



Enabling Grids for E-scienceE

Computing Element installation & configuration

Giuseppe Platania

INFN Catania

EMBRACE Tutorial

Clermont-Ferrand, 07-13.10.2006

www.eu-egee.org



Information Society



- **OVERVIEW**
- **INSTALLATION & CONFIGURATION**
- **TESTING**
- **FIREWALL SETUP**
- **TROUBLESHOOTING**

- The **Computing Element** is the central service of a site.
- Its main functionalities are:
 - manage the jobs (job submission, job control)
 - update to WMS the status of the jobs
 - publish all site informations (site location, queues, about the CPUs status, and so on) via ldap (site BDII service)
- It can run several kinds of batch system:
 - **Torque + MAUI**
 - LSF
 - Condor

- The **Torque server** is composed by a:
 - *pbs_server* which provides the basic batch services such as receiving/creating a batch job.
- The **Torque client** is composed by a:
 - *pbs_mom* which places the job into execution. It is also responsible for returning the job's output to the user
- The **MAUI** system is composed by a:
 - *job_scheduler* which contains the site's policy to decide which job must be executed.



Computing Element installation & configuration using YAIM

- **Request host certificate for CE.**
 - <https://gilda.ct.infn.it/CA/mgt/restricted/srvreq.php>
- **Install host certificate (**hostcert.pem and hostkey.pem**) in **/etc/grid-security**.**
 - `mkdir /etc/grid-security`
 - `chmod 644 hostcert.pem`
 - `chmod 400 hostkey.pem`

- **Because of SUN licence used for Java SDK, it is not possible to redistribute it with the middleware.**
- **You have to download Java SDK 1.4.2 from Sun web site: <http://java.sun.com/j2se/1.4.2/download.html>**
- **Select "Download J2SE SDK", and download the "RPM in self-extracting file". Follow the instruction on the pages to extract the rpm.**

- **Download and install latest version of glite-yaim-3.0.0 -* on all your grid nodes:**

<http://glitesoft.cern.ch/EGEE/gLite/APT/R3.0/rhel30/RPMS.Release3.0/>

- **Download and install the latest version of gilda_ig-yaim-3.0.0 -* on all your grid nodes:**

http://grid018.ct.infn.it/apt/gilda_app-i386/utils/gilda_ig-yaim.latest

- Copy **gilda_ig-site-info.def** template file provided by gilda_ig_yaim in to the root dir and customize it

```
cp /opt/glite/yaim/examples/gilda_ig-site-info.def \
/root/my-site-info.def
```

- Open `/root/my-site-info.def` file using a text editor and set the following values according to your grid environment:

```
MY_DOMAIN=<your DOMAIN>
```

```
CE_HOST=<write the CE hostname you are installing>
```

```
NTP_HOSTS="193.206.144.10"
```

```
TORQUE_SERVER=$CE_HOST
```

