



*Nordic Testbed for Wide Area
Computing and Data Handling*

2003-08-25

The NorduGrid toolkit user interface User's manual

Mattias Ellert

Abstract

This is the user's manual to the NorduGrid toolkit user interface.

Table of Contents

1 The NorduGrid user interface commands.....	4
1.1 ngsb.....	4
1.2 ngstat.....	5
1.3 ngcat.....	6
1.4 ngget.....	7
1.5 ngkill.....	8
1.6 ngclean.....	9
1.7 ngrenew.....	10
1.8 ngsync.....	10
2 User configuration files.....	11
2.1 .ngrc.....	11
2.2 .nggiislist.....	12
2.3 .nghistory.....	12
2.4 .ngjobs.....	12

1 The NorduGrid user interface commands

The user interface part of the NorduGrid toolkit consists of a set of command line commands. These are described in the following sections.

1.1 *ngsub*

The *ngsub* command is used for submitting jobs to the NorduGrid. The jobs are described using the extended resource description language (xrsl).

Usage: *ngsub* [options] [xrsl]

Options:

<i>-c, -cluster</i>	<i>[-]name</i>	explicitly select or reject a specific cluster
<i>-g, -giisurl</i>	<i>url</i>	url to the central giis
<i>-f, -file</i>	<i>filename</i>	xrslfile describing the job to be submitted
<i>-o, -joblist</i>	<i>filename</i>	file where the jobids will be stored
<i>-dryrun</i>		add dryrun option to the xrsl
<i>-dumpxrsl</i>		do not submit – dump transformed xrsl to stdout
<i>-t, -timeout</i>	<i>time</i>	timeout for MDS queries in seconds (default 40)
<i>-d, -debug</i>	<i>debuglevel</i>	0 = none, 1 = some, 2 = more, 3 = a lot
<i>-x, -anonymous</i>		use anonymous bind for MDS queries
<i>-X, -gsi</i>		use gsi-gssapi bind for MDS queries
<i>-v, -version</i>		print version information
<i>-h, -help</i>		print this help

Arguments:

xrsl ... xrslstrings describing the jobs to be submitted

The *ngsub* command locates the available clusters by querying a globus information index server (GIIS). Which GIIS to query can be specified by giving the *-giisurl* option to the command. The url to the GIIS should be given in the format

```
ldap://hostname[:port]/basedn
```

If the port number is omitted the default 2135 is used. Two different syntax for the base dn of the ldap query are accepted, e.g. can the NorduGrid GIIS in Copenhagen be specified in either of the two following ways:

```
ldap://grid.nbi.dk:2135/O=Grid/Mds-Vo-name=NorduGrid
ldap://grid.nbi.dk:2135/Mds-Vo-name=NorduGrid,O=Grid
```

Several GIISes can be specified by repeating the *-giisurl* option:

```
ngsub -giisurl url1 -giisurl url2 ...
```

In this case the GIIS specified as *url1* will be queried first, and if no list of clusters

can be retrieved the GIIS specified as *url2* will be queried.

If no `-giisurl` option is given a list of urls is read from a `giislist` file. The first existing file in the following list is used:

1. `$(HOME)/nggiislist`
2. `$(NORDUGRID_LOCATION)/etc/giislist`
or `/opt/nordugrid/etc/giislist` if `$(NORDUGRID_LOCATION)` is undefined
3. `/etc/giislist`

The `-cluster` option can be used to force the job to be submitted to a particular cluster, or to reject submission to a particular cluster. The `-cluster` option can be repeated several times.

Examples:

```
ngsub -c grid.nbi.dk -c grid.tsl.uu.se ...  
submit the job to either grid.nbi.dk or grid.tsl.uu.se
```

```
ngsub -c -grid.quark.lu.se -c -grid.nbi.dk ...  
don't submit the job to grid.quark.lu.se or grid.nbi.dk
```

The `xrsl` that describes the job can be given either as an argument on the command line or can be read from a file by using the `-file` option. Several jobs can be requested at the same time by giving more than one `xrsl` argument or by repeating the `-file` option. It is also possible to mix `xrsl` arguments and `-file` options in the same `ngsub` command.

