

Best Practices @OLCF (and more)

Bill Renaud
OLCF User Support



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

OAK RIDGE National Laboratory

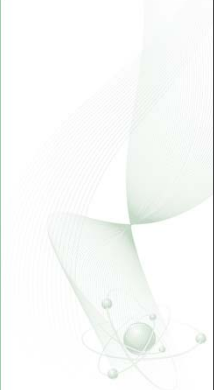
General Information

- This presentation covers some helpful information for users of OLCF
 - Staying informed
 - Aspects of system usage that may differ from your past experience
 - Common error messages
 - Common questions
 - Quick tips & best practices
- This is by no means an all-inclusive presentation
- Feel free to ask questions

Best Practices @OLCF (& more)

OAK RIDGE National Laboratory

Staying Informed



OAK RIDGE National Laboratory

Staying Informed


- OLCF provides multiple layers of user notifications about system status and downtimes
 - OLCF Weekly Update
 - Status indicators on olcf.ornl.gov
 - Opt-in email lists (*subscription management currently unavailable*)
 - Twitter (@OLCFStatus)
 - StatusCast (*coming soon*)
- A summary of these items can also be found at http://www.olcf.ornl.gov/kb_articles/communications-to-users/

Best Practices @OLCF (& more)

OAK RIDGE National Laboratory

Staying Informed-Weekly Update

- Sent weekly (Thu/Fri)
- Contains
 - Announcements about upcoming training
 - Announcements about system changes
 - Planned outages
- **All OLCF users should receive this email**
 - Let us know if you're not receiving it



Best Practices @OLCF (& more)

OAK RIDGE National Laboratory


Staying Informed-System Status

- Logs from monitoring software parsed to make educated guess on system status
- Status is sent to multiple destinations
 - Websites
 - Twitter
 - Email lists
- Fairly accurate, but still a fully automated process
 - Possibility of both false positives and false negative
 - We do take some measures to mitigate this

Best Practices @OLCF (& more)

OAK RIDGE National Laboratory

Staying Informed-System Status (via www.olcf.ornl.gov)



7 Best Practices @OLCF (8 more)



Staying Informed-System Status (via Twitter)

@OLCFstatus




8 Best Practices @OLCF (6 more)



Staying Informed-System Status (via Opt-In Email Lists)

- We also send status up/down notices via email
 - These area available on an opt-in basis
 - Online sign up (*coming soon*) at http://www.olcf.ornl.gov/kb_articles/system-notification-lists/
 - Sign up now by emailing help@olcf.ornl.gov
- Lists aren't just for status; you may want to sign up

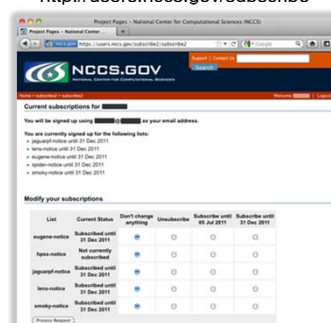


9 Best Practices @OLCF (3 more)



Staying Informed-System Status (via Opt-in Email Lists)

<http://users.nccs.gov/subscribe>



10 Best Practices @OLCF (6 more)



Staying Informed-System Status (via StatusCast)

- Users have told us via our survey that they'd like more status information online
- Work is ongoing on a new tool to provide that information
 - Current status
 - Outages (upcoming, current, and historical)
 - Announcements
- Stay tuned to the OLCF Weekly Update for more information

11 Best Practices @OLCF (6 more)



System Differences: Compiling & Finding Software



Finding Software

- Some is part of the default environment
 - Basic commands
 - Text editors
- Other software typically managed via ‘modules’
 - Software is actually installed in /sw
 - “module avail”, “module load”, etc.
 - More information is available on the OLCF website
- Compilers, MPI and GPU libraries, etc. are also provided via modules (but are not in /sw)

