

## GRACE Project

### Feedback and requirements on the usage of LCG2

*Last update: November 2004*

The comments and requirements formulated below refer to the experience acquired in the context of the GRACE project (more information at <http://www.grace-ist.org>). GRACE has been using LCG2.0.0 and is currently using LCG2.2.0 for experiencing IR (Information Retrieval) based on Grid technology.

<b>List of possible values per field</b>
<b>Requirement type:</b> (one or more of the following)
Function Interoperability Performance Reliability Security Portability Accessibility Metrics Production Accounting
<b>Related to:</b>
User interface – (description) Job submission – (description) Data management – (description) Information system – (description) Storage – (description) Network – (description) Operation – (description)
<b>Priority:</b>
Very High High Helpful Low

Selected requirements are being submitted to EGEE through a Web interface at:  
<http://egee-na4.ct.infn.it/requirements/>

#	1
Date:	July 2004
<b>Contacts:</b>	
Roberta Faggian Marque (Roberta.Faggian@cern.ch)	
<b>Difficulty encountered title:</b>	Usage of Grid personal certificates
<b>Description :</b>	
At present, each user who wants to access to the EGEE Grid service is required to own and use its personal Grid user certificate. The procedure of getting and using this private certificate limits any Grid-based application (portal) usage by anonymous users.	
<b>Requirement type:</b>	Security, Accessibility
<b>Related to:</b>	User Interface - "Application level" certification
<b>Priority:</b>	Very high
<b>Description:</b>	
In the context of GRACE application, and of other applications of similar nature (Grid portals), anonymous access to restricted services is required. Access and usage of Grid resources should be granted at "application level" instead of at "user level". This would mean for the application to be able to access to the Grid services by using a unique general GRACE user certificate for all GRACE requests. In this way, GRACE users will not be forced to register for a Grid certificate to perform searches on publicly available resources.	

#	2
Date:	June 2004
<b>Contacts:</b>	
Roberta Faggian Marque (Roberta.Faggian@cern.ch)	
<b>Difficulty encountered title:</b>	new VO configuration
<b>Description :</b>	
After installation of LCG2.0.0 using LCFG server, GRACE Grid resources access rights were configured for a list of "standard" pre-configured VOs. No documentation or automatic procedure was made available for the set up a new VO for GRACE and for the configuration of the existing Grid nodes for the new VO.	
<b>Requirement type:</b>	Function
<b>Related to:</b>	Operation - VO creation
<b>Priority:</b>	Helpful
<b>Description :</b>	
Creation of new VOs and configuration of existing Grid nodes for new VOs should be facilitated by an automatic procedure or documented in a simple list of instructions.	

#	3
Date:	September 2004
<b>Contacts:</b>	
Roberta Faggian Marque (Roberta.Faggian@cern.ch)	
<b>Difficulty encountered title:</b>	Real-time performances
<b>Description :</b>	
It is not possible to achieve real-time response with short jobs execution due to the time overhead introduced by the use of the middleware.	
<b>Requirement type:</b>	Function
<b>Related to:</b>	Job submission – short job execution

<b>Priority:</b>	High
<b>Description :</b>	
It should be possible to execute short jobs without introducing significant overhead.	

<b>#</b>	<b>4</b>
<b>Date:</b>	August 2004
<b>Contacts:</b>	
Enrico DiMarco (enrico.dimarco@gl2006europe.com), Srinivasan (srini@gl2006europe.com)	
<b>Difficulty encountered title:</b>	Information system interface
<b>Description :</b>	
<b>Requirement type:</b>	Function
<b>Related to:</b>	Information system – API interface
<b>Priority:</b>	Helpful
<b>Description :</b>	
It should be possible to be notified by the Grid about the job status progress by a programmable interface (application) through an API.	

<b>#</b>	<b>5</b>
<b>Date:</b>	June 2004
<b>Contacts:</b>	
Enrico DiMarco (enrico.dimarco@gl2006europe.com), Srinivasan (srini@gl2006europe.com)	
<b>Difficulty encountered title:</b>	Usage of Replica Catalogue
<b>Description :</b>	
Application deployment through Replica Catalogue doesn't work with LCG2.0.0 on GILDA. Registration works successfully. First usage of the replicated file works, second use doesn't work.	
<b>Requirement type:</b>	
<b>Related to:</b>	
<b>Priority:</b>	
<b>Description :</b>	
More stable Replica Service. To be tested if this has been improved in LCG2.2.0.	

