Introduction to the Open Science Grid

OSG Software Carpentry Workshop

Duke University

Oct 29th, 2015

Emelie Harstad <eharstad@unl.edu>

OSG User Support

Outline

- What is the OSG?
- Virtual Organizations
- Jobs that are good for OSG/DHTC
- OSG Entry Paths (Duke Ci-Connect)
- Job Submission
- User Tools

The Open Science Grid

A framework for large scale distributed resource sharing addressing the technology, policy, and social requirements of sharing computing resources.

OSG is a consortium of software, service and resource providers and researchers, from universities, national laboratories and computing centers across the U.S., who together build and operate the OSG project. The project is funded by the NSF and DOE, and provides staff for managing various aspects of the OSG.

>50 research communities
>120 sites
>100,000 cores accessible

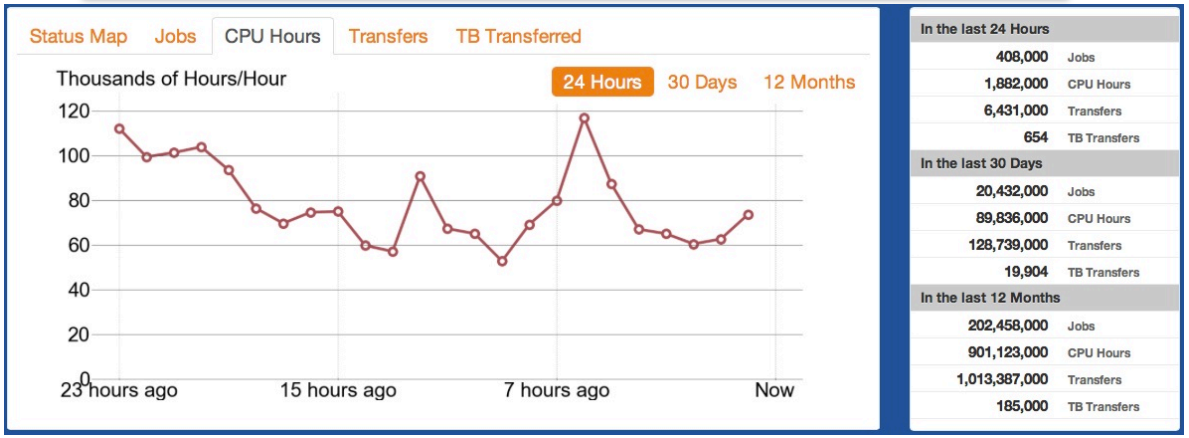

Virtual Organizations

- The OSG environment is VO based.
  - Resource usage accounting
  - VOs can be science communities (e.g. ATLAS, CMS) or
  - Multi-disciplinary Campus based [e.g. U-Nebraska(HCC), U-Wisconsin(GLOW)]
- Each OSG user is a member of a VO
- Users can be members of multiple VOs
- Site resources are owned by one or more VOs.

Virtual Organizations

- The **OSG VO** provides access to US researchers who are *not already affiliated* with an existing community in OSG.
- OSG VO is "Opportunistic" VO: users take advantage of unused cycles on resources owned by others.
- Opportunistic VOs: OSG, GLOW, Engage, HCC, SBGrid, Gluex