13 Best Practices @OLCF (& more)



Software Updates

- Cray-provided software (compilers, MPI, CUDA, etc.) goes through testing before we receive it
 - Usually not available to us immediately after vendor releases it
- OLCF provided software is moving to a model of updates at certain intervals
 - Not all minor versions will be installed
 - We'll provide build “recipes” in case you want/need a minor version

14 Best Practices @OLCF (& more)



Software Installation

- You are free to install software in your directories
 - Subject to export control, license agreements, etc.
- If you think software would be of general interest, you might ask us to install it for general use
 - Via email to help@olcf.ornl.gov or the request form at <http://www.olcf.ornl.gov/support/software/software-request/>
 - Requests are reviewed by our software council

15 Best Practices @OLCF (& more)



Compiling for the XK7 and XC30

- May differ (greatly) from your previous experience
- Mix of `?-asynce (craype)` & `PrgEnv-?` modules
 - `PrgEnv` loads compiler, math, MPI, etc. modules
 - `craype/asynce` loads compiler wrapper scripts
- Regardless of actual compiler (PGI, Intel, etc), invoke with `cc`, `CC`, or `ftn`
- MPI, math, & scientific libraries automatically linked
 - No `-lmpi`, `-lblas`, etc.
 - Can be challenging for some build processes

16 Best Practices @OLCF (& more)



Compiling for the XK7 and XC30

- You are actually cross-compiling
 - Processors/instructions sets differ between login & compute nodes
 - It's very important to remember this...utilities like `configure` don't expect this
- Compiling for batch/login nodes is occasionally necessary
 - See §7.2 of the Titan User Guide for examples <https://www.olcf.ornl.gov/support/system-user-guides/titan-user-guide/>

17 Best Practices @OLCF (& more)



Running Jobs

- Batch job information is on the web site
- Note that XK7 queue policy favors large jobs
- Jobs are charged based on what you make unavailable to others, not what you use
- Requests for high priority/quick turnaround are considered

18 Best Practices @OLCF (& more)



Debugging/Optimization

- Several tools available: DDT, Vampir, ScoreP, CrayPAT
- Some documentation is available on the OLCF website; you can also contact User Assistance

19 Best Practices @OLCF (& more)



Common Runtime Error Messages



Common Runtime Errors

- Illegal Instruction
 - A code compiled for compute nodes was executed on login/batch/service nodes
 - (Remember, you're cross-compiling)
- request exceeds max nodes alloc
 - The `aprun` request requires more nodes than the job requested
 - or–
 - Request was correct, but at launch time some nodes were discovered to be down (see potential fix later in this presentation)

21 Best Practices @OLCF (& more)



Common Runtime Errors (2)

- relocation truncated to fit: R_X86_64_PC32
 - Program is using too much static memory
 - Limit on Titan is 2GB
 - Solution: use dynamic memory allocation as much as possible

22 Best Practices @OLCF (& more)



Common Runtime Errors (3)

- `aprun: [NID 94]Exec a.out failed: chdir /autofs/na1_home/user1 No such file or directory`
 - Compute nodes only mount Lustre directories
 - When you launch `aprun`, you must be in a directory visible to compute nodes (in this example, it was launched from `$HOME`)
 - Any files used by the processes on compute nodes must also be in Lustre
 - The executable itself need not be in Lustre

23 Best Practices @OLCF (& more)



Common Questions



Why isn't my job running?

- Can be a number of reasons
 - Insufficient nodes available
 - Pending downtime
 - Queue policy
 - Dependency on another job not met
- Use `checkjob <jobid>` to diagnose
 - Use `-v` for verbose mode
 - Reason for job not running is usually near the end
- Use `showres` to see upcoming reservations/outages
 - Also shows running jobs...look carefully

25 Best Practices @OLCF (& more)



Why isn't my job running? (2)

