

OLCF Overview for New Users

Bill Renaud
OLCF User Support

ORNL is managed by UT-Battelle, LLC for the US 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 presentation:
<https://vimeo.com/channels/olcftraining/254494153>
 - Titan/Eos/Lustre info has been removed
 - Some new info for Summit

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

Oak Ridge Leadership Computing Facility – The OLCF was established at Oak Ridge National Laboratory in 2004 with the mission of standing up a supercomputer 100 times...

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


OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

ABOUT OLCF ▾ OLCF RESOURCES ▾ R&D ACTIVITIES ▾ SCIENCE AT OLCF ▾ FOR USERS ▾ OLCF MEDIA ▾




2020 INCITE CALL FOR PROPOSALS

Open from April 15 to June 21, 2019, INCITE's call provides an opportunity for researchers to pursue supercomputer allocations at ALCF and OLCF.



NO SCALING BACK: EXASCALE AT ORNL

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



SYSTEM DECOMMISSIONS AND CHANGES

As the OLCF begins preparations for Frontier, be aware of plans that will impact Titan, Eos, Rhea, and the Atlas filesystem in 2019.

CENTER UPDATES

- Weekly Update: June 5, 2019
- Weekly Update: May 29, 2019

MORE UPDATES

UPCOMING EVENTS

| | | |
|--------|---|-----------------|
| 07 JUN | INTRODUCTION TO AMD GPU PROGRAMMING WITH HIP ● | EVENT DETAIL |
| 11 JUN | LINUX COMMAND LINE PRODUCTIVITY TOOLS ● | EVENT DETAIL |
| 19 JUN | JUNE 2019 OLCF USER CONFERENCE CALL-OLCF BEST PRACTICES ● | EVENT DETAIL |

OLCF: Weekly Update: June 5, 2019 – Announce and Notice Lists

Message

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

OLCF: Weekly Update: June 5, 2019

OLCF Announcements via CCS-Announce <ccs-announce@email.ornl.gov>

ccs-announce@email.ornl.gov; OLCF Announcements

Wednesday, June 5, 2019 at 11:55 AM

Show Details

*** IN THIS MESSAGE ***

* Center Announcements

- Frontier CAAR Call for Proposals (Closes Jun 7)
- 2020 INCITE Call for Proposals (Closes Jun 21)
- Notable System Changes in 2019

* Meetings & Workshops

- Introduction to AMD GPU Programming With HIP Webinar (Jun 7)
- Linux Command Line Productivity Tools Tutorial (Jun 11)
- June OLCF User Conference Call (Jun 19)
- OLCF/ECP OpenMP Hackathon (Jul 22-26)
- Introduction to CMake Training (Jul 29-30)
- 2019 Petascale Computing Institute (Aug 19-23)
- OLCF GPU Hackathon Series (Open Calls; Multiple Dates)
- Want to Be a GPU Hackathon Mentor?

* Upcoming Scheduled Outages

- HPSS, DTNs (Jun 11)
- Alpine, Summit (Jun 25)

*** CENTER ANNOUNCEMENTS ***

FRONTIER CAAR CALL FOR PROPOSALS (CLOSES JUN 7)

The OLCF is now accepting proposals for Center for Accelerated Application Readiness partnerships to prepare modeling and simulation, data-intensive, and machine learning applications for highly effective use of the OLCF Frontier system that will become available to OLCF users in 2022. Applications must be received by 11:59 PM (Eastern Time) on Friday, June 7. For more information about the Frontier CAAR program or to apply, please visit <https://www.olcf.ornl.gov/frontier-center-for-accelerated-application-readiness/>

2020 INCITE CALL FOR PROPOSALS (CLOSES JUN 21)

The 2020 Innovative and Novel Computational Impact on Theory and Experiment (INCITE) Call for Proposals is now open and closes June 21. The INCITE program promotes transformational advances in science and technology for computationally and/or data-intensive, large-scale research projects through large allocations of

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 Leadership Computing Facility – The OLCF was established at Oak Ridge National Laboratory in 2004 with the mission of standing up a supercomputer 100 times more powerful than the leading...

OLCF Status (@olcfstatus) | Twitter

OLCF Status (@olcfstatus) 194 TWEETS 28 FOLLOWERS

OLCF Status @olcfstatus · Nov 18
System status change (18 Nov @ 13:40 EST): titan is up

OLCF Status @olcfstatus · Nov 18
System status change (18 Nov @ 11:45 EST): titan is down

OLCF Status @olcfstatus · Nov 18
System status change (18 Nov @ 11:05 EST): atlas2 is degraded

OLCF Titan XK7 Notice: Up (Dec 04) - deleted nagios rt status

Message

OLCF Titan XK7 Notice: Up (Dec 04)

OLCF Cray XK7 Notifications
Friday, December 4, 2015 at 6:40 PM
To: OLCF Cray XK7 Notifications

