



LCG 2 INSTALL NOTES

Document identifier:

EDMS id: **n/a**

Version: **v2.2.0**

Date: **August 9, 2004**

Section: **GIS**

Document status: **Final**

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

File: **LCG2InstallNotes**

Abstract: These notes will assist you in installing the latest LCG-2 tag and upgrading from the previous tag.



CONTENTS

1. INTRODUCTION	3
1.1. WHAT IS LCG?	3
1.2. HOW TO JOIN LCG2?	3
1.3. HOW TO REPORT PROBLEMS	5
1.4. HOW TO SETUP YOUR SITE	5
1.5. NETWORK ACCESS	5
1.6. GENERAL NOTE ON SECURITY	5
1.7. SITES MOVING FROM LCG1 TO LCG2	6
1.8. CHANGES FROM LCG-2_1_0 TO LCG-2_2_0	6
1.9. CHANGES FROM LCG-2 ₀₀ to LCG-2 ₁₀	7
1.10. DOCUMENTATION	8
1.11. REGISTRATION	9
2. INTRODUCTION AND OVERALL SETUP	10
3. PREPARING THE INSTALLATION OF CURRENT TAG	12
4. STORAGE	14
5. FIREWALL CONFIGURATION	16
6. NODE INSTALLATION AND CONFIGURATION	17
6.1. COMMON STEPS	17
6.2. USERINTERFACE	17
6.3. RESOURCEBROKER	18
6.4. COMPUTINGELEMENT	18
6.5. MONNODE	19
6.6. WORKERNODE	19
6.7. STORAGEELEMENT	20
6.8. PLAINGRIS	20
6.9. BDII NODE	20
6.10. LCG-BDII NODE	20
6.11. REGIONAL MDS NODE	21
6.12. MYPROXY NODE	21
6.13. TESTING	21
7. APPENDIX A	22
7.1. SYNTAX FOR THE MDS_HOST_LIST VARIABLE	22
8. APPENDIX B	23
8.1. HOW TO CONFIGURE THE PBS SERVER ON A COMPUTINGELEMENT	23
9. APPENDIX C	25
9.1. HOW TO CONFIGURE THE MYSQL DATABASE ON A RESOURCEBROKER	25
10. APPENDIX D	26
10.1. PUBLISHING WN INFORMATION FROM THE CE	26
11. APPENDIX E	27
11.1. MODIFICATIONS TO YOUR SITE-CFG.H FILE	27



12. APPENDIX F	30
12.1. TESTING THE UI	30
12.2. TESTING THE CE AND WNS	36
12.3. TESTING THE SE	48
12.4. TESTING THE R-GMA	49
13. APPENDIX G	50
13.1. SITE INFORMATION NEEDED FOR THE CONTACT DATA BASE	50
14. APPENDIX H	52
14.1. LCG SITE CONFIGURATION DATABASE AND GRID OPERATION CENTER (GOC) ...	52
15. APPENDIX I	53
15.1. LCG-2 ₂₀	53
16. CHANGE HISTORY	59



1. INTRODUCTION

These notes will assist you in installing the latest LCG-2 tag and upgrading from the previous tag. The current tag is: LCG-2₂₀

The document is not a typical release note. It covers in addition some general aspects related to LCG2 operation and testing. If you are new to LCG or EGEE you should read th This document is intended for:

1. Sites that run LCG2 and need to upgrade to the current version
2. Sites that move from LCG1 to LCG2
3. Sites that join the LCG
4. Sites that operate LCG2

We realised that not many people read the large parts of the document. This is understandable since some of the material presented here is quite detailed. However, we recommend that you read at least the Introduction section and then then decide on the installation method. If you go for the manual installation you still have to read in addition the appendix on testing and the section related to the configuration of the firewall. The appendix on the required changes to the site-cfg file contains additional useful information and should be read by everyone starting and it is advisable to look for changes between versions. This appendix is describing an LCFGng configuration file, but the descriptions of the various parameters are important for the manual installation, too.

1.1. WHAT IS LCG?

This is best answered by material found on the projects web site <http://lcg.web.cern.ch/LCG/> . From there you can find information about the nature of the project and its goals. At the end of the introduction you can find a section that collects most of the references.

subsectionWhat is EGEE? EGEE and LCG are two project that are in many aspects closely related. Until the new flavour of software from the EGEE project is released LCG is used as the production platform for EGEE. For more information go to: <http://egee-intranet.web.cern.ch/egee-intranet/gateway.html>.

1.2. HOW TO JOIN LCG2?

If you want to join LCG and add resources to it you should contact the LCG deployment manager Ian Bird (<Ian.Bird@cern.ch>) to establish the contact with the project.

If you only want to use LCG you can follow the steps described in the LCG User Overview (http://lcg.web.cern.ch/LCG/peb/grid_deployment/user_intro.htm). The registration and initial training using the LCG-2 Users Guide (<https://edms.cern.ch/file/454439//LCG-2-Userguide.pdf>) should take about a week. However only 8 hours is related to working with the system, while the majority is waiting for the registration process with the VOs and the CA.

If you are interested in adding resources to the system you should first register as a user and subscribe to the LCG Rollout mailing list (<http://www.listserv.rl.ac.uk/archives/lcg-rollout.html>). In addition you need to contact the Grid Operation Centre (GOC) (<http://goc.grid-support.ac.uk/gridsite/gocmain/>) and get access to the GOC-DB for registering your resources with them. This registration is the basis for your system being present in their monitoring. It is mandatory to register



at least your service nodes in the GOC DB. It is not necessary to register all farm nodes. Please see Appendix H for a detailed description.

LCG has introduced a hierarchical support model for sites. Regions have primary sites (P-sites) that support the smaller centers in this region. If you do not know who is your primary site, please contact the LCG deployment manager Ian Bird. If you have identified your primary site you should fill the form that you find at the end of the guide in Appendix G and send it to your primary site AND to the deployment team at CERN (<support-lcg-deployment@cern.ch>). The site security contacts and sysadmins will receive material from the LCG security team that describes the security policies of LCG.

Discuss with the grid deployment team or with your primary site a suitable layout for your site. Various configurations are possible. Experience has shown that using at the beginning a standardised small setup and evolve from this to a larger more complex system is highly advisable. Typical layout for a minimal site is a user interface node (UI) which allows to submit jobs to the grid. This node will use the information system and resource broker either from the primary site, or the CERN site. A site that can provide resources will add a computing element (CE), that acts as a gateway to the computing resources and a storage element (SE), that acts as a gateway to the local storage. In addition a few worker nodes (WN) to provide the computing power can be added. Smaller sites will most likely add the RGMA monitoring node functionality to their SE, while medium to large sites should add a separate node as the MON node.

Large sites with many users that submit a large number of jobs will add a resource broker (RB). The resource broker distributes the jobs to the sites that are available to run jobs and keeps track of the status of the jobs. The RB uses for the resource discovery an information index (BDII). It is good practise to setup a BDII on each site that operates a RB. A complete site will add a Proxy server node that allows the renewal of proxy certificates. To save nodes while having a complete setup the manual and LCFGng based installation guides contain now the description of nodes that integrate the function of several nodes in one.

In case you don't find a setup described in this installation guide that meets your needs you should contact your primary site for further help. Another place to look for alternative configurations is the administration FAQs at <http://goc.grid.sinica.edu.tw/gocwiki/FrontPage>.

During the last few weeks we received several requests about how sites can add support for VOs that are not in the list of standard VOs that we support. The steps involved in adding a new VO are described on this web page: <http://grid-deployment.web.cern.ch/grid-deployment/cgi-bin/index.cgi?var=gis/vo-deploy>. In addition sites that support additional VOs have to add these VOs to their configuration files. Currently this is a slightly tedious operation because many steps are involved. However the tasks are conceptually simple and can be summarised for both the manual and the LCFGng based version by selecting one of the existing VOs and repeating the operation for the new VO. Some hints can be found in addition on the FAQ and gocwiki pages referenced at the end of the chapter. The procedure to setup a file catalogue service for a new VO is described on the gocwiki page (<http://goc.grid.sinica.edu.tw/gocwiki/FrontPage>).

The LCG middle ware has only very modest requirements on the hardware on which it can be installed. But keep in mind that a minimal configuration can be quite slow under load. The requirements of the HEP experiments for their productions are much more demanding and can be seen here <http://ibird.home.cern.ch/ibird/LCGMinResources.doc>. The minimal configuration is:

- CPU: For all nodes a Pentium with more than 500MHz
- Memory: More than 254MB
- Disc: The middleware uses about 1GB, in addition the RBs require at least 20GB for the storage of the sandboxes. The WNs need adequate scratch space of at least 5GB for each job run at the same time. The SEs storage size depends on the applications
- Shared Filesystem: A small shared file system is currently required for the storage of the experiments software.



- Network: A network card with at least 10Mbit is required.

After a site has been setup the site manager, or the support persons of the primary sites should run the initial tests that are described in the first part of the chapter on testing.

If these tests have been run successful the site should contact the deployment team via e-mail. The mail should contain the sites GIIS name and the hostname of the GIIS. To allow further testing the site will be added to a LCG-BDII which is used for testing new sites. Then the primary site, or the site managers can run the additional tests described.

When a site has passed these tests the site, or the primary site will announce this to the deployment team which then after a final round of testing will add the site to the list of production sites.

1.3. HOW TO REPORT PROBLEMS

The way problems are reported is currently changing. On the LCG user introduction page (http://lcg.web.cern.ch/LCG/peb/grid_deployment/user_intro.htm) you can find information on the current appropriate way to report problems. Before reporting a problem you should first try to consult your primary site. Many problems are currently reported to the rollout list. Internally we still use a Savannah based bug tracking tool that can be accessed via this link <https://savannah.cern.ch/bugs/?group=lcgoperation>.

1.4. HOW TO SETUP YOUR SITE

With this release you have the option to either install and configure your site using LCFGng, a fabric management tool that is supported by LCG, or to install the nodes following a manual step by step description which can be used as a basis to configure your local fabric management system.

For very small sites the manual approach has the advantage that no learning of the tool is required and no extra node needs to be maintained. In addition no re-installation of your nodes is required. However, the maintenance of the nodes will require more work and it is more likely to introduce hidden misconfiguration.

For medium to larger sites without their own fabric management tools using LCFGng can be an advantage. It is up to a site to decide which method is preferred.

The documentation for the manual installation can be found here:

<http://grid-deployment.web.cern.ch/grid-deployment/gis/release-docs/MIG-index.html>

All node types are supported. In case you decide to use the manual setup you should nevertheless have a look at parts of this document. For example the section about firewalls and testing are valid for both installation methods.

1.5. NETWORK ACCESS

The current software requires outgoing network access from all the nodes. And incoming on the RB, CE, and SE and the MyProxy server.

Some sites have gained experience with running their sites through a NAT. We can provide contact information of sites with experience of this setup.

To configure your firewall you should use the port table that we provide as a reference. Please have a look at the chapter on firewall configuration.

1.6. GENERAL NOTE ON SECURITY

While we provide in our repositories Kernel RPMs and use for the configuration certain versions it has to be pointed out that you have to make sure that you consider the kernel that you install as safe. If the provided default is not what you want please replace it.



We expect site manager to be aware of the relevant security related policies of LCG. A page that summarises this information has been prepared and can be accessed under: http://proj-lcg-security.web.cern.ch/proj-lcg-security/sites/for_sites.htm.

1.7. SITES MOVING FROM LCG1 TO LCG2

Since LCG2 is significantly different from both LCG1 and EDG, it is mandatory to study this guide even for administrators with considerable experience. In case you see the need to deviate from the described procedures please contact us.

Due to the many substantial changes w.r.t LCG1, updating a site from any of the LCG1 releases to LCG-2 is not possible in a reliable way.

A complete re-installation of the site is the only supported procedure.

