



NORDUGRID-MANUAL-13

21/3/2012

ARC CLIENTS

User Manual for client versions 1.0.0 and above

Contents

1	Introduction	5
2	Commands	7
2.1	Proxy utilities	7
2.1.1	arcproxy	7
2.1.2	arcslcs	9
2.2	Job submission and management	9
2.2.1	arcsb	9
2.2.2	arcstat	13
2.2.3	arccat	15
2.2.4	arcget	16
2.2.5	arcsync	17
2.2.6	arcinfo	18
2.2.7	arckill	18
2.2.8	arcclean	19
2.2.9	arc renew	20
2.2.10	arc resume	21
2.2.11	arc sub	22
2.2.12	arc migrate	23
2.3	Data manipulation	24
2.3.1	arcls	24
2.3.2	arccp	25
2.3.3	arcrm	27
2.3.4	arc mkdir	27
2.3.5	arc rmping	28
2.3.6	chelonia	28
2.4	Test suite	35
2.4.1	arctest	35
3	URLs	37
4	ARC Client Configuration	41
4.1	Block [common]	41
	defaultservices	41

rejectservices	42
verbosity	42
timeout	42
brokername	42
brokerarguments	42
joblist	43
bartender	43
proxypath	43
keypath	43
certificatepath	43
cacertificatesdirectory	43
cacertificatepath	44
vomsserverpath	44
username	44
password	44
keypassword	44
keysize	44
certificatelifetime	44
slcs	44
storedirectory	45
jobdownloaddirectory	45
idpname	45
4.2 Block [alias]	45
4.3 srms.conf	45
4.4 Deprecated configuration files	46

Chapter 1

Introduction

The command line user interface of ARC consists of a set of commands necessary for job submission and manipulation and data management. This manual replaces the older version in `NORDUGRID-MANUAL-1` and is valid for ARC versions 0.9 and above. Command line tools semantics is the same as in earlier versions of ARC, roughly following that of basic Linux commands and most common batch system commands. One obvious difference is change of the legacy prefix from “ng” to the more appropriate “arc”. This is not only a cosmetic change: **behaviour of the commands also have changed**, as did their functionalities and options.

<p>Users are strongly discouraged from modifying their old scripts by simply replacing “ng” with “arc” – results may be unpredictable.</p>
--

Chapter 2

Commands

2.1 Proxy utilities

ARC now comes complete with a set of utilities to create temporary user credentials (proxies) used to access Grid services.

2.1.1 arcproxy

In order to contact Grid services (submit jobs, copy data, check information etc), one has to present valid credentials. These are commonly formalized as so-called “proxy” certificates. There are many different types of proxy certificates, with different Grids and different services having own preferences. `arcproxy` is a powerful tool that can be used to generate most commonly used proxies. It supports the following types:

- pre-RFC GSI proxy
- RFC-compliant proxy (default)
- VOMS-extended proxy
- MyProxy delegation

`arcproxy` requires presence of user’s private key and public certificate, as well as the public certificate of their issuer CA.

arcproxy [options]

Options:

<code>-P, --proxy</code>	<i>path</i>	path to the proxy file
<code>-C, --cert</code>	<i>path</i>	path to the certificate file
<code>-K, --key</code>	<i>path</i>	path to the key file
<code>-T, --cadir</code>	<i>path</i>	path to the trusted certificate directory, only needed for VOMS client functionality
<code>-V, --vomses</code>	<i>path</i>	path to the VOMS server configuration file
<code>-S, --voms</code>	<i>voms[:command]</i>	Specify VOMS server (more than one VOMS server can be specified like this: -voms VOa:command1 -voms VOb:command2) :command is optional, and is used to ask for specific attributes(e.g. roles). Command options are:

		all – put all of this DN’s attributes into AC;
		list – list all of the DN’s attribute,will not create AC extension;
		/Role=yourRole – specify the role, if this DN has such a role, the role will be put into AC
		/voname/groupname/Role=yourRole – specify the VO,group and role; if this DN has such a role, the role will be put into AC
-o, --order	<i>group[:role]</i>	Specify ordering of attributes, Examples: -order /knowarc.eu/coredev:Developer,/knowarc.eu/testers:Tester -order /knowarc.eu/coredev:Developer,/knowarc.eu/testers:Tester
-G, --gsicom		use GSI communication protocol for contacting VOMS services
-O, --old		use GSI proxy (default is RFC 3820 compliant proxy)
-I, --info		print all information about this proxy. In order to show the Identity (DN without CN as suffix for proxy) of the certificate, the 'trusted certdir' is needed.
-r, --remove		removes the proxy file
-U, --user	<i>string</i>	username for MyProxy server
-L, --myproxysrv	<i>URL</i>	URL of MyProxy server
-M, --myproxycmd	<i>PUT GET</i>	command to MyProxy server. The command can be PUT and GET. PUT/put – put a delegated credential to MyProxy server; GET/get – get a delegated credential from MyProxy server, credential (certificate and key) is not needed in this case.
-c, --constraint	<i>string</i>	proxy constraints
-t, --timeout	<i>seconds</i>	timeout for network communication, in seconds (default 20 seconds)
-d, --debug	<i>verbosity</i>	verbosity level is one of FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG
-z, --conffile	<i>filename</i>	configuration file (on Linux, default \$HOME/.arc/client.conf)
-v, --version		print version information
-h, --help		print help page

Note that it does not make sense to specify the order of attributes if you have two or more different VOMS servers specified.

MyProxy functionality can be used together with VOMS functionality.

Supported constraints are:

- **validityStart=time** – e.g. 2008-05-29T10:20:30Z; time when certificate becomes valid. Default is now.
- **validityEnd=time** – time when certificate becomes invalid. Default is 43200 (12 hours) from start.
- **validityPeriod=time** – e.g. 43200 or 12h or 12H; for how long certificate is valid. If neither **validityPeriod** nor **validityEnd** are specified, default is 12 hours

- `vomsACvalidityPeriod=time` – e.g. 43200 or 12h or 12H; for how long the AC is valid. Default is the same as `validityPeriod`.
- `proxyPolicy=policy content` – assigns specified string to proxy policy to limit it's functionality.
- `proxyPolicyFile=policy file`

2.1.2 arslcs

This utility generates short-lived credential based on the credential to IdP in SAML2SSO profile (normally the username/password to Shibboleth IdP).

arslcs [options]

Options:

<code>-S, --ur;</code>	<i>URL</i>	URL of SLCS Service (e.g. <code>https://127.0.0.1:60000/slcs</code>)
<code>-I, --idp</code>	<i>URL</i>	the name of IdP (e.g. <code>https://idp.testshib.org/idp/shibboleth</code>)
<code>-U, --user</code>	<i>string</i>	User account to IdP
<code>-P, --password</code>	<i>string</i>	password for user account to IdP
<code>-Z, --keysize</code>	<i>integer</i>	size of the private key, default is 1024
<code>-K, --keypass</code>		passphrase for protecting the private key; if not set, the private key file will not be protected by the passphrase.
<code>-L, --lifetime</code>	<i>hours</i>	life time of the credential (hours), starting with current time
<code>-D, --storedir</code>	<i>path</i>	store directory of the credential
		communication, in seconds (default 20 seconds)
<code>-d, --debug</code>	<i>verbosity</i>	verbosity level is one of FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG
<code>-z, --conffile</code>	<i>filename</i>	configuration file (on Linux, default <code>\$HOME/.arc/client.conf</code>)
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page

2.2 Job submission and management

The following commands are used for job submission and management, such as status check, results retrieval, cancellation, re-submission and such. The jobs must be described using a job description language. ARC supports the following languages: JSDL [2], xRSL [9] and JDL [7].

2.2.1 arcsub

The `arcsub` command is the most essential one, as it is used for submitting jobs to the Grid resources. `arcsub` matches user's job description to the information collected from the Grid, and the optimal site is being selected for job submission. The job description is then being submitted to that site, and usually is then forwarded to a local Resource Management System (LRMS), which can be, e.g., PBS or Condor or SGE etc.

arcsub [options] [filename ...]

Options:

<code>-c, --cluster</code>	<code>[-] designator</code>	explicitly select or reject (-) a specific site
<code>-g, --index</code>	<code>[-] designator</code>	explicitly select or reject (-) a specific index server
<code>-e, --jobdescrstring</code>	<i>filename</i>	string describing the job to be submitted
<code>-f, --jobdescrfile</code>	<i>filename</i>	file describing the job to be submitted
<code>-j, --joblist</code>	<i>filename</i>	the file storing information about active jobs (on Linux, default <code>\$/arc/jobs.xml</code>)
<code>-o, --jobids-to-file</code>	<i>filename</i>	the IDs of the submitted jobs will be appended to this file
<code>-D, --dryrun</code>		add dryrun option to the job description
<code>-x, --dumpdescription</code>		do not submit – dump transformed job description to stdout
<code>-b, --broker</code>	<i>string</i>	select broker method (default is Random)
<code>-P, --listplugins</code>		list the available plugins
<code>-t, --timeout</code>	<i>seconds</i>	timeout for network communication, in seconds (default 20)
<code>-d, --debug</code>	<i>verbosity</i>	verbosity level, FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG - default WARNING
<code>-z, --conffile</code>	<i>filename</i>	configuration file (on Linux, default <code>\$HOME/.arc/client.conf</code>)
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page

Arguments:

filename ... file(s) describing the job(s) to be submitted

A typical Grid job submission looks like:

```
arcsub myjob.jsdl
```

Here `myjob.jsdl` is a file containing job description. Note that in this example `-f` is omitted since the job description file is the last item on the command line.

Please remember that you must have **valid credentials** (see Section 2.1) and be authorised at at least one Grid site.

The job must be described using one of the supported job description languages. The description can be entered either as an argument on the command line, or can be read from a file, as in the example above. Several jobs can be requested at the same time by giving more than one file name as an argument, or by repeating the `-f` or `-e` options. It is possible to mix `-e` and `-f` options in the same `arcsub` command.

A simple "Hello World" job description `myjob.jsdl` using the standard JSDL language is shown below.

```
<?xml version="1.0" encoding="UTF-8"?>
<JobDefinition
  xmlns="http://schemas.ggf.org/jSDL/2005/11/jSDL"
  xmlns:posix="http://schemas.ggf.org/jSDL/2005/11/jSDL-posix">
  <JobDescription>
    <JobIdentification>
      <JobName>Hello World job</JobName>
```

```

</JobIdentification>
<Application>
  <posix:POSIXApplication>
    <posix:Executable>/bin/echo</posix:Executable>
    <posix:Argument>'Hello World'</posix:Argument>
    <posix:Output>out.txt</posix:Output>
    <posix:Error>err.txt</posix:Error>
  </posix:POSIXApplication>
</Application>
</JobDescription>
</JobDefinition>

```

If a job is successfully submitted, a **job identifier** (*job ID*) is printed to standard output.