<b>#</b>	<b>6</b>
<b>Date:</b>	September 2004
<b>Contacts:</b>	
Jan Fiete Grosse-Oetringhaus (Jan.Fiete.Grosse-Oetringhaus@cern.ch)	
<b>Difficulty encountered title:</b>	No limit to execution time
<b>Description :</b>	
If a job is sent to the Grid and due to a bug in the executed application it does not terminate, it will be cancelled by the Grid only if the user's proxy expires. If this is not the case it will run forever.	
<b>Requirement type:</b>	Function
<b>Related to:</b>	Job submission – Limiting execution time
<b>Priority:</b>	Helpful
<b>Description :</b>	
It should be possible or maybe even compulsory to specify a maximum execution	

time in the JDL file, after this time the Grid is allowed to cancel the job.  
 The Resource Broker could also use this information to improve the selection of the computing element for a job, because it could select the one which will be idle next.

#	7
<b>Date:</b>	September 2004
<b>Contacts:</b>	
Roberta Faggian Marque	
<b>Difficulty encountered title:</b>	Automatic Job cancellation
<b>Description :</b>	
Sometimes scheduled jobs are executed after quite long time from their submission (after waiting in a queue). In some cases their “late” execution can be cancelled.	
<b>Requirement type:</b>	Function
<b>Related to:</b>	Job submission - expiry
<b>Priority:</b>	Low
<b>Description :</b>	
It would be useful to be able specify a sort of “expiry date/time” for the job execution.	

#	8
<b>Date:</b>	September 2004
<b>Contacts:</b>	
Jan Fiete Grosse-Oetringhaus (Jan.Fiete.Grosse-Oetringhaus@cern.ch)	
<b>Difficulty encountered title:</b>	No dependency model for jobs
<b>Description :</b>	
At the moment the execution of a Grid job depends on its input data and the specified requirements. Dependencies between different jobs cannot be specified.	
<b>Requirement type:</b>	Function
<b>Related to:</b>	Job submission – Dependency model
<b>Priority:</b>	Very high
<b>Description :</b>	
<p>The GRACE application uses natural language processing which is very computing-intensive. Therefore the input data is split into parts, processed by parallel jobs and merged afterwards. Furthermore the application consists of different components which are executed sequentially.</p> <p>It should be possible to schedule jobs which are not executed before other jobs have finished. The first group of jobs provides the input for the second group of jobs in this case.</p> <p>A job dependency model should be developed to provide the possibility to send a group of jobs and dependency information to the system.</p>	

#	9
<b>Date:</b>	September 2004
<b>Contacts:</b>	
Jan Fiete Grosse-Oetringhaus (Jan.Fiete.Grosse-Oetringhaus@cern.ch)	
<b>Difficulty encountered title:</b>	No output from cancelled jobs
<b>Description :</b>	
No standard and error output is retrieved from cancelled jobs.	
<b>Requirement type:</b>	Function
<b>Related to:</b>	Job submission – Canceling of jobs

<b>Priority:</b>	High
<b>Description :</b>	
It should be possible to retrieve the standard output and error output from cancelled jobs. In cases where the user decides to cancel a job because it needs too much time or something else seems fishy, this could be very helpful to debug the application.	

<b>#</b>	<b>10</b>
<b>Date:</b>	April 2004
<b>Contacts:</b>	
Giuseppe Sisto (Giuseppe.Sisto@TILAB.COM), Srinivasan (srini@gl2006europe.com)	
<b>Difficulty encountered title:</b>	LCG2 nodes installation
<b>Description :</b>	
Software difficult to install and configure. Although simpler than the manual installation, the software automated installation (LCGFGng) is very tricky and requires high level skills. Once installed, it was possible to configure it and to make it to run only with the external support of grid experts (in our case those who run GILDA infrastructure, which were really very helpful). Moreover, with the automated installation it is possible to install only grid elements (UI, CE, SE, WN), and not centralized services (such as RB, BDII, RLS...), it would be helpful to have the possibility to create easily a full grid testbed.	
<b>Requirement type:</b>	Function
<b>Related to:</b>	Operation - Installation
<b>Priority:</b>	Helpful
<b>Description :</b>	
Improve the installation and the configuration procedures by making them more user-friendly and documented. In particular document better the possible problems and errors that may occur.	