The Cray XK7 (titan) returned to service at approximately 18:25 on Dec 04

[titan-notice mailing list](#)

As an active Titan user, you have been automatically subscribed to this list.
To be removed from this list, contact help@olcf.ornl.gov

www.olcf.ornl.gov/kb_articles/communications-to-users/
www.olcf.ornl.gov/support/system-user-guides/titan-user-guide/

Center Status

Summit

Operational

Up since May 14, 2019

Next 10 days scheduled downtimes

Titan

Operational

Up since Apr 9,

Next 10 days scheduled downtimes

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 :

Miscellaneous 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)

Finding 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

Account & Project Policies



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/projects

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      Allocation      Project Totals      user1
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
 - 30 day for 100-125% of allocation
 - 365 day 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*)
 - If lots of jobs were affected, let us know (we can delay priority reduction)
 - 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

Data Management



Storage Areas

- Divided by scope
 - User
 - Project
- 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 |

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

Data Storage Locations

| Storage Area | Location ¹ | Default Permissions | Can Permissions Be Changed? | Backed Up? |
|-----------------|-------------------------|---------------------|-----------------------------|------------|
| User home | /ccs/home/\$USER | 0750 | Yes | Yes |
| Project home | /ccs/proj/project_id | 0770 | No | Yes |
| User work | \$MEMBERWORK/project_id | 0700 | No ² | No |
| Project work | \$PROJWORK/project_id | 0770 | No ² | No |
| Global work | \$WORLDWORK/project_id | 0775 | No ² | No |
| User archive | /home/\$USER | 0700 | Yes | No |
| Project archive | /home/project_id | 0770 | No | No |

¹*These are the recommended ways to reference these directories & not necessarily absolute path names.*

²*Instead of changing permissions in the work areas, move the files in question to the work area with appropriate permissions*

Data Backups

- NFS directories are backed up to a limited extent
 - Accessed via the `.snapshot` subdirectory of your user/project home directory or any subdirectory thereof
 - 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 an `"ls -a1"` of your user home/project home
- GPFS **is not** backed up
- 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
```

- 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 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 `chown` 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 (`dtm.ccs.ornl.gov`) are the preferred place for external data transfer
- 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

Using HPSS



The High Performance Storage System

- HPSS is the proper location for long-term storage
- Project home and project work areas offer a common area for files shared among project members, but neither are long-term storage
 - Home areas are somewhat space constrained
 - GPFS areas are subject to the purge
 - There's still a need for a backup location
- HPSS is accessed via the `hsi` and `htar` commands
- HPSS is also accessible via Globus

HPSS Best Practices

- Don't run multiple simultaneous transfers (very little benefit)
- Avoid numerous consecutive `hsi get` calls
- File size best practices
 - For optimal transfer performance, use files $\geq 768\text{GB}$
 - Minimum recommended file size is 512MB
 - Smaller files will be handled but read/write performance may be negatively affected
 - If you have numerous small files, we recommend bundling with `htar` to achieve the 512MB threshold
 - When using `htar`, no individual member file can be $\geq 64\text{GiB}$ (but the archive itself can be)

HPSS Best Practices

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

```
$ hsi get file1  
  
$ hsi get file2  
  
$ hsi get file3
```

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

```
$ cat getfiles.lst  
get <<EOF  
file1  
file2  
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 with `#PBS -q dtn`)
- 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>

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
- Processor Architecture
 - POWER9™
 - May have issues with precompiled software

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/>

Getting Help



Where do I find documentation?

- OLCF Website/System User Guides
<https://www.olcf.ornl.gov/>
<https://www.olcf.ornl.gov/support/system-user-guides/>
- OLCF Vimeo site (training videos)
<https://vimeo.com/channels/olcftraining>
- NVIDIA hosted documentation
<http://docs.nvidia.com>

Training Opportunities

- We host numerous training events through the year
 - Monthly User Conference Call
 - Software-specific courses
 - GPU Hackathons
<https://www.olcf.ornl.gov/for-users/training/gpu-hackathons/>
- Watch for announcements in the Weekly Update
- Check the OLCF Training Calendar
<https://www.olcf.ornl.gov/for-users/training/training-calendar/>

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
- Start a new ticket for new issues instead of replying to an old ticket
 - Helps us in classifying/searching through old tickets
 - Gives it greater visibility

Requesting a priority boost/higher walltime limit/purge exemption/etc

- Request can be made from our Documents & Forms page (in the third section of the page)
<https://www.olcf.ornl.gov/support/documents-forms/>
 - Reviewed by Resource Utilization Council, so 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

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

- <https://www.olcf.ornl.gov>
- <https://lmod.readthedocs.io/>
- <https://www.olcf.ornl.gov/support/software/software-request>
- <https://www.olcf.ornl.gov/support/documents-forms/>
- Documentation
 - <https://www.olcf.ornl.gov/support/system-user-guides/>
 - <http://docs.nvidia.com>

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/for-users/training/gpu-hackathons/>