The user interface transforms the `xrsl` into a format that can be understood by the grid-manager to which it is being submitted. By specifying the `-dumpxrsl` option the transformed `xrsl` is written to `stdout` instead of being submitted to the remote cluster.

If the job is successfully submitted a `jobid` is printed to `stdout`. This `jobid` uniquely identifies the job while it is being executed. If the `-joblist` option is given the `jobid` is also written to a file with the specified filename. This file can later be used with the `-joblist` option of the other user interface commands.

For setting defaults, see the section about `.ngrc` (section 2.1).

1.2 ngstat

The `ngstat` command is used for obtaining the status of jobs that have been submitted to NorduGrid and the status of the clusters in the grid.

Usage: `ngstat [options] [job ...]`

Options:

-a, -all		all jobs
-i, -joblist	<i>filename</i>	file containing a list of jobids
-c, -cluster	<i>[-]name</i>	explicitly select or reject a specific cluster
-s, -status	<i>statusstr</i>	only select jobs whose status is <i>statusstr</i>
-g, -giisurl	<i>url</i>	url to the central giis
-q, -queues		show information about clusters and queues
-l, -long		long format (more information)
-t, -timeout	<i>time</i>	timeout for MDS queries in seconds (default 40)
-d, -debug	<i>debuglevel</i>	0 = none, 1 = some, 2 = more, 3 = a lot
-x, -anonymous		use anonymous bind for MDS queries
-X, -gsi		use gsi-gssapi bind for MDS queries
-v, -version		print version information
-h, -help		print this help

Arguments:

job ... list of jobids and/or jobnames

The `-giisurl` and `-cluster` options have been described in the `ngsub` section (section 1.1).

The `ngstat` command gives the status of a job submitted to NorduGrid. The job can be referred to either by the jobid that was returned by `ngsub` at submission time or by its jobname if the `xrsl` that was submitted contained a jobname attribute.

More than one jobid and/or jobname can be given. If several jobs were submitted with the same jobname the status of all those jobs are shown. If the `-joblist` option is used the list of jobids is read from a file with the specified filename. By specifying the `-all` option, the status of all active jobs will be shown. If the `-long` option is given more detailed information is given.

If the `-queues` option is used, `ngstat` gives the status of the clusters and queues available rather than of the jobs submitted. Also in this case the `-long` option can be used to obtain more detailed information.

For setting defaults, see the section about `.ngrc` (section 2.1).

1.3 *ngcat*

The `ngcat` command can be used to view the stdout or stderr of a running job. It can also be used to show the grid manager's error log of a job.

Usage: `ngcat [options] [job ...]`

Options:

-a, -all		all jobs
-i, -joblist	<i>filename</i>	file containing a list of jobids

Options:

-c, -cluster	<i>[-]name</i>	explicitly select or reject a specific cluster
-s, -status	<i>statusstr</i>	only select jobs whose status is <i>statusstr</i>
-o, -stdout		show the stdout of the job (default)
-e, -stderr		show the stderr of the job
-l, -gmlog		show the grid manager's error log of the job
-t, -timeout	<i>time</i>	timeout for MDS queries in seconds (default 40)
-d, -debug	<i>debuglevel</i>	0 = none, 1 = some, 2 = more, 3 = a lot
-x, -anonymous		use anonymous bind for MDS queries
-X, -gsi		use gsi-gssapi bind for MDS queries
-v, -version		print version information
-h, -help		print this help

Arguments:

job ... list of jobids and/or jobnames

The `ngcat` command displays the stdout or stderr of running jobs. It can also display the grid manager's error log of a job. The job can be referred to either by the jobid that was returned by `ngsub` at submission time or by its jobname if the xrsl that was submitted contained a jobname attribute.

