



LUND UNIVERSITY

Usage statistics and usage patterns on the NorduGrid: Analyzing the logging information collected on one of the largest production Grids of the world

Pajchel, K.; Eerola, Paula; Konya, Balazs; Smirnova, Oxana; Erkarlan, U.; Ould-Saada, F.; Ellert, M.; Renner Hansen, J.; Langgard Nielsen, J.; Wäänänen, A.; Ekelöf, T.; Konstantinov, A.; ATLAS Collaboration

Published in:
CERN-2005-002

2005

[Link to publication](#)

Citation for published version (APA):

Pajchel, K., Eerola, P., Konya, B., Smirnova, O., Erkarlan, U., Ould-Saada, F., Ellert, M., Renner Hansen, J., Langgard Nielsen, J., Wäänänen, A., Ekelöf, T., Konstantinov, A., & ATLAS Collaboration (2005). Usage statistics and usage patterns on the NorduGrid: Analyzing the logging information collected on one of the largest production Grids of the world. In A. Aimar, J. Harvey, & N. Knoors (Eds.), *CERN-2005-002* (Vol. 2, pp. 711-714). CERN.

Total number of authors:
13

General rights

Unless other specific re-use rights are stated the following general rights apply:
Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

USAGE STATISTICS AND USAGE PATTERNS ON THE NORDUGRID ANALYZING THE LOGGING INFORMATION COLLECTED ON ONE OF THE LARGEST PRODUCTION GRIDS IN THE WORLD.

K. Pajchel*, Dr. U. Erkarlan, Prof. F. Ould-Saada, University of Oslo, Norway
Dr. M. Ellert, CERN, Geneva, Switzerland

Prof. P. Eerola, Dr. B. Kónya, Dr. O. Smirnova, Lund University, Sweden

Prof. J. R. Hansen, Dr. J. L. Nielsen, Dr. A. Wäänänen, Niels Bohr Institute, Copenhagen,
Denmark

Prof. T. Ekelöf, Uppsala University, Sweden

Dr. A. Konstantinov, Vilnius University, Lithuania and University of Oslo, Norway

Abstract

The Nordic Grid facility (NorduGrid [1]) 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 [2] and the ongoing Data Challenge 2.

- The increasing number of resources has made NorduGrid one of the largest production Grids in 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 [3]) 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,

heterogeneous 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, gridftpd (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 gridftpd and the information service.
- Indexing services for the resources and data.
- Client making 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.

* katarina.pajchel@fys.uio.no

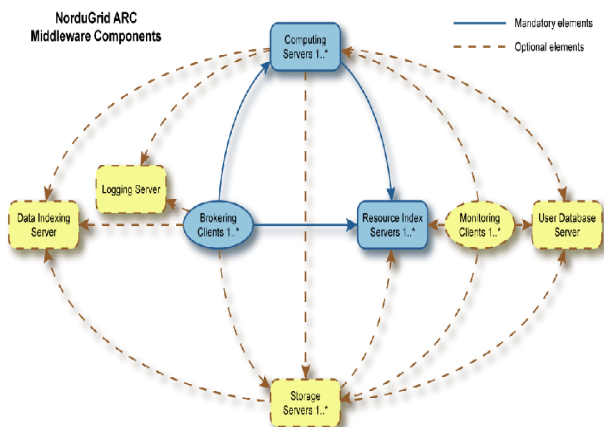


Figure 1: NorduGrid's ARC middleware components.

LOGGING

A Logger service is one of the Web services implemented by NorduGrid and based on gSOAP and Globus IO API. It provides a frontend to the underlying MySQL database to store and retrieve information about computing resources' usage (jobs).

This database can be accessed through a graphical Web interface implemented using PHP4, JavaScript and JGraph based on GD library.

The main goals for the Logger:

- Provide information about the development and usage of NorduGrid over time.
- Provide 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/CN=host/lscf.nbi.dk
user       Name of the user
id         46236674234684697
name       name of the job
failure    Job exit code is 265 != 0
lrms       pbs, condor
queue      fastq, slowq, gridlong
rsl        content of the .xrsl file
ui         129.240.86.35:57818;susy.uio.no
```

Figure 2: Content of the Logger database. The columns shown are temporary. The database is being extended.

THE GRID MONITOR

The Grid Monitor allows the user and system administrators 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 objects:

- 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 to a given queue

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

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")
)
```

Once a proxy is obtained, jobs can be submitted to NorduGrid using simple commands and the user can keep track of the job through the whole process.

- 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 and can later be displayed using the Web based Logger.

