

An Overview of a Grid Architecture for Scientific Computing

A. Wäänänen^a, M. Ellert^b, A. Konstantinov^c, B. Kónya^d, O. Smirnova^d

^aNiels Bohr Institutet for Astronomi, Fysik og Geofysik,
Blegdamsvej 17, DK-2100, Copenhagen Ø, Denmark

^bDepartment of Radiation Sciences, Uppsala University,
Box 535, 751 21 Uppsala, Sweden

^cUniversity of Oslo, Department of Physics,
P. O. Box 1048, Blindern, 0316 Oslo, Norway

^dParticle Physics, Institute of Physics, Lund University,
Box 118, 22100 Lund, Sweden

Abstract

This document gives an overview of a Grid testbed architecture proposal for the NorduGrid project. The aim of the project is to establish an inter-Nordic testbed facility for implementation of wide area computing and data handling. The architecture is supposed to define a Grid system suitable for solving data intensive problems at the Large Hadron Collider at CERN. We present the various architecture components needed for such a system. After that we go on to give a description of the dynamics by showing the task flow.

1 Introduction

This document assumes basic knowledge of the computing *Grid* concept, which is a paradigm for the modern distributed computing and data handling. For a general introduction to Grid computing the reader is referred to eg. [1]. The most common starting point for constructing a computing Grid is the Globus Toolkit¹ [2]. This toolkit provides a Grid API and developing libraries as well as basic Grid service implementations.

The NorduGrid [3] project is a common effort by the Nordic² countries to create a Grid infrastructure, making use of the available middleware. Through the European DataGrid project (EDG) [4] the NorduGrid project has had extensive experience with the Globus Toolkit and with deploying and using a Grid Testbed. During this we have found some shortcomings in the Globus Toolkit and some problems with the EDG Testbed architecture that we would like to address on a Grid testbed in the Nordic countries. In this paper we present a proposal for a Grid architecture for a production testbed at the LHC [5] experiments. It is not the intent to define a general Grid system, but rather a system specific for batch processing suitable for problems encountered in High Energy Physics. Interactive and parallel applications have not been considered. It is our goal to deploy a Grid testbed based on this architecture for the LHC Data challenges in the Nordic countries.

We will start by describing the various components of the system. Focus will be given to the components which we have developed ourselves or taken from others, but heavily modified. After the description we will give a review of the task flow and the communication between the various components.

It is assumed that the reader is familiar with the major components of the Globus Toolkit since the NorduGrid Testbed architecture uses this as the foundation. An example is the

¹Globus Project and Globus Toolkit are trademarks held by the University of Chicago.

²The term *Nordic* covers the countries: Denmark, Norway, Sweden and Finland.

underlying security model - the Grid Security Infrastructure (GSI) - which is based on the X.509 certificate system and takes care of authentication and authorisation and the GRIS and GIIS of the information system (for an explanation of the GRIS, GIIS, VO terms see [6] and [7]).

One of the aims of the architecture was to make the system able to install on top of an existing Globus installation. Rather than modifying the Globus Toolkit, we have made a clear boundary between the two architectures, thus enabling them to coexist and try out this new system without destroying an already working Globus installation.

2 The NorduGrid Architecture Components

2.1 Computing Element - CE

The *computing element* (CE) is the backend of the system. It is where the execution of the application is performed. This is a very general element which can be anything from a single processor to complex computing clusters. In our case it is at present limited to simple PBS (Portable Batch System) clusters.

One of our basic ideas about Grid implementations is that it should not impose any restriction on the local site. Therefore it should be possible to use existing computing clusters and place them on the Grid with minor or no special reconfiguration. One thing that we have stressed, which should not be imposed on the site, is that the individual worker nodes (WN's) of the cluster can not be required to have direct access to the Internet. This excludes dependence on eg. global filesystems like AFS and direct download from the Internet from any WN. Grid services will thus only be run from a front-end machine. This scenario does not however exclude the use of local network filesystems and NFS is actually often used for cluster-wide application distribution.

