OSG Installation on the FLTECH Cluster - Summer 2008

Patrick Ford

This guide is a basic reference for installing and configuring the OSG software on the FLTECH cluster.
Below is a checklist of the minimum required steps to pass site-verification.

1. Install Pacman
2. Install Condor batch system
3. Install Compute Element packages
4. Host and User certificates
5. Monitoring
6. Run Configuration script
7. Start VDT services
8. Create user accounts for VOs
9. Update grid-mapfile
10. Update IPTables
11. Test fork/condor queue and GSIFTP
12. Run site_verify script

1. Pacman

To install Pacman we created a directory: /usr/local/pacman – and then ran the following commands:

```
# wget http://physics.bu.edu/pacman/sample_cache/tarballs/pacman-3.19.tar.gz
```

This can (and should) be substituted with the latest version of pacman, or OSG packages may not install properly.

```
# tar –no-same-owner –xzvf pacman-3.19.tar.gz
# cd pacman-3.19
# ./setup.sh
```

2. Condor

If you already have condor installed (as we did) you can skip this command, but be sure to set the environmental variables below.

```
# pacman –get http://www.cs.wisc.edu/vdt/vdt_161_cache:Condor
```

Set the following variables either manually or in /etc/profile

```
# export VDTSETUP_CONDOR_LOCATION=/opt/condor/
```

```
        # export VDTSETUP_CONDOR_CONFIG=$VDTSETUP_CONDOR_LOCATION/etc/
condor_config
```

3. OSG:CE

At the time of writing this, we have OSG installed and working. First create a directory for OSG, in our case it was /opt/osg

Set # umask 0022

Now install OSG services for the compute element.

```
        # pacman –get OSG:ce
        $ ./setup.sh
```

Make sure that the user has execution rights for the setup script. Also, you may want to add the script to /etc/profile as you will need to run it a lot during the setup of OSG. Now make sure that the condor environmental variables have set correctly and run the following command:

```
        # pacman –get OSG:Globus-Condor-Setup
```

Before moving on to the next steps, your /etc/profile should look like this:

```
export VDT_LOCATION=/opt/osg
export VDTSETUP_CONDOR_LOCATION=/opt/condor
export VDTSETUP_CONDOR_CONFIG=$VDTSETUP_CONDOR_LOCATION/etc/condor_config
export GLOBUS_TCP_PORT_RANGE=40000,40200
export GLOBUS_TCP_SOURCE_RANGE=40000,40200
source $VDT_LOCATION/setup.sh
source $VDT_LOCATION/globus/etc/globus-user-env.sh
source $VDT_LOCATION/monitoring/osg-attributes.conf
```

4. Certificates

A host machine needs a certificate on the root account to authenticate the gatekeeper, and the admin needs a user certificate in order to run jobs via globus. If you intend to run RSV you will also need the rsv and http service certificates.

To get the host certificate, simply run the following command and follow instructions:

```
        # cert-request –ou s –dir . –label uscms1
```

To get the service certificates, run the following commands.

```
        # cert-request -ou s -service rsv -host uscms1.fltech-grid3.fit.edu -label rsv-uscms1
        # cert-request -ou s -service http -host uscms1.fltech-grid.3.fit.edu -label http-uscms1
```

When you receive your host certificate and produced a private key, you must login as root and put them both into ~/.globus/ naming them usercert.pem and userkey.pem
Do the same for your own user cert in your own user account. Keep a backup copy of all certificates and keys. Also put a copy of the host certificate and key in /etc/grid-security/ named hostcert.pem and hostkey.pem (with root owner), and also named containercert.pem and containerkey.pem (with globus owner).
When you do the same for the service certificates, place the rsvcert and key into /etc/grid-security (owned by rsvuser, create it) and the httpcert and key into /etc/grid-security/http.

5. Configuring OSG

You must now set up the configuration python script.

```
# cd $VDT_LOCATION
# ./setup.sh (if you have not added it to profile)
# cd monitoring
# vim config.ini
```

Set up this file with the information related to your site. Most of it is self-explanatory and documentation is included in comments, but there are detailed instructions on the OSG TWiki.
Run the following command:

```
# configure-osg.py -c -f config.ini
```

Note: I exported the $VDT_LOCATION so that it could be mounted on all compute nodes. This directory should be the $OSG_GRID variable.

6. Start VDT Control

Ensure that all standard services are enabled (all default except gums):

```
$ vdt-control --list
```

Run the following command as root:

```
# vdt-control –on
```

Now make sure that condor is running. If you are using the VDT Cache version then you need to run # vdt-control –on condor and re-run the OSG setup script.