GRAPHICAL INTERFACES

Nordugrid offers intuitive access to the system data through two different Web-based graphical interfaces. The users can easily follow their jobs and obtain information about the performance of the system and job statistics.


Grid Monitor

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.



Country	Site	CPU's	Load (processes: Grid+local)	Queueing
Australia	Atlas (UniMelb)	28	0+2	0+0
	Alfred (UniMelb)	88	0+35	2+1
	DistLab (DIKU)	10	0+0	0+0
Denmark	Aalborg Grid Gateway	46	2+0	0+0
	Horseshoe (DCSC/SDU)	1172	0+0	0+75
	HEPAKI	1	0+0	0+0
	Morpheus	18	10+0	2+0
	VER (VideoRecorder)	1	0+0	0+0
Estonia	UT IMCB Anakonda clus>	15	0+0	0+0
	UT CS Antarctica Clus>	20	0+0 (queue down)	0+0
	CMS on CERN Linux	1	0+0	0+0
	CMS Production server	5	0+0	0+0
	UT DOUG Cluster	2	0+0	0+0
Finland	CMS test cluster	1	0+0	0+0
	UT Physics Cluster	6	0+0	0+0
	UT Physics Cluster	17	0+0	0+0
	CSC Kippu	1	1+0	0+0
	Mill (Physicum)	60	0+0 (queue down)	0+4
Germany	Alpha (HIP)	1	0+0	0+0
	Testbed0 (HIP)	3	0+0	4+1
	FZK cluster	740	30+253	0+0
Norway	LRZ cluster	234	0+218	0+487
	Oslo Temp Cluster	13	0+0	2+0
	Parallab IBM Cluster	58	0+0	0+0
Slovakia	Oslo Grid Cluster	44	1+0	23+0
	Uio Grid	88	0+79	0+1
	UPJS GRID	1	0+0	0+0
Slovenia	SIGNET	51	0+0	9+0
	BTU Smoke (SweGrid, NS>	98	0+0	378+0
	Kosofy farm	60	0+0	44+0
Sweden	ISV	4	0+0	0+0
	Hagrid (SweGrid, Uppm>	100	100+0	158+0
	Ingrid (SweGrid, HPC2N)	101	101+0	51+0
Switzerland	Quark Cluster	7	0+0	0+0
	Beppe (SweGrid PDC KT>	95	91+0	31+0
	Toto7/Whenim64 (Lunar>	192	0+143	0+2
TOTAL	37 sites	3304	501 + 1660	708 + 571

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



Attribute	Value
Distinguished name	nordugrid-cluster-name=grid.uio.no,Mds-Vo-name=local,ou=grid
objectClass	Mds
Front-end domain name	nordugrid-cluster
Cluster alias	grid.uio.no
Owner	University of Oslo
Postal code	N-0316
Certificate issuer	/O=Grid/O=NorduGrid/CN=NorduGrid Certification Authority
Contact string	gsiftp://grid.uio.no:2811/jobs
E-mail contact	grid.support@fys.uio.no
LRMS type	Condor
LRMS version	6.5.5
LRMS details	vanilla, single job per processor
Architecture	i686
Operating system	Red Hat Enterprise Linux WS release 3 (Taroon Update 2)
Homogeneous cluster	False
Memory (MB, smallest)	640
Node IP connectivity	outbound
CPU's, total	44
CPU's, occupied	18
CPU:machines	1cpu:16 2cpu:14
Jobs, queued	22
Jobs, total amount	16
Storage Element, local	gsiftp://dc1.uio.no/dc1/
	gsiftp://dc2.uio.no/dc1/
	gsiftp://dc3.uio.no/dc1/
	gsiftp://dc4.uio.no/dc1/
Disk space, available (MB)	24711
Disk space, total (MB)	38646
Cache size, available (MB)	5335
Cache size, total (MB)	9536
Grid middleware	nordugrid-untagged
Runtime environment	globus-2.4.3-10ng
	ATLAS-8.0.5
	DS2003
	BIO-GEUJER-0.0.1
	BIO-GEUJER-0.0.2
Mds-validfrom	2004-09-11 10:17:49
Mds-validto	2004-09-11 10:18:19

Figure 4: Example of cluster information

- 1) Attributes: dump of all the attributes of the nordugrid-cluster object.
- 2) Queues at a given cluster.

NGLogger

The NGLogger 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.




Figure 5: NGLogger

The user can get access to information about his or her usage of NorduGrid through a search form.

- 1) Search form.
- 2) Number of jobs.
- 3) Time from submission of the job until it is finished.

