

GT 4.2.0 RLS: System Administrator's Guide

DRAFT

GT 4.2.0 RLS: System Administrator's Guide

Introduction

This guide contains advanced configuration information for system administrators working with RLS. It provides references to information on procedures typically performed by system administrators, including installation, configuring, deploying, and testing the installation.

Important

This information is in addition to the basic Globus Toolkit prerequisite, overview, installation, security configuration instructions in the [Installing GT 4.2.0](#). Read through this guide before continuing!

DRAFT

Table of Contents

| | |
|--|----|
| 1. Building and Installing | 1 |
| 1. Prerequisites | 1 |
| 2. Installation of RLS | 5 |
| 3. Database Setup | 6 |
| 2. Configuring RLS | 11 |
| 1. Configuration overview | 11 |
| 2. Server configuration file (globus-rls-server.conf) | 11 |
| 3. Basic configuration | 11 |
| 4. Host key and certificate configuration | 12 |
| 5. Configuring LRC to RLI updates | 13 |
| 6. Configuring the RLS Server for the WS MDS Index Service | 14 |
| 7. Configuring the RLS Server for the MDS2 GRIS | 15 |
| 8. Complete RLS Server settings (globus-rls-server.conf) | 15 |
| 3. Starting the RLS Server | 21 |
| 4. Stopping the RLS Server | 22 |
| 5. Testing | 23 |
| 6. Security Considerations | 24 |
| 1. Replica Location Service (RLS) Security Considerations | 24 |
| 7. Troubleshooting | 25 |
| 1. Errors | 25 |
| 2. I need help installing MySQL DBMS | 26 |
| 3. How can I check if MySQL is installed correctly? | 26 |
| 4. I need help installing PostgreSQL DBMS | 27 |
| 5. How can I check if PostgreSQL is installed correctly? | 27 |
| 6. I cannot connect to PostgreSQL | 27 |
| 7. Performance of RLS with PostgreSQL has degraded over time | 28 |
| 8. I need help installing SQLite embedded database | 28 |
| 9. How can I check if SQLite is installed correctly? | 28 |
| 10. I need help installing iODBC | 28 |
| 11. I need help installing unixODBC | 29 |
| 12. Building unixODBC failed on OS X | 29 |
| 13. I need help installing MySQL Connector/ODBC | 30 |
| 14. I need help installing psqlODBC | 32 |
| 15. I need help installing SQLite ODBC Driver | 33 |
| 16. Which libraries do you recommend to new users? | 34 |
| 17. I need help configuring my DSNs | 34 |
| 18. How can I tell which odbc.ini is being used? | 35 |
| 19. Testing your ODBC configuration | 35 |
| 20. Important notes on RLS initialization | 37 |
| 21. Linux kernel and glibc incompatibility | 37 |
| 8. Usage statistics collection by the Globus Alliance | 39 |
| 1. RLS-specific usage statistics | 39 |
| Glossary | 41 |
| Index | 42 |

List of Tables

| | |
|--|----|
| 1.1. Database & ODBC Observations | 4 |
| 2.1. Complete RLS Server settings (globus-rls-server.conf) | 16 |
| 7.1. Replica Locator Service (RLS) Errors | 25 |

DRAFT

Chapter 1. Building and Installing

RLS is built and installed as part of the Globus Toolkit installation. This section provides instructions specific to building and installing RLS. Most of this section concerns the optional setup and configuration of a database management system (DBMS) and database connectivity (ODBC) software for the RLS server. By default, the RLS is installed and automatically configured to use an embedded database. No additional steps are required unless you wish to configure RLS to use your own ODBC-compliant database management system.

1. Prerequisites

By default, the RLS is installed and automatically configured to use an embedded database. If you wish to use the default configuration, you may skip this section and proceed to [Section 2, “Installation of RLS”](#). If you wish to use your own ODBC-compliant database management system, please review this section carefully.

The RLS depends on DBMS and ODBC software. If you are not familiar with DBMS and ODBC software, you may need to familiarize yourself with these tools before proceeding with RLS installation. Most of the difficulty users experience when installing RLS is actually due to DBMS and ODBC issues, which are often beyond our control.

The following list summarizes the prerequisite database software:

- [Section 1.1, “Database Management System”](#): for persistent storage of replica location information
- [Section 1.2, “ODBC Manager”](#): for vendor-neutral connectivity to DBMS
- [Section 1.3, “ODBC Driver”](#): for vendor-specific connectivity to DBMS from ODBC Managers

When using this guide, please consult the vendor and official project sites for authoritative information on DBMS and ODBC software. Any additional information provided on this site is intended as helpful advice and only reflects our limited experiences as fellow users of the software.

1.1. Database Management System

The RLS depends on a DBMS for persistent storage of replica location information. The RLS is designed to work with MySQL, PostgreSQL, SQLite, and Oracle DBMS software. In the case of MySQL, it is important to use the transaction-capable InnoDB database engine. For complete information on installation and configuration of third-party DBMS software, please refer to:

- [MySQL](#)¹
- [PostgreSQL](#)²
- [SQLite](#)³
- [Oracle](#)⁴

We provide the following tips from the troubleshooting section of this guide:

- [Section 4, “I need help installing PostgreSQL DBMS”](#)

¹ <http://www.mysql.com>

² <http://www.postgresql.org>

³ <http://www.sqlite.org>

⁴ <http://www.oracle.com>

- [Section 5, “How can I check if PostgreSQL is installed correctly?”](#)
- [Section 6, “I cannot connect to PostgreSQL”](#)
- [Section 7, “Performance of RLS with PostgreSQL has degraded over time”](#)
- [Section 2, “I need help installing MySQL DBMS”](#)
- [Section 3, “How can I check if MySQL is installed correctly?”](#)
- [Section 8, “I need help installing SQLite embedded database”](#)
- [Section 9, “How can I check if SQLite is installed correctly?”](#)

1.2. ODBC Manager

The RLS depends on an ODBC manager for standardised database connectivity. The unixODBC or iODBC libraries may be used for this purpose. Please confirm that one of these libraries is installed and configured before proceeding with RLS installation. If you need more information on the ODBC managers, consult their respective project home pages.

