

# VO Infrastructure

## **Application Discovery and Launch**

Steve Johnson  
Texas A&M University  
[steve@math.tamu.edu](mailto:steve@math.tamu.edu)



# Overview

- Where are we now?
- Where do we want to be?
- Info for Resource Providers.
- Running Applications using SG Apps info page.
  - Preliminary tasks
  - Traditional Globus
  - Condor-G
- Moving forward.
- A challenge.

# Where are we now?

<http://www.math.tamu.edu/osg/sgstatus.php>

Resource Group	Globus Resource	Free Slots	Max Jobs	Walltime (minutes)
TTU-ANTAEUS	antaeus.hpcc.ttu.edu:2119/jobmanager-sge-suragrid	208	208	9999999
TTU-ANTAEUS	antaeus.hpcc.ttu.edu:2119/jobmanager-sge-serial	46	432	9999999
TAMU_Calclab	calclab-ce.math.tamu.edu:2119/jobmanager-pbs-weekend	0	894	2520
TAMU_Calclab	calclab-ce.math.tamu.edu:2119/jobmanager-pbs-night	0	894	720
GridUNESP_CENTRAL	ce.grid.unesp.br:2119/jobmanager-pbs-medium	224	1992	1440
FNAL_FERMIGRID	fermigridosg1.fnal.gov:2119/jobmanager-condor-default	16	18432	1440
FNAL_FERMIGRID_ITB	fgitbgkp2.fnal.gov:2119/jobmanager-pbs-batch	2	2	1440
SWT2_CPB	gk01.atlas-swt2.org:2119/jobmanager-pbs-input_q	11	1388	1440
UTA_SWT2	gk04.swt2.uta.edu:2119/jobmanager-pbs-input_q	2	1144	1440
Hampton PPCF ATLAS	hugrid02.hamptonu.edu:2119/jobmanager-condor-default	40	200	1440
TAMU_BRAZOS	hurr.tamu.edu:2119/jobmanager-pbs-grid	256	256	480
UGA_SGrid	schuster.rcc.uga.edu:2119/jobmanager-sge-suragrid	4	4	86400
UGA_SGrid	sgrid-test.rcc.uga.edu:2119/jobmanager-sge-all_q	2	2	9999999

# Where are we now?

- 8 distinct resources on Compute Element Resource Status Page, half from traditional SURAGRID sites (*italic* below)

FNAL_FERMIGRID	GridUNESP_CENTRAL	Hampton PPCF ATLAS
Hampton PPCF ATLAS	<i>TAMU_BRAZOS</i>	<i>TAMU_Calclab</i>
<i>TTU_ANTAEUS</i>	<i>UGA_SGrid</i>	

