to join a project that has been granted the desired access.

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
  - System access & unix groups are disabled on systems
  - Even if you can still log in via another project, you can't access the old project's directories (since groups are removed)
- Projects are given an extra month for data retrieval
  - But only on `home.ccs.ornl.gov` and `dtn.ccs.ornl.gov`
  - Don't wait to move data, though…

OAK RIDGE National Laboratory

Data Management

Storage Areas

By Scope	By Intended Use
• User	• Semi-Permanent/Home Directory
• Share among project members	• Scratch
• Shared with (possibly) all users	• 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

- Backed up to a limited extent via the `.snapshot` subdirectory
  - Subdirectories are copies of the directory as of the snapshot time
  - Note that `.snapshot` won't show up in "`ls -al`"

```
$ ls –al /ccs/proj/abc123/important_data |grep .snapshot

$ ls /ccs/proj/abc123/important_data/.snapshot
daily.2021-01-29_0010   hourly.2021-02-03_1905  hourly.2021-02-04_0605
daily.2021-01-30_0010   hourly.2021-02-03_2005  hourly.2021-02-04_0705
daily.2021-01-31_0010   hourly.2021-02-03_2105  hourly.2021-02-04_0805
daily.2021-02-01_0010   hourly.2021-02-03_2205  hourly.2021-02-04_0905
daily.2021-02-02_0010   hourly.2021-02-03_2305  hourly.2021-02-04_1005
daily.2021-02-03_0010   hourly.2021-02-04_0005  hourly.2021-02-04_1105
daily.2021-02-04_0010   hourly.2021-02-04_0105  hourly.2021-02-04_1205
hourly.2021-02-03_1505  hourly.2021-02-04_0205  hourly.2021-02-04_1305
hourly.2021-02-03_1605  hourly.2021-02-04_0305  hourly.2021-02-04_1405
hourly.2021-02-03_1705  hourly.2021-02-04_0405  weekly.2021-01-31_0015
hourly.2021-02-03_1805  hourly.2021-02-04_0505
```

Data Backups – GPFS and HPSS

- GPFS ***is not*** backed up (*AND* is purged)
  - GPFS directories are scratch areas (not intended for long-term storage)
  - To ensure available space, files not accessed/modified recently are purged (see Data Management Policy for the purge threshold)
  - *Archive data in GPFS to a more permanent location as soon as possible after creation!*
- HPSS ***is not*** backed up

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
```

- No good tool for GPFS at present…ask help@olcf.ornl.gov if you need to know.

- 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 generate lots of data very quickly; projects should develop a data strategy ***as soon as possible***. (It's easier to fix things w/100 files than thousands of files/directories!)
- Several things to consider
  - How are files/directories shared among project members?
    - Where will project members store data?
    - What file attributes (permissions, group, etc.) are needed?
  - What happens when someone leaves the project?
  - What happens when the project ends?
    - Where does the data need to go?
    - How much data is there, who's moving it, and how long will it take?

Whose Data Is It?

- Data in the User Home is considered the user's data
  - Normally, user must approve transfer to others
  - Anything that's part of the project needs to be elsewhere
- Data in other directories is considered project data
  - *This includes User Scratch & User Archive areas*…yes, they're for things the user doesn't want to share, but they still are part of the project's data
  - PI approval is sufficient to change permissions/move this data

Transferring Data

- https://docs.olcf.ornl.gov/data/transferring.html
- Data transfer nodes (`dtn.ccs.ornl.gov`) are the preferred place for internal & external data transfers
- Several ways to transfer your data depending on your requirements, including `hsi/htar`, `scp/sftp` & Globus
- Start early/transfer data as it's generated

OAK RIDGE National Laboratory

Using HPSS

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
- Accessing HPSS:
  - From OLCF systems: via the `hsi` and `htar` commands
  - Externally: via Globus

HPSS Directory Structure Changes - Overview

- Old structure (prior to last year) had two main storage areas
  - User home directories: `/home/$USER`
  - Project directories: `/proj/proj_id`
- New structure mimics GPFS
  - Project based: `/hpss/prod/proj_id`
  - Each project directory has `users`, `proj-shared`, and `world-shared` subdirectories; each project member has a `/hpss/prod/proj_id/users/$USER` directory

HPSS Directory Structure Changes - Migrating

- Project directories (`/proj/proj_id`) have been migrated
  - Moved to `/hpss/prod/proj_id/proj-shared`
  - New `/proj/proj_id` set up as link to `/hpss/proj/proj_id/proj-shared`
- User directories (`/home/$USER`) are (or will be) link farm to `/hpss/prod/proj_id` for each of your projects
  - Already done for new users/existing users with no data in the directory
  - Will be done automatically for existing users w/data there once data have been moved/directory is empty
- *Please begin using the new directories & if you have data in `/home`, please migrate to an appropriate project directory ASAP*

HPSS Best Practices

- Multiple simultaneous transfers won't necessarily yield benefits
- Avoid numerous consecutive **`hsi get`** calls (see next slide)
- File size guidelines
  - No `htar` member file can be ≥ 64GiB (the archive itself can be)
  - For optimal transfer performance, use files ≥ 768GiB
  - Minimum recommended file size is 512MiB
  - Bundle lots of small files into one (or more) files
    - Via `tar` then **`hsi put`**
    - Via `htar`

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:
  - DTNs (interactively or via batch job on `dtn.ccs.ornl.gov`)
  - Summit/Andes Login nodes; Summit batch nodes

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 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
  - You may need to build from source for Summit
- Other software
  - If you think something might be of general interest, you can request via email to 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 with an NFS directory as the installation prefix
  - Be sure to set appropriate permissions
- Review may be needed prior to adding software that wasn't listed on your project application

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 (see https://jobstepviewer.olcf.ornl.gov/)
- Processor Architecture
  - POWER9™
  - Not compatible with x86_64 binaries, so you may have issues with precompiled software (including python modules via conda as noted earlier)

Job Step Viewer

- The `jsrun` command can be tricky
  - It's a different approach than other parallel launchers
  - A full description is well beyond the scope of this talk
- We provide a tool, Job Step Viewer, that runs a `jsrun` command & displays job placement on nodes via the web
  - To use it, simply load the module (`module load job-step-viewer`) and run your command. It'll echo the URL that shows job/task placement
- For more info, see slide presentation on https://www.olcf.ornl.gov/calendar/userconcall-mar2020/

Summit Training Resources

- https://github.com/olcf-tutorials
- https://docs.olcf.ornl.gov/training/training_archive.html, 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

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)
- It's not necessary to send us codes/reproducers via email
  - More efficient to tell us where they are on the filesystem
  - If appropriate, set up a test directory or collect everything in a .tar file
- 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 exceptions via email to help@olcf.ornl.gov.
  - These are reviewed by the RUC, so make them well in advance
  - Let us know why the existing limits are an issue
  - If requesting a job priority boost…submit the job! It may 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
- Computer operations staff available 24x7 for limited troubleshooting via phone at (865) 241-6536

OAK RIDGE National Laboratory

Links

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
  - https://docs.olcf.ornl.gov/data/transferring.html

List of Links

- Training
  - https://vimeo.com/channels/olcftraining
  - https://www.olcf.ornl.gov/calendar/userconcall-mar2020/
  - https://github.com/olcf-tutorials
  - https://docs.olcf.ornl.gov/training/training_archive.html
  - 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://jobstepviewer.olcf.ornl.gov/
  - 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/