- [unixODBC](#)⁵
- [iODBC](#)⁶

You may also find that unixODBC or (less likely) iODBC are included along with a vendor's odbc driver. This appears to be the case with MySQL Connector/ODBC drivers at the time of writing this document.

We provide the following tips from the troubleshooting section of this guide:

- [Section 11, “I need help installing unixODBC”](#)
- [Section 10, “I need help installing iODBC”](#)



Note

If you are installing RLS from a Globus Toolkit binary bundle, you do not need to install the unixODBC or iODBC software. The binary installation of RLS is linked to iODBC version 3.51.2 included with the GT bundle.

1.3. ODBC Driver

The RLS depends on an ODBC driver for vendor-specific database connectivity via the vendor neutral ODBC manager. For each vendor you will find a related ODBC driver. You may find relevant information at the following sites:

- [psqlODBC](#)⁷ the official PostgreSQL ODBC Driver
- [MySQL Connector/ODBC](#)⁸ offered directly by MySQL (note: use the 3.51 branch)
- [SQLite ODBC Driver](#)⁹ a freely available ODBC driver for SQLite

⁵ <http://www.unixodbc.org>

⁶ <http://www.iodbc.org>

⁷ <http://pgfoundry.org/projects/psqlodbc/>

⁸ <http://www.mysql.com/products/connector/odbc/index.html>

⁹ <http://www.ch-werner.de/sqliteodbc/>

- Consult Oracle¹⁰ for details on ODBC Driver for Oracle DBMS

We provide the following tips from the troubleshooting section of this guide:

- Section 14, “I need help installing psqlODBC”
- Section 13, “I need help installing MySQL Connector/ODBC”
- Section 15, “I need help installing SQLite ODBC Driver”

1.4. Choosing the right DBMS & ODBC software

We have used or know of users who have used the following combinations of software packages on the following platforms. What we provide in this section is a small subset of available options when choosing DBMS & ODBC software.

¹⁰ <http://www.oracle.com>

Table 1.1. Database & ODBC Observations

| Architecture | OS | Distribution | ODBC Manager | ODBC Driver | DBMS |
|---------------|---------|--------------|-------------------|------------------------------|------------------------|
| PPC | OS X | Apple | unixODBC-2.2.11* | TBD | TBD |
| SPARC, 32-bit | Solaris | Sun | libiodbc-3.51.2 | MyODBC-3.51.06 | mysql-standard-4.0.18 |
| SPARC, 32-bit | Solaris | Sun | libiodbc-3.51.2 | MyODBC-3.51.12 | mysql-4.0.16 |
| SPARC, 32-bit | Solaris | Sun | unixODBC-2.2.11 | MyODBC-3.51.12 | mysql-4.0.16 |
| SPARC, 32-bit | Solaris | Sun | unixODBC-2.2.11 | TBD | TBD |
| x86 | Linux | RedHat | libiodbc-3.51.1 | MyODBC-3.51.06 | mysql-standard-4.0.18 |
| x86 | Linux | RedHat | libiodbc-3.51.2 | MyODBC-3.51.06 | mysql-standard-4.0.18 |
| x86 | Linux | RedHat | libiodbc-3.52.4 | MyODBC-3.51.06 | mysql-standard-4.0.18 |
| x86 | Linux | RedHat | libiodbc-3.52.4 | psqlodbc-7.2.5 | postgresql-7.4.6 |
| x86 | Linux | RedHat | libiodbc-3.51.2 | psqlodbc-08.00.0102 | postgresql-7.4.6 |
| x86 | Linux | RedHat | libiodbc-3.51.2 | sqliteodbc-0.69 | sqlite-3.3.6 |
| x86 | Linux | RedHat | unixODBC-2.2.11 | MyODBC-3.51.06 | mysql-standard-4.0.18 |
| x86 | Linux | RedHat | unixODBC-2.2.11 | mysql-connector-odbc-3.51.12 | mysql-standard-4.0.18 |
| x86 | Linux | RedHat | unixODBC-2.2.11 | mysql-connector-odbc-3.51.12 | mysql-standard-5.0.21 |
| x86 | Linux | RedHat | unixODBC-2.2.11 | psqlodbc-7.2.5 | postgresql-7.4.6 |
| x86 | Linux | RedHat | unixODBC-2.2.11 | psqlodbc-08.00.0102 | postgresql-7.4.6 |
| x86 | OS X | Apple | TBD | TBD | TBD |
| x86_64 | Linux | SuSE | unixODBC-2.2.8-35 | MyODBC-unixODBC-3.51.06-121 | mysql-client-4.0.18-25 |
| x86_64 | Linux | SuSE | unixODBC-2.2.8-35 | psqlodbc-08.00.0102 | postgresql-7.4.9 |

(*) See troubleshooting tip [Section 12, “Building unixODBC failed on OS X”](#)

If you build from source, be aware that some versions of a vendor's drivers do not build against some versions of the same vendor's clients. For instance, MySQL Connector/ODBC 3.51.06 does not appear to build against MySQL 4