Another change is related to the CVS repository used. For CERN internal reasons we had to move to a different server and switch to a different authorisation scheme. See <http://grid-deployment.web.cern.ch/grid-deployment/documentation/cvs-guide/> for details about getting access to the CVS repository.

For web based browsing the access to CVS is via <http://lcgdeploy.cvs.cern.ch/cgi-bin/lcgdeploy.cgi/>

As described later we changed for LCG2 the directory structure in CVS. There are now two relevant directories `lcg2` and `lcg2-sites`. The first contains common elements while the later contains the site specific information.

In addition to the installation via LCFGng all node types are now supported to be installed manually. In addition we started to provide descriptions for installing combined services on the same node to allow smaller sites a more economic installation.

If you move from LCG1 to LCG2 you should note that the structure of the information system has been changed significantly. The regional MDSs have disappeared and we introduced a complete rewritten BDII that uses a web based configuration.

1.8. CHANGES FROM LCG-2_1_0 TO LCG-2_2_0

The major change for this release is that R-GMA has been included. For those who want to learn a bit about how the Relational Grid Monitoring Architecture (R-GMA) is working and what it can be used for the following link is the reference: <http://www.r-gma.org/>.

R-GMA has a client and server. The client has been added the WN, CE, SE, RB and UI. The server should be installed on a new node, called a MON node. All large sites should install a MON node. Small sites (total machines 20) could install the R-GMA server on the SE, however it is advised to use a separate machine if possible. A detailed description can be found in the FAQ list at <http://goc.grid.sinica.edu.tw/gocwiki/AdministrationFaq>.

Both the RB and the Proxy nodes should now publish their service in the information system. Which should make it easier to locate them.

The BDII has been updated to fix problems found when submitting large sets of jobs. Another problem fixed by this release is related to the timeout behaviour for the http server that serves the configuration file.

The Workload Management System has been update to address several issues that have been found during the extensive usage by the HEP experiments during the last month.

The GridICE Monitoring has been updated and now has the capability to list details of jobs per VO. If upgrading from the previous version, the `edg-fmon-agent` on the monitored node must be restarted. The same applies to the `edg-fmon-server` on the local collector node.

From this release on we will add the full internal release notes in Appendix I.

Since the `edg-rm` tools for data management are being gradually replaced by the LCG Data Management Tools, in this release we added a set of environment variables that define for each site the default SE

per VO. This `WN_DEFAULT_SE_;``VOname;` parameter will point to the (close) SE that the VO prefers at the site.

In addition there are a few changes that should make the system more consistent. These changes are corrections to unclear documentation in the previous releases.

Please read, when upgrading to this release, the note on the correct normalisation of the queue length and CPU time that your site is publishing. This is described in Appendix E.

Since several sites have configured the MDS URL for the `edg-rm` tools in such a way that they point to a BDII that does not include the site itself, some consistency problems have affected these sites. We suggest that sites that use the `testZone` BDII at CERN (`lxn1189.cern.ch`) do not change their configuration. All other sites either switch for the configuration of this parameter on the WNs and UIs to this node or use BDIIs that are configured to use the site list published in <http://grid-deployment.web.cern.ch/grid-deployment/gis/lcg2-bdii/dteam/lcg2-all-sites.conf>.

Larger sites are requested to change the layout of their queues. While in the past all queues have been setup with the access enabled for all supported VOs this mode of operation is now creating problems for the HEP VOs. We suggest to setup the queues for the longest jobs on a per VO basis. The FAQ [Adding one queue per VO](http://goc.grid.sinica.edu.tw/gocwiki/AdministrationFaq) at <http://goc.grid.sinica.edu.tw/gocwiki/AdministrationFaq> describes the procedure.

Another problem observed in larger sites that have been handling large number of jobs is that the site GIIS is not reliable reporting the full information of the site. This has been tracked down to an interplay between MDS and some information providers. We suggest to run in addition to the site GIIS a BDII that publishes the same information. A configuration guide for this mode of operation can be found at the above mentioned FAQ list under the title: [Using the BDII as a site GIIS](#)

Please note that for the time being we recommend to run both in parallel. If you do this, please send us the new ldap URL for the BDII serving as a site GIIS.

For sites that control the outgoing high numbered ports with their fire wall we added the ports used by the HEP experiment software to the port table. Please have a look at the section on Firewall configuration.

We would like to point out that the HEP experiments can not use sites that do not provide space on a shared filesystem for the installation of their software. If your site currently configures the `VO_VOname_SW_DIRto.yoush`

1.9. CHANGES FROM LCG-2₀₀ to LCG-2₁₀

Again a major change was needed for the BDII. We encountered scalability problems when we added more than 40 sites to the system. The new version of the BDII will allow to scale the system to significantly more sites and a security problem regarding the information in the BDII has been solved with the new version. Simplified tests indicate the system should be capable to operate with a few hundred sites. Apart from installing the new BDII the web based configuration files have changed slightly. As a sample file the `testZone` one can be used it can be accessed via <http://grid-deployment.web.cern.ch/grid-deployment/gis/lcg2-testZone-new.conf>

The work load management system has seen some changes that have removed some of the scalability problems seen during the growing of LCG2 and includes several bug fixes.

VDT has been upgraded to the current release. The changes are mainly bug fixes.

In addition to the updated replica management software several utilities to improve the performance of registering files and moving files have been added. The documentation of these new packages should appear soon in the LCG2 Users Guide.

As a more transparent way to access data on the grid the GFAL library is now included. Basic documentation for this is available via the man pages, which include a small sample job. A small test job using can GFAL can be found in the chapter on testing too.

The man pages are available on the web at the following location: <http://grid-deployment.web.cern.ch/grid-deployment/gis/GFAL/GFALindex.html>

Realising that smaller sites have problems to justify the large number of service nodes we started to support nodes that integrate several services on one machine.

For sites that use LCFGng we support nodes that merge the RB/BDII and UI on one node. For very small sites that just want to get a first look at LCG we added a manual installation guide that for a node that integrates a UI, WN, SE and CE.

As usual, we have tried to improve the documentation. As part of this effort we started to collect symptoms of frequently problems and questions. Please visit the pages at the GOCs at RAL and Taipei. <http://goc.grid.sinica.edu.tw/gocwiki/FrontPage> contains links to troubleshooting guides and FAQs.

Since the diversity of the sites in LCG is steadily increasing we can't cover all variants in this guide. Several alternative configurations will be covered by entries in the FAQ pages.

There is a very important change concerning the configuration of the local batch systems and their queues. In the past the default settings have been sufficient for the experiments to run short and medium long jobs. There is now an extra section on configuration of the queues in the CE and PBS configuration section. Please read this carefully and configure your systems parameter correctly.

In addition the experiments have put forward their requests for memory, local scratch space and storage on the local SEs. This is summarised in this document: <http://ibird.home.cern.ch/ibird/LCGMinResources.doc>.

Changes from LCG-2 beta to LCG-2.0.0

In the previous beta-release the new LCG-BDII node type has been introduced. And for some time the two information system structures have been operated in parallel. Since we expect many sites to move from LCG1 to LCG2 we will switch now permanently to the new layout which we describe later in some detail.

The new LCG-BDII does not use any more on the Regional MDSes but collects information directly from the Site GIISes. The list of existing sites and their addresses are downloaded from a pre-defined web location. See notes in the BDII specific section in this document for installation and configuration. This layout will allow sites and VOs to configure their own super- or subset of the LCG2 resources.

A new Replica Manager client has also been introduced in the previous version. This is the only client which is compatible with the current version of the RLS server, so file replication at your site will not work till you have updated to this release.

1.10. DOCUMENTATION

- [D1] LCG Project Homepage:
<http://lcg.web.cern.ch/LCG/>
- [D2] Starting point for users of the LCG infrastructure:
http://lcg.web.cern.ch/LCG/peb/grid_deployment/user_intro.htm
- [D3] LCG-2 User's Guide:
<https://edms.cern.ch/file/454439//LCG-2-Userguide.pdf>
- [D4] LCFGng server installation guide:
http://lcgdeploy.cvs.cern.ch/cgi-bin/lcgdeploy.cgi/lcg2/docs/LCFGng_server_install.txt
- [D5] LCG-2 Manual Installation Guide:
<http://grid-deployment.web.cern.ch/grid-deployment/documentation/manual-installation/>
- [D6] LCG GOC Mainpage:
<http://goc.grid-support.ac.uk/gridsite/gocmain/>
- [D7] CVS User's Guide:
<http://grid-deployment.web.cern.ch/grid-deployment/documentation/cvs-guide/>



1.11. REGISTRATION

[R1] LCG rollout list:

<http://www.listserv.rl.ac.uk/archives/lcg-rollout.html>

- join the list

[R2] Get the Certificate and register in VO:

<http://lcg-registrar.cern.ch/>

- read LCG Usage Rules
- choose your CA and contact them to get USER certificate (for some CAs online certificate request is possible)
- load your certificate into web browser (read instructions)
- choose your VO and register (LCG Registration Form)

[R3] GOC Database:

<http://goc.grid-support.ac.uk/gridsite/db-auth-request/>

- apply for access to the GOCDB

[R4] CVS read-write access and site directory setup:

Send a mail to Louis Poncet (<Louis.Poncet@cern.ch>)

- prepare and send a NAME for your site following the schema <domain>-<organization>[-<section>] (e.g. es-Barcelona-PIC, ch-CERN, it-INFN-CNAF)

[R5] Site contact database:

Send a mail to the Support Group (<support-lcg-deployment@cern.ch>)

- fill in the form in Appendix G and send it

[R6] Report bugs and problems with installation:

<https://savannah.cern.ch/bugs/?group=lcgoperation>



2. INTRODUCTION AND OVERALL SETUP

In this text we will assume that you are already familiar with the LCFGng server installation and management.

Access to the manual installation guides is given via the following link: <http://grid-deployment.web.cern.ch/grid-deployment/gis/release-docs/MIG-index.html>

The sources for the html and pdf files are available from the CVS repository in the documentation directory.

Note for sites which are already running LCG1: due to the incompatible update of several configuration objects, a LCFG server cannot support both LCG1 and LCG-2 nodes. If you are planning to re-install your LCG1 nodes with LCG-2, then the correct way to proceed is:

1. kill the rdxprof process on all your nodes (or just switch your nodes off if you do not care about the extra down-time at your site);
2. update your LCFG server using the objects listed in the LCG-2 release;
3. prepare the new configuration files for your site as described in this document;
4. re-install all your nodes.

If you plan to keep your LCG1 site up while installing a new LCG-2 site, then you will need a second LCFG server. This is a matter of choice. The LCG1 installation is of very limited use if you setup the LCG-2 site since several core components are not compatible anymore.

Files needed for the current LCG-2 release are available from a CVS server at CERN. This CVS server contains the list of rpms to install and the LCFGng configuration files for each node type. The CVS area, called "lcg2", can be reached from <http://lcgdeploy.cvs.cern.ch/cgi-bin/lcgdeploy.cgi/>

Note1: at the same location there is another directory called "lcg-release": this area is used for the integration and certification software, NOT for production. Please ignore it!

Note2: documentation about access to this CVS repository can be found in <http://grid-deployment.web.cern.ch/grid-deployment/documentation/cvs-guide/>

In the same CVS location we created an area, called lcg2-sites, where all sites participating to LCG-2 should store the configuration files used to install and configure their nodes. Each site manager will find there a directory for their site with a name in the format

```
<domain>-<city>-<institute>
```

or

```
<domain>-<organization>[-<section>]
```

(e.g. es-Barcelona-PIC, ch-CERN, it-INFN-CNAF): this is where all site configuration files should be uploaded. Site managers that install a site

Site managers are kindly asked to keep these directories up-to-date by committing all changes they do to their configuration files back to CVS so that we will be able to keep track of the status of each site at any given moment. Once a site reaches a consistent working configuration, site managers should create a CVS tag which will allow them to easily recover configuration information if needed. Tag names should follow the following convention: The tags of the LCG-2 modules are:

```
LCG-<RELEASE>
```

e.g. LCG-1_1_1 for software release 1.1.1

If you tag your local configuration files, the tag name must contain a reference to the lcg2 release in use at the time. The format to use is:



LCG-<RELEASE>_<SITENAME>_<DATE>_<TIME>

e.g. LCG-1.1.1_CERN_20031107_0857 for configuration files in use at CERN on November 7th, 2003, at 8:57 AM. The lcg2 release used for this example is 1.1.1.