The job ID uniquely identifies the job. Job IDs differ strongly for different computing service flavours, but basically they have a form of a URL. You should use job ID as a handle to refer to the job when doing other job manipulations, such as querying job status (`arcstat`), killing it (`arckill`), re-submitting (`arcresub`), or retrieving the result (`arcget`).

Usually job ID is a valid URL for the job session directory. You can almost always use it to access the files related to the job, by using data management tools (see Chapter 2.3).

There may be exceptions for some computing service flavours like CREAM which do not support listing job session directory.

The `-c` option can be used to manually select known computing sites, for example:

```
arcsub -c alias1 -c alias2 job.xrsl
```

This will submit a job to either of the sites known by the aliases `alias1` or `alias2`. To submit a job to any site except `badsite`, use `-` sign in front of the name:

```
arcsub -c -badsite job.xrsl
```

See below for description of different kinds of designators which can be used with `-c` and `-g` options.

The `arcsub` command locates the available sites by querying the information system (unless option `-c` is used, in which case only the listed sites are queried). Default index services for the information system are specified in the configuration template distributed with the middleware, and can be overwritten both in the user's configuration (see Section 4) and from the command line using option `-g`. Different Grid flavours may use different notation for such index services.

The designators for `-c` and `-g` are either alias names or meta-URLs of the format `interface:URL`, where `interface:` is optional, specifying the computing service flavour (and the corresponding plugin) to be used when handling the URL. Currently possible flavours are:

ARC0 Legacy ARC execution and index services (requires the `nordugrid-arc-plugins-globus` package to be installed)

ARC1 Web service ARC execution service derived from OGSA-BES standard

CREAM CREAM BES-compliant execution service

BES Generic BES plugin consistent with the OGSA-BES standard

EMIES Web service following EMI Execution Service specifications

Here are examples of full designators for ARC legacy index services:

```
ARC0:ldap://ce.ng.eu:2135/nordugrid-cluster-name=ce.ng.eu,Mds-Vo-name=local,o=grid
ARC0:ldap://index.ng.org:2135/mds-vo-name=sweden,O=grid
```

and for CREAM

```
CREAM:ldap://cream.glite.org:2170/o=grid
```

In case *interface:* part is missing every communication protocol/interface corresponding to supported flavours and matching URL will be tried. Because `arcsub` supports multiple Grid flavours and this number is continuously increasing it is strongly advisable not to skip *interface:* part. Example of such designator is

```
ldap://ce.ng.eu:2135/nordugrid-cluster-name=ce.ng.eu,Mds-Vo-name=local,o=grid
```

For convenience it is possible to shorten designator even more by skipping protocol and path parts of URL. So designator may be as simple as hostname of service to be contacted. Here is an example of such shorthand designator for index server

```
index.ng.org/mds-vo-name=sweden
```

and those suitable both for `-c` and `-g` options:

```
cream.glite.org
ce.ng.eu
```

If such short designators are used then rest of the URL is automatically generated according to the flavour which is currently tried. For example in the case of ARC1, `https` communication protocol is assumed.

If You are using some services frequently enough it is recommended to use aliases for these URLs. Aliases are specified in the configuration file (see Section 4).

In order to keep track of submitted jobs, ARC client stores information in a dedicated file, on Linux platforms by default located in `$HOME/.arc/jobs.xml`. It is sometimes convenient to keep separate lists (e.g., for different kinds of jobs), to be used later with e.g. `arcstat`. This is achieved with the help of the `-j` command line option.

The ARC client transforms input job description into a format that can be understood by the Grid services to which it is being submitted. By specifying the `--dumpdescription` option, such transformed description is written to the standard output instead of being submitted for execution.

Possible broker values for the `arcsub` command line option `-b` are:

- `Random` – ranks targets randomly (default)
- `FastestQueue` – ranks targets according to their queue length
- `Benchmark[:name]` – ranks targets according to a given benchmark, as specified by the `name`. If no benchmark is specified, `CINT2000 *` is used
- `Data` – ranks targets according the amount of megabytes of the requested input files that are already in the computing resources cache.
- `Python:<module>.<class>[:arguments]` – ranks targets using any user-supplied custom Python broker module, optionally with broker arguments. Such module can reside anywhere in user's `PYTHONPATH`

*<http://www.spec.org/cpu2000/CINT2000/>

- `<otherbroker>[:arguments]` - ranks targets using any user-supplied custom C++ broker plugin, optionally with broker arguments. Default location for broker plugins on Linux systems is `/usr/lib/arc` (may depend on the operating system), or the one specified by the `ARC_PLUGIN_PATH`.

To write a custom broker in C++ one has to write a new specialization of the `Broker` base class and implement the `SortTargets` method in the new class. The class should be compiled as a loadable module that has the proper ARC plugin descriptor for the new broker. For example, to build a broker plugin “MyBroker” one executes:

```
g++ -I /arc-install/include \
-L /arc-install/lib \
'pkg-config --cflags glibmm-2.4 libxml-2.0' \
-o libacmybroker.so -shared MyBroker.cpp
```

For more details, refer to *libarclib* documentation [4].

It often happens that some sites that `arcsb` has to contact are slow to answer, or are down altogether. This will not prevent you from submitting a job, but will slow down the submission. To speed it up, you may want to specify a shorter timeout (default is 20 seconds) with the `-t` option:

```
arcsb -t 5 myjob.jsdl
```

Default value for the timeout can be set in the user’s configuration file (see Section 4).

If you would like to get diagnostics of the process of resource discovery and requirements matching, a very useful option is `-d`. The following command:

```
arcsb -d VERBOSE myjob.xrsl
```

will print out the steps taken by the ARC client to find the best cluster satisfying your job requirements. Possible diagnostics levels, in the order of increasing verbosity, are: `FATAL`, `ERROR`, `WARNING`, `INFO`, `VERBOSE` and `DEBUG`. Default is `WARNING`, and it can be set to another value in the user’s configuration file.

Default configuration file on Linux platforms is `$HOME/.arc/client.conf`. However, a user can choose any other pre-defined configuration through option `-z`.

Command line option `-v` prints out version of the installed ARC client package, and option `-h` provides a short help text.

For certain advanced computational jobs which may need to communicate their status to some external services, there may be a need for knowing internal job ID. For jobs accepted by ARC computational services this information can be found in the local (for the job executable) environment variable `GRID.GLOBAL.JOBID`. One needs to take into account that this ID is probably different from the one provided by `arcsb`. An example is an ID provided by the A-REX computing service. That service provides OGSA-BES compatible interface for job management and the ID contains an XML document compliant with OGSA-BES specifications.

2.2.2 arcstat

arcstat [options] [job ...]

Options:

<code>-a, --all</code>		all jobs
<code>-j, --joblist</code>	<i>filename</i>	the file storing information about active jobs (on Linux, default <code>\$HOME/.arc/jobs.xml</code>)
<code>-i, --jobids-from-file</code>	<i>filename</i>	file containing a list of job IDs

<code>-c, --cluster</code>	<code>[-]name</code>	explicitly select or reject a specific site
<code>-s, --status</code>	<code>statusstr</code>	only select jobs whose status is <i>statusstr</i>
<code>-l, --long</code>		long format (extended information)
<code>-S, --sort</code>	<code>criterion</code>	sort jobs according to job ID (criterion <i>jobid</i>), submission time (<i>submissiontime</i>) or job name (<i>jobname</i>)
<code>-R, --rsort</code>	<code>criterion</code>	reverse sorting of jobs according to job ID, submission time or job name
<code>-p, --print-jobids</code>		instead of the status only the IDs of the selected jobs will be printed
<code>-P, --listplugins</code>		list the available plugins
<code>-t, --timeout</code>	<code>time</code>	timeout for queries (default 20 sec)
<code>-d, --debug</code>	<code>verbosity</code>	verbosity level is one of FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG
<code>-z, --conffile</code>	<code>filename</code>	configuration file (on Linux, default <code>\$HOME/.arc/client.conf</code>)
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page
Arguments:		
<code>job ...</code>		list of job IDs and/or job names

The `arcstat` command returns the status of jobs in the Grid, and is typically issued with a job ID (as returned by `arcsub`) as an argument. It is also possible to use job name instead of ID, but if several jobs have identical names, information will be collected about all of them. More than one job ID and/or name can be given.

When several of the `-c <cluster>`, `-i <jobidfile>` and `[job...]` command line options are specified, the command returns information about all jobs listed on the command line plus all jobs on the specified clusters plus all jobs from the specified `jobidfile`.

However the `-c <cluster>` and `-s <status>` options will filter the jobs selected by the above mentioned options, or if none of those are specified, then these will filter all the jobs.

For example, `arcstat -s Finished -c mycluster <jobid>` will return information about the finished jobs on `mycluster` plus about `<jobid>` but only if it is finished. Or, `arcstat -i jobidfile -c -mycluster` will return information about jobs which are in the `jobidfile` but not on `mycluster`.

If the `-l` option is given, extended information is printed.

Jobs can be sorted according to the job ID, submission time or job name, either in normal or reverse order. By using the `--sort` or `--rsort` option followed by the desired ordering (`jobid`, `submissiontime` or `jobname`, respectively), jobs will be sorted in normal or reverse order. Note that the options `--sort` and `--rsort` cannot be used at the same time.

Options `-a`, `-c`, `-s` and `-j` do not use job ID or names. By specifying the `-a` option, the status of all active jobs will be shown. If the `-j` option is used, the list of jobs is read from a file with the specified filename, instead of the default one (`$HOME/.arc/jobs.xml`) (on Linux).

Option `-c` accepts arguments in the `GRID:URL` notation, as explained in the description of `arcsub`, or their aliases as specified in the configuration file.

Different sites may report different job states, depending on the installed grid middleware version. Typical values can be e.g. "Accepted", "Preparing", "Running", "Finished" or "Deleted". Please refer to the respective middleware documentation for job state model description.

Command line option `-s` will instruct the client to display information of only those jobs which status matches the instruction. This option must be given together with either `-a` or `-c` ones, e.g.:

```
arcstat -as Finished
```

Other command line options are identical to those of `arcsub`.

2.2.3 arccat

It is often useful to monitor the job progress by checking what it prints on the standard output or error. The command `arccat` assists here, extracting the corresponding information from the execution cluster and dumping it on the user's screen. It works both for running tasks and for the finished ones. This allows a user to check the output of the finished task without actually retrieving it.

arccat [options] [job ...]

