


Improving User Notification on Frequently Changing HPC Environments

Chris Fuson

William Renaud

James Wynne III

CUG2016
May 8-12, 2016
London



Oak Ridge Leadership Computing Facility

- Users cover multiple science domains and experience levels
- New projects and users are added to system throughout year
- Average over 1,000 active users yearly
- Projects receive 100x more hours than at other generally available centers.
- Average 250 projects per year
- OLCF users are spread out all over the globe and come from academia, national laboratories, other government agencies, and industry



Titan

Cray XK7
#2 HPC Top 500
27 petaflops
299,008 Processor Cores
18,688 GPUs

Eos

Cray XC30
11,776 Processor Cores
47 TB
Aries interconnect

DTNs

Data Transfer Nodes
Interactive nodes
Batch scheduled nodes
Tuned for wide-area

Rhea

Linux cluster
8,192 Processor cores
9 1TB big memory nodes
9 nodes: 2 NVIDIA® K80

Spider

Center-wide Lustre
32 PB
2 Filesystems:
- atlas1
- atlas2

HPSS

High Performance
Storage System
Longer-term storage
Tape and Disk storage

Frequently Changing Environment

Changes

- System user environment
 - Software, libraries, compilers
 - Batch system
 - Environment variables
 - Login resources, compute resources
- Adding, removing, changing defaults
- OS upgrades
- Center policies


Impacts

- Changing center resources impacts use of system
- Recompile often required
- Software, library, compiler, driver versions dependencies
- Compute time allocations are limited

Complications

- Many environment variables, libraries, and similar exist behind the scenes
 - Programming environment used to hide complexity also hide many libraries, versions, variable triggers
 - Build processes hide libraries and dependencies
- Changes made to systems during production
 - Project allocations start and end throughout year
 - Can not wait until end of project allocation
- Not all change impacts all users
 - multiple science categories
 - variety of codes and workflows
- Many have access to multiple centers with varying policies and configurations
- Varying levels of experience
- Systems available and accessible 24x7

Average Hourly Logins



Notification Methods

Passive

Interrupt



web

email

wall

NOTICES

- 24 MAY 2016 User Meeting - Agenda Now Available
- 23 MAY OLCF Tutorial Day

To: [ccs-announce](#)
Subject: OLCF: Weekly Update, March 23, 2016

*** IN THIS MESSAGE ***

- * Center Announcements
 - Gordon Bell Award Nominations (due Apr 15)
 - Allinea Forge (MAP + DDT) Now Available
 - Removal of Old Software Versions on Eos (Mar 29)
- * Meetings & Workshops
 - INCITE Proposal Writing Webinar (Apr 13 & May 19)
 - 2016 User Meeting: Registration is Open
 - 2016 GPU Hackathons
 - March OLCF User Conference Call (Mar 23)
- * Scheduled Outages Through April 1 (none)

*** CENTER ANNOUNCEMENTS ***

```
Broadcast Message from root@titan-login2
at 07:43 ...

System will be shutting down at 08:00 for
scheduled maintenance.
```

Email

- Most common communication method
 - Center events
 - System change
 - Policy change
- Management systems
 - Dissemination to large groups
 - Allows staff to send message on behalf of center
- Issues
 - Easy to loose in email mass
 - “I’ll read later”

Too Much vs. Not Enough Email

Only the Necessities

- Low volume
 - Mandatory
 - Contains all valid OLCF users
 - Weekly email

More, Please

- High volume
 - Optional, can choose to opt-out
 - Contains active users
 - Automated system state change

List for each resource

- Titan
- Rhea
- HPSS
- Eos
- DTNs
- Lustre

Email Automation

Active Users

- High volume list membership changed from opt-in to opt-out
- Automatically populate each resource list with active users
- Active user defined:
 - Successful authentication
 - Job in the batch queue

System Status Change

- Automatically email resource's high volume list when resource state changes
- Based on resource log monitoring

```
From: OLCF Rhea Notifications
Reply-To: help@olcf.ornl.gov
Date: Sunday, April 24, 2016 at 8:00 PM
To: rhea-notice
Subject: OLCF Rhea Notice: Up (Apr 24)
```

```
The Data Analysis System (rhea) returned to
service at approximately 19:45 on Apr 24
```

Email Format and Size

- Emails are more useful if they are read
- Short, concise
 - Details on web
- Bulleted topics
 - Allows quick scan

```
From: OLCF Announcements
Reply-To: "help@olcf.ornl.gov"
Date: Wednesday, March 23, 2016 at 11:26 AM
To: ccs-announce
Subject: OLCF: Weekly Update, March 23, 2016

*** IN THIS MESSAGE ***
* Center Announcements
  - Gordon Bell Award Nominations (due Apr 15)
  - Allinea Forge (MAP + DDT) Now Available
  - Removal of Old Software Versions on Eos (Mar 29)

* Meetings & Workshops
  - INCITE Proposal Writing Webinar (Apr 13 & May 19)
  - 2016 User Meeting: Registration is Open
  - 2016 GPU Hackathons
  - March OLCF User Conference Call (Mar 23)

* Scheduled Outages Through April 1 (none)

*** CENTER ANNOUNCEMENTS ***
GORDON BELL AWARD NOMINATIONS DUE APRIL 15, 2016
The deadline for submissions for the annual Gordon Bell
Friday, April 15, 2016. If you are planning on submitt
```