OSG Usage

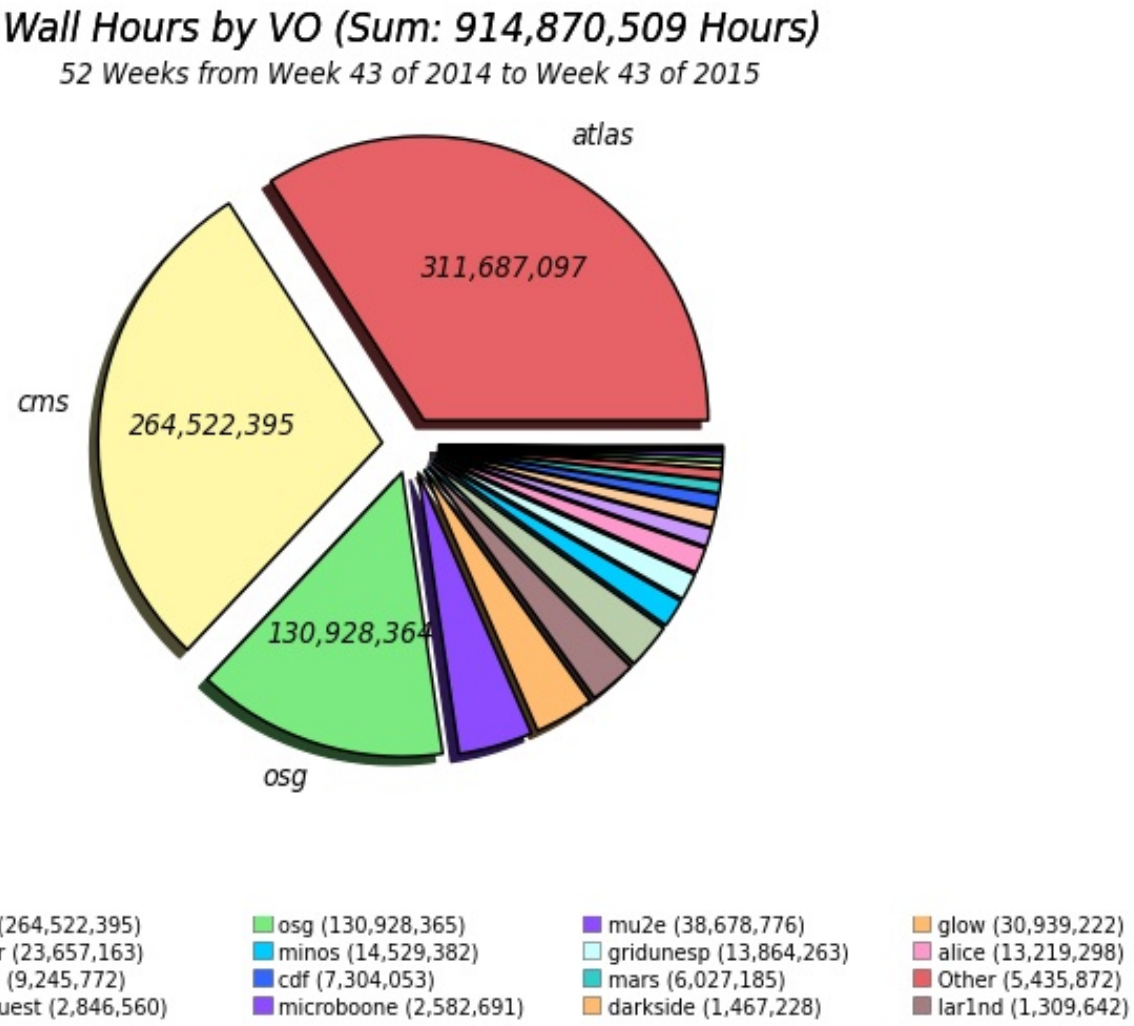
Wall Hours by VO (Sum: 914,870,509 Hours)
52 Weeks from Week 43 of 2014 to Week 43 of 2015
atlas
311,687,097
cms
264,522,395
130,928,364
osg
atlas (311,687,097)
cms (264,522,395)
osg (130,928,365)
mu2e (38,678,776)
glow (30,939,222)
nova (24,436,003)
dosar (23,657,163)
minos (14,529,382)
gridunesp (13,864,263)
alice (13,219,298)
minerva (9,319,821)
dune (9,245,772)
cdf (7,304,053)
mars (6,027,185)
Other (5,435,872)
sbgrid (2,869,721)
seaquest (2,846,560)
microboone (2,582,691)
darkside (1,467,228)
lar1nd (1,309,642)

