

# USER'S AND ADMINISTRATOR'S MANUAL OF AMGA METADATA CATALOG V 2.3.0 (EMI-2)

S. Ahn, A. Calanducci, T. Huh, S. Hwang, G. Park, N. Kim, B. Koblitz, N. Santos

April 16, 2012

## **Abstract**

This is the manual for users and administrators of AMGA. It intends to give an overview on the installation of the client and server packages as well as the client-api packages. Examples of the usage of the command line clients and the client APIs are given.



# Contents

<b>1</b>	<b>Overview</b>	<b>5</b>
<b>2</b>	<b>Installation</b>	<b>5</b>
2.1	Client installation . . . . .	5
2.2	Server installation . . . . .	6
2.3	Server upgrade . . . . .	8
2.4	Installation from Source . . . . .	9
2.5	Creating a script for the site-BDII . . . . .	10
<b>3</b>	<b>Configuration of the C++ and Java command line clients</b>	<b>10</b>
3.1	Configuration of the C++ command line client . . . . .	10
3.2	Configuration of the Java command line client . . . . .	12
<b>4</b>	<b>Metadata Access from the Shell</b>	<b>13</b>
4.1	Manipulating Collections . . . . .	14
4.2	Commands for Manipulating Attributes . . . . .	15
4.3	Commands for entry manipulation . . . . .	16
4.4	Finding and Updating Entries . . . . .	18
4.5	Native SQL Query . . . . .	19
4.6	Permission Handling . . . . .	19
4.7	Capabilities . . . . .	20
4.8	Index Management . . . . .	20
4.9	Sudo . . . . .	20
4.10	Table Constraints . . . . .	20
4.11	Views . . . . .	21
4.12	Sequences . . . . .	21
4.13	Backing Up Data . . . . .	21
4.14	Site management . . . . .	22
4.15	Various administrative commands . . . . .	23
4.16	Replication . . . . .	23
4.16.1	Commands for Slave Nodes . . . . .	23
4.16.2	Commands for Master Nodes . . . . .	24
4.17	Federation . . . . .	25
4.18	Commands for Replica Management . . . . .	26
<b>5</b>	<b>Metadata Queries</b>	<b>27</b>
<b>6</b>	<b>Native SQL Queries</b>	<b>29</b>
6.1	SQL-92 Entry Level <direct_data_statements> . . . . .	29
6.2	LIMIT, OFFSET and JOIN . . . . .	30
6.3	Current Limitation . . . . .	30
<b>7</b>	<b>AMGA data types</b>	<b>30</b>
<b>8</b>	<b>Configuring the AMGA Server and the Replicatin Daemon</b>	<b>31</b>
8.1	Format of the Configuration File . . . . .	31
8.2	Example of a server configuration file . . . . .	31
8.3	Description of the server options . . . . .	32
8.3.1	Replication Settings . . . . .	34
8.3.2	Federation Settings . . . . .	35
<b>9</b>	<b>Replication in AMGA</b>	<b>35</b>
9.1	Overview . . . . .	35
9.1.1	Features Overview . . . . .	35
9.1.2	Concepts . . . . .	35
9.1.3	Architecture Overview . . . . .	36
9.1.4	Operation . . . . .	36
9.2	Setup . . . . .	37

9.2.1	Setting up a Master node . . . . .	37
9.2.2	Setting up a Slave node . . . . .	38
9.3	Security . . . . .	38
9.3.1	Replicating Security Information . . . . .	38
9.3.2	Security at Slave Nodes . . . . .	39
9.3.3	Connections from the Slave to the Master . . . . .	39
9.3.4	Granting the replication right on the master . . . . .	39
9.4	Limitations . . . . .	39
9.4.1	Concurrent Updates on Master . . . . .	39
9.5	Tutorial - Setting up a Master to Single Slave Setup . . . . .	40
9.5.1	Configuring the master . . . . .	40
9.5.2	Configuring the slave . . . . .	41
<b>10</b>	<b>Federation in AMGA</b>	<b>42</b>
10.1	Overview . . . . .	42
10.2	Setup . . . . .	43
10.3	Authentication . . . . .	43
10.4	Accessing Patterns . . . . .	43
10.5	Fault Handling . . . . .	43
10.6	Limitations . . . . .	44
<b>11</b>	<b>User Management</b>	<b>44</b>
11.1	Configuration . . . . .	44
11.2	Management via a Grid Map file . . . . .	45
11.3	Management of Users using the database backend . . . . .	45
11.4	Management via a VOMS . . . . .	46
11.5	Management via VO-Certificates . . . . .	46
<b>12</b>	<b>Users, Groups and ACLs</b>	<b>46</b>
12.1	Users . . . . .	46
12.2	Groups . . . . .	46
12.3	Access Control Lists . . . . .	47
12.4	The sudo command . . . . .	47
12.5	The sticky bit . . . . .	47
<b>13</b>	<b>Accessing preexisting databases in AMGA</b>	<b>47</b>
<b>14</b>	<b>Using the C++ Client API</b>	<b>49</b>
<b>15</b>	<b>Using the Java Client API</b>	<b>51</b>
<b>16</b>	<b>Using the Python Client API</b>	<b>52</b>
<b>17</b>	<b>Definition of the Client Server Protocol</b>	<b>53</b>
<b>18</b>	<b>Monitoring support in AMGA</b>	<b>55</b>
<b>19</b>	<b>Multi-threaded AMGA Server</b>	<b>56</b>
19.1	Configuration . . . . .	56
<b>20</b>	<b>SOAP Interface</b>	<b>56</b>
20.1	Installation . . . . .	57
20.1.1	SOAP client installation . . . . .	57
20.1.2	SOAP server installation . . . . .	57
20.2	How to use SOAP client tools . . . . .	58
20.3	WS-DAIR Architecture . . . . .	59
20.3.1	Data access service and data resources . . . . .	59
20.3.2	Data Resource Abstract Name & Data Resource Address . . . . .	60
20.3.3	Supported Dataset and Language . . . . .	60
20.3.4	Installation from Source . . . . .	60

20.4 Howto make WSDAIR clients . . . . .	61
<b>A Changelog</b>	<b>64</b>

# 1 Overview

AMGA is a metadata service for the Grid. In a more general way this is a database access service for Grid applications which allows user jobs running on the Grid to access databases by providing a Grid style authentication as well as an opaque layer which hides the differences of the different underlying database systems from the user. To achieve this, AMGA is a service sitting between the RDBMS and the user's client application.

In addition to this database translation layer, AMGA intends to solve another problem database services face on the Grid which is latencies. AMGA intends to provide a replication layer which makes databases locally available to user jobs and replicate the changes between the different participating databases. A simple implementation based on PostgreSQL asynchronous replication is already working.

All examples given in this manual as well as additional documentation in particular the reference manuals of all APIs are given on the projects home page at <http://project-arda-dev.web.cern.ch/project-arda-dev/metadata/>

## 2 Installation

The server and the clients (C++, Java and Python) are provided as RPMs packages. RPMs are currently only supported for CERN Scientific Linux or RedHat Enterprise Linux. For users running other operating systems (including Windows and other unix flavours), the Java client is also provided as a platform-independent package (a tar ball), including the Java API, a command line client, some examples and the documentation. You can find the packages in the [download directory](#). Download the latest version. As of writing this it is 2.3.0

You can install the `emi.amga.amga-cli` and the `python` and `java api` packages independently of the server package, however the server packages depend on the client package.

### 2.1 Client installation

To install the command line client and C++ api, you will only need to download the `emi.amga.amga-cli` rpm itself. By default the package will be installed into `/usr`, so you will need to have write permission for this directory. Install the package via

```
$ rpm -i emi.amga.amga-cli-{version}.{platform}.rpm
or
$ sudo yum install emi.amga.amga-cli
```

Copy the `/etc/mdclient.config` client configuration file into the directory from which you intend to work or into `~/mdclient.config` and customize it according to the instructions in **Configuration of the C++ and Java command line clients (p. 10)**.

If the amga client was built against the editline library instead of the standard unix readline library, you will need to also install that RPM.

To install the RPM with Java API, install the file named `glite-amga-api-java-0.X.Y.rpm`. This RPM is architecture-independent and will install in any Unix platform that supports RPM packages. This RPM contains only the Java API (no command line client or documentation). It installs the file `glite-amga-api-java.jar` in `/opt/glite/share/java/`. To use the Java API, this file must be included on the classpath.

The tar ball with the Java Client API does not need any special installation procedure, apart from unpacking it. It provides both the documentation, the Java API, and two command line utilities similar to the ones installed by the C++ RPM, with the advantage of running in any platform with a Java virtual machine implementation (must be at least Java 1.4 compliant). These command line tools are:

- **mdjavaclient** - an interactive command line shell with the server.
- **mdjavacli** - an utility to submit a single command to the server. For use in shell scripts.

The tar ball includes two scripts to start these command line tools:

- **mdjavaclient.sh** : for Linux

- **mdjavacli.sh** : for Windows The interactive command line client can also be started directly by executing the class `arda.md.javaclient.ConsoleClient`.

There is also a Python Client API module available as an RPM package, which will install the modules under `/opt/glite/python2.2/site-packages/amga`.

- **RPM package name** : `glite.amga.api-python-1.1.0-1.noarch.rpm`

Alternatively you can use the Python2.3 RPM which works nicely on Debian Sarge. If you prefer a custom installation you can download the source tarball `glite.amga.api-python-1.1.0.tar.gz` and install it via the Python installation mechanism (if you don't know how this works, `./setup.py -help` should get you started).

## 2.2 Server installation

The AMGA server depends on some external dependencies which are provided by the EMI team for yum based installation.

The following packages are necessary to use AMGA and must be installed first: These are standard SL packages and can be also found in the usual EMI repositories.

- boost
- editline
- gdbm
- libxml2
- openssl
- unixodbc
- libuuid(SL6 only,SL5 is OS Default)

If AMGA was compiled against the Globus environment and not the SL environment in order to use the Globus versions of several SL system libraries like openssl, then you will also need to install the Globus RPM. The EMI project currently only provides versions of AMGA which does not depend on Globus.

Now install the server RPMs for AMGA (the server needs the client):

```
$ wget
http://amga.web.cern.ch/amga/downloads/${version}/${platform}/emi.amga.amga-server-${version}.${p
$ rpm -i emi.amga.amga-server-${version}.${platform}.rpm
or
$ sudo yum install emi.amga.amga-server
```

You also need a database and the appropriate ODBC driver. AMGA currently supports 4 different database backends via ODBC drivers. You will need to get at least one of the currently supported 4 database backends installed, including their ODBC driver. You have the choice among PostgreSQL, MySQL, Oracle and SQLite. The default is PostgreSQL and this database should be set-up correctly when installing the AMGA RPM if you have PostgreSQL and its ODBC driver installed:

```
$ sudo yum install postgresql-server postgresql
```

If you want to use a different database or you want to have a different setup of the ODBC driver, have a look at the more detailed instruction in the **Installation from Source (p. 5)**.

Note: For AMGA to work properly, you need at least version 08.01.0200 of the PostgreSQL ODBC driver. On the default SL, the driver is too old, so that it is necessary to install the latest driver as follows.

```
$ sudo yum install postgresql-odbc
```

You will need to get at least one of these 4 database backends in order to use the service. ODBC should have been installed correctly by the AMGA rpm, if you want to use a standalone installation using PostgreSQL. If you don't know about ODBC, some more hints can be found in the **Installation from Source (p. 5)** or on a general page on the Unix ODBC User Manual.

For the configuration, you first have to create a DB user, make sure he can connect via a TCP/IP connection (even locally this is required since ODBC does not work via a Unix Domain socket) and then set up an ODBC data source. This can be done with the single script as follows.

```

$ su root

$ /etc/rc.d/init.d/mdservice init

The steps in manual are:

$ su root

$ service postgresql initdb (If you are using postgresQL 8.4 or higher)

$ /etc/init.d/postgresql start
$ /etc/init.d/postgresql stop

$ cat /var/lib/pgsql/data/pg_hba.conf
# ....
local  metadata      arda      trust
host   metadata      arda      127.0.0.1/24    trust
host   metadata      arda      ::1/128         trust

$ /etc/init.d/postgresql start
$ su postgres
$ createuser -S -d arda      # Create DB user arda, allow him new DBs
$ createdb -O arda metadata # New DB metadata owner is arda
$ createlang plpgsql metadata # Allow stored procedures

```

The ODBC data source is created by appending the following lines to `/etc/odbc.ini`:

```

[PSQL]
Description      = AMGA metadata catalogue database
Driver           = PostgreSQL_AMGA
Trace            = No
TraceFile        = /tmp/metadata/odbc.log
Database         = metadata
Servername       = localhost
Port             = 5432
ReadOnly         = No
Address          = localhost:5432
User             = arda

```

The Driver needs to name a valid driver description in the `/etc/odbcinst.ini` file:

```

[PostgreSQL_AMGA]
Description      = ODBC for AMGA PostgreSQL
FileUsage       = 1
CPOutput        = 120
CPReuse         =
Debug           = 0
Threading       = 1
CommLog         = 0
DontDLClose    = 1
Driver          = /usr/lib/psqlodbc.so

```

This setup should be done automatically correctly when installing the `unixODBC` or `postgresql-odbc` RPMs on other distributions. If you want to check your ODBC installation, you can use the `DataManagerII` application which can be found in the `unixODBC-kde` package on SLC.

Now you can initialize the database using the `createInitialXXX.sql` scripts, where `XXX` has to be replaced with the DB of choice. You can find the `createInitialXXX.sql` file in the `/usr/share/doc/glite-amga-server/` after the installation of the server RPM:

Initialization for PostgreSQL:

```
$ psql -Uarda metadata < /usr/share/doc/glite-amga-server/createInitialPG.sql
```

Finally activate this data source and configure the AMGA server in the `amgad.config` in `/etc/` file as described in **Configuring the AMGA Server and the Replicatin Daemon (p. 31)**.

You should now be able to start the service and verify whether it is running by doing

```
$ /etc/init.d/mdservice start
$ tail /var/log/amgad.log
```

If the server is complaining about a missing boost-library (starting `amgad: /usr/bin/amgad: error while loading shared libraries: libboost_thread-gcc-mt-1_32.so.1.32.0: cannot open shared object file: No such file or directory`), create one with:

```
$ su
$ cd /usr/lib
$ ln -s libboost_thread.so.1.32.0 libboost_thread-gcc-mt-1_32.so.1.32.0
```

If you want to set up AMGA and postgresQL to start at system boot on Linux, you can use follow steps

```
$ sudo chkconfig postgresql on
$ sudo chkconfig --add mdservice
```

To confirm whether everything was registered properly, run the following :

```
$ chkconfig --list postgresql
$ chkconfig --list mdservice
```

If you want to change the configuration of the metadata server, you should go on and read **Configuring the AMGA Server and the Replicatin Daemon (p. 31)**.

## 2.3 Server upgrade

To upgrade from an old version(1.3, 1.9) to the 2.x version, The first thing you need is to to backup metadata.

```
$ mdcli "dump /" > amga.backup
```

If you finished to backup your metadata, install the AMGA 2.x rpm after deleting the old version.

```
$ rpm -e glite-amga-server
$ rpm -e glite-amga-cli
$ rpm -i emi.amga.amga-server-${version}.${platform}.rpm
```

After upgrading the amga software, it is necessary to upgrade the backend DB schema. Here, let us assume that you were using AMGA 1.3 and trying to upgrade to AMGA 2.1 and the backend DB used was PostgreSQL. Then, upgrading can be easily done with the following command.

```
$ psql -Uarda metadata < /usr/share/doc/glite-amga-server/upgradePG_13_20.sql
$ psql -Uarda metadata < /usr/share/doc/glite-amga-server/upgradePG_20_21.sql
```

Currently we provide various upgrade scripts for Oracle, MySQL5, and PostgreSQL. There are also scripts upgrading from 1.9 to 2.x version at `/usr/share/doc/glite-amga-server/` directory.

If you have a failure on upgrade, it is possilbe to go back to the old version with the data backup. To recover data, you need to reinstall the old version after deleting new version and restore data as follows.

```
$ mdclient < amga.backup
```



## 2.4 Installation from Source

To install the ARDA metadata server from source you will need to first download the source distribution from the `download` directory.

For compilation you need to install a development package for ODBC (e.g. `unixodbc`) This should be part of any standard distribution. On CERN SLC you simply should be able to do:

```
$ yum install unixODBC unixODBC-devel
```

You will then need the `libxml2` development library, which should also be part of any distribution. On SLC you can simply install it using

```
$ yum install libxml2-devel
```

For the server you will finally need to install the boost libraries:

```
$ yum install boost-devel
```

Now you should be ready to compile and install the AMGA server:

```
$ tar xvfz emi.amga.amga-server-2.x.x.tar.gz
$ cd emi.amga.amga-server-2.x.x
$ ./configure
$ make
$ su
$ make install
```

If you want to make rpm packages for your architecture and install them so that you will be able to deinstall them easily later or install them also on other machines, do

```
$ tar xvfz emi.amga.amga-server-2.x.x.tar.gz
$ cd emi.amga.amga-server-2.x.x
$ ./configure --with-readline --enable-rpm-rules --without-globus
$ make rpm
```

You will need to get at least one of the currently supported 4 database backends installed, including their ODBC driver. You have the choice among PostgreSQL, MySQL, Oracle and SQLite:

- **PostgreSQL**. This is the easiest solution, since ODBC drivers are included in any distribution package. On SLC you can install the Postgres ODBC driver using

```
$ yum install postgresql-odbc
```

PostgreSQL needs a little setup: You need to create a database and a user. Access should be made possible via TCP/IP from localhost. The `scripts/init-arda-psqldb.sh` script should be able to do this for you.

- **MySQL**. Again, ODBC drivers are included in any distribution. On SLC you can install the MySQL ODBC driver using

```
$ yum install MyODBC
```

Again, make sure you have a user created which gets access to the database on MySQL.

- **SQLite** is a file-based database. You need to get the ODBC driver and compile if it is not part of your distribution.
- **Oracle** can be used from CERN. You need to get the instant client and ODBC driver from Oracle. Which you can freely download after a registration [here](#). . You will need to set up Oracle for the service names at you place. At CERN you might try to learn something from the Oracle Linux pages.

Examples of ODBC configuration files can be found in the `scripts` directory. Copy the `odbc.ini` and `odbcinst.ini` configurations into `/etc` or into your home directory (but then called `.odbc.ini` and `.odbcinst.ini`). In `odbc.ini` you need to configure the database used by the server and which server is being connected to. Examples are given for all 4 databases. The users (and passwords if required) must then be setup in the `amgad.config` file. The ODBC configuration can be checked with e.g. `gODBCConfig` or `DataManagerII` which are probably installed along the ODBC package or other ODBC clients like OpenOffice. If you don't know about ODBC, some more hints can be found in <http://www.unixodbc.org/doc/> the Unix ODBC User Manual .

Some initial tables need to be setup in the database. You need to run one of followings to setup the database.

```
$ sqlplus user/passwd < createInitialOracle.sql
$ psql -User database < createInitialPG.sql
$ sqlite3 dbfile.db < createInitial.sql
$ mysql database < createInitialMySQL5.sql
```

You should now proceed to configure the server (**Configuring the AMGA Server and the Replicatin Daemon (p. 31)**) and start it up.