Options:

<code>-a, --all</code>		all jobs
<code>-j, --joblist</code>	<i>filename</i>	the file storing information about active jobs (default <code>/.arc/jobs.xml</code>)
<code>-i, --jobids-from-file</code>	<i>filename</i>	file containing a list of job IDs
<code>-c, --cluster</code>	<code>[-] url</code>	explicitly select or reject (-) a specific site
<code>-s, --status</code>	<i>statusstr</i>	only select jobs whose status is <i>statusstr</i>
<code>-o, --stdout</code>		show the stdout of the job (default)
<code>-e, --stderr</code>		show the stderr of the job
<code>-l, --joblog</code>		show the CE's error log of the job
<code>-P, --listplugins</code>		list the available plugins
<code>-t, --timeout</code>	<i>time</i>	timeout for queries (default 20 sec)
<code>-d, --debug</code>	<i>verbosity</i>	verbosity level is one of FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG
<code>-z, --conffile</code>	<i>filename</i>	configuration file (on Linux, default <code>\$HOME/.arc/client.conf</code>)
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page

Arguments:

<code>job ...</code>	list of job IDs and/or job names
----------------------	----------------------------------

The `arccat` command returns the standard output of a job (`-o` option), the standard error (`-e` option) or errors reported by either Grid Manager or A-REX (`-l` option).

Other command line options have the same meaning as in `arcstat`.

When several of the `-c <cluster>`, `-i <jobidfile>` and `[job...]` command line options are specified, the command prints logs of all jobs listed on the command line plus all jobs on the specified clusters plus all jobs from the specified `jobidfile`.

However the `-c <cluster>` and `-s <status>` options will filter the jobs selected by the above mentioned options, or if none of those are specified, then these will filter all the jobs.

For example, `arccat -s Finished -c mycluster <jobid>` will print logs of the finished jobs on

`mycluster` plus of `<jobid>` but only if it is finished. Or, `arccat -i jobidfile -c -mycluster` will print logs of jobs which are in the `jobidfile` but not on `mycluster`.

2.2.4 arcget

To retrieve the results of a finished job, the `arcget` command should be used. It will transfer the files specified for download in job description to the user's computer.

arcget [options] [job ...]

Options:

<code>-a, --all</code>		all jobs
<code>-j, --joblist</code>	<i>filename</i>	the file storing information about active jobs (default <code>/.arc/jobs.xml</code>)
<code>-i, --jobids-from-file</code>	<i>filename</i>	file containing a list of job IDs
<code>-c, --cluster</code>	<code>[-]name</code>	explicitly select or reject a specific site (cluster)
<code>-s, --status</code>	<i>statusstr</i>	only select jobs whose status is <i>statusstr</i>
<code>-D, --dir</code>	<i>dirname</i>	download path (the job directory will be created in that location)
<code>-J, --usejobname</code>	<i>dirname</i>	use the job name instead of the short ID as the job directory name
<code>-k, --keep</code>		keep files in the Grid (do not clean)
<code>-f, --force</code>		force download (overwrite existing job directory)
<code>-P, --listplugins</code>		list the available plugins
<code>-t, --timeout</code>	<i>time</i>	timeout for queries (default 20 sec)
<code>-d, --debug</code>	<i>verbosity</i>	verbosity level is one of FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG
<code>-z, --conffile</code>	<i>filename</i>	configuration file (on Linux, default <code>\$HOME/.arc/client.conf</code>)
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page

Arguments:

<code>job ...</code>	list of job IDs and/or job names
----------------------	----------------------------------

Only the results of jobs that have finished can be downloaded. Just like in `arcstat` and `arccat` cases, the job can be referred to either by the job ID that was returned by `arcsub` at submission time, or by its name, if the job description contained a job name attribute.

By default, the job is downloaded into a newly created directory in the current path, with the name typically being a large random number. In order to instruct `arcget` to use another path, use option `-D` (note the capital "D"), e.g.

```
arcget -D /tmp/myjobs "Test job nr 1"
```

After downloading, your jobs will be erased from the execution site! Use command line option `-k` to keep finished jobs in the Grid.

Other command line options are identical to those of e.g. `arcstat`.

When several of the `-c <cluster>`, `-i <jobidfile>` and `[job...]` command line options are specified, the command retrieves all jobs listed on the command line plus all jobs on the specified clusters plus all jobs from the specified `jobidfile`.

However the `-c <cluster>` and `-s <status>` options will filter the jobs selected by the above mentioned options, or if none of those are specified, then these will filter all the jobs.

For example, `arcget -s Finished -c mycluster <jobid>` will retrieve the finished jobs on `mycluster` plus `<jobid>` but only if it is finished. Or, `arcget -i jobidfile -c -mycluster` will retrieve jobs which are in the `jobidfile` but not on `mycluster`.

2.2.5 arcsync

It is advised to start every grid session by running `arcsync`, especially when changing workstations. The reason is that your job submission history is cached on your machine, and if you are using ARC client installations on different machines, your local lists of submitted jobs will be different. To synchronise these lists with the information in the Information System, use the `arcsync` command.

arcsync [options]

Options:

<code>-c, --cluster</code>	<code>[-] name</code>	explicitly select or reject a specific site
<code>-g, --index</code>	<code>url</code>	explicitly select or reject (-) a specific index server
<code>-j, --joblist</code>	<code>filename</code>	the file storing information about active jobs (default <code>/.arc/jobs.xml</code>)
<code>-f, --force</code>		don't ask for confirmation
<code>-T, --truncate</code>		truncate the job list before synchronising
<code>-P, --listplugins</code>		list the available plugins
<code>-t, --timeout</code>	<code>seconds</code>	timeout for network communication, in seconds (default 20)
<code>-d, --debug</code>	<code>verbosity</code>	verbosity level, FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG - default WARNING
<code>-z, --conffile</code>	<code>filename</code>	configuration file (on Linux, default <code>\$HOME/.arc/client.conf</code>)
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page

The ARC client keeps a local list of jobs in the user's home directory. If this file is lost, corrupt, or the user wants to recreate the file on a different workstation, the `arcsync` command will recreate this file from the information available in the Information System.

Since the information about a job retrieved from a cluster 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 `-f` option disables this warning.

If the job list is not empty when invoking synchronisation, the old jobs will be merged with the new jobs, unless the `-T` option is given (note the capital "T"), in which case the job list will first be truncated and then the new jobs will be added.

2.2.6 arcinfo

The `arcinfo` command is used to obtain status information about clusters on the Grid.

arcinfo [options]

Options:

<code>-c, --cluster</code>	<code>[-]name</code>	explicitly select or reject a specific site
<code>-g, --index</code>	<code>url</code>	explicitly select or reject (-) a specific index server
<code>-l, --long</code>		long format (extended information)
<code>-P, --listplugins</code>		list the available plugins
<code>-t, --timeout</code>	<code>seconds</code>	timeout for network communication, in seconds (default 20)
<code>-d, --debug</code>	<code>verbosity</code>	verbosity level, FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG - default WARNING
<code>-z, --conffile</code>	<code>filename</code>	configuration file (on Linux, default \$HOME/.arc/client.conf)
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page

The `arcinfo` command is used to obtain information about clusters and queues (*targets*) available on the Grid. Either the `--cluster` or `--index` flag should be used to specify the target(s) which should be queried for information. Both of these flags take a service endpoint as argument. See `arcsub` and the configuration notes in Section 4 for description of these.

Detailed information about queried computing services can be obtained by specifying the `--long` flag.

When specifying the `--index` flag, the information about the computing services registered at the index server will be queried, rather than the status of the index server itself.

2.2.7 arckill

It happens that a user may wish to cancel a job. This is done by using the `arckill` command. A job can be killed almost at any stage of processing through the Grid.

arckill [options] [job ...]

Options:

<code>-a, --all</code>		all jobs
<code>-j, --joblist</code>	<code>filename</code>	the file storing information about active jobs (default /.arc/jobs.xml)
<code>-i, --jobids-from-file</code>	<code>filename</code>	file containing a list of job IDs
<code>-c, --cluster</code>	<code>[-]url</code>	explicitly select or reject (-) a specific site
<code>-s, --status</code>	<code>statusstr</code>	only select jobs whose status is <i>statusstr</i>
<code>-k, --keep</code>		keep files in the Grid (do not clean)
<code>-P, --listplugins</code>		list the available plugins
<code>-t, --timeout</code>	<code>time</code>	timeout for queries (default 20 sec)
<code>-d, --debug</code>	<code>verbosity</code>	verbosity level is one of FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG

<code>-z, --conffile</code>	<i>filename</i>	configuration file (on Linux, default \$HOME/.arc/client.conf)
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page
Arguments:		
<code>job ...</code>		list of job IDs and/or job names

If a job is killed, its traces are being cleaned from the Grid. If you wish to keep the killed job in the system, e.g. for a post-mortem analysis, use the `-k` option.

Job cancellation is an asynchronous process, such that it may take a few minutes before the job is actually cancelled.

Command line options have the same meaning as the corresponding ones of `arcstat` and others.

When several of the `-c <cluster>`, `-i <jobidfile>` and `[job...]` command line options are specified, the command kills all jobs listed on the command line plus all jobs on the specified clusters plus all jobs from the specified `jobidfile`.

However the `-c <cluster>` and `-s <status>` options will filter the jobs selected by the above mentioned options, or if none of those are specified, then these will filter all the jobs.

For example, `arckill -s INLRMS:R -c mycluster <jobid>` will kill the running jobs on `mycluster` plus `<jobid>` but only if it is running. Or, `arckill -i jobidfile -c -mycluster` will kill all jobs which are in the `jobidfile` but not on `mycluster`.

2.2.8 arcclean

If a job fails or gets killed with `-k` option, or when you are not willing to retrieve the results for some reasons, a good practice for users is not to wait for the system to clean up the job leftovers, but to use `arcclean` to release the disk space and to remove the job ID from the list of submitted jobs and from the Information System.

arcclean [options] [job ...]

Options:		
<code>-a, --all</code>		all jobs
<code>-j, --joblist</code>	<i>filename</i>	the file storing information about active jobs (default /.arc/jobs.xml)
<code>-i, --jobids-from-file</code>	<i>filename</i>	file containing a list of job IDs
<code>-c, --cluster</code>	<code>[-] name</code>	explicitly select or reject a specific site (cluster)
<code>-s, --status</code>	<i>statusstr</i>	only select jobs whose status is <i>statusstr</i>
<code>-f, --force</code>		removes the job ID from the local list even if the job is not found on the Grid
<code>-P, --listplugins</code>		list the available plugins
<code>-t, --timeout</code>	<i>time</i>	timeout for queries (default 20 sec)
<code>-d, --debug</code>	<i>verbosity</i>	verbosity level is one of FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG

<code>-z, --conffile</code>	<i>filename</i>	configuration file (on Linux, default <code>\$HOME/.arc/client.conf</code>)
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page
Arguments:		
<code>job ...</code>		list of job IDs and/or job names

Only jobs that have finished or were cancelled can be cleaned.

It happens ever so often that the job is cleaned by the system, or is otherwise unreachable, and yet your local job list file still has it listed. Use `-s` option with value `Undefined` to remove such stale job information from the local list. Note that specifying `-a` and `-f` options together also removes such stale job information, while also removing finished and cancelled jobs.

