



LHC COMPUTING GRID

LCG2 SITE SETUP

Document identifier:

EDMS id: **n/a**

Version: **v2.7.0**

Date: **January 31, 2006**

Section: **gis**

Document status: **final**

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

File: **LCG2-Site-Setup**

Abstract: These notes will tell you how to join LCG as a site



CONTENTS

1. INTRODUCTION.....	3
1.1. WHAT IS LCG?	3
1.2. WHAT IS EGEE?.....	3
1.3. WHICH OS VERSIONS ARE SUPPORTED?	3
2. HOW TO JOIN THE EGEE/LCG2 PRODUCTION SERVICE.....	4
2.1. WHAT TO SETUP	5
2.2. HARDWARE.....	6
2.3. HOW TO JOIN AS A USER.....	6
2.4. HOW TO REPORT PROBLEMS	6
2.5. HOW TO SETUP YOUR SITE	7
2.6. NETWORK ACCESS	7
2.7. GENERAL NOTE ON SECURITY	7
2.8. DOCUMENTATION.....	7
2.9. FIREWALL CONFIGURATION	7
2.10. HOW TO PUBLISH YOUR SITES SPECINT VALUES AND QUEUE LENGTH CORRECTLY	8
2.11. LCG SITE CONFIGURATION DATABASE AND GRID OPERATION CENTRE (GOC) ..	8



1. INTRODUCTION

Please see the following page for details of the current release <http://lcg.web.cern.ch/LCG/Sites/releases.html>.

These notes are provided to describe the process of setting up and registering a grid site using the middleware packaged by LCG. This middleware represents the current middleware stack used in the LCG-2 and EGEE production grid. This information is relevant for site managers or sysadmins that want to setup a EGEE/LCG-2 production site or upgrade their site to the latest release.

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.

1.2. WHAT 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>.

1.3. WHICH OS VERSIONS ARE SUPPORTED?

This release is available for Scientific Linux 3 (SL3) and compatible distributions.



2. HOW TO JOIN THE EGEE/LCG2 PRODUCTION SERVICE

If you want to join 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.

The support for sites is organized in a hierarchical way. Please contact the managers at the regional operations centres (ROCs) for your region. In case your site is not covered by the following list you should contact Ian Bird and your site will be either connected to one of the existing ROCs or the CERN deployment team will provide the required support.

The ROC managers mailing list and a look at the EGEE project web page and there especially the SA1 page might help to find a matching ROC: <project-egee-roc-managers@cern.ch>

- France: grid.support@cc.in2p3.fr
- Italy: it-roc@infn.it
- UK/Ireland: support@grid-support.ac.uk
- North: grid.support@sara.nl
- South-west: grid.support@lip.pt
- Germany + Switzerland: support@ggus.org
- South-east: grid-support@egee-see.org
- Central Europe: Egee-cegc-admin-support@cyf-kr.edu.pl
- Russia: roc.support@ihep.ru
- Asia-Pacific: roc@lists.grid.sinica.edu.tw
- CERN: hep-project-grid-cern-testbed-managers@cern.ch

The formal process to become a site in LCG2/EGEE is currently adapted to the new structure given by the EGEE project. Until this is finalized the following steps should be followed. Please include in your mail exchanges with your ROC the deployment team <support-lcg-deployment@cern.ch> After the initial step described above you should follow these steps:

1. send the following information to you ROC:
 - The full name of the participating institute and site:
 - The name of the site to be published in the information system:
 - The name, e-mail address and phone number of the site manager:
 - The name, e-mail address and phone number of the site security contact:
 - The e-mail address of a managed list to contact your incident response team:
 - The name of your regional operations centre



2. The ROC will sent this information to the Grid Operations Centre (GOC) and to the security team.
3. The site security contacts and sysadmins will receive material from the LCG security team that describes the security policies
4. The ROC will help you to setup your site
5. Register your site domain with the service at RAL, send mail to Steve Traylen <s.traylen@rl.ac.uk>
6. The ROC will help you using the tests described in the LCG2-Site-Testing document to get your site working
7. The ROC will forward the node name and LDAP contact string of your site GIIS to the GOC and the deployment team to add your site to the information system.
8. Your site will be monitored and certified every 24 hours. Problems will be reported to your ROC.

In addition you have to go through some additional steps after you have started the setup of your site:

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>). On this list new releases are announced and it is the common place to exchange information.

It is quite useful to have a look at the user guide:<https://edms.cern.ch/file/454439//LCG-2-Userguide.pdf>. 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.

2.1. WHAT TO SETUP

Discuss with your ROC or the grid deployment team a suitable layout for your site. Various configurations are possible. Experience has shown that using at the beginning a standardized 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 ROC or CIC 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 practice 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.

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



The process to add additional VOs is described in the installation guides. 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. 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>.

2.2. HARDWARE

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.

2.3. HOW TO JOIN AS A USER

If you want to use the grid as a user you are currently reading the wrong document. Please go the EGEE NA4 page and get into contact with a VO. To learn more about using LCG you can follow the steps described in the LCG User Overview (<http://lcg.web.cern.ch/LCG/peb/grid\protect\T1\textunderscoredeployment/user\protect\T1\textunderscoreintro.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.

2.4. HOW TO REPORT PROBLEMS

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. Always report problems first to your ROC. 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>.



2.5. HOW TO SETUP YOUR SITE

A set of scripts are provided to ease the installation. Sometimes we refer to this as YAIM (Yet Another Installation Method). For worker nodes and user interface nodes we have prepared in addition releases that are based on tar-balls. Both methods are described in the Manual Installation Guide [1].

2.6. 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 and using dual network interfaces on the service nodes. The ROC in Italy has compiled some information about this. Please contact them for details.

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.

2.7. 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.

2.8. DOCUMENTATION

- [D0] EGEE Project Hoempage:
<http://egee-intranet.web.cern.ch/egee-intranet/gateway.html>
- [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>
- [D6] LCG GOC Mainpage:
<http://goc.grid-support.ac.uk/gridsite/gocmain/>

2.9. 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.

2.10. HOW TO PUBLISH YOUR SITES SPECINT VALUES AND QUEUE LENGTH CORRECTLY

Your queues lengths are currently adjusted to CPU time and wall clock time. To allow the users a proper match between the local resources and their jobs, some care has to be taken to configure in the information system of the CE the parameters that describe the speed of your nodes.

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, SpecFloat and SpecInt, 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 (batch systems) 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.specbench.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,

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

2.11. LCG SITE CONFIGURATION DATABASE AND GRID OPERATION CENTRE (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/
```



REFERENCES

- [1] O. K. A. R. A. U. Guillermo Diez-Andino, Laurence Field. Lcg generic installation guide, 2005.
<http://grid-deployment.web.cern.ch/grid-deployment/documentation/LCG2-Manual-Install>.