The following database and ODBC versions are known to work:

- SQLite 3.2.1, but not 3.2.7 with the 0.64 and 0.65 ODBC drivers. In 3.2.7 the ODBC driver is incompatible to the library.
- MySQL 4.0.x and 4.1.x work, 3.x does not.
- PostgreSQL works starting from version 7.2, however only one attribute can be added at a time in versions before 8.0. Explicitely tested were versions 7.2, 7.3, 7.4 and 8.0 allways in their latest sub-versions.

## 2.5 Creating a script for the site-BDII

To create script for publishing version information for AMGA server, you can use the single script as follows.

```
$ su root
$ /etc/init.d/mdservice site <SITENAME>
```

This script will install `glite-info-provider-service` (If it is not installed) and make a script file named `glite-data-metadata-amga`. If you already have the `site-BDII` then this new script file will be created in the `/var/lib/bdii/gip/provider` directory, otherwise, the file will be create in your home directory.

This script can be used for publishing GLUE2 information via LDAP.

## 3 Configuration of the C++ and Java command line clients

### 3.1 Configuration of the C++ command line client

The AMGA C++ command line clients as well as all other executables using the `MDClient (p. ??)` class of the C++ client package need to read an `mdclient.config` file for their configuration. This configuration file is searched first in the current directory, then as `$HOME/.mdclient.config` and finally as `/usr/etc/mdclient.config`. Only the port and server of the configuration file can be overridden on the command line of the clients:

```
$ mdclient [-p port] [hostname]
```

The following is an example configuration file:

```

$ cat mdclient.config
# Connection options
Host = localhost
Port = 8822

# User settings
Login = koblitz
PermissionMask = rwx
GroupMask = r-x
Home = /

# Security options
UseSSL = yes
# AuthenticateWithCertificate = 1
# CertFile=/home/koblitz/.globus/usercert.pem
# KeyFile=/home/koblitz/.globus/userkey.pem
# UseGridProxy = 1
# Password = secret
# VerifyServerCert = 1
# IgnoreCertificateNameMismatch = 0

# If server certificates are verified, local certificates need to be loaded:
# TrustedCertDir = /etc/grid-security/certificates

```

The following options are supported:

- **Host:** The name of the host to connect to. This option can be overridden on the command line of mdclient. (Default: localhost)
- **Port:** Port of the mdserver to connect to. Can be overridden on the command line of mdclient using the -p option. (Default: 8822)
- **Login:** The login name of the user on the AMGA server. All entries created in the catalogue will have this owner. This is also the user which you need to authenticate to the AMGA server if authentication is enabled. (Default: NULL which gives the default role when authenticating with a VO certificate)
- **PermissionMask:** A 3 character string giving the owner permissions of newly created entries in the metadata catalogue.
- **GroupMask:** A 3 character string giving the group permissions of newly created entries in the metadata catalogue.
- **Home:** The home-directory. The default is "/".
- **Name:** Can be used to give a full name to the server, comparable to the comment in an /etc/passwd file. Currently only used for information in the server.
- **UseSSL:** Possible values are no, try, require (synonyme is yes). Default is no. Needed for any authentication using certificates (also proxy certificate). You want this if you intend to use passwords which are not sent in plain text. If you use SSL the entire session will be encrypted. Some servers may require you to use SSL to connect. If you want to be sure that SSL is always used you need to set this to require or yes.
- **AuthenticateWithCertificate:** Set this to 1 to enable certificate based authentication, also grid-proxy certificates. You will need to either enable normal certificates via a Cert-File, KeyFile option pair, or use a grid proxy certificate via the UseGridProxy option. If you specify both, then the grid proxy gets precedence. This option does not work if you have not enabled SSL via UseSSL! In fact you will try to authenticate as the user named in Login, so you need to have this field set as well. However, the Password field is ignored unless certificate based authentication fails in which case the password is tried.

- **CertFile:** Path to your x509 certificate in .pem format. For this option to have any effect, you need also options UseSSL and AuthenticateWithCertificate enabled and UseGridProxy disabled!
- **KeyFile:** Path to your x509 private key in .pem format. If the key is encrypted you will be asked for the passphrase on client startup. For this option to have any effect, you need also options UseSSL and AuthenticateWithCertificate enabled and UseGridProxy disabled!
- **UseGridProxy:** Tries to use the a grid proxy certificate in /tmp/x509up\_u[user-id].
- **Password:** If a password is given, password based authentication will be tried. You should want to use an SSL connection with this. This is *\_not\_* the password for your private key. For that you will always be prompted on the command line. For a discussion on how to escape special characters (#\ ) see **Configuring the AMGA Server and the Replicatin Daemon (p. 31)** .
- **VerifyServerCert:** Verifies the server certificate against CA certifiates in TrustedCertDir, default is 1: yes.
- **IgnoreCertificateNameMismatch:** Do not try to match the DN in the server certificate to the canonical name of the server. Useful for services that have an alias or are multi-homed. Default is off.
- **TrustedCertDir:** A directory with certificate authority certificates to verify the server certificate.

### 3.2 Configuration of the Java command line client

The Java client also reads its configuration from a file, called by default `mdjavaclient.config`. This file is searched in the same places as the C++ API looks for the `mdclient.config`, that is, first in the current directory, then in `$HOME/.mdclient.config` and finally on `/usr/etc/mdjavaclient.config`. Like the C++ client, only the port and server of the configuration file can be overridden on the command line of the clients:

```
$ mdjavaclient.sh [-p port] [hostname]
```

Next are the properties recognized on the `mdjavaclient.config` file. The following have the same syntax as in the C++ configuration.

- **Host**
- **Port**
- **Login**
- **PermissionMask**
- **GroupMask**
- **Name:**

The following exist only on the the Java configuration file or else have a slightly different syntax:

- **AuthMode:** Possible values are GridProxy, Certificate, Password, None. Default is None.
- **Password:** The user password. Only used when the AuthMode is set to Password. If not provided here, it must be provided at runtime by the user.
- **UseSSL:** Possible values are 1 for enable, 0 for disable. Default is 0. Needed for any authentication using certificates and grid proxies. If you use SSL the entire session will be encrypted. Some servers may require you to use SSL to connect.
- **CertFile:** Path to your x509 certificate in .pem format. For this option to have any effect, you need also options UseSSL and the AuthMode set to Certificate.

- **KeyFile:** Path to your x509 private key in .pem format. If the key is encrypted you will be asked for the passphrase on client startup. For this option to have any effect, you need also options UseSSL and the AuthMode set to Certificate.
- **PrivateKeyPassword:** Password for private key. If not defined, the password must be provided at runtime.
- **VerifyServerCert:** Verifies the server certificate against CA certificates in TrustedCertDir.
- **TrustedCertDir:** A directory with certificate authority certificates to verify the server certificate.

## 4 Metadata Access from the Shell

You can access the metadata catalogue either with the `mdclient` metadata terminal tool as configured in the last section, or via the `mdcli` (or `mdjavacli` for those using the java package) command line tool which allows you to directly issue metadata commands on the shell, it's output is intended to be easily parseable by scripting languages:

```
$ mdcli -p8822 -slocalhost listattr /
t
text
f
float
```

Metadata commands are parsed into pieces which are each separated by white space similarly to shell commands. If you want the white space to be part of one piece of the command itself, for example when you want to set an attribute to a string which contains white space, you must enclose it in single quotes: `' '`. Single quotes are part of the command syntax and used when parsing the commands into parts. You need them every time a part shall contain spaces. Double quotes however are used in queries **Metadata Queries (p. 27)** (expressions evaluated by the database backend) to distinguish strings from variable references and common values. In order to put a single quote into a command argument or any other character, you can use an octal code, e.g. to get a `"`.

Note that quotes may be removed by the shell when parsing a shell command, so if you are using the `mdcli` tool where the AMGA command is given on the command line, you will need to protect these single quotes from being removed by the shell with double quotes: `" "`. The various APIs will in contrast usually (that is the Python and Java APIs do so, but not the C-API) automatically quote any arguments you pass to them with single quotes so they are not to be used in those APIs. The following is an example with `mdclient`, `mdcli` and the Python API showing how quotes are being used:

```
$ mdclient
Query> find /files '/files:producer="CERN"'

$ mdcli find /files "'/files:producer="CERN"'"
```

An in python:

```
mdclient.find('/files', '/files:producer="CERN"');
```

The metadata server uses a streaming protocol. Some APIs (for example the Java one) allow to interrupt the streaming of a response. The same is true for the `mdclient`. Pressing **CTRL-C once during the transmission of the result** will interrupt the streaming of the result. Only pressing **CTRL-C** a second time will terminate the client.

In the following is given a list of metadata commands. Additionally commands may be available for group or user access management, as described in **Users, Groups and ACLs (p. 46)** or **Management of Users using the database backend (p. 45)** depending on the server setup. To find out which commands are available on the server you are connected to, use the `help` command.

## 4.1 Manipulating Collections

- **createdir /parentdir/dir [option]:** Creates the directory dir if it does not yet exist but parentdir already exists. The directory is created with the current owner and the same ACLs as the parent directory. The option field is a comma separated list of options (no spaces allowed). The following options are available as of AMGA 1.1:
  - shared: Subdirectories created under this directory share the same schema and database table of the parent directory
  - acls: Creates a directory with acls for every entry, this is currently only supported by PostgreSQL and MySQL5 and only if the necessary supporting stored procedures have been installed first.
  - type=<type> Specifies the data type of the entry column. Explicitly supported are int, float, date.
  - table=<TableType>: Only MySQL. Allows to specify the storage type of the table, e.g. InnoDB. Necessary because not all tables are alike. Specify e.g. MyIsam, this allows for example GIS functionality.
- **ls [options] [directory]:** Returns the name of all subdirectories, files, indexes, sequences, and views in the directory in a very similar way to the UNIX ls command. The following options can be used with ls.
  - l : use a long listing format
  - s : returns subdirectories only
- **dir [directory]:** Returns the name of all subdirectories and files in the directory. The result is returned in the following way:

```
0
entry 1
entry-type-1
...
entry n
entry-type-n
EOT
```

where the entry-type is either 'entry' or 'collection'. If AMGA collaborates with a file catalogue this command will effectively show the content of the file catalogue. If you want to see which entries have already been attached to a schema in the AMGA part use the listentries command.

- **stat [dir/entry]:** Returns information on a given entry or directory. They can also be a pattern in the case of several entries. You need read permission to get this information. For an entry the following is returned (if the information is not stored in the table, it

```
0
name
type (=entry)
owner-permissions
group-permissions
owner
size
creation time (Format: 1999-12-22 04:05:06)
guid
expires
link

EOT
```

For a directory the following is returned:

```
0
name
type [options]
owner-permissions
owner
acls
EOT
```

- **rmdir path:** Removes all directories matching path. Directories are only deleted if they are empty and they have no attributes defined.
- **pwd:** Prints the current directory which you can change with cd.
- **cd path:** Changes the directory to the given path. Possible errors are:
  - 1 No such file or directory
  - 4 Permission denied

## 4.2 Commands for Manipulating Attributes

- **addattr dir (key type)+:** Adds new keys to the list of keys of a directory. In a relational database backend these keys become the columns of a table associated to a directory. You should only use one key/type pair currently for compatibility reasons because some older backends like PostgreSQL <=7.4 do not allow to alter a table adding several columns. Possible types are explained in **AMGA data types (p. 30)**. The type is only used as a hint for the back end to store the data efficiently and allow efficient queries. The type may be ignored by the implementation (e.g. if the back end is a filesystem). In a filesystem the types and defined keys could be stored as attributes of directories. Some storage backend may allow you to define keys on a per-entry basis. Possible errors are:
  - 1 No such file or directory
  - 4 Permission denied. You need write permission on the directory.
  - 7 Illegal Key: Keys must be alphanumeric. The following keys are reserved: ACL, CREATED, FILE, GROUP\_RIGHTS, GUID, LINK, MD5, OWNER, PERMISSIONS, SIZE.
  - 9 Internal error: All kinds of errors like duplicate keys
  - 23 Not a directory. You can only add attributes to a directory.
- **removeattr dir/entry key+:** Removes the attributes or keys from the list of attributes of the directory dir or the directory of the given entry, or if the implementation allows it from the list of attributes of a given entry. Attributes can only be removed if they are not used by any entry. So you either have to remove all entries for which this key is set or use clearattr to set the value of the attributes to NULL. In order to remove an attribute, write permissions on the directory are necessary. Possible Errors are:
  - 1 No such file or directory
  - 4 Permission denied. You need write permission on the directory.
  - 10 No such key
  - 14 Attribute in use
- **schema\_create dir (attr type)+ [option]:** Creates a new directory with a given schema. This is an atomic replacement for a sequence of mkdir and addattr. The meaning of the optional option argument is backend dependent and you should not use it if you want to retain this independence. With a MySQL backend you can give here the name of the table engine, for PostgreSQL the keyword **inherit** will make the table inherit its schema from the parent directory.

- **setattr file (attribute value)+:** Sets a list of attributes of a file to given values. The attributes must exist.
- **getattr pattern (attribute)+:** Returns the filename and all attributes in turn for every file matching pattern. The following output is returned on success:

```

0
file 1
attr 1
...
attr n
file 2
attr 1
...
attr n
...
file n
attr 1
...
attr n
EOT

```

- **listattr file:** Returns a list of all attributes and their types in the following format:

```

0
attr 1
type 1
attr 2
.
.
.
type n
EOT

```

- **clearattr path attribute:** Sets the attributes of all files matching path to NULL. Path may currently contain wildcards only in the file-part. On success returns 0.

### 4.3 Commands for entry manipulation

- **addentry entry (attr value)+:** Creates a new entry and assigns the given values to the provided attributes. Examples:

```

Query> addentry /testdir/a id 10
Query> addentry b id 10 finished '0ct-10-2004'

```

Possible errors are:

- 3 Illegal command: Syntax error
  - 4 Permission denied: You need write permission on the directory to create a file in it.
  - 7 Illegal Key
  - 10 No such key
  - 15 Entry exists
- **insert entry (attr value)+:** Creates a new entry and assigns the given evaluated values to the provided attributes. Same as addentry, except that the values are evaluated first, similar to an INSERT statement in SQL.



- **addentries (entry)+:** Creates the given entries in the catalogue. This command is done in a transaction, that is either all entries are inserted or none. Entries can be spread over several directories. Example:

```
Query> addentries a
Query> addentries /test1/a /test2/a /test1/b
```

Possible errors are:

- 1 No such file or directory: You tried to insert into a directory which did not exist.
  - 3 Illegal command: Syntax error
  - 4 Permission denied
- **rm [-options] pattern [condition]:** Removes all files matching pattern, where pattern may only contain wild cards in the file part. If a condition is given, than that condition needs to be met by the entry's metadata. In order to remove an entry you need write permissions on the parent directory. The pattern can also be the name of a directory if a condition is given, if the table is a plain table, the pattern `_must_` be a directory name and there must be a condition given. Valid flags are: `-r` (recursively delete). Possible errors are:
    - 1 No such file or directory: You tried to insert into a directory which did not exist.
    - 4 Permission denied
    - 20 Not an entry.
    - 29 Operation not permitted on plain table.
  - **listentries [directory/schema]:** Returns the name of all entries in a given directory or schema. This differs from the `dir` command in that it will not show any directories but also that it shows only entries attached to a schema in the case of an AMGA catalogue collaborating with a file-catalogue. The result is returned in the following way:

```
0
entry 1
...
entry n
EOT
```

- **transaction:** Starts a transaction. Any changes to the backend of AMGA are done only when committed. To cancel a transaction use `abort`.
- **upload dir (attribute)+:** Starts an upload of entries into the catalogue. Currently a static restriction of the prototype is that there can be only up to 98 attributes assigned like this. After the upload is initialized, the `put`, `abort` and `commit` commands are allowed. Errors are returned by the call immediately, the OK code is delayed till the entire upload is successfully committed.
- **put file (values)+:** Inserts a new entry during upload. Errors are returned by the call immediately, OK is delayed until upload is committed.
- **abort:** Aborts upload or transaction. Errors and OK are returned by the call immediately.
- **commit:** Commits upload or transaction. Errors and OK are returned by the call immediately.

## 4.4 Finding and Updating Entries

- **find path query:** Returns a list of filenames matching path and fulfilling the query with their attributes. The path may currently contain only wild cards in the file name part. Query must be enclosed in single quotes. Strings in the query must be quoted with double quotes. For the supported syntax see the chapter about **Metadata Queries (p. 27)**. **WARNING:** Be careful with patterns which also match a subdirectory, the result is undefined. The result is returned in the following way:

```
0
file 1
.
.
.
file n
EOT
```

The following errors may occur: 8 Illegal query: The query could not be parsed or violates security rules. No further information is currently provided on the reasons for this error. 2 Connection to DB failed or syntax error in path or query

- **updateattr pattern (attribute value)+ condition:** Updates attributes of entries matching a pattern in a single collection based on a condition. The values to which the attributes are updated can contain attributes as variables. Complex expressions are allowed as values. The condition may reference attributes of other collections. Updates are atomic. Examples:

```
Query> updateattr /testdir1/* events events+1 'events>100'
Query> updateattr /testdir1/* events events+1 '/testdir2:key > 0'
```

The first example increases the number of events of every file in /testdir1 which has more than 100 events by one. The second example increases the number of events of every file in /testdir1 provided there is an entry in the collection /testdir2 which has the attribute "key" set to anything larger than 1 (usefull to do locking by clients: putting such an entry into /testdir2 would lock /testdir1). For the supported syntax for the queries see the chapter about **Metadata Queries (p. 27)**.

- **updateattr\_single pattern (attribute value)+ condition:** Almost same as updateattr, but this updates attributes of a random single entry matching a pattern in a single collection based on a condition. This command is useful when AMGA is used as a placeholder for task retrieval. Example:

```
Query> updateattr_single /task status "running" 'status="waiting"'
>> /task/simulration-005
```

- **update pattern (attribute value)+ condition:** Same as updateattr, but the values are not evaluated prior to insertion into the table. This command works with bound variables, which will also fix problems with SQL command length limitations in Oracle.
- **selectattr (attribute)+ condition:** Selects attributes from several collections based on a condition doing an inner join on the collections based on a join condition. The FILE attribute is used to select the entry name of an entry. Example:

```
Query> selectattr /jobdir:FILE /configdir:id /jobdir:eventGen /configdir:id
'/jobdir:events>1000 and /configdir:key=/jobdir:key'
```

This selects the entry-name of a job, the id in the configuration, the event generator name of a job and the id in a configuration for all jobs and configurations where the job has more than 1000 events and the keys attributes of the jobs and configurations match. For the supported syntax for the queries see the chapter about **Metadata Queries (p. 27)**.

As of AMGA 1.2 `selectattr` also supports constraints to the query similar to a `SELECT` clause. Queries can now take the form:

```
query [distinct] [limit xx [offset yy]] [order exp] [group_by exp]
```

where the `distinct` keyword translates into a `SELECT DISTINCT`, the `limit` and `offset` clause limits the number of rows returned and with the `order` clause rows can be ordered according to the given expression. The `group_by` construct corresponds to the `GROUP BY` SQL statement.

## 4.5 Native SQL Query

Supporting native SQL query greatly eases the work needed to port existing SQL-based database applications to the Grid using AMGA and complements the existing metadata query language in AMGA.

The followings are currently supported commands for native SQL queries.

- **SELECT**
- **INSERT**
- **UPDATE**
- **DELETE**

They conforms to the SQL-92 Entry Level `<direct_data_statements>`. All the keywords in native SQL queries should be provided in capital letters. However table names, aliases, and column names can be either of capital or small. For the details, refer **Native SQL Queries (p. 29)**.

## 4.6 Permission Handling

- **whoami**: Prints out the name of the current user. Note that this command does not need any connections of the AMGA server and can thus be also used to do a test on whether an AMGA server is alive and what response time it has.
- **chown [-R] entry/dir new\_owner**: Changes the owner of a directory or entry. Only the owner of an entry is allowed to execute this, or the root-user. `chown` does not check whether the user exists, since user management is considered to be handled outside of AMGA (ideally). `[-R]` option changes owner of all the subdirectories recursively. Possible errors are:
  - 1 No such file or directory
  - 4 Permission denied
- **chmod entry/dir new\_permissions**: Changes the access permissions of an entry or directory. Entries have owner and group-permissions, while directories have owner permissions and group permissions are handled via ACLs. Group permissions for entries allow you to remove privileges granted for all entries in a directory via the directories ACLs. The format of `new_permissions` is `rw-rwx` for entries and `rw-x` for directories where "-"-signs can be substituted for the letters if you do not want to give a certain privilege. The permissions for entries are the concatenation of first the user and then the group rights. The `x`-Flag allows a user to enter a directory or respectively list an entry. `r`-and `w`-flags allow users to read/write metadata while the `w`-flag for directories allows users to create or delete entries in the given directory. Users cannot list directories for which they don't have read permissions. The command works also for patterns and uses a transaction. Possible errors are:
  - 1 No such file/directory
  - 4 Permission denied

## 4.7 Capabilities

Capabilities are additional attributes assigned to individual users. They are used for example to allow a user to replicate login information. Currently no mapping of VOMS capabilities is done, but this could be a future use-case.

- **capabilities\_add <user> <capability>**: Adds the given capability to the user's capabilities. Only root can do this.
- **capabilities\_remove <user> <capability>**: Removes the given capability from the user's capabilities. Only root can do this.
- **capabilities\_list [user]**: Lists all capabilities of a user, default is the current user.

## 4.8 Index Management

- **index\_create name dir 'key1...keyn' [type]**: Creates an index name on a collection directory using several attributes and a given algorithm. Algorithms depend on the backend. The index is later referred to /collection/name in **index\_remove**.
- **index\_remove /path**: Removes an index.
- **index\_list /path**: lists all indexes.

## 4.9 Sudo

Since AMGA 2.0 a sudo command exists, which allows the root user to become any other user. The syntax is: `sudo <user>`

## 4.10 Table Constraints

- **constraint\_add\_not\_null directory attribute name**: Adds a not NULL constraint for the given attribute of the directory. Name is the name used to refer to the name of the constraint. It must be unique for that directory. Write permissions on the directory are necessary for this operation.
- **constraint\_add\_unique directory attribute name**: Adds a UNIQUE constraint for the given attribute of the directory. Name is the name used to refer to the name of the constraint. It must be unique for that directory. Write permissions on the directory are necessary for this operation. NOTE: On MySQL you can use attribute(length) to set the length of indexed columns.
- **constraint\_add\_reference directory attribute reffered\_attr name**: Adds a foreign key constraint for the given attribute of the directory. The foreign key is given by the referenced attribute which must fully qualify that attribute including the table part, e.g. /dir:attr. Name is the name used to refer to the name of the constraint. It must be unique for that directory. Write permissions on the directory are necessary for this operation.
- **constraint\_add\_check directory check name**: Adds a check constraint to the directory. Check constraints are boolean expression which must be true for all entries inserted into the directory. An example would be `events > 0` requiring the value assigned to the events attribute to be positive. Name is the name used to refer to the name of the constraint. It must be unique for that directory. Write permissions on the directory are necessary for this operation.
- **constraint\_add\_primary\_key directory key(s)**: Adds primary key constraint of the given directory, which by default is on the entry name only. Example:

```
Query> constraint_add_primary_key . t,i
```

In the given example a the primary key becomes the pair of the t, and i attributes.

- **constraint\_drop directory p\_key:** Drops the constraint with the given name from the directory. Write permissions on the directory are necessary for this operation. The primary key of a table can be removed by specifying the name "p\_key".
- **constraint\_list directory:** Prints all constraints of a directory. You need read permissions on the concerned directory.

## 4.11 Views

Views allow you to create virtual new tables (directories) that combine the information of other tables, similar to what `selectattr` does. AMGA uses the native support of the database to provide views, so the actual behaviour depends on the database backend. For example some backends (like PostgreSQL) allow you to update an existing view, which actually updates the tables behind it.

An important use-case of views are to support access restrictions to attributes (the columns of the underlying table). This is a typical use-case for views also in normal database usage.

Views can be accessed and deleted like normal directories.

- **view\_create name maindir attr\_1 ... attr\_n condition:** Creates a view with the given name based on the entries in the given directory, attaching the attributes given in the list to these entries, based on the join condition.

In the following examples, the first one shows a use case where a view (view1) is created using all the entries in the current directory, but using only the attr1 or attr2 columns. After assigning the right permissions to the resulting new directory (./view1), this can be made readable for users who need to read these attributes, while they will not have access to the rest of the attributes in the . directory. In the second example a view (view2) is created combining attributes from the current directory and the dir subdirectory.

```
Query> view_create view1 . attr1 attr2 ''
Query> view_create view2 . attr1 ./dir:attr2 'dir:FILE = FILE'
```

## 4.12 Sequences

Sequences allow the creation of a sequence of integer numbers, which are guaranteed to be unique. They are also monotonically increasing at least during a single AMGA connection. The exact implementation depends on the database backend, which can optimize handing out parts of the sequence in batches, so that two consecutive connections not necessarily get first a smaller number in the sequence and then the larger. Sequences are not supported by MySQL <5.0 and SQLite. On MySQL and Oracle sequences are implemented through stored procedures. In PostgreSQL the native mechanism is used.

Sequences behave like another directory in a directory. They cannot be deleted with `rmdir`, however, instead `sequence_remove` must be used. The name of the sequence must be lower case due to limitations in some backends.

- **sequence\_create name dir [increment] [start value]:** Creates a new sequences with the given name in the given directory. The name of the sequence is then /dir/name. It is possible to define the increment as well as the start value. Note that backends may not necessarily follow this behaviour strictly if multiple connections are being used.
- **sequence\_next sequence:** Gets the next value from a sequence.
- **sequence\_remove sequence:** Deletes a sequence.

## 4.13 Backing Up Data

- **dump [-sec\_all|sec\_none] [dir]:** Recursively dumps the contents of a directory and all subdirectories so that they can be recreated by calling the sequence of AMGA commands printed out. If no directory is specified, it uses "/", making a full catalogue dump. The first option controls whether entry permissions and ACLs are included in the dump: `-sec_all` includes them, while `-sec_none` only dumps the metadata.

Only root is allowed to use this command. Possible errors are:

- 4 Permission denied
- **user\_dump [dir]:** Dumps the contents of a the user database such that it can be recreated from the sequence of AMGA commands printed out.  
Only root is allowed to use this command. Possible errors are:
  - 4 Permission denied
- **grp\_dump [dir]:** Dumps out the information on all existing groups so that they can be recreated by calling the sequence of AMGA commands printed out.  
Only root is allowed to use this command. Possible errors are:
  - 4 Permission denied

#### 4.14 Site management

For using replication, each AMGA server needs to know about the other servers that take part in the system, in order to communicate with them. These will be referred to as *<it>sites</it>*. The information about other sites is stored in the backend. Sites have the following configuration properties:

```

id
name
hostname
port
login
password
use_ssl
authenticate_with_certificate
cert_file
key_file
use_grid_proxy
verify_server_cert
trusted_cert_dir
require_data_encryption

```

`id` is a numeric identifier internal to each AMGA instance and generated automatically by AMGA when the site is inserted in the configuration. `name` is an human-readable identifier of the site, which can be freely chosen by the administrator. `hostname` and `port` is the network address of the remote site. The rest of the properties control the security settings of the connection to the master and are similar to the ones defined in the `mdclient.config` file, having a similar meaning. Section **Configuration of the C++ and Java command line clients (p. 10)** describes their usage.

The following commands can be used to manage the sites and their configuration:

- **site\_list:** Lists all sites and their IDs. Shows only a summary for each site, consisting of `id`, `name`, `hostname` and `port`.
- **site\_add site\_name [hostname:port]:** Registers a site. Only root can perform this operation.
  - 4 Permission denied
- **site\_remove site:** Removes a site. Only root can perform this operation.
  - 4 Permission denied
- **site\_set\_properties site property value [property value]\*:** Updates one or more configuration properties of a site. Only root can perform this operation.
  - 4 Permission denied

- **site\_get\_properties site property [property]\*:** Gets one or more configuration properties of a site. Only root can perform this operation.
  - 4 Permission denied
- **site\_list\_properties site:** List all configuration properties of a site, as a list of key/value pairs. Only root can perform this operation.
  - 4 Permission denied
- **site\_dump [siteName]\*:** Outputs the configuration of sites as a list of metadata commands that can be executed by AMGA to recreate the configuration. If no `siteName` is provided, it dumps all sites in the configuration. Only root can perform this operation.
  - 4 Permission denied

## 4.15 Various administrative commands

- **import tablename dir:** Makes the table given by `tablename` available under the given `dir` as a plain table. Only root can perform this operation.
  - 4 Permission denied
- **execute script [Options]...:** Executes a script in the directory given by the `ExecRoot` option in the server configuration file. The script is run with the user-id given in the `ExecUser` configuration variable, if the `amgad` is started as root. The script has access to the `USER` and `DN` environment variables giving the user name of the `amga` user and the distinguished name if a certificate was used to authenticate to AMGA. To enable this feature, `ExecRoot` must be set. The administrator of AMGA is responsible to set this up carefully and make sure that scripts are checking the permissions in `USER/DN`. The stdout of the script is returned through AMGA to the client. This gives AMGA a functionality similar to a `cgi-bin` in a web-server. The following error codes exist:
  - 21 Function not implemented: `ExecRoot` not set.
  - 4 Permission denied: Cannot execute anything outside of `ExecRoot`.
- **backend:** Returns the name of the current database backend.
- **version:** Returns a version number of the connected AMGA server.

## 4.16 Replication

The following are the commands used to control replication from the AMGA command line interface. Some of them are for nodes acting as slaves, others for nodes acting as masters. Nodes acting both as slave and master can use all of them. Section **Replication in AMGA (p. 35)** provides the background information required to understand these commands.

### 4.16.1 Commands for Slave Nodes

Slave nodes are responsible for initiating replication, by contacting the master and requesting the replication of the directories they are interested on. This is done using the following commands:

- **rep\_list\_mounts:** Lists all local mounts and their current state. It prints out the following information for each mount

```
<nodeID>:<directory> - <currentXID>, <state>
```

where:

- `<nodeID>` is the master from where the directory is being replicated.
- `<directory>` is the root directory of this mount.

- <currentXID> is the `xid` of the last log that was applied to this mount
  - <current state> is either `Disconnected` if there is no connection to this master, `Reconnecting` (after a broken connection, trying to reconnect to master) or `Receiving` (connected to master and receiving logs).
- **rep\_mount [-nosync | -noperms] nodeID remoteDirectory:** Mounts `remoteDirectory` from the master with id `nodeID`. After mounted, `remoteDirectory` is used to refer to this mount and is therefore also called `mountPoint`. This command fails if `remoteDirectory` already exists locally or if its parent directory does not exist. For instance, to mount `/a/b`, directory `/a` must be created and `/a/b` should not exist. It also synchronizes the local directory by copying the contents of the mount point from the master to the slave, but it does not make the slave start to receive logs. For that, use `rep_start_receive`. After this command executes successfully, the mount is on the `Inactive` state. By default, the permissions and acls associated with the metadata are also replicated. If `-noperms` is specified, the metadata will be replicated without any security information. If `-nosync` option is specified, existing metadata is not synchronized but changes are replicated to the slave after `rep_mount`. On the slave, the metadata will be owned by the local user performing replication (which currently must be root).
  - **rep\_umount [-noerase] mountPoint:** Unmounts the given mount point. The local contents of `mountPoint` is deleted, as well as the root directory of the mount point. This command only works if the connection to this master is inactive, that is, it should not be receiving logs. See `rep_stop_receive`. If `-noerase` option is specified, the existing metadata are not erased with `res_umount`. The `nodeID` has to be assigned via `site_add` with the port of the site pointing to the **replication daemon**, and where at least the `login` `site-property` has to be set.
  - **rep\_mount\_users nodeID:** Replicates user and group information from the given master. The root user and any groups owned by root are not replicated. The slave must not have any user or group defined, except root or groups owned by root. Each slave can only replicate users and groups from a single master. Apart from these restrictions, the replication of users and groups is handled by any other mount.
  - **rep\_umount\_users:** Stops replicating users and deletes all users and groups imported from the master. This command only works if the connection to the master is inactive, that is, it should not be receiving logs. See `rep_stop_receive`.
  - **rep\_users\_allow user:** Allows the given user to replicate user (`login`) information. This needs to be allowed done on the master as well as on the slave.
  - **rep\_users\_disallow user:** Disallows the given user to replicate user information
  - **rep\_start\_receive master:** Connects to `master` starts receiving the replication logs for all mount points originating from this node. This opens a TCP connection that is kept open while the slave is running. If the slave is shutdown, the next time it restarts it will try to reestablish the connection to the master to continue receiving logs. The state of all mounts from this master is changed to `Receiving`.
  - **rep\_stop\_receive master:** Stops receiving logs from `master`. The TCP connection is closed and the state of all mount points originating from this node is changed to `Disconnected`. In this state, connections will not be reestablished when the slave is restarted.

#### 4.16.2 Commands for Master Nodes

The main responsibility of master nodes is to configure the access control rights of slaves. Slaves connect to the master using the standard AMGA users and authenticate in the same way. Access control is done using the replication right, which is granted to users to control the directories they are allowed to replicate. The following commands allow granting and removing this right:



- **rep\_allow directory user:** Grants to `user` the right to replicate `directory` and all its subdirectories. Having this permission is a necessary and sufficient condition for a user to be allowed to replicate a directory. In particular, the normal user permissions play no role in deciding to allow a user to replicate a directory.
- **rep\_disallow directory user:** Remove from `user` the right to replicate `directory` and all its subdirectories.
- **rep\_users\_allow user:** Allows the given user to replicate user (login) information. This needs to be allowed done on the master as well as on the slave.
- **rep\_users\_disallow user:** Disallows the given user to replicate user information.
- **rep\_show\_permissions:** Show the replication rights for all directories.
- **rep\_list\_subscribers:** Show the list of the current subscribers, with their subscriptions. This includes information about whether users/groups and permissions are being replicated and the last xid acknowledge by the subscribers.

AMGA also has special commands for user and group management. They are optional and may not be available on your installation for example if it collaborates with a file catalogue and uses the permission system of that catalogue. For more information see **Users, Groups and ACLs (p. 46)** and **User Management (p. 44)**.

#### 4.17 Federation

The following are the commands used to control federation from the AMGA command line interface. Section **Federation in AMGA (p. 42)** provides the background information required to understand these commands. One AMGA server may initiate federation, by contacting other AMGA server and requesting the federation of the directories it is interested on. This is done using the following commands:

- **fed\_list\_mounts:** Lists all local mounts and their current state. It prints out the following information for each mount

```
Query> fed_list_mounts
<directory1>
<site1>
<directory1>
<site2>
...
```

where:

- `<directory>` is the root directory of this mount.
- `<site>` is the site name connected to this mount
- **fed\_mount site directory:** Mounts `directory` from other AMGA server named with `site`. After mounted, `directory` is used to refer to this mount and is therefore also called `mountPoint`. This command fails if `remoteDirectory` already exists locally or if its parent directory does not exist. For instance, to mount `/a/b`, `directory /a` must be created and `/a/b` should not exist.
- **fed\_umount directory:** Unmounts the given mount point.
- **fed\_policy :** shows current policy about federation. The first line shows whether connection failures are allowed in federation, and the second line shows whether execution failures are allowed.
- **fed\_check\_failure** shows whether last command had failures. The first line shows whether connection failures happened, and the second line shows whether execution failures happened.

## 4.18 Commands for Replica Management

The set of commands below provide functionality for the management of replicas, in terms of a global storage index, as well as a local replica lookup system.

In order to have support for replica management, the entries in the table need to have support for GUIDs. This is enabled, by using the `makedir` command instead of the `createdir` command, which takes exactly the same options, except, that permissions and the `guid` feature are switched on by default.

GUIDs are inserted as root via

```
Query> addentry /fdir/entry GUID beda818d-1090-489c-b9f9-6f3156a81828
```

If you want to play with this, use the `uuidgen` UNIX-command to create new guides. To list the `guid` values, use the `-g` switch in `ls` or `stat`.

The following commands allow the management of replicas of files, similar to a file catalogue local to a site. This functionality provides the possibility to map GUIDs to URLs (also known as PFNs, Physical File Names). Note that AMGA only supports one URL per `guid`. This means that if you want to support multiple replicas (e.g. on multiple Storage Elements), you need to have one AMGA catalogue per SE:

- **replica\_add guid url <expires>**: Adds a URL as a replica to the `guid`. That `GUID` must exist. Optionally it is possible to define an expiration time (e.g. for volatile storage). That time must be a date parseable by the database (this functionality is not implemented). The `guid` must already be known to the system, e.g. by adding an entry in a `guid`-enabled table with `addentry`. This is a root-only command.
- **replica\_remove url [guid]**: Removes a replica URL of all files. It is possible to restrict the operation to `guid` only. This command requires root privileges as the permissions on the entry itself are not checked.
- **replica\_get [guid ... | -l lfn ...]**: Returns the URLs of all `guids` or `lfns` given.
- **lfn\_lookup guid** looks up the LFN and the URL of the `guid`.

The next set of commands allows to assign sites as locations of replicas to `guids`. This is useful if the catalogue is being used in a global mode. In that case it can easily provide information on sites with replicas of a given file. This functionality is also referred to as a Site-Index (SI) and AMGA stores the information in a highly-optimized bitfield. The **Site (p. ??)** Index functionality can be combined with the AMGA functionality for local replica management above. Note that the ideal mode of operation would be if local AMGA catalogues store URLs and can be locally updated, replicating the information back to a global SI. However, this mode of operation is currently not supported by AMGA's replication functionality.

The SI functionality is highly optimized in AMGA for the usage with Oracle, but also PostgreSQL was successfully tested and the performance is very good. The `guids` are stored in binary format as well as the storage index bitmap, so the information is stored in the most compact format. Indexes are used for row-lookup.

- **replica\_register guid site**: Registers a replica for `guid` at the given `site`. The `site` can be either the `site id` (number) as returned by `site_list`, or the name of the `site`. This command requires root privileges.
- **replica\_unregister guid site**: Unregisters a replica `guid` from a `site`. The `site` can be either the numerical id, or its name. This command requires root privileges.
- **replica\_list [guid ... | -l lfn ...]**: Returns the sites which hold replicas of all `guids` or `lfns` given.

The function `lfn_lookup guid` is used to retrieve the LFN and URL from a given `guid`, however in AMGA 2.0 it only returns the directory in which a file resides.

## 5 Metadata Queries

AMGA provides its own query language which is similar to the SQL query language. It tries to offer a large subset of the common functionality of database systems in a transparent way to the user. The biggest difference to SQL is that in AMGA's query language tables are referred to as references to directories. AMGA will ensure access restrictions on the data so that users cannot infer data from queries if they have no read-access to that data.

Queries are performed in the following AMGA commands:

```
Query> find entry_pattern query_condition
Query> selectattr column_1_query ... column_n_query query_condition
Query> updateattr attr_1 update_query_1 .... attr_n update_query_n query_condition
Query> updateattr_single attr_1 update_query_1 .... attr_n update_query_n query_condition
```

A query condition is a query which returns a boolean in order to select or not select an entry for retrieval or update. Examples are

```
/jobdir:events>1000 and /configdir:key=/jobdir:key
like(/jobdir:FILE, "t%")
```

Query conditions are used in the WHERE statements of the SQL queries which are passed to the backends.

The other queries used in the AMGA commands return general values which are returned to the user in the `selectattr` command or which are used to update attributes in the `updateattr` command.

Queries can contain either literal values like numbers or strings which are marked by double quotes. Make sure to use single quotes around double quotes if the string contains spaces. Queries can also contain attributes which are evaluated in the query by filling in the values of the attributes for the current value. In AMGA all queries are in fact inner joins over all tables mentioned in any of the queries of command, that is the all possible combinations of all entries of all tables are made and those selected matching the query condition. Inner joins are the most common type of join. Some database systems provide no other kind.

References to attributes take the form:

```
<directory>:<attribute>
```

where relative paths to the directory (which is synonyme for table or schema, here) are allowed. Examples are:

```
Query> selectattr /test:t 'like(t, "Test%")'
Query> selectattr count(/test:t) 'like(t, "Test%")'
```

From the above example you can see that also functions are allowed. Function names are case-sensitive and lowercase. The following functions are available:

- `lower(string)`: Converts string to lower case.
- `upper(string)`: Converts string to upper case.
- `count(x)`: Aggregate function. Counts how often the attribute is set (not = NULL)
- `abs(x)`: Absolute value of x.
- `sin(x)`: The sine of x.
- `cos(x)`: The cosine of x.
- `tan(x)`: The tangens of x.
- `atan(x)`: The arc-tangens of x.
- `sqrt(x)`: The square root of x.
- `ln(x)`: The natural logarithm of x.

- `log(x)`: The base 10 logarithm of x.
- `rnd()`: A random number between 0 and 1.
- `sum(x)`: The aggregate sum of x.
- `max(x)`: The aggregate maximum of x.
- `min(x)`: The aggregate minimum of x.
- `avg(x)`: The aggregate average of x.
- `length(string)`: The length of the string.
- `pow(x, y)`: x to the power of y.
- `mod(x, y)`: x modulo y.
- `concat(str1, str2)`: The concatenation of str1 and str2.
- `like(str, pattern)`: Whether str is like pattern. The pattern is an SQL90 pattern.
- `substr(str, n, m)`: The substring of length m of str starting at n.
- `isnull(arg)`: Checks that the argument is NULL;
- `notnull(arg)`: Checks that the argument is not NULL;
- `is(condition, t, e)`: Evaluates to t if condition is fulfilled, otherwise to e.

Queries can contain the following operators: `+`, `-`, `*`, `/`, `=`, `and`, `or`, `not`, `>=`, `<=`, `<>` or `!=`

Special attribute names refer to the properties of an entry:

- `FILE`: The name of the entry.
- `LINK`: The link pointed to.
- `OWNER`: The owner of the entry.
- `PERMISSIONS`: The owner's permission.
- `GROUP_RIGHTS`: The group-rights. These names could e.g. be used to restrict access to entries using a `VIEW`.

As of version 1.2.11, AMGA supports other table joins apart from the implicate cross join. The supported joins are left and right outer joins and the inner join. The following shows some examples

```
Query> selectattr /t1:num /t1:name /t2:num /t2:value '/t1:num = 1 join_left_on(/t1:, /t2:, /t1:num = /t2:num) limit 1'
```

```
Query> selectattr /t1:name /t2:value 'join_right_on(/t1:, /t2:, /t1:num = /t2:num)'
```

The following is a list of the supported joins and their translation into SQL:

- `join_left_on(<left>, <right>, <condition>)` left LEFT OUTER JOIN right ON condition
- `join_right_on(<left>, <right>, <condition>)` left RIGHT OUTER JOIN right ON condition
- `join_inner_on(<left>, <right>, <condition>)` left INNER JOIN right ON condition
- `join_cross(<left>, <right>)` left C JOIN right ON condition

The exact syntax of AMGA queries is described in the annotated `parser.y++` and `lexer.l++` sourcecodes.

## 6 Native SQL Queries

As of version 1.9, AMGA supports the native SQL query, which greatly eases the work needed to port existing SQL-based database applications to the Grid using AMGA and complements the existing metadata query language in AMGA. This feature also makes full use of the advanced security features of AMGA, namely schema ACLs and access through the GSI.

The followings are currently supported native SQL queries.

- SQL-92 Entry Level <direct\_data\_statements>
  - SELECT, INSERT, UPDATE, and DELETE
- Part of SQL-92 Intermediate Level
  - OUTER JOIN
- Not defined in the SQL Standard
  - LIMIT and OFFSET

**IMPORTANT NOTE: All the keywords in native SQL queries should be provided in capital letters.** However table names, aliases, and column names can be either of capital or small.

Table names are referred to as references to directories. And reference to an attribute can take one of the forms as follows:

```
<attribute>
<directory>:<attribute>
<directory>.<attribute>
```

Special attribute such as FILE, OWNER, PERMISSIONS, GROUP\_RIGHT refer to the properties of an entry. For detail, refer metadata\_queries .

### 6.1 SQL-92 Entry Level <direct\_data\_statements>

AMGA supports <direct\_data\_statements> in the SQL-92 entry level that is supported by most DBMS such as PostgreSQL, Oracle, MySQL, SQLite and etc. The SQL-92 Entry Level <direct\_data\_statements> includes 4 basic statements: SELECT, INSERT, UPDATE, and DELETE. The brief syntax of supported statements is as follows.

- SELECT [ALL | DISTINCT] select\_list from\_clause [where\_clause] [group\_by\_clause] [having\_clause] [order\_by\_clause] [;]
- INSERT INTO table [column\_list] VALUES ( { value\_list | query\_specification } ) [;]
- UPDATE table SET assignment\_list [where\_clause] [;]
- DELETE FROM table [where\_clause] [;]

Examples are as follows.

- ```
- Query> SELECT EMPNUM,HOURS FROM WORKS WHERE PNUM='P2' ORDER BY 2 ASC

- Query> SELECT SUM(HOURS),AVG(HOURS),MIN(HOURS),MAX(HOURS) FROM WORKS
WHERE EMPNUM='E1';

- Query> SELECT * FROM STAFF FIRST1, STAFF SECOND2
WHERE FIRST1.CITY = SECOND2.CITY AND FIRST1.EMPNUM < SECOND2.EMPNUM

- Query> INSERT INTO WORKS VALUES ('P22','E22',NULL);

- Query> UPDATE STAFF SET GRADE = NULL
WHERE EMPNUM = 'E1' OR EMPNUM = 'E3' OR EMPNUM = 'E5';

- Query> DELETE FROM WORKS WHERE PNUM='P18';
```

The exact syntax of SQL-92 entry level is described in the following site.

- <http://www.contrib.andrew.cmu.edu/~shadow/sql/sql1992.txt>.

The output of `SELECT` statement prints attribute names followed by attributes values. Other statements return how many rows have been changed as a result.

Possible Errors are:

- 4 Permission denied : You need proper permission on the table
- 88 SQL Parsing Error : SQL syntax is wrong

## 6.2 LIMIT, OFFSET and JOIN

`LIMIT`, `OFFSET` and `JOIN` are supported in AMGA even though they are not defined in the SQL-92 Entry Level, because there were many requirements from user communities. With `LIMIT` and `OFFSET`, it is possible to specify a range of rows to return. The SQL `JOIN` combines records from two tables, resulting in a new, temporary table, sometimes called a "joined table".

Examples are as follows.

```
- Query> SELECT * FROM WORKS RIGHT OUTER JOIN STAFF USING (EMPNUM)
ORDER BY WORKS.EMPNUM LIMIT 1 OFFSET 2
```

## 6.3 Current Limitation

- The native SQL queries do not provide means to modify ACLs on tables or entries, which can be set properly by other AMGA operations.

## 7 AMGA data types

The AMGA metadata server supports a set of generic datatypes for all the possible backends. It is guaranteed that if you add an attribute of with one of these types that it is translated into a type supported by the database. AMGA also guarantees that a `listattr` command will return the same data type.

Other datatypes supported by a given back end can be used, however they will be not portable and it is not guaranteed that `listattr` will return the type you specified with `addattr`.

The following table lists the generic metadata data types supported by AMGA and their internal representation in the back ends:

|                     | PostgreSQL           | MySQL                | Oracle       | SQLite       | Python         |
|---------------------|----------------------|----------------------|--------------|--------------|----------------|
| <b>int</b>          | integer              | int                  | number(38)   | int          | int            |
| <b>float</b>        | double precision     | double precision     | float        | float        | float          |
| <b>varchar(n)</b>   | character varying(n) | character varying(n) | varchar2(n)  | varchar(n)   | string         |
| <b>timestamp</b>    | timestamp w/o TZ     | datetime             | timestamp(6) | unsupported  | time (unsupp.) |
| <b>text</b>         | text                 | text                 | long         | text         | string         |
| <b>numeric(p,s)</b> | numeric(p,s)         | numeric(p,s)         | numeric(p,s) | numeric(p,s) | float          |

The datatypes have the following properties:

- **int**: at least 32 bit integer value.
- **float**: 64 bits IEEE double precision floating point number.
- **varchar(n)**: A string of up to at least  $n = 254$  characters. The low limit of 254 is imposed by MySQL 4.0 and smaller version. MySQL may also truncate any trailing white space, see MySQL documentation.
- **timestamp**: Timestamp with format 'YYYY-MM-DD HH:MI:SS'.
- **text**: Long text string (2GB size limit).
- **numeric(p,s)**: SQL numeric type with precision  $p$  and scale  $s$ .

Note that SQLite not really has a strong typing system. SQLite is in fact typeless, so you can try to store anything into a given column, although internally it distinguishes between text strings and 64 bit double precision values. A column is of type text if the column type contains any of the following substrings: BLOB, CHAR, CLOB, TEXT.

The pure Python back end distinguishes strings, floats (64 bit), integers (64 bit) and timestamps.

Both pure python and SQLite will accept any type and guarantee you that the same type is returned by `listattr` with the default datatypes being numeric and string respectively.

## 8 Configuring the AMGA Server and the Replicatin Daemon

The AMGA metadata server `amgad` the AMGA replication daemon `mdservermt.config` are configured using the `amgad.config` file. By default, they share the same configuration file, but by using the `-c` command line option it is possible to specify a different configuration file for each. Nevertheless, the format of the configuration is similar for both programs.

### 8.1 Format of the Configuration File

Every line of the file contains a key value pair seperated by an equal (=) sign. Blank space is ignored around the equal sign and the end of the line. Keys may not contain any white space but white space in the value is allowed. To have white space at the beginning of a value (after the =) or at the end (at the end of the line), escape it with a backslash (\). Lines can be continued with a backslash (\) at the end of a line (white space after the \ is ignored). Comments can be put at the beginning or end of a line after a hash-sign (#). The escape character is the backslash (\). The following escape sequences exist:

- \ : Gives \
- # : Gives #
- \ () : Insert a space at beginning/end of option

The configuration file is divided into two sections. The first, consisting of settings defined in the global section, contains generic server options. The second section is called **Replication** and contains all the settings controlling replication. Sections are opened like this:

```
[SectionA]
...
```

```
[SectionB]
...
```

### 8.2 Example of a server configuration file

```
# Server options
Port=8822 # Default 8822
MinProcesses = 2
MaxProcesses = 50 # Default: 50
Mode = process # process (default) or thread mode
# initThreadNumber = 8
# MaxConnectsPerProcess = 1000000

Report = enable

# Database options
# The DataSource option must match a data source name in the odbc.ini file
#DataSource=sqlitedb
#DataSource=Oracle10g
```

```

DataSource=PSQL
DBUser=arda

# Session options
Sessions = allow          # Values are: no allow force
IdleTimeout = 1200       # Timeout for a connection/query in sec [20m default]
SessionTimeout = 86400   # Timeout for a session in the cache in sec
MaxSessionSize = 3000    # The maximum size that persistent session info can have

# Connection limits
#MaxConnectsPerUser = -1 # Maximum concurrent connections per user
#ReservedConnections = -1 # Number of connections reserved for root

# Secure Connections
UseSSL = 1

# Authentication options
RequireAuthentication = 0 # If this is off, no authentication is done!
AllowCertificateAuthentication = 1
AllowPasswordAuthentication = 1
# If you use SSL, you need to load server certificates:
CertFile = cert.pem
KeyFile = key.pem
# If Certificate based authentication allowed, you need to load server certs
# TrustedCertDir = /etc/grid-security/certificates
# AllowGridProxyLogin = 1          # Requires also AllowCertificateAuthentication

# Authorization options, choose 0 or more
#MapFile = /etc/grid-map-file      # Authorization based on certs
# Authorization based on certificates, put a list of VOMS URL and assigned users, here:
#VOMSGroups = https://lcg-voms.cern.ch:8443/voms/lhcb/services/VOMSAdmin?method=listMembers, lh
#          https://kuiken.nikhef.nl:8443/voms/picard/services/VOMSAdmin?method=listMembers,
UserDB = 1 # Authorization based on certs & passwords
#VOGroupMap =
#VirtualOrganizations = gildav(gilda) # vo1(defaultUser1), vo2, vo3(defaultUser2)...
#VOGroupMap = gildav:/gildav(gildav:users)
#VOUserMap = gildav:/gildav/Role=TrailersManager(gildav)
#MyProxyHack = 1 # Allow roles in MyProxy certificates for login

#####
# Replication settings
#####
[Replication]
# Settings for the master
EnableMaster = 0
ReplicationDaemonPort=8833

# The remaining slaves are for slave nodes
EnableSlave = 1
NodeName=gridpc1

```

### 8.3 Description of the server options

The following options are supported:

- **Port:** The number of the port the server will listen on. This can be overridden on the command line with the `-p` switch. The default port is 8822.



- **MinProcesses:** This is the minimum number of processes waiting for client connections the server must offer. When the server starts up or there are no client connections for some time, **MinProcesses** is the number of processes spawned waiting for connections.
- **MaxProcesses:** This is the maximum number of processes the server will spawn in total. The server always tries to have 1/3 of the processes in the awaiting connection state. To achieve this, the server will spawn new processes until the number of **MaxProcess** is reached. Please make sure that your database backend can support as many client connections.
- **Mode :** This define the execute mode of AMGA daemon. In process mode, each AMGA process is created and responses one client request. In thread mode, each AMGA process several has threads to response requests. The default value is process.
- **initThreadNumber :** The number of threads each process has. When the server start up **initThreadNumber** is the number of threads spawned waiting for connection if **Mode** value is thread. The default is 4.
- **Report :** AMGA contains statistics collection functionality. When the server start up AMGA sends statistics such as AMGA version, Daemon mode and so on to the collection server on kisti.re.kr. This is enabled by default.
- **MaxConnectsPerProcess:** To prevent any very rare memory leaks or other resource leaks to reduce the stability of the service, server processes can be asked to terminate themselves after serving a certain number of connections. The default is not to do this.
- **DataSource:** An ODBC database source which you need to have configured in an `odbc.ini` file. Note: you can use programs like `gODBCConfig` or `iodbcadm-gtk` to configure ODBC. The underlying database needs to be prepared for AMGA by running the `createInitial.sql` script in the `scripts` directory of the AMGA distribution.
- **DBUser:** The user with which the server will contact the database backend. The user name given in the `odbc.ini` file is the default.
- **DBPass:** The password the server will give when contacting to the database backend. The password in `odbc.ini` is the default.
- **DBSchema:** The schema within the database AMGA shall use for its own tables. Via the `import` feature it is possible to then access tables in different schemas. The default depends on the database backend and is the user-name on Oracle, the database name on MySQL and for SQLite and PostgreSQL it is the default schema of the database.
- **Sessions:** This defines whether you want to allow sessions. Sessions create an overhead on the protocol if they are enforced, so the performance of individual clients may reduce while you will be able to support more clients which share the available connections (there is a maximum of **MaxProcesses** connections, if they are all hogged by a client, then no new clients will be able to connect). Such a denial-of-service situation can be prevented by forcing sessions. Values are: `no`, `allow`, `force`. Default is `allow`.
- **IdleTimeout:** Timeout for an idle connection (that is a connection that waits for a client command) in seconds. There are no timeouts currently for database queries apart from how the database is configured. The default is 20 minutes. This is what will make your `mdclient` command line tool time out.
- **SessionTimeout:** Timeouts for session. The lifetime of a session in seconds. Default is 1 day.
- **MaxSessionSize:** Maximum size a session can have when persistently stored on the server between connections. Default is 3000 bytes. This should be sufficient for most certificates used. Rember that the total shared memory needed by the `amgad` is `Sessions*MaxSessionSize`.
- **MaxConnectsPerUser:** Maximum number of concurrent connections for a user. Default is -1: No limits.

- **ReservedConnections:** Number of connections reserved for the root user. Default is -1: No reserved connections.
- **UseSSL:** Whether the server will offer SSL as a connection protocol. This is also required to allow certificate based authentication and if you want to use passwords this is recommended if you want to be sure no one listens in. Values are 0 and 1, default is 1. Note that you cannot force the client to use an SSL connection.
- **RequireAuthentication:** Whether users need to be authenticated. Default is 0: no.
- **AllowCertificateAuthentication:** Whether you allow users to authenticate with their certificate. You will need to have CA certificates loaded for the server for this to work in order to be able to verify the client's certificate. See **TrustedCertDir**. You should also look at **User Management (p. 44)** because you will need to have a user manager module running for this to work. Default is 0: no.
- **AllowPasswordAuthentication:** Allow authentication with a password. You need a user manager module running for this to work. See **User Management (p. 44)**. Default is 0: no.
- **CertFile:** The path to the server certificate in PEM format. Preferably without encryption. Necessary to use SSL.
- **KeyFile:** The path to the private key of the server in PEM file format. Necessary to use SSL.
- **TrustedCertDir:** Path to a directory with trusted CA certificates. Needed to verify a client certificate.
- **AllowGridProxyLogin:** Whether you allow users to authenticate with a proxy certificate. Default is 0: no.

### 8.3.1 Replication Settings

The following options control the replication system. They should be defined on the replication section of the configuration file, otherwise AMGA will either not recognize them or will understand them as referring to one of the generic settings, since some names are equal.

These two settings are used by master nodes:

- **EnableMaster:** Allow this node to act as a master for replication. In particular, this setting activates saving of replication logs to the hard disk. Therefore, if the node is not going to be used as a master, this setting should be disabled for performance reasons. Default: 0. (Disabled)
- **ReplicationDaemonPort:** The port used by the replication daemon to listen for connections from slaves. Default: 8823.

The following settings are used by slave nodes:

- **EnableSlave:** Allow this node to act as a slave for replication. Leaving this setting active has no impact on performance, the only different is that it allows the use of the client side replication commands. Default: 0. (Disabled)
- **NodeName:** The identifier of the local node. Should be unique among all the nodes on the system.

For the rest of the options see **User Management (p. 44)**.

### 8.3.2 Federation Settings

The following options control the federation system.

- **ServerSideFed:** Allow the AMGA server to do actual federation work. Otherwise, the AMGA server just sends information of federated hosts to clients, and the client does the actual federation. For the detail, refer **Federation in AMGA (p. 42)**. Default: 1. (enabled)
- **AllowFedConnFail:** Allows connection failures while federating data. Default: 1 (enabled).
- **AllowFedExecFail:** Allows command execution failures while federating data. Default: 1 (enabled).

## 9 Replication in AMGA

This section presents the Replication mechanisms that are part of AMGA since version 1.2.

### 9.1 Overview

In this section we provide the technical background necessary to understand replication in AMGA. Reading this section is highly recommended for anyone interested in using this feature. Further details can be found in the following article, available on the AMGA web page:

N. Santos and B. Koblitz, *Distributed Metadata with the AMGA Metadata Catalog*, In *Workshop on Next-Generation Distributed Data Management - HPDC-15*, Paris, France, June 2006

Some more background is also explained here.

#### 9.1.1 Features Overview

Replication in AMGA follows an **asynchronous, master-slave** model, and supports **partial replication** of the directory hierarchy.

In **master-slave** replication only one of the replicas, the master, is writable, with all other replicas being always read-only. This model is sufficient for all applications with read-only metadata or where the metadata is written only at a single geographical location, which covers most of the Grid applications.

Slaves can replicate any sub-tree of the metadata hierarchy. This allows slaves to copy only the data they are interested on, reducing the load both on the slave and on the master, as well as the bandwidth requirements. If desired, a full replica of the master can be obtained by replicating the root directory.

#### 9.1.2 Concepts

Next is the description of the main concepts used in AMGA Replication.

- **Master:** Any node that exports part of its metadata hierarchy for replication.
- **Slave:** Any node that replicates metadata from a master. A node can be both a slave and a master at the same time, or neither if working standalone.
- **Replication log:** Shipped from a master to a slave, containing a metadata command plus some context information that allows the slave to replay the original command executed at the master.
- **Subscriptions:** To replicate a directory, the slave creates a subscription for that directory with the master that will allow it to receive the replication logs for that particular directory.
- **Mount:** On the slave side, the root of the sub-tree subscribed from the master is called the mount. There is a mount for each subscription and it corresponds to the path of the directory that is replicated. Each mount contains also all the sub-directories of the mounted directory.

- **Initial Synchronization:** When a slave subscribes for the first time to a directory, it must copy the existing data for which there are no replication logs.

### 9.1.3 Architecture Overview

In AMGA replication, the master ships to the slaves logs containing the metadata commands executed at the master, much in the same way as it is done in Oracle Streams and MySQL replication. When updated by a client, the master saves a replication log containing the command and some context information to the `logs` table on its database backend. A separate program, the replication daemon, queries the database periodically looking for new logs and shipping them to the slaves with subscriptions to the directories updated by the command of the log. This replication daemon is also responsible for managing the list of slaves that are replicating from the master. At the slave, the logs are replayed to update the slave's metadata. The logs contain only metadata commands, and they are totally independent of the underlying database backend.

Logs are assigned a unique sequence number at creation time, the log `xid`. Each AMGA instance generates its own identifiers in an independent way, so the `xid` is unique only inside an AMGA instance. The `xid` is necessary for synchronization between the slave and the master. For each directory replicated, the slave keeps the `xid` of the last log it received and applied, so that it knows from what point to resume after connection failures or shutdowns. The master also keeps track of the subscriptions by storing them persistently on its database backend, on the `subscribers` table. For each subscription, it stores the most recent `xid` that the slave has acknowledge. This is necessary to know when there are no more slaves waiting for a particular log, so it can be removed from the `logs` table.

**Replication Daemon** The bulk of the operations necessary for replication on master nodes are implemented outside the `amgad` process, by the `mdrepdaemon` program, also called the replication daemon. The `amgad` is only responsible for writing the replication logs to the database, while the replication daemon does everything else related to replication, including:

- Managing subscriptions
- Accepting replication requests from slaves
- Sending the initial snapshot of replicated directories to slaves
- Polling the logs table and shipping logs to the interested subscribers
- Cleaning up the logs table, removing any unnecessary logs.

The replication daemon is independent of the AMGA server it is working for, and needs only to connect to the database backend used by the AMGA server. Apart from that, there is no communication between them and they can run separately, even on different machines for better load balancing.

### 9.1.4 Operation

Nodes interested in replicating from another node must subscribe to the directories they wish to receive by contacting the replication daemon of that node. After connecting, the slave informs the master of the directories it is interested on, and begins copying the contents of the database. This is done using the dump feature of AMGA, which generates the commands that must be executed on another AMGA instance to recreate a directory hierarchy. The replication daemon internally executes a dump and forwards the commands to the slave, that replays them. Each directory is shipped using a database transaction to isolate it from updates that may be happening concurrently. It is also tagged with the `xid` of the last log generated for that directory before the synchronization started, so the slave will know from what point to start receiving logs after finishing the initial synchronization. Updates generated during the synchronization will be saved as logs and shipped to the slave after the synchronization is done.

The initial synchronization might be a lengthy process and currently there are no provisions for resuming from failed synchronizations. Nevertheless, if the synchronization is interrupted,

the slave will reestablish a consistent state by discarding all the information received. In the future, we plan to implement mechanisms to allow resuming partial synchronizations.

After having the initial snapshot, the slave can start receiving and applying logs. In the current implementation the slave connects to the master using a TCP connection, sends the `xid` of the earliest log that it wants to receive and waits for incoming logs. When the slave receives a log, it executes the log locally, and after making sure that the log is safely committed to the database, it sends an acknowledge to the master. After receiving the acknowledge, the master is free to delete the log. To ensure good performance over high-latency connections, this communication between the master and the slave is asynchronous, that is, the master sends the logs without waiting for acknowledges of previous logs, and polls the socket periodically with a non-blocking operation looking for incoming acknowledgments.

The replication daemon stores the information about the subscriptions in the local database in order for it to survive eventual crashes of the master node. This information includes the slave's id, address, directories subscribed, and the id of the last log acknowledged.

Logs must be deleted when they are no longer needed. This is done also in the replication daemon, by the same thread that monitors the logs. After polling for new logs, shipping them to subscribers, and polling for new acknowledges from clients, this thread goes over the logs table and deletes the logs that are no longer needed by any subscriber. Under normal conditions, that is, when all subscribers are connected, logs are deleted shortly after being generated. A log is kept for a longer time only when a subscriber is disconnected.

To tolerate failures resulting in disconnections of the subscribers, subscriptions are persistent, in the sense that if a subscriber disconnects without having first requested to be unsubscribed, the master will preserve the subscription and continue saving logs for the directories subscribed by the slave. When the subscriber reconnects, the subscription is resumed from the point it was interrupted. Currently, there is no provision for dealing with slaves that are disconnected for too long. In this case, the logs will accumulate on the logs table, eventually causing problems on the database backend. In the future, we plan to implement mechanisms for controlling the growth of the log table by removing old subscriptions when certain conditions are met, like the log table exceeding a maximum size or a subscriber being disconnected for too long.

## 9.2 Setup

### 9.2.1 Setting up a Master node

A master node consists of an `amga` daemon configured to save replication logs on the database plus an associated replication daemon. To simplify the setup, both processes use exactly the same format for the configuration file, allowing them to share the same configuration file.

The `amgad` is configured to save logs by setting the following property on the configuration file:

```
EnableMaster = 1
```

As an optimization, logs are saved only if there is at least one slave interested on the directory updated by the log, so if there are no subscriptions, the cost of having replication enabled is only an extra database query (to look for active subscriptions).

The replication daemon must be configured to point to the same database as the `amga` daemon. The simplest way of doing this is to reuse the same configuration file for both programs, although it is possible to have separate files. Additionally, the replication daemon needs to be assigned to a port where it will listen for incoming connections. This is configured by the `Port` property on the `Replication` section of the configuration file:

```
[Replication]
...
ReplicationDaemonPort=8823
```

This is all that needs to be configured for the replication daemon. After this, it can be started with:

```
$ mdrepdaemon [-c <amgad.config>]
```

The accepted options are:

```
-p <port>      : listen port
-d            : Activate debug output. Very verbose.
-c <configFile> : configuration file
```

If no options are given, it will look for `amgad.config` in the same way as `amgad` does.

After starting up, the replication daemon waits for connections from slaves and polls the database periodically looking for new logs. It will also go over the `logs` table periodically to delete logs that are no longer needed.

### 9.2.2 Setting up a Slave node

The slave functionality is implemented fully on the `mdserver`, and therefore, there is no need to run an external program like the replication daemon. To activate replication as a slave, the following setting must be enabled on the configuration of the `mdserver`:

```
EnableSlave = 1
```

Activating the slave mode enables the `rep_*` commands for the slave, and has no impact on the performance.

A slave node needs also to be configured with the settings of master nodes that it might be required to contact. This configuration is defined on a per-node basis, which is done through the `site_*` commands, which can be used to define sites and set their parameters like the login name the slave will use or the port of the replication daemon to contact on this site. The minimum configuration necessary is (pointing to a replication daemon listening on port 8832 and logging into it as root):

```
Query> site_add SiteA localhost:8832
Query> site_set_properties SiteA login root
```

Please consult Section **Replication (p. 23)** for further information on how to define the master nodes and **Site management (p. 22)** for how to configure sites. Once configured as a slave node, the server is started normally. There are no special command line options related to replication.

On startup, the server will consult its database backend to check if there are any subscriptions on the active state. If so, it will try to restart receiving logs from these masters, by reopening a TCP connection to each master and asking for the logs after the last one it received. If the connection can't be established immediately, it will continue to try, waiting 60 seconds between attempts, until either it succeeds or the subscriptions is stopped using the `rep_stop_receive` command.

## 9.3 Security

### 9.3.1 Replicating Security Information

There are two types of security information managed by AMGA:

- **permissions and acls** - This is associated with the metadata.
- **groups and users** - Global to the catalogue instance, and not associated directly with the metadata.

Both types of information can be replicated in AMGA. The first is replicated as part of the metadata replication. The `rep_mount` command by default imports this information from the master. In this case, the permissions and owner of the metadata on the slave will be exactly the same as in the master and will be kept synchronized with the master. This behavior can be disabled by specifying the `-noperms` option, so that the slave replicates the metadata without any security information. The resulting replica will be created using the security settings in effect at the slave at the time of replication. After a directory being mounted in a slave, it is not possible to change its configuration concerning replication of permissions and acls.

Groups and users are replicated on their own, using the commands `rep_mount_users` and `rep_umount_users`. They are handled like a virtual mount, in the sense that they have to be

synchronized initially by executing `rep_mount_users` and will be updated after establishing a connection to the master with `rep_start_receive`. The hash of the password of the user is also replicated. This allows the slave to authenticate the replicated users locally. A slave can only replicate users and groups from a single master and it should have no local users or groups of its own. The exception is the root user and any group owned by root, which are always considered local users and are not replicated.

### 9.3.2 Security at Slave Nodes

Only the root user is allowed to initiate or to control replication, that is, to execute the `rep_*` commands. For any other users, the replicated directories will be readable depending on their permissions and access control lists, which may or may not have been replicated from the master depending on the options given to `rep_mount`. A replicated directory is always read-only, regardless of the write permission or access control right.

### 9.3.3 Connections from the Slave to the Master

Connections from slaves to the replication daemon use a similar protocol as the ones from the clients to the `mdserver`, including the support for authentication and encryption. Therefore, the connection settings (e.g., use SSL, authenticate with password, certificates or grid proxies,...) used by the `mdrepdaemon` can be configured in the same way as the `mdserver`. In fact, it's even possible to use the same configuration file, in which case the connection settings will be the same. To have different settings, it's enough to specify different configuration files using the command line options of each program.

When connecting to the replication daemon, slaves may need to authenticate. Once again, this is done just like if the slave were a client connecting directly to the master. The replication daemon will accept the same credentials and users as the `mdserver` of that node. The slave node is configured in the configuration file of the slave's `mdserver`, on the [Replication] section. The slave can specify several different master nodes, with different connection settings for each. It is possible to check the list of known nodes at a slave during an interactive client session using the `rep_list_nodes` command.

### 9.3.4 Granting the replication right on the master

The master controls what groups can replicate which directories by granting the replication right. This right is granted to groups, allowing them to replicate the specified directory, including the full sub-tree rooted on it and all their entries. The right is granted using an interactive session to the master with the commands `rep_allow` and `rep_disallow`.

An important point about this right is that for replication it is the only access control performed by the master. This is easier to illustrate with an example. Suppose that user `joe` does not have permission to read or write directory `/jobs` when connected as a client, but has the replicate permission over that directory. Then, if a slave connects to this master authenticating as `joe` it will be able to replicate the full contents of `/jobs`. Once this information is on the slave, it is completely exposed to the slave's administrator. If the administrator is trustworthy, he will allow the system to enforce the original access permissions. But malicious administrators can easily expose the information to any user.

## 9.4 Limitations

### 9.4.1 Concurrent Updates on Master

When two or more clients are updating the master concurrently, there is a small probability that the slaves become inconsistent with the master. For this to happen, the clients must be writing (updating, deleting or inserting) to the same collection and to intersecting sets of rows. This is a concurrency problem, so it will happen randomly depending on the interleaving of the execution of two or more concurrent commands. This problem can be avoided in two ways:

- Prevent concurrent updates to the same directory.
- Set the transaction isolation mode of the database to serializable.

The second solution is generic but has a high overhead on the database backend. Another disadvantage is that the database backend will abort the commands that cannot be serialized due to conflicts with other concurrent clients, and this will create a new failure mode that clients might not be expecting.

#### Technical Details

The replication mechanisms of AMGA are based on the assumption that it is possible to order the write commands received by the master into a sequence that if replayed in the slave will produce the exact same results. But this is not necessarily true, depending on the transaction isolation level used by the database back-end. Most database systems use an isolation mode (Read Committed) that allows transactions to interleave in a way that makes them non-serializable. This is, there is no ordering of transactions that if replayed sequentially on a different database would produce the same results. This basically means that with the default isolation levels used by databases, there is no way of guaranteeing consistency between the master and the slave in all possible situations. Databases support stricter isolation levels, including a Serializable level. This mode will provide the guarantees needed by AMGA, but has a high overhead and therefore should be used only when the server needs to support concurrent updates.

## 9.5 Tutorial - Setting up a Master to Single Slave Setup

This section provides a step by step tutorial of how to prepare a simple replication setup. It will show how to setup replication from a master to a single slave. These are the names of the hosts that will be used:

- master - gridpc1.cern.ch, port 8822
- slave - gridpc2.cern.ch, port 8822

The tutorial assumes that the master and the slave nodes have already valid AMGA instances configured and operational, but with replication disabled. On the master, we assume the following directory hierarchy:

```
\users
\files
\files\2005
\files\2006
```

We assume that the slave is interested in replicating the `\files\2006` directory.

### 9.5.1 Configuring the master

- Configure the AMGA instance at `gridpc1` to act as a master and configure the port where the replication daemon will listen. This is done on the `mdserver.config` file, with the following properties:

```
[Replication]
EnableMaster = 1
ReplicationDaemonPort=8823
```

- Setup authentication and security for connections from slaves. For this tutorial we will require no security and accept plain text connections. Edit `amgad.config` at the server adding the following lines to the main section:

```
UseSSL = 0
RequireAuthentication = 0
```

- Configure access permissions to replication. This is done by connecting to the `mdserver` instance of the master over a normal interactive session using `mdclient`. First we create a group for users allowed to replicate. We will call it `replicators`, but any group name will do:



```
$ mdclient gridpc1
Connecting to localhost:8822...
ARDA Metadata Server 1.2.0
Query> grp_create replicators
```

- Now we must add to this group the users allowed to replicate. For this tutorial, we will create a new user called `joe` but any existing user can be used:

```
Query> user_create joe
Query> grp_adduser replicators joe
```

- Add the `replicators` group to the list of groups allowed to replicate the `/files` directory.

```
Query> rep_allow /files replicators
```

This will allow the `replicators` group to replicate all the contents and subdirectories of the `/files` directory.

- Back at the shell, start the replication daemon:

```
$ mdrepdaemon -c amgad.config
```

### 9.5.2 Configuring the slave

Now we are going to configure the AMGA instance at `gridpc2` to act as a slave. We start by editing the `mdserver.config` file.

- Activate the slave functionality:

```
[Replication]
...
EnableSlave = 1
...
```

- Configure the list of master nodes. In this case, it's only a single master, you need to be root to set up the sites:

```
Query> site_add GridPC1 gridpc1.cern.ch:8823
Query> site_set_properties GridPC1 login joe
Query> site_set_properties GridPC1 use_ssl 0
Query> site_set_properties GridPC1 authenticate_with_certificate 0
```

Note that only the login property is mandatory, the other settings are the default, anyway.

We can now start the AMGA server on the slave and connect to it using the `mdclient`. The rest of the tutorial is done from inside the `mdclient` shell.

- Create the parent directory of the directory we are going to mount.

```
Query> createdir /files
```

The slave expects this directory to exist, otherwise it will fail. We should also make sure that there is no directory with the same name as the one we are going to mount, in this case `/files/2006`.

- Mount the directory.

```
Query> rep_mount GridPC1 /files/2006
```

The slave connects to the master, subscribes to this directory and copies its contents. When the command finishes executing, we should have a local copy of the remote directory. We can check the status of the mount:

```
Query> rep_list_mounts
>> GridPC1:/files/2006 - 23628, Disconnected
```

The `Disconnected` state means that we are currently not receiving logs. The number before the state is the xid corresponding to the position of our local snapshot on the log sequence.

- Start receiving logs. The slave now must connect to the master and wait for logs. This is done with the command:

```
Query> rep_start_receive GridPC1
Query> rep_list_mounts
>> gridpc11:/files/2006 - 23645, Receiving
```

The slave will keep the connection to the Master open until it is shutdown or until we stop the subscription. If the slave is restarted, it will resume automatically the subscriptions that were active when it was shutdown. It will also try to reconnect if the connection fails for some external reason.

- Removing the subscriptions. If we want to cancel the subscription we have first to stop receiving the logs and then unmount the remote directory:

```
Query> rep_stop_receive GridPC1
Query> rep_umount /files/2006
Query> rep_list_mounts
Query>
```

## 10 Federation in AMGA

This section presents the Federation mechanisms that are part of AMGA since version 2.1.

### 10.1 Overview

The federation mechanism in AMGA provides a user a virtualized view on metadata as if one metadata server has all data which are actually distributed at multiple sites.

Federation in AMGA is very similar to mounting in NFS. That is, if a remote directory is mounted to a local directory, a user may access it as if it is located at local site. Federation actually redirects user's query to other multiple AMGA nodes and integrates the results. The list of other AMGA nodes to redirect user's query should be set properly on a per-directory basis.

AMGA provides two types of federation methods; server-side federation and client-side federation. In server-side federation, an AMGA server does the actual federation, which redirects a query and integrates results. Meanwhile, in client-side federation, an AMGA server provides a client with a list of AMGA nodes to redirect user's query and a client side does the actual federation. The client-side federation puts little overhead on the AMGA server caused by federation, but it only works with C++ APIs currently. The server-side federation may have high overhead caused by federation at the server-side, but it works with all the clients APIs. This type can be set at the server configuration as follows.

```
$ cat amgad.config
...
ServerSideFed = 1 # Federation at Server Sider
...
```

## 10.2 Setup

It is necessary to have one central AMGA server which maintains all the information about federation. That is, it needs to be configured with the settings of other AMGA nodes that it might be required to contact for federation. The configuration is defined on a per-node basis, which is done through the `site_*` commands, which can be used to define sites and set their parameters like the login name the central AMGA server will use. Please consult Section **Site management (p. 22)** for how to configure sites.

```
Query> site_add SiteA hostA:8822
Query> site_set_properties SiteA login root
Query> site_add SiteB hostB:8822
Query> site_set_properties SiteB login root
```

The next step is to initiate federation, that is, to execute the `fed_mount` commands. Only the root user is allowed to initiate or to control federation.

```
Query> fed_mount SiteA /dir
Query> fed_mount SiteB /dir
```

## 10.3 Authentication

In federation, connections to other AMGA nodes use the same protocol as the ones from clients to an AMGA server, including the support for authentication and encryption. With the client-side federation, all the necessary configuration settings should be stored at the `mdclient.config` file. And with the server-side federation, site information should hold necessary configuration settings (e.g., use SSL, authenticate with password, certificates or grid proxies,...).

## 10.4 Accessing Patterns

Accessing patterns differ depending on a command type. For example, if the "listattr" command accesses a federated directory, it is redirected to one of random federated hosts (access random type). However, in case of the "SELECT" command, a query is redirected into all the federated hosts (access all type). The following list types and corresponding commands. Other commands not listed below may result an error if they access a federated directory.

- Access random type : listAttr, addentry, addentries, stat, insert, INSERT
- Access all type : getAttr, find, listentries, ls, dir, chmod, chown, acl\_\*, removeEntries, selectattr, SELECT, DELETE, setAttr, updateattr, update, UPDATE, rm

## 10.5 Fault Handling

It is possible to have a failure while connecting and executing a command for federation. In case of "access random type" commands, it reports a success if there is any host which executed the given command successfully. If no hosts have a success, a failure is reported to a user. In case of "access all type" commands, it is allowed to define policies. It is possible to set two policies; AllowFedConnFail and AllowFedExecFail. This policies can be set at the server configuration file as follows.

```
$ cat amgad.config
...
AllowFedConnFail = 1   # Allows connection failures in federation
AllowFedExecFail = 1  # Allows execution failures in federation
...
```

AllowFedConnFail allows connection failures in federation. If AllowFedConnFail is set to 1, then it reports a success if there is any one host having a successful connection. If it is set to 0, then all federated hosts should be connectable in order not to have a failure. AllowFedExecFail is very similar to AllowFedConnFail, but it is related to execution failures. By default, these policies are set to 1.

Even though there was a success to execute a command, a user may want to know there was failures with federation. There are two commands supporting this purpose. With these commands, it is possible to check partial failures when the policy is set to ignore the connection or execution failures. The `fed_policy` command shows the current policy about federation. The first line of the result shows whether connection failures are allowed in federation, and the second line shows whether execution failures are allowed. The `fed_check_failure` command shows whether last command had failures. The first line of the result shows whether connection failures happened, and the second line shows whether execution failures happened.

```
Query> fed_policy
>> 1
>> 1
Query> SELECT * FROM /dir
>> ...
Query> fed_check_failure
>> 1
>> 0
```

## 10.6 Limitations

The current limitations related to federations are as follows

- No Schema Heterogeneity is allowed
- Federating Directories with different Names is not allowed
- No Distributed Join is allowed
- No Transaction is Supported

## 11 User Management

The standalone AMGA server comes with a powerful system to manage users as well as to control access to entries and metadata. If AMGA is run as an add on to a file catalogue, however, these features are not available and the access controls of the file catalogue is used instead.

To understand the user management of the AMGA server it is necessary to know that the server does not really manage users but only their authentication and authorization. When changing the owner of an entry for example, the server does not check that this owner exists. Users are only relevant for logging in. This allows to manage users outside of the server, e.g. in a VOMS.

**IMPORTANT NOTE:** When you use sessions or connections, then changes to the way a user is authenticated or how the authorization is done through the mapping to an AMGA user, will not affect active sessions or connections. **You must restart the server to propagate changes to the user management to active sessions or connections.**

### 11.1 Configuration

To use the metadata service, a user must be authenticated and authorized. Authentication can be done via a certificate or a password, see **Configuring the AMGA Server and the Replicatin Daemon (p. 31)**. After the authenticity of a user is established in the handshaking of the client with the server, the client needs to be authorized to use the role of a certain user. Authorization is optional, if authorization is not enabled for the server, any authenticated user can assume any role he wishes. Authorization is controlled via the `mdserver.config` configuration file:

```
# Authorization options, choose 0 or more
# MapFile = /etc/grid-map-file          # Authorization based on certs
# Authorization based on certificates, put a list of VOMS URL and assigned users, here:
VMSGGroups = https://lcg-voms.cern.ch:8443/voms/lhcb/services/VOMSAdmin?method~=listMembers, lh
              https://kuiken.nikhef.nl:8443/voms/picard/services/VOMSAdmin?method~=listMembers,
UserDB = 1 # Authorization based on certs & passwords
```

Authorization can be done via certificates or passwords (password authentication actually includes authorization), both must be explicitly enabled. **For authentication via certificates to work, both the server and the client must have SSL enabled (UseSSL).** Four ways are foreseen to accessing the necessary information to match user names with their credentials, one or more must be enabled for the `RequireUserAuthorization` to work:

- A grid-map file mapping certificate subjects, that is distinguished names (DN of users) to users. This is a static setup and no new users can be added at runtime. No password authorization is possible via a grid map file. Option `MapFile`.
- A user database using the database backend. This allows creation of users and the management of their credential at runtime. This is the only option which allows password based login. Option `UseDB`.
- Authorization using a VOMS. All users registered with a VO will be assigned to the user specified here. You can give several VOMS-URL user pairs here. Option `VOMSGroups`.
- Authorization via VOMS certificates. All users connecting with a VO information-enriched certificate obtained via `voms-proxy-init` will be assigned to specific AMGA users depending on the role within the VO. Option `VirtualOrganizations`. Note that only the DB based user management module is able to make changes to the user setup. If you have several user management modules activated at the same time, then listing users and checking their credentials for authorization will go through the users in all of the modules. A user is authorized as soon as he has been found in **any** of the modules.

## 11.2 Management via a Grid Map file

You can give a location of a grid map file using the `MapFile` option for user authorization. This file contains pairs of distinguished names and user names. The DN must be enclosed in double quotes and must be in the form where its fields are separated by commas on one line (output of `openssl x509 -subject -in usercert.pem -nameopt oneline -noout`):

```
$ cat mapfile
"/C=CH/O=CERN/OU=GRID/CN=Birger Koblitz 9904" koblitz
```

There are no wild cards currently allowed. The map file will be read only once at server startup. It is not possible to add or change users using the command line tool.

## 11.3 Management of Users using the database backend

To enable user management using your database backend, you need to enable this feature by setting `UserDB = 1`. If you have run the `createInitial.sql` script, to set up you database, the necessary tables have already been created. You can now manage users via the `mdclient` command line tool:

- **user\_list:** Lists all users known to the authentication subsystem.
- **user\_listcred user:** Lists the credentials with which the user can be authenticated. Returns first the user name, then whether there is a password and then the certificates which are mapped in a Grid-Mapfile and via the user database. Finally the different VO and VO roles which allow you to become that user are listed. Only root is allowed to see the credentials of other users.
- **user\_create user [password]:** Creates a new user and assigns a password if give. This command is for the root user only. Only hashes of passwords are stored in the database backend. Charaters that should be avoided in the user name are `'`, `"`, `\`, `:`, `|`, and `<space>`
- **user\_remove user:** Deletes a user. This command is for root only. It does not check whether there are still files or directories owned by that user.
- **user\_password\_change user password:** Changes the password of a user. Only root can change the password of any user. Non-privileged users may only change their own passwords.

- **user\_subject\_add user subject:**Adds a certificate identified by its subject line to be used to authenticate a user. While every user can only have one password. Several certificates can point to the same user. Remember that in order to have spaces in the subject, you need to enclose it by single quotes ("). See **Definition of the Client Server Protocol (p. 53)**. The form of the subject needs to be the one where parts are separated by commas as in the output of e.g. `openssl x509 -subject -in usercert.pem -nameopt oneline -noout`.
- **user\_subject\_remove user subject:**Removes a certificate from the list of certificates which allow to login as a certain user.

## 11.4 Management via a VOMS

Giving pairs of VOMS member list URLs and user names in the `VOMSGroups` option, you can assign all members of a VO to a user (role would be the better word here).

## 11.5 Management via VO-Certificates

You can allow users to log in with VO-enabled certificates by using the `VirtualOrganizations` option and assigning it a list of `VO(default_user)` definitions. By enabling `MyProxyHack` this works also with certificates issued by a MyProxy server. The `VOGroupMap` and `VOUserMap` options allow to map VO groups to AMGA groups and special VO roles to AMGA users with the syntax used by `VirtualOrganizations`. You can also manage mapping between a user/group & a VOMS Role/group in the command line tools.

- **user\_voms\_add** - Allows user to log in with a certificate with the given VOMS Role
- **user\_voms\_list** - Lists all possible VOMS Role
- **user\_voms\_remove** - Removes VOMS Role allowed for user
- **grp\_voms\_add** - Creates mapping between a group and a VOMS group
- **grp\_voms\_list** - Lists all possible VOMS group mapping
- **grp\_voms\_remove** - Removes group mapping allowed for VOMS Group

## 12 Users, Groups and ACLs

The standalone AMGA server comes with a powerful system to manage users as well as to control access to entries and metadata. If AMGA is run as an add on to a file catalogue, however, these features are not available and the access controls of the file catalogue is used instead.

The permission schema tries to copy the semantics of POSIX APIs. Some of the semantics are different from the POSIX semantics for a file system as AMGA is a metadata catalogue. As an example, it is necessary to have the 'x' permission for a directory to read the attribute list, while 'r' permissions for any file are necessary to read the values of the attributes for a file. The exact behaviour is described together with the respective commands.

### 12.1 Users

The size of a username is limited to 64 lower-case latin alphabet characters.

### 12.2 Groups

Any user can create groups. Group names are scoped with the name of the user creating them. A fully qualified group name has the form `user:groupname`. If the user scope of the group is the current user, it does not need to be specified in a command. The size of groupname is limited to 64 lower-case latin alphabet characters.

A special group exists and is maintained by AMGA internally, the `system:anyuser` group which contains automatically any user which is authenticated to the system. Using this group it

is possible to emulate the permissions for 'other'-users in a Unix filesystem which are missing in AMGA.

The following commands can be used to manage groups:

- **grp\_create groupname:** Creates a new group with name groupname. It is not possible to create groups belonging to others.
- **grp\_delete groupname:** Deletes a group with name groupname. Only root can delete groups of other users.
- **grp\_show groupname:** Shows all the members belonging to group groupname. You can only look into groups of which you are a member or your own groups. Root can list all groups.
- **grp\_adduser groupname user:** Adds a user to a group. Only owners of a group or root can change group membership.
- **grp\_removeuser groupname user:** Removes a user from a group. Only owners of a group or root can change group membership.
- **grp\_member [user]:** Shows to which groups a user belongs. Only root can ask this question for other users.
- **grp\_list [-a] [user]:** Shows the groups owned by user, by default the current user. If the -a option is given, all groups are shown.

### 12.3 Access Control Lists

ACLs (Access Control Lists) can be assigned to any directory.

The following commands exist to manipulate ACLs of a directory.

- **acl\_add directory group rights**
- **acl\_remove directory group:** You can use the "\*" to remove all ACLs of a directory.
- **acl\_show directory**

On MySQL5 or PostgreSQL you can create directories with the "acls" option, which will allow you to put ACLs also on individual files.

### 12.4 The sudo command

Since AMGA 2.0 a sudo command exists, which allows the root user to become any other user. The syntax is: `sudo <user>`

### 12.5 The sticky bit

Since AMGA 2.0 a sticky bit exists, which allows an entry inside a directory to be renamed or deleted only by the entry's owner, or superuser when it is set. To set a sticky bit on a directory, use `chmod` command with permission "t".

```
Query> chmod /test rwt
```

## 13 Accessing preexisting databases in AMGA

It is possible to access previously existing databases through AMGA. This is done by importing them into AMGA's information schema. The importation will just create links to the imported table, and does not copy the table.

The first step is to configure an Access through ODBC to your existing database. Please follow the instructions for UnixODBC or tools like Tora for this.

The next step is to install the AMGA server on a system with direct access to the backend network. You will need to add the AMGA support tables to the existing database by applying

the suitable createInitial.sql script for your database, you can find in the config directory of the source distribution or in PREFIX/share/doc of the binary distributions. Should there be clashes of the namespace, you can try to solve them using different schema names, provided your RDBMS supports this.

Once this has been achieved, set up the amgad server to point to your database (see **Configuring the AMGA Server and the Replicatin Daemon (p. 31)**). Then connect the client using the AMGA root-user (see **Configuration of the C++ and Java command line clients (p. 10)**) and start importing the tables using the `import tablename dir` command (see **Metadata Access from the Shell (p. 13)**). This will make the table with the given tablename known to AMGA as the name given in the "dir" option. You can set access permissions to the new table via the `chmod/chown` commands or via the ACL system. You can also now use AMGA to alter the table-schema.

After reconnecting as an ordinary user, you can start using the imported tables. Since they do not have the supporting information normal AMGA tables have, only the standard SQL functionality works, namely insertion of entries, selecting (including joins to other tables) and updating and deleting. The `ls`-command is (currently) not supported.

The following is an example using a PostgreSQL database in which we create a simple table called test.

```
$ psql -U arda metadata
Welcome to psql 8.1.8, the PostgreSQL interactive terminal.
metadata=> CREATE TABLE test(item integer, name varchar(256));
CREATE TABLE
metadata=> INSERT INTO test (item, name) VALUES (1, 'apple');
INSERT 0 1
metadata=> INSERT INTO test (item, name) VALUES (2, 'pear');
INSERT 0 1
```

Setting the user to "root" in mdclient.config, we can connect to the AMGA server to import the generated table:

```
$ mdclient
Connecting to localhost:8822...
ARDA Metadata Server 1.2.10
Query> whoami
>> root
Query> import test /atest plain:main=item
Query> ls -ld /atest
>> rwxrwxr-x root /atest
Query> addattr /atest i int
Query> listattr /atest
>> item
>> int
>> name
>> varchar(256)
>> i
>> int
Query> quit
```

As shown, it is possible to investigate the table schema, and even alter it.

Now edit mdclient.config so that you log in as an ordinary user, and you can use the imported table:

```
Query> selectattr /atest:name /atest:item ''
>> apple
>> 1
>> pear
>> 2
Query> selectattr /atest:name /atest:item '/atest:item=2'
>> pear
```



```
>> 2
Query> rm /atest 'item=2'
Error: 4: Permission denied.
```

As you can see, AMGA has protected the table using its access control mechanisms.

When you import tables, AMGA allows you to use options to describe which properties the table has. AMGA will assume that the table has the following properties, which must be present in a way that AMGA can understand and use them, that is they must be implemented in the standard way AMGA does: shared: The table can be later shared by other directories created "under" it.

- `acl`: The table has ACLs.
- `plain`: The table is "plain" that is there are no permissions, or other AMGA specific extensions.
- `file`: This is a table with the file extensions: `guids`, `sizes` etc. `inherit`: The table inherits its structure from a parent table (only PostgreSQL).
- `perms`: The table has posix permissions.
- `type=< type >`: The data type of the main column (ignored)
- `main=< main >`: The main column.
- `table=< table >`: Ignored, this is the name of the table itself.

## 14 Using the C++ Client API

There are two different C++ client APIs available for the AMGA metadata service. One is through the `md_api` which provides several api functions, the other is directly through the `MDClient (p. ??)` class which also serves as a backend to the `md_api`.

The `MDClient (p. ??)` class offers an interface which allows to issue AMGA commands directly but does not understand the semantics of the commands and thus does not parse the responses of the server into suitable structures, while this is done by the `md_api`. However, the control on the connection to the server is much better in the case of the `MDClient (p. ??)` class, for example it allows you to abort a query easily. It may also happen that some commands are not available in the `md_api` yet.

In any case, both ways to access the metadata service from C++ depend on an existing and accessible `mdclient.config` file being either in the current working directory or in the home directory as `~/mdclient.config`. See **Configuration of the C++ and Java command line clients (p. 10)** for explanations how to set up the client configuration.

The following is an example of a program using the `md_api` to

```
#include "client/md_api.h"
#include <iostream>

int main (int argc, char *argv[])
{
    std::cout << "Listing attributes of /test\>";
    std::list< std::string > attrList;
    std::list< std::string > types;
    if( (res=listAttr("/test", attrList, types)) == 0){
        std::cout << "  Result:" << std::endl;
        std::list< std::string >::iterator I=attrList.begin();
        while(I != attrList.end())
            std::cout << "  >" << (*I++) << "<" << std::endl;
    } else {
        std::cout << "  Error: " << res << std::endl;
    }
}
```

```

std::cout << "Getting gen and events attributes of /test/*\n";
AttributeList attributeList(2);
std::list< std::string > attributes;
attributes.push_back("gen");
attributes.push_back("events");
if( (res=getAttr("/test/*", attributes, attributeList)) == 0){
    std::cout << " Result:" << std::endl;
    while(!attributeList.lastRow()){
        std::vector< std::string > attrs;
        std::string filename;
        attributeList.getRow(filename, attrs);
        std::cout << "File: >" << filename << "<" << std::endl;
        for(size_t i=0; i< attrs.size(); i++)
            std::cout << " >" << attrs[i] << "<" << std::endl;
        std::cout << std::endl;
    }
} else {
    std::cout << " Error: " << res << std::endl;
}

return 0;
}

```

A full overview of the available API functions is given at the following url.

[http://project-arda-dev.web.cern.ch/project-arda-dev/metadata/md\\_\\_api\\_8cc.html](http://project-arda-dev.web.cern.ch/project-arda-dev/metadata/md__api_8cc.html)

```

#include <MDClient.h>
#include <iostream>

int main (int argc, char *argv[])
{
    int res;

    MDClient client;
    // client.setDebug(true);

    if(client.connectToServer()){
        std::cout << client.getError() << std::endl;
        return 5;
    }

    std::string command="pwd";
    if( ( res=client.execute(command)) ){
        std::cout << " ERROR: execute failed"
            << " (" << res << "): "
            << client.getError() << std::endl;
        return res;
    }

    while(!client.eot()) {
        std::string row;
        if(res=client.fetchRow(row)){
            std::cout << "Error fetching: " << res << std::endl;
            return res;
        }
        std::cout << row << std::endl;
    }
}

```

```

return 0;
}

```

All capabilities of the **MDClient** (p. ??) like cancellation of requests or the catching of CTRL\_C are explained in the reference at <http://project-arda-dev.web.cern.ch/project-arda-dev/metadata/class\MDCClient.html> a short(!) example of how to make use of them is the `mdclient.cc` program itself.

## 15 Using the Java Client API

The AMGA Java API is distributed in two forms. As an RPM and as a tar ball. The tar ball is provided so that the Java API can be used in other platforms other than Linux, including Windows and MacOS.

To use the Java API it is necessary to include the `glite-amga-api-java.jar` file in the classpath. If the Java API was installed from the RPM, then this file is typically located at `<GLITE_HOME>/share/java`, where `<GLITE_HOME>` is the base directory where the gLite software is installed (typically, `/usr`). If the Java API was installed directly from the tar ball available on the AMGA Web Site (p. ??), the `glite-amga-api-java.jar` is located on the top level directory to where the tar ball was unpacked.

A jar can be included in the classpath in two ways: by setting the `CLASSPATH` environment variable or by using the `-classpath` option in the command line arguments when running `java`.

To set the classpath variable:

- Unix (bash): `export CLASSPATH=./glite-amga-api-java.jar`
- Windows: `set CLASSPATH=./glite-amga-api-java.jar`

After setting the `CLASSPATH`, to run a Java program it is only necessary to do the following to run a class called, for instance, `QueryMetadata`:

```
java QueryMetadata
```

To specify the jar file directly on the command, one must do:

Unix:

```
java -classpath ./glite-amga-api-java.jar QueryMetadata
```

On Windows the command line is similar, except the path separator is `;` instead of `:`.

The Javadocs for the Java API can be found here <http://project-arda-dev.web.cern.ch/project-arda-dev/metadata/java>.

Like the C++ API, the Java API can be used in two ways. Either through the higher-level interface exposed by the class `arda.md.javaclient.MDClient` or by sending the commands directly to the server using the low-level API in `arda.md.javaclient.MDServerConnection`. Next are the two examples given for the C++ client API rewritten using the Java API. The first uses the higher-level Java API.

```

import java.io.IOException;
import arda.md.javaclient.*;

public class MDJavaAPI {

    public static void main(String[] args) throws IOException {
        MDServerConnection serverConn = new MDServerConnection(
            MDServerConnectionContext.loadDefaultConfiguration());
        MDClient mdClient = new MDClient(serverConn);

        System.out.println("Listing attributes of /test");
        try {
            AttributeDef[] attrs = mdClient.listAttr("/test");
            System.out.println("Result: ");

```

```

        for (int i = 0; i < attrs.length; i++) {
            System.out.println(" >" + attrs[i].name + ":" + attrs[i].type);
        }
    } catch (CommandException e) {
        System.out.println("Error: " + e.getMessage());
    }
}

System.out.println("Getting gen and events attributes of /test");
try {
    String[] keys = {"gen", "events"};
    NamedAttributesIterator attrs = mdClient.getAttr("/test", keys);
    while (attrs.hasNext()) {
        NamedAttributes entry = attrs.next();
        System.out.println("File: " + entry.getEntryName());
        String[] keys1 = entry.getKeys();
        for (int i = 0; i < keys1.length; i++) {
            System.out.println(" >" + keys1[i] + "=" + entry.getValue(keys1[i]));
        }
    }
} catch (CommandException e) {
    System.out.println("Error: " + e.getMessage());
}
}
}
}

```

The following example uses the low-level API directly:

```

import java.io.IOException;
import arda.md.javaclient.*;

public class DirectServerConnection {

    public static void main(String[] args) throws IOException
    {
        // Loads default configuration and connects to server
        MDServerConnection serverConn = new MDServerConnection(
            MDServerConnectionContext.loadDefaultConfiguration());
        try {
            serverConn.execute("pwd");
            while (!serverConn.eot()) {
                String row = serverConn.fetchRow();
                System.out.println(">" + row);
            }
        } catch (CommandException e) {
            System.out.println("Error executing command: " + e.getMessage());
        }
    }
}
}

```

## 16 Using the Python Client API

The Python client for AMGA is distributed in the in the `glite.amga.api-python` RPM. After installation the `amga` package is available with the `mdclient` and `mdinterface` modules. The `mdclient` class offers an interface similar to the `MDClient` (`p. ??`) interface in C++ plus methods for most AMGA command. All arguments are automatically quoted before being sent to the server.

The following is an example script which creates a directory, cd's into it and then gets the "sin" and "events" attributes of all entries in the directory.

```
#!/usr/bin/env python

#import amga classes
from amga import mdclient, mdinterface

#instantiate an AMGA client connecting to localhost:8822 as 'guest'
client = mdclient.MDClient('localhost', 8822, 'guest')

try:
    print "Creating directory /pytest ..."
    client.createDir("/pytest")
except mdinterface.CommandException, ex:
    print "Error:", ex

try:
    print "cd /pytest"
    client.cd("/pytest")
except mdinterface.CommandException, ex:
    print "Error:", ex

try:
    print "Getting all attributes of the files in /pytest..."
    client.getattr('/pytest', ['sin', 'events'])
    while not client.eot():
        file, values=client.getEntry()
        print "->",file, values
except mdinterface.CommandException, ex:
    print "Error:", ex
```

## 17 Definition of the Client Server Protocol

The protocol is a streamed ASCII protocol which is line oriented. Three bytes are special control characters: \n (012) is the line-ending byte which needs to be attached to any line, EOT (004) is the end of transmission sent by the server after any response because server responses can have many (also empty) lines and CAN (030) is the cancel byte which can be sent out-of-band by the client to abort the request or inline in the servers response if during response processing an error occurs.

This defines the full client server protocol including the handshaking with four example commands(first one OK, second cancelled by user, third has an execution error, fourth is cancelled by the server):

| SERVER                                  | CLIENT                    |
|-----------------------------------------|---------------------------|
| Greeting\n                              |                           |
| Protocol <protocol number>\n            |                           |
| <space sep. list of serv opts>\n        | <requested serv option>\n |
|                                         | <more opts>\n             |
|                                         | <...>\n                   |
|                                         | \n                        |
| OK\n                                    |                           |
| ----- SSL handshaking if required ----- |                           |
| SSL_accept()                            | SSL_connect()             |
| -----                                   |                           |
|                                         | <command>                 |
| 0 <number of columns>\n                 |                           |
| <line 1 of response>\n                  |                           |
| ...                                     |                           |

```

<line n of response>\n
EOT[<session-id>]EOT
<command>

<error-code> <literal err.>\n
...
CAN (Out-of-Band!)
CAN
<abort error code>\n
EOT[<session-id>]EOT
<command>

<err-code> <comment>\n
EOT[<session-id>]EOT
<command>

0\n
<line 1 of response>\n
...
CAN (timeouts, back end-error...)
<err-code> <comment>\n
EOT[<session-id>]EOT

```

Remarks: The greeting is a free-form greeting string sent by the server. The client cannot make any assumptions about the content, apart that it ends in a version number, which may contain dots. The intention is to allow the client to display this greeting. The client should only depend on the protocol version number, which is integer for any assumptions on the protocol or the available calls.

As of protocol version 2 the server options can include "plain", "ssl" and "statistics" giving the possible ways for the connection security/encryption. The statistics option, used for monitoring is explained in the **Monitoring support in AMGA (p. 55)** section.

The requested server option sent by the client is either "plain" or "ssl", requesting an ssl connection. Alternatively the client may resume a session by responding "resumeSSL<session-ID>" or "resume<sessionID>". The other options sent one by one on a line consist of an integer number and the value for that option:

- 0 <login name>="">
- 1 <Plain text="" full="" name>="">
- 2
- 3 <group mask>="">
- 4 <current directory>="">
- 5 <password> Where the login name is the requested user id string, the full name is entirely optional and should contain the full user name of the user as in the comment field of /etc/passwd. The group and user permission masks are 3 character fields like umask. All of these fields are optional, unless e.g. the server requires that particular user to use a password for authentication.

In case a session is resumed, the server ignores any additional options requested by the client, and will expect a command next if no SSL session is requested, that is the "OK\\n" will only be sent if SSL is being used. This reduces round-trips in case of No-SSL, but the OK is necessary to synchronize SSL startup in case of an SSL session.

Any authentication failure is sent by the server as the response to the first command, to save round-trips in the case of sessions (see below) and because the result of the authentication may only be known after the initialization of the SSL session if the certificate is being used for authentication.

Commands are strings terminated by a newline as explained in the **Metadata Access from the Shell (p. 13)** section . Sending a wrong command does not violate the protocol, but results in an error "3 Illegal command". Three special commands are part of the protocol:

- quit Asks the server to close the connection.

- exit The same as quit.
- close Ask the server to send a session ID and then close the connection.

Every command is answered by an integer error code, which is 0 in case of everything OK. Error codes are not part of the protocol, but part of the responses to commands.

If sessions are forced on the server, then the server will always send a sessionID at the end of any response to the command and then simply close the connection. This also saves round-trips.

The CAN bytes to interrupt the streaming of a result sent by the server are sent in a TCP out-of-band message. However if the server decides to interrupt the streaming of a response the CAN byte is sent in-band.

## 18 Monitoring support in AMGA

AMGA supports monitoring via the <http://sourceforge.net/projects/monami/> MonAMI monitoring service which can also forward information into MonALISA. The following is an example script in Python which requests information from an AMGA server. The server is contacted on the standard port:

```
#!/usr/bin/env python
import socket

# Open TCP socket to AMGA server
s = socket.socket(socket.AF_INET, socket.SOCK_STREAM)
s.connect(("localhost", 8822))

# Request statistics
s.send('statistics\n\n')

# Read response
result = ''
while 1:
    r = s.recv(1000)
    result = result + r
    if not r:
        break

start = result.find('<Service>') # Skip response header
result = result[start:]
print result
```

The server is contacted and monitoring is requested via the "statistics\n\n" command. In the response of the server the server greeting needs to be skipped. The rest of the result is then an XML encoded status report according to the GLUE schemas as described [here](#).

The following is an example result:

```
<Service>
<Name>AMGA</Name>
<Version>1.1.0</Version>
<Data>
  <Key>MaxConnections</Key>
  <Value>50</Value>
  <Key>PreparedConnections</Key>
  <Value>1</Value>
  <Key>UsedConnections</Key>
  <Value>1</Value>
  <Key>MaxSessions</Key>
  <Value>1024</Value>
  <Key>UsedSessions</Key>
```

```
<Value>0</Value>
<Key>SessionStorage</Key>
<Value>shm</Value>
</Data>
</Service>
```

You can use a Nagios probe for AMGA. You can find this probe named `check_amga` in the `scripts` directory. This probe needs a `perl-XML-Simple` package. You must install this package first.

```
$ sudo yum install perl-XML-Simple
```

You can test as follow.

```
$ cd org.glite.amga.server/scripts
```

```
$ ./check_amga
```

Return :

```
OK: Success|0.001s;10;0
```

or

```
CRITICAL: Failed to connect AMGA Server (localhost:8822)
```

## 19 Multi-threaded AMGA Server

AMGA is implemented as a multi-process server, each process with one backend DB connection. On receiving a request from clients, if there is a pre-forked AMGA server process that is idle, it will pick up the request and take care of it. If there is no idle process available, another process is supposed to be forked and create a new backend DB connection to handle the request. This one process per client approach works well unless thousands of clients try to have access to the server at the same time. In case of that thousands of server processes should be able to run on the server machine, with each holding one DB connection. Due to the fact that DB connection is a system resource that is very expensive for some backend database systems such as Oracle, thousands of DB connections are certainly not permitted to open. To deal with this kind of situations, in AMGA 2.0, users are allowed to configure a multi-threaded mode that allows one server process to hold multiple threads that share one DB connection to its backend database system. With this approach, when a client request comes to the server, a thread is supposed to pick up the request. Since one AMGA process is set up to have multiple threads, there is no need for thousands of process with thousands of DB connection to be forked to handle thousands of concurrent requests from clients on the Grid.

### 19.1 Configuration

In order to configure the AMGA server in multi-threaded mode, you should set the following option to the `amgad.conf` file. More detailed information about AMGA configuration is described in the chapter "Configuring the AMGA Server and the Replication Daemon".

```
Mode = thread
initThreadNumber = 4
```

The "Mode" is the option variable for users to choose process or thread mode of AMGA daemon. The "initThreadNumber" is used to indicate the number of threads a process maintains. In the above example, each process will have 4 threads that share one DB connection. This number should be configured appropriately on an application basis.

## 20 SOAP Interface

The AMGA implementation uses TCP streaming to communicate between client and server which shows a very promising performance. To meet the EGEE requirements, we have also implemented an alternative SOAP-based frontend.



The interface provide by the SOAP frontend is compliant to OGF WS-DAI & WS-DAIR specification.

- **WS-DAI** : "Web Services Data Access and Integration - The Core (WS-DAI) Specification, Version 1.0", <http://www.ogf.org/documents/GFD.74.pdf>
- **WS-DAIR** : "Web Services Data Access and Integration - The Relational Realization (WSDAIR) Specification, Version 1.0", <http://www.ogf.org/documents/GFD.76.pdf>

## 20.1 Installation

### 20.1.1 SOAP client installation

To install the command line soap client tools and C++ api, you need to install `glite-amga-cli` and `glite-amga-soapclient` packages. By default the packages will be installed in `/usr`.

```
$ rpm -i glite-amga-cli-{version}.{platform}.rpm
$ rpm -i glite-amga-soapclient-{version}.{platform}.rpm
```

The configuration file(`mdclient.conf`) used in the soap client tools are the same as one used in the `mdclient`. If you want to change the configuration, you should go on and read **Configuring Configuration of the C++ and Java command line clients (p. 10)** .

### 20.1.2 SOAP server installation

The AMGA SOAP server depends on the AMGA server package. To install the soap server, you need to install `glite-amga-server` and `glite-amga-soapserver` packages.

```
$ rpm -i glite-amga-server-{version}.{platform}.rpm
$ rpm -i glite-amga-soapserver-{version}.{platform}.rpm
```

To setup a database and ODBC driver, refers **Server installation (p. 6)** . You should now be able to start the service and verify whether it is running by doing

```
$ /etc/init.d/mdsoapservice start
$ tail /var/log/amgasoap.log
```

The AMGA SOAP server is configured using the `mdsoapserver.config` file. The options are almost same as the options in `amgad.config` file. For the configuration on `amgad.config` file, you should go on and read **Configuring the AMGA Server and the Replicatin Daemon (p. 31)** . For a soap server, the following additional options are supported.

- **MaxConnections** - This is the maximum number of DB connections shared by AMGA soap threads
- **WorkerThreads** - This is the number of threads waiting for client connections
- **SoapKeepAlive** - HTTP KeepAlive enabled/disabled
- **IndirectServiceLimit** - This is the maximum number of indirect services/resouces a single user can create. If it is set to -1, then there is no limitation
- **SoapMaxBuffer** - This is the maximum buffer space assigned for a single soap message. If the user's queried result exceeds the limit, the rest of data are just ignored
- **CertKeyFile** - This is just a concatenated file of two files (CertFile and KeyFile)

## 20.2 How to use SOAP client tools

You can access the metadata catalogue with the soap client tools. These tools are:

- **wsdair\_access\_execute** - an interactive command line shell with the soap server
- **wsdair\_access\_cli** - an utility to submit a single command to the server
- **wsdair\_access\_factory** - an utility to create an indirect data resource with the given command
- **wsdair\_list** - an utility to list all the data resources including indirect data resources created by **wsdair\_access\_factory**
- **wsdair\_destory** - an utility to remove an indirect data resource
- **wsdair\_rowset** - an utility to retrieve dataset from the given indirect data resource
- **wsdair\_property** - an utility to show the property of a data resource

The **wsdair\_access\_execute** tool is very similar to **mdclient** tool. The following is an example with the the **wsdair\_access\_execute**. The default port number SOAP server uses is 8833. However, you can change the port number you want to connect with "-p" option.

```
$ wsdair_access_execute -p 8833
Query> ls
/STAFF
/PROJ
/WORKS
Query> SELECT * FROM STAFF
STAFF.FILE
STAFF.EMPNUM
STAFF.EMPNAME
STAFF.GRADE
STAFF.CITY
S1
E1
Alice
12
Deale
S2
E2
Betty
10
Vienna
```

The **wsdair\_access\_cli** tool is also similar to **mdcli** tool, which allows you to directly issue metadata commands on the shell. The following is an example with the the **wsdair\_access\_cli**.

```
$ wsdair_access_cli -p 8833 "SELECT * FROM WORKS"
WORKS.FILE
WORKS.EMPNUM
WORKS.PNUM
WORKS.HOURS
W1
E1
P1
40
W2
E1
P2
20
```

W4  
E1  
P4  
20  
W5  
E1  
P5  
12

The `wsdair_access_factory` tool creates an indirect data resource with the given command. When creating "indirect data resource", a user does not receive the results in the response to a request made, and the results is made available to the user indirectly as a new data resource. The following is an example with the the `wsdair_access_execute`.

```
$ wsdair_access_factory -p 8833 "SELECT * FROM WORKS"  
http://localhost:8833/SQLResponse/responses186_4  
DRAN : http://localhost:8833/SQLResponse/responses186_4
```

The newly created indirect data resource can be accessed with `wsdair_rowset` command. The following example retrieves two rows from the indirect data resource with `wsdair_rowset` command. "-P" option defines the location where the retrieval starts and "-C" option is the number of rows to be retrieved.

```
$ wsdair_rows -p 8833 -P 0 -C 2 http://localhost:8833/SQLResponse/responses186_4  
/SQLResponse/responses186_4.FILE  
/SQLResponse/responses186_4.EMPNUM  
/SQLResponse/responses186_4.PNUM  
/SQLResponse/responses186_4.HOURS  
W1  
E1  
P1  
40  
W2  
E1  
P2  
20
```

The `wsdair_list` lists all the data resources including indirect data resources created by `wsdair_access_factory`.

```
$ wsdair_list -p 8833  
Address : http://localhost:8833/SQLResponse/responses186_4  
DRAN : http://localhost:8833/SQLResponse/responses186_4  
Address : http://localhost:8833/SQLAccess/Metadata  
DRAN : http://localhost:8833/SQLAccess/Metadata
```

The `wsdair_destory` utility removes an indirect data resource. The following is an example with the the `wsdair_destory`.

```
$ wsdair_destory -p 8833 http://localhost:8833/SQLResponse/responses186_4  
  
$ wsdair_list -p 8833  
Address : http://localhost:8833/SQLAccess/Metadata  
DRAN : http://localhost:8833/SQLAccess/Metadata
```

## 20.3 WS-DAIR Architecture

### 20.3.1 Data access service and data resources

WS-DAI defines that a data access service is a web service that implements one or more of the WS-DAI specified interfaces to provide access to data resources. A single standalone service manages all the WS-DAI & WS-DAIR interfaces in AMGA WS-DAIR.

There are two types of data accesses defined in WS-DAI. The first one is "direct data access", which means that a consumer receives a direct response, containing the requested data, following a request made to a data access service. The second is "indirect data access". It means that a consumer does not receive the results in the response to a request made to a data access service. The request to access data will be processed by the data access service and data resource, with the results being made available to the consumer indirectly as a new data resource, often through a different data access service that may support a different set of interfaces.

To support these two data access types, there are two types of resources in AMGA WS-DAIR. The first one is **SQLAccess (p. ??)** type, which is the data source to support "direct data access." In AMGA WS-DAIR, it is possible to have only one **SQLAccess (p. ??)** resource; a relational database which is managed by AMGA for a metadata catalogue service. A consumer accesses it through the **SQLAccess (p. ??)** interface. And the second one is a **SQLRowset (p. ??)** type. It is a data resource newly created by "indirect data access." A **SQLRowset (p. ??)** type resource is a data source that AMGA WS-DAIR service accesses to handle requests from both SQLResponse interface and **SQLRowset (p. ??)** interface. Because AMGA does not support a stored procedure, the queried result is always a single dataset or number of updated rows. Therefore, it would be redundant if we separate **SQLRowset (p. ??)** resource from SQLResponse resource.

WS-DAI consider two kinds of data resources: externally managed data resource and service managed data resource. In AMGA WS-DAIR, data resources are externally managed. This means that the managed data resources exist outside the WS-DAIR services and their lifetime can be managed in ways that are not specified in the WS-DAI specifications. AMGA WS-DAIR services do not support WSRF, which is optional in WS-DAI.

### 20.3.2 Data Resource Abstract Name & Data Resource Address

A data resource abstract name is defined as a unique and persistent name for a data resource. The WS-DAI specification mandates an abstract name to use the form of a URI. A URI has two specializations known as URL and URN. The AMGA WS-DAIR implementation uses the form of a URL. A data resource address is a name that specifies the locations of a data resource as accessed through a data access service. A data resource address is the same as a data resource abstract name in the AMGA WS-DAIR implementation. A data resource abstract name on the initial data resource for "direct data access" can be configured in configuration file, which is usually located in etc directory of installed root.

### 20.3.3 Supported Dataset and Language

The WS-DAIR specification mandates the **SQLAccess (p. ??)** interface provide at least WebRowSet format. The AMGA WS-DAIR implementation supports WebRowSet format. Supported Languages are <direct\_data\_statements> in SQL-92 entry level & and some of AMGA metadata languages

- **Native SQL Language** : Direct data statements in SQL-92 entry level : DELETE, INSERT, SELECT, UPDATE (JOIN, LIMIT, OFFSET keywords are supported)
- **AMGA Language** : AclAdd, AclRemove, AclShow, AddAttr, AddEntries, AddEntry, Backend, CapabilitiesAdd, CapabilitiesList, CapabilitiesRemove, Chmod, Chown, ClearAttr, CreateDir, Dir, Find, GetAttr, GrpAddUser, GrpCreate, GrpDelete, GrpList, GrpMember, GrpRemoveUser, GrpShow, Help, Import, IndexCreate, IndexRemove, Insert, Link, ListAttr, Ls, Mkdir, RemoveAttr, RemoveEntries, RenameAttr, Rm, Rmdir, SelectAttr, SetAttr, Update, UpdateAttr, UpdateSingle, UserCreate, UserList, UserPassChange, UserRemove, ViewCreate, ViewCreateUnion, ViewRemove, WhoAmI

### 20.3.4 Installation from Source

To install the SOAP server from source, you will need to first download the source distribution from the download directory.

For compilation you need to install gsoap package.

```

$ wget http://sourceforge.net/projects/gsoap2/files/gSOAP%20/gsoap_2.7.7.tar.gz
$ tar xzvf gsoap_2.7.7.tar.gz
$ cd gsoap-2.7
$ ./configure --prefix=$PWD/out
$ make
$ make install
$ cd ..

```

Now you should be ready to compile and install the AMGA server

```

$ wget
http://amga.web.cern.ch/amga/downloads/glite-amga-server- $\{version\}$ .tar.gz
$ tar xzvf glite-amga-server- $\{version\}$ .tar.gz
$ cd org.glite.amga.server
$ sh bootstrap
$ ./configure --with-readline --with-editline-prefix=/usr --with-server
--with-soapserver --with-gsoap-location=./gsoap-2.7/out --enable-rpm-rules
$ make
$ make install

```

You also need a database and the appropriate ODBC driver To setup a database and ODBC driver, refers **Server installation (p. 6)** .

## 20.4 Howto make WSDAIR clients

WS-DAIR is a soap-based interface, that allows you to create clients by creating interface bindings in any language out of the so called wsdl files, that describe the soap interface.

In the following an example is given how to create a client using the Ruby programming language for AMGA from the wsdl. Creating other interfaces for other languages is similar.

You need the 1.8 version of Ruby and the 1.5.8 version of soap4r, that provides the soap support for Ruby. The instructions are for debian, but this should work similarly on other OSs.

Install the gem support package

```

$ apt-get install rubygems
$ gem install soap4r
$ export PATH=/var/lib/gems/1.8/bin/:$PATH

$ mkdir wsdair
$ cd wsdair
$ cp org.glite.amga.server/src/soap/base/*.xsd .
$ cp amga/org.glite.amga.server/src/soap/base/*.wsdl .

```

Unfortunately there is an incompatibility between the gsoap soap implementation and the python/ruby implementations that were studied. You will therefore need to slightly modify the imple.wsdl file: Please insert the folling lines before the part <binding> starts in imple.wsdl

```

<wsdl:import namespace="http://www.ggf.org/namespaces/2005/12/WS-DAIR"
location="./sqlaccess.wsdl"/>
<wsdl:import namespace="http://www.ggf.org/namespaces/2005/12/WS-DAIR"
location="./sqlresponse.wsdl"/>
<wsdl:import namespace="http://www.ggf.org/namespaces/2005/12/WS-DAIR"
location="./sqlrowset.wsdl"/>
<wsdl:import namespace="http://www.ggf.org/namespaces/2005/12/WS-DAIR"
location="./wsdai_core_porttypes.wsdl"/>

```

Now the Ruby client can be created

```

$ wsdl2ruby.rb -d --type client --force --wsdl imple.wsdl

```

Several .rb ruby interface files are being created. If you have to use a gem for the 1.5.8 soap4r packages, like on debian, you need to edit wsdairDriver.rb and add the following 2 lines at the top

```
require 'rubygems'
gem 'soap4r'
```

so that the soap4r packages are being found.  
Create the following small demonstration client:

```
$ cat amgaclient.rb
#!/usr/bin/ruby
# Start me with -d to get debug output
#$DEBUG = true

require 'wsdairDriver.rb'

endpoint_url = 'http://localhost:8833/'
obj = SQLAccessPT.new(endpoint_url)
obj.wiredump_dev = STDERR if $DEBUG

expression = SQLExpressionType.new("SELECT * FROM /tutorial/cities")
request = SQLExecuteRequest.new("AMGA_Metadata",
"http://java.sun.com/xml/ns/jdbc", expression)

response = obj.sQLExecute(request)
print response

dataset=response.sQLDataset() # Type: SQLDatasetType
print "Dataset format URL >#{dataset.datasetFormatURI()}<\n"
print "Dataset data >#{dataset.datasetData()}<\n"
print "Update count >#{dataset.sQLOUpdateCount()}<\n"
print "Output parameter >#{dataset.sQLOutputParameter()}<\n"
print "Return value >#{dataset.sQLReturnValue()}<\n"
print "Communication >#{dataset.sQLCommunicationsArea()}<\n"

returnValue = dataset.sQLReturnValue()
data = dataset.datasetData()
print "Return value >#{returnValue}<\n"
print "Data retrieved >#{data}<\n"

wrs = data.webRowSet()
print "Web Row set >#{wrs}<\n"

wrsMetadata = wrs.metadata()
print "Web Row set metadata >#{wrsMetadata}<\n"
wrsColumnCount = wrsMetadata.column_count()
print "Web Row set columns >#{wrsColumnCount}<\n"

for i in 0..wrsColumnCount.to_i()
  print "#{wrsMetadata.column_definition[i].table_name} :
#{wrsMetadata.column_definition[i].column_name}\n"
end

currentRow = wrs.data().currentRow()

print "WRS Data >#{wrs.data()}<\n"
currentRow.each { |v| p v }
```

And change the query in the expression variable as well as the endpoint\_url according to your setup.

Start the soap-server:

```
$ ./mdsoapserver -c mdsoapserver.config
```

now you can run the client:

```
$ chmod +x amgaclient.rb  
$ ./amgaclient.rb
```

# A Changelog

## Changes in 2.3.0: (emi-2)

- Added support for SL6 x86\_64
- Added a Nagios probe for AMGA
- Added mdservice site option for Version Information Publish

## Changes in 2.2.0:

- Added support for openssl 1.0.0.0 or higher

## Changes in 2.1.2: (emi-1)

- fixed a bug which chmod fails with \*
- Added R option with chown : recursive
- fixed a bug that rep\_mount fails in replicating a table with too many rows because of memory problem.
- changed the default installation directory from /opt/glite to /usr

## Changes in 2.1.1:

- Added nosync option with rep\_mount : nosync option makes rep\_mount at the slave not to synchronize to the master. However, changes at the master are replicated to the slave after rep\_mount.
- Added noerase option with rep\_umount : noerase option makes rep\_umount not to remove the contents of mounted directory.
- fixed a bug which rep\_umount does not umount a directory cleanly after some replication failure
- fixed a bug which rep\_mount lets a master send dumped commands continuously even after there is a failure at the slave
- fixed a bug which user\_create inside a transaction can not be replicated.
- limitations of group name are removed. Now only characters which should be avoided are ' ', \, :, |, <space> at the group name.
- Added s option with ls : ls -s prints only subdirectories.
- Added version command which prints a version of the connected amga server

## Changes in 2.1.0:

- Added support for federation
- Added support for replicating plain tables
- Added support for replicating indexes
- Fixed a bug which WS-DAIR does not replicate properly
- Fixed a bug which native SQL does not replicate properly

## Changes in 2.0.0 :

- Added support for WS-DAIR (SOAP) Interface : Integration of AMGA into the DAIS framework of OGF standardized Grid Data Access Services
- Added support for sticky bit
- Added support for dynamic mapping from VOMS Role/Group to AMGA user/group



- Added support for Entry level ACL & permission check in native SQL query
- Added support for run-time configurable AMGA server : Existing Multi Process model or new Multi-threaded & Multi-process Model can be chosen through configuration
- Added support for native SQL Query Support on Oracle
- Changed command name mkdir to mkdir
- Added support for accessing FILE, OWNER, PERMISSIONS, GROUP\_RIGHTS attributes in native SQL query
- Fixed "SELECT \*" to print FILE followed by user's attributes
- Fixed a bug which a deleted group can not be removed from an ACL
- Fixed a bug which permission is not checked when creating sequence
- Fixed a bug which acl\_show generates error 9 when it accesses an entry
- Fixed a bug which index is not removed automatically with "rm -r"
- Fixed a bug which "rm -r" and "rmdir" do not require permission check
- Fixed a bug which permissions are checked in a wrong way with setattr, clearattr, getattr and find when a target in an entry
- Removed showing Group Permission with "ls -l" command, because it is useless information
- Fixed a bug which sequence\_next does not produce correct sequence in MySQL5
- Fixed a bug which configure does not fail when boost-thread library is not found
- Fixed a bug of rm command in Oracle
- Fixed a bug of view\_create command
- Added support for varbit (only works with PostgreSQL)
- Fixed a bug that rep\_allow command produces "Error: 4: Directory not found" in Oracle
- Fixed a bug that rep\_mount produces "Local execution failed. 4 Directory not found" in Oracle
- Fixed a bug that rep\_umount produces "'Error: 4: Permission denied ...'"
- Fixed a bug that inherit option is not shown with stat and dump command
- Fixed a bug that native sql does not compile in gentoo (-tags=CXX)
- Fixed a bug that very long SQL executes very slow (200K length query took about 1 minute)
- Fixed a bug that very long input data is truncated when inserted

#### **Changes in 1.9.0:**

- Added support for native SQL query
- Added support for multi-threaded multi-process service model with DB connection sharing
- Added support for more flexible masterindex schema

#### **Changes in 1.3 (glite 3.1):**

- Added support for admin flag in ACLs
- Added support for database schemas
- Added support for outer and inner joins

- Fixed nasty bug in stream buffer corrupting output under load when transactions were used
- Made directory traversal use bound parameters (DB performance)

#### **Changes in 1.2.10:**

- Fixed user replication
- Added monitoring of user connections numbers
- Added limits on number of connections per user
- Added possibility to reserve connections for root
- Fixed constraints in MySQL, in particular for plain directories
- Added capabilities (only used for replication), to use, please update schema: ALTER TABLE users ADD COLUMN "capabilities" VARCHAR(256); See "help capabilities"
- Added the possibility to allow ordinary users to replicate login information via rep\_users\_ - allow, rep\_users\_ disallow
- Added new type of view to create a union of several (replicated) directories, see: create\_ - view\_ union
- Made create\_ union more powerful, please read the docs
- Added custom primary keys: constraint\_ add\_ primary\_ key
- Fixed username/password default from odbc.ini to work (thanks to Tamas Hauer).
- Fixed several bugs on 64 bit architecture (thanks to Tamas for reporting them).
- Added isnull() and notnull() function to give IS [NOT] NULL syntax in SQL.
- Added if() function for conditional evaluation.
- Loosened permission checking on ACL changes for entry-ACLs

#### **Changes in 1.2.9:**

- Fixed dump command in MySQL < 5.0 and for SQLite
- Fixed parsing problem with capital "AND"
- Problems with table creation fixed for SQLite, MySQL < 5.0 and PostgreSQL
- Separated MySQL initialization files for version < 5.0 from those for 5.x
- Stat shows possible directory options
- New regression testing tool for the query parser
- Sequences implemented for MySQL5.x

#### **Changes in 1.2.8:**

- Added type=<TYPE> option in createdir/mkdir to create entry column with custom type
- Added execute command to run server-side scripts
- Initial implementation of import feature for existing tables
- Added update command, which updates using bound variable to circumvent Oracle problems with SQL command length limitations
- fixes for rm command
- fixed constraint\_ add\_ unique and constraint\_ add\_ not\_ null on MySQL

**Changes in 1.2.7:**

- Reimplemented mdserver as amgad.
- addattr checks for system columns and rejects them
- addattr works only on directories
- Support for preexisting tables as plain tables
- Added backend command to allow user to figure out the database backend.
- Added documentation for views.
- Fixed command option check for index\_create.
- Support for ACLs for sequences, indices, views.
- Changed listattr to use 'x' flag for permission checking.
- Fixed getattr to allow showing of entries when default ACL allows it.
- Changed selectattr to allow access to individual entries.
- Initial version of renameattr (works only on Postgres and Oracle)

**Changes in 1.2.6:**

- Fixing CRL checking, adding CRLDir option to amgad.config
- Create all groups successfully mapped by VOGroupMap automatically on login
- Cleaning up include files of C++ API
- Adding several functions to C++ API to make it more Posix-IO like
- Several Oracle bugs found by Viktor fixed
- rm -r added
- Added -a option to grp\_list. Users can now see other user's groups.
- Fixed constraint\_add\_reference and constraint\_add\_check

**Changes in 1.2.5:**

- Replication of users and groups
- Fixed stat output, now it is doing ls -lP
- Fixed bug in user\_listcred if no certificate was assigned to a user
- updateattr reports error if no update performed
- Improved mdclient behaviour if upload command is messed up by the user

**Changes in 1.2.4:**

- Support for GLITE\_LOCATION environment variable for the client configuration file
- Support for X509\_USER\_PROXY environment variable to locate the grid-proxy
- Support for limited proxies

**Changes in 1.2.3:**

- Fix login problem with VO certificates and Groups-on-the-fly
- Replication layer fixes

- Increased listen queue in mdserver to 500: Much better performance with high number of clients

#### **Changes in 1.2.2:**

- Replication layer refactored
- Standalone (without gLite build system) RPMs now possible

#### **Changes in 1.2.1:**

- Java API: Implemented on the **MDClient** (p. ??) the methods for constraints, groups, acls and sequences
- Java API: Fixed a bug in the **MDClient** (p. ??) constructor where the username and password were not being properly initialized on th
- Adding date in log messages
- Nicer debug messages
- Fixed constraint removal when removing attribute
- Fixed reading of configuration file in home-directory on SLC3
- Fixed request cancellation in C++ client
- Fixed uninitialized memory in command encoding routine problem
- Fixed several bugs of replication with Oracle
- Fixed replication of some commands where the logs were not being saved with the correct directory information.
- Improved dump performance
- Implemented replication of indices and of the upload/put/commit commands

#### **Changes in 1.2.0:**

- Added support for distinct, order and limit clauses in selectattr
- New parser for configuration files. New escape char is \. Read the documentation!
- Support for MonAMI (and from there to MonALISA) monitoring of the service.
- Support for explicite transaction handling.
- AMGA commands added to set up table constraints.
- BDII support including fallback to another AMGA server for the client

#### **Changes in 1.1.0:**

- Support for MySQL5
- Support for GIS on MySQL5
- Directories with shared schema may share a table (createdir parent shared)
- Per-entry ACLs in a directory (Postgres only!): createdir dir acls
- Inheriting schemas has been removed from the PostgreSQL version (was never supported elsewhere and shared directories make up for it)
- mdclient silent mode (-S) fixed
- mdclient command completion added

#### **Changes in 1.0.0 (gLite 1.5):**

- Dependency on readline optional (`./configure --with-readline`) by default editline is used
- Fixes to work with openssl 0.9.6 to 0.9.8, not only 0.9.7
- `Log()` function returns always natural logarithm
- `dir()`, `stat()`, `listentries()`, `find()` take `r-flag` into account
- `listattr()` and `getattr()` honour `r-flag` of directory
- Fixed group-existence checking in group commands
- Separating `user_list` into `user_list` and `user_listcred`
- Dump method improved
- Directory matching operation fixed in SQLite
- Fixing rollback mechanism in uploads
- Several checks on put-command arguments added
- SQL query added to error output
- Disallow creation of subdirectories of entries
- Adding `grp_dump`, `user_dump`
- Fixed ACL-inheritance of new directories
- Cleaned up certificate checking
- Package can compile against Globus

#### **Changes in 0.9.9:**

- Move to gLite build environment
- Removed ZThreads dependency, Boost is not used instead
- External dependency on gSOAP removed, instead generated code is included
- Bug with unary minus fixed in parser
- Change to use 3 RPMs: client, server, soapserver
- Move to `GLITE_LOCATION` home-directory, default is `/opt/glite`
- Use `datetime` as timestamp datatype on MySQL
- `listschemas` command added to support PTF-API
- Move to flex 2.5.4 for improved compatibility, flexer object
- Several new functions in parser
- `vomsattrusermanager` added
- MyProxy works
- `like()` and `count()` added
- Added man-pages for client and server executables
- Lots of documentation improvements
- New `grp_list` command

#### **Changes in 0.9.6:**

- gLite integration

- New authentication and authentication module vomsinfo based on VOMS attribute certificates
- Introduction of default system:anyuser group
- Enabled 'x' entry visibility flag
- Delete() operation performance problem solved
- Separation into client, server and soapserver RPMs
- like() and count() functions in queries

**Changes in 0.9.4:**

- Support for LFC running with an Oracle back-end
- Fixes for several problems with the new Oracle ODBC driver
- New dump command to backup the database
- Fixes for wildcards in filenames

**Changes in 0.9.3:**

- Introduction of basic datatypes which are supported on all backends.
- Initial support of MySQL as a standalone backend, everything except sequences, views and indices should work.
- Moving to the new Oracle ODBC driver: The new release need the Oracle 10.2 ODBC driver for Oracle access. The EasySoft driver will not work.
- Many bugs fixed for the Oracle backend, more performant access to large metadata made possible through new ODBC driver.

Acknowledgement : This work is co-funded by the EC EMI project under the FP7 Collaborative Projects Grant Agreement Nr. 261611.