



LHC COMPUTING GRID

INSTALLING THE RELEASE

Document identifier:

EDMS id: **n/a**

Version: **v2.4.0**

Date: **April 6, 2005**

Section: **gis**

Document status: **final**

Author(s): GRID Deployment Group
(<support-lcg-deployment@cern.ch>)

File: **LCG2-LCFG-Install**

Abstract: These notes will assist you in installing the new release using LCFGng



CONTENTS

1. INSTALLING A NEW RELEASE USING LCFGNG	3
1.1. DOWNLOADING THE LCFGNG CONFIGURATION FILES.....	3
1.2. DOWNLOAD THE RPMS	3
1.3. UPDATING THE LCFGNG SERVER.....	4
1.4. SETTING THE ROOT PASSWORD.....	4
1.5. SITE CONFIGURATION.....	4
1.6. NODE PROFILE CREATION	5
2. MANUAL STEPS.....	6
2.1. COMMON STEPS	6
2.2. USERINTERFACE	6
2.3. RESOURCEBROKER	6
2.4. COMPUTINGELEMENT	7
2.5. MONNODE	8
2.6. WORKERNODE	9
2.7. STORAGEELEMENT	10
2.8. PLAINGRIS	10
2.9. BDII NODE	10
2.10. MYPROXY NODE	10
3. APPENDIX A	11
3.1. HOW TO CONFIGURE THE PBS SERVER ON A COMPUTINGELEMENT	11



1. INSTALLING A NEW RELEASE USING LCFGNG

If you do not already have an LCFGng server installed, please use the LCFGng Server Installation Guide to install an LCFGng server. The files used by the LCFGng server can be located in different places. There are three important locations. The rpmlist directory contains the files related to rpm lists. The source directory contains files related to configuration. The profile directory contains the site-cfg.h file and the node profiles. If you do not know where these directories are or how they are used, you should read the LCFGng Server Installation Guide. The example tag LCG-x.y.z will be used through out the guide, this should be replaced with the current tag.

1.1. DOWNLOADING THE LCFGNG CONFIGURATION FILES.

All the LCFGng configuration files needed for the release can be found in the LCG CVS repository. These files should be checked out from the CVS repository onto your LCFGng Server. To do this, first set the CVS environment variables.

```
> export CVS_RSH=ssh
> export CVSROOT=:pserver:anonymous@lcgdeploy.cvs.cern.ch:/cvs/lcgdeploy
```

Checkout the tag from CVS.

```
> cvs checkout -r LCG-x.y.z -d LCG-x.y.z lcg2
```

This will create a directory LCG-x.y.z that contains the new configuration files. The files in the LCG-x.y.z/rpmlist and LCG-x.y.z/source directories should be copied to the locations that you use for rpmlist and source.

1.2. DOWNLOAD THE RPMS

The rpms needed for the release need be copied on to the LCFGng server. To do this copy the file LCG-x.y.z/tools/updaterep.conf to /etc/updaterep.conf and run the updaterep to download the rpms.

```
> cp LCG-x.y.z/tools/updaterep.conf /etc/updaterep.conf
> updaterep
```

By default all rpms will be copied to the /opt/local/linux/7.3/RPMS area, which is visible from the client nodes as /export/local/linux/7.3/RPMS. You can change this location by editing /etc/updaterep.conf and modifying the REPOSITORY_BASE variable.

IMPORTANT NOTICE: As the list and structure of Certification Authorities (CA) accepted by the LCG project can change independently from the middle-ware releases, the rpm list related to the CAs certificates and URLs has been decoupled from the standard LCG release procedure. This means that the version of the security-rpm.h file contained in the rpmlist directory associated to the current tag might be incomplete or obsolete. Please go to the URL <http://markusw.home.cern.ch/markusw/lcg2CAlist.html> and follow the instructions there to update the CAs. Changes and updates of the CAs will be announced on the LCG-Rollout mailing list.



1.3. UPDATING THE LCFGNG SERVER

To ensure that all the LCFGng server-side object rpms are installed on your LCFGng server, run the command,

```
> LCG-x_y_z/tools/lcfgng_server_update.pl LCG-x_y_z/rpmlist/lcfgng-server-rpm.h
```

This script will report which rpms are missing or have the wrong version and will create a script `/tmp/lcfgng_server_update_script.sh`. The script will update all the rpms needed on the LCFGng server. Please check and verify that all the commands look reasonable before running it.