To activate a write-enabled account to the CVS repository at CERN please get in touch with Louis Poncet (<Louis.Poncet@cern.ch>).

Judit Novak (<Judit.Novak@cern.ch>) or Markus Schulz (<Markus.Schulz@cern.ch>) are the persons to contact if you do not find a directory for your site or if you have problems uploading your configuration files to CVS.

If you just want to install a site, but not join LCG, you can get anonymous read access to the repository. As described in the CVS access guide set the CVS environment variables.

- Set CVS_RSH to :

```
> setenv CVS\_RSH ssh
```

- Set CVSROOT to :

```
> setenv CVSROOT :pserver:anonymous@lcgdeploy.cvs.cern.ch:/cvs/lcgdeploy
```

All site managers have in any case to subscribe to and monitor the LCG-Rollout mailing list. Here all issues related to the LCG deployment, including announcements of updates and security patches, are discussed. You can subscribe from the following site: <http://cclrlsv.RL.AC.UK/archives/lcg-rollout.html> and click on the "Join or leave the list"

This is the main source for communicating problems and changes.



3. PREPARING THE INSTALLATION OF CURRENT TAG

The current LCG tag is ---> LCG-2.1.0 <---

In the following instructions/examples, when you see the <CURRENT_TAG> string, you should replace it with the name of the tag defined above.

To install it, check it out on your LCFG server with

```
> cvs checkout -r <CURRENT_TAG> -d <TAG_DIRECTORY> lcg2
```

Note: the "-d <TAG_DIRECTORY> " will create a directory named <TAG_DIRECTORY> and copy there all the files. If you do not specify the -d parameter, the directory will be a subdirectory of the current directory named lcg2.

The default way to install the tag is to copy the content of the rpmlist subdirectory to the /opt/local/linux/7.3/rpms directory on the LCFG server. This directory is NFS-mounted by all client nodes and is visible as /export/local/linux/7.3/rpms

Go to the directory where you keep your local configuration files. If you want to create a new one, you can check out from CVS any of the previous tags with:

```
> cvs checkout -r <YOUR_TAG> -d <LOCAL_DIR> lcg2/<YOUR_SITE>
```

If you have not committed any configuration file yet or if you want to use the latest (HEAD) versions, just omit the "-r <YOUR_TAG> " parameter.

Now cd to <LOCAL_DIR> and copy there the files from <TAG_DIRECTORY>/examples: following the instructions in the 00README file, those in the example files themselves, and those reported below in this document you should be able to create an initial version of the configuration files for your site. If you have problems, please contact your reference primary site.

NOTE: if you already have localised versions of these files, just compare them with the new templates to verify that no new parameter needs to be set. Be aware that there are several critical differences between LCG1 and LCG-2 site-cfg.h files, so apply extra care when updating this file.

IMPORTANT NOTICE: If you have a CE configuration file from LCG1, it probably includes the definition of the secondary regional MDS for your region. This is now handled by the ComputingElement-cfg.h configuration file and can be configured directly from the site-cfg.h file. See Appendix E for details.

To download all the rpms needed to install this version you can use the updaterep command. In <TAG_DIRECTORY>/tools you can find 2 configuration files for this script: updaterep.conf and updaterep_full.conf. The first will tell updaterep to only download the rpms which are actually needed to install the current tag, while updaterep_full.conf will do a full mirror of the LCG rpm repository. Copy updaterep.conf to /etc/updaterep.conf and run the updaterep command. 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 the repository area 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 all CA-related settings. Changes and updates of these settings will be announced on the LCG-Rollout mailing list.

To make sure that all the needed object rpms are installed on your LCFG server, you should use the lcfng_server_update.pl script, also located in <TAG_DIRECTORY>/tools. This script will report which rpms are missing or have the wrong version and will create the /tmp/lcfng_server_update_script.sh script which you can then use to fix the server configuration. Run it in the following way:



```
lcfgng_server_update.pl <TAG_DIRECTORY>/rpmlist/lcfgng-common-rpm.h  
/tmp/lcfgng_server_update_script.sh  
lcfgng_server_update.pl <TAG_DIRECTORY>/rpmlist/lcfgng-server-rpm.h  
/tmp/lcfgng_server_update_script.sh
```

WARNING: please always give a look to `/tmp/lcfgng_server_update_script.sh` and verify that all rpm update commands look reasonable before running it.

In the source directory you should give a look to the `redhat73-cfg.h` file and see if the location of the rpm lists (`updaterpms.rpmdir`) and of the rpm repository (`updaterpms.rpmdir`) are correct for your site (the defaults are consistent with the instructions in this document). If needed, you can redefine these paths from the `local-cfg.h` file.

In `private-cfg.h` you can (must!) replace the default root password with the one you want to use for your site:

```
+auth.rootpwd <CRYPTED_PWD> <— replace with your own crypted password
```

To obtain `<CRYPTED_PWD>` using the MD5 encryption algorithm (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. E.g.

```
> openssl passwd -1  
Password: <- write clear text password here  
$1$iPJJehjc$rtV/651890BaPinzkb58z1 <- <CRYPTED_PWD> string
```

To finalize the adaptation of the current tag to your site you should edit your `site-cfg.h` file. If you already have a `site-cfg.h` file that you used to install any of the LCG1 releases, you can find a detailed description of the modifications to this file needed for the new tag in Appendix E below.

WARNING: the template file `site-cfg.h.template` assumes you want to run the PBS batch system without sharing the `/home` directory between the CE and all the WNs. This is the recommended setup.

There may be situations when you have to run PBS in traditional mode, i.e. with the CE exporting `/home` with NFS and all the WNs mounting it. This is the case, e.g., if your site does not allow for host based authentication. To revert to the traditional PBS configuration you can edit your `site-cfg.h` file and comment out the following two lines:

```
#define NO_HOME_SHARE  
...  
#define CE_JM_TYPE lcgpbs
```

In addition to this, your WN configuration file should include this line:

```
#include CFGDIR/UsersNoHome-cfg.h"
```

just after including `Users-cfg.h` (please note that BOTH `Users-cfg.h` AND `UsersNoHome-cfg.h` must be included).

4. STORAGE

In the current version LCG still uses the "Classical SE" model. This consists into a storage system (either a real MSS or just a node connected to some disks) which exports a GridFTP interface. Information about the SE must be published by a GRIS registered to the Site GIIS.

If your SE is a completely independent node connected to a bunch of disks (these can either be local or mounted from a disk server) then you can install this node using the example SE_node file: this will install and configure on the node all needed services (GridFTP server, GRIS, authentication system).

If you plan to use a local disk as the main storage area, you can include the flatfiles-dirs-SECLASSIC-cfg.h file: LCFG will take care of creating all needed directories with the right access privileges.

If on the other hand your SE node mounts the storage area from a disk server, then you will have to create all needed directories and set their privileges by hand. Also, you will have to add to the SE node configuration file the correct commands to NFS-mount the area from the disk server.

As an example, let's assume that your disk server node is called <server> and that it exports area <diskarea> for use by LCG. On your SE you want to mount this area as /storage and then allow access to it via GridFTP.

To this end you have to go through the following steps:

1. in site-cfg.h define

```
#define CE_CLOSE_SE_MOUNTPOINT /storage
```

2. in the SE_node configuration file add the lines to mount this area from <server>:

```
EXTRA(nfsmount.nfsmount) storage
nfsmount.nfsdetails_storage /storage <server>:<diskarea> rw
```

3. once the SE node is installed and /storage has been mounted, create all VO directories, one per supported VO, giving read/write access to the corresponding group. For VO <vo>:

```
> mkdir /storage/<vo>
> chgrp <vo> /storage/<vo>
> chmod g+w /storage/<vo>
```

A final possibility is that at your site a real mass storage system with a GridFTP interface is already available (this is the case for the CASTOR MSS at CERN). In this case, instead of installing a full SE, you will need to install a node which act as a front-end GRIS for the MSS, publishing to the LCG information system all information related to the MSS.

This node is called a PlainGRIS and can be installed using the PG_node file from the examples directory. Also, a few changes are needed in the site-cfg.h file. Citing from site-cfg.h.template:

```
/* For your storage to be visible from the grid you must have a GRIS which
 * publishes information about it. If you installed your SE using the classical
 * SE configuration file provided by LCG (StorageElementClassic-cfg.h) then a
 * GRIS is automatically started on that node and you can leave the default
 * settings below. If your storage is based on a external MSS system which
 * only provides a GridFTP interface (an example is the GridFTP-enabled CASTOR
 * service at CERN), then you will have to install an external GRIS server
 * using the provided PlainGRIS-cfg.h profile. In this case you must define
 * SE_GRIS_HOSTNAME to point to this node and define the SE_DYNAMIC_CASTOR
 * variable instead of SE_DYNAMIC_CLASSIC (Warning: defining both variables at
 * the same time is WRONG!).
```

```
*
* Currently the only supported external MSS is the GridFTP-enabled CASTOR used
* at CERN.
*/
#define SE_GRIS_HOSTNAME      SE_HOSTNAME
#define SE_DYNAMIC_CLASSIC
/* #define SE_DYNAMIC_CASTOR */
```



5. FIREWALL CONFIGURATION

If your LCG nodes are behind a firewall, you will have to ask your network manager to open a few "holes" to allow external access to some LCG service nodes.

A complete map of which port has to be accessible for each service node is provided in file `lcg-port-table.pdf` in the `lcg2/docs` directory. <http://lcgdeploy.cvs.cern.ch/cgi-bin/lcgdeploy.cgi/lcg2/docs/lcg-port-table.pdf>.

If possible don't allow ssh access to your nodes from outside your site.

6. NODE INSTALLATION AND CONFIGURATION

In the `<TAG_DIRECTORY>/tools` you can find a new version of the `do_mkxprof.sh` script. A detailed description of how this script works is contained in the script itself. You are of course free to use your preferred call to the `mkxprof` command but note that running `mkxprof` as a daemon is NOT recommended and can easily lead to massive catastrophes if not used with extreme care: do it at your own risk.

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://<Your_LCFGng_Server>/status/
```

and clicking on the name of the node with a faulty configuration (a small red bug should be shown beside the node name).

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" mentioned above (`LCFGng_server_install.txt`).

When the initial installation completes (expect two automatic reboots in the process), each node type requires a few manual steps, detailed below, to be completely configured. After completing these steps, some of the nodes need a final reboot which will bring them up with all the needed services active. The need for this final reboot is explicitly stated among the node configuration steps below.

6.1. 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`.

6.2. USERINTERFACE

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

6.3. RESOURCEBROKER

- Configure the MySQL database. See detailed recipe in Appendix C at the end of this document
- Reboot the node

6.4. 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;_;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
```

6.5. 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
```

6.6. 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.

6.7. 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.

6.8. PLAINGRIS

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

6.9. BDII NODE

The BDII node using the regional GIISes is no longer supported. It has been replaced by the LCG-BDII.

6.10. LCG-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>

If you don't want to maintain your own sites file you can use this URL to start.

Change the URL to the the URL of the file. Add or remove sites. To make the BDIIs realize the change you have to change the Date field. Don't forget this.



6.11. REGIONAL MDS NODE

No more regional MDS nodes are installed since the system based on the LCG-BDII doesn't require them any more.

6.12. 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/O
```

6.13. TESTING

IMPORTANT NOTICE: if `/home` is NOT shared between CE and WNs (this is the default configuration) due to the way the new jobmanager works, a `globus-job-run` command will take at least 2 minutes. Even in the configuration with shared `/home` the execution time of `globus-job-run` will be slightly longer than before. Keep this in mind when testing your system.

