

Usage statistics and usage patterns on the NorduGrid

Analyzing the logging information collected on one of the largest production Grids of the world.



Dr. Kónya, Balázs (1) Prof. Eerola, Paula (1) Prof. Ekelöf, Tord (2) Dr. Ellert, Mattias (3) Prof. Hansen, John Renner (4) Dr. Konstantinov, Aleksandr (5) Dr. Nielsen, Jakob Langgaard (4) Prof. Ould-Saada, Farid (6) Dr. Smirnova, Oxana (1) Dr. Wäänänen, Anders (4) Dr. Erkarslan, Ugur (6) Pajchel, Katarina (6) (1) Lund University, Sweden (2) Uppsala University, Sweden (3) CERN (4) Niels Bohr Institute, Copenhagen, Denmark (5) Vilnius University, Lithuania (6) University of Oslo, Norway

ABSTRACT

The Nordic Grid facility (Nordugrid) came into operation during summer 2002 when the Scandinavian ATLAS HEP group started to use the Grid for the ATLAS Data Challenges (DC) and was thus the first Grid ever contributing to an ATLAS production. Since then, the Grid facility has been in continuous 24/7 operation.

Nordugrid is being used by a growing set of active users from various scientific areas including physics, chemistry, biology and informatics. It has given major contributions to the ATLAS Data Challenge 1 [1] and the ongoing Data Challenge 2.

- The increasing number of resources has made Nordugrid one of the largest production Grids of the world continuously running on more than 30 sites more than 3000 CPUs.
- The resources range from small test clusters at academic institutions to large farms at several supercomputer centers and the Nordugrid software runs on clusters with very different Linux distributions.

This presentation gives a short overview of the design and implementation of the Nordugrid middleware, logging and monitoring facilities. It will be followed by a description of a typical job on Nordugrid and the information about its parameters which are monitored on-line and persisted in the logging service.

MIDDLEWARE

The Nordugrid Middleware (or Advanced Resource Connector, ARC) is an open source software solution distributed under the GPL license, enabling production quality computational and data Grids.

It is a light-weight Grid solution, designed to support a dynamic, heterogenous Grid facility, spanning different computing resources and user communities.

Emphasis is put on scalability, stability, reliability, robustness, and performance of the middleware. These goals of the architecture of the ARC imply concrete decentralization and functional independence of the components.

Nordugrid provides innovative solutions essential for a production quality middleware:

- the Grid Manager, gridftp (the ARC/Nordugrid GridFTP server),
- the information model and providers (Nordugrid schema),
- User Interface and broker (a "personal" broker integrated into the user interface),
- extended Resource Specification Language (xRSL),
- logging service,
- and the monitoring system.

Most of these services are provided through the security layer of the GSI. The middleware builds upon standard open source solutions like the OpenLDAP, OpenSSL, SASL and Globus Toolkit 2 (GT2) libraries.

ARC main components are:

- Grid services running on the resources: the Grid Manager, the gridftp and the information service.
- Indexing services for the resources and data.
- Redirecting intelligent use of the distributed information and data available on the Grid.

Distribution:

The middleware is available as source and binary in tarballs and RPMs.

Logging

A Logger Service is one of web services implemented by Nordugrid and based on gSOAP and Globus IO API. It provides frontend to underlying MySQL database to store and retrieve information about computing resources' usage (jobs). This database is accessed through a graphical Web interface implemented using PHP4, JavaScript and JGraph based on GD library.

- The main goals for the logger:
- development and usage of Nordugrid over time
- statistics for different clusters, time intervals, applications and users

```
columns:
start      2004-06-04 16:32:15
end        2004-06-05 00:07:16
cluster    /O=Grid/O=Nordugrid/O=host/locf.nbi.dk
user       4623674234684697
name       Name of the job
failure    Job exit code is 265 != 0
lms        facty, slow, gridlong
zone       content of the_xrsl file
ral        129.240.66.30:57818:easy:uso.no
ul
```

The shown columns are temporary. The database is being extended.

A schematic example of a typical Nordugrid job:

```
{
(executable="myprog")
(inputFiles=
("myprog" "http://www.myserver.org/myfiles/myprog")
("myinputfile" "gsiftp://www.mystorage.org/data/file007.dat")
)
(outputFiles=
("myoutput" "gsiftp://www.mystorage.org/results/file007.res")
)
(disk=1000)
(notify="e myname@mydomain.org")
}
```

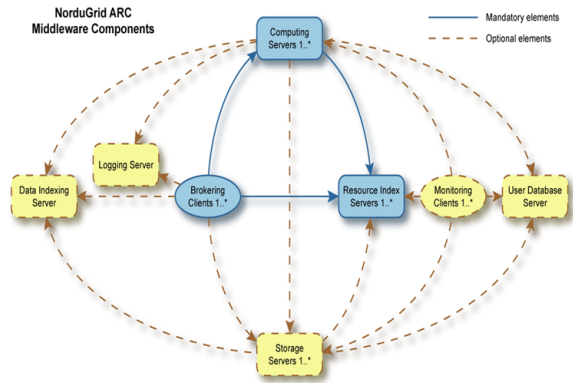
Submitting:
> ngsb -f myjob.xrsl

The user can now follow the job on the Monitor

Check the status of the job.
> ngstat [options] [jobs]

Retrieve the results of a finished job.
> ngget [options] [jobs]

The job usage record gets logged in the logger database by the grid manager.



