

NorduGrid Tutorial

Client Installation and Job Examples

Linux Clusters for Super Computing Conference
Linköping, Sweden
October 18, 2004

Arto Teräs
arto.teras@csc.fi

Steps to Start Using NorduGrid

- 1) Install the client software
- 2) Request a certificate from a Certificate Authority (CA)
- 3) Install the certificate
- 4) Log in to the Grid
- 5) Test the installation
- 6) Write a job description using xRSL language
- 7) Submit the job
- 8) Monitor the progress of the job
- 9) Fetch the results

Installing the NorduGrid Client

- Required to submit jobs to NorduGrid
- Download from <http://ftp.nordugrid.org/download/>
 - Binaries for various Linux distributions, source code also available
- Easiest way to get started is to install the standalone client
 - Uncompress in a directory (no root privileges required):

```
$ tar zxvf nordugrid-standalone-0.4.3-1.i386.tgz
```
 - Run the environment setup script:

```
$ cd nordugrid-standalone-0.4.3  
$ . ./setup.sh
```
- RPM packages are recommended for multi-user installations

Requesting and Installing the Certificate

- Create a certificate request
 - \$ `grid-cert-request -int`
- This generates directory `.globus` in your home directory and inside it a file named `usercert_request.pem` which should be sent to a Certification Authority (CA)
 - Check the instructions at your local site / country which CA to contact
- Wait for an answer from the CA
 - Signed certificate sent by the Certificate Authority should be saved as file `.globus/usercert.pem`

What Does a Certificate Look Like?

- Consists of two files:
 - Private key is protected by a password and kept secret
 - Public key is given out to third parties
 - Certificate Authorities sign the public key, even they never see the private key
- Look like a string of random numbers and letters, but tools can be used to convert the information in readable form
 - \$ `grid-cert-info -file <certificate file>`
 - For example, my identity stored in my NorduGrid certificate is "O=Grid, O=NorduGrid, OU=csc.fi, CN=Arto Teras"

Logging in and Testing the Installation

- Log in to the Grid
 - \$ `grid-proxy-init`
- Use command `ngtest` to test the installation
 - \$ `ngtest 1 -d 1` (send test job 1, show level 1 debug info)
 - \$ `ngget ngtest-job-1` (fetch result files of the test job)
- In case of problems, read the manual and frequently asked questions list (FAQ), ask the mailing list ...

NorduGrid User Interface

- Set of command line utilities:
 - **ngsub** to submit a task
 - **ngstat** to obtain the status of jobs and clusters
 - **ngcat** to display the stdout or stderr of a running job
 - **ngget** to retrieve the result from a finished job
 - **ngkill** to cancel a job request
 - **ngclean** to delete a job from a remote cluster
 - **ngrenew** to renew user's proxy
 - **ngsync** to synchronize the local job info with the MDS
 - **ngcopy** to transfer files to, from and between clusters
 - **ngremove** to remove files

Writing a Job Description File

- Resource Specification Language (RSL) files are used to specify job requirements and parameters for submission
 - NorduGrid uses an extended language (xRSL) based on the Globus RSL
- Similar to scripts for local queueing systems, but include some additional attributes
 - Job name
 - Executable location and parameters
 - Location of input and output files of the job
 - Architecture, memory, disk and CPU time requirements
 - Library dependencies and version requirements

xRSL example

```
& (executable=hellogrid.sh)
  (jobname=hellogrid)
  (stdout=hello.out)
  (stderr=hello.err)
  (gmlog=gridlog)
  (architecture=i686)
  (cputime=10)
  (memory=32)
  (disk=1)
```

Submitting a Job

- Submit the job

```
$ ngsbub -f hellogrid.xrsl
```

```
=> Job submitted with jobid gsiftp://morpheus.  
dcgc.dk:2811/jobs/1757591474592630108
```

- Fetch the results

```
$ ngget hellogrid
```

```
=> ngget: downloading files to  
/home/ajt/1757591474592630108  
ngget: download successful - deleting job  
from gatekeeper.
```