More than one jobid and/or jobname can be given. If several jobs were submitted with the same jobname the stdout, stderr or grid manager error log of all those jobs are shown. If the `-joblist` option is used the list of jobids is read from a file with the specified filename. By specifying the `-all` option, the stdout, stderr or grid manager error log of all active jobs will be shown.

Only jobs where the stdout, stderr and gmlog arguments were given in the xrsl can display the contents of those files.

For setting defaults, see the section about `.ngrc` (section 2.1).

1.4 ngget

The `ngget` command is used for retrieving the results from a job

Usage: `ngget [options] [job ...]`

Options:

-a, -all		all jobs
-i, -joblist	<i>filename</i>	file containing a list of jobids
-c, -cluster	<i>[-]name</i>	explicitly select or reject a specific cluster
-s, -status	<i>statusstr</i>	only select jobs whose status is <i>statusstr</i>
-k, -keep		keep files on gatekeeper (do not clean)
-dir	<i>dirname</i>	download directory (the job directory will be created in this directory)

Options:

-j, -usejobname		use the jobname instead of the short ID as the job directory name
-t, -timeout	<i>time</i>	timeout for MDS queries in seconds (default 40)
-d, -debug	<i>debuglevel</i>	0 = none, 1 = some, 2 = more, 3 = a lot
-x, -anonymous		use anonymous bind for MDS queries
-X, -gsi		use gsi-gssapi bind for MDS queries
-v, -version		print version information
-h, -help		print this help

Arguments:

job ...	list of jobids and/or jobnames
---------	--------------------------------

This command downloads the results after a job has completed on the NorduGrid. Only the results of jobs that have finished can be downloaded. The job can be referred to either by the jobid that was returned by `ngsub` at submission time or by its jobname if the `xrsl` that was submitted contained a jobname attribute.

More than one jobid and/or jobname can be given. If several jobs were submitted with the same jobname the results of all those jobs are downloaded. If the `-joblist` option is used the list of jobids is read from a file with the specified filename. By specifying the `-all` option, the results of all active jobs are downloaded.

For each job that is downloaded a subdirectory will be created in the download directory that will contain the downloaded files.

If the download was successful the job will be removed from the remote cluster unless the `-keep` option was specified.

For setting defaults, see the section about `.ngrc` (section 2.1).

1.5 *ngkill*

The `ngkill` command is used to kill running jobs in the NorduGrid.

Usage: `ngkill [options] [job ...]`

Options:

-a, -all		all jobs
-i, -joblist	<i>filename</i>	file containing a list of jobids
-c, -cluster	<i>[-]name</i>	explicitly select or reject a specific cluster
-s, -status	<i>statusstr</i>	only select jobs whose status is <i>statusstr</i>
-t, -timeout	<i>time</i>	timeout for MDS queries in seconds (default 40)
-d, -debug	<i>debuglevel</i>	0 = none, 1 = some, 2 = more, 3 = a lot
-x, -anonymous		use anonymous bind for MDS queries
-X, -gsi		use gsi-gssapi bind for MDS queries
-v, -version		print version information

of jobids is read from a file with the specified filename. By specifying the `-all` option, all active jobs can be removed.

The `-force` option removes the job from your local list of jobs (see section 2.4) even if the job can not be found in the MDS.

For setting defaults, see the section about `.ngrc` (section 2.1).

1.7 *ngrenew*

The `ngrenew` command can be used to renew the proxies of NorduGrid jobs.