To perform the standard tests (`edg-job-submit` & co.) you need to have your certificate registered in one VO and to sign the LCG usage guidelines.

Detailed information on how to do these two steps can be found in: <http://lcg-registrar.cern.ch/> If you are working in one of the four LHC experiments, then ask for registration in the corresponding VO, otherwise you can choose the "LCG Deployment Team" (aka DTeam) VO.

A test suite which will help you in making sure your site is correctly configured is now available. This software provides basic functionality tests and various utilities to run automated sequences of tests and to present results in a common HTML format.

Extensive on-line documentation about this test suite can be found in

- <http://grid-deployment.web.cern.ch/grid-deployment/tstg/docs/>
LCG-Certification-help

All tests related to job submission should work out of the box.

In Appendix F you can find some core tests that should be run certify that the site is providing the core functionality.

7. APPENDIX A

7.1. SYNTAX FOR THE MDS_HOST_LIST VARIABLE

This appendix is no longer needed since with the introduction of the LCG-BDII no configuration related to regional MDSs is needed.

8. APPENDIX B

8.1. 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
```

9. APPENDIX C

9.1. HOW TO CONFIGURE THE MYSQL DATABASE ON A 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).

10. APPENDIX D

10.1. PUBLISHING WN INFORMATION FROM THE CE

When submitting a job, users of LCG are supposed to state in their jdl the minimal hardware resources (memory, scratch disk space, CPU time) required to run the job. These requirements are matched by the RB with the information on the BDII to select a set of available CEs where the job can run.

For this schema to work, each CE must publish some information about the hardware configuration of the WNs connected to it. This means that site managers must collect information about WNs available at the site and insert it in the information published by the local CE.

The procedure to do this is the following:

- choose a WN which is "representative" of your batch system (see below for a definition of "representative") and make sure that the chosen node is fully installed and configured. In particular, check if all expected NFS partitions are correctly mounted.
- on the chosen WN run the following script as root, saving the output to a file.

```
@-----  
#!/bin/bash  
echo -n 'hostname: '  
host `hostname -f` | sed -e 's/ has address.*//'  
echo "Dummy: `uname -a`"  
echo "OS_release: `uname -r`"  
echo "OS_version: `uname -v`"  
cat /proc/cpuinfo /proc/meminfo /proc/mounts  
df  
@-----
```

- copy the obtained file to /opt/edg/var/info/edg-scl-desc.txt on your CE, replacing any pre-existing version.
- restart the GRIS on the CE with

```
> /etc/rc.d/init.d/globus-mds restart
```

Definition of "representative WN": in general, WNs are added to a batch system at different times and with heterogeneous hardware configurations. All these WNs often end up being part of a single queue, so that when an LCG job is sent to the batch system, there is no way to ask for a specific hardware configuration (note: LSF and other batch systems offer ways to do this but the current version of the Globus gatekeeper is not able to take advantage of this possibility). This means that the site manager has to choose a single WN as "representative" of the whole batch cluster. In general it is recommended that this node is chosen among the "least powerful" ones, to avoid sending jobs with heavy hardware requirements to under-spec nodes.



11. APPENDIX E

11.1. MODIFICATIONS TO YOUR SITE-CFG.H FILE

In this appendix we describe changes between the previous version and the current that have to be applied to the site-cfg.h file. However we have left in the instructions for the adjustment of some additional parameters since they are frequently misconfigured. If your site is still at LCG-1 and you want to upgrade to LCG-2 follow first the instructions of the notes included in the LCG-2.1.0 release. These can be found on the release page <http://grid-deployment.web.cern.ch/grid-deployment/cgi-bin/index.cgi?var=releases>. If you start from the site-cfg.h template in the examples directory all the parameters are already included. It is nevertheless worthwhile to read this appendix since it adds additional comments to clarify the correct settings.

1. change the published version to LCG-2.2.0

```
#define SITE_EDG_VERSION LCG-2_2_0
```

2. the BDII configuration section includes the URL to the LCG-BDII configuration file. The below mentioned URL should be used on all UIs and WNs. If you setup an application specific RB you can override this parameter in the node configuration file by adding the following line.

```
#define SITE_BDII_URL          http://YOUR_WEB_SERVER/YOUR_BDII_CONFIGURATION_FILE.conf

#define SITE_BDII_URL http://grid-deployment.web.cern.ch/grid-deployment/\
    gis/lcg2-bdii/dteam/lcg2-all-sites.conf
```

3. the CE_QUEUES parameter is a space-separated list (in the past it was a comma-separated list):

```
#define CE_QUEUES              short long infinite
```

If you decide, as suggested in the notes for upgrading to LCG-2.2.0, to setup individual queues as described in the FAQ [Adding a Queue for a VO](http://goc.grid.sinica.edu.tw/gocwiki/AdministrationFaq) (<http://goc.grid.sinica.edu.tw/gocwiki/AdministrationFaq>) you have to change this list.

4. all VO software related parameters have been removed from the CE_IP_RUNTIMEENV parameter definition:

```
#define CE_IP_RUNTIMEENV      LCG-2 LCG-2\_1\_0 LCG-2\_2\_0 R-GMA
```

5. A section is needed to configure the disk areas where VO managers can install VO-related software. Please make sure that you setup these areas if you support one of the HEP experiments. For other VOs you should get into contact with them and verify whether they require this or not.

```
/* Area on the WN for the installation of the experiment software */
/* If on your WNs you have predefined shared areas where VO managers can
   pre-install software, then these variables should point to these areas.
   If you do not have shared areas and each job must install the software,
   then these variables should contain a dot ( . )
*/
/* #define WN_AREA_ALICE    /opt/exp_software/alice */
/* #define WN_AREA_ATLAS   /opt/exp_software/atlas */
```

```
/* #define WN_AREA_CMS      /opt/exp_software/cms */
/* #define WN_AREA_LHCB    /opt/exp_software/lhcb */
/* #define WN_AREA_DTEAM   /opt/exp_software/dteam */
#define WN_AREA_ALICE     .
#define WN_AREA_ATLAS     .
#define WN_AREA_CMS       .
#define WN_AREA_LHCB      .
#define WN_AREA_DTEAM     .
```

6. The new monitoring system based on GridICE is now included in the default setup. To configure it add to your site-cfg.h file the "GRIDICE MONITORING" section from site-cfg.h.template and edit it (if needed) for your site.
7. To allow the experiments a proper match between the local resources and their jobs, some care has to be taken to configure the following values correctly.

```
/* CE InformationProviders: SpecInt 2000 */
#define CE_IP_SI00          380
/* CE InformationProviders: SpecFloat 2000 */
#define CE_IP_SF00         400
```

The whole issue is a bit complicated and we have put together the following as a guideline for selecting the right values. Since we can't set both values correctly we suggest to set the SpecFloat to 0.

- Sites that have a homogeneous farm:
 - SpecInt: Use the correct value, or take a number that is close to your node specification from the list at the end
- Sites that have a heterogeneous farm and use internal scaling. This means scaling at the LRMS level of the allowed execution time on the CPU:
 - SpecInt: The specInt of the reference machine that corresponds to the published time
- Sites with heterogeneous farms, but without any internal scaling
 - SpecInt: The SpecInt of the slowest machine in the farm

If you have very different nodes (factor of 5 or more) consider splitting the farm .

The SpecInt value can be taken either from <http://www.spechbench.org/osg/cpu2000/results/cint2000.html>, or from this short list:

		SI2K
P4	2.4 GHz	852
P3	1.0 GHz	461
P3	0.8 GHz	340
P3	0.6 GHz	270

Please note that some of the HEP experiments run very long jobs. If you support them your longest queue should be able to handle 48 hours jobs on a node correspondin to a 1GHz PIV,

8. In this release we include the support for R-GMA. The following parameters have to be set:

```
#define SITE_RGMA_SERVER      mon-box.host.invalid
#define GRID_RGMA_INFO_CATALOG  lcgic01.gridpp.rl.ac.uk
```

Do not change the second parameter. For the first one you have to select a MON node either on your site or a remote one. Smaller sites are advised to integrate the MON box functionality on your local SE.

9. To add the functionality of an individual default SE for each of the VOs the following parameters have to be set.

```
#define WN_DEFAULT_SE_ALICE  alice.host.invalid
#define WN_DEFAULT_SE_ATLAS  atlas.host.invalid
#define WN_DEFAULT_SE_CMS    cms.host.invalid
#define WN_DEFAULT_SE_LHCB   lhcb.host.invalid
#define WN_DEFAULT_SE_DTEAM  dteam.host.invalid
#define WN_DEFAULT_SE_SIXT   dteam.host.invalid
```

If you have only one close SE at your site set the node name to the hostname of your SE. This will only affect the lcg data management tools. If you are in doubt contact the experiments that you support to learn about their preferences.

12. APPENDIX F

This is a collection of basic commands that can be run to test the correct setup of a site. These tests are not meant to be a replacement of the test tools provided by LCG test team. Extensive documentation covering this can be found here:

<http://grid-deployment.web.cern.ch/grid-deployment/tstg/docs/>
LCG-Certification-help

The material in this chapter should enable the site administrator to verify the basic functionality of the site.

- Testing the UI
- Testing the CE and WNs
- Testing the SE

Not included in this release:

- Testing the RB
- Testing the BDII
- Testing the Proxy

12.1. TESTING THE UI

The main tools used on a UI are:

1. Tools to manage certificates and create proxies
2. Tools to deal with the submission and status retrieval of jobs
3. Client tools of the data management. These include tools to transport data and to query the replica location service

1. Create a proxy

The `grid-proxy-init` command and the other commands used here should be in your path.

```
[adc0014] ~ > grid-proxy-init
Your identity: /C=CH/O=CERN/OU=GRID/CN=Markus Schulz 1319
Enter GRID pass phrase for this identity:
Creating proxy ..... Done
Your proxy is valid until: Mon Apr  5 20:53:38 2004
```

2. Run simple jobs

Check that `globus-job-run` works. First select a CE that is known to work. Have a look at the GOC DB and select the CE at CERN.

```
[adc0014] ~ > globus-job-run lxn1181.cern.ch /bin/pwd
/home/dteam002
```



What can go wrong with this most basic test? If your VO membership is not correct you might be not in the grid-mapfile. In this case you will see some errors that refer to grid security.

Next is to see if the UI is correctly configured to access a RB. Create the following files for these tests:

testJob.jdl this contains a very basic job description.