- **Set the repositories:**

```
INSTALL_SERVER_HOST=training50d.$MY_DOMAIN
```

```
OS_REPOSITORY="rpm http://$INSTALL_SERVER_HOST slc306-i386 os  
updates extras localrpms"
```

```
LCG_REPOSITORY="rpm http://$INSTALL_SERVER_HOST glite_sl3-  
i386 3_0_0 3_0_0 externals 3_0_0 updates"
```

```
IG_REPOSITORY="rpm http://$INSTALL_SERVER_HOST ig_sl3-i386  
3_0_0 utils"
```

```
GILDA_REPOSITORY="rpm http://$INSTALL_SERVER_HOST gilda_app-  
i386 app 3_0_0"
```

```
CA_REPOSITORY="rpm http://$INSTALL_SERVER_HOST glite_sl3-  
i386 security"
```

```
JAVA_LOCATION="/usr/java/j2sdk1.4.2_12"
```

```
MYSQL_PASSWORD=set_this_to_a_good_password
```

```
APEL_DB_PASSWORD="APELDB_PWD"
```

```
SITE_EMAIL=grid-prod@healthgrid.org
```

```
SITE_NAME=<EMBRACE-151 .. EMBRACE-161>
```

```
SITE_LOC="Clermont, France"
```

```
SITE_LAT=45.7
```

```
SITE_LONG=3.08
```

```
SITE_WEB="http://www.healthgrid.org."
```

```
SITE_TIER="EMBRACE Testbed"
```

```
SITE_SUPPORT_SITE="grid-prod@healthgrid.org"
```

```
JOB_MANAGER=lcgpbs
CE_BATCH_SYS=pbs
BATCH_BIN_DIR=/usr/bin
BATCH_VERSION=torque-1.0.1b
CE_CPU_MODEL=PIII
CE_CPU_VENDOR=intel
CE_CPU_SPEED=1400
CE_OS="Scientific Linux CERN"
CE_OS_RELEASE=3.0.6
CE_OS_VERSION="SLC"
CE_MINPHYSMEM=1024
CE_MINVIRTMEM=2048
CE_SMP_SIZE=2
CE_SI00=1000
CE_SF00=1200
CE_OUTBOUNDIP=TRUE
CE_INBOUNDIP=TRUE
CE_RUNTIMEENV="list of tags to publish"
```

```
CLASSIC_HOST="classic_SE_hostname"  
DPM_HOST="dpm_hostname"  
SE_LIST="$DPM_HOST $CLASSIC_HOST"  
BDII_REGIONS="CE SE"  
BDII_CE_URL="ldap://$CE_HOST:2135/mds-vo-name=local,o=grid"  
BDII_SE_URL="ldap://$CLASSIC_HOST:2135/mds-vo-  
name=local,o=grid"  
BDII_SE1_URL="ldap://$DPM_HOST:2135/mds-vo-name=local,o=grid"  
  
VOS="write here the VOs you want to support"  
ALL_VOMS="write here the VOs supported that have a VOMS"  
  
QUEUES="short long infinite"
```

```
WN_LIST=/opt/glite/yaim/examples/gilda_wn-list.conf
```

The file specified in **WN_LIST** has to be set with the list of all your WNs's hostname.

WARNING: It's important to setup it before to run the configure command

There are several kind of metapackages to install:

GILDA_ig_CE

- LCG ComputingElement without batch system packages.

GILDA_ig_CE_LSF

- LCG ComputingElement with LSF. **IMPORTANT:** provided for consistency, it does not install LSF but it apply some fixes via `GILDA_ig_configure_node`.

GILDA_ig_CE_torque

- LCG ComputingElement with Torque+MAUI.

GILDA_ig_glite_CE gLite

- ComputingElement without batch system packages.

GILDA_ig_glite_CE_LSF

- gLite ComputingElement with LSF. **IMPORTANT:** provided for consistency, it does not install LSF

GILDA_ig_glite_CE_torque

- gLite ComputingElement with Torque+MAUI.

- **This command will download and install all the needed packages:**

```
/opt/glite/bin/gilda_ig_install_node  
/root/my-site-info.def GILDA_ig_CE_torque
```

- **Now we can configure the node:**

```
/opt/glite/bin/gilda_ig_configure_node  
/root/my-site-info.def GILDA_ig_CE_torque
```




Computing Element testing

- **Edit a file and write:**

```
#!/bin/sh
```

```
sleep 10 #(it's useful to see the job status)
```

```
hostname
```

- **Save it and set the permission of execution:**

```
chmod 700 test.sh
```

```
[gilda003@ce gilda003]$ qsub -q short test.sh
```

```
3.ce-wn.localdomain
```

```
[gilda003@ce gilda003]$ qstat -a
```

```
ce.localdomain:
```

| Job ID | Username | Queue | Jobname | SessID | NDS | TSK | Req'd Memory | Req'd Time | Elap S | Time |
|-----------------|----------|-------|---------|--------|-----|-----|-----------------|---------------|-----------|------|
| ----- | ----- | ----- | ----- | ----- | --- | --- | ----- | ----- | - | ---- |
| 3.ce-wn.localdo | gilda003 | short | test.sh | 5839 | -- | -- | -- | 00:15 | R | -- |

```
[gilda003@ce gilda003]$ qstat -a
```

```
[gilda003@ce gilda003]$
```

- **The job execution has finished and we have to list the output file:**

```
[gilda003@ce gilda003]$ ls
```

```
test.sh.e3 test.sh.o3
```