2.2 Storage Element - SE

Similar to the computing element, the *storage element* (SE) is the common term for another of the basic Grid resources - storage. The storage can be as simple as a standard filesystem or eg. a database. The authorisation is handled using Grid credentials (certificates). A SE can be *local* or *remote* to a CE. Here *local* means that access is done via standard filesystem calls (eg. open, read and close) and is usually realised by a NFS server. A remote SE is usually a stand-alone machine running eg. GridFTP [8] server with local file storage. Data replication is done by services running on the SE.

A dedicated pluggable GridFTP server is being developed for use on the SE. At the moment a simple file access plugin exists. The main reason for this is to have a way to provide a consistent certificate-based data access to the data. At least one other Grid solution to a certificate-based filesystem exists [9]. One advantage of the GridFTP approach is that it is done entirely in user space and thus is very portable.

2.3 Replica Catalog - RC

The information about replicated data is contained in the *Replica Catalog* (RC). This is an entirely add-on component to the system and as such is not a requirement.

2.4 Information System - IS

A stable, robust, scalable and reliable information system is the cornerstone of any kind of Grid system. Without a properly working information system it is not possible to construct a functional Grid. The Globus Project has laid down the foundation of a Grid information system with their LDAP-based Metacomputing Directory Service (MDS) [6]. The NorduGrid information system is built upon the Globus MDS.

The information system described below forms an integral part of the NorduGrid Testbed Architecture. In our Testbed, the NorduGrid MDS plays a central role: all the information related tasks, like resource-discovery, Grid-monitoring, authorised user information, job status monitoring, are exclusively implemented on top of the MDS. This has the advantage that all the Grid information is provided through a uniform interface in an inherently scalable and distributed way due to the Globus MDS. Moreover, it is sufficient to run a single MDS service per resource in order to build the entire system. In the NorduGrid Testbed a resource does not need to run dozens of different (often centralized) services speaking different protocols: the NorduGrid Information System is purely Globus MDS built using only the LDAP protocol.

The design of a Grid information system is always deals with questions like how to represent the Grid resources (or services), what kind of information should be there, what is the best structure of presenting this information to the Grid users and to the Grid agents (i.e. Brokers). These questions have their technical answers in the so-called LDAP schema files. The Globus Project provides an information model together with the Globus MDS. We found their model unsuitable for representing computing clusters, since the Globus schema is rather single machine oriented. The EDG suggested a different CE model which we have evaluated [10]. The EDG's CE-based schema fits better for computing clusters. However, its practical usability was found to be questionable due to improper implementation.

Because of the lack of a suitable schema, NorduGrid decided to design its own information model (schema) in order to properly represent and serve its Testbed. A working information system, as part of the NorduGrid Architecture, has been built around the schema. Nevertheless, NorduGrid hopes that in the not-so-far future a usable common Grid information model will emerge. We think that the experience of the NorduGrid users gained with our working system will provide a useful feedback for this process.

2.4.1 The Information Model.

The NorduGrid Testbed consists of different resources (they can be referred as services), located at different sites. The list of implemented Grid services contains computing resources (Linux clusters operated by PBS), SEs (at the moment basically disk space with a GridFTP server) and RCs. The designed information model naturally maps these resources onto a LDAP-based MDS tree, where each resource is represented by an MDS entry. In this tree, each NorduGrid resource operates as a separate GRIS. The various resources (GRIS's) can be grouped together to form Virtual Organisations (VO) which are served by a GIIS (i.e. in our present Testbed configuration, the resources within a country are grouped together to form a VO). The structure created this way is called a *hierarchical MDS tree*.

The NorduGrid schema is a true mirror of our architecture: it contains information about computing clusters (`nordugrid-cluster` objectclass), storage elements (`nordugrid-se` objectclass) and replica catalogs (`nordugrid-rc` objectclass). In Fig. 1, a part of the NorduGrid MDS tree [11] is shown. The clusters provide access to different PBS queues which are described by the `nordugrid-pbsqueue` objectclass (Fig. 2 shows an example queue entry). Under the queue entries, the `nordugrid-authuser` and the `nordugrid-pbsjob` entries can be found grouped in two distinct sub-trees (the branching is accomplished with the `nordugrid-info-group` objectclass). The `nordugrid-authuser` entry contains all the user-dependent information of a specific authorised Grid user. Within the user entries, the Grid users can find, among other things, how many CPUs are available for them in that specific queue, what is the disk space a job can consume, what is an effective queue length (taking into account the local UNIX mappings), etc. The `nordugrid-pbsjob` entries (see Fig. 3 an example) describe the Grid jobs submitted to a queue. The detailed job information includes the job's unique Grid ID, the certificate subject of the job's owner and the job status.