Targeting Smaller Groups of Users

- Notifying entire user community not always most effective
 - Not all change impacts all users
 - Users may not have access to all center resources
- Staying out of the spam folder
 - Too much irrelevant email
- Contact only those impacted by change
 - System access, software/library users
 - Contact at time of use
- Methods to be discussed
 - Wrapping command-line tools
 - Web authentication
 - Email groups
 - Write to ALL

Email, Web, Write to All

- Write to All
 - Still effective
 - Message to screen of all connected
 - Provide warning of resource outage
- Email
 - Individual/small group emails
 - Multiple email lists
 - Group by resource access
 - Enabled users for system or resource
 - Titan, Eos, Rhea, HPSS, Spider
 - Active user lists containing users who have authenticated recently
- Web
 - Main site, can only target users based on page/section
 - Authentication, allows altering page content based on viewing user

```
Broadcast Message from root@titan-login2  
at 07:43 ...
```

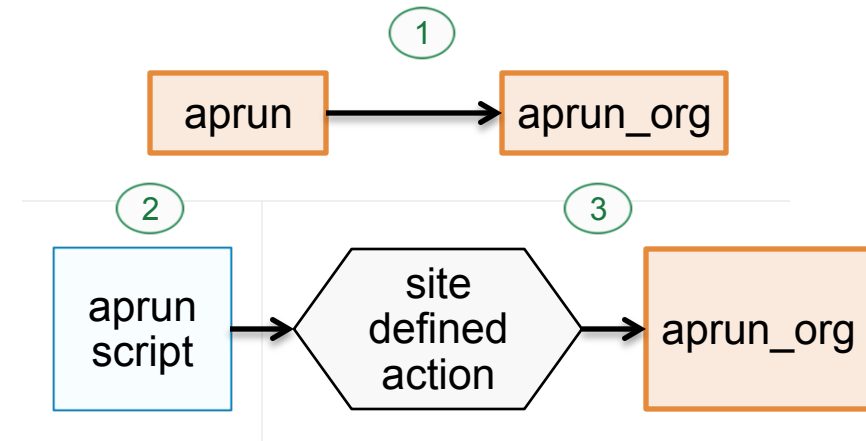
```
System will be shutting down at 08:00 for  
scheduled maintenance.
```

```
zcmehqjefq wazupfenuance:  
zlgfem wllj pe zpuuffrjg qomw of 08:00 fo4
```

Wrapping Command-Line Tools

- **Goal:** Insert center created action between user and existing tool

1. Rename tool
2. Create script with same name as tool
3. Script performs center-defined action before calling the original tool



- **Example actions**

- Verify given arguments
- Enforce center policy
- Log tool execution
- Provide user friendly notice
- Reject before executing original tool

- **Example command line tools**

- qsub
- aprun
- module

Environment Modules

- Environment modules used on all OLCF systems
- Diverse user community requiring multiple versions of software, libraries, compilers
- Center must provide multiple packages and versions of each
- Control defaults, future and previous versions, parings
- Multiple maintainers
 - Cray provides as part of Programming Environment
 - Center maintained

Why Wrap Environment Modules

- Provides opportunity to target users who are using a tool, library, or compiler
- Load tool specific module prior to use
- Change versions for non-default

```
eos-ext1 1003> module load cray-petsc/3.4.2.0

#####
##                               NOTICE                               ##
#####

The following cray-petsc versions will be removed on
Tuesday, March 29. Please use a more recent version.
  3.4.2.0
  3.4.3.1
  3.4.4.0
  3.5.1.0

If you still require a version from the list,
please contact help@olcf.ornl.gov

#####

eos-ext1 1004>
```


Wrapping Environment Module

```
titan> module load ddt/5.1
```

```
module ()  
{  
  eval `modulecmd bash load ddt/5.1`  
}
```

- Module is a function/alias
- Can not simply wrap module
- Modulefiles maintained by Cray and center
- Need a solution that supports both

/sw/modulefiles/ddt/5.1

##%Module

```
set ddt_version 5.1  
set apppath /sw/xk7/ddt/5.1  
setenv DDT_HOME $appath  
prepend-path PATH $appath/bin
```

Wrapping Environment Module

- Replacing modulecmd with wrapper allows interaction with both Cray and center maintained modulefiles without modifying modulefile
- Using environment variable keys allows users to disable features

modulecmd_wrap

```
### Perform center defined actions
if (defined $ENV{'OLCF_MODULE_MSG'})
{
    $module_msg = $ENV{'OLCF_MODULE_MSG'}.'/bin/module_msg';

    if (-e $module_msg)
    {
        $results = `timeout 2 $module_msg @ARGV 2>&1`;
        print STDERR "$results";
    }
}

### Call the real binary
$return = system("$modulecmd @ARGV");
```

Conclusion

- Our systems will continue to change
- Change will likely be no less impactful
- Centers must continue to notify
- Notifications evolve to reach user base

Acknowledgements

- Many people contributed to this effort: Cathy Willis, Ross Miller, Scott Atchley, Robert French.
- This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

Questions?

Thank you!