Overview of the OLCF

Presented by:

Fernanda Foertter

Oak Ridge Leadership Computing Facility (OLCF)

ORNL's mission

Deliver scientific discoveries and technical breakthroughs that will accelerate the development and deployment of solutions in clean energy and global security, and in doing so create economic opportunity for the nation

Leadership Computing

Leadership computing capability is required for scientists to tackle the high-resolution, multi-scale/multi-physics simulations of greatest interest and impact to both science and the nation

Leadership Computing capability is typically 10-100X more than available at other computational centers

The work at the OLCF is mission critical to inform policy decisions and advance innovation in far reaching topics such as

- **energy assurance**
- **ecological sustainability**
- **scientific discovery**
- **global security**

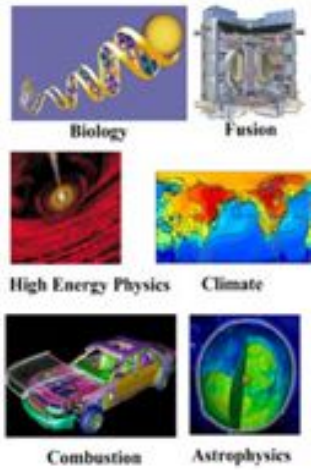

What is the Leadership Computing Facility?

- Collaborative, multi-lab, DOE/SC initiative ranked top national priority in *Facilities for the Future of Science: A Twenty-Year Outlook*
- Mission: Provide the computational and data science resources required to solve the most important scientific & engineering problems in the world
- Highly competitive user allocation program (INCITE, ALCC)
- Projects receive 100x more capability than at other generally available centers
- The OLCF partners with users to enable science & engineering breakthroughs. (Liaisons)

OAK RIDGE National Laboratory

What is Titan?

- The next phase of the Leadership Computing Facility program at ORNL
- An upgrade of Jaguar from 2.3 PF to 27PF
- Built with Cray's newest XK6 compute blades