The Grid Monitor

The Grid Monitor allows the user and sysadmins alike to monitor and check in an intuitive way the status of the ARC Grid facility and the information stored in the Information System. The structure of the Monitor is built upon the following objectclasses:

- nordugrid-cluster: a cluster
- nordugrid-queue: a queue at the cluster, accessible by the ARC users
- nordugrid-job: a Grid job, associated with a queue
- nordugrid-authuser: a user, authorized to submit jobs on a given queue

It also displays information about storage facilities, authorised users, and groups of users (Virtual Organisations), using the corresponding objectclasses.

Country	Site	CPUs	Load (processes)	Grid-local	Queueing
Australia	Atlas (UniMelb)	28	0+0	0+0	0+0
	Alfred (UniMelb)	86	0+0	2+1	0+0
	Dixlab (ORNL)	10	0+0	0+0	0+0
	Aalborg Grid Gateway	46	0+0	0+0	0+0
Denmark	Horseshoe (DCS/GSL)	1172	6+094	6+75	0+0
	HEPAX1	1	0+0	0+0	0+0
	Morpheus	18	10+0	2+0	0+0
	VCR (VideoRecorder)	1	0+0	0+0	0+0
Estonia	UT IMCR Anakonda clus>	15	0+0	0+0	0+0
	UT CS Antarctica Ches>	20	0+0 (Queue down)	0+0	0+0
	CMS on CERN Linux	1	0+0	0+0	0+0
	CMS Production server	5	0+0	0+0	0+0
Finland	UT DOUG Cluster	2	0+0	0+0	0+0
	CMS test cluster	1	0+0	0+0	0+0
	HEMET Cluster	6	0+0	0+0	0+0
	UT Physics Cluster	17	0+0	0+0	0+0
Germany	LSC Kirppu	1	1+0	0+0	0+0
	Mill (Physicum)	60	0+0 (Queue down)	6+4	0+0
	Alpha (HIP)	1	0+0	0+0	0+0
	Testbed0 (HIP)	3	0+0	4+1	0+0
Norway	FZK cluster	740	30+253	0+0	0+0
	LRZ cluster	234	12+10	6+07	0+0
	Oslo Temp Cluster	13	0+0	2+0	0+0
	Parallab IBM Cluster	58	0+0	0+0	0+0
Slovakia	Oslo Grid Cluster	44	1+0	21+0	0+0
	LJO Grid	88	0+0	0+1	0+0
	UP JS GRID	1	0+0	0+0	0+0
	SIGNET	51	4+2	9+0	0+0
Sweden	Bluesmoke (SweGrid_NS>	98	30+0	37+0	0+0
	Kosufy farm	60	0+0	4+0	0+0
	ISV	4	0+0	0+0	0+0
	Hagrid (SweGrid_Uppm>	100	100+0	158+0	0+0
Switzerland	Ingrid (SweGrid_HPC2N)	101	101+0	51+0	0+0
	Quark Cluster	7	0+0	0+0	0+0
	Beppe (SweGrid_PDC KT>	95	0+0	31+0	0+0
	Toto7/Whemim64 (Lunar>	192	0+0	0+2	0+0
TOTAL		3384	501 + 1660	708 + 571	



The Monitor is the starting point for browsing the system information. It shows the current state of the system. Grid jobs are not kept infinitely in the information system. Those whose lifetime has expired are removed from clusters and thus they disappear from the monitor too. All the elements in the Monitor are links to the information about the different objects in the system.

Examples of cluster information

1) Attributes: dump of all the attributes of the nordugrid-cluster object

2) Queues at a given cluster

Number of occupied CPUs

General information

Job status information

1) Cluster alias
2) Cluster load showing both Grid- and non-Grid jobs
3) Queuing jobs

The NLogger is a graphical interface to the underlying MySQL database. The logger service is the persistent database for job usage records. It provides information which is complementary to the Monitor, showing the history of the Nordugrid usage. The queries are based on different basic approaches to the database like cluster, time period, application or user.

Data Challenge 2 on Nordugrid

Parameters to be set by the user

- 1) time period
- 2) datasets
- 3) cluster

Statistics for a specific cluster

A user can get access to information about his or her usage of Nordugrid

Select - time period - cluster - jobs

Current year summary for a specific cluster on Nordugrid.

References:

Nordugrid live: www.nordugrid.org - Grid Monitor and NLogger

Eerola, P.; Ekelöf, T.; Ellert, M.; Hansen, J R.; Helman, S.; Konstantinov, A.; Konya, B.; Mykkestub, T.; Nielsen, J L.; Ould-Saada, F.; Smirnova, O.; Waananen, A. *Atlas Data-Challenge I on Nordugrid* [arXiv:physics/0306013]

The Nordugrid/ARC User Guide, Available: <http://www.nordugrid.org/documents/userguide.pdf>

Conclusion

Nordugrid offers a user friendly monitoring service showing the present situation of the system as well as a complementary tool, the logger service, which provides the history of the activity on the grid. Both services are web based offer graphical interfaces to the usage data.