

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

- This presentation gives 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>

<https://vimeo.com/channels/olcftraining/519216250>

General Information

- Video/Slides will be uploaded to the event page (hopefully by the end of the week)
- All URLs are summarized at the end of the presentation

Overview

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

Staying Informed



Staying Informed

- We 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 16, 2022 - Inbox

Message

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

[EXTERNAL] Weekly Update: February 16, 2022

OR Oak Ridge Leadership Computing Facility <olcf@olcf-communications.org> Today at 3:03 PM

To: Renaud, William

To protect your privacy, some pictures in this message were not downloaded. [Download pictures](#)

[View this email in your browser](#)

OLCF Weekly Update February 16, 2022

In This Message

- Did You Know
- Center Announcements
 - Machine Learning Survey (Open Through Feb 25)
- Meetings and Workshops
 - February OLCF User Conference Call: New User Training (Feb 23)
 - An Introduction to Programming with SYCL (Mar 1)
- Upcoming Downtimes

Oak Ridge Leadership Computing Facility

https://www.olcf.ornl.gov

NO SCALING BACK: EXASCALE AT ORNL

Frontier was delivered in 2021 and is being integrated. It's expected to be the world's most powerful computer for science and innovation when it comes online this year.

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.

FEBRUARY 23, NEW USER TRAINING

This month's user call will feature a hands-on training opportunity, providing tips and best practices for both new and existing users to more effectively utilize OLCF resources.

CENTER UPDATES

- Weekly Update: February 16, 2022
- Weekly Update: February 9, 2022

[MORE UPDATES](#)

UPCOMING EVENTS

23 FEB	FEBRUARY 2022 OLCF USER CONFERENCE CALL: NEW USER TRAINING/BEST PRACTICES	EVENT DETAIL
01 MAR	AN INTRODUCTION TO PROGRAMMING WITH	EVENT DETAIL

Wednesday

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*

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.

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
 - Tokencodes can only be used once
 - Other methods (password, public key, etc.) are not permitted
- RSA terminology
 - *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

Common Login Issues: Incorrect Username

- SSH (typically) doesn't prompt for a username
 - Nor does it tell you which one it's using
 - Unless you request verbose output (`-v[v[v]]`)
 - By default, it uses your username on the client system; you must tell it if your username differs
 - Command line: `ssh olcfusername@home.ccs.ornl.gov`
 - Via settings in `~/.ssh/config` file
 - Various ways to do this in graphical SSH clients

Common Login Issues: Password Prompt

- SSH prompts for a password after PASSCODE fails
 - Typically happens after three PASSCODE failures...this is a fallback behavior of SSH (PASSCODE didn't work, so it's trying something else)
 - Your only option here is to kill the process (`ctrl-c`) & try again.
 - *Tip: If your PASSCODE has failed twice, let the tokencode change before you try again (to avoid token being locked out)*

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

Common Login Issues: Token Synchronization

- RSA token gets out of sync with the server
 - Sometimes you may be prompted for the 'next tokencode'

Enter PASSCODE:

Wait for the tokencode to change, then enter the new tokencode :