1.4. SETTING THE ROOT PASSWORD

You must replace the default root password in the file, `private-cfg.h` with the one you want to use for your site:

```
+auth.rootpwd <CRYPTED_PWD>
```

To obtain the encrypted MD5 encryption password (stronger than the standard crypt method) you can use the following command:

```
> openssl passwd -1
```

This command will prompt you to insert the clear text version of the password and then print the encrypted version.

1.5. SITE CONFIGURATION

The directory `LCG-x_y_z/examples` contains an example of the `site-cfg.h` file. This file contains all the information needed about your site. You should copy this file to the profiles directory and edit it. The description of the values are well documented within the file.

Additional VOs on the LCFGng Server

In the directory `LCG-x_y_z/examples`, there is an example VO configuration file `new-vo.conf`. For every additional VO that you want to support, copy this file to the sources directory. Edit this file so that it contains all the required information about the VO that you want to support. In the directory `LCG-x_y_z/sources`, there is a script called `addvo.py`. Copy this script to the sources directory. Create the VO configuration file by running the command.

```
> ./addvo.py -i new-vo.conf > vo-newco.cfg.h
```

To include this file you will need to add the following line to the `vos-cfg.h` file.

```
#include "new-vo-cfg.h"
```



1.6. NODE PROFILE CREATION

The directory `LCG-x.y.z/examples` contains example profiles that can be used to create the initial node profiles needed for your site. In the `LCG-x.y.z/tools` there is a script `do_mkxprof.sh`. A detailed description of how this script works is contained in the script itself. To create the LCFG configuration for one or more nodes you can do

```
> do_mkxprof.sh node1 [node2 node3, ...]
```

If you get an error status for one or more of the configurations, you can get a detailed report on the nature of the error by looking into URL

```
> http://<LCFGng_Server>/status
```

Once all node configurations are correctly published, you can proceed and install your nodes following any one of the installation procedures described in the LCFGng Server Installation Guide.



2. MANUAL STEPS

Each node type requires a few manual steps to be completely configured.

2.1. COMMON STEPS

- On the ResourceBroker, MyProxy, StorageElement, and ComputingElement nodes you must install the host certificate/key files in `/etc/grid-security` with names `hostcert.pem` and `hostkey.pem`. Also make sure that `hostkey.pem` is only readable by root with

```
> chmod 400 /etc/grid-security/hostkey.pem
```

- All Globus services grant access to LCG users according to the certificates listed in the `/etc/grid-security/grid-mapfile` file. The list of VOs included in `grid-mapfile` is defined in `/opt/edg/etc/edg-mkgridmap.conf`. This file is now handled automatically by the `mkgridmap` LCFG object. This object takes care of enabling only the VOs accepted at each site according to the `SE_VO.<VO>` definitions in `site-cfg.h`. If you need to modify the default configuration for your site, e.g. by adding users to `grid-mapfile-local`, you can do this from your `local-cfg.h` file by following the examples in `<TAG_DIRECTORY>/source/mkgridmap-cfg.h`.

After installing a ResourceBroker, StorageElement, or ComputingElement node you should force a first creation of the `grid-mapfile` by running

```
> /opt/edg/sbin/edg-mkgridmap --output=/etc/grid-security/grid-mapfile
--safe !
```

Every 6 hours a cron job will repeat this procedure and update `grid-mapfile`.

2.2. USERINTERFACE

No additional configuration steps are currently needed on a UserInterface node.

2.3. RESOURCEBROKER

Log as root on your RB node, represented by `<rb_node>` in the example, and make sure that the `mysql` server is up and running:

```
> /etc/rc.d/init.d/mysql start
```

If it was already running you will just get notified of the fact.

Now you can choose a DB management `<password>` you like (write it down somewhere!) and then configure the server with the following commands:



```
> mysqladmin password <password>
> mysql --password=<password> \
    --exec "set password for root@<rb_node>=password('<password>')" mysql
> mysqladmin --password=<password> create lbserver20
> mysql --password=<password> lbserver20 < /opt/edg/etc/server.sql
> mysql --password=<password> \
    --exec "grant all on lbserver20.* to lbserver@localhost" lbserver20
```

Note that the database name "lbserver20" is hardwired in the LB server code and cannot be changed so use it exactly as shown in the commands.

Make sure that /var/lib/mysql has the right permissions set (755).

- Reboot the node

2.4. COMPUTINGELEMENT

Don't forget after upgrading the CE to make sure that the experiment specific runtime environment tags can still be published. For this move the /opt/edg/var/info/;VO-NAME;/.ldif files to ;VO-NAME;.list.