```
Executable = "testJob.sh";
StdOutput = "testJob.out";
StdError = "testJob.err";
InputSandbox = {"/testJob.sh"};
OutputSandbox = {"testJob.out", "testJob.err"};
#Requirements = other.GlueCEUniqueID == "lxn1181.cern.ch:2119/jobmanager-lcgpbs-short";
```

testJob.sh contains a very basic test script

```
#!/bin/bash
date
hostname
echo "*****"
echo "env | sort"
echo "*****"
env | sort
echo "*****"
echo "mount"
echo "*****"
mount
echo "*****"
echo "rpm -q -a | sort"
echo "*****"
/bin/rpm -q -a | sort

sleep 20
date
```

run the following command to see which sites can run your job

```
adc0014] ~/TEST > edg-job-list-match --vo dteam testJob.jdl
```

the output should look like:

```
Selected Virtual Organisation name (from --vo option): dteam
Connecting to host lxn1177.cern.ch, port 7772

*****
COMPUTING ELEMENT IDs LIST
The following CE(s) matching your job requirements have been found:

*CEId*
hik-lcg-ce.fzk.de:2119/jobmanager-pbspro-lcg
hotdog46.fnal.gov:2119/jobmanager-pbs-infinite
hotdog46.fnal.gov:2119/jobmanager-pbs-long
hotdog46.fnal.gov:2119/jobmanager-pbs-short
lcg00125.grid.sinica.edu.tw:2119/jobmanager-lcgpbs-infinite
```

```
lcg00125.grid.sinica.edu.tw:2119/jobmanager-lcgpbs-long
lcg00125.grid.sinica.edu.tw:2119/jobmanager-lcgpbs-short
lcgce02.ifaes.es:2119/jobmanager-lcgpbs-infinite
lcgce02.ifaes.es:2119/jobmanager-lcgpbs-long
lcgce02.ifaes.es:2119/jobmanager-lcgpbs-short
lxn1181.cern.ch:2119/jobmanager-lcgpbs-infinite
lxn1181.cern.ch:2119/jobmanager-lcgpbs-long
lxn1184.cern.ch:2119/jobmanager-lcglstf-grid
tbn18.nikhef.nl:2119/jobmanager-pbs-qshort
wn-04-07-02-a.cr.cnae.infn.it:2119/jobmanager-lcgpbs-dteam
tbn18.nikhef.nl:2119/jobmanager-pbs-qlong
lxn1181.cern.ch:2119/jobmanager-lcgpbs-short
*****
```

If an error is reported rerun the command using the `-debug` option. Common problems are related to the RB that has been configured to be used as the default RB for the node. To test if the UI works with a different UI you can run the command using configuration files that overwrite the default settings. Configure the two files to use for the test a known working RB. The RB at CERN that can be used is: `lxn1177.cern.ch` The file that contains the VO dependent configuration has to contain the following:

```
lxn1177.vo.conf

[
VirtualOrganisation = "dteam";
NSAddresses = "lxn1177.cern.ch:7772";
LBAddresses = "lxn1177.cern.ch:9000";
## HLR location is optional. Uncomment and fill correctly for
## enabling accounting
#HLRLocation = "fake HLR Location"
## MyProxyServer is optional. Uncomment and fill correctly for
## enabling proxy renewal. This field should be set equal to
## MYPROXY_SERVER environment variable
MyProxyServer = "lxn1179.cern.ch"
]
```

and the common one:

```
lxn1177.conf

[
rank = - other.GlueCEStateEstimatedResponseTime;
requirements = other.GlueCEStateStatus == "Production";
RetryCount = 3;
ErrorStorage = "/tmp";
OutputStorage = "/tmp/jobOutput";
ListenerPort = 44000;
ListenerStorage = "/tmp";
LoggingTimeout = 30;
LoggingSyncTimeout = 30;
LoggingDestination = "lxn1177.cern.ch:9002";
# Default NS logger level is set to 0 (null)
# max value is 6 (very ugly)
NSLoggerLevel = 0;
```

```
DefaultLogInfoLevel = 0;
DefaultStatusLevel = 0;
DefaultVo = "dteam";
]
```

Then run the list match with the following options:

```
edg-job-list-match -c `pwd`/lxn1177.conf --config-vo `pwd`/lxn1177.vo.conf \
testJob.jdl
```

If this works you should have investigate the configuration of the RB that is selected by default from your UI or the associated configuration files.

If the job-list-match is working you can submit the test job using:

```
edg-job-submit --vo dteam testJob.jdl
```

The command returns some output like:

```
Selected Virtual Organisation name (from --vo option): dteam
Connecting to host lxn1177.cern.ch, port 7772
Logging to host lxn1177.cern.ch, 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://lxn1177.cern.ch:9000/0b6EdeF6dJlnHkKByTkc_g
*****
```

In case the output of the command has a significant different structure you should rerun it and add the `-debug` option. Save the output for further analysis.

Now wait some minutes and try to verify the status of the job using the command:

```
edg-job-status https://lxn1177.cern.ch:9000/0b6EdeF6dJlnHkKByTkc_g
```

repeat this until the job is in the status: Done (Success)

If the job doesn't reach this state, or gets stuck for longer periods in the same state you should run a command to access the logging information. Please save the output.

```
edg-job-get-logging-info -v 1 \
https://lxn1177.cern.ch:9000/0b6EdeF6dJlnHkKByTkc_g
```

Assuming that the job has reached the desired status please try to retrieve the output:

```
edg-job-get-output https://lxn1177.cern.ch:9000/0b6EdeF6dJlnHkKByTkc_g
```

```
Retrieving files from host: lxn1177.cern.ch ( for https://lxn1177.cern.ch:9000/
0b6EdeF6dJlnHkKByTkc_g )
```

```
*****
                                JOB GET OUTPUT OUTCOME

Output sandbox files for the job:
- https://lxn1177.cern.ch:9000/0b6EdeF6dJlnHkKByTkc_g
have been successfully retrieved and stored in the directory:
/tmp/jobOutput/markusw_0b6EdeF6dJlnHkKByTkc_g
*****
```

Check that the given directory contains the output and error files.

One common reason for this command to fail is that the access privileges for the jobOutput directory are not correct, or the directory has not been created.

If you encounter a problem rerun the command using the `-debug` option.

3. Data management tools

Test that you can reach an external SE. Run the following simple command to list a directory at one of the CERN SEs.

```
edg-gridftp-ls gsiftp://castorgrid.cern.ch/castor/cern.ch/grid/dteam
```

You should get a long list of files.

If this command fails it is very likely that your firewall setting is wrong.

Next see which resources you can see via the information system you should run:

```
[adc0014] ~/TEST/STORAGE > edg-replica-manager -v --vo dteam pi
edg-replica-manager starting..
Issuing command : pi
Parameters:
Call replica catalog printInfo function
VO used          : dteam
default SE       : lxn1183.cern.ch
default CE       : lxn1181.cern.ch
Info Service     : MDS
```

..... and a long list of CEs and SEs and their parameters.

Verify that the default SE and CE are the nodes that you want to use. Make sure that these nodes are installed and configured before you conduct the tests of more advanced data management functions.

If you get almost nothing back you should check the configuration of the replica manager. Use the following command to get the BDII that you are using: `grep mds.url /opt/edg/var/etc/edg-replica-manager/edg-replica-manager.conf` this should return the name and port of the BDII that you intended to use. For the CERN UIs you would get:

```
mds.url=ldap://lxn1178.cern.ch:2170
```

Convince yourself that this is the address of a working BDII that you can reach.

```
ldapsearch -LLL -x -H ldap://<node specified above>:2170 \
-b "mds-vo-name=local,o=grid"
```

this should return something starting like this:

```
dn: mds-vo-name=local,o=grid
objectClass: GlobusStub

dn: Mds-Vo-name=cernlcg2,mds-vo-name=local,o=grid
objectClass: GlobusStub

dn: Mds-Vo-name=nikheflcgprod,mds-vo-name=local,o=grid
objectClass: GlobusStub

dn: GlueSEUniqueID=lxn1183.cern.ch,Mds-Vo-name=cernlcg2,mds-vo-name=local,o=gr
id
objectClass: GlueSETop
objectClass: GlueSE
objectClass: GlueInformationService
objectClass: Gluekey
objectClass: GlueSchemaVersion
GlueSEUniqueID: lxn1183.cern.ch
GlueSEName: CERN-LCG2:disk
GlueSEPort: 2811
GlueInformationServiceURL: ldap://lxn1183.cern.ch :2135/Mds-Vo-name=local,o=gr
id
GlueForeignKey: GlueSLUniqueID=lxn1183.cern.ch
.....
```

In case the query doesn't return the expected output verify that the node specified is a BDII and that the node is running the service.

As a crosscheck you can try to repeat the test with one of the BDII's at CERN. In the GOC DB you can identify the BDII for the production and the test zone. Currently these are lxn1178.cern.ch for the production system and lxn1189.cern.ch for the test Zone.

Before the `edg-replica-manager -v -vo dteam pi` command and the `edg-gridftp-ls` commands are not working it makes no sense to conduct further tests.

Assuming that this functionality is well established the next test is to move a local file from the UI to the default SE and register the file with the replica location service.

Create a file in your home directory. To make tracing this file easy the file should be named according to the scheme:

```
testFile.<SITE-NAME>.txt
```

the file should be generated using the following script:

```
#!/bin/bash
echo "*****"
echo "hostname: " `hostname` " date: " `date`
echo "*****"
```

the command to move the file to the default SE is:

```
edg-replica-manager -v --vo dteam cr file://`pwd`/testFile.<SiteName>.txt \
-l lfn:testFile.<SiteName>.`date +%m.%d.%y:%H:%M:%S`
```

The command returns if everything is setup correctly a line with:



guid:98ef70d6-874d-11d8-b575-8de631cc17af

Save the guid for further reference and the expanded lfn. We will refer to these as YourGUID and YourLFN.

In case this command failed you should keep the output and analyze it with your support contact. There are various reasons why this command has failed.

Now we check that the RLS knows about your file. This is done by using the listReplicas (lr) option.

```
edg-replica-manager -v --vo dteam lr lfn:YourLFN
```

this command should return a string with a format similar to:

```
sfn://lxn1183.cern.ch/storage/dteam/generated/2004-04-06/file92c9f455-874d-11d8-b575-8de631cc17af
ListReplicas successful.
```

as before, report problems to your primary site.

If the RLS knows about the file the next test is to transport the file back to your UI. For this we use the cp option.

```
edg-replica-manager -v --vo dteam cp lfn:YourLFN file://`pwd`/testBack.txt
```

this should create in the current working directory a file named testBack.txt. List this file.

With this you tested most of the core functions of your UI. Many of these functions will be used to verify the other components of your site.

12.2. TESTING THE CE AND WNS

We assume that you have setup a local CE running a batch system. On most sites the CE provides two major services. For the information system the CE runs the site GIIS. The site GIIS is the top node in the hierarchy of the site and via this service the other resources of the site are published to the grid.

To test the working of the site GIIS you can run an ldap query of the following form. Inspect the output with some care. Are the computing resources (queues, etc.) correctly reported? Can you find the local SE?. Do these numbers make sense?

```
ldapsearch -LLL -x -H ldap://lxn1181.cern.ch:2135 -b "mds-vo-name=cernlcg2,o=grid"
```

replace lxn1181.cern.ch with your site's GIIS hostname and cernlcg2 with the name that you have assigned to your site GIIS.

If nothing is reported try to restart the MDS service on the CE.

Now verify that the GRIS on the CE is operating correctly: Here again the command for the CE at CERN.

```
ldapsearch -LLL -x -H ldap://lxn1181.cern.ch:2135 -b "mds-vo-name=local,o=grid"
```

One common reason for this to fail is that the information provider on the CE has as problem. Convince yourself that MDS on the CE is up and running. Run on the CE the qstat command. If this command doesn't return there might be a problem with one of the worker nodes WNs, or PBS. Have a look at the following link that covers some aspects on trouble shooting PBS on the GRID. <http://goc.grid.sinica.edu.tw/gocwiki/TroubleShootingHistory>

The next step is to verify that you can run jobs on the CE. For the most basic test no registration with the information system is needed. However tests can be run much easier if the resource is registered in the information system. For these tests the testZone BDII and RB have been setup at CERN. Forward your site GIIS name and host name to the deployment team for registration.

Initial tests that work without registration.

First tests from a UI of your choice:

As described in the subsection covering the UI tests the first test is a test of the fork jobmanger.

```
adc0014] ~ > globus-job-run <YourCE> /bin/pwd
```

Frequent problems that have been observed are related to the authentication. Check that the CE has a valid host certificate and that your DN can be found in the grid-mapfile.

Next logon to your CE and run a local PBS job to verify that PBS is working. Change your id to a user like dteam001. In the home directory create the following file:

```
test.sh
```

```
#!/bin/bash
```

```
echo "Hello Grid"
```

run: `qsub test.sh` this will return a job ID of the form: 16478.lxn1181.cern.ch you can use `qstat` to monitor the job. However it is very likely that the job has finished before your have queried the status. PBS will place two files in your directory:

```
test.sh.o16478 and test.sh.e16478 These contain the stdout and stderr
```

Now try to submit to one of your PBS queues that are available on the CE. The following command is an example for a site that runs a PBS without shared home directories. The short queue is used. It can take some minutes until the command returns.