Monitoring the Jobs

- Status of jobs can be queried with command `ngstat`

```
$ ngstat hellogrid
```

```
=> Job
```

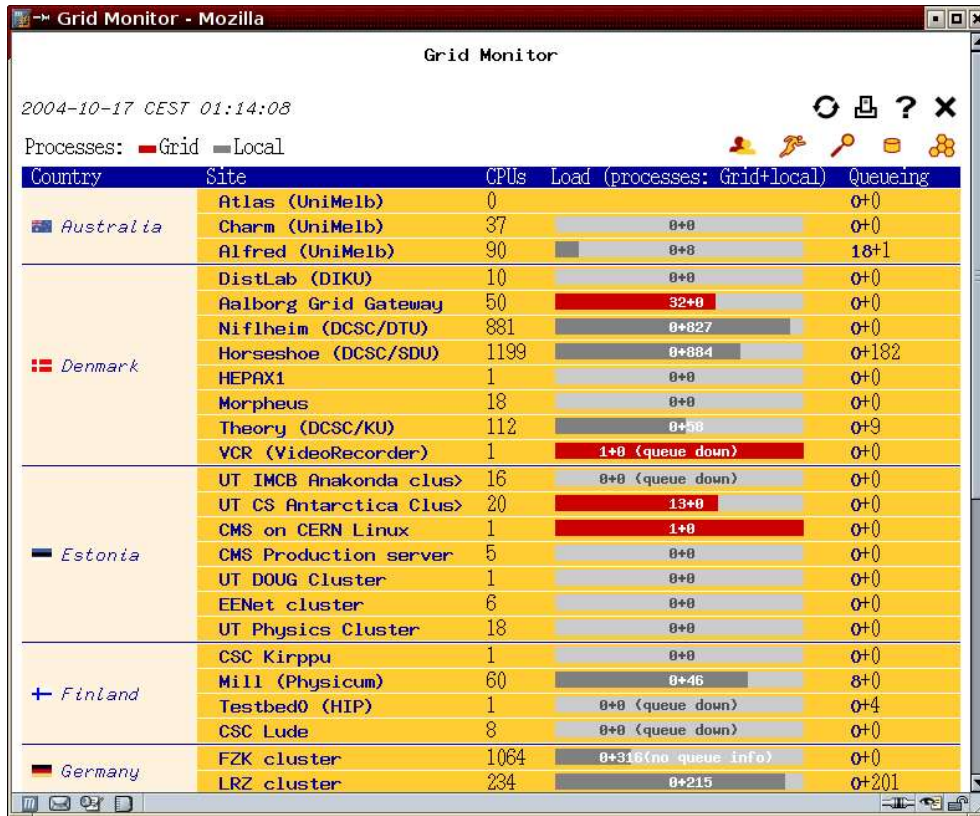
```
gsiftp://ingvar.nsc.liu.se:2811/jobs/5436235811735113812  
  Jobname: hellogrid  
  Status: FINISHING
```

```
Job
```

```
gsiftp://datagrid3.csc.fi:2811/jobs/1593889897762957743  
  Jobname: hellogrid  
  Status: ACCEPTED
```