7. Set up VO user accounts

You will need to make a list of VOs that you want the cluster to service. Examples of VOs are: osg, uscms01, fermilab, mis…
The VO "mis" is very important, as it just so happens to be the VO of the VORS site-verify script. You will not pass tests without it.
Create the user accounts with the same name as the VOs and make sure they have home directories that are automounted to /home

8. Grid-mapfile, CRL & Crontab

The grid mapfile contains a list of all DNs that are permitted to run commands on the grid, there should be no host DNs on the list – only user DNs.

First, go to the $VDT_LOCATION/edg directory.
Inside the etc folder you will need to edit the edg-mkgridmap.conf file. Add the following line to the bottom of the file:

```
gmf_local   /usr/local/osg/edg/etc/grid-mapfile-local
```

Then edit the grid-mapfile-local file and add your local user identities. You can get your user identity by running grid-proxy-info (if you have run grid-proxy-init first of course) and looking at the identity line. An example of a DN to add to this file is below:

```
"/DC=org/DC=doegrids/OU=People/CN=Patrick Ford ######" uscms01
```

You will also need to add the RSV identity.

```
"/DC=org/DC=doegrids/OU=Services/CN=rsv/uscms1.fltech-grid3.fit.edu" rsvuser
```

Note that a VO needs to be associated with the DN by putting it at the end of the line. Editing these files ensures that you will be included in the grid-mapfile.

Now make sure that the edg-mkgridmap-upgrade.conf file has the following line:

```
GRIDMAP_LOCAL_FILE=$VDT_LOCATION/edg/etc/grid-mapfile-local
```

Run the following command:

```
# edg-mkgridmap
```

The grid-mapfile in /etc/grid-security/ should update automatically after a certain amount of time.

The gridmap and certificate revocation (crl) should be added by default to the crontab when you run vdt-control --on

Our crontab contains the following:

```
            15 5 * * * /opt/osg/fetch-crl/share/doc/fetch-crl-2.6.6/fetch-crl.cron
41 14,20,2,8 * * * /opt/osg/edg/sbin/edg-mkgridmap >> /opt/osg/edg/log/edg-mkgridmap.log 2>&1
8,18,28,38,48,58 * * * * /opt/osg/gratia/probe/condor/condor_meter.cron.sh > /opt/osg/gratia/var/
logs/gratia-probe-condor.log 2>&1
33 * * * * /opt/osg/vdt/sbin/vdt-update-certs-wrapper --vdt-install /opt/osg --called-from-cron
0 0 * * * /opt/osg/vdt/bin/vdt-rotate-logs
```

Notice that edg-mkgridmap and fetch-crl.cron are updating regularly.

Verify that the grid-mapfile is extremely large, it should be by this point, and it must be before getting external verification.

9. IPTables

Several OSG services need certain ports opened and available. Ensure that your iptables file has the following ports open:

```
## INCOMING ##
tcp 2119, tcp 2811, tcp 2135, tcp 9443, udp 9000, tcp 40000 to 40200, tcp 9000 to 9010,
tcp 39281, tcp 7512, tcp 2136, tcp 8443

## OUTGOING ##
tcp 40000 to 40200, tcp 9443
```

10. Running Tests

Up to this point, you should have the environmental variables I listed above set correctly, and an updated grid-mapfile, CRL, and ports opened in IPTables. You should be logged into your own user account and have your usercert and userkey in your ~/.globus/ directory. Create a proxy by running the following commands.

```
$ grid-proxy-init
$ grid-proxy-info –identity
```

Now test the fork and condor queue

```
$ globus-job-run $(hostname –f):2119/jobmanager-fork /usr/bin/id
$ globus-job-run $(hostname –f):2119/jobmanager-condor /usr/bin/id
```

These commands may take a while and should return your VO as the ID (in my case it was uscms01)

Now test the Grid FTP: GSIFTP – this includes creating a test file, copying it to the location on your cluster reserved for grid data (/mnt/nas0/OSG/DATA) through globus, and then verifying that the file copied.

```
$ echo “My test gsiftp file” > /tmp/gsiftp.test
$ source $VDT_LOCATION/monitoring/osg-attributes.conf (for data variable)
$ globus-url-copy file:/tmp/gsiftp.test gsiftp://$(hostname)${OSG_DATA}/gsiftp.test
$ ls –l $OSG_DATA/gsiftp.test
```

This file should exist in the directory /mnt/nas0/OSG/DATA/

11. Site Verify

Run the site verification script as your user with an open proxy.

```
$ $VDT_LOCATION/verify/site_verify.pl --host=uscms1.fltech-grid3.fit.edu
```

You should pass all tests if you followed the instructions completely. Now you need to register your site with the Grid Operations Center in order to get external verification and therefore get on the VORS map. The following are links that involve verification:

http://scan.grid.iu.edu/cgi-bin/show_results?grid=2
http://vors.grid.iu.edu/cgi-bin/index.cgi