# Where are we now?

<http://www.math.tamu.edu/osg/sgapps.php>

Application	Version	Arch	OS	Path	Resource Group	Globus Resource
<a href="#">gcc</a>	4.5.0	x86_64	sles11	/usr/bin/gcc	TAMU_Calclab	calclab-ce.math.tamu.edu:2119/jobmanager-pbs-weekend calclab-ce.math.tamu.edu:2119/jobmanager-pbs-night
<a href="#">gcc</a>	4.1.2	x86_64	rhel5	/usr/bin/gcc	UGA_SGrid	schuster.rcc.uga.edu:2119/jobmanager-sge-suragrid sgrid-test.rcc.uga.edu:2119/jobmanager-sge-all.q
<a href="#">gcc</a>	4.1.2-51.el5	x86_64	centos	/usr/bin/gcc	UGA_SGrid	schuster.rcc.uga.edu:2119/jobmanager-sge-suragrid sgrid-test.rcc.uga.edu:2119/jobmanager-sge-all.q
<a href="#">R</a>	2.13.0	x86_64	sles11	/usr/bin/R	TAMU_Calclab	calclab-ce.math.tamu.edu:2119/jobmanager-pbs-weekend calclab-ce.math.tamu.edu:2119/jobmanager-pbs-night
<a href="#">R</a>	2.12.2	x86_64	rhel5	module load r/2.12.2/intel/64	TAMU_BRAZOS	hurr.tamu.edu:2119/jobmanager-pbs-grid
<a href="#">octave</a>	3.4.2	x86_64	sles11	/usr/bin/octave	TAMU_Calclab	calclab-ce.math.tamu.edu:2119/jobmanager-pbs-weekend calclab-ce.math.tamu.edu:2119/jobmanager-pbs-night
<a href="#">octave</a>	3.2.4	x86_64	rhel5	module load octave/3.2.4/gcc/64	TAMU_BRAZOS	hurr.tamu.edu:2119/jobmanager-pbs-grid
<a href="#">openmpi</a>	1.4.5	x86_64	sles11	module load openmpi/1.4.5	TAMU_Calclab	calclab-ce.math.tamu.edu:2119/jobmanager-pbs-weekend calclab-ce.math.tamu.edu:2119/jobmanager-pbs-night
<a href="#">openmpi</a>	1.4.5	x86_64	rhel5	module load openmpi/1.4.5/gcc/64	TAMU_BRAZOS	hurr.tamu.edu:2119/jobmanager-pbs-grid
<a href="#">poy</a>	4.1.1	x86_64	rhel5	module load poy/4.1.1/gcc/64	TAMU_BRAZOS	hurr.tamu.edu:2119/jobmanager-pbs-grid
<a href="#">maxima</a>	5.21.0	x86_64	rhel5	module load maxima/5.21.0/gcc/64	TAMU_BRAZOS	hurr.tamu.edu:2119/jobmanager-pbs-grid
<a href="#">singular</a>	3.1.3	x86_64	rhel5	module load singular/3-1-3	TAMU_BRAZOS	hurr.tamu.edu:2119/jobmanager-pbs-grid
<a href="#">garli</a>	0.96b8	x86_64	rhel5	module load garli/0.96b8/gcc/64	TAMU_BRAZOS	hurr.tamu.edu:2119/jobmanager-pbs-grid
<a href="#">mrbayes</a>	3.2.1	x86_64	rhel5	module load mrbayes/3.2.1/gcc/64	TAMU_BRAZOS	hurr.tamu.edu:2119/jobmanager-pbs-grid
<a href="#">qhull</a>	2012.1	x86_64	rhel5	module load qhull/2012.1/gcc/64	TAMU_BRAZOS	hurr.tamu.edu:2119/jobmanager-pbs-grid
<a href="#">gromacs</a>	4.5.5	x86_64	rhel5	module load gromacs/4.5.5/pgi/10/64	TAMU_BRAZOS	hurr.tamu.edu:2119/jobmanager-pbs-grid
<a href="#">raxml</a>	7.0.4	x86_64	rhel5	module load raxml/7.0.4/gcc/64	TAMU_BRAZOS	hurr.tamu.edu:2119/jobmanager-pbs-grid
<a href="#">uptime</a>	3.2.7	x86_64	rhel5	/usr/bin/uptime	UGA_SGrid	schuster.rcc.uga.edu:2119/jobmanager-sge-suragrid sgrid-test.rcc.uga.edu:2119/jobmanager-sge-all.q
<a href="#">date</a>	1.0	x86_64	centos	/bin/date	UGA_SGrid	schuster.rcc.uga.edu:2119/jobmanager-sge-suragrid sgrid-test.rcc.uga.edu:2119/jobmanager-sge-all.q

# Where are we now?

- 11 real applications on Application Availability Page  
<http://www.math.tamu.edu/osg/sgapps.php>

garli	gromacs	maxima	mrbayes
octave	openmpi	poy	qhull
R	raxml	singular	

# Where do we want to be?

- Enable more SG contributed resources to run OSG stack.
- Define useful set of applications installed on these resources.
- Document best practices for each application.
- Move resource status and applications to well known location under [suragrid.org](http://suragrid.org).
- Bring on the users!

# For Resource Providers

- Install and configure the OSG stack.
- Register your resource with OSG and it will appear in OSG BDII and SG status page.  
*It's that simple.*
- Define your applications in  
`app_dir/etc/grid3-locations.txt`  
where `app_dir` is defined in  
`$VDT_LOCATION/osg/etc/config.ini`

## Example:

```
app_dir = /net/apps/osg
```



# For Resource Providers

- **Format of the grid3-locations.txt fields**

```
# Name                                     Version      Path
VO-suragrid-sles11-x86_64-R               2.13.0       /usr/bin/R
VO-suragrid-sles11-x86_64-openmpi        1.4.5        \
                                           /module%20load%20openmpi/1.4.5
```

- **Name:** VO-suragrid-sles11-x86\_64-R
  - VO identifier for our apps: VO-suragrid
  - Operating system: rhel5, rhel6, sles10, sles11
    - Keep it simple: centos = sl = rhel; opensuse = sles
  - Architecture: x86\_64, i386, power5
  - Application Name: R, octave, gromacs
- **Version:** 2.13.0
  - Just the version of the application

# For Resource Providers

- Format of grid3-locations.txt (cont'd)
  - Path
    - Path to the executable: `/usr/bin/R`
    - or –
    - Module command: `/module%20load%20openmpi/1.4.5`
      - LDAP schema for Path requires leading '/' and no spaces (replace with '%20')
      - Scripts can rewrite this to  
`module load openmpi/1.4.5`
      - Module command on resource must load any prerequisite modules
- With these definitions your applications will be published to BDII and will appear on the sgapps page.

