
A Tour of Services Provided by the Open Science Grid

Jeff Dost (UC San Diego)



Collaborating E-Infrastructures 2018

Service Categories



Open Science Grid

- Software repositories
- User submit infrastructure
- User data management
- Hosted CEs
- Resource registration
- Accounting / monitoring
- Other services

Software repositories



Open Science Grid

- Yum based repos for
 - site grid middleware
 - user client tools
 - OSG Service infrastructure
- Packaging, testing, releasing done by UW Madison Software team
- Koji instance hosted @ UW as rpm build platform
- Jira ticketing system for internal issue tracking

User submit infrastructure



Open Science Grid

- Software repositories
- **User submit infrastructure**
- User data management
- Hosted CEs
- Resource registration
- Accounting / monitoring
- Other services

OSG Connect



Open Science Grid

- A collection of services for end users hosted at UChicago
- Web portal used to register / join project to submit jobs to OSG
 - contains links to documentation, support, job monitoring, globus data transfer interface
 - Users log in using local campus CILogon authentication
- Login host which contains the job scheduler and end user command line tools
 - Job scheduler connected to OSG Flock HTCondor pool
 - credentials managed via portal
- StashCache origin server for exporting user data to the grid

GlideinWMS Factory



Open Science Grid

- GlideinWMS is a pilot based system containing distributed components
- The OSG Glidein Factory @ UCSD submits Pilots (aka glideins) based on user demand
 - pilots reserve machines across the grid matching user requirements
- Glidein Frontends connect to the factory to request pilots on user's behalf
 - Different frontends are owned and operated by different Virtual Organizations
- Each frontend creates a VO HTCondor Pool overlay across sites
 - Independent of underlying site batch system
- End users submit jobs to VO HTCondor pool, which run on machines that have been reserved by pilots and have joined the pool

OSG Flock



Open Science Grid

- GlideinWMS Frontend maintained for the OSG VO
- OSG VO is a catch all VO for projects that aren't large enough to justify starting their own VO but would like to run on the OSG
- OSG Flock maintains the HTCondor pool backend for the OSG Connect infrastructure
- The scheduler on the OSG Connect login is just a condor schedd connected to the Flock pool

User data management



Open Science Grid

- Software repositories
- User submit infrastructure
- **User data management**
- Hosted CEs
- Resource registration
- Accounting / monitoring
- Other services

OASIS



Open Science Grid

- CVMFS based storage for applications, optimized for GB sized data sets
- allows data access over http
- CVMFS takes advantage of Squid caching at sites
- shows up as FUSE mounts on WNs across the grid as
 - /cvmfs/oasis.opensciencegrid.org
- OSG provides 2 options:
 - direct access to OASIS stratum 0 host to publish data
 - VOs that already have a stratum 0 can replicate to OASIS stratum 1 making repo available at OSG sites
- OASIS stratum 0 and 1 are hosted at UNL

StashCache



Open Science Grid

- XRootD read-only store for data sets GB to TB range
- OSG Connect origin server based at UChicago
- 9 XrootD Caches deployed across the US for improved performance
- posix file namespace accessed via CVMFS @ stash.osgstorage.org
 - GeoIP locates closest cache to where user job is running
- XRootD redirector is hosted at UNL

Hosted CEs



Open Science Grid

- Solution for sites who don't have the effort / expertise to run their own HTCondor Compute Element and grid middleware
- OSG runs a remote CE hosted in VM at UChicago
- Connects to site batch via BOSCO (HTCondor over ssh)
- Site only needs to supply an ssh login to submit host, OSG takes care of the rest

Resource registration



Open Science Grid

- Software repositories
- User submit infrastructure
- User data management
- Hosted CEs
- **Resource registration**
- Accounting / monitoring
- Other services

Topology (replaced OIM)



Open Science Grid

- OSG registry of site resources (CE, Storage), resource downtimes, VOs, projects
- Contains contact info, config info for accounting (GRACC), resource availability
- YAML based and stored in OSG GitHub
 - Admins can make changes by pull request
- Topology information is published in xml format
 - <https://my.opensciencegrid.org/>
 - WLCG parses this for resource information

CE Collector



Open Science Grid

- HTCondor CEs on OSG advertise attributes about their resource constraints (max walltime, memory, cores, allowed VOs)
- Is just a HTCondor collector daemon hosted @ UNL, which stores key = value attributes in classads for each CE
 - can be queried with `condor_status`
- Information is queried by pilot factories and used to submit properly configured pilots to sites
- Replaced BDII as central information system for site clusters on OSG

Accounting / monitoring



Open Science Grid

- Software repositories
- User submit infrastructure
- User data management
- Hosted CEs
- Resource registration
- **Accounting / monitoring**
- Other services

GRACC



Open Science Grid

- accounting for grid resource usage including walltime hrs, cpu hrs, data transfers
- Usage stats broken down per-site, VO, project, pilot, payload
- Elasticsearch backend, Grafana and Kibana frontends, all hosted @ UNL
- Probes gather stats at various levels, send to Elasticsearch:
 - Site – CEs, gridftp endpoints
 - VO – user submit hosts

Network Data Pipeline



Open Science Grid

- network monitoring is collected via perfSONAR instances across the grid
- network metrics are stored in Nebraska Elasticsearch and also sent to CERN
- MaDDash monitoring page to display network connections between perfSONAR instances hosted at UMich/AGLT2
- metrics include: latency, packet loss, bandwidth, traceroute

Other Services



Open Science Grid

- OSG Documentation hosted on GitHub Pages:
 - <http://opensciencegrid.org/docs/>
- OSG Helpdesk ticketing system powered by Freshdesk:
 - <https://support.opensciencegrid.org/support/home>
 - Initially only for OSG Connect support
 - support extended to all OSG services, and site support
 - * GGUS is preferred for submitting WLCG site tickets



Discussion