



Open Science Grid

Increasing the Number of Opportunistic Resources for VOs

Mats Rynge
OSG User Support



Purpose

- Opportunistic VOs looking for cycles
 - 51 VOs / 8 new
 - Can take a long time to get new VOs accepted on enough sites
- Why we ask sites to support more VOs
 - no automatically support for VOs
- What can OSG do better?



In a perfect world...

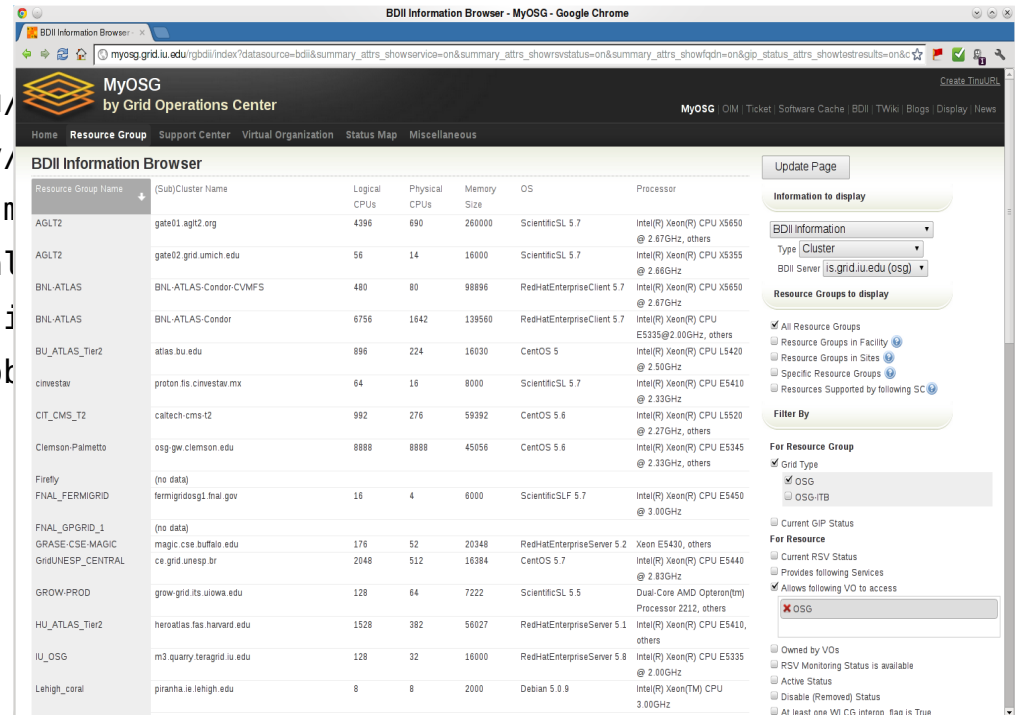
- 1) New VO is created
- 2) The VO is added to the VO package
- 3) Sites install the VO package
- 4) Sites set VO as supported to GUMS or edgmkgridmap
- 5) Sites add accounts / directories / priorities / ...
- 6) Support is advertised in BDII/ReSS
- 7) The VO finds out where they are supported
- 8) The VO verifies the sites
- 9) Verified sites are added to workload management system
- 10) Jobs start flowing!

Which sites support the my VO?

- Query ReSS or BDII

```
condor_status -pool osg-ress-1.fnal.gov \
  -constraint 'StringlistIMember("VO:osg";GlueCEAccessControlBaseRule)' \
  -format '%-20s' GlueSiteUniqueID -format '%s\n' GlueCEInfoContactString \
  | sort | uniq
```

AGLT2	gate02.grid.umich.edu/
BNL-ATLAS	gridgk01.racf.bnl.gov/
cinvestav	proton.fis.cinvestav.m
CIT_CMS_T2	cit-gatekeeper2.ultra
CIT_CMS_T2	cit-gatekeeper.ultra
FNAL_FERMIGRID	d0cabosg1.fnal.gov/jok



The screenshot shows the 'BDII Information Browser - MyOSG - Google Chrome' interface. The main content is a table with columns: Resource Group Name, (Sub)Cluster Name, Logical CPUs, Physical CPUs, Memory Size, OS, and Processor. The table lists various resource groups like AGLT2, BNL-ATLAS, BU_ATLAS_Tier2, etc., with their respective specifications.