The NorduGrid information system has been designed to be able to effectively serve the User Interface (UI) (job status query commands, free resource discovery utilities), the brokering agent (in the present implementation it is integrated with the UI's job submission

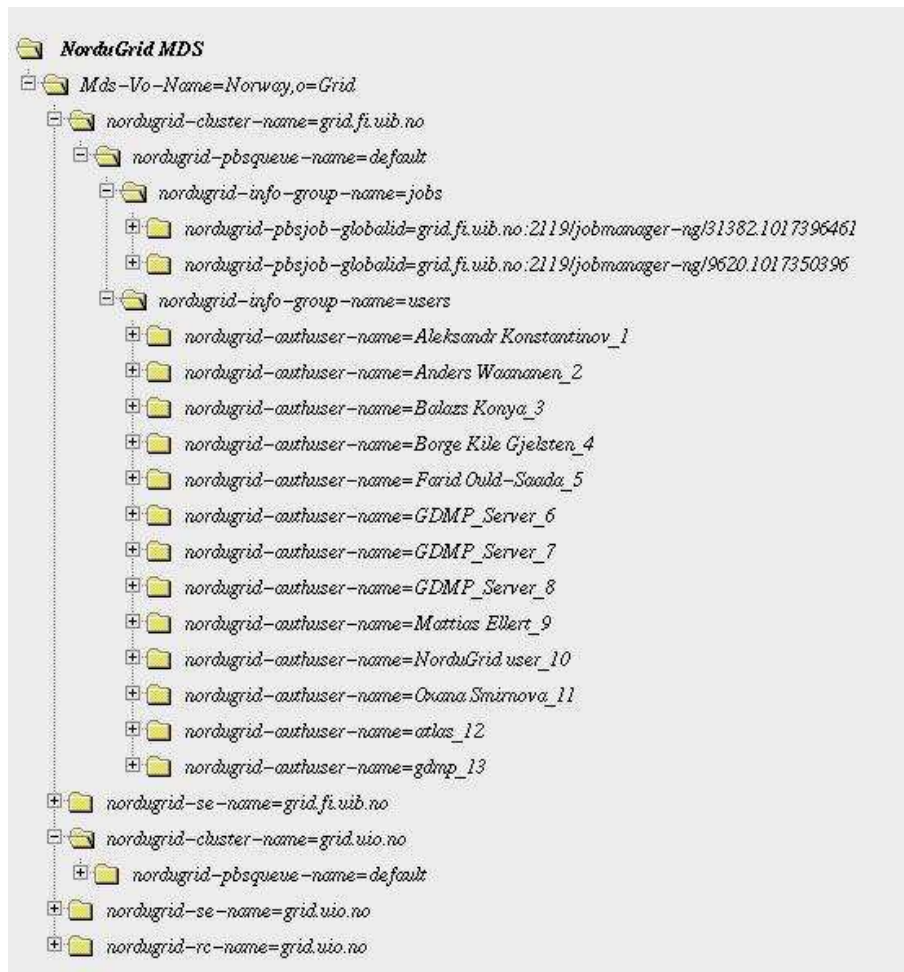


Figure 1: The Norway branch of the NorduGrid MDS tree

command) and a general Grid user who can access the Grid information either directly with a simple `ldapsearch` or through an LDAP-enabled Grid portal.

2.5 Grid Manager - GM

In our model, job management is handled by a single entity which we call the *Grid Manager* (GM). It is the job of the GM to process user requests and prepare them for execution on the CE. It also takes care of post-processing the jobs before they leave the CE. In the Globus Toolkit context, the GM takes care of what is normally done by the Globus job-manager backend scripts. In fact it is installed in a similar way to a standard Globus jobmanager and can work perfectly together with already existing jobmanagers. Authorisation and authentication is still done by the Globus gatekeeper.