OSG Opportunistic Use

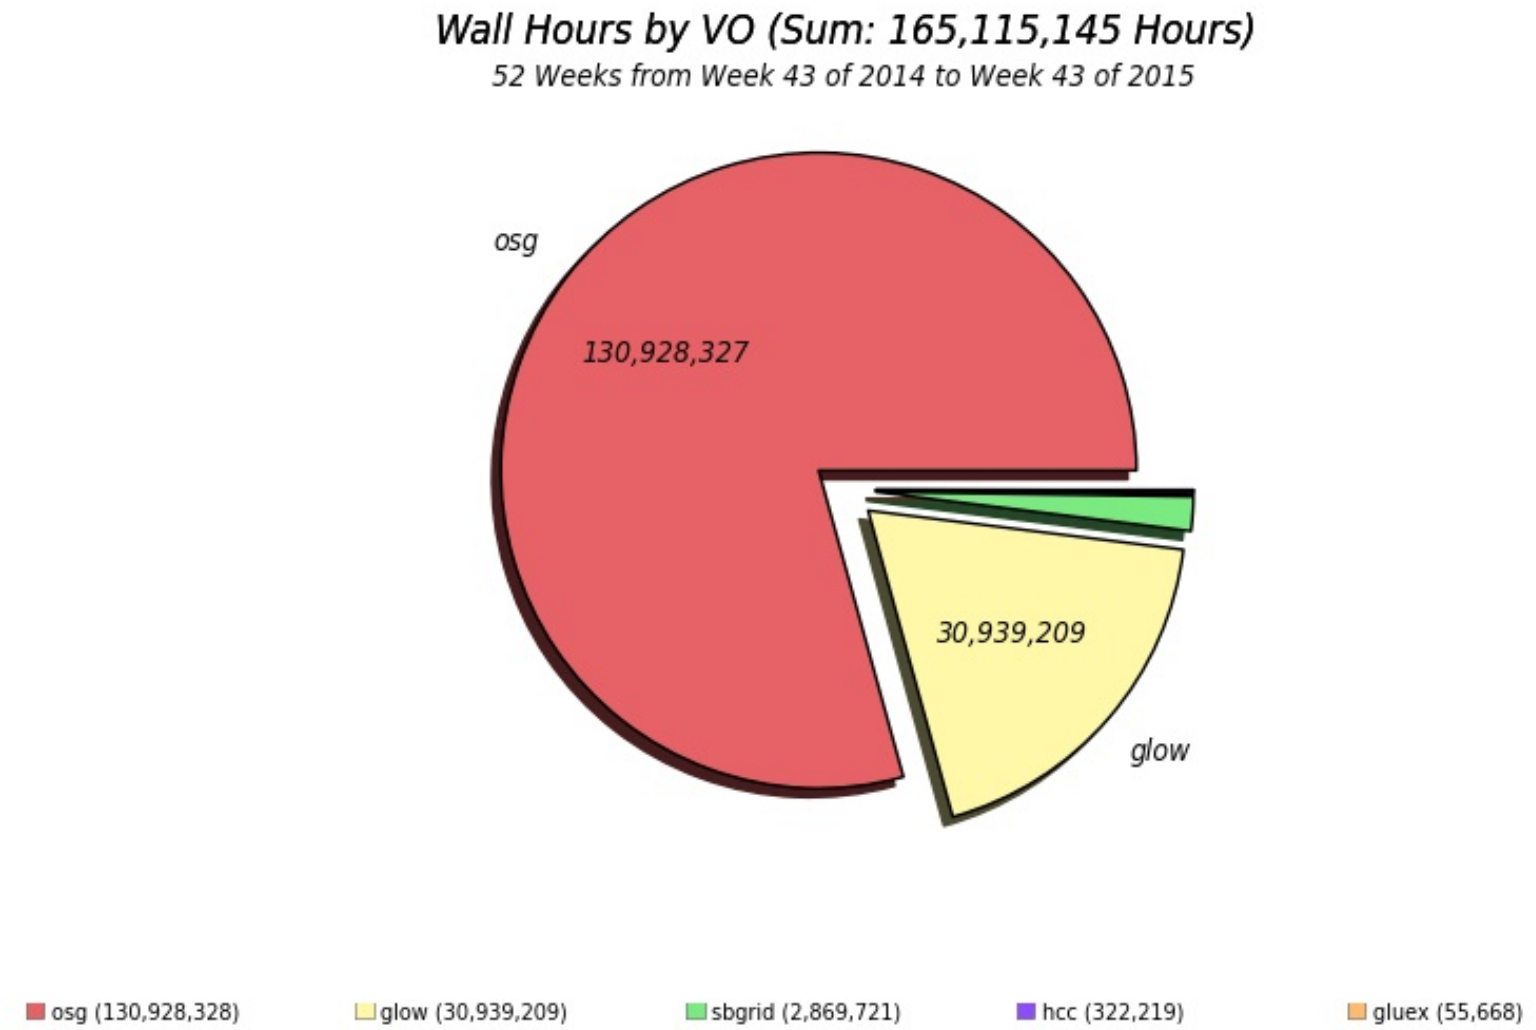

Who is Using the OSG?

- Astrophysics
- Biochemistry
- Bioinformatics
- Earthquake Engineering
- Genetics
- Gravitational-wave physics
- Mathematics
- Nanotechnology
- Nuclear and particle physics
- Text mining
- And more…

ATLAS Detector
Copyright CERN
Permission Information

SDSS Telescope
Image Credit Fermilab
Permission Information

CDMS photo
Image Credit Fermilab
Permission Information

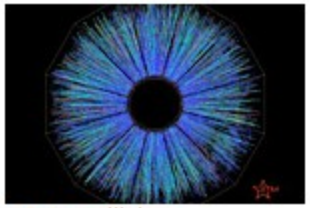

STAR Collision
Image Credit Brookhaven National Laboratory/STAR Collaboration
Permission Information

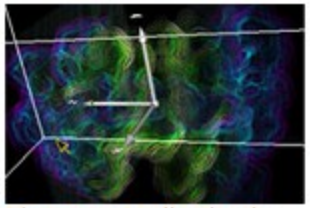

BioMOCA Application in nanoHUB
Image Credit Shawn Rice, Purdue University
Permission Information

CMS Detector
Copyright CERN
Permission Information

Auger photo
Image Credit Pierre Auger Observatory
Permission Information

MiniBooNE photo
Image Credit Fermilab
Permission Information

DZero Detector
Image Credit Fermilab
Permission Information

OSG Entry Paths

There are three simple ways:

- OSG Connect
  - Connects directly to the OSG
- Duke CI-Connect
  - Connect to OSG, Duke Compute
  - Cluster (DCC), plus the “UC3” cluster at Uchicago
- Connect Client
  - Use Duke Ci-Connect from the DCC login server