<b>#</b>	<b>11</b>
<b>Date:</b>	April 2004
<b>Contacts:</b>	
Giuseppe Sisto (Giuseppe.Sisto@TILAB.COM), Srinivasan (srini@gl2006europe.com)	
<b>Difficulty encountered title:</b>	LCG2 nodes installation
<b>Description :</b>	
During the first attempts to install the test bed, we had a number of spare machines, each one with different HW characteristics. It soon turned out that most of them were not "standard hardware", in terms of NICs (some of them were not taken into account by the compiled linux kernel) disks (it was not possible to install on a SCSI disk, but only on IDE), etc. ... To simplify things, it was necessary to get rid of such machines, and to adopt a set of PCs all alike, for which the same configuration applied.	
<b>Requirement type:</b>	Function
<b>Related to:</b>	Operation - Installation
<b>Priority:</b>	Helpful
<b>Description :</b>	
It would be very helpful to have a definition of what is intended for "standard hardware", providing minimal and suggested HW requirements for the machines where the services will be installed, also according to the kind of service that will go	

on each machine (i.e. requirements for SE, CE, WN, etc....)

#	12
Date:	April 2004
<b>Contacts:</b>	
Giuseppe Sisto (Giuseppe.Sisto@TILAB.COM)	
<b>Difficulty encountered title:</b>	LCG2 nodes installation
<b>Description :</b>	
With the automatic installation through LCFG, each grid element requires a dedicated machine. This may imply to invest in a significant number of machines, which can be hardly used for any other purpose.	
<b>Requirement type:</b>	Function
<b>Related to:</b>	Operation - Installation
<b>Priority:</b>	Helpful
<b>Description :</b>	
It may be helpful to allow different services to run on the same machine. This is possible only with the manual installation (although this kind of installation is scarcely documented and unaffordable for the great majority of “new” grid site managers). It would help if this were possible with the automatic installation too, so that the HW could be used more efficiently.	

#	13
Date:	September 2004
<b>Contacts:</b>	
Giuseppe Sisto (Giuseppe.Sisto@TILAB.COM)	
<b>Difficulty encountered title:</b>	LCG2 upgrade
<b>Description :</b>	
The upgrade procedure from LCG2.0.0 to LCG2.2.0, like any other version upgrade, although well documented, is very tricky and implies some risk to introduce instability. A lot of errors appear during the upgrade: the instructions say that some of them are “normal” and can be ignored; however they are really annoying, as someone cannot know how serious they may be. For example, in our case, after the upgrade, one WN is no longer monitored by GridICE, another WN was not able to reboot, as some partition configuration was lost, so the node had to be fully re-installed with PXE.	
<b>Requirement type:</b>	Function
<b>Related to:</b>	Operation - Deployment
<b>Priority:</b>	Helpful
<b>Description :</b>	
Upgrade to new middleware releases should be a simple process, completely automated and as far as possible transparent to the user.	

#	14
Date:	April 2004
<b>Contacts:</b>	
Giuseppe Sisto (Giuseppe.Sisto@TILAB.COM)	
<b>Difficulty encountered title:</b>	“full IP” machines
<b>Description :</b>	
Integrating in a grid infrastructure requires “full-ip” machines and this may introduce	

vulnerability problems in an environment with high security requirements, such as an enterprise. In our case, we had to install the grid machines on an isolated sub-network, and open the firewall to allow interconnection with the external grid infrastructure. This prevents the use of the grid from the company's intranet.	
<b>Requirement type:</b>	Function
<b>Related to:</b>	Security
<b>Priority:</b>	Helpful
<b>Description :</b>	
As part of the installation instruction, it would be desirable to have a precise and up-to-date mapping of the ports to be opened in the firewall, and avoid the use of random port numbers.	

<b>#</b>	<b>15</b>
<b>Date:</b>	August 2004
<b>Contacts:</b>	
Enrico DiMarco (enrico.dimarco@gl2006europe.com), Srinivasan (srini@gl2006europe.com)	
<b>Difficulty encountered title:</b>	Usage of LCG2 Java API calls
<b>Description :</b>	
In LCG2.0.0 high level Java API's wrapping native libraries are not as stable as expected. Libraries' loading errors/exceptions prevent from further operations on job submission/output retrieval.	
<b>Requirement type:</b>	Function
<b>Related to:</b>	Job submission - Java APIs
<b>Priority:</b>	High
<b>Description :</b>	
More stable job submission system though APIs, otherwise, instead of having native libraries for job submission it would be better to have Web services (WSDL). To be tested if this has been improved in LCG2.2.0.	