

LCG-0 installation instructions

Feb/28, 2003

Common Prerequisites:

It is expected that all machines where the LCG Grid software will be installed already have the Linux operating system running. We tested the installation on the vanilla RedHat 7.3 ('server' configuration) or CERN RH 7.3.1. If you are installing the LCG Grid software on a machine which already has some parts of the EDG and/or Globus software installed, it is recommended to completely remove it before starting the LCG-0 installation.

To install the software the super user (or "root") privileges are required.

All LCG-0 software has been compiled with gcc-2.95-2, so we distribute also this compiler and its libraries. As a general rule it is mandatory to use all packages we supply in the LCG-0 distribution (some of them could differ from similar packages in a vanilla RH distribution).

A generic user (example: lcg) should exist for the execution of Grid jobs. This user identity and the corresponding home directory must be shared between the machines which will serve as CE, SE and WNs (uid and gid must be the same on all the machines).

A batch system must be installed, so that jobs may be submitted from the CE machine and run on any WN machines. Currently OpenPBS is supported for the batch system. As an example details of how to setup an OpenPBS system are included in the Appendix.

The batch system must be setup with a queue to which the generic grid user is allowed to submit jobs. This is the queue to which all grid jobs will be submitted.

If you are going to have a shared NFS area for your SE (example: /flatfiles/lcg) you need to mount this area on all your machines (CE, SE, WNs). This will allow you to use the "file" data access protocol to access files replicated on your SE.

The LCG-0 software is installed in all or some of the following directories (depending on what is being installed):

- /opt/edg
- /opt/edt
- /opt/vdt
- /opt/globus
- /opt/globus-24
- /usr/local/gcc-alt-2.95.2
- /etc
- /etc/grid-security

`/etc/rc.d/init.d/`

Sufficient free space (approximately 100Mbytes) should be available under `/opt`.

The file `/etc/services` is modified.

It is not yet possible to install different kinds of services (CE, SE, UI) on the same machine.

1.Download installation scripts:

You can download the installation scripts from the Grid Deployment Group web page: <http://cern.ch/grid-deployment/> following the LCG release link.

2.Install UI (User Interface):

The UI has to be installed on every site. It allows access to the grid infrastructure. At least one UI is required on a site for LCG-0 testing purposes. User accounts will be created on this machine.

1) Download the installation script `install_UI.sh` from the Grid Deployment web site:

```
wget http://grid-deployment.web.cern.ch/grid-deployment/bindir/LCG-0/install_UI.sh
```

2a) Check the allowed arguments:

```
sh install_UI.sh -help
```

2b) Run from bash (supply any necessary arguments):

```
sh install_UI.sh arguments > /root/UI_installation.log 2> &1
```

(the default download directory will be `/root/rpm_lcg` and will not be removed automatically).

3.Install CE (Computing Element):

The installation of the CE is required if the site wants to share local computing resources across the grid. In this case one CE is enough for LCG-0 testing purposes.

The CE requires a machine certificate before you start the installation. Contact your national certification authority to understand how to obtain a machine certificate if you do not have one already (certificate public and private keys will be automatically copied to

```
/etc/grid-security/hostcert.pem  
/etc/grid-security/hostkey.pem
```

the latter with permission set to 0400).

We assume the CE machine already has been configured as the master node of a batch system.

The following extra requirements are needed for the CE.

A user (example: edginfo) must be setup on the CE machine. It is not necessary that the home directory for this user be shared with other machines. This user must be different from the one under which Grid jobs will run.

Installation is by way of a perl script, available from <http://cern.ch/grid-deployment/> following the links for the LCG-0 installation to get the install script for the CE.

1) Download it from the page and install it in some convenient place, eg /root/ :

```
wget http://grid-deployment.web.cern.ch/grid-deployment/bindir/LCG-0/install_CE.pl
```

The installation process modifies various directories and therefore must be run as root. (See below for some details of the process of installation).