Duke Ci-Connect

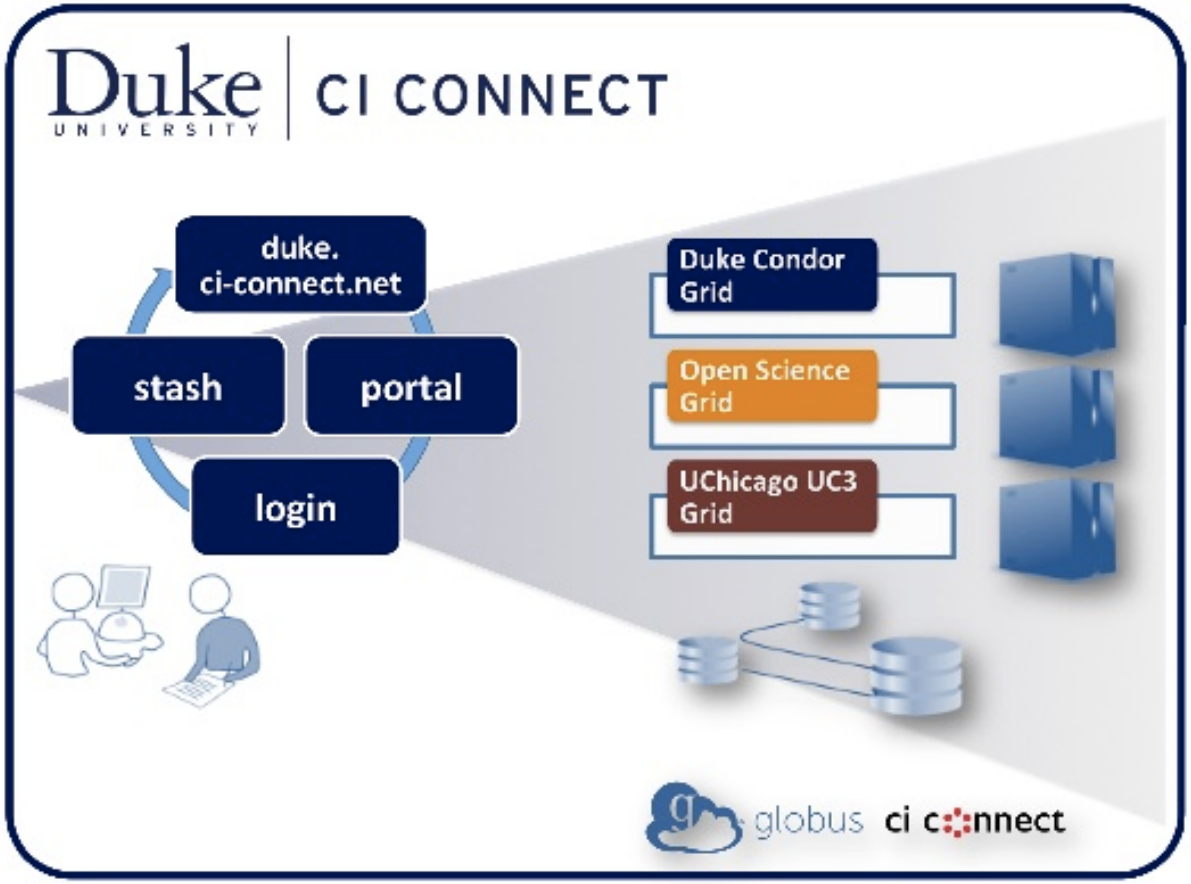

Special version of OSG Connect specifically for Duke users and resources.

Use the Duke Compute Cluster, OSG and UC3 together.

OSG Jobs

- **High Throughput Computing**
  - Sustained computing over long periods of time. Usually serial codes, or low number of cores threaded/MPI.

  vs. High Performance Computing
  - Great performance over relative short periods of time. Large scale MPI.

- ***Distributed* HTC**
  - No shared file system
  - Users ship input files and (some) software packages with their jobs.

- **Opportunistic Use**
  - Applications (esp. with long run times) can be *preempted* (or killed) by resource owner's jobs.
  - Applications should be relatively short or support being restarted.

OSG Jobs

- **High Throughput Computing**
  - Sustained computing over long periods of time. Usually serial codes, or low number of cores treaded/MPI.

  vs. High Performance Computing
  - Great performance over relative short periods of time. Large scale MPI.

- ***Distributed* HTC**
  - No shared file system
  - Users ship input files and (some) software packages with their jobs.

- **Opportunistic Use**
  - Applications (esp. with long run times) can be *preempted* (or killed) by resource owner's jobs.
  - Applications should be relatively short or support being restarted.

OSG Jobs

- **High Throughput Computing**
  - Sustained computing over long periods of time. Usually serial codes, or low number of cores treaded/MPI.

  vs. High Performance Computing
  - Great performance over relative short periods of time. Large scale MPI.

- ***Distributed* HTC**
  - No shared file system
  - Users ship input files and (some) software packages with their jobs.

- **Opportunistic Use**
  - Applications (esp. with long run times) can be *preempted* (or killed) by resource owner's jobs.
  - Applications should be relatively short or support being restarted.