- **And show them:**

```
[gilda003@ce gilda003]$ cat test.sh.e3      (error file)
```

```
[gilda003@ce gilda003]$
```

```
[gilda003@ce gilda003]$ cat test.sh.o3      (output file)
```

wn.localdomain

```
[plt@ui plt]$ voms-proxy-init -voms gilda
```

```
[plt@ui plt]$ globus-job-run ce.localdomain:2119/jobmanager-lcgpbs -q short /bin/hostname  
wn.localdomain
```

```
[plt@ui plt]$ edg-job-submit -r ce-wn.localdomain:2119/jobmanager-lcgpbs-short hostname.jdl
```

```
Selected Virtual Organisation name (from proxy certificate extension): gilda
```

```
Connecting to host ui-rb-bdii.localdomain, port 7772
```

```
Logging to host ui-rb-bdii.localdomain, port 9002
```

```
*****
```

JOB SUBMIT OUTCOME

The job has been successfully submitted to the Network Server.

Use edg-job-status command to check job current status. Your job identifier (edg_jobId) is:

- <https://ui-rb-bdii.localdomain:9000/Vo-4Ih1s-iDbBPr3rs69GQ>

```
*****
```

- Check if the local GRIS and the site BDII are running on CE and are publishing the right informations (CPU, site name and so on)

```
ldapsearch -x -h <ce_hostname> -p 2135 -b mds-vo-name=local,o=grid
```

```
ldapsearch -x -h <ce_hostname> -p 2170 -b mds-vo-name=<site_name>,o=grid
```



FIREWALL SETUP

```

*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [0:0]
:RH-Firewall-1-INPUT - [0:0]
-A INPUT -j RH-Firewall-1-INPUT
-A FORWARD -j RH-Firewall-1-INPUT
-A RH-Firewall-1-INPUT -i lo -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state ESTABLISHED,RELATED -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 22 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 2135 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 2119 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 2170 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 2811 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport maui -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport pbs_mom -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport pbs_resmon -j
ACCEPT
    
```



```

-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport pbs -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 3878 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 3879 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m udp -p udp --dport 3879 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 3882 -j ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m udp -p udp --dport 1020:1023 -j
  ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 20000:25000 -j
  ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 32768:65535 -j
  ACCEPT
-A RH-Firewall-1-INPUT -m state --state NEW -m udp -p udp --dport 32768:65535 -j
  ACCEPT
-A RH-Firewall-1-INPUT -p tcp -m tcp --syn -j REJECT
-A RH-Firewall-1-INPUT -j REJECT --reject-with icmp-host-prohibited
COMMIT

```

```
/sbin/chkconfig iptables on
```

```
/etc/init.d/iptables start
```



Troubleshooting

```
[plt@ui plt]$ globus-job-run ce-wn.localdomain:2119/jobmanager-lcgpbs -q short /bin/hostname
GRAM Job submission failed because the connection to the server failed (check host and port) (error
code 12)
```

solution: check if the globus-gatekeeper daemon is up and running on CE

```
[plt@ui plt]$ globus-job-run ce-wn.localdomain:2119/jobmanager-lcgpbs -q short /bin/hostname
GRAM Job submission failed because authentication failed:
GSS Major Status: Authentication Failed
GSS Minor Status Error Chain:

init.c:499: globus_gss_assist_init_sec_context_async: Error during context initialization
init_sec_context.c:171: gss_init_sec_context: SSLv3 handshake problems
globus_i_gsi_gss_utils.c:888: globus_i_gsi_gss_handshake: Unable to verify remote side's credentials
globus_i_gsi_gss_utils.c:847: globus_i_gsi_gss_handshake: Unable to verify remote side's credentials:
    Couldn't verify the remote certificate
OpenSSL Error: s3_pkt.c:1046: in library: SSL routines, function SSL3_READ_BYTES: sslv3 alert bad
certificate (error code 7)
```

solution: probably there is no GILDA CA rpm installed on CE

```
[plt@ui plt]$ edg-gridftp-ls gsiftp://ce.localdomain/  
error the server sent an error response: 530 530 LCMAPS  
credential mapping NOT successful
```

```
error the server sent an error response: 530 530 LCMAPS  
credential mapping NOT successful
```

solution: check on CE the VO mapping in
`/opt/edg/etc/lcmaps/gridmapfile`
`/opt/edg/etc/lcmaps/groupmapfile`