2) Ensure the CE install script is executable:

```
chmod a+x install_CE.pl
```

You can invoke *./install_CE.pl* without arguments to see the usage:

Usage: *install_CE.pl*

<hostname.domain>	The fully qualified name of this machine
<lcguname>	The username of the grid user
<hostcert.pem>	Location of a copy of the host certificate
<hostkey.pem>	Location of a copy of the host certificate key
<closese.domain>	The fully qualified name of the close storage element
<infouname>	The username of the info user
<Batch system>	Type of batch system: currently only type PBS is supported
<Batch system path>	Path under which to find the batch system commands
<Batch queue>	The batch queue name that grid jobs should be submitted to

The *install_CE.pl* script needs to be invoked with 9 parameters. Here we describe each of them in a little more detail than is reported in the command usage.

hostname.domain	The fully qualified host name of the machine. (deployment onto multihomed machines is not currently supported)
lcgname	The user name under which jobs received over the Grid will be run. All jobs will use this single id.
hostcert.pem	Location of a file containing the host certificate which will be used to identify the CE. The host certificate should be in 'pem' format.
hostkey.pem	Similar to hostcert.pem but this file should contain the private key corresponding to the host certificate.
closese.domain	The fully qualified host name of the machine which will be treated as the 'close' SE for this CE.
infouname	A username required by the Grid information system. Must be different from <lcgname>.
Batch system	Defines the type of batch system to be used. Currently PBS is supported.
Batch system path	The path under which the batch system commands can be found.
Batch queue	This should be name of the batch queue to which all jobs received over the Grid will be submitted.

Before starting the install you should note that install_CE.pl will create a temporary directory 'rpm_lcg' inside the directory from which it is started. The temporary directory will hold all the LCG-0 components needed for the CE, which are about 50 Mbytes in size. You should therefore choose a location with sufficient space.

3) The installation process usually takes less than 2 minutes, but this depends on the time taken to retrieve the LCG components from the distribution site.

Included below is an example of the installation of a CE. The machine being installed as the CE was 'lxshare0240.cern.ch' and the close SE was 'lxshare0241.cern.ch'.

```
cd /root
./install_CE.pl lxshare0240.cern.ch lcg hostcert.pem hostkey.pem lxshare0241.cern.ch
edginfo PBS /usr/pbs/bin workq
```

++ Starting the installation of LCG Compute Element ++

```

Fetching list of LCG components from distribution site... Done
Fetching globus_VDT_CE.tgz... Now installing
Fetching GNU.LANG_gcc-alt-2.95.2-6.i386.rpm... Now installing
Fetching BrokerInfo-gcc32-3.2-0.i386.rpm... Now installing
Fetching ReplicaCatalogue-gcc32-3.2-3.i386.rpm... Now installing
Fetching edg-replica-manager-gcc32-2.0-6.i386.rpm... Now installing
Fetching workload-profile-1.2.19-1.i386.rpm... Now installing
Fetching locallogger-profile-1.2.21-1.i386.rpm... Now installing
Fetching locallogger-1.2.21-1.i386.rpm... Now installing
Fetching globus_gatekeeper-edgconfig-0.17-nodep.1.noarch.rpm... Now installing
Fetching globus_gsi_wuftpd-edgconfig-0.17-nodep2.noarch.rpm... Now installing
Fetching globus_profile-edgconfig-0.17-nodep.noarch.rpm... Now installing
Fetching edg-user-env-0.3-1.noarch.rpm... Now installing
Fetching edg-profile-0.3-1.noarch.rpm... Now installing
Fetching perl-Convert-ASN1-0.16-7.i386.rpm... Now installing
Fetching perl-Net_SSLeay-1.21-7.i386.rpm... Now installing
Fetching perl-IO-Socket-SSL-0.92-7.i386.rpm... Now installing
Fetching perl-perl-ldap-0.26-7.i386.rpm... Now installing
Fetching edg-mkgridmap-1.0.9-2.i386.rpm... Now installing
Fetching edg-utils-system-1.3.2-1.noarch.rpm... Now installing
Executing globus postinstall scripts... Done
Starting globus-gatekeeper:                [ OK ]
Starting globus-gsi_wuftpd:                [ OK ]
Starting up Openldap 2.0 SLAPD server for the GRIS
Starting LocalLogger: interlogger and dglogd
LCG CE install completed
Leaving temporary directory /root/rpm_lcg for reference
--

```

(the temporary download directory, /root/rpm_lcg in this case, will not be removed automatically but may be removed by hand after installation, if desired)

After installation the required services should already be running, there is no need to restart the machine.