Other command line options have the same meaning as the corresponding ones of `arcstat` and others.

When several of the `-c <cluster>`, `-i <jobidfile>` and `[job...]` command line options are specified, the command cleans all jobs listed on the command line plus all jobs on the specified clusters plus all jobs from the specified `jobidfile`.

However the `-c <cluster>` and `-s <status>` options will filter the jobs selected by the above mentioned options, or if none of those are specified, then these will filter all the jobs.

For example, `arcclean -s FAILED -c mycluster <jobid>` will clean the failed jobs on `mycluster` plus `<jobid>` but only if it is failed. Or, `arcclean -i jobidfile -c -mycluster` will clean all jobs which are in the `jobidfile` but not on `mycluster`.

2.2.9 arc renew

Quite often, the user proxy expires while the job is still running (or waiting in a queue). In case such job has to upload output files to a Grid location (a storage element), it will fail. By using the `arc renew` command, users can upload a new proxy to the job. This can be done while a job is still running, thus preventing it from failing

If a job has failed in file upload due to expired proxy, `arc renew` can be issued within 24 hours (or whatever is the expiration time set by the site) after the job end, which must be followed by `arc resume`. The Grid Manager or A-REX will then attempt to finalize the job by uploading the output files to the desired location.

arc renew [options] [job ...]

Options:

<code>-a, --all</code>		all jobs
<code>-j, --joblist</code>	<i>filename</i>	the file storing information about active jobs (default <code>/.arc/jobs.xml</code>)
<code>-i, --jobids-from-file</code>	<i>filename</i>	file containing a list of job IDs
<code>-c, --cluster</code>	<code>[-] name</code>	explicitly select or reject a specific site (cluster)
<code>-s, --status</code>	<i>statusstr</i>	only select jobs whose status is <i>statusstr</i>
<code>-P, --listplugins</code>		list the available plugins
<code>-t, --timeout</code>	<i>time</i>	timeout for queries (default 20 sec)
<code>-d, --debug</code>	<i>verbosity</i>	verbosity level is one of FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG

<code>-z, --conffile</code>	<i>filename</i>	configuration file (on Linux, default \$HOME/.arc/client.conf)
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page
Arguments:		
<code>job ...</code>		list of job IDs and/or job names

Prior to using `arc renew`, be sure to actually create the new proxy by running `arc proxy`!

Command line options have the same meaning as the corresponding ones of `arc stat` and others.

When several of the `-c <cluster>`, `-i <jobidfile>` and `[job...]` command line options are specified, the command renews proxies of all jobs listed on the command line plus all jobs on the specified clusters plus all jobs from the specified `jobidfile`.

However the `-c -<cluster>` and `-s <status>` options will filter the jobs selected by the above mentioned options, or if none of those are specified, then these will filter all the jobs.

For example, `arc renew -s FAILED -c mycluster <jobid>` will renew proxies of the failed jobs on `mycluster` plus `<jobid>` but only if it is failed. Or, `arc renew -i jobidfile -c -mycluster` will renew proxies of all jobs which are in the `jobidfile` but not on `mycluster`.

2.2.10 arc resume

In some cases a user may want to restart a failed job, for example, when input files become available, or the storage element for the output files came back online, or when a proxy is renewed with `arc renew`. This can be done using the `arc resume` command.

Make sure your proxy is still valid, or when uncertain, run `arc proxy` followed by `arc renew` before `arc resume`. The job will be resumed from the state where it has failed.

arc resume [options] [job ...]

<code>-a, --all</code>		all jobs
<code>-j, --joblist</code>	<i>filename</i>	the file storing information about active jobs (default /.arc/jobs.xml)
<code>-i, --jobids-from-file</code>	<i>filename</i>	file containing a list of job IDs
<code>-c, --cluster</code>	<code>[-]name</code>	explicitly select or reject a specific site (cluster)
<code>-s, --status</code>	<i>statusstr</i>	only select jobs whose status is <i>statusstr</i>
<code>-P, --listplugins</code>		list the available plugins
<code>-t, --timeout</code>	<i>time</i>	timeout for queries (default 20 sec)
<code>-d, --debug</code>	<i>verbosity</i>	verbosity level is one of FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG
<code>-z, --conffile</code>	<i>filename</i>	configuration file (on Linux, default \$HOME/.arc/client.conf)
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page

Arguments:

`job ...` list of job IDs and/or job names

Command line options have the same meaning as the corresponding ones of `arcstat` and others.

When several of the `-c <cluster>`, `-i <jobidfile>` and `[job...]` command line options are specified, the command resumes all jobs listed on the command line plus all jobs on the specified clusters plus all jobs from the specified `jobidfile`.

However the `-c <cluster>` and `-s <status>` options will filter the jobs selected by the above mentioned options, or if none of those are specified, then these will filter all the jobs.

For example, `arcresume -s FAILED -c mycluster <jobid>` will resume the failed jobs on `mycluster` plus `<jobid>` but only if it is failed. Or, `arcresume -i jobidfile -c -mycluster` will resume all jobs which are in the `jobidfile` but not on `mycluster`.

2.2.11 arcresub

Quite often it happens that a user would like to re-submit a job, but has difficulties recovering the original job description file (e.g. the xRSL file). This happens when job description files are created by scripts on-fly, and matching of job description to the job ID is not straightforward. The utility called `arcresub` helps in such situations, allowing users to resubmit jobs.

arcresub [options] [job ...]

Options:

<code>-a, --all</code>		all jobs
<code>-g, --index</code>	<i>url</i>	explicitly select or reject (-) a specific index server
<code>-j, --joblist</code>	<i>filename</i>	the file storing information about active jobs (default <code>/.arc/jobs.xml</code>)
<code>-i, --jobids-from-file</code>	<i>filename</i>	file containing a list of job IDs
<code>-o, --jobids-to-file</code>	<i>filename</i>	the IDs of the submitted jobs will be appended to this file
<code>-c, --cluster</code>	<code>[-]name</code>	explicitly select or reject a specific source site
<code>-q, --qluster</code>	<code>[-]name</code>	explicitly select or reject a specific site as re-submission target
<code>-m, --same</code>		re-submit to the same site
<code>-M, --not-same</code>		do not resubmit to the same cluster
<code>-s, --status</code>	<i>statusstr</i>	only select jobs whose status is <i>statusstr</i>
<code>-x, --dumpdescription</code>		do not submit – dump transformed job description to stdout
<code>-k, --keep</code>		keep files in the Grid (do not clean)
<code>-b, --broker</code>	<i>string</i>	select broker method (default is Random)
<code>-P, --listplugins</code>		list the available plugins
<code>-t, --timeout</code>	<i>time</i>	timeout for queries (default 20 sec)
<code>-d, --debug</code>	<i>verbosity</i>	verbosity level is one of FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG
<code>-z, --conffile</code>	<i>filename</i>	configuration file (on Linux, default <code>\$HOME/.arc/client.conf</code>)

```

-v, --version          print version information
-h, --help            print help page
Arguments:
job ...               list of job IDs and/or job names

```

More than one job ID and/or job name can be given. If several jobs were submitted with the same job name all those jobs will be resubmitted.

If the job description of a job to be resubmitted, contained any local input files, checksums of these was calculated and stored in the job list, and those will be used to check whether the files has changed. If local input files has changed the job will not be resubmitted.

In case the job description is not found in the job list, an attempt will be made to retrieve it from the cluster holding the original job. This however may fail, since both the submission client and the cluster can have made modifications to the job description.

Upon resubmission the job will receive a new job ID, and the old job ID will be stored in the local job list file, enabling future back tracing of the resubmitted job.

Upon resubmission the job will receive a new job ID. The old job ID will be kept in the local job list file, enabling future back tracing of the resubmitted job.

Regarding command line options, `arcresub` behaves much like `arcsub`, except that `-c` in this case indicates not the submission target site, but on the contrary, the **site from which the jobs will be resubmitted**. Submission target site is specified with option `-q`. If you wish to re-submit each job to the same site, use option `-m`.

If the original job was successfully killed, its traces will be removed from the execution site, unless the `-k` option is specified.

When several of the `-c <cluster>`, `-i <jobidfile>` and `[job...]` command line options are specified, the command resubmits all jobs listed on the command line plus all jobs on the specified clusters plus all jobs from the specified `jobidfile`.

However the `-c <cluster>` and `-s <status>` options will filter the jobs selected by the above mentioned options, or if none of those are specified, then these will filter all the jobs.

For example, `arcresub -s FAILED -c mycluster <jobid>` will resubmit the running jobs on `mycluster` plus `<jobid>` but only if it is failed. Or, `arcresub -i jobidfile -c -mycluster` will resubmit all jobs which are in the `jobidfile` but not on `mycluster`.

2.2.12 arcigrate

Quite often jobs end up stuck in long queues, and users wish to migrate them to a better resource. Command `arcigrate` is triggering this migration. It applies only to jobs submitted to the web service interface of A-REX, as other Grid execution services do not support this functionality.

`arcigrate [options] [job ...]`

```

Options:
-c, --cluster          [-]name  explicitly select or reject a specific site (cluster)
-q, --qluster         [-]name  explicitly select or reject a specific site as re-
                               submission target
-f, --force           force migration, ignoring kill failure
-k, --keep            keep the files on the server (do not clean)
-g, --index          url      explicitly select or reject (-) a specific index server

```

<code>-a, --all</code>		all jobs
<code>-b, --broker</code>	<i>string</i>	select broker method (default is Random)
<code>-o, --jobids-to-file</code>	<i>filename</i>	the IDs of the submitted jobs will be appended to this file
<code>-i, --jobids-from-file</code>	<i>filename</i>	file containing a list of job IDs
<code>-j, --joblist</code>	<i>filename</i>	the file storing information about active jobs (default <code>/.arc/jobs.xml</code>)
<code>-z, --conffile</code>	<i>filename</i>	configuration file (on Linux, default <code>\$HOME/.arc/client.conf</code>)
<code>-t, --timeout</code>	<i>time</i>	timeout for queries (default 20 sec)
<code>-P, --listplugins</code>		list the available plugins
<code>-d, --debug</code>	<i>verbosity</i>	verbosity level is one of FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page
Arguments:		
<code>job ...</code>		list of job IDs and/or job names

Currently only jobs having the status “Queuing” can be migrated

Command line options `-c` and `-q` are interpreted in the same way as in `arcresub`, namely, `-c` indicates “from” and `-q` – “to” which site the job will be migrated.

If the job(s) is successfully migrated, a new job ID(s) is printed out. This job ID uniquely identifies the job while it is being executed.

When several of the `-c <cluster>`, `-i <jobidfile>` and `[job...]` command line options are specified, the command migrates all jobs listed on the command line plus all jobs on the specified clusters plus all jobs from the specified `jobidfile`.

