ARC Usage Record Logger

Configuration Guide

Henrik Thostrup Jensen `<htj@ndgf.org>`

1 Configuration

From version 0.8.1 ARC includes a usage record generator and a registrator to SGAS. This document describes how to configure these. The generator and registrator should come with the nordugrid-arc-arex package.

1.1 Requirements

- Python 2.4 or later

- Twisted Core and Web (http://twistedmatrix.com/)

- PyOpenSSL (https://launchpad.net/pyopenssl)

- ElementTree (http://effbot.org/zone/element-index.htm - only needed with Python 2.4)

Debian/Ubuntu package names:

```
python-twisted-core, python-twisted-web, python-openssl
python-elementtree (only needed with Python 2.4)
```

RPM based distributions (e.g., RHEL, CentOS, SL, Fedora, etc.):

```
python-twisted-core python-twisted-web pyOpenSSL
python-elementtree (only needed with Python 2.4)
```

Other Linux distributions: You are on your own :-)

1.2 arc.conf

To invoke the UR generator, an authplugin line must be set in the `[grid-manager]` section:

```
[grid-manager]
authplugin="FINISHED timeout=10,onfailure=pass /usr/libexec/arc/arc-ur-logger %C %I %S %U"
```

Change the path if you've installed ARC elsewhere than /usr/libexec/arc/. You'll need to restart the grid-manager for this to take effect (but wait until you finished reading this document).

The plugin will log to the file: /var/log/arc-ur-logger.log. This file will not appear until a job has finished.

Additional logger configuration happens in a seperate logger section. Example:

```
[logger]
log_dir=/var/spool/nordugrid/usagerecords/
log_all="https://sgas.ndgf.org:6143/sgas"
log_vo="bio.ndgf.org https://biosgas.ndgf.org:6143/sgas"
ur_lifetime=30
```

The `log_dir` option will set the top directory for the generated usage records. The option will default to /var/spool/nordugrid/usagerecords/ and can be left out. We suggest you leave out this option, unless you have a reason not to.

The `log_all` and `log_vo` options configure where to usage records are registered. All usage records will be registered to the URLs specified with the `log_all` option. It is possible specifiy multiple URLs be having them space seperated, e.g.:

```
log_all="https://sgas.ndgf.org:6143/sgas https://sgas.grid.dk:6143/sgas"
```

The `log_vo` option makes it possible to only register usage records run with certain VO users to given URL. The above example will register all usage records where the VO information includes the VO bio.ndgf.org to the url `https://biosgas.ndgf.org:6143/sgas`. It is possible to have multiple entries, by seperating entries with comma, e.g.,

```
log_vo="vo1 url1, vo2 url2"
```

It is NOT possible to have multiple `log_all` or `log_vo` lines.

The `ur_lifetime` option specifies how many days usage records are kept after being archived. The default is 30, and the option can be left out.

It is possible to set the logfiles for both the logger and the registrant. Furthermore it is possible to specify logging level for the logger (not possible or needed for the registrant, as it is not that verbose). Examples:

```
urlogger_logfile="/tmp/arc-ur-logger.log"
urlogger_loglevel="info"
registrant_logfile="/tmp/arc-ur-registrant.log"
```

By default the ur logger will write to `/var/log/arc-ur-logger.log` and the registrant to `/var/log/arc-ur-registration.log`. Valid options for logger log level are: `debug`, `info`, `warning`. The default is `info` which will write one line log per job, assuming everything goes as planned.

1.3   cron

To register the usage records and registrant should be invoked regularly be CRON. We suggest every hour. The crontab entry would typically look like this:

```
0 * * * * /usr/libexec/arc/arc-ur-registrant
```

To ensure that the registrant is working, you can run the script from the command line first. Note that the script will still write its log to to /var/log/arc-ur-registration.log. By running the script with -s its ouput will be directed to stdout.