Ask the LCG-0 administrators on <http://cern.ch/grid-deployment/> to include your CE in the information index.

----- Appendix - Installation of PBS masternode (eg. on the CE machine) -

*Install and configure a PBS master node (extracted from EDG guide)
(please refer to PBS documentation to obtain detailed informations)*

```

download from our repository openpbs-2.3pl2-1
rpm -ivh --nodeps openpbs-2.3pl2-1

```

a) Set the PBS server name:

```
echo "yourCE.yourDomain" > /usr/spool/PBS/server_name  
where "yourCE.yourDomain" is your CE.
```

b) Add PBS ports in /etc/services as following:

```
# PBS  
pbs 15001/tcp  
pbs_mom 15002/tcp  
pbs_remom 15003/tcp  
pbs_remom 15003/udp  
pbs_sched 15004/tcp
```

c) /usr/pbs/sbin/pbs_server -t create

d) Create the WN list /usr/spool/PBS/server_priv/nodes. The format is:

```
someWN.yourDomain np=2 lcgqueue  
"lcgqueue" is an arbitrary name which has been used to configure the server.  
"np" sets the number of concurrent jobs which can be run on the WN.  
" someWN.yourDomain" is one of your WN.
```

e) /sbin/chkconfig pbs on

```
/etc/rc.d/init.d/pbs stop  
/etc/rc.d/init.d/pbs start
```

f) /usr/pbs/bin/qmgr < /usr/spool/PBS/pbs_server.conf

Please refer to PBS documentation for accurate configuration on the batch system

4.Install SE (Storage Element):

The installation of the SE is required if the site wants to share local storage resources across the grid. In this case one SE is enough for LCG-0 testing purposes.

The SE requires a machine certificate before you start the installation. Contact your national certification authority to understand how to obtain a machine certificate if you do not have one already.

1) Download the installation script *install_SE.sh* from the Grid Deployment web site:

```
wget http://grid-deployment.web.cern.ch/grid-deployment/bindir/LCG-0/install_SE.sh
```

2) Consult your Certification Authority to obtain the certificate for the machine.

Put machine certificates (public and private key) in

```
/etc/grid-security/hostcert.pem  
/etc/grid-security/hostkey.pem
```

the latter with permission set to 0400.

3a) Check the allowed arguments:

```
sh install_SE.sh -help
```

3b) Run from bash (supply any necessary arguments):

```
sh install_SE.sh arguments > /root/SE_installation.log 2>&1
```

(the default download directory will be /root/rpm_lcg and will not be removed automatically)

4) Check `/opt/edt/mds/infoprovider/se/se.config`:

If needed, modify the protocols that you use on the SE (examples: file,gridftp,...) and the data directory (example: /flatfiles/lcg).

Remember that the storage area (/flatfiles/lcg) needs to be exported to the CE and WN for the "file" access protocol.

5) Create an account needed for gridftp (example: lcg).

N.B. This account MUST be the same on the CE, WN and SE with the same uid and gid. ("groupadd -g 1000 lcg" , "useradd -u 1000 -g lcg lcg")

6) In the storage area you should have a dedicated area for each VO; for this release we will support only LCG. The owner of this area (/flatfiles/lcg) must be the same user used in item (5) ("lcg").

7) Ask the LCG-0 administrators on <http://cern.ch/grid-deployment/> to include your SE in the information index.

5.Install WN (Worker Nodes):

Currently the WNs must reside on a public network. As many WNs can be installed as necessary. For the WN we assume a batch system ("slave node") has already been installed and configured. We will give some instructions as an example for OpenPBS configuration.