```
globus-job-run <YourCE>/jobmanager-lcgpbs -queue short /bin/hostname  
lxshare0372.cern.ch
```

The next test submits a job to your CE by forcing the broker to select the queue that your have chosen. You can use the testJob JDL and script that has been used before for the UI tests.

```
edg-job-submit --debug --vo dteam -r <YourCE>:2119/jobmanager-lcgpbs-short \  
testJob.jdl
```

The `--debug` option should only be used if you have been confronted with problems.

Follow the status of the job and as before try to retrieve the output. A quite common problem is that the output can't be retrieved. This problem is related to some inconsistency of ssh keys between the CE and the WN. See <http://goc.grid.sinica.edu.tw/gocwiki/TroubleShootingHistory> and the CE/WN configuration.

If your UI is not configured to use a working RB you can, as described in the UI testing subsection use configuration files to use the testZone RB.

For further tests get registered with the testZone BDII. As described in the subsection on joining LCG2 you should send your CE's hostname and the site GIIS name to the deployment team.

The next step is to take the testJob.jdl that you have created for the verification of your UI. Remove the comment from the last line of the file and modify it to reflect your CE.

```
Requirements = other.GlueCEUniqueID == "<YourCE>:2119/jobmanager-lcgpbs-short";
```

Now repeat the `edg-job-list-match -vo dteam testJob.jdl` command known from the UI tests. This output should just show one resource.

The remaining tests verify that core of the data management is working from the WN and that the support for the experiment software installation as described in <https://edms.cern.ch/file/412781/SoftwareInstallation.pdf> is working correctly. The tests you can do to verify the later are limited if you are not mapped to software manager for your VO. To test the data management functions your local default SE has to be setup and tested. Of course you can assume the SE working and run the tests before testing the SE.

Add an argument to the JDL that allows to identify the site. The jdl file should look like:

```
testJob_SW.jdl
```

```
Executable = "testJob.sh";
StdOutput = "testJob.out";
StdError = "testJob.err";
InputSandbox = {"/testJob.sh"};
OutputSandbox = {"testJob.out", "testJob.err"};
Requirements = other.GlueCEUniqueID == "lxn1181.cern.ch:2119/jobmanager-lcgpbs-short";
Arguments = "CERNPBS" ;
```

replace the name of the site and the CE and queue names to reflect your settings.

The first script to run collects some configuration information from the WN and test the user software installation area.

```
testJob.sh
```

```
#!/bin/bash
echo "+++++"
echo "          $1          " `hostname` " " `date`
echo "+++++"
echo "the environment on the node"
echo " "
env | sort
echo "+++++"
echo "+++++"
echo "+++++"
echo "software path for the experiments"
env | sort | grep _SW_DIR
echo "+++++"
echo "+++++"
echo "+++++"
echo "mount"
mount
echo "+++++"
echo "=====
echo "veryfiy that the software managers of the supported VOs can \
write and the users read"
echo "DTEAM ls -l " $VO_DTEAM_SW_DIR
ls -dl $VO_DTEAM_SW_DIR
echo "ALICE ls -l " $VO_ALICE_SW_DIR
ls -dl $VO_ALICE_SW_DIR
echo "CMS ls -l " $VO_CMS_SW_DIR
```

```
ls -dl $VO_CMS_SW_DIR
echo "ATLAS ls -l " $VO_ATLAS_SW_DIR
ls -dl $VO_ATLAS_SW_DIR
echo "LHCB ls -l " $VO_LHCB_SW_DIR
ls -dl $VO_LHCB_SW_DIR
echo "====="
echo "====="
echo "====="
echo "====="
echo "cat /opt/edg/var/etc/edg-replica-manager/edg-replica-manager.conf"
echo "====="
cat /opt/edg/var/etc/edg-replica-manager/edg-replica-manager.conf
echo "====="
echo "====="
echo "====="
echo "====="
echo "rpm -q -a | sort "
rpm -q -a | sort
echo "====="
date
```

Run this job as described in the subsection on testing UIs. Retrieve the output and verify that the environment variables for the experiment software installation is correctly set and that the directories for the VOs that you support are mounted and accessible.

In the `edg-replica-manager.conf` file reasonable default CEs and SEs should be specified: The output for the CERN PBS might serve as an example:

```
localDomain=cern.ch
defaultCE=lxn1181.cern.ch
defaultSE=wacdr002d.cern.ch
```

Then a working BDII node has to be specified as the MDS top node: For the CERN production this is currently:

```
mds.url=ldap://lxn1178.cern.ch:2170
mds.root=mds-vo-name=local,o=grid
```

Please keep the output of this job as a reference. It can be helpful if problems have to be located.

Next we test the data management. For this the default SE should be working. The following script will do some operations similar to those used on the UI.

LCG currently provides two sets of data management tools. For testing them we provide two test scripts. The first one listed is using the `edg-replica-manager` tools. The second script uses the `lcg` tools which are based on the `GFAL` library.

We first test that we can access a remote SE via simple `gridftp` commands. Then we test that the replica manager tools have access to the information system. This is followed by exercising the data moving capabilities between the WN, the local SE and between a remote SE and the local SE. Between the commands we run small commands to verify that the RLS service knows about the location of the files.

Submit the job via `edg-job-submit` and retrieve the output. Read the file containing `stdout` and `stderr`. Keep the files for reference.

Here now a listing of `testJob.sh`:



```
#!/bin/bash

TEST_ID=`hostname -f`-`date +%y%m%d%H%M`
REPORT_FILE=report
rm -f $REPORT_FILE
FAIL=0
user=`id -un`
echo "Test Id: $TEST_ID"
echo "Running as user: $user"

if [ "x$1" == "x" ]; then
    echo "Usage: $0 <VO>"
    exit 1
else
    VO=$1
fi

grep mds.url= /opt/edg/var/etc/edg-replica-manager/edg-replica-manager.conf
echo
echo "Can we see the SE at CERN?"
set -x
edg-gridftp-ls --verbose gsiftp://castorgrid.cern.ch/castor/cern.ch/grid/$VO > /dev/null
result=$?
set +x

if [ $result == 0 ]; then
    echo "We can see the SE at CERN."
    echo "ls CERN SE: PASS" >> $REPORT_FILE
else
    echo "Error: Can not see the SE at CERN."
    echo "ls CERN SE: FAIL" >> $REPORT_FILE
    FAIL=1
fi

echo
echo "Can we see the information system?"
set -x
edg-replica-manager -v --vo $VO pi
result=$?
set +x

if [ $result == 0 ]; then
    echo "We can see the Information System."
    echo "RM Print Info: PASS" >> $REPORT_FILE
else
    echo "Error: Can not see the Information System."
    echo "RM Print Info: FAIL" >> $REPORT_FILE
    FAIL=1
fi

lname=testFile.$TEST_ID.txt
```

```
rm -rf $lfname
cat <<EOF > $lfname
*****
Test Id: $TEST_ID

File used for the replica manager test
*****

EOF
myLFN="rep-man-test-$TEST_ID"

echo
echo "Move a local file to the default SE and register it with an lfn."
set -x
edg-replica-manager -v --vo $VO cr file://`pwd`/$lfname -l lfn:$myLFN
result=$?
set +x

if [ $result == 0 ]; then
    echo "Local file moved the the SE."
    echo "Move file to SE: PASS" >> $REPORT_FILE
else
    echo "Error: Could not move the local file to the SE."
    echo "Move file to SE: FAIL" >> $REPORT_FILE
    FAIL=1
fi

echo
echo "List the replicas."
set -x
edg-replica-manager -v --vo $VO lr lfn:$myLFN
result=$?
set +x

if [ $result == 0 ]; then
    echo "Replica listed."
    echo "RM List Replica: PASS" >> $REPORT_FILE
else
    echo "Error: Can not list replicas."
    echo "RM List Replica: FAIL" >> $REPORT_FILE
    FAIL=1
fi

lf2=$lfname.2
rm -rf $lf2

echo
echo "Get the file back and store it with a different name."
set -x
edg-replica-manager -v --vo $VO cp lfn:$myLFN file://`pwd`/$lf2
```

```
result=$?
diff $lfname $lf2
set +x

if [ $result == 0 ]; then
    echo "Got get file."
    echo "RM copy: PASS" >> $REPORT_FILE
else
    echo "Error: Could not get the file."
    echo "RM copy: FAIL" >> $REPORT_FILE
    FAIL=1
fi

if [ "x`diff $lfname $lf2`" == "x" ]; then
    echo "Files are the same."
else
    echo "Error: Files are different."
    FAIL=1
fi

echo
echo "Replicate the file from the default SE to the CASTOR service at CERN."
set -x
edg-replica-manager -v --vo $VO replicateFile lfn:$myLFN -d castorgrid.cern.ch
result=$?
edg-replica-manager -v --vo $VO lr lfn:$myLFN
set +x

if [ $result == 0 ]; then
    echo "File replicated to Castor."
    echo "RM Replicate: PASS" >> $REPORT_FILE

else
    echo "Error: Could not replicate file to Castor."
    echo "RM Replicate: FAIL" >> $REPORT_FILE
    FAIL=1
fi

echo
echo "3rd party replicate from castorgrid.cern.ch to the default SE."
set -x
ufilesfn=`edg-rm --vo $VO lr lfn:TheUniversalFile.txt | grep lxn1183`
edg-replica-manager -v --vo $VO replicateFile $ufilesfn
result=$?
edg-replica-manager -v --vo $VO lr lfn:TheUniversalFile.txt
set +x

if [ $result == 0 ]; then
    echo "3rd party replicate succeeded."
    echo "RM 3rd party replicate: PASS" >> $REPORT_FILE
else
```

```
    echo "Error: Could not do 3rd party replicate."
    echo "RM 3rd party replicate: FAIL" >> $REPORT_FILE
    FAIL=1
fi

rm -rf TheUniversalFile.txt

echo
echo "Get this file on the WN."
set -x
edg-replica-manager -v --vo $VO cp lfn:TheUniversalFile.txt file://`pwd`/TheUniversalFile.txt
result=$?
set +x

if [ $result == 0 ]; then
    echo "Copy file succeeded."
    echo "RM copy: PASS" >> $REPORT_FILE
else
    echo "Error: Could not copy file."
    echo "RM copy: FAIL" >> $REPORT_FILE
    FAIL=1
fi

defaultSE=`grep defaultSE /opt/edg/var/etc/edg-replica-manager/edg-replica-manager.conf | cut -d= -f2`

# Here we have to use a small hack. In case that we are at CERN we will never remove the
if [ $defaultSE = lxn1183.cern.ch ]
then
    echo "I will NOT remove the master copy from: " $defaultSE
else
    echo
    echo "Remove the replica from the default SE."
    set -x
    edg-replica-manager -v --vo $VO del lfn:TheUniversalFile.txt -s $defaultSE
    result=$?
    edg-replica-manager -v --vo $VO lr lfn:TheUniversalFile.txt
    set +x

    if [ $result == 0 ]; then
    echo "Deleted file."
    echo "RM delete: PASS" >> $REPORT_FILE
    else
    echo "Error: Could not do Delete."
    echo "RM delete: FAIL" >> $REPORT_FILE
    FAIL=1
    fi
fi

echo "Cleaning Up"
rm -f $lfname $lf2 TheUniversalFile.txt
```




```
if [ $FAIL = 1 ]; then
    echo "Replica Manager Test Failed."
    exit 1
else
    echo "Replica Manager Test Passed."
    exit 0
fi
```

The following test is using the lcg data management tools.