However the `-c -<cluster>` and `-s <status>` options will filter the jobs selected by the above mentioned options, or if none of those are specified, then these will filter all the jobs.

For example, `arcmigrate -s FAILED -c mycluster <jobid>` will migrate the running jobs on `mycluster` plus `<jobid>` but only if it is failed. Or, `arcmigrate -i jobidfile -c -mycluster` will resubmit all jobs which are in the `jobidfile` but not on `mycluster`.

2.3 Data manipulation

ARC provides basic data management tools to copy, create, list and remove files and directories to, from and between Grid storage elements and index services.

2.3.1 arcls

`arcls` is a simple utility that allows to list contents and view some attributes of objects of a specified (by a URL) remote directory.

arcls [options] <URL>

Options:	
<code>-l, --long</code>	detailed listing
<code>-L, --locations</code>	detailed listing including URLs from which files can be downloaded
<code>-m, --metadata</code>	display all available metadata
<code>-r, --recursive</code>	<i>recursion_level</i> operate recursively (if possible) up to specified level (0 - no recursion)
<code>-n, --nolist</code>	show only description of requested object, do not list content of directories (like <code>ls -d</code>).
<code>-c, --checkaccess</code>	check readability of object. Retrieving and showing information about object is suppressed.
<code>-t, --timeout</code>	<i>seconds</i> timeout for network communication, in seconds (default 20)
<code>-P, --plugins</code>	list the available plugins (protocols supported)
<code>-d, --debug</code>	<i>verbosity</i> verbosity level, FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG - default WARNING
<code>-z, --conffile</code>	<i>filename</i> configuration file (on Linux, default <code>\$HOME/.arc/client.conf</code>)
<code>-v, --version</code>	print version information
<code>-h, --help</code>	print help page
Arguments:	
URL	file or directory URL

This tool is very convenient not only because it allows to list files at a Storage Element or records in an indexing service, but also because it can give a quick overview of a job's working directory, which is explicitly given by job ID.

Usage examples can be as follows:

```
arccls -L rls://rls.nordugrid.org:38203/logical_file_name
arccls -l gsiftp://lscf.nbi.dk:2811/jobs/1323842831451666535
arccls srm://grid.uio.no:8446/srm/managerv2?SFN=/johndoe/log2
```

Examples of URLs accepted by this tool can be found in Section 3, though `arccls` won't be able to list a directory at an HTTP server, as they normally do not return directory listings.

2.3.2 arccp

`arccp` is a powerful tool to copy files over the Grid. It is a part of the A-REX, but can be used by the User Interface as well.

arccp [options] <source> <destination>

Options:	
<code>-p, --passive</code>	use passive transfer (does not work if secure is on, default if secure is not requested)
<code>-n, --nopassive</code>	do not try to force passive transfer

<code>-f, --force</code>		if the destination is an indexing service and not the same as the source and the destination is already registered, then the copy is normally not done. However, if this option is specified the source is assumed to be a replica of the destination created in an uncontrolled way and the copy is done like in case of replication. Using this option also skips validation of completed transfers.
<code>-i, --indicate</code>		show progress indicator. If the transfer time is short then there may be no indicator.
<code>-T, --nottransfer</code>		do not transfer file, just register it - destination must be non-existing meta-url
<code>-u, --secure</code>		use secure transfer (insecure by default)
<code>-y, --cache</code>	<i>path</i>	path to local cache (use to put file into cache). See [6] for information on caching.
<code>-r, --recursive</code>	<i>recursion_level</i>	operate recursively (if possible) up to specified level (0 - no recursion)
<code>-R, --retries</code>	<i>number</i>	how many times to retry transfer of every file before failing
<code>-L, --location</code>	<i>URL</i>	physical file to write to when destination is an indexing service. Must be specified for indexing services which do not automatically generate physical locations. Can be specified multiple times - locations will be tried in order until one succeeds.
<code>-t, --timeout</code>	<i>seconds</i>	timeout for network communication, in seconds (default 20)
<code>-P, --plugins</code>		list the available plugins (protocols supported)
<code>-d, --debug</code>	<i>verbosity</i>	verbosity level, FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG - default WARNING
<code>-z, --conffile</code>	<i>filename</i>	configuration file (on Linux, default \$HOME/.arc/client.conf)
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page
Arguments:		
<code>source</code>		source URL
<code>destination</code>		destination URL

This command transfers contents of a file between 2 end-points. End-points are represented by URLs or meta-URLs or local file paths. For supported endpoints please refer to Section 3.

`arccp` can perform multi-stream transfers if `threads` URL option is specified and server supports it.

Source can end with `"/`. In that case, the set of files under source will be copied into destination and destination must also end with `"/`. Destination will be created if it does not exist. If copying deeper than one level is required then `-r` must be used. If destination alone ends with `"/`, it is extended with the part of source after last `"/`, thus allowing users to skip the destination file or directory name if it is meant to be identical to the source.

Usage examples of `arccp` are:

```
arccp gsiftp://lscf.nbi.dk:2811/jobs/1323842831451666535/job.out job.out
arccp http://www.nordugrid.org/data/somefile gsiftp://hathi.hep.lu.se/data/
arccp gsiftp://pgs02.grid.upjs.sk:2811/jobs/13331297786445657047863/ output/
arccp -L srm://srm.host.org;spacetoken=MYTOKEN/my.file.1 my.file \
    lfc://lfc.host.org/grid/my.file
```

2.3.3 arcrm

The `arcrm` command allows users to erase files and directories at any location specified by a valid URL.

arcrm [options] <URL>

Options:

<code>-f, --force</code>		remove logical file name registration even if not all physical instances were removed
<code>-t, --timeout</code>	<i>seconds</i>	timeout for network communication, in seconds (default 20)
<code>-P, --plugins</code>		list the available plugins (protocols supported)
<code>-d, --debug</code>	<i>verbosity</i>	verbosity level, FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG - default WARNING
<code>-z, --conffile</code>	<i>filename</i>	configuration file (on Linux, default \$HOME/.arc/client.conf)
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page

Arguments:

URL	file or directory URL
-----	-----------------------

A convenient use for `arcrm` is to erase the files in a data indexing catalog (LFC, RLS or such), as it will not only remove the physical instance, but also will clean up the database record.

Here is an `arcrm` example:

```
arcrm lfc://grid.uio.no/grid/atlas/A0D_0947.pool.root
```

2.3.4 arcmkdir

The `arcmkdir` command allows users to create directories, if the protocol of the specified URL supports it.

arcmkdir [options] <URL>

Options:

<code>-p, --parents</code>		make parent directories as needed
<code>-t, --timeout</code>	<i>seconds</i>	timeout for network communication, in seconds (default 20)
<code>-P, --plugins</code>		list the available plugins (protocols supported)
<code>-d, --debug</code>	<i>verbosity</i>	verbosity level, FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG - default WARNING
<code>-z, --conffile</code>	<i>filename</i>	configuration file (on Linux, default \$HOME/.arc/client.conf)
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page

Arguments:

URL	directory to create
-----	---------------------

`arcrm` creates directories on grid storage elements and indexing services. If the parent directory does not exist and `-p` is not specified, then `arcmkdir` will probably fail, but it depends on the protocol. The permissions on the new directory are the default of the server, or if the protocol requires them to be specified then the directory is only readable/writable/searchable by the user (the equivalent of 700 on a file system).

Example:

```
arcmkdir lfc://grid.uio.no/grid/atlas/newdir
```

2.3.5 arcsrmping

The `arcsrmping` command is used to quickly test availability of an SRM service, similarly to the `ping` tool in Unix.

arcsrmping [options] <service>

Options:

<code>-t, --timeout</code>	<i>seconds</i>	timeout for network communication, in seconds (default 20)
<code>-d, --debug</code>	<i>verbosity</i>	verbosity level, FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG - default WARNING
<code>-z, --conffile</code>	<i>filename</i>	configuration file (on Linux, default \$HOME/.arc/client.conf)
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page

Arguments:

<code>service</code>	A URL to an SRM service
----------------------	-------------------------

The `arcsrmping` command is a `ping` client for the SRM service. It sends an SRM ping request to the SRM service and displays the result.

2.3.6 chelonia

`chelonia` is a client tool for accessing the Chelonia storage system. It is capable of creating, removing and listing collections, uploading, downloading and removing files and moving and stating both, using Logical Names (LN). Collections contain files and other collections, creating a hierarchical namespace.

chelonia [options] <method> [arguments]

(ARC 0.9)

Options:

<code>-b</code>	<i>URL</i>	URL of Bartender to connect
<code>-x</code>		print SOAP XML messages
<code>-v</code>		verbose mode
<code>-z</code>	<i>filename</i>	configuration file (on Linux, default \$HOME/.arc/client.conf)
<code>-w</code>		allow to run without the ARC python client libraries (with limited functionality)

Methods:

<code>stat</code>	<i>LN [LN ...]</i>	get detailed information about an entry or several
<code>makeCollection, make, mkdir</code>	<i>LN</i>	create a collection
<code>unmakeCollection, unmake, rmdir</code>	<i>LN</i>	remove an empty collection
<code>list, ls</code>	<i>LN</i>	list the content of a collection
<code>move, mv</code>	<i>source target</i>	move entries within the namespace (both LNs)
<code>putFile, put</code>	<i>source target</i>	upload a file from a <i>source</i> to a <i>target</i> (both specified as LNs)
<code>getFile, get</code>	<i>source [target]</i>	download a file from a <i>source</i> to a <i>target</i>
<code>delFile, del, rm</code>	<i>LN [LN ...]</i>	remove file(s)
<code>modify, mod</code>	<i>string</i>	modify metadata
<code>policy, pol</code>	<i>string</i>	modify access policy rules
<code>unlink</code>	<i>string</i>	remove a link to an entry from a collection without removing the entry itself
<code>credentialsDelegation, cre</code>	<i>string</i>	delegate credentials for using gateway
<code>removeCredentials, rem</code>	<i>string</i>	remove previously delegated credentials
<code>makeMountPoint, makemount</code>	<i>string</i>	create a mount point

Without arguments, each method prints its own help. Detailed explanation of each method is given below.

Examples:

```
chelonia list /
chelonia put orange /
chelonia stat /orange
chelonia get /orange /tmp
chelonia mkdir /fruits
chelonia mkdir /fruits/apple
chelonia mv /orange /fruits
chelonia ls /fruits
chelonia rmdir /fruits/apple
chelonia rmdir /fruits
chelonia rm /fruits/orange
chelonia policy / change ALL +read +addEntry
chelonia modify /pennys-orange set states neededReplicas 2
```

stat

With the `stat` method it is possible to get all the metadata about one or more entry (file, collection, etc.). The entries are specified with their Logical Name (LN).

```
chelonia stat <LN> [<LN> ...]
```

The output contains key-value pairs grouped in sections. The ‘states’ section contains the size and the checksum of a file, the number of needed replicas, and whether a collection is closed or not (a ‘closed’ collection should not be modified anymore, if it gets modified, its state becomes ‘broken’); the ‘entry’ section contains the DN of the owner, the globally unique ID (GUID) of the entry, and the type of the entry (file, collection, etc.); the ‘parents’ section contains the GUID of the parent collection(s) of this entry, and the name of this entry in that collection separated with a ‘/’; the ‘locations’ sections contains the location of the replicas of a file, which contains of the ID (the URL) of the Storage Element, the ID of the replica within the Storage Element, and the state of the replica; the ‘timestamps’ section contains the creation time of the entry; the ‘entries’ section contains the name and GUID of the entries of a collection. Example `stat` of a file:

```
$ chelonia stat /thing
```

```

'/thing': found
  states
    checksumType: md5
    neededReplicas: 3
    size: 6
    checksum: a0186a90393bd4a639a1ce35d8ef85f6
  entry
    owner: /C=HU/O=NIIF CA/OU=GRID/OU=NIIF/CN=Nagy Zsombor
    GUID: 398CBDEA-E282-4735-8DF6-2464CD00BE2D
    type: file
  parents
    0/thing: parent
  locations
    https://localhost:60000/Shepherd D519F687-EF65-4AEA-9766-E6E2D42166C4: alive
  timestamps
    created: 1257351119.3

```

Example stat of a collection:

```

$ chelonia stat /
'/': found
  states
    closed: no
  entry
    owner: /C=HU/O=NIIF CA/OU=GRID/OU=NIIF/CN=Nagy Zsombor
    GUID: 0
    type: collection
  timestamps
    created: 1257351114.37
  entries
    thing: 398CBDEA-E282-4735-8DF6-2464CD00BE2D

```

makeCollection

With the `makeCollection` or `mkdir` method it is possible to create a new empty collection. The requested Logical Name (LN) should be specified.

```
chelonia makeCollection <LN>
```

The parent collection of the requested Logical Name must exist.

Example output of the method:

```

$ chelonia mkdir /newcoll
Creating collection '/newcoll': done

$ chelonia mkdir /nonexistent/newcoll
Creating collection '/nonexistent/newcoll': parent does not exist

```

unmakeCollection

With the `unmakeCollection` or `rmdir` method it is possible to delete an empty collection which is specified by its Logical Name (LN).

```
chelonia unmakeCollection <LN>
```

Example output of the method:

```

$ chelonia rmdir /newcoll
Removing collection '/newcoll': removed

```

```
$ chelonia rmdir /dir
Removing collection '/dir': collection is not empty
```

list

With the `list` or `ls` method it is possible to list the contents of one or more collections which are specified by their Logical Name (LN).

```
chelonia list <LN> [<LN> ...]
```

Example output of the method:

```
$ chelonia list / /newcoll
'/newcoll': collection
  empty.
'/': collection
  thing <file>
  dir <collection>
  newcoll <collection>
```

move

With the `move` or `mv` method it is possible to move a file or collection within the namespace of `chelonia` (including renaming the entry). The source path and the target path should be specified as Logical Names

```
chelonia move <sourceLN> <targetLN>
```

Example output of the method:

```
$ chelonia mv /thing /newcoll/
Moving '/thing' to '/newcoll/': moved

$ chelonia mv /newcoll/thing /newcoll/othername
Moving '/newcoll/thing' to '/newcoll/othername': moved
```

putFile

With the `putFile` or `put` method it is possible to upload a new file into the system creating a new Logical Name (LN). It can upload directories recursively using the `-r` flag. It is also possible to specify the number of needed replicas.

```
chelonia putFile [-r] <source filename> <target LN> [<number of replicas needed>]
```

Example output of the method:

```
$ chelonia put thing /newcoll/
'thing' (6 bytes) uploaded as '/newcoll/thing'.
```

getFile

With the `getFile` or `get` method it is possible to download a file specified with its Logical Name (LN). If the target local path is not given, then the file will be put into the local directory. It can download collections recursively using the `-r` flag.

```
chelonia getFile [-r] <source LN> [<target filename>]
```

Example output of the method:

```
$ chelonia get /newcoll/thing newlocalname
'/newcoll/thing' (6 bytes) downloaded as 'newlocalname'.
```

delFile

With the delFile or rm method it is possible to delete one or more files from the system.

```
chelonia delFile <LN> [<LN> ...]
```

Example output of the method:

```
$ chelonia rm /newcoll/othername
/newcoll/othername: deleted
```

modify

With the modify or mod method it is possible to modify some metadata of an entry.

```
chelonia modify <LN> <changeType> <section> <property> <value>
```

The possible values of 'changeType' are 'set' (sets the property to value within the given section), 'unset' (removes the property from the given section - the 'value' does not matter) and 'add' (sets the property to value within the given section only if it does not exist yet).

To change the number of needed replicas for a file:

```
chelonia modify <LN> set states neededReplicas <number of needed replicas>
```

To close a collection:

```
chelonia modify <LN> set states closed yes
```

A closed collection should not be modified later. If it gets modified its state becomes 'broken'.

To change metadata key-value pairs:

```
chelonia modify <LN> set|unset|add metadata <key> <value>
```

policy

With the policy or pol method it is possible to modify the policy of the entry

```
chelonia policy <LN> <changeType> <identity> <action list>
```

The possible values of 'changeType' are 'set' (sets the action list to the given user overwriting the old one), 'change' (modify the current action list with adding and removing actions) and 'clear' (clear the action list of the given user).

The 'identity' could be currently three things: the DN of a user; the name of a VO (with the syntax: 'VOMS:<VO name>'); or 'ALL' for all users.

The 'action list' is a list of actions prefixed with '+' or '-', e.g. '+read +addEntry -delete'.

These are the actions which can be used for access control:

- *read*: user can get the list of entries in the collection; user can download the file
- *addEntry*: user can add a new entry to the collection;
- *removeEntry*: user can remove any entry from the collection
- *delete*: user can delete the collection if it is empty; user can delete a file
- *modifyPolicy*: user can modify the policy of the file/collection
- *modifyStates*: user can modify some special metadata of the file/collection (close the collection, change the number of needed replica of the file)
- *modifyMetadata*: user can modify the arbitrary metadata section of the file/collection (these are property-value pairs)

There is an implicit default policy: the owner always has all the rights. Checking the 'stat' of new collections:

```
$ chelonia stat /newcoll
'/newcoll': found
states
  closed: no
entry
  owner: /C=HU/O=NIIF CA/OU=GRID/OU=NIIF/CN=Nagy Zsombor
  GUID: 41CBD461-09BE-46FD-8A1B-767C7D427AF9
  type: collection
parents
  0/newcoll: parent
timestamps
  created: 1257435820.26
entries
  thing: A63658B4-2C6E-46A3-8238-7D291F8F81C2
```

shows no policies, but it shows the owner. This collection has no additional policies just the default one: the owner can do anything, noone else can do anything.

Let's set it in a way that all users can read the contents of this collection:

```
$ chelonia policy /newcoll change ALL +read
Setting action list of '/newcoll' for user ALL to +read: set.
$ chelonia stat /newcoll
'/newcoll': found
[...]
policy
  ALL: +read
[...]
```

Then we can set that all the members of the knowarc VO would be able to add entries to this collection:

```
$ chelonia policy /newcoll change VOMS:knowarc +addEntry
Setting action list of '/newcoll' for user VOMS:knowarc to +addEntry: set.
$ chelonia stat /newcoll
'/newcoll': found
[...]
policy
  ALL: +read
  VOMS:knowarc: +addEntry
[...]
```

And for example we can set a specific user to be able to remove entries from this collections:

```
$ chelonia policy /newcoll change \
  "/C=HU/O=NIIF CA/OU=GRID/OU=NIIF/CN=TestUser" +removeEntry
```

```
Setting action list of '/newcoll'
  for user /C=HU/O=NIIF CA/OU=GRID/OU=NIIF/CN=TestUser to +removeEntry: set.
$ chelonia stat '/newcoll'/newcoll': found
[...]
policy
  /C=HU/O=NIIF CA/OU=GRID/OU=NIIF/CN=TestUser: +removeEntry
  ALL: +read
  VOMS:knowarc: +addEntry
[...]
```

unlink

With the `unlink` method it is possible to remove a file or collection just from its parent collection without removing the file or collection itself. This means that the file or collection wouldn't be part of its former parent collection anymore. It would be still possible to access it with its GUID, or if it was linked from another collection too. **NOTE:** If we don't know the GUID or the logical name of any other link, then we cannot access the file or collection anymore.

chelonia unlink <LN>

If there is a file called `'/newcoll/thing'`, it is in the listing of the `'/newcoll'` collection:

```
$ chelonia list /newcoll
'/newcoll': collection
  thing <file>
```

The file is in the entries of the collection:

```
$ chelonia stat /newcoll
'/newcoll': found
  entries
    thing: A63658B4-2C6E-46A3-8238-7D291F8F81C2
[...]
```

It is possible the `'stat'` the file with the Logical Name `'/newcoll/thing'`:

```
jim:~ zsombor$ chelonia stat /newcoll/thing
'/newcoll/thing': found
  states
    checksumType: md5
    neededReplicas: 3
    size: 6
    checksum: a0186a90393bd4a639a1ce35d8ef85f6
[...]
```

Now with the `'unlink'` method it is possible to remove the file from the `'/newcoll'` collection, but not from the system:

```
$ chelonia unlink /newcoll/thing
Unlinking '/newcoll/thing': unset
```

Now the file is not in the collection anymore:

```
$ chelonia list /newcoll
'/newcoll': collection
  empty.
$ chelonia stat /newcoll/thing
'/newcoll/thing': not found
```

But with the GUID of the file, it can still be accessed:

```
$ chelonia stat A63658B4-2C6E-46A3-8238-7D291F8F81C2
'A63658B4-2C6E-46A3-8238-7D291F8F81C2': found
  states
    checksumType: md5
    neededReplicas: 3
    size: 6
    checksum: a0186a90393bd4a639a1ce35d8ef85f6
  [...]
```

credentialDelegation

With the `credentialDelegation` or `cre` method it is possible to delegate credentials to the Bartender.

chelonia credentialDelegation

removeCredentials

With the `removeCredentials` or `rem` method it is possible to remove the previously delegated credentials.

chelonia removeCredentials

makeMountPoint

With the `makeMountPoint` or `makemount` method it is possible to create a mount point within the namespace of Chelonia which points to a GridFTP server.

chelonia makeMountPoint <LN> <URL>

The 'LN' is the requested Logical Name for the mount point, the 'URL' points to the GridFTP server.

unmakeMountPoint

With the `unmakeMountPoint` or `unmount` method it is possible to remove a previously created mount point.

chelonia unmakeMountPoint <LN>

The 'LN' is the Logical Name of the mount point.

2.4 Test suite

2.4.1 arctest

`arctest` is a simple utility that tests very basic functionalities of the middleware. It is convenient for:

- first-time users who do not know job description languages and yet want to test e.g. their credentials or client setup,
- system administrators who'd like to quickly test their installations without having to learn job description languages.

The `arctest` utility contains pre-defined test jobs which can be submitted either to a specific test site or to a regular Grid infrastructure. In addition, `arctest` can provide basic information about available user credentials (proxy certificate).

arctest [options]

Options:

<code>-c, --cluster</code>	<code>[-]url</code>	explicitly select or reject (-) a specific site
<code>-g, --index</code>	<code>[-]url</code>	explicitly select or reject (-) a specific index server
<code>-J, --jobid</code>	<i>integer</i>	test job number
<code>-j, --joblist</code>	<i>filename</i>	the file storing information about active jobs (on Linux, default <code>\$/arc/jobs.xml</code>)
<code>-o, --jobids-to-file</code>	<i>filename</i>	the IDs of the submitted jobs will be appended to this file
<code>-D, --dryrun</code>		add dryrun option to the job description
<code>-x, --dumpdescription</code>		do not submit – dump transformed job description to stdout
<code>-E, --certificate</code>		prints information about available user credentials
<code>-b, --broker</code>	<i>string</i>	select broker method (default is Random)
<code>-t, --timeout</code>	<i>seconds</i>	timeout for network communication, in seconds (default 20)
<code>-d, --debug</code>	<i>verbosity</i>	verbosity level, FATAL, ERROR, WARNING, INFO, VERBOSE or DEBUG - default WARNING
<code>-z, --conffile</code>	<i>filename</i>	configuration file (on Linux, default <code>\$HOME/arc/client.conf</code>)
<code>-v, --version</code>		print version information
<code>-h, --help</code>		print help page

There are currently three test jobs defined. Once submitted, their results can be inspected and retrieved in a usual manner, using `arccat`, `arcstat`, `arcget` etc..

Test job descriptions:

1. A classical “Hello World” job, printing *hello, grid* text to standard output at a remote execution site.
2. Lists all environment variables defined for the Grid user at the remote site (using standard output).
3. Downloads the pre-defined input file from the HTTP server, and produces an output file by copying input with a new name. This job thus demonstrates usage of `inputfiles` and `outputfiles` attributes of job description languages.

Pre-defined test jobs are Grid-flavor specific and are used depending of the selected execution site.

`arctest` is complementary to the `arcinfo` utility, which extracts information about Grid resources without submitting test jobs, and to `arcproxy -I`, which provides more detailed information about user credentials.

Chapter 3

URLs

File locations in ARC can be specified both as local file names, and as Internet standard *Uniform Resource Locators (URL)*. There are also some additional URL *options* that can be used.

The following transfer protocols and metadata servers are supported:

<code>ftp</code>	ordinary <i>File Transfer Protocol (FTP)</i>
<code>gsiftp</code>	GridFTP, the Globus [®] -enhanced FTP protocol with security, encryption, etc. developed by The Globus Alliance [5]
<code>http</code>	ordinary <i>Hyper-Text Transfer Protocol (HTTP)</i> with PUT and GET methods using multiple streams
<code>https</code>	HTTP with SSL v3
<code>httpg</code>	HTTP with Globus [®] GSI
<code>ldap</code>	ordinary <i>Lightweight Data Access Protocol (LDAP)</i> [10]
<code>rls</code>	Globus [®] <i>Replica Location Service (RLS)</i> [3]
<code>lfc</code>	LFC catalog and indexing service of EGEE gLite [1]
<code>srm</code>	Storage Resource Manager (SRM) service [8]
<code>arc</code>	for the Chelonia storage service, communicates with Bartenders, the path should be a Logical Name (LN) (supported only in <code>arcls</code> , <code>arcsub</code> and other <code>arc*</code> tools.)
<code>file</code>	local to the host file name with a full path

An URL can be used in a standard form, i.e.

```
protocol://[host[:port]]/file
```

Or, to enhance the performance, it can have additional options:

```
protocol://[host[:port]] [;option[;option[...]]/file
```

For a metadata service URL, construction is the following:

```
protocol://[url[|url[...]]@]host[:port] [;option[;option[...]]/  
lfn[:metadataoption[:metadataoption[...]]]
```

In user-level tools, URLs may be expressed in this syntax, or there may simpler ways to construct complex URLs. In particular, command line tools such as `arccp` and the xRSL and JSDL job description languages provide methods to express URLs in a simpler way.

For Chelonia, the syntax is

```
arc://<LogicalName>[?BartenderURL=<URL>]
```

where the `BartenderURL` may come from the `bartender` parameter of the client configuration file. Note that Chelonia URLs are **not** supported in `ngls`, `ngsub` and other `ng*` tools.

For the SRM service, the syntax is

```
srm://host[:port][;options]/[service_path?SFN=]file
```

Versions 1.1 and 2.2 of the SRM protocol are supported. The default `service_path` is `srm/managerv2` when the server supports v2.2, `srm/managerv1` otherwise.

The URL components are:

<code>host[:port]</code>	Hostname or IP address [and port] of a server
<code>lfn</code>	Logical File Name
<code>url</code>	URL of the file as registered in indexing service
<code>service_path</code>	End-point path of the web service
<code>file</code>	File name with full path
<code>option</code>	URL option
<code>metadataoption</code>	Metadata option for indexing service

The following options are supported for location URLs:

<code>threads=<number></code>	specifies number of parallel streams to be used by GridFTP or HTTP(s,g); default value is 1, maximal value is 10
<code>cache=yes no renew copy check</code>	indicates whether the GM should cache the file; default for input files is yes . renew forces a download of the file, even if the cached copy is still valid. copy forces the cached file to be copied (rather than linked) to the session dir, this is useful if for example the file is to be modified. check forces a check of the permission and modification time against the original source. (copy option is available in ARC 0.8.1 and above, check option is available in ARC 2.0.0 and above).
<code>readonly=yes no</code>	for transfers to <code>file://</code> destinations, specifies whether the file should be read-only (unmodifiable) or not; default is yes
<code>secure=yes no</code>	indicates whether the GridFTP data channel should be encrypted; default is no
<code>blocksize=<number></code>	specifies size of chunks/blocks/buffers used in GridFTP or HTTP(s,g) transactions; default is protocol dependent
<code>checksum=cksum md5 adler32 no</code>	specifies the algorithm for checksum to be computed (for transfer verification or provided to the indexing server). This is overridden by any metadata options specified (see below). If this option is not provided, the default for the protocol is used. checksum=no disables checksum calculation (adler32 and no options are available in ARC 0.8.1 and above).
<code>exec=yes no</code>	means the file should be treated as executable
<code>preserve=yes no</code>	specify if file must be uploaded to this destination even if job processing failed (default is no)
<code>guid=yes no</code>	make software use GUIDs instead of LFNs while communicating to indexing services; meaningful for <code>r1s://</code> only
<code>overwrite=yes no</code>	make software try to overwrite existing file(s), i.e. before writing to destination, tools will try to remove any information/content associated with specified URL

<code>protocol=gsi gssapi ssl tls ssl3</code>	to distinguish between different kinds of <code>https/httpg</code> and <code>srn</code> protocols. Here <code>gssapi</code> stands for <code>httpg</code> implementation using only GSSAPI functions to wrap data and <code>gsi</code> uses additional headers as implemented in Globus IO. The <code>ssl</code> and <code>tls</code> stand for usual <code>https</code> and especially usable only if used with <code>srn</code> protocol. The <code>ssl3</code> is mostly same as <code>ssl</code> but uses SSLv3 hadshake while establishing <code>https</code> connection. The default is <code>gssapi</code> for <code>srn</code> connections, <code>tls</code> for <code>https</code> and <code>gssapi</code> for <code>httpg</code> . In case of <code>srn</code> if default fails, <code>gsi</code> is then tried.
<code>spacetoken=<pattern></code>	specify the space token to be used for uploads to SRM storage elements supporting SRM version 2.2 or higher
<code>autodir=yes no</code>	specify if before writing to specified location software should try to create all directories mentioned in specified URL. Currently this applies to FTP and GridFTP only. Default for those protocols is <code>yes</code>
<code>tcpnodelay=yes no</code>	controls the use of the TCP_NODELAY socket option (which disables the Nagle algorithm). Applies to <code>http(s)</code> only. Default is <code>no</code> (supported only in <code>arc1s</code> and other <code>arc*</code> tools)
<code>transferprotocol=protocols</code>	specify transfer protocols for meta-URLs such as SRM. Multiple protocols can be specified as a comma-separated list in order of preference.

Local files are referred to by specifying either a location relative to the job submission working directory, or by an absolute path (the one that starts with `"/`), preceded with a `file://` prefix.

URLs also support metadata options which can be used for registering additional metadata attributes or querying the service using metadata attributes. These options are specified at the end of the LFN and consist of name and value pairs separated by colons. The following attributes are supported:

<code>guid</code>	GUID of the file in the metadata service
<code>checksumtype</code>	Type of checksum. Supported values are <code>cksum</code> (default), <code>md5</code> and <code>adler32</code>
<code>checksumvalue</code>	The checksum of the file

The checksum attributes may also be used to validate files that were uploaded to remote storage.

Examples of URLs are:

```
http://grid.domain.org/dir/script.sh
gsiftp://grid.domain.org:2811;threads=10;secure=yes/dir/input_12378.dat
ldap://grid.domain.org:389/lc=collection1,rc=Nordugrid,dc=nordugrid,dc=org
rls://gsiftp://se.domain.org/datapath/file25.dat@grid.domain.org:61238/myfile02.dat1
file:///home/auser/griddir/steer.cra
lfc://srm://srm.domain.org/griddir@lfc.domain.org/user/file1:guid=\
bc68cdd0-bf94-41ce-ab5a-06a1512764dc:checksumtype=adler32:checksumvalue=123456782
lfc://lfc.domain.org;cache=no/:guid=bc68cdd0-bf94-41ce-ab5a-06a1512764d3
```

¹This is a destination URL. The file will be copied to the GridFTP server at `se.domain.org` with the path `datapath/file25.dat` and registered in the RLS indexing service at `grid.domain.org` with the LFN `myfile02.dat`.

²This is a destination URL. The file will be copied to `srm.domain.org` at the path `griddir/file1` and registered to the LFC service at `lfc.domain.org` with the LFN `/user/file1`. The checksum will be compared to what is reported by the SRM service after the transfer. The given GUID and checksum attributes will also be registered in the LFC.

³This is a source URL. The file is registered in the LFC service at `lfc.domain.org` with the given GUID and can be copied or queried by this URL.

Chapter 4

ARC Client Configuration

The default behaviour of an ARC client can be configured by specifying alternative values for some parameters in the client configuration file. The file is called `client.conf` and is located in directory `.arc` in user's home area, e.g., on Linux:

```
$HOME/.arc/client.conf
```

If this file is not present or does not contain the relevant configuration information, the global configuration files (if exist) or default values are used instead. Some client tools may be able to create the default `$HOME/.arc/client.conf` on Linux, if it does not exist.

The ARC configuration file consists of several configuration blocks. Each configuration block is identified by a keyword and contains configuration options for a specific part of the ARC middleware.

The configuration file is written in a plain text format known as INI. Configuration blocks start with identifying keywords inside square brackets. Typically, first comes a common block: `[common]`. Thereafter follows one or more attribute-value pairs written one on each line in the following format:

```
[common]
attribute1=value1
attribute2=value2
attribute3=value3 value4
# comment line 1
# comment line 2
...
```

Most attributes have counterpart command line options. Command line options always overwrite configuration attributes.

Two blocks are currently recognized, `[common]` and `[alias]`. Following sections describe supported attributes per block.

4.1 Block `[common]`

defaultservices

This attribute is multi-valued.

This attribute is used to specify default services to be used. Defining such in the user configuration file will override the default services set in the system configuration.

The value of this attribute should follow the format:

```
service_type:grid:service_url
```

where `service_type` is type of service (e.g. `computing` or `index`), `grid` specifies type of middleware plugin to use when contacting the service (e.g. `ARC0`, `ARC1`, `CREAM`, etc.) and `service_url` is the URL used to contact the service. Several services can be listed, separated with a blank space (no line breaks allowed).

Example:

```
defaultservices=index:ARC0:ldap://index1.ng.org:2135/Mds-Vo-name=testvo,o=grid
└─index:ARC1:https://index2.ng.org:50000/isis
└─computing:ARC1:https://ce.arc.org:60000/arex
└─computing:CREAM:ldap://ce.glite.org:2170/o=grid
```

rejectservices

This attribute is multi-valued.

This attribute can be used to indicate that a certain service should be rejected (“blacklisted”). Several services can be listed, separated with a blank space (no line breaks allowed).

Example: `rejectservices=computing:ARC1:https://bad.service.org/arex`

verbosity

Default verbosity (debug) level to use for the ARC clients. Corresponds to the `-d` command line option of the clients. Default value is `WARNING`, possible values are `FATAL`, `ERROR`, `WARNING`, `INFO`, `VERBOSE` or `DEBUG`.

Example: `verbosity=INFO`

timeout

Sets the period of time the client should wait for a service (information, computing, storage etc) to respond when communicating with it. The period should be given in seconds. Default value is 20 seconds. This attribute corresponds to the `-t` command line option.

Example: `timeout=10`

brokername

Configures which brokering algorithm to use during job submission. This attribute corresponds to the `-b` command line option. The default one is the `Random` broker that chooses targets randomly. Another possibility is, for example, the `FastestQueue` broker that chooses the target with the shortest estimated queue waiting time. For an overview of brokers, please refer to Section 2.2.1.

Example: `brokername=Data`

brokerarguments

This attribute is used in case a broker comes with arguments. This corresponds to the parameter that follows column in the `-b` command line option.

Example: `brokerarguments=cow`

joblist

The file storing information about active jobs. This file will be used by commands such as `arcsub`, `arcstat`, `arcsync` etc. to read and write information about jobs. This attribute corresponds to the `-j` command line option. The default location of the file on Linux platforms is in the `$HOME/.arc/client.conf` directory with the name `jobs.xml`.

Example:

```
joblist=/home/user/run/jobs.xml
joblist=C:\\run\\jobs.xml
```

bartender

Specifies default *Bartender* services. Multiple Bartender URLs should be separated with a blank space. These URLs are used by the `chelonia` command line tool, the Chelonia FUSE plugin and by the data tool commands `arccp`, `arcls`, `arcrm`, etc..

Example: `bartender=http://my.bar.com/tender`

proxypath

Specifies a non-standard location of proxy certificate. It is used by `arcproxy` or similar tools during proxy generation, and all other tools during establishing of a secure connection. This attribute corresponds to the `-P` command line option of `arcproxy`.

Example: `proxypath=/tmp/my-proxy`

keypath

Specifies a non-standard location of user's private key. It is used by `arcproxy` or similar tools during proxy generation. This attribute corresponds to the `-K` command line option of `arcproxy`.

Example: `keypath=/home/username/key.pem`

certificatepath

Specifies a non-standard location of user's public certificate. It is used by `arcproxy` or similar tools during proxy generation. This attribute corresponds to the `-C` command line option of `arcproxy`.

Example: `certificatepath=/home/username/cert.pem`

cacertificatesdirectory

Specifies non-standard location of the directory containing CA-certificates. This attribute corresponds to the `-T` command line option of `arcproxy`.

Example: `cacertificatesdirectory=/home/user/cacertificates`

cacertificatepath

Specifies an explicit path to the certificate of the CA that issued user's credentials.

Example: `cacertificatepath=/home/user/myCA.0`

vomserverpath

Specifies non-standard path to the file which contains list of VOMS services and associated configuration parameters. This attribute corresponds to the `-V` command line option of `arcproxy`.

Example: `vomserverpath=/etc/voms/vomses`

username

Sets default username to be used for requesting credentials from Short Lived Credentials Service. This attribute corresponds to the `-U` command line option of `arcslcs`.

Example: `username=johndoe`

password

Sets default password to be used for requesting credentials from Short Lived Credentials Service. This attribute corresponds to the `-P` command line option of `arcslcs`.

Example: `password=secret`

keypassword

Sets default password to be used to encode the private key of credentials obtained from a Short Lived Credentials Service. This attribute corresponds to the `-K` command line option of `arcslcs`.

Example: `keypassword=secret2`

keysize

Sets size (strength) of the private key of credentials obtained from a Short Lived Credentials Service. Default value is 1024. This attribute corresponds to the `-Z` command line option of `arcslcs`.

Example: `keysize=2048`

certificatelifetime

Sets lifetime (in hours, starting from current time) of user certificate which will be obtained from a Short Lived Credentials Service. This attribute corresponds to the `-L` command line option of `arcslcs`.

Example: `certificatelifetime=12`

slcs

Sets the URL to the Short Lived Certificate Service. This attribute corresponds to the `-S` command line option of `arcslcs`.

Example: `slcs=https://127.0.0.1:60000/slcs`

storedirectory

Sets directory which will be used to store credentials obtained from a Short Lived Credential Service. This attribute corresponds to the `-D` command line option of `arcslcs`.

Example: `storedirectory=/home/mycredentials`

jobdownloaddirectory

Sets directory which will be used as the default job download directory. This attribute corresponds to the `-D` command line option of `arcget`.

Example: `jobdownloaddirectory=/home/myjobs`

idpname

Sets Identity Provider name (Shibboleth) to which user belongs. It is used for contacting Short Lived Certificate Services. This attribute corresponds to the `-I` command line option of `arcslcs`.

Example: `idpname=https://idp.testshib.org/idp/shibboleth`

4.2 Block [alias]

Users often prefer to submit jobs to a specific site; since contact URLs (and especially end-point references) are very long, it is very convenient to replace them with aliases. Block `[alias]` simply contains a list of alias-value pairs.

Alias substitutions is performed in connection with the `-c` command line switch of the ARC clients.

Aliases can refer to a list of services (separated by a blank space).

Alias definitions can be recursive. Any alias defined in a list that is read before a given list can be used in alias definitions in that list. An alias defined in a list can also be used in alias definitions later in the same list.

Examples:

```
[alias]
```

```
arc0=computing:ARC0:ldap://ce.ng.org:2135/nordugrid-cluster-name=ce.ng.org,Mds-Vo-name=local,o=grid
arc1=computing:ARC1:https://arex.ng.org:60000/arex
cream=computing:CREAM:ldap://cream.glite.org:2170/o=grid
crossbrokering=arc0 arc1 cream
```

4.3 srms.conf

If any data management commands are used with the Storage Resource Management (SRM) [8] protocol, the file

```
$HOME/.arc/srms.conf
```

(or its analogue on non-Linux platforms) may be created to store cached information on these services. For more information see the description inside this file.

4.4 Deprecated configuration files

ARC configuration file in releases 0.6 and 0.8 has the same name and the same format. Only one attribute is preserved (`timeout`); other attributes unknown to newer ARC versions are ignored.

In $ARC \leq 0.5.48$, configuration on Linux platforms was done via files `$HOME/.ngrc`, `$HOME/.nggiislist` and `$HOME/.ngalias`.

The main configuration file `$HOME/.ngrc` could contain user's default settings for the verbosity level, the information system query timeout and the download directory used by `ngget`. A sample file could be the following:

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

If the environment variables `NGDEBUG`, `NGTIMEOUT` or `NGDOWNLOAD` were defined, these took precedence over the values defined in this configuration. Any command line options override the defaults.

The file `$HOME/.nggiislist` was used to keep the list of default GIIS server URLs, one line per GIIS (see `giis` attribute description above).

The file `$HOME/.ngalias` was used to keep the list of site aliases, one line per alias (see `alias` attribute description above).

Acknowledgements

This work was supported in parts by: the Nordunet 2 program, the Nordic DataGrid Facility, the EU KnowARC project (Contract nr. 032691), the EU EMI project (Grant agreement nr. 261611) and the Swedish Research council via the eSENCE strategic research program.

getFile	31	arcsub	9	jobdownloaddirectory	45
list	31	arcsync	17	joblist	43
makeCollection	30	arctest	36	keypassword	44
makeMountPoint	35	chelonia	28	keypath	43
modify	32	credentialDelegation	35	keysize	44
move	31	delFile	32	password	44
policy	32	getFile	31	proxypath	43
putFile	31	list	31	rejectservices	42
removeCredentials	35	makeCollection	30	slcs	44
stat	29	makeMountPoint	35	srms.conf	45
unlink	34	modify	32	storedirectory	45
unmakeCollection	30	move	31	timeout	42
unmakeMountPoint	35	policy	32	username	44
commands		putFile	31	verbosity	42
arccat	15	removeCredentials	35	vomsserverpath	44
arcclean	19	stat	29		
arccp	25	unlink	34	D	
arcget	16	unmakeCollection	30	data management	24
arcinfo	18	unmakeMountPoint	35		
arckill	18	configuration		J	
arcls	24	bartender	43	job ID	11
arcmigrate	23	brokerarguments	42	job management	9
arcmkdir	27	brokername	42		
arcproxy	7	cacertificatepath	44	S	
arcrenew	20	cacertificatesdirectory	43	security	7
arcresub	22	certificatelifetime	44	submit job	9
arcresume	21	certificatepath	43		
arcrm	27	defaultservices	41	U	
arcslcs	9	deprecated files	46	URL	37
arcsrmping	28	idpname	45	options	38
				URLs	37