Resource Group Name	(Sub)Cluster Name	Logical CPUs	Physical CPUs	Memory Size	OS	Processor
AGLT2	gate01.agl2.org	4396	690	260000	ScientificSL 5.7	Intel(R) Xeon(R) CPU X5650 @ 2.87GHz, others
AGLT2	gate02.grid.umich.edu	56	14	16000	ScientificSL 5.7	Intel(R) Xeon(R) CPU X5355 @ 2.66GHz
BNL-ATLAS	BNL-ATLAS-Condor-CVMS	480	80	98896	RedHatEnterpriseClient 5.7	Intel(R) Xeon(R) CPU X5650 @ 2.87GHz
BNL-ATLAS	BNL-ATLAS-Condor	6756	1842	139560	RedHatEnterpriseClient 5.7	Intel(R) Xeon(R) CPU E5335@2.00GHz, others
BU_ATLAS_Tier2	atlas.bu.edu	896	224	16030	CentOS 5	Intel(R) Xeon(R) CPU L5420 @ 2.50GHz
cinvestav	proton.fis.cinvestav.mx	64	16	8000	ScientificSL 5.7	Intel(R) Xeon(R) CPU E5410 @ 2.33GHz
CIT_CMS_T2	caltech-cms-t2	992	276	59392	CentOS 5.6	Intel(R) Xeon(R) CPU L5520 @ 2.27GHz, others
Clemson-Palmetto	osg-gw.clemson.edu	8888	8888	45056	CentOS 5.6	Intel(R) Xeon(R) CPU E5345 @ 2.33GHz, others
Firstly	(no data)					
FNAL_FERMIGRID	fermigridosg1.fnal.gov	16	4	6000	ScientificSLF 5.7	Intel(R) Xeon(R) CPU E5450 @ 3.00GHz
FNAL_GPGRID_1	(no data)					
MAGIC_CSE-MAGIC	magic.cse.buffalo.edu	176	52	20348	RedHatEnterpriseServer 5.2	Xeon E5430, others
GIUNESP_CENTRAL	ce.grid.unesp.br	2048	512	16384	CentOS 5.7	Intel(R) Xeon(R) CPU E5440 @ 2.83GHz
GROW-PROD	grow.grid.its.uiowa.edu	128	64	7222	ScientificSL 5.5	Dual-Core AMD Opteron(tm) Processor 2212, others
HU_ATLAS_Tier2	herodias.fas.harvard.edu	1528	382	56027	RedHatEnterpriseServer 5.1	Intel(R) Xeon(R) CPU E5410, others
IU_OSG	m3.quarry.teragrid.iu.edu	128	32	16000	RedHatEnterpriseServer 5.8	Intel(R) Xeon(R) CPU E5335 @ 2.00GHz
Lehigh_coral	piranha.ie.lehigh.edu	8	8	2000	Debian 5.0.9	Intel(R) Xeon(TM) CPU 3.00GHz