- If this happens, enter (only) the next tokencode your RSA token generates

Projects & User Accounts



System & Group Access

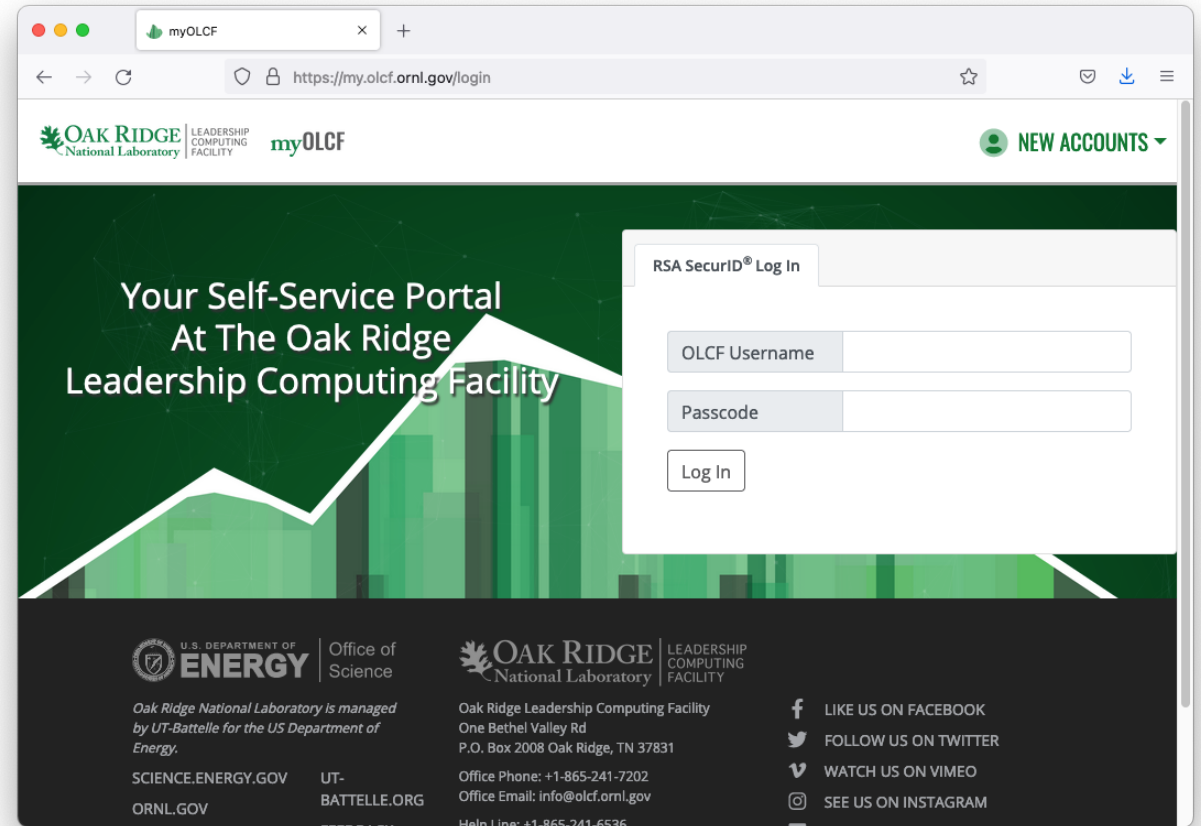
- Projects
 - Are granted system access/assigned to Unix groups
- Users
 - Are assigned to projects
 - Inherit the project's groups & system accesses (they're not directly added to systems/groups)
- So...
 - For access to a group/system, you must join a corresponding project
 - If your project ends, your system access ends (unless you're on another project w/access to that system)

Managing Your Allocation

- Projects are not disabled for going over allocation; they keep running, albeit at reduced priority
 - 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)
 - If you notice jobs being impacted by system issues, try to keep track of affected jobs

myOLCF

- myOLCF allows you to
 - apply for projects (initial/renewal)
 - view your projects
 - get usage information
 - update your contact information
 - And more!
- <https://my.olcf.ornl.gov>



https://docs.olcf.ornl.gov/services_and_applications/myolcf/index.html

Project Closeout

- When your project ends:
 - You'll no longer be able to access OLCF resources associated with that project *unless* you have access via another project
 - You'll lose access to that project's Unix groups (and therefore the project's directories)
 - You'll have 1 extra month for data retrieval
 - But only on `home.ccs.ornl.gov` and `dtn.ccs.ornl.gov`
 - Don't wait to move data, though...

Data Management



Storage Areas

By Scope

- User
- Share among project members
- Shared with (possibly) all users

By Intended Use

- Semi-Permanent/Home Directory
- 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