The status of each job is recorded in a special status directory which also contains informational files needed by the GM.

2.5.1 Grid Manager tasks.

The Grid Manager uses a job ID assigned by the Globus gatekeeper/jobmanager to distinguish between job requests. It starts by creating a session directory reachable by the WN and parsing the job request. The job request is passed in a resource specification language.

<i>Distinguished Name</i> =	<input type="text" value="jnorugrid-pbsqueue-name=default,nordugrid-cluster-name=grid.uio.no,mds-vo-name=Norway,o=grid"/>
<i>objectClass</i> =	<input type="text" value="jnds"/>
<i>objectClass</i> =	<input type="text" value="jnorugrid-pbsqueue"/>
<i>nordugrid-pbsqueue-name</i> =	<input type="text" value="default"/>
<i>nordugrid-pbsqueue-status</i> =	<input type="text" value="active"/>
<i>nordugrid-pbsqueue-running</i> =	<input type="text" value="1"/>
<i>nordugrid-pbsqueue-queued</i> =	<input type="text" value="0"/>
<i>nordugrid-pbsqueue-maxrunning</i> =	<input type="text" value="3"/>
<i>nordugrid-pbsqueue-maxqueueable</i> =	<input type="text" value="1000"/>
<i>nordugrid-pbsqueue-schedulingpolicy</i> =	<input type="text" value="strict FIFO"/>
<i>nordugrid-pbsqueue-assignedqtype</i> =	<input type="text" value="mixed"/>
<i>nordugrid-pbsqueue-assignedqnumber</i> =	<input type="text" value="3"/>
<i>Mds-validfrom</i> =	<input type="text" value="20020829155451Z"/>
<i>Mds-validto</i> =	<input type="text" value="20020829155481Z"/>

Figure 2: Example nordugrid-pbsqueue entry

This specifies, among other things, the list of input files needed for the job execution. The Grid Manager then proceeds to download files to the session directory from remote SEs or copy files from local SEs as specified by the job request. Files can also be uploaded by the user submitting the job. Once all files are in place the GM submits the job to the local scheduling system (PBS). All information about job status must be retrieved through the MDS. Communication to the job such as cancellation is handled through the GM as a specific RSL requests. When the job leaves the WN(s) it is the responsibility of the GM to clean up afterwards. The GM also manages the output files and uploads them to local or remote storage elements and registers uploaded files to the Replica Catalog.

2.6 User Interface (UI)

In contrast to eg. the Grid implementation in the EDG, the NorduGrid *User Interface (UI)* has significantly more responsibility. This is mainly due to the choice of having the resource broker placed at the UI rather than having it as a central service. We found that the EDG implementation with a central broker, where all job requests as well as data payload had to pass through, would be a single point of failure and non-scalable.

The NorduGrid UI is at present command-line driven, while a web based solution is foreseen in the future. The UI is responsible for generating the user request in a *Resource Specification Language (RSL)* based on the user input. The RSL we use has additional attributes to those provided by the Globus Toolkit [12]. This xRSL has been enhanced to support enriched input/output capabilities and more specification of PBS requirements. All unneeded Globus attributes has been deprecated.

The UI matches the request or job options with the available resources as reported by the MDS, and returns a list of matching resources. Job options can reflect eg.:

- required CPU time
- required system type
- required disk space

<i>Distinguished Name</i> =	/nordugrid-pbsjob-globalid=grid.uio.no:2119/jobmanager-ng/31633.1017417271,
<i>objectClass</i> =	Mds
<i>objectClass</i> =	nordugrid-pbsjob
<i>nordugrid-pbsjob-globalid</i> =	grid.uio.no:2119/jobmanager-ng/31633.1017417271
<i>nordugrid-pbsjob-globalowner</i> =	/o=Grid/ou=Nordugrid/ou=quark.lu.se/cn=Balazs Konya
<i>nordugrid-pbsjob-jobname</i> =	Sleepy job
<i>nordugrid-pbsjob-submissiontime</i> =	20020329165433Z
<i>nordugrid-pbsjob-execqueue</i> =	default
<i>nordugrid-pbsjob-execcluster</i> =	grid.uio.no
<i>nordugrid-pbsjob-status</i> =	INLEMS: R
<i>nordugrid-pbsjob-usedcputime</i> =	00:00:00
<i>nordugrid-pbsjob-usedwalltime</i> =	00:06:42
<i>nordugrid-pbsjob-usedmem</i> =	2388
<i>nordugrid-pbsjob-comment</i> =	Job started on Fri Mar 29 at 16:54
<i>nordugrid-pbsjob-stdout</i> =	out.txt
<i>nordugrid-pbsjob-stderr</i> =	err.txt
<i>nordugrid-pbsjob-submissionid</i> =	Not yet Implemented
<i>Mds-validfrom</i> =	20020329160226Z
<i>Mds-validto</i> =	20020329160256Z