```
#!/bin/bash
echo "*****"
echo "    Test of the LCG Data Management Tools          "
echo "*****"
TEST_ID=`hostname -f`-`date +%y%m%d%H%M`
REPORT_FILE=report
FAIL=0
user=`id -un`
echo "Test Id: $TEST_ID"
echo "Running as user: $user"

if [ "x$1" == "x" ]; then
    echo "Usage: $0 <VO>"
    exit 1
else
    VO=$1
fi

echo "LCG_GFAL_INFOSYS      " $LCG_GFAL_INFOSYS
echo "VO_DTEAM_DEFAULT_SE   " $VO_DTEAM_DEFAULT_SE

echo
echo "Can we see the information system?"
set -x
ldapsearch -H ldap://$LCG_GFAL_INFOSYS -x -b o=grid | grep numE
result=$?
set +x

if [ $result == 0 ]; then
    echo "We can see the Information System."
    echo "LDAP query:  PASS" >> $REPORT_FILE
else
    echo "Error: Can not see the Information System."
    echo "LDAP query:  FAIL" >> $REPORT_FILE
    FAIL=1
fi

lfname=testFile.$TEST_ID.txt
rm -rf $lfname
cat << EOF > $lfname
*****
Test Id: $TEST_ID
```



File used for the replica manager test

EOF

```
myLFN="rep-man-test-$(TEST_ID)"
echo
echo "Move a local file to the default SE and register it with an lfn."
set -x
lcg-cr -v --vo $VO -d $VO_DTEAM_DEFAULT_SE -l lfn:$myLFN file://`pwd`/$lfname
result=$?
set +x

if [ $result == 0 ]; then
    echo "Local file moved the the default SE."
    echo "Move file to default SE: PASS" >> $REPORT_FILE
else
    echo "Error: Could not move the local file to the default SE."
    echo "Move file to default SE: FAIL" >> $REPORT_FILE
    FAIL=1
fi

echo
echo "List the replicas."
set -x
lcg-lr --vo $VO lfn:$myLFN
result=$?
set +x

if [ $result == 0 ]; then
    echo "Replica listed."
    echo "LCG List Replica: PASS" >> $REPORT_FILE
else
    echo "Error: Can not list replicas."
    echo "LCG List Replica: FAIL" >> $REPORT_FILE
    FAIL=1
fi

lf2=$lfname.2
rm -rf $lf2

echo
echo "Get the file back and store it with a different name."
set -x
lcg-cp -v --vo $VO lfn:$myLFN file://`pwd`/$lf2
result=$?
diff $lfname $lf2
set +x

if [ $result == 0 ]; then
```

```
    echo "Got get file."
    echo "LCG copy: PASS" >> $REPORT_FILE
else
    echo "Error: Could not get the file."
    echo "LCG copy: FAIL" >> $REPORT_FILE
    FAIL=1
fi

if [ "x`diff $lfname $lf2`" == "x" ]; then
    echo "Files are the same."
else
    echo "Error: Files are different."
    FAIL=1
fi

echo
echo "Replicate the file from the default SE to the CASTOR service at CERN."
set -x
lcg-rep -v --vo $VO -d castorgrid.cern.ch lfn:$myLFN
result=$?
lcg-lr --vo $VO lfn:$myLFN
set +x

if [ $result == 0 ]; then
    echo "File replicated to Castor."
    echo "LCG Replicate: PASS" >> $REPORT_FILE
else
    echo "Error: Could not replicate file to Castor."
    echo "LCG Replicate: FAIL" >> $REPORT_FILE
    FAIL=1
fi

echo
echo "3rd party replicate from castorgrid.cern.ch to the default SE."
set -x
ufilesfn=`lcg-lr --vo $VO lfn:TheUniversalFile.txt | grep lxn1183`
lcg-rep -v --vo $VO -d $VO_DTEAM_DEFAULT_SE $ufilesfn
result=$?
lcg-lr --vo $VO lfn:TheUniversalFile.txt
set +x

if [ $result == 0 ]; then
    echo "3rd party replicate succeeded."
    echo "LCG 3rd party replicate: PASS" >> $REPORT_FILE
else
    echo "Error: Could not do 3rd party replicate."
    echo "LCG 3rd party replicate: FAIL" >> $REPORT_FILE
    FAIL=1
fi
```

```
rm -rf TheUniversalFile.txt

echo
echo "Get this file on the WN."
set -x
lcg-cp -v --vo $VO lfn:TheUniversalFile.txt file://`pwd`/TheUniversalFile.txt
result=$?
set +x

if [ $result == 0 ]; then
    echo "Copy file succeeded."
    echo "LCG copy: PASS" >> $REPORT_FILE
else
    echo "Error: Could not copy file."
    echo "LCG copy: FAIL" >> $REPORT_FILE
    FAIL=1
fi

defaultSE=VO_DTEAM_DEFAULT_SE

# Here we have to use a small hack. In case that we are at CERN we will never remove the
if [ $defaultSE = lxn1183.cern.ch ]
then
    echo "I will NOT remove the master copy from: " $defaultSE
else
    echo
    echo "Remove the replica from the default SE."
    set -x
    lcg-del -v --vo $VO -s $defaultSE lfn:TheUniversalFile.txt
    result=$?
    lcg-lr --vo $VO lfn:TheUniversalFile.txt
    set +x

    if [ $result == 0 ]; then
        echo "Deleted file."
        echo "LCG delete: PASS" >> $REPORT_FILE
    else
        echo "Error: Could not do Delete."
        echo "LCG delete: FAIL" >> $REPORT_FILE
        FAIL=1
    fi
fi

echo "Cleaning Up"
rm -f $lfn1 $lfn2 TheUniversalFile.txt

if [ $FAIL = 1 ]; then
    echo "LCG Data Manager Test Failed."
    exit 1
else
    echo "LCG Data Test Passed."
```

```
exit 0
fi
```

12.3. TESTING THE SE

If the tests described to test the UI and the CE on a site have run successful then there is no additional test for the SE needed. We describe here some of the common problems that have been observed related to SEs.

In case the SE can't be found by the edg-replica-manager tools the SE GRIS might be not working, or not registered with the site GIIS.

To verify that the SE GRIS is working you should run the following ldapsearch. Note that the hostname that you use should be the one of the node where the GRIS is located. For mass storage SEs it is quite common that this is not the the SE itself.

```
ldapsearch -LLL -x -H ldap://lxn1183.cern.ch:2135 -b "mds-vo-name=local,o=grid"
```

If this returns nothing or very little the MDS service on the SE should be restarted. If the SE returns some information you should carefully check that the VOs that require access to the resource are listed in the GlueSAAccessControlBaseRule field. Does the information published in the GlueSEAccessProtocolType fields reflect your intention? Is the GlueSEName: carrying the extra "type" information?

The next major problem that has been observed with SEs is due to a mismatch with what is published in the information system and what has been implemented on the SE.

Check that the gridmap-file on the SE is configured to support the VOs that are published the GlueSAAccessControlBaseRule fields.

Run a ldapsearch on your site GIIS and compare the information published by the local CE with what you can find on the SE. Interesting fields are: GlueSEName, GlueCESEBindSEUniqueID, GlueCESEBindCEAccesspoint

Are the access-points for all the supported VOs created and is the access control correctly configured?

The edg-replica-manager command printInfo summarizes this quite well. Here is an example for a report generated for a classic SE at CERN.

```
SE at CERN-LCG2 :
      name : CERN-LCG2
      host : lxn1183.cern.ch
      type : disk
      accesspoint : /storage
      VOs : dteam
      VO directories : dteam:dteam
      protocols : gsiftp,rfio
```

to test the gsiftp protocol in a convenient way you can use the edg-gridftp-ls and edg-gridftp-mkdir commands. You can use the globus-url-copy command instead. The -help option describes the syntax to be used.

Run on your UI and replace the host and accesspoint according to the report for your SE:

```
edg-gridftp-ls --verbose gsiftp://lxn1183.cern.ch/storage
drwxrwxr-x    3 root      dteam          4096 Feb 26 14:22 dteam
```

and:

```
edg-gridftp-ls --verbose gsiftp://lxn1183.cern.ch/storage/dteam
drwxrwxr-x   17 dteam003 dteam          4096 Apr  6 00:07 generated
```

if the globus-gridftp service is not running on the SE you get the following message back: error a system call failed (Connection refused)

If this happens restart the globus-gridftp service on your SE.

Now create a directory on your SE.

```
edg-gridftp-mkdir gsiftp://lxn1183.cern.ch/storage/dteam/t1
```

Verify that the command ran successful with:

```
edg-gridftp-mkdir gsiftp://lxn1183.cern.ch/storage/dteam/t1
edg-gridftp-ls --verbose gsiftp://lxn1183.cern.ch/storage/dteam/
```

Verify that the access permissions for all the supported VOs are correctly set.

12.4. TESTING THE R-GMA

R-GMA comes with two testing scripts. These can be used on any node that has the R-GMA client installed (CE, SE, RB, UI, MON). On any of these nodes run the following two commands.

```
/opt/edg/sbin/test/edg-rgma-check
/opt/edg/sbin/test/edg-rgma-run-examples
```

The first script will check that R-GMA has been configured correctly. The most important thing is that it can connect to the servlets.

```
Checking the status of the servlets...
```

```
SchemaServlet OK
RegistryServlet OK
StreamProducerServlet OK
LatestProducerServlet OK
ConsumerServlet OK
```

The second script will test the operation of R-GMA. It will create a producer, try to insert some data and check that this has been inserted correctly.

If either test fails. The common reasons are that the servlets are down or that there is a firewall problem. Check that the values for the Servlets and Registry in are correct in the file `/opt/edg/var/edg-rgma/rgma.props`. Take the URL from this file, append `/getStatus` to the end and use a browser to connect to the servlet. If you can not connect to the servlet, check the URL again. Try to restart the tomcat server on the MON node

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

If there are problems starting the servlets some information can be found in the tomcat logs.

```
/var/tomcat4/logs/catalina.out
```



13. APPENDIX G

13.1. SITE INFORMATION NEEDED FOR THE CONTACT DATA BASE

Please fill and send to your primary site and the CERN deployment team (<support-lcg-deployment@cern.ch>).

===== START =====

0) Preferred name of your site

I. Communication:

=====

a) Contact email for the site

b) Contact phone for the site

c) Reachable during which hours

d) Emergency phone for the site

e) Site (computer/network) security contact for your site

f0) Official name of your institute

f1) Name and title/role of individual(s) responsible for computer/network security at your site

f2) Personal email for f1)



f3) Telephone for f1)

f4) Telephone for emergency security incident response
(if different from f3)

f5) Email for emergency security incident response (listbox preferred)

g) Write access to CVS

The LCG CVS repository is currently moved to a different CVS server.
To access this server a CERN AFS account is required. If you have none
please contact Louis Poncet (Louis.Poncet@cern.ch)

AFS account at CERN:

II) Site specific information

a) Domain

e) CA that issued host certificates for your site

===== END =====



14. APPENDIX H

This has been provided by David Kant (<D.Kant@rl.ac.uk>).

14.1. LCG SITE CONFIGURATION DATABASE AND GRID OPERATION CENTER (GOC)

The GOC will be responsible for monitoring the grid services deployed through the LCG middleware at your site.

Information about the site is managed by the local site administrator. The information we require are the site contact details, list of nodes and IP addresses, and the middleware deployed on those machines (EDG, LCG1, LCG2 etc)

Access to the database is done through a web browser (https) via the use of an X.509 certificate issued by a trusted LCG CA .

GOC monitoring is done hourly and begins with an SQL query of the database to extract your site details. Therefore, it is important to ensure that the information in the database is ACCURATE and UP-TO-DATE.

To request access to the database, load your certificate into your browser and go to:

```
http://goc.grid-support.ac.uk/gridsite/db-auth-request/
```

The GOC team will then create a customised page for your site and give you access rights to these pages. This process should take less than a day and you will receive an email confirmation. Finally, you can enter your site details:

```
https://goc.grid-support.ac.uk/gridsite/db/index.php
```

The GOC monitoring pages displaying current status information about LCG2:

```
http://goc.grid-support.ac.uk/gridsite/gocmain/
```



15. APPENDIX I

For each of the releases we keep a detailed internal release note. These are given here for reference. Since we maintain them in a text file we add them verbatim.

15.1. LCG-2₂₀

From: Laurence <Laurence.Field@cern.ch>
Date: Fri Aug 6, 2004 14:16:16 Europe/Zurich
To: Markus Schulz <markus.schulz@cern.ch>
Subject: release notes

Major points:

=====

- R-GMA has been included.
- VOMS has been worked on a lot but it isn't considered as solid enough for inclusion in the release. In particular there are remaining problems with synchronization of the LDAP and VOMS servers (gridmapfile generation) that have to be resolved before it can be released.
- dCache situation has not changed. We have received new RPMs a few days ago which should fix a few problems. RPMs have been installed and tests started. Since some new problems have been found already, we will be releasing a new version soon.
- Monitoring, GridICE Now has the capability to list details of jobs per VO. To present the monitoring information the GridICE server has to be updated.
Note: If upgrading from the previous version, the edg-fmon-agent on the monitored node must be restarted. The same applies to the edg-fmon-server on the local collector node.
- Bug fixes for WMS
No major changes in functionality, but numerous bug fixes.

Summary of changes with respect to the LCG-2_1_0

=====

- Globus:

No change, remains VDT 1.1.14-4.lcg1
- Condor G:

Upgrade from condor-6.6.0-2.edg5 to condor-6.6.0-2.edg6 to fix the following bug:
3716 - 'gridmanager can repeatedly try to restart jobmanagers'
- Information System:

new lcfg components to configure bdii (1-1.0.9-1)
- Information Providers:



patch 188 - lcg-info-dynamic-dcache updated

- Workload Management System:

- Workload Management System

Upgraded from lcg2.1.25 -> lcg2.1.32 with the following changes:

- 2716 - 'network server can become unresponsive'
- 3252 - 'Some suggestions for the edg-job-status '--all'' query'
- 3546 - 'edg-job-get-output for many files always fails'
- 3566 - 'UIutils.py truncates hostname when port not specified in config'
- 3589 - change of JAVA version
- 2682 - 'workload manager ranking queries' (minor revision to previous patch)
- 3258 - 'edg-wl-ns start takes a long time due to unneeded chown -R' (revision to previous patch)
- 3546 - 'edg-job-get-output for many files always fails' (previous patch needed correcting)
- 3666 - 'Too many files matched in bash by /tmp/dglogd_sock_*
- 3807 - 'LogMonitor must not crash on bad Condor-G log files'
- 3880 - 'renewal daemon - minimum condor limit sometimes not checked'
- 3882 - 'renewal daemon - proxy sharing between jobs' (LCG change)
- 3891 - 'Partially closed connections kept open by gatekeeper/jobmanager' (workaround)
- 3895 - 'Job canceled on receipt of ULOG_GLOBUS_RESOURCE_DOWN' (workaround)
- 3896 - 'renewd memory usage'
- 3897 - 'edg-wl-logd credential watch'
- 3900 - 'WM needs fully qualified HOSTNAME in environment' (workaround)
- 3905 - 'LHCb jobs failing at NIKHEF because of exhausted disk space' (LCG change)
- 3923 - 'LM sometimes not resubmitting jobs'
- 3931 - 'Suggest a local proxy expiration check for WMS jobs' (LCG change)
- 3973 - 'Initial working directory for WMS jobs & cleanup of job directories' (LCG change)
- 3990 - 'bug checking rm in configure of workload'
- 4039 - 'File descriptor leak in job submission API + fix'
- 4047 - 'logd sometimes logs connection failures' (LCG change)
- 4070 - 'FuzzyRanking (stochasticRankSelector)' (LCG change)
- 4109 - 'Certain forms of 'rank' fail'
- 4126 - 'NS can crash with 'Pipe Closed'' (workaround)

The wlconfig object was also updated:

```
edg-lcfg-wlconfig-1.0.20-lcg1      -->      edg-lcfg-wlconfig-1.0.21-lcg1
edg-lcfg-wlconfig-defaults-s1-1.0.20-lcg1  -->      edg-lcfg-wlconfig-defaults-s1-1.0.21-lcg1
```

to speed up startup. Remember that if upgrading one should ensure that the WMS services are restarted, to allow the new version to become active.

known problems:

- The network server (NS) service threads can exit when there are connection problems. Currently the period cron check on the WMS services will restart the NS in case all the threads have exited.



- Monitoring (Grid ICE):

2004-07-28: patches 201, 211 and 212
support for job monitoring and PBS Torque job

edt_sensor release 1.4.20-pl2

New functionalities:

this release improves the 1.4.19ADC as follows:

- several bug fix in the ChekJobs.pl (job monitoring)
- support for PBS Torque
- fix for edg-fmon-agent so that when an output greater than 64Kbytes is generated (edg-fmon limit), no log info is written; the only effect is that jobs won't be able to be published until the output is lower than this limit; no other effects; this limitation will be shortly solved, see the following bug:
https://savannah.cern.ch/bugs/?func=detailitem&item_id=4255

Configuration:

in the ComputingElement-GridICE-cfg.h you can add the following option to the CheckJobs.pl command: `--queued-jobs=off`
this will disable the publishing of queued jobs; it is recommended if farms managing 500 jobs or more (queueud/running/executed in the last 2 hours); this is the amount of jobs generating around 60Kbytes of output;

since the 1.4.20 release, a new option has been also added (`--lrms-path`) for pbs in order to specify the path where to find the accounting log of pbs; read comments in ComputingElement-GridICE-cfg.h; a new option in the site-cfg.h is required

known problems:

- when the output of job monitoring on a single CE is greater than 64 Kbytes (around 500 jobs), this cannot be sent to the fmon-server; to alleviate the problem, activate the `--queued-jobs=off` option

- Data Management:

The focus of this release is internal build systems changes to support new OS versions and architectures (in particular cel3 and ia64)

Also, the catalogs have been tested and validated on MySQL/Tomcat4 again.

Finally, unneeded code that interfaced with the RLI (and gave us complicated build dependencies) has been commented out/removed.

Fixed Bugs

=====



- 3005 - compile WP2 on IA64: build & spec files
- 3264 - edg-rm should not put port numbers for SURLS in catalog. We remove port numbers now at the canonicalization stage
- 3285 - LRC configuration incorrect for MySQL manual install put default for PFN.LENGTH in values file
- 3484 - remove dependency on edg-se-webservice for edg-rm.
- 3485 - 2.2.7 catalogs don't work on MySQL. Extra comma was being handled by Oracle, but not by MySQL
- 3701 - rpmbuild fails on WP2 packages due to files in install dir not packaged.
- 3708 - upgrade internal devo version of ant to 1.6.1 for WP2 code
- 3718 - replica manager 'rep' with no specific SE throws NPE
- 3861 - edg-rm 1.7.4 does not load jars properly

Module Versions

edg-java-tools:v1_0_37
edg-java-security:v1_5_11
edg-java-data-util:v1_3_22
edg-rls-server:v2_2_9
edg-metadata-catalog:v2_2_9
edg-ros:v2_2_2
edg-reptor:v1_7_4

edg-gsoap-base:v1_1_0
edg-rls-client:v2_3_1
edg-metadata-client:v2_3_1
edg-ros-client:v2_3_1
edg-reptor-client:v2_3_1

REMOVED MODULES

edg-se-webservice
edg-replica-location-index
edg-replica-location-index-client

- GFAL:

- add support for EDG-SE and new dCache release

- lcg-util:

New feature:

- publish filesize in LRC (ALICE)
- use different uuid when generating target filename for lcg-cr and lcg-rep to allow for parallel jobs doing the same replication
- do not replicate if replica is already at target host.

Bug fixes:

- 3549 - add method lcg-la (list aliases)



3586 - remove incomplete file after failed GridFTP transfer (CMS)

- R-GMA:

RGMA version is now 3.4.28.

Bugs fixed:

3625 - ``/opt/edg/etc/profile.d/edg-rgma-env.csh still contains wrong java version``

3645 - ``/tmp is not the best place to put logs``.
They now log into /opt/edg/var/log/

3647 - ``rgma default log level is debug`` fixed even if now Rob Byrom says it would be right to have it switched on (see the thread in Savannah)

3648 - ``confusing configuration file`` fixed

3655 - ``edg-rgma-servlets overwrite configuration file`` still opened;

3720 - ``/opt/edg/libexec/edg-rgma-restart-all in release 3.4.28-1``
the script assumes that the user tomcat4 exists on any machine, not only on the server.

3779 - ``/opt/edg/sbin/test/edg-rgma-run-examples logs``
If an error occurs, R-GMA developers suggest to have look into a directory where there is no log file or error file.

3780 - ``log for /opt/edg/libexec/edg-rgma-restart-all``
This script is used by the rgma object that must run on all the machines. But the log file is created only if gin and service-status run; we do not run them, so the log file is not created and, if a problem occurs (as in bug 3720) one does not know where to look.

3798 - ``edg-info-service-1.0.0-1`` closed; we removed this useless rpm from our tag

Update to edg-lcfg-javasecurity-1.5.1 (patch 210)

Note: all known RGMA bugs were fixed.

New bugs against RGMA can be opened in Savannah but they will be sent directly to RGMA developers, C&T does not do analysis of them, or, in other words, RGMA is not directly supported by LCG, we only include it in the release.

New R-GMA configuration files included into LCFGng to support R-GMA installation:

Monitoring-cfg.h

Monitoring_over_CESE-cfg.h
RGMA-cfg.h
CE-rgma-cfg.h
SE-rgma-cfg.h
UI-rgma-cfg.h
WN-rgma-cfg.h

The following files had to be modified to support R-GMA installation:

site-cfg.h
tomcat-cfg.h
StorageElementClassic-cfg.h
ComputingElement-<....>-cfg.h
UserInterface-cfg.h
WorkerNode-cfg.h

- Other:

patches:

158 - openpbs 2.3.16-lcg1
175 - fedora legacy linux - security
178 - fedora legacy linux - security
182 - pine - security
184 - auditlog - latest version of auditlog

- Others:

patch 205:

- IUCC and BEGrid root certs and updated Russia CA

Cleanup of unused configuration settings for the registration
to TOP GIIS from SITE GIIS



16. CHANGE HISTORY

LCG-2_1_0 added information on queue length and general references for external documentation

-merged the document with the how2start guide and added additional material to it. This is the last text based version.

Release LCG-2_0_0 (XX/02/2004):

Major release: please see release notes for details.

Release LCG1-1_1_3 (04/12/2003):

- Updated kernel to version 2.4.20-24.7 to fix a critical security bug
- Removed ca_CERN-old-0.19-1 and ca_GermanGrid-0.19-1 rpms as the corresponding CAs have recently expired
- On user request, added zsh back to the UI rpm list
- Updated myproxy-server-config-static-lcg rpm to recognize the new CERN CA
- Added oscar-dar rpm from CMS to WN

Release LCG1-1_1_2 (25/11/2003):

- Added LHCb software to WN
- Introduced private-cfg.h.template file to handle sensible settings for the site (only the encrypted root password, for the moment)
- Added instructions on how to use MD5 encryption for root password
- Added instructions on how to configure http server on the LCFG node to be accessible only from nodes on site
- Fixed TCP port range setting for Globus on UI
- Removed CERN libraries installation on the UI (added by mistake in release LCG1-1_1_1)
- Added instructions to increase maximum number of open files on WNs
- Added instructions to correctly set the root password for the MySQL server on the RB
- Added instructions to configure WNs to use a web proxy for CRL download