¹HPSS 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

- NFS areas are 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.2022-02-10_0010    hourly.2022-02-15_1605    hourly.2022-02-16_0005    hourly.2022-02-16_0805
daily.2022-02-11_0010    hourly.2022-02-15_1705    hourly.2022-02-16_0105    hourly.2022-02-16_0905
daily.2022-02-12_0010    hourly.2022-02-15_1805    hourly.2022-02-16_0205    hourly.2022-02-16_1005
daily.2022-02-13_0010    hourly.2022-02-15_1905    hourly.2022-02-16_0305    hourly.2022-02-16_1105
daily.2022-02-14_0010    hourly.2022-02-15_2005    hourly.2022-02-16_0405    hourly.2022-02-16_1205
daily.2022-02-15_0010    hourly.2022-02-15_2105    hourly.2022-02-16_0505    hourly.2022-02-16_1305
daily.2022-02-16_0010    hourly.2022-02-15_2205    hourly.2022-02-16_0605    hourly.2022-02-16_1405
hourly.2022-02-15_1505    hourly.2022-02-15_2305    hourly.2022-02-16_0705    weekly.2022-02-13_0015
```

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      Project Totals      user1
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 with 100 files than with 100,000!)
- Some 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
 - *Including 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 (`dtm.ccs.ornl.gov`) are the preferred place for internal & external data transfers
- Command-line transfer available with `hsi/htar` & `scp/sftp`
- Globus transfers available via two endpoints
 - OLCF DTN: for NFS and GPFS
 - OLCF HPSS: for HPSS
- Start early/transfer data as it's generated

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
 - Preferably from `dtm.ccs.ornl.gov`; `hsi` from Summit/Andes logins should be OK
 - Externally: via Globus ("OLCF HPSS" endpoint)

HPSS Directory Structure Changes - Overview

- Old structure had `/home/$USER` and `/proj/proj_id` storage areas
- New structure is project-based and mimics GPFS
 - `/hpss/prod/proj_id/[users|proj-shared|world-shared]`
(old `/proj/proj_id` areas point to `/hpss/proj/proj_id/proj-shared`)
 - New users: Home directories are a link farm to proj areas
 - Existing users: Old home areas remain; users should migrate files to appropriate project areas & system will then convert to link farm)

More HPSS Information

- Great presentation from Gregg Gawinski last month:
<https://www.olcf.ornl.gov/calendar/userconcall-jan2022/>
- 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>

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 + conda/venv to build modules you need
 - You may need to build from source for Summit
 - Good presentation from OLCF's Matt Belhorn:
 - https://www.olcf.ornl.gov/wp-content/uploads/2020/06/20200603_summit_workshop_python.pdf
 - <https://vimeo.com/427794043>
- Other software
 - If you think something might be of general interest, you can request via email to help@olcf.ornl.gov

Building Software

- Consider 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; you must also comply with licensing/export control requirements

Using Summit



Summit Considerations

- Running Jobs
 - Summit uses LSF for the batch queue system
 - Summit uses `jsrun` as its parallel launcher
- Processor Architecture
 - POWER9™
 - **Not compatible with x86_64 binaries!** You may have issues with prebuilt software (including python modules via conda as noted earlier)
- Filesystems
 - NFS areas are read-only on compute nodes
 - NVMe areas only available if you submit with `#BSUB -alloc_flags "nvme"`

Job Step Viewer

- The `jsrun` command can be tricky...it's a different approach than other parallel launchers
- We provide a tool, Job Step Viewer, that runs a `jsrun` command & displays job placement on nodes via the web
 - <https://jobstepviewer.olcf.ornl.gov/>)
 - 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
 - Must be run from within a batch job
- 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/>

Training Opportunities



Training Opportunities

- We support numerous training events: concalls, software-specific courses, hackathons (<https://gpuhackathons.org>)
- Recent training has been virtual; we anticipate hybrid training for future events
- What training is coming up?
 - Watch for announcements in the Weekly Update
 - <https://www.olcf.ornl.gov/for-users/training/training-calendar/>
- Past training:
 - <https://vimeo.com/channels/olcftraining>
 - https://docs.olcf.ornl.gov/training/training_archive.html

Getting Help



Where do I find documentation?

- OLCF Website/User Documentation: <https://docs.olcf.ornl.gov/>
 - Summit/Andes User Guides
 - Slate/Jupyter documentation
 - Data storage/transfer
 - Training archive
- 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.

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

Links



Links

- General

- <https://www.olcf.ornl.gov>
- <https://www.olcf.ornl.gov/for-users/user-assistance/>
- https://docs.olcf.ornl.gov/accounts/olcf_policy_guide.html
- <https://my.olcf.ornl.gov>
- https://docs.olcf.ornl.gov/services_and_applications/myolcf/index.html

- Documentation

- <https://docs.olcf.ornl.gov/>
- <http://docs.nvidia.com>
- <https://lmod.readthedocs.io/>

Links

- Data Transfer

- <https://docs.olcf.ornl.gov/data/transferring.html>
- <https://www.olcf.ornl.gov/calendar/userconcall-jan2022/>
- https://www.olcf.ornl.gov/wp-content/uploads/2018/12/storage_areas_summit_links.pdf
- <https://vimeo.com/306433952>

Links

- Summit Topics

- <https://jobstepviewer.olcf.ornl.gov/>
- <https://www.olcf.ornl.gov/calendar/userconcall-mar2020/>
- <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/>
- https://www.olcf.ornl.gov/wp-content/uploads/2020/06/20200603_summit_workshop_python.pdf
- <https://vimeo.com/427794043>

Links

- Training

- <https://gpuhackathons.org>
- <https://github.com/olcf-tutorials>
- <https://www.olcf.ornl.gov/for-users/training/training-calendar/>
- <https://vimeo.com/channels/olcftraining>
- https://docs.olcf.ornl.gov/training/training_archive.html
- Previous Best Practices
 - <https://vimeo.com/channels/olcftraining/254494153>
 - <https://vimeo.com/channels/olcftraining/343636411>
 - <https://vimeo.com/channels/olcftraining/405885960>
 - <https://vimeo.com/channels/olcftraining/427792537>
 - <https://vimeo.com/channels/olcftraining/519216250>