Figure 3: Example nordugrid-pbsjob entry

- required runtime environment (eg. application software)
- required memory
- required data from SE's

3 Task flow

In this section we describe the life of a job and expose the functions of the various components of the NorduGrid architecture. An overview of the system with task flow can be seen in Fig. 4.

The numbers on the figure refers to tasks which we describe below.

1. The User Interface does a filtered query against the GHS and query the Replica Catalog to get the location of input data. Based on these responses and its brokering algorithm the Broker within the UI selects a remote resource.
2. The UI contacts the selected resource by submitting the xRSL file to the Gatekeeper of the resource, along with the (resolved by the Broker) physical file names.

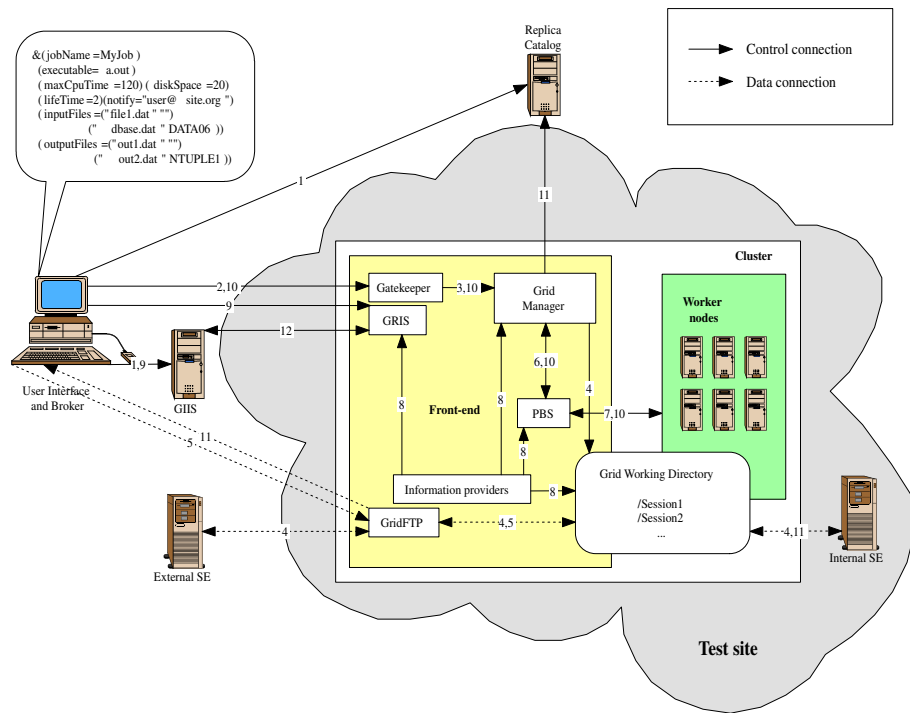


Figure 4: NorduGrid task flow

3. The Gatekeeper does the required authentication and authorisation and passes the job request to the Grid Manager.
4. Grid Manager creates the session directory and downloads or copies files from Storage Elements to this directory
5. User Interface uploads input files and executables via GridFTP to the session directory of the job
6. After all files are pre-staged, the Grid Manager submits the job to the local resource management system (PBS)
7. PBS schedules and executes the job on the Worker Nodes in the normal way
8. On request, Information Providers collect job , user, queue and cluster information, disk usage information and writes the information to the MDS
9. Job status information is retrieved from the MDS, the User Interface monitors the job status by querying the Information System. E-mail notification to the user during the various job stages by the Grid Manager is also possible.
10. User Interface may cancel and clean jobs by sending RSL request with special attribute set through the Gatekeeper to the Grid Manager. The Grid Manager will then take care of stopping the job and removing the job session directory if requested
11. When the job finishes the Grid Manager moves requested output results to Storage Elements and does registration in the Replica Catalog as specified by the user. The Grid Manager takes care of cleaning up the session directory after its lifetime has exceeded, too While the directory exists, the User Interface can download the specified output files produced by the job.