OSG Jobs

- Run-time: 1-12 hours
- Single-threaded
- Require <2 GB Ram
- Statically compiled executables (transferred with jobs)

OSG Jobs

- Run-time: 1-12 hours
- Single-threaded
- Require <2 GB Ram
- Statically compiled executables (transferred with jobs)

These are not hard limits!

- Checkpointing – for long jobs that are preempted
  - Many applications support built-in checkpointing
  - Job image is saved periodically so that it can be restarted on a new host after it is killed (without losing the progress that was made on the first host)
- Limited support for larger memory jobs
- "Partitionable" slots – for parallel applications using up to 8 cores
- OASIS modules – a collection of pre-installed software packages

Submitting Jobs

OSG Job Scheduler: HTCondor

Basic job parts:

- Condor Submit Script
- Executable and (optionally) job wrapper
- Input files and helper scripts
- Statically linked software packages and libraries
- DAG (optional – for managing workflow)

Submit script: Text file; Provides information about the job, specific requirements (memory, etc.), executable, input arguments, file names: input, output, log. Many other optional arguments/directives.

Steps:

Copy job files, scripts, and software to login node (or create on login node)

ssh to login node:
```
$ ssh login.osgconnect.net
```
Submit the job(s)
```
$ condor_submit helloworld.submit
```
Monitor job:
```
$ condor_q <userid>
```

Learn more about HTCondor: HTCondor Manual

Job Wrappers

- Optional, but can be useful
- Wrap the science code in bash/perl/python/…
  - Job success check – surprisingly many science codes do not follow zero/non-zero exit code convention
  - Data staging
  - Tar / select executable / check dependencies / …
  - Load Modules (preinstalled software packages)

User Tools

- **HTCondor**
  - Job Scheduler used by the OSG
- **Dagman/Pegasus**
  - Workflow managing tools

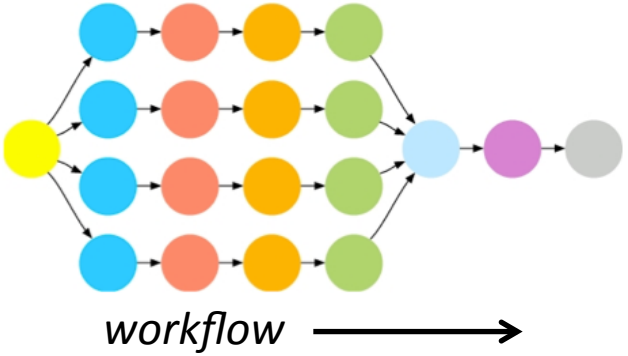

- **Stash**
  - Temporary storage space for job staging (larger I/O jobs).
  - Available through Globus Online, http, OSG/Duke Ci-Connect login nodes and worker nodes (via Parrot/Chirp).

User Tools

- **Globus**
  - OSG Connect portal to Globus Online for managing projects and data transfers (access to stash endpoint)
- **OASIS Modules**
  - A set of pre-installed software packages (more next slide)
- **Connect Client**
  - Submit jobs to OSG directly from your home institution.
- **Tutorials**
  - Hands-on learning for different science applications

User Tools

For more information and tutorials, see:
https://
support.opensciencegrid.org

- **Globus**
  - OSG C… …ects and data
- **OASIS …**
  - A set of pre-installed software packages (more next slide)
- **Connect Client**
  - Submit jobs to OSG directly from your home institution.
- **Tutorials**
  - Hands-on learning for different science applications

User Support

Helpdesk:
https://
support.opensciencegrid.org

- Knowledge Base
  - User guides/tutorials
  - HTC Recipes
- Forums
- "How do I…?" articles
- Interactive online chat