- Configure the PBS server. See detailed recipe in Appendix B at the end of this document.
- Create the first version of the /etc/ssh/ssh_known_hosts file by running

```
> /opt/edg/sbin/edg-pbs-knownhosts
```

A cron job will update this file every 6 hours.

- If your CE is NOT sharing the /home directory with your WNs (this is the LCG-2 default configuration: if you have modified site-cfg.h to run PBS in traditional mode as described in a previous chapter, just ignore the following instructions) then you have to configure sshd to allow WNs to copy job output back to the CE using scp. This requires the following two steps:

1. modify the sshd configuration. Edit the /etc/ssh/sshd_config file and add these lines at the end:

```
HostbasedAuthentication yes
IgnoreUserKnownHosts yes
IgnoreRhosts yes
```

and then restart the server with

```
> /etc/rc.d/init.d/sshd restart
```

2. configure the script enabling WNs to copy output back to the CE.



- in `/opt/edg/etc`, copy `edg-pbs-shostsequiv.conf.template` to `edg-pbs-shostsequiv.conf` then edit this file and change parameters to your needs. Most sites will only have to set `NODES` to an empty string.
- create the first version of the `/etc/ssh/shosts.equiv` file by running

```
> /opt/edg/sbin/edg-pbs-shostsequiv
```

A cron job will update this file every 6 hours.

Note: every time you will add or remove WNs, do not forget to run

- > `/opt/edg/sbin/edg-pbs-shostsequiv` <— only if you do not share `/home`
- > `/opt/edg/sbin/edg-pbs-knownhosts`

on the CE or the new WNs will not work correctly till the next time cron runs them for you.

- The CE is supposed to export information about the hardware configuration (i.e. CPU power, memory, disk space) of the WNs. The procedure to collect these informations and publish them is described in Appendix D of this document.
- Reboot the node
- If your CE exports the `/home` area to all WNs, then after rebooting it make sure that all WNs can still see this area. If this is not the case, execute this command on all WNs:

```
> /etc/obj/nfsmount restart
```

2.5. MONNODE

The MySQL database used by the servlets needs to be configured.

```
> mysql -u root < /opt/edg/var/edg-rgma/rgma-db-setup.sql
```

Restart the R-GMA servlets.

```
/etc/rc.d/init.d/edg-tomcat4 restart
```

For the accounting package, `apel`, you will need to setup the database.

```
> mysql -u root --exec "create database accounting"
> mysql -u root accounting < /opt/edg/var/edg-rgma/apel-schema.sql
> mysql -u root --exec "grant all on accounting.* to accounting@$MON_HOST identified by APELDB_"
> mysql -u root --exec "grant all on accounting.* to accounting@localhost identified by APELDB_"
> mysql -u root --exec "grant all on accounting.* to accounting@$CE_HOST identified by APELDB_"
```




2.6. WORKERNODE

- The default allowed maximum number of open file on a RedHat node is only 26213. This number might be too small if users submit file-hungry jobs (we already had one case) so you may want to increase it on your WNs. At CERN we currently use 256000. To set this parameter you can use this command:

```
> echo 256000 > /proc/sys/fs/file-max
```

You can make this setting reboot-proof by adding the following code at the end of your `/etc/rc.d/rc.local` file:

```
# Increase max number of open files
if [ -f /proc/sys/fs/file-max ]; then
    echo 256000 > /proc/sys/fs/file-max
fi
```

- Every 6 hours each WN needs to connect to the web sites of all known CAs to check if a new CRL (Certificate Revocation List) is available. As the script which handles this functionality uses `wget` to retrieve the new CRL, you can direct your WNs to use a web proxy. This is mandatory if your WNs sit on a hidden network with no direct external connectivity.

To redirect your WNs to use a web proxy you should edit the `/etc/wgetrc` file and add a line like:

```
http_proxy = http://web_proxy.cern.ch:8080/  where you should replace the node
name and the port to match those of your web proxy.
```

Note: I could not test this recipe directly as I am not aware of a web proxy at CERN. If you try it and find problems, please post a message on the `lcg-rollout` list.

- If your WNs are NOT sharing the `/home` directory with your CE (this is the default configuration) then you have to configure `ssh` to enable them to copy job output back to the CE using `scp`. To this end you have to modify the `ssh` client configuration file `/etc/ssh/ssh_config` adding these lines at the end:

```
Host *
    HostbasedAuthentication yes
```

Note: the "Host *" line might already exist. In this case, just add the second line after it.