12. GIIS queries GRIS on demand based on cache timeout values in order to provide fresh enough information. The single GIIS in the figure really represents a whole Virtual Organisation hierarchy. In NorduGrid testbed the hierarchy has a central NorduGrid GIIS which connects to country level GIIS's. The country level GIIS's either connects directly to resources or have an institutional layer in between.

In the figure data and control connections are specified. In designing the system we have been very conscious about the location of data bottlenecks and places where the system will have scalability problems. The only data intensive transfers occur between the Computing Element and the Storage Elements and the User Interface. In contrast to the EDG there is no single point where all data has to pass through. All transfers are truly peer-to-peer.

4 Conclusion and Outlook

A very early prototype implementation of the architecture exists and is being tested and further developed. Scalability tests still has to be performed.

In order to have a functional testbed within a short time-frame, there are some areas which we have not given much attention and have taken simple but still secure and functional solutions. In the future we plan to include, for example a more advanced authorisation and accounting system.

Once we get experience with the system we are going to make the Broker more sophisticated and make it able to perform user dependent choices such as RC location and preferred resource selection.

5 Acknowledgment

We are very grateful for the developers of the Globus Project for their willingness to answer questions and providing an open development environment with access to early version of their toolkit.

As participants in the European DataGrid we have been provided access to a Testbed and the collaboration has been very fruitful.

The NorduGrid project - *A Nordic testbed for wide area computing and data handling* (<http://www.nordugrid.org/>) - is funded by the Nordunet2 program (<http://www.nordunet2.org/>).

References

- [1] I. Foster and C. Kesselman: *The Grid: Blueprint for a New Computing Infrastructure* (1998)
- [2] I. Foster, et al.: The Globus Toolkit: <http://www.globus.org/>
- [3] The NorduGrid Project: <http://www.nordugrid.org/>
- [4] The European DataGrid: <http://www.eu-datagrid.org/>
- [5] The Large Hadron Collider: <http://lhc.web.cern.ch/lhc/>
- [6] K. Czajkowski, S. Fitzgerald, I. Foster, C. Kesselman: *Grid Information Services for Distributed Resource Sharing*. Proceedings of the Tenth IEEE International Symposium on High-Performance Distributed Computing (HPDC-10), IEEE Press, August 2001.
- [7] I. Foster, C. Kesselman, S. Tuecke: *The Anatomy of the Grid: Enabling Scalable Virtual Organisations*. International J. Supercomputer Applications, 15(3), 2001.

- [8] W. Allcock, J. Bester, J. Breshnahan, A. Chervenak, L. Liming, S. Tuecke: GridFTP: Protocol Extensions to FTP for the Grid
<http://www.globus.org/datagrid/deliverables/GridFTP-Overview-200201.pdf>
- [9] A. McNabb: *SlashGrid - a framework for Grid aware filesystems*.
<http://www.gridpp.ac.uk/slashgrid/>
- [10] B. Kónya: *Comments on the Computing Element Information Provider of the EU DataGrid Testbed 1*.
http://www.nordugrid.org/documents/comments_on_ceinfo.pdf
- [11] T. Miao: *LDAPExplorer*. <http://igloo.its.unimelb.edu.au/LDAPExplorer/>
- [12] K. Czajkowski, I. Foster, N. Karonis, C. Kesselman, S. Martin, W. Smith, S. Tuecke: *A Resource Management Architecture for Metacomputing Systems*. Proc. IPPS/SPDP '98 Workshop on Job Scheduling Strategies for Parallel Processing, pp. 62-82, 1998.