1) Download the installation script `install_WN.sh` from the Grid Deployment web site:

wget http://grid-deployment.web.cern.ch/grid-deployment/bindir/LCG-0/install_WN.sh

2a) Check the allowed arguments:

```
sh install_WN.sh -help
```

2b) Run from bash (supply any necessary arguments):

```
sh install_WN.sh arguments > /root/WN_installation.log 2>&1
```

(the default download directory will be /root/rpm_lcg and will not be removed automatically)

3) Remember to create on the WN the same accounts used by the batch system on your CE (example: lcg).

N.B. This account MUST be the same on the CE, WN and SE with the same uid and gid ("groupadd -g 1000 lcg" , "useradd -u 1000 -g lcg lcg").

4) Home directories must be shared between CE and WNs.

You could use a disk server or you could export /home from the CE.

(example: mount -t nfs -o bg,intr,hard your_CE_machine:/home /home)

5) If you have a SE and you would like to use the "file" data access protocol, the SE storage area (example: /flatfiles/lcg) must be mounted.

6) If you have multiple WN machines to install, you may prefer using a "tar ball" for the other ones, instead of letting the installation script always download the necessary files from the LCG web server at CERN. In this case, after the successful installation of the first WN machine, create a tar ball from the files contained in the temporary installation directory and save it in a convenient location. For example:

```
cd /root/rpm_lcg
tar cvf /home/admin/WN.tar *
```

Then on the other WN machines you can run the installation as follows (assuming /home/admin/WN.tar is present now):

```
sh install_WN.sh --tarball /home/admin/WN.tar > /root/WN_installation.log 2>&1
```

-- Appendix - Installation of a PBS worknode -----

Install and configure PBS

(please refer to PBS documentation to obtain detailed informations)

```
download from our repository openpbs-exechost-2.3pl2-1.i386.rpm
rpm -ivh --nodeps openpbs-exechost-2.3pl2-1.i386.rpm
```


a. Edit `/usr/spool/PBS/server_name` to put your CE machine name (yourCE.yourDomain).

Edit `/usr/spool/PBS/mom_priv/config` to insert

```
$clienthost localhost
$clienthost your_CE_machine
$restricted your_CE_machine
$logevent 255
$ideal_load 1.6
$max_load 2.1
$usecp your_CE_machine:/home /home
```

where you are sharing home directories between CE and WNs.
Moreover in this example your machine has 2 CPU.

b. Remember to add the hostname of your WN to `/usr/spool/server_priv/nodes` and restart pbs daemon on your CE.

c. On WNs, add PBS ports in `/etc/services` as following:

```
# PBS
pbs 15001/tcp
pbs_mom 15002/tcp
pbs_remom 15003/tcp
pbs_remom 15003/udp
pbs_sched 15004/tcp
```

d. `/sbin/chkconfig pbs on, /etc/rc.d/init.d/pbs start`

Network requirements:

It is assumed that all machines (CE, SE, UI and WNs) will have unrestricted outbound TCP connectivity. Some incoming TCP connect requests also need to be allowed. As a guide the CE, SE and UI should be able to accept incoming TCP connections to the following ports:

2119	Globus Gatekeeper
2135	MDS info port
2169	FTree info port
2170	Information Index
2171	FTree info port
2811	GSI ftp server
6375	SE services
6376	SE services
7846	Logging & Bookkeeping

8881	Job Sub. Service (client)
9991	Job Sub. Service (server)
x -> y	Globus Job Manager*
15830	Locallogger

* example 14000 -> 15000 :

in addition to these fixed port numbers an open range (x -> y) for inbound connections is also needed.

These ports are required for some of the Globus services (at least for GSI-Ftp).

If at you site you are using a firewall the range x->y should match your firewall open range.

You can set this range on the CE and/or SE machines editing the file `/proc/sys/net/ipv4/ip_local_port_range`.

Contact:

Contact <http://cern.ch/grid-deployment/> with problems.