Compare list of potential VOs with list of sites the VO is already running on

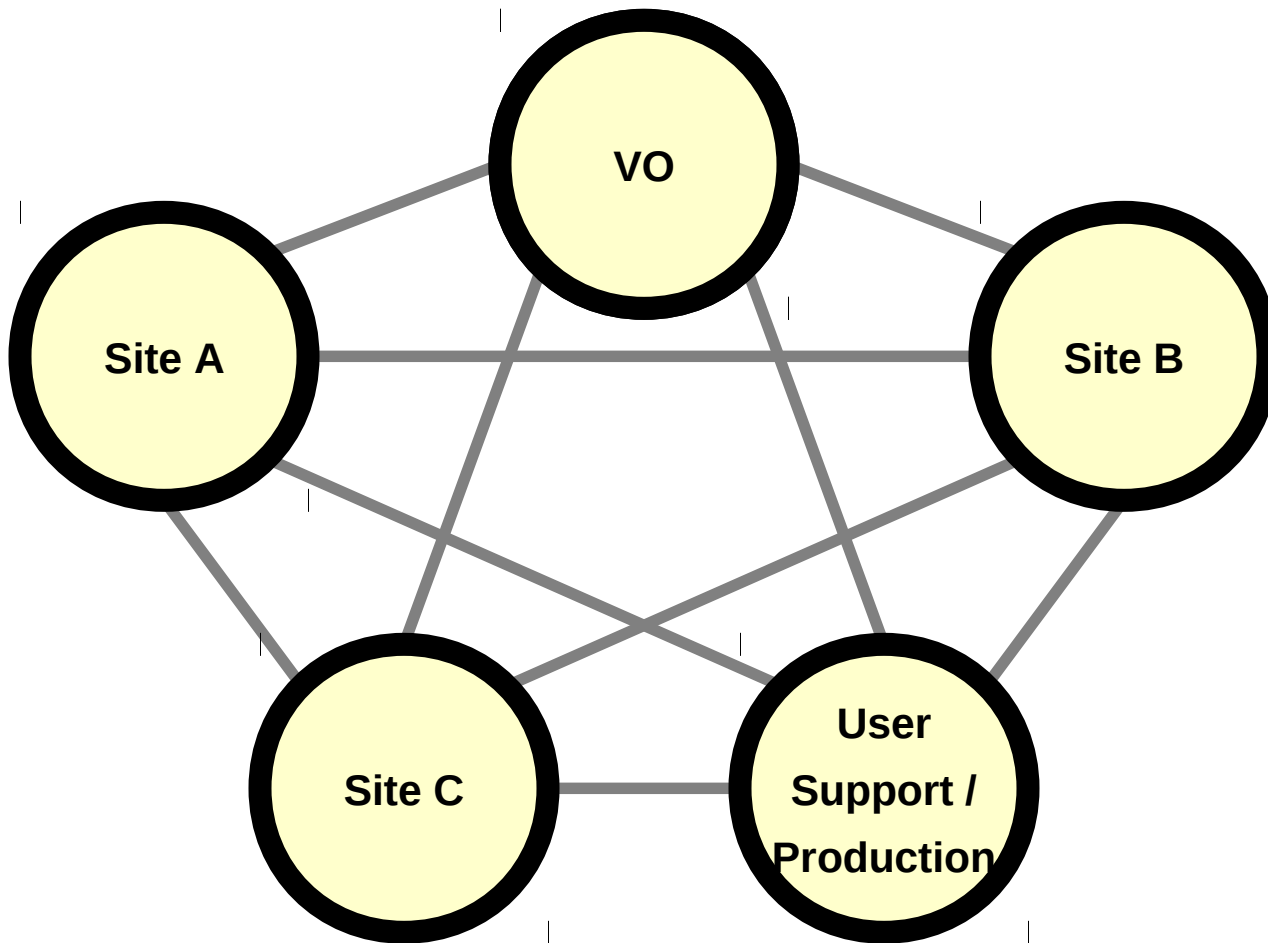
```
condor_status -any \  
    -pool glidein-1.t2.ucsd.edu \  
    -constraints 'stringListMember("OSGVO",Glidein_Supported_VO)' \  
    -format '%s\n' GLIDEIN_ResourceName \  
| sort | uniq
```

```
AGLT2  
cinvestav  
Firefly  
FNAL_FERMIGRID  
GridUNESP_CENTRAL  
GROW-PROD  
MWT2_IU  
...
```

Sites advertise support for my VO is supported, but do they work?

- Advertised support is **not** verified by the infrastructure
- Good idea to at least run simple globus-job-run jobs to verify the site before adding it to workload management system

What about sites which do not support my VO – can I convince them to support us?





Open Science Grid

User Support / VO Forum

The goal of User Support is to enable new communities to quickly adopt the OSG DHTC model and to improve productivity for all VOs as OSG service capabilities evolve.

<https://twiki.grid.iu.edu/bin/view/VirtualOrganizations/WebHome>

user-support@opensciencegrid.org

Call every other Tuesday at 1pm central



Production Calls

OSG Production is the culmination of many OSG efforts including testing, integration, services, support, operations, hardware, software, and security -- VO's get their jobs run, their data analyzed, and in so doing enable the advancement of science.

<https://twiki.grid.iu.edu/bin/view/Production/WebHome>

Call every Tuesdays 3:00 Central



Wishlist

- Sites to support more VOs by default
 - Use priorities between VOs
- More automatic everything
 - Discovery / job flow
 - VO propagation
 - Site setups