- Insufficient resources available


```
NOTE: job cannot run (insufficient available procs: 3264 available)
```
- Unresolved dependency


```
NOTE: job cannot run (dependency 1851375 jobcomplete not met)
```
- Queue policy issue


```
NOTE: job violates constraints for partition titan (job 1854474
violates active HARD MAXJOB limit of 2 for qos smallmaxjobs user
partition ALL (Req: 1 InUse: 2))
BLOCK MSG: job 1854474 violates active HARD MAXJOB limit of 2 for qos
smallmaxjobs user partition ALL (Req: 1 InUse: 2) (recorded at last
scheduling iteration)
```

26 Best Practices @OLCF (& more)



Why isn't my job running? (3)

- Pending maintenance period/reservation
 - This is example output from `showres`
 - Look for alphanumeric ReservationIDs (those without letters are likely currently running jobs)
 - Note the full system reservation starting in 15 hours...am I asking for more than 15 hours in my job?

```
ReservationID Type S Start End Duration N/P StartTime
INTNOutage.127 User - 15:05:13 1:03:05:13 12:00:00 11/176 Tue Jan 28 08:00:00
RM_128 User - 15:05:13 1:03:05:13 12:00:00 18688/299008 Tue Jan 28 08:00:00
RM_Login.129 User - 15:05:13 1:03:05:13 15:00:00 16/2048 Tue Jan 28 08:00:00
```

27 Best Practices @OLCF (& more)



When will my job start?

- Can be difficult to tell since jobs submitted in the future can affect it
- `showstart <jobid>` is not always reliable
- `showq` output is sorted by priority (can be helpful)
- `mdiag -p` gives detailed priority info
 - Shows boosts based on job size
 - Shows overallocation penalties

28 Best Practices @OLCF (& more)



What is my project ID and what is the project allocation?

- Use `showproj` and `showusage` to list projects and usage, respectively
- Both commands have a help option (run with `-h`)

```
$ showproj
User1 is a member of the following project(s) on titan:
atf007
$ showusage
titan usage for the project's current allocation period:
Project Totals User1
Project Allocation Usage Remaining Usage
-----
atf007 7350000 11138 7338862 12
```

29 Best Practices @OLCF (& more)



What happens when my project overruns its allocation?

- Jobs will run with a lower priority to facilitate some degree of "fairshare" with projects that have allocation remaining
 - 30 day penalty for slightly over (usage 100-125% of allocation)
 - 365 day penalty for usage >125% of allocation
- Projects are **NOT** disabled for going over allocation

30 Best Practices @OLCF (& more)



My project lost X hours due to system issues...can I get a refund?

- Since we don't disable projects for going over allocation, we don't issue refunds (*per se*)
- If many jobs were affected, we can delay the priority reduction
 - This is the same thing as a refund but is much easier to manage

31 Best Practices @OLCF (& more)



Can I get a priority boost/higher walltime limit/purge exemption/etc?

- Perhaps
- Make request via <https://www.olcf.ornl.gov/support/documents-forms/> (Forms to Request Changes to Computers, Jobs or Accounts)
- Reviewed by Resource Utilization Council
- Please 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

32 Best Practices @OLCF (& more)



Are my files backed up?

- NFS directories: Yes, to a limited extent

```
$ cd ~/SomeDir/SomeOtherDir
$ ls
file1 file2
$ rm file1
$ ls
file2
$ cp -/.yesterday/SomeDir/SomeOtherDir/file1 .
$ ls
file1 file2
```

- Lustre: No
- HPSS: No
(*You might use it as a backup, but HPSS itself is not backed up. It's a good idea to have another level of backup at another site.*)

33 Best Practices @OLCF (& more)



Where do I archive data? Is there a mass storage system available?

- Yes, HPSS
- Accessed via `hsi` and `htar`
- See the "Storage @ OLCF" presentation for details https://www.olcf.ornl.gov/wp-content/uploads/2013/02/Storage_at_OLCF_2-MG.pdf

34 Best Practices @OLCF (& more)



How do I transfer my data? Can I access my data when Titan is down?