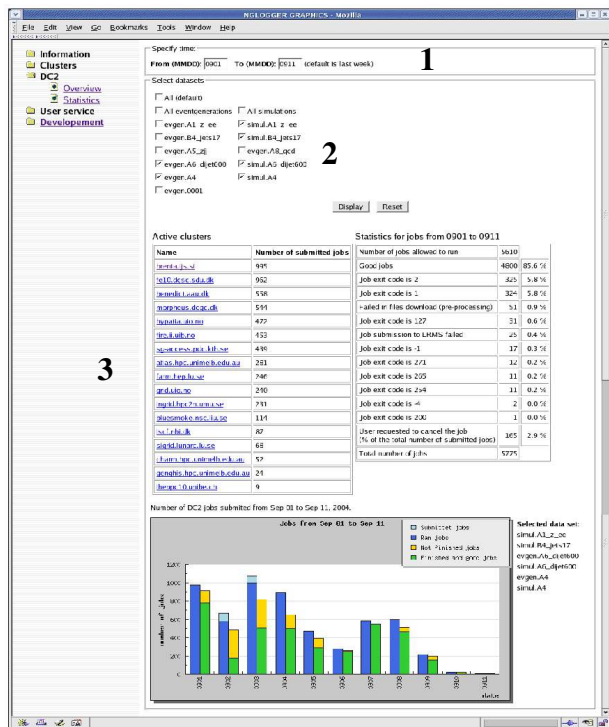


Figure 6: Data Challenge 2 on NorduGrid
Parameters to be set by the user:

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

CONCLUSION

NorduGrid offers a user-friendly monitoring service showing the present status 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 and offer graphical interfaces to the usage data.

REFERENCES

- [1] NorduGrid live: <http://www.nordugrid.org> - Grid Monitor and NGLogger
- [2] Eerola, P; Ekelof, T; Ellert, M; Hansen, J R; Hellman, S; Konstantinov, A; Konya, B; Myklebust, T; Nielsen, J L; Ould-Saada, F; Smirnova, O; Wäänänen, A, *Atlas Data-Challenge 1 on NorduGrid* [arXiv:physics/0306013]
- [3] The NorduGrid/ARC User Guide, Available: <http://www.nordugrid.org/documents/userguide.pdf>

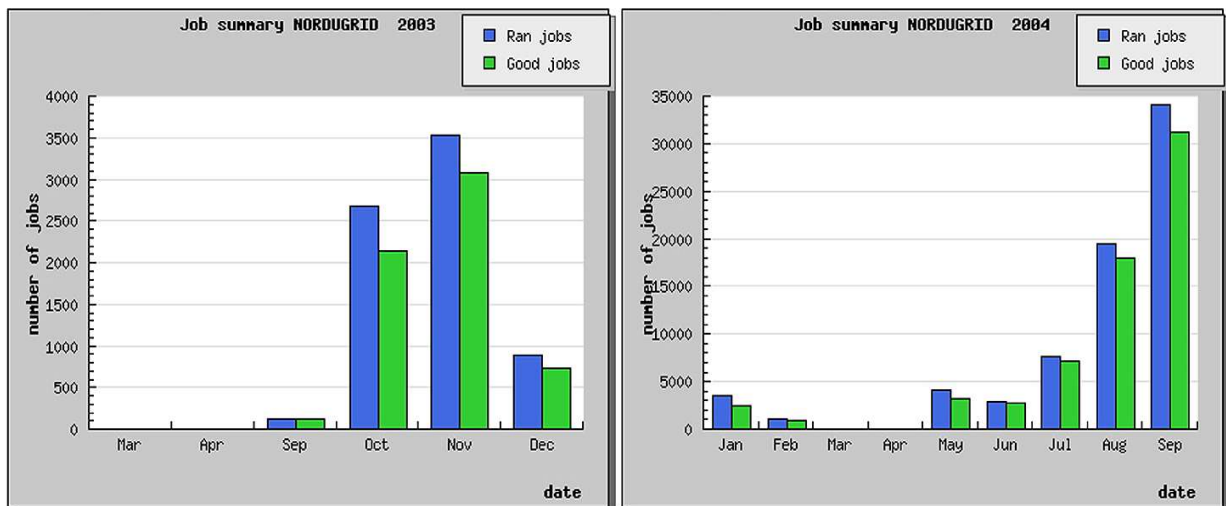


Figure 7: Number of jobs recorded in the logger. Over 30 000 jobs were run successfully in September 2004. Notice that it is not the Grid, but the experimental logger service which was down in March and April. The early information may be incomplete.