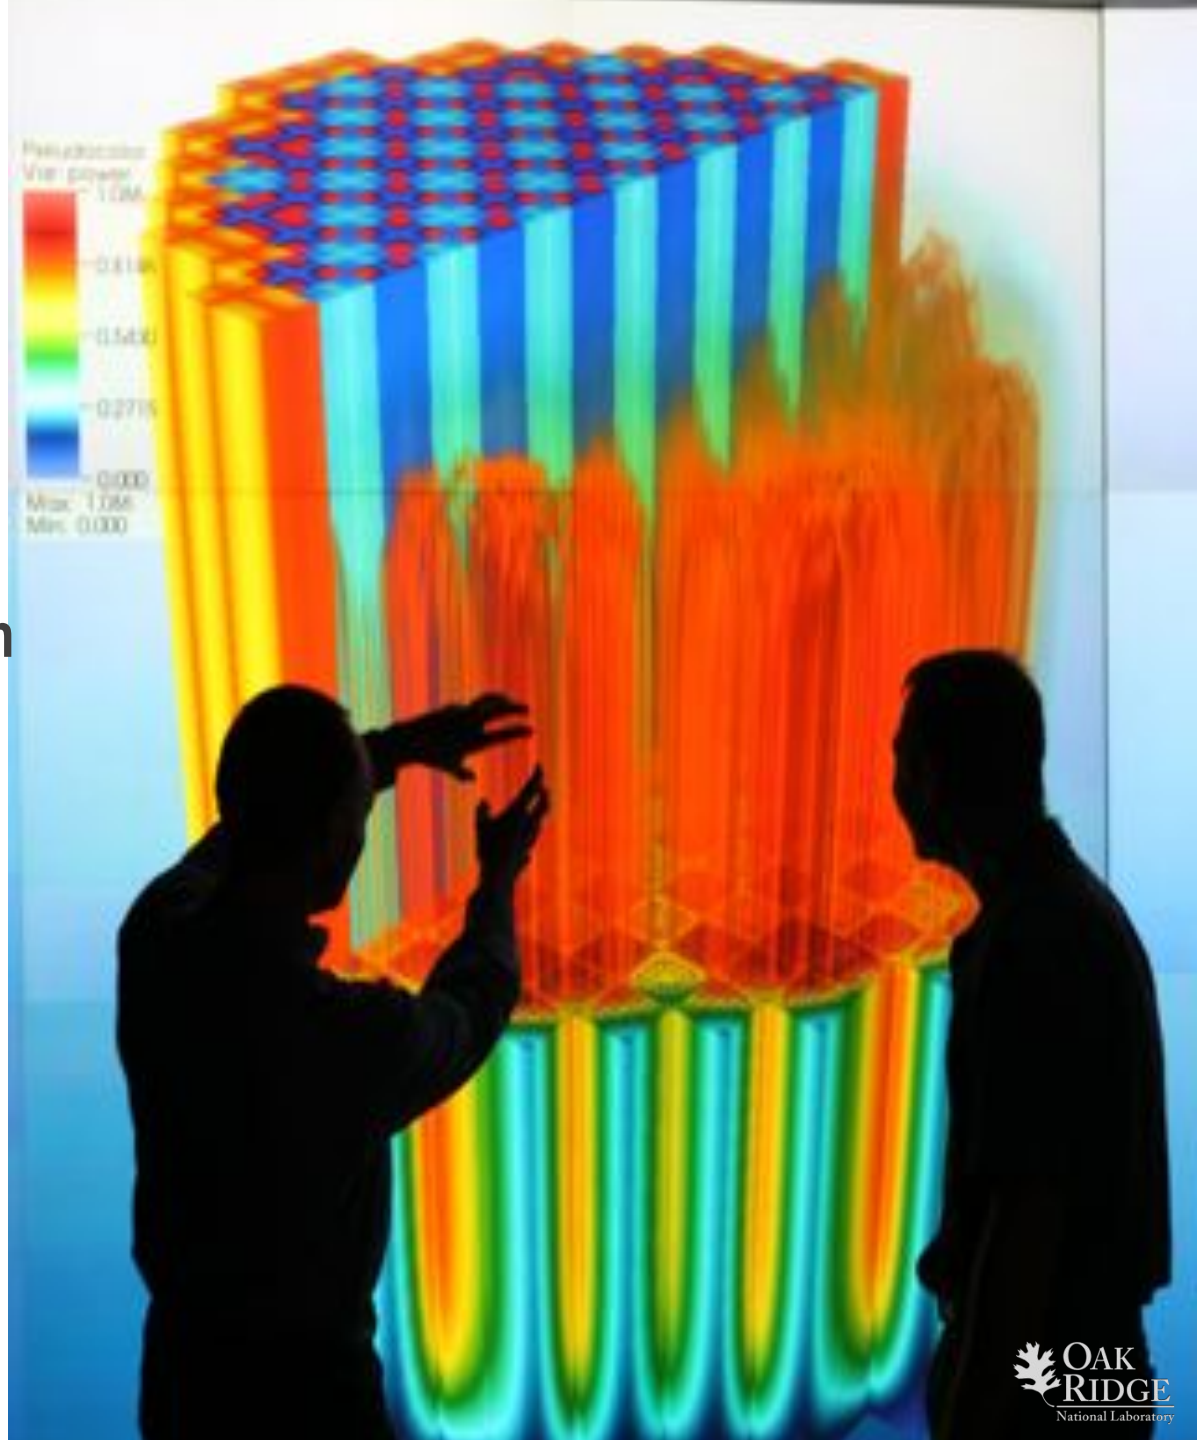

Cray XK7 Compute Node

XK7 Compute Node Characteristics
AMD Opteron 6200 Interlagos 16 core processor @ 2.2GHz
Tesla M2090 @ 665 GF with 6GB GDDR5 memory
Host Memory 32GB 1600 MHz DDR3
Gemini High Speed Interconnect
Upgraded to NVIDIA's next generation KEPLER processor in 2012
Four compute nodes per XK7 blade. 24 blades per rack

Titan breakthrough performance

Jaguar specs (2011)	
Compute nodes	18,688
Login & I/O nodes	512
Memory per node	24 GB
# of Opteron cores	224,256
# of NVIDIA K20 “Kepler” processors (2013)	NA
Total system memory	450 TB
Total system peak performance	2.3 petaflops

Titan specs (2013)	
Compute nodes	18,688
Login & I/O nodes	512
Memory per node	32 + 6 GB
# of Opteron cores	299,008
# of NVIDIA K20 “Kepler” processors (2013)	18,688
Total system memory	710 TB
Total system peak performance	27 petaflops

Hybrid Programming Model

- On Jaguar with 299,008 cores, we were seeing the limits of a single level of MPI scaling for most applications
- To take advantage of the vastly larger parallelism in Titan, users need to use hierarchical parallelism in their codes
  - Distributed memory: MPI, SHMEM, PGAS
  - Node Local: OpenMP, Pthreads, local MPI communicators
  - Within threads: Vector constructs on GPU, libraries, CPU SIMD
- ***These are the same types of constructs needed on all multi-PFLOPS computers to scale to the full size of the systems!***

Running on Titan

Leadership Metric and Scheduling Policy

As a DOE Leadership Computing Facility, the OLCF has a mandate to be used for large, *leadership-class* (aka *capability*) jobs.

To that end, the OLCF implements queue policies that enable large jobs to run in a timely fashion.

- Basic queue priority is set by the time a job has been waiting relative to other jobs in the queue.
- However, we use several factors to modify the *apparent* time a job has been waiting. These factors include:
  - The job's processor core request size.
  - The queue to which the job is submitted.
  - The 8-week history of usage for the project associated with the job.
  - The 8-week history of usage for the user associated with the job.

Leadership Usage Metric:

35% of the CPU time used on the system will be accumulated by jobs using 20% or more of the available processors (60,000 cores)

OLCF Scheduling Policy

Bin	Min Nodes	Max Nodes	Max Walltime (Hours)	Aging Boost (Days)
1	11,250	- - - - -	24	15
2	3,750	11,249	24	5
3	313	3,749	12	0
4	125	312	6	0
5	1	124	2	0

Bin 2 is the leadership mark.

OLCF Allocation Overuse Policy

Projects that overrun their allocation are still allowed to run on LCF systems, although at a reduced priority.

- For projects that have used between 100% and 125% of their allocations, the following rules apply:
  - Jobs have their priority reduced by 30 days.
- For projects that have used greater than 125% of their allocation, the following rules apply:
  - Jobs have their priority reduced by 365 days.

To view the entire scheduling policy please see:
http://www.olcf.ornl.gov/kb_articles/scheduling-policy-olcf/

Job Resource Accounting on Titan

•Each Titan XK7 node will be *defined* as possessing 30 “Titan core-hours” - 16 CPU cores + 14 GPU core equivalents.

•Titan will be scheduled in full node increments; a node’s cores cannot be allocated to multiple jobs.

•Notably, codes that do not take advantage of GPUs will have only (16) CPU cores available per node; however, allocation requests–and units charged–will be based on (30) cores per node.

•For more information, visit http://www.olcf.ornl.gov/kb_articles/job-resource-accounting/.

Two Ways to Check Utilization

Graphical

https://users.nccs.gov

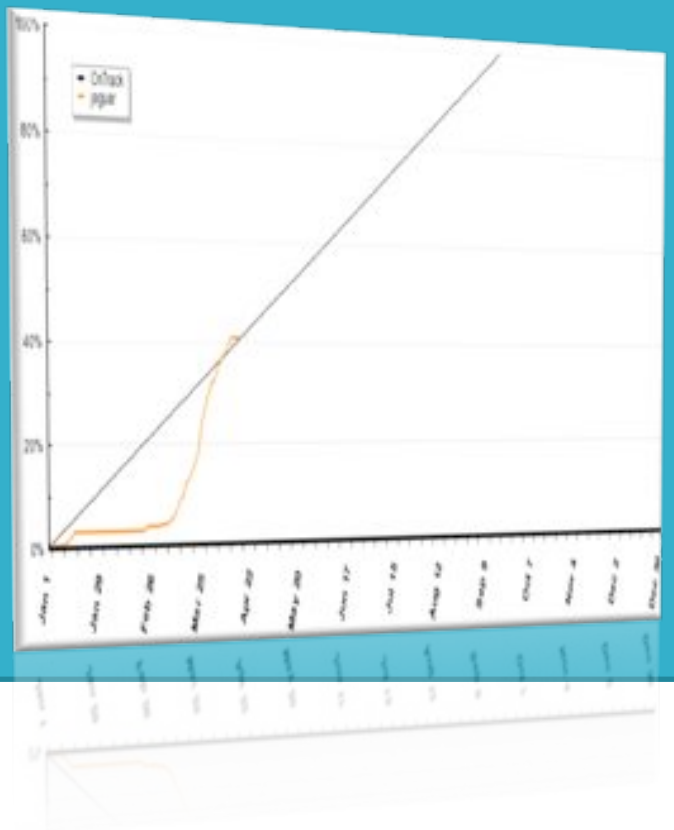

Command Line

> showusage

Can be executed from any OLCF system:

• jaguar	• dtn01
• home	• dtn02
• lens	• frost

showusage

Without arguments the utility returns usage for each project and subproject on which the user has an account.

```
> showusage
jaguar usage in CPU hours:
                              Project Totals             userid
 Project     Allocation       Usage    Remaining          Usage
____________________|________________________|___________
 PRJ001       600001    |   26255.27    573745.73  |   155.96
 PRJ001sub    500000    |       0.00    500000.00  |     0.00
>
```

showusage

Usage:

showusage [-h] [-help]

showusage [-s <*system*>] [-p <*projectid*>] [-f]

Options:

-h, -help	showusage options
-s <*system*>	display usage for specified system
-p <*projectid*>	display usage for specified project
-f	list usage for all project members

Users Web Site

- users.nccs.gov
- Updated daily
- Access to the site is limited to current OLCF users with a valid SecurID fob

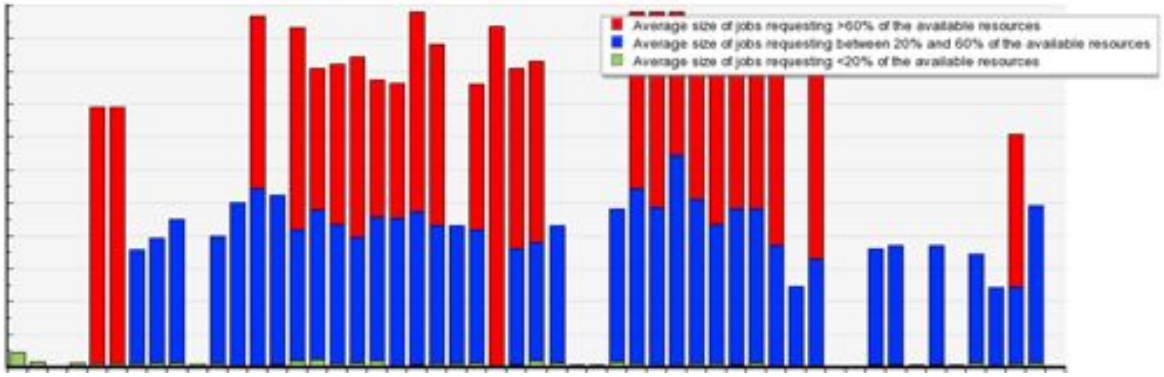

- Provides the following information by Project, Subproject, and System:

	Alloc Usage	Remaining Allocation	Usage by User	Usage by Job Size	User Usage by Job Size	Batch Priority	Enabled Users
Monthly	✔	✔	✔	✔	✔		
YTD	✔	✔	✔	✔	✔	✔	✔

Quarterly Reports

- Principal Investigators of current Titan projects must submit a quarterly progress report.
- Due to the importance of these reports, the OLCF imposes the following penalties for late submission:
  - **1 Month Late** - Job submissions against offending project will be suspended
  - **3 Months Late** - Login privileges will be suspended for all OLCF resources for all users associated with offending project

Acknowledgement of OLCF Resources

Please use the following acknowledgement in publications based on work performed using OLCF systems.

"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."

The acknowledgement can be found on the OLCF website at https://www.olcf.ornl.gov/media-center/media-kit/.

Policies

- To view all OLCF policies in one place, visit https://www.olcf.ornl.gov/support/system-user-guides/olcf-policy-guide/

OLCF Organization


Director of Science
J. Wells
Director of Operations
J. Rogers
Chief Technology Officer
A. Geist
National Center for Computational Sciences
Oak Ridge Leadership Computing Facility
J. Hack, Director
A. Bland, OLCF Project Director
L. Gregg, Division Secretary
Industrial Partnerships
S. Tichenor
INCITE Program
J. White
Chief Computational Scientist
R. Kendall
Cray Center for Excellence
J. Levesque
J. Larkin
N. Wichmann
K. Seymour
Deputy Project Director
K. Boudwin
B. Hammontree, Site Preparation
J. Rogers, Computer Acquisition
R. Kendall, (Acting) Pre-commissioning & Accept. Test Dev.
A. Baker, (Acting) Commissioning
D. Hudson, Project Management
A. Barker, Training & Support Development
Operations Council
W. McCrosky, Finance Officer
S. Hummel, HR Rep.
J. Presley, Recruiting
S. Milliken, Facility Mgmt.
D. Edds, ES&H Officer
R. Adamson, K Bivens, Cyber Security
Computer Science Research[5]
D. Bernholdt
C. Sonewald
S. Boehm[8]
M. Brimm
N. Forrington[6]
M. Gorentla-Venkata
O. Hernandez
T. Jones
C. Kartsaklis
T. Mintz[7]
T. Naughton
D. Schubert[6]
J. Schuchart[8]
P. Shamis
G. Vallee
J. Vasnier[6]
Scientific Computing
B. Messer[9]
A. Fields
M. Abbasi[5]
S. Ahern[3,5]
V. Anantharaj
M. Brown
J. Daniel
M. Eisenbach
M. Fahey[3,5]
J. Hill
W. Joubert[#]
S. Klasky[5]
R. Kumar[7]
Y. Li[7]
M. Matheson
B. Messer
H. Nam
T. Nguyen[7]
M. Norman
S. Pannala[5]
S. Parete-Koon
N. Podhorszki[5]
D. Pugmire
M. Rodriguez-Rodriguez[7]
D. Rosenberg
R. Sankaran
R. Tchoua[5]
A. Tharrington[#]
Y. Tian[5]
R. Toedte
Z. Wang[7]
User Assistance and Outreach
A. Barker
S. Mowery
A. Carlyle
A. DiGirolamo
F. Foertter
C. Fuson
B. Gajus[5]
M. Griffith
J. Hines[#]
S. Jones[5]
M. Keele[1]
D. Levy[5]
M. Miller
L. Moore[1]
T. Nguyen[7]
S. Parete-Koon
L. Rael
S. Ray
B. Renaud
D. Rose
J. Rumsey[1]
A. Simpson
J. Smith
R. Whitten[3,5]
L. Williams[5]
Technology Integration
S. Vazhkudai
S. Mowery
S. Atchley
T. Barron
D. Dillow
D. Fuller
R. Gunasekaran
J. Harney
S. Hicks[5]
Y. Kim
R. Miller
S. Oral
B. Settlemyer[5]
J. Simmons
B. Smith
D. Stansberry
D. Steinert
C. Wang[2]
F. Wang
V. White
High-Performance Computing Operations
K. Thach
S. Allen
R. Adamson
J. Anderson
M. Bast
J. Becklehimer[4]
K. Bivens
J. Breazeale[6]
J. Brown[6]
B. Caldwell
A. Enger[4]
C. England
M. Ezell
J. Evanko[4]
A. Funk[4]
D. Garman[4]
D. Giles
M. Gyurgyik
M. Hermanson
J. Hill
J. Kincl
S. Koch
C. Layton
C. Leach[6]
D. Leverman
J. Lewis[4]
D. Londo[4]
D. Maxwell[@]
M. McDonough[4]
J. Miller[6]
D. Pelfrey
G. Phipps, Jr.[6]
S. Shpanskiy
B. Tennessen[4]
J. Walsh[4]
S. White
C. Willis[4]
T. Wilson[6]
[1]Intern
[2]Post Graduate
[3]JICS
[4]Cray, Inc.
[5]Matrixed
[6]Subcontractor
[7] Post Doc
[8] Post Master
[9] Acting Group Leader
[#] Task Lead
[@] Technical Coordinator