Support email:
user-support@opensciencegrid.org

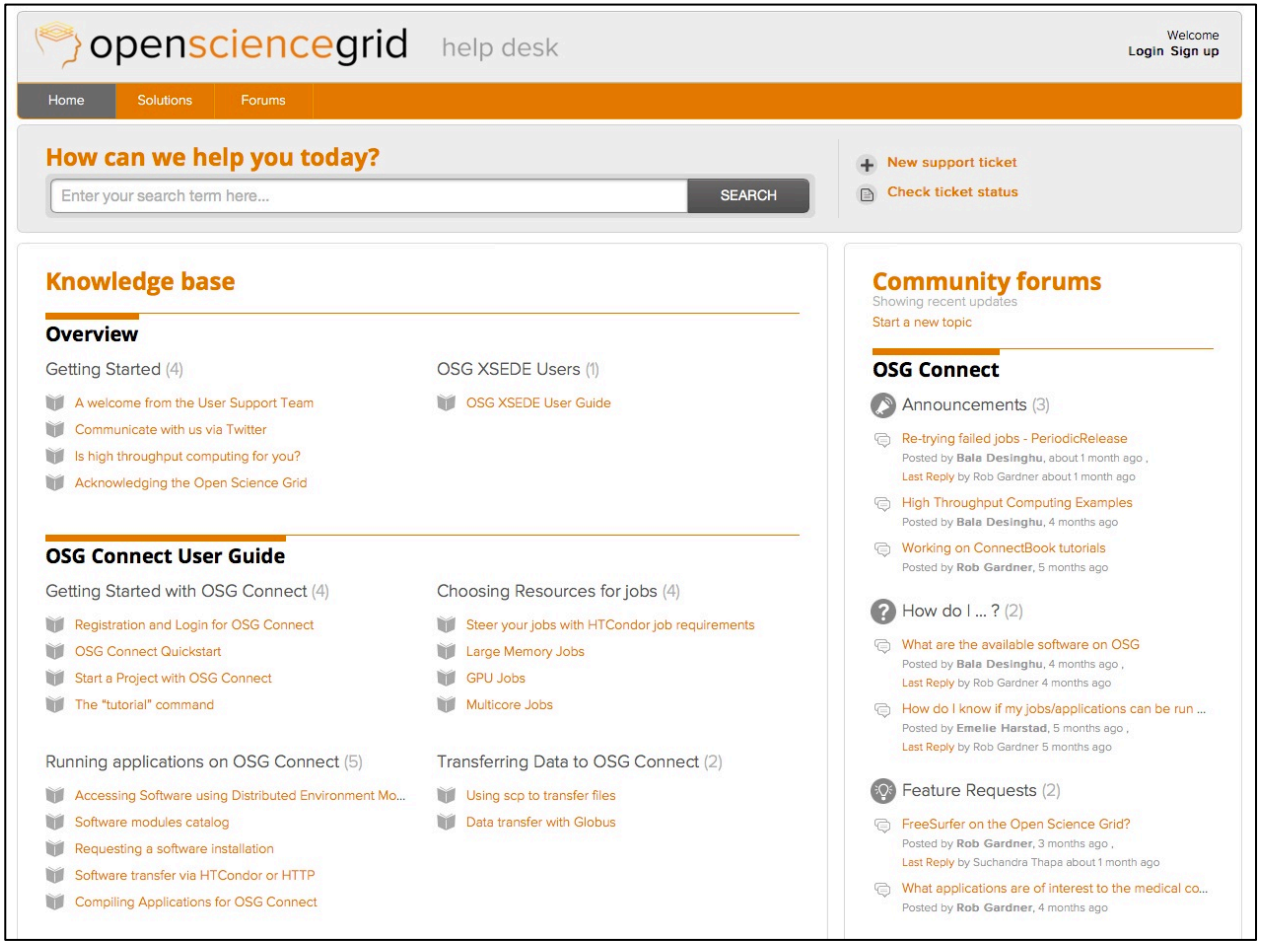

Modules Environment

- Popular (and some requested) software packages and libraries are made available via OASIS repository
- Users don't have to transfer software with their jobs
- Modules address ease of use issues
- Many users are already familiar with modules environment
- Example: `module load python/3.4`