- System user guides discuss data transfer
- Data user guide: <https://www.olcf.ornl.gov/computing-resources/data-management/data-management-user-guide/>
- Data transfer nodes (`dtn.ccs.ornl.gov`) are the preferred place for internal and external data transfer
 - Often remain available during system outages
 - All users should have access

35 Best Practices @OLCF (& more)



Where do I store data?

- Can be complicated
 - NFS, Lustre, HPSS
 - User directories, Project directories
- See the Data Management Policy https://www.olcf.ornl.gov/kb_articles/data-management-policy

36 Best Practices @OLCF (& more)



How do I check my disk usage?

- For home directories, use the quota command

```
$ quota
quota: error while getting quota from master:/cm/shared for user1 (id 98765):
Connection refused
quota: error while getting quota from master:/home for user1 (id 98765): Connection
refused
Disk quotas for user user1(uid 98765):
Filesystem blocks quota limit grace files quota limit grace
nccafiler3.ccs.ornl.gov:/vol/home1
$ 10485760 10485760 2 4294967295 4294967295
nccafiler4.ccs.ornl.gov:/vol/home2
8669108 15728640 15728640 51766 4294967295 4294967295
```

- For lustre directories, use lustredu

```
$ lustredu $MEMBERWORK/stf007
Last Collected Date Size File Count Directory
2014-01-10 12:30:08 31.21 MB 9 /lustre/Atlas1/stf007/scratch/user1
```

27 Best Practices @OLCF (& more)



Where do I find documentation?

- <http://www.olcf.ornl.gov>
 - User Guides available at <https://www.olcf.ornl.gov/support/system-user-guides/>
- <http://docs.cray.com> (CrayDocs)
- <http://docs.nvidia.com> (Nvidia hosted documentation)

33 Best Practices @OLCF (& more)



Quick Tips, Best Practices, and Other Useful Information



Quick Tip-Dealing With the Lustre Purge

- Problem: Needed files were deleted by the purge
- Several “fixes”
 - Pulling from HPSS every time: inefficient
 - Multiple scripts (data in/hot in place): cumbersome
 - touch: not ideal
 - Conditional transfer: good

```
#!/bin/bash
...
if [[ ! -a $MEMBERWORK/stf007/some_important_file ]]; then
    hsi -q get /home/user1/data/some_important_file
fi
...
aprun -n 4096 ./a.out
...

```

43 Best Practices @OLCF (& more)



Quick Tip-Dealing With Failed Nodes

- Sometimes nodes fail when your job starts
 - Since node allocation has already happened, the batch system can't replace them
 - You can work around this by requesting “extra” nodes

```
#!/bin/bash
#PBS -lnodes=104
...
APRUN_RETURN_VALUE=1
while [[ $APRUN_RETURN_VALUE -ne 0 ]]; do
    aprun -n1600 ./a.out
    APRUN_RETURN_VALUE=$?
done
...

```

41 Best Practices @OLCF (& more)



Best Practices: Data Storage

- HPSS is the proper location for long-term storage
- Project areas offer a common area for files shared among project members, but is not long-term storage
 - Monitor disk usage
 - Should be backed up
 - May be purged

62 Best Practices @OLCF (& more)



Best Practices: Lustre Filesystem

- Codes may need revisions to prevent large metadata operations
- Compiling in Lustre can be slow
- There is a purge, but please delete files when no longer needed

43 Best Practices @OLCF (& more)



Best Practices: Using HPSS

- HPSS prefers a few large files rather than many small files (`htar` is your friend)
- Running many simultaneous transfers can be problematic
- Planned upgrades (larger disk cache) should make a notable difference in performance

44 Best Practices @OLCF (& more)



Best Practices: User Support

- Email is often useful if there are long/complicated error messages
- Send as many error messages as possible
 - Or place them in a file & direct us to it
- When sending codes, create a `.tar` file & direct us to it
 - Include all files necessary to run
 - More efficient than sending via email
- Send new issues in new tickets, not replies to old ones

45 Best Practices @OLCF (& more)



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

46 Best Practices @OLCF (& more)