# Running Applications

## Preliminary Tasks

- Install the OSG client on your submit host  
<https://twiki.grid.iu.edu/bin/view/ReleaseDocumentation/ClientInstallationGuide>
- Store your grid certificate in  
`$HOME/.globus/usercert.pem` and `userkey.pem`
- Setup the environment  
`source $VDT_LOCATION/setup.sh`
- Grab a VOMS proxy  
`voms-proxy-init -voms suragrid`
- Read on and run your application
- Destroy your VOMS proxy before logging out  
`voms-proxy-destroy`

# Application Detail View

<http://www.math.tamu.edu/osg/sgapps.php?details>

**R**

Resource Group: **TAMU\_BRAZOS**

Globus Resources:

hurr.tamu.edu:2119/jobmanager-pbs-grid

Operating System: rhel5

Architecture: x86\_64

Hostname: hurr.tamu.edu

Version: 2.12.2

Path: module load r/2.12.2/intel/64

This provides information for running R on the resource

- Resource is named TAMU\_BRAZOS
- CE hostname is in the Globus Resource: hurr.tamu.edu
- One PBS queues is available: grid
- OS is rhel5 (RedHat Enterprise Linux 5 or derivative)
- Architecture is x86\_64
- Version is 2.12.2
- Executable is found using shell modules: module load r/2.12.2/intel/64

# Running Applications Old School Globus

- Use information from sgapps page to construct a script, `rbench.sh`
- Use shell `case` statement on `$OSG_SITE_NAME` to set `PATH` and other vars for use on multiple resources.
- Copy script/input files to the remote resource using hostname from SG apps page

```
globus-url-copy file:///`pwd`/rbench.sh gsiftp://hurr.tamu.edu/~rbench.sh
globus-url-copy file:///`pwd`/R-benchmark25.R \
  gsiftp://hurr.tamu.edu/~R-benchmark25.R
globus-url-copy file:///`pwd`/Rlibs.tgz gsiftp://hurr.tamu.edu/~Rlibs.tgz
globus-job-run hurr.tamu.edu/jobmanager-fork /bin/tar zxf Rlibs.tgz
globus-job-run hurr.tamu.edu/jobmanager-fork /bin/chmod u+x rbench.sh
```
- Submit the job using resource name from SG apps page

```
globus-job-submit hurr.tamu.edu/jobmanager-pbs-grid rbench.sh
```

*Note the job's URL ID.*
- Check status

```
globus-job-status jobID
```


# Running Applications Old School Globus

- Copy output files back to your system  

```
globus-url-copy gsiftp://hurr.tamu.edu/~myoutput.txt \  
file:///`pwd`/myoutput.txt
```
- Get stdout/stderr  

```
globus-job-get-output -out jobID  
globus-job-get-output -err jobID
```
- Remove the job  

```
globus-job-clean jobID
```
- It's cumbersome, but it works and can be scripted.



# Running Applications Condor-G



- Can use the same script, `rbench.sh`, or something specific to the target resource
- Condor-G handles the staging of the script and input files.
- Libraries in `tgz/zip` formats may still need to be staged manually.
- Condor-G will keep a log file on your submit host and stages the `stdout/stderr` files.
- Can utilize Condor's internal variables.

# Scripts and Input Files

- Files located at <http://www.math.tamu.edu/osg/R/>
- rsiteprep.{condor,sh} – scripts for siteprep
- rbench.{condor,sh} – scripts to run benchmark
- R-benchmark-25.R – R input
- lattice\_0.19-33.tar.gz  
Matrix\_0.9996875-3.tar.gz  
SuppDists\_1.1-8.tar.gz – libraries



# Example Condor Script

```
# rbench.condor
Universe           = grid
Grid_Type          = gt2
Notification       = never
Executable         = rbench.sh
# unused
# arguments        = "--slave --quiet --file=lognormal.r --args $(Process)"
Transfer_Executable = True
Output             = rbench-$(Cluster)-$(Process).out
Error              = rbench-$(Cluster)-$(Process).err
Log                = rbench-$(Cluster)-$(Process).log
Should_Transfer_Files = YES
When_To_Transfer_Output = ON_EXIT
Transfer_Input_Files = R-benchmark-25.R
GlobusRSL          = (maxWalltime=60)

GlobusScheduler    = calclab-ce.math.tamu.edu/jobmanager-pbs
queue 8
GlobusScheduler    = hurr.tamu.edu/jobmanager-pbs
queue 8
GlobusScheduler    = antaeus.hpcc.ttu.edu/jobmanager-sge
queue 8
```



# Running Applications Condor-G

- Submit `rbench.condor` script from your submit host

```
condor_submit rbench.condor
```

*Note the Condor job ID.*

- Check job status and logfile

```
condor_q jobID
```

```
cat rbench-*-*.log
```

- At job completion check stdout/stderr files

```
more rbench-*-*.{out,err}
```

- Condor-G takes care of job cleanup

# Moving Forward

- Test resource availability, ala old SG portal
  - GRAM authentication
  - GRAM job submission, turn around time
  - GRIDFTP
  - Specific applications: R, octave, gromacs
- Evaluate other WMS options
  - GlideinWMS w/ Condor startd on compute nodes
    - Uses setuid glexec – may conflict with site policies
    - Requires Internet access on all nodes
    - Non-trivial installation and management
  - Future of OSGMM?



# A Challenge...



It's time for SURAGRID to turn the corner...



*I challenge each institution represented here today to have a new user obtain a grid certificate, join the SURAGrid VO, and run an application on at least one remote SURAGrid resource by the time of the Fall All-Hands Meeting.*

*Q?*