Usage: `ngrenew [options] [job ...]`

Options:

<code>-a, -all</code>		all jobs
<code>-i, -joblist</code>	<i>filename</i>	file containing a list of jobids
<code>-c, -cluster</code>	<code>[-]name</code>	explicitly select or reject a specific cluster
<code>-s, -status</code>	<i>statusstr</i>	only select jobs whose status is <i>statusstr</i>
<code>-t, -timeout</code>	<i>time</i>	timeout for MDS queries in seconds (default 40)
<code>-d, -debug</code>	<i>debuglevel</i>	0 = none, 1 = some, 2 = more, 3 = a lot
<code>-x, -anonymous</code>		use anonymous bind for MDS queries
<code>-X, -gsi</code>		use <code>gsi-gssapi</code> bind for MDS queries
<code>-v, -version</code>		print version information
<code>-h, -help</code>		print this help

Arguments:

<code>job ...</code>	list of jobids and/or jobnames
----------------------	--------------------------------

This command renews the proxy of a job on the NorduGrid. The job can be referred to either by the jobid that was returned by `ngsub` at submission time or by its jobname if the `xrsl` that was submitted contained a `jobname` attribute.

More than one jobid and/or jobname can be given. If several jobs were submitted with the same jobname the proxies of all those jobs are renewed. If the `-joblist` option is used the list of jobids is read from a file with the specified filename. By specifying the `-all` option, the proxies of all jobs can be renewed.

For setting defaults, see the section about `.ngrc` (section 2.1).

1.8 *ngsync*

The `ngsync` command synchronizes your local list of jobs with the information in the MDS.

Usage: `ngsync [options]`

Options:

<code>-c, -cluster</code>	<code>[-]name</code>	explicitly select or reject a specific cluster
<code>-g, -giisurl</code>	<code>url</code>	url to the central giis
<code>-f, -force</code>		don't ask for verification
<code>-t, -timeout</code>	<code>time</code>	timeout for MDS queries in seconds (default 40)
<code>-d, -debug</code>	<code>debuglevel</code>	0 = none, 1 = some, 2 = more, 3 = a lot
<code>-x, -anonymous</code>		use anonymous bind for MDS queries
<code>-X, -gsi</code>		use gsi-gssapi bind for MDS queries
<code>-v, -version</code>		print version information
<code>-h, -help</code>		print this help

The `-giisurl` and `-cluster` options have been described in the `ngsub` section (section 1.1).

The NorduGrid toolkit user interface keeps a local list of active jobs in the users home directory (see section 2.4). If this file is lost, or the user wants to recreate the file on a different computer, the `ngsync` command can be used to recreate the file from the information available in the MDS.

Since the information about a job in the MDS can be slightly out of date if the user very recently submitted or removed a job a warning is issued when this command is run. The `-force` option disables this warning.

For setting defaults, see the section about `.ngrc` (section 2.1).

2 User configuration files

There are a few files located in the user's home directory that are used for configuring the NorduGrid user interface. These files all start with `.ng`.

2.1 `.ngrc`

This file contains the users default settings for the debug level, the MDS query timeout and the download directory used by `ngget` (see section 1.4). A sample file could be the following.

```
# Sample .ngrc file
# Comments starts with #
NGDEBUG=1
NGTIMEOUT=60
NGDOWNLOAD=/tmp
```

If the file is missing or some options are not specified in the file, the default defaults are used. The default debug level is 0, the default MDS timeout is 40 seconds and the default download directory is the current working directory.

If the environment variables `NGDEBUG`, `NGTIMEOUT` or `NGDOWNLOAD` are defined these take precedence over the values defined in `.ngrc`. Any command line

options overrides the defaults.

2.2 .nggiislist

The user can specify a personal giislist file different from the system wide default by creating the file `.nggiislist` in the home directory. See the discussion in the `ngsub` section (section 1.1) for the format of this file.

2.3 .nghistory

This file contains the jobids of the jobs the user has submitted together with the time of submission. This file is purely informational.

2.4 .ngjobs

This a local list of the user's active jobs. When a job is successfully submitted it is added to this list and when it is removed from the remote cluster it is removed from this list. This list is used as the list of all active jobs when the user specifies the `-all` option to the various NorduGrid toolkit user interface commands. For information about how to reconstruct this file if it is lost or on a different computer see the section about the `ngsync` command (section 1.8).