- Create the first version of the `/etc/ssh/ssh_known_hosts` file by running

```
> /opt/edg/sbin/edg-pbs-knownhosts
```

A cron job will update this file every 6 hours.



2.7. STORAGEELEMENT

- Make sure that the storage area defined with `CE_CLOSE_SE_MOUNTPOINT` exists and contains the VO specific sub-directories with the correct access privileges (group=`<VO>` and r/w access for the group).
- Reboot the node.

2.8. PLAINGRIS

No additional configuration steps are currently needed on a PlainGRIS node.

2.9. BDII NODE

This is the current version of the BDII service which does not rely on Regional MDSes. If you want to install the new service then you should use the `LCG-BDII_node` example file from the "examples" directory. After installation the new LCG-BDII service does not need any further configuration: the list of available sites will be automatically downloaded from the default web location defined by `SITE_BDII_URL` in `site-cfg.h` and the initial population of the database will be started. Expect a delay of a couple of minutes from when the machine is up and when the database is fully populated.

If for some reason you want to use a static list of sites, then you should copy the static configuration file to `/opt/lcg/var/bdii/lcg-bdii-update.conf` and add this line at the end of your LCG-BDII node configuration file:

```
+lcgbdii.auto    no
```

If you need a group of BDIIs being centrally managed and see a different set of sites than those defined by URL above you can setup a web-server and publish the web page containing the sites. The URL for this file has to be used to configure the `SITE_BDII_URL` in the `site-cfg.h`. Leave the `lcgbdii.auto` to yes.

This file has the following structure: `http://grid-deployment.web.cern.ch/grid-deployment/gis/lcg2-bdii/dteam/lcg2-all-sites.conf`

2.10. MYPROXY NODE

- Reboot the node after installing the host certificates (see "Common Steps" above).

Make sure that in the `site-cfg.h` file you have included all Resource Brokers that your users want to use. This is done in the following line:

```
#define GRID_TRUSTED_BROKERS    "/C=CH/O=CERN/OU=GRID/CN=host/BROKER1.Domain.ch \  
                                /C=CH/O=CERN/OU=GRID/CN=host/Broker2.Domain.ch"
```



3. APPENDIX A

3.1. HOW TO CONFIGURE THE PBS SERVER ON A COMPUTINGELEMENT

Note that queues short, long, and infinite are those defined in the site-cfg.h file and the time limits are those in use at CERN. Feel free to add/remove/modify them to your liking but do not forget to modify site-cfg.h accordingly.

The values given in this example are only reference values. Make sure that the requirements of the experiment as stated here: <http://ibird.home.cern.ch/ibird/LCGMinResources.doc> are satisfied by your configuration.

1. load the server configuration with this command (replace <CEhostname> with the hostname of the CE you are installing):

```
@-----  
/usr/bin/qmgr <<EOF  
  
set server scheduling = True  
set server acl_host_enable = False  
set server managers = root@<CEhostname>  
set server operators = root@<CEhostname>  
set server default_queue = short  
set server log_events = 511  
set server mail_from = adm  
set server query_other_jobs = True  
set server scheduler_iteration = 600  
set server default_node = lcgpro  
set server node_pack = False  
  
create queue short  
set queue short queue_type = Execution  
set queue short resources_max.cput = 00:15:00  
set queue short resources_max.walltime = 02:00:00  
set queue short enabled = True  
set queue short started = True  
  
create queue long  
set queue long queue_type = Execution  
set queue long resources_max.cput = 12:00:00  
set queue long resources_max.walltime = 24:00:00  
set queue long enabled = True  
set queue long started = True  
  
create queue infinite
```



```
set queue infinite queue_type = Execution
set queue infinite resources_max.cput = 80:00:00
set queue infinite resources_max.walltime = 100:00:00
set queue infinite enabled = True
set queue infinite started = True
EOF
@-----
```

2. edit file `/var/spool/pbs/server_priv/nodes` to add the list of WorkerNodes you plan to use. An example setup for CERN could be:

```
@-----
lxshare0223.cern.ch np=2 lcgpro
lxshare0224.cern.ch np=2 lcgpro
lxshare0225.cern.ch np=2 lcgpro
lxshare0226.cern.ch np=2 lcgpro
@-----
```

where `np=2` gives the number of job slots (usually equal to #CPUs) available on the node, and `lcgpro` is the group name as defined in the `default_node` parameter in the server configuration.

3. Restart the PBS server

```
> /etc/rc.d/init.d/pbs_server restart
```