- Grid monitor on the NorduGrid website is also a useful monitoring tool

Grid Monitor on NorduGrid Website



2004-10-17 CEST 01:14:08

Processes: ■ Grid ■ Local

Country	Site	CPUs	Load (processes: Grid+Local)	Queueing
Australia	Atlas (UniMelb)	0		0+0
	Charm (UniMelb)	37	0+0	0+0
	Alfred (UniMelb)	90	0+0	18+1
Denmark	DistLab (DIKU)	10	0+0	0+0
	Aalborg Grid Gateway	50	32+0	0+0
	Niflheim (DCSC/DTU)	881	0+827	0+0
	Horseshoe (DCSC/SDU)	1199	0+884	0+182
	HEPAX1	1	0+0	0+0
	Morpheus	18	0+0	0+0
	Theory (DCSC/KU)	112	0+58	0+9
	VCR (VideoRecorder)	1	1+0 (queue down)	0+0
	UT IMCB Anakonda clus>	16	0+0 (queue down)	0+0
Estonia	UT CS Antarctica Clus>	20	13+0	0+0
	CMS on CERN Linux	1	1+0	0+0
	CMS Production server	5	0+0	0+0
	UT DOUG Cluster	1	0+0	0+0
	EENet cluster	6	0+0	0+0
	UT Physics Cluster	18	0+0	0+0
Finland	CSC Kirppu	1	0+0	0+0
	Mill (Physicum)	60	0+46	8+0
	Testbed0 (HIP)	1	0+0 (queue down)	0+4
	CSC Lude	8	0+0 (queue down)	0+0
Germany	FZK cluster	1064	0+316 (no queue info)	0+0
	LRZ cluster	234	0+215	0+201

- Shows currently connected resources
- Almost all elements "clickable"
 - browse queues and job states by cluster
 - list jobs belonging to a certain user
- No authentication, anyone can browse the info

Using a Storage Element

- Storage Elements are disk servers accessible via the Grid
- Allows to store input files close to the cluster where the program is executed, on a high bandwidth network
- Possibility to upload output files at a desired place:

```
(inputFiles=  
  ('input1', '/home/user/myexperiment'  
  ('input2',  
  'gsiftp://se.somewhere.ee/files/commondata'))  
  
(outputFiles=  
  ('output', 'gsiftp://se.somewhere.com/mydir/result1')  
  ('prog.out', 'gsiftp://se.somewhere.com/mydir/stdout'))  
  
(stdout='prog.out')
```

xRSL Example Using a Storage Element

- xRSL file for the `hellogrid` example, uploading the job results to a storage element:

```
& (executable=hellogrid.sh)
(jobname=hellogrid-se)
(stdout=gsiftp://grid.tsl.uu.se/tutorial/hello.out)
(stderr=gsiftp://grid.tsl.uu.se/tutorial/hello.err)
(gmlog=gridlog)
(architecture=i686)
(cputime=10)
(memory=32)
(disk=1)
```

Gsincftp

- Can be used to transfer files to and from storage elements
 - Based on the popular `ncftp` ftp client, but uses certificate based authentication instead of standard ftp authentication
- Example session:

```
$ gsincftp grid.tsl.uu.se
...Logged in to grid.tsl.uu.se.

$ cd tutorial

$ get hello.out
```
- Already deprecated by the Globus project, does not work with their newest GridFTP server
 - replacement: UberFTP (<http://dims.ncsa.uiuc.edu/set/uberftp/>)

Runtime Environments

- Software packages which are preinstalled on a computing resource
 - Avoid the need of sending the binary at the start of executing a job
 - Allow local optimizations (e.g. compiling to the installed architecture using optimized compiler flags)
- Very useful if there are many users of the same software or if the same program is used frequently
- Required runtime environment(s) can be specified in the job description file (xRSL file), for example:

(runtimeenvironment=povray-3.5)

Real Jobs

- Real jobs usually send several subjobs to the Grid to solve a larger problem
- Parallel MPI jobs to a single cluster supported (if correct runtime environment installed), but no MPI between clusters
- Splitting the job to suitable parts and gathering the parts together is left to the user
 - More error prone environment than traditional local systems => error checking and recovery important
 - Fault reporting and debugging has room for improvements
- Leif Nixon's example: Rendering an image in slices using the **povray** tool

References

- NorduGrid website: <http://www.nordugrid.org>
- The NorduGrid User Guide:
<http://www.nordugrid.org/documents/userguide.pdf>
- Balazs Konya's presentation at the 4th International Workshop on Grid Computing:
<http://www.nordugrid.org/slides/20031117-balazs.pdf>
- Povray example by Leif Nixon:
<http://www.nsc.liu.se/~nixon/ng-povray/>

Thank you!