User Support at the OLCF

Customer Support Interfaces
Scientific liaisons
User assistance analysts
Problem resolution
User satisfaction surveys

Outreach to the Next Generation
Learning opportunities for students and faculty

Communication with Users
Multiple channels; easy and frequent access

Training Current and Future Users
Workshops tailored to user needs

User Support Analysts

- Help users access OLCF resources
- Help users compile and debug large science and engineering applications
- Identify and resolve system-level bugs in conjunction with other technical staff and vendors
- Install third-party applications and provide documentation for usage
- Engage center staff to ensure users have up-to-date information
- Research, develop, and maintain support materials for users
- Communicate with users
- User advocates
- Develop and deliver training

OLCF Liaisons

- All INCITE projects are assigned a liaison
- PhD-level scientists with expertise in computation
  - Astrophysics, biology, chemistry, climate, computer science, engineering, materials science, mathematics, nuclear physics, plasma physics, etc.
  - Experienced computational scientists with one thing in common
- Liaisons matched with INCITE projects based on science, mathematical, and algorithmic expertise
  - Can't always match for science first, e.g., I am not a chemist, but I am familiar with their math and algorithms
- Our motto: Whatever it takes!

Liaison Role

- Liaisons are collaborators whose unique expertise with leadership-level computers will enhance your experience and help you get more science done
- Levels of liaison support
  - Level 1: User support +
  - Level 2: Paratrooper – fix a specific problem in your code, O(1 month)
  - Level 3: Embedded member of code development team and science collaborator

Liaison Role

- Typical liaison activities
  - Profile code performance, providing feedback to code team
  - Code porting
  - Implement solutions to problems experienced by application scientists
  - Advocate for users regarding tools, libraries, etc.
  - Collaborate scientifically

Training

- Training targets:
  - OLCF Users
  - OLCF Staff
  - OLCF Future users
- Training Formats
  - In person
  - Webcast
  - Online Tutorials
  - Planned: Video channel, frequent short webcasts, expand tutorials

Request Assistance from the OLCF

- Send email to help@olcf.ornl.gov
- Call (865) 241-6536 (9-5 ET M-F)
- Contact your liaison directly

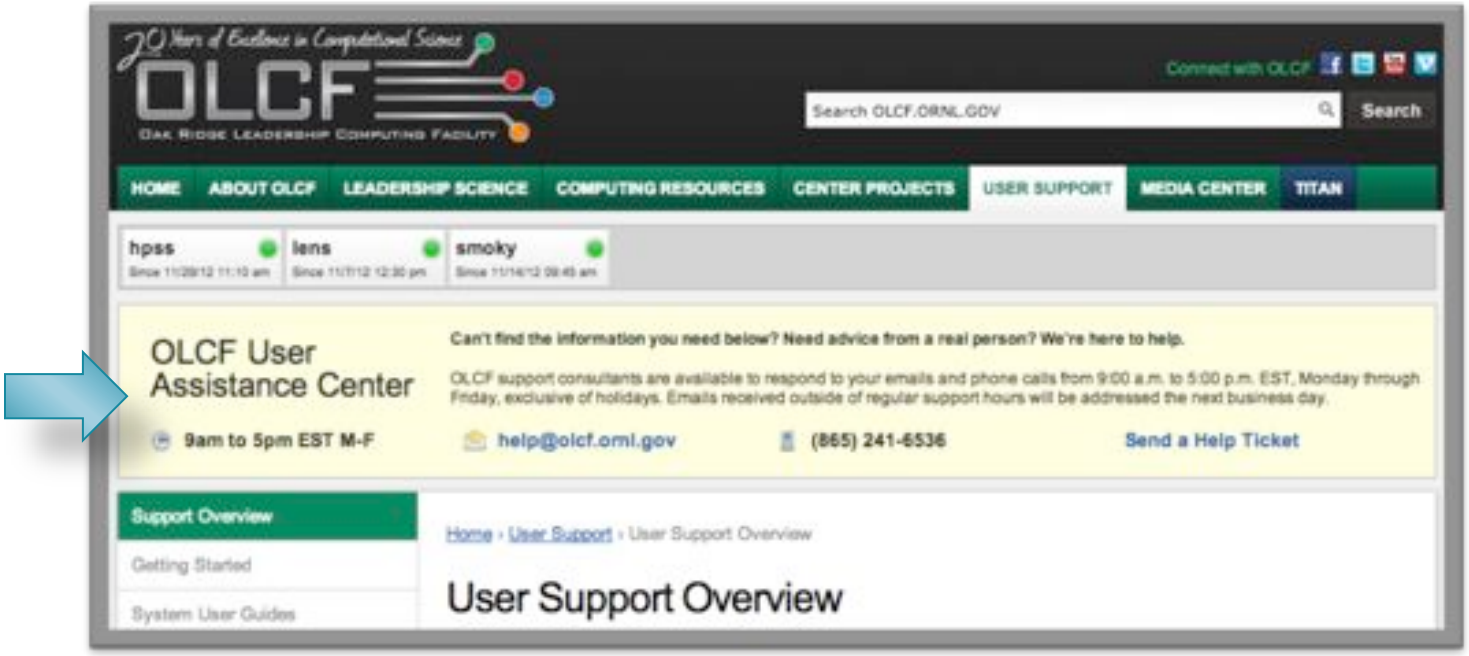

OLCF Support Site

OLCF Support Site

- All support information available at:

`www.olcf.ornl.gov/support/`

  - Titan user guide
  - GPU programming tutorials
  - Software inventory
  - Knowledge base
  - Known issues
  - Official policies

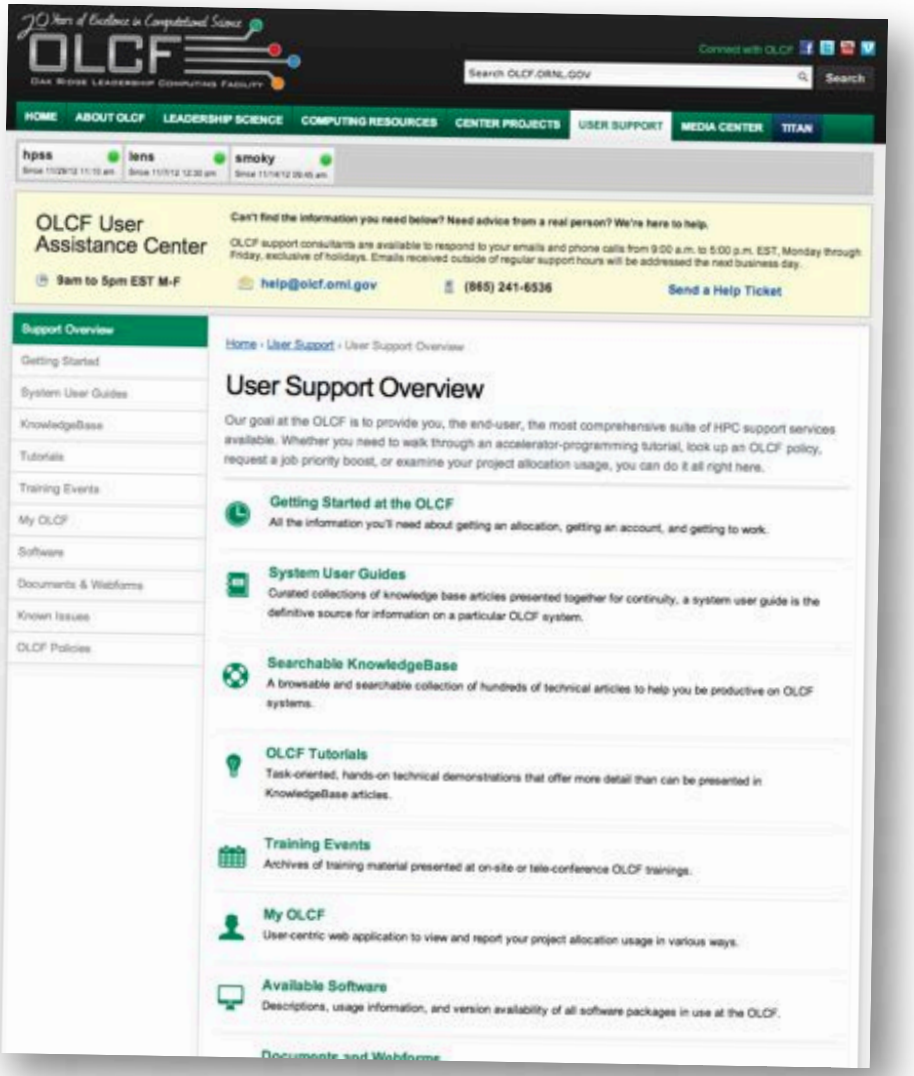

Titan User Guide

- Definitive source for support information on Titan

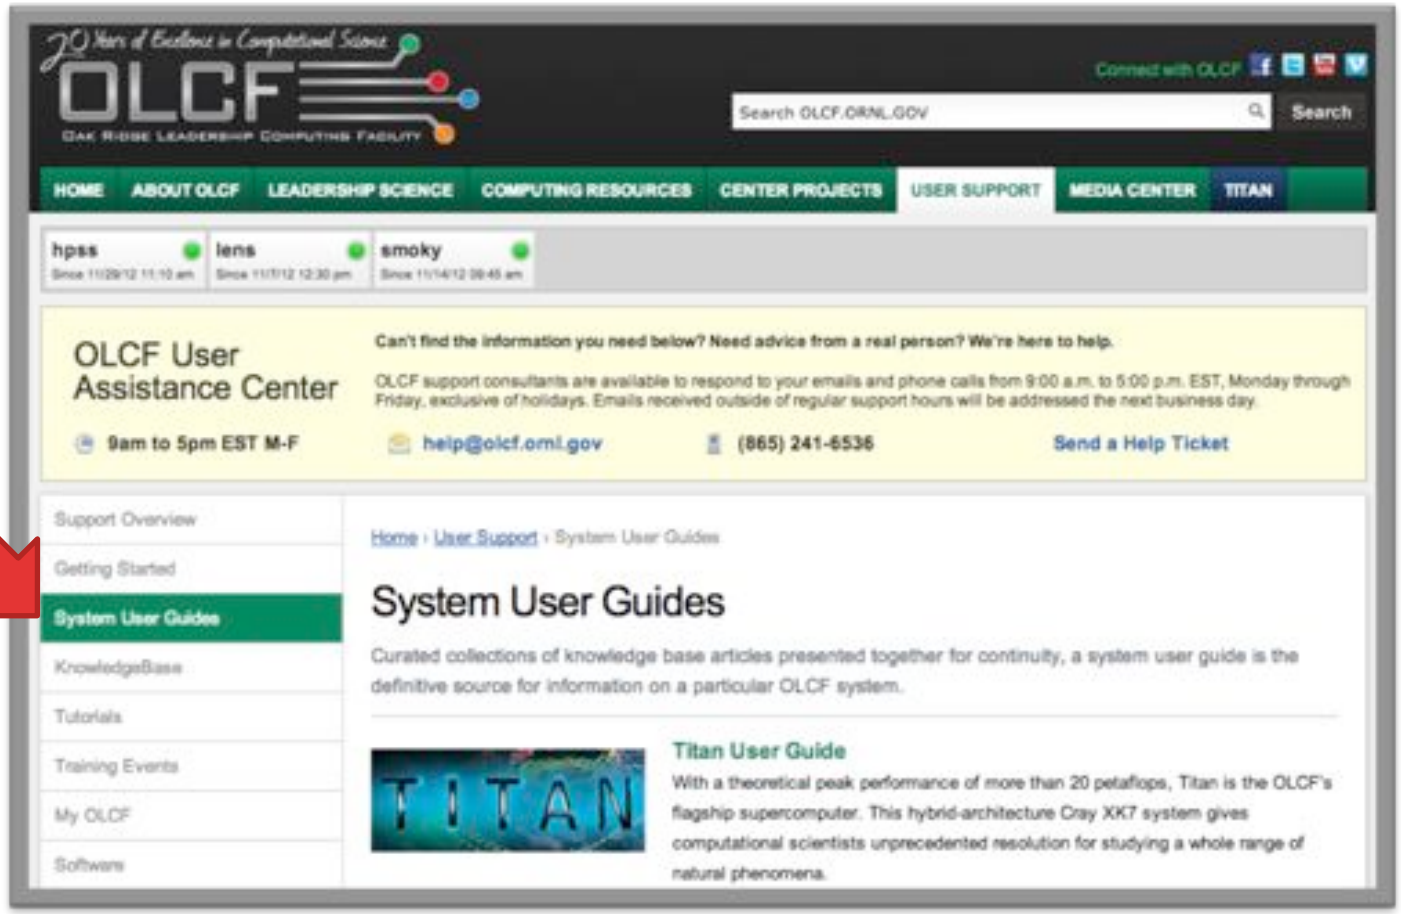

Titan User Guide

- User guide content is aggregated from searchable OLCF knowledge base.

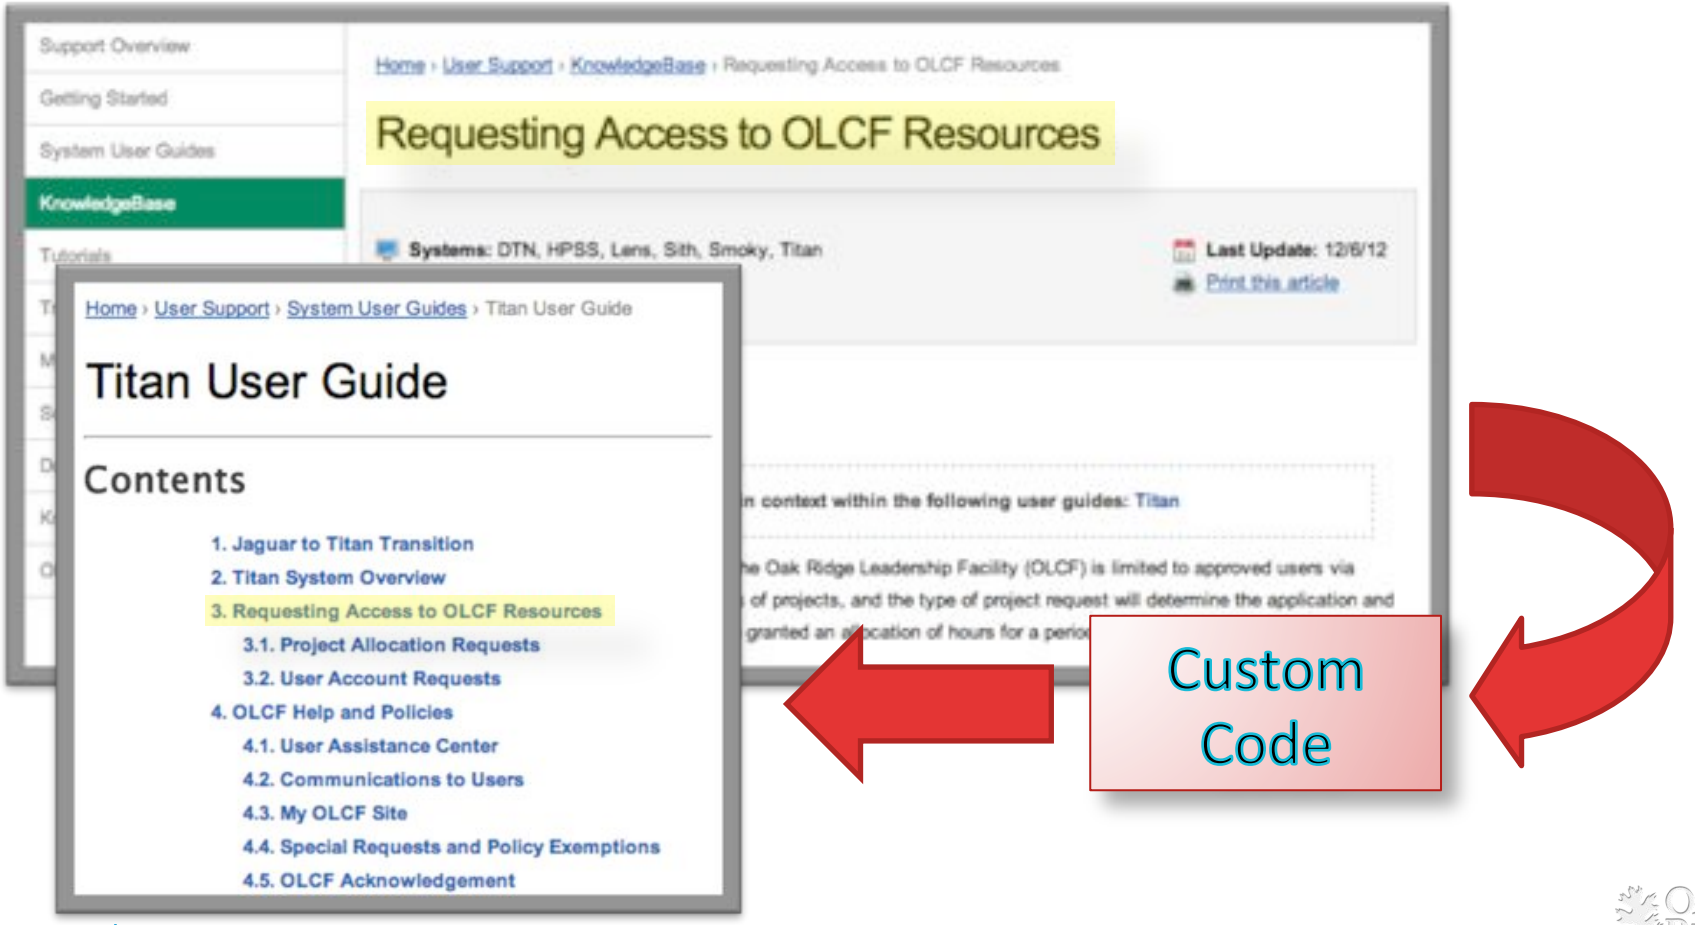

Tutorials

Task-oriented technical demonstrations that offer more detail than can be presented in knowledge base articles.

Software Installation Requests

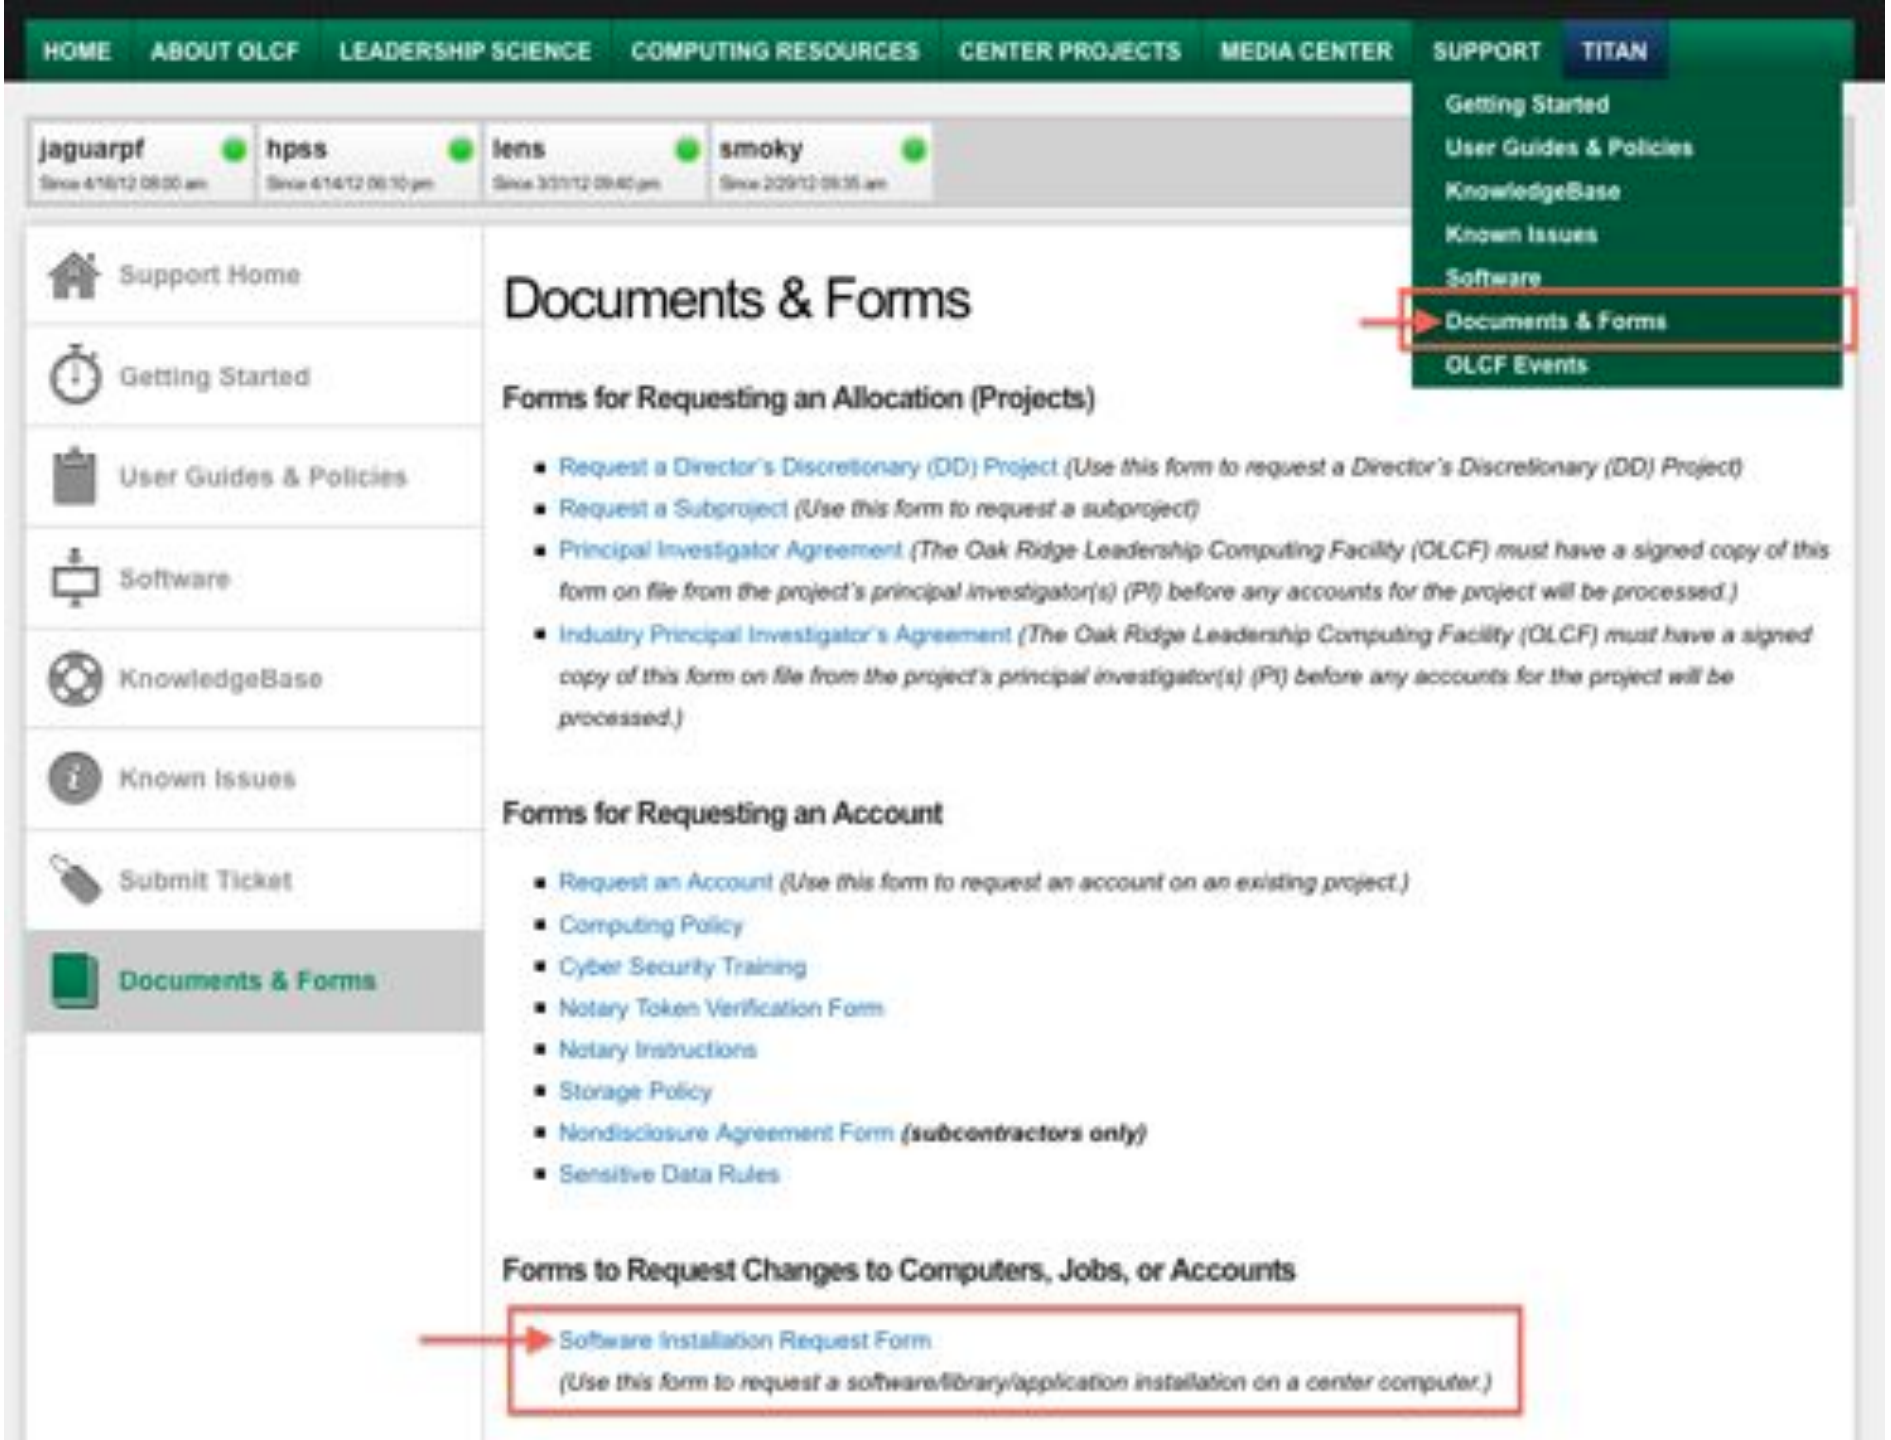
HOME
ABOUT OLCF
LEADERSHIP SCIENCE
COMPUTING RESOURCES
CENTER PROJECTS
MEDIA CENTER
SUPPORT
TITAN
Getting Started
User Guides & Policies
KnowledgeBase
Known Issues
Software
Documents & Forms
OLCF Events
jaguarpf
hpss
lens
smoky
Support Home
Getting Started
User Guides & Policies
Software
KnowledgeBase
Known Issues
Submit Ticket
Documents & Forms
Documents & Forms
Forms for Requesting an Allocation (Projects)
Request a Director's Discretionary (DD) Project (Use this form to request a Director's Discretionary (DD) Project)
Request a Subproject (Use this form to request a subproject)
Principal Investigator Agreement (The Oak Ridge Leadership Computing Facility (OLCF) must have a signed copy of this form on file from the project's principal investigator(s) (PI) before any accounts for the project will be processed.)
Industry Principal Investigator's Agreement (The Oak Ridge Leadership Computing Facility (OLCF) must have a signed copy of this form on file from the project's principal investigator(s) (PI) before any accounts for the project will be processed.)
Forms for Requesting an Account
Request an Account (Use this form to request an account on an existing project.)
Computing Policy
Cyber Security Training
Notary Token Verification Form
Notary Instructions
Storage Policy
Nondisclosure Agreement Form (subcontractors only)
Sensitive Data Rules
Forms to Request Changes to Computers, Jobs, or Accounts
Software Installation Request Form
(Use this form to request a software/library/application installation on a center computer.)

Software Requests

Please provide as much information as possible on the form

http://www.olcf.ornl.gov/support/software/software-request

Requests are reviewed within two weeks:

- Accepted the software will be installed for all users.
- Licensed software will take longer to install.
- If rejected, don't panic! We can help you build it in your home or scratch directories.

OLCF Communications

Multiple OLCF Communication Channels

- Weekly email messages
- Website announcements
- Social media
- Smart phone apps
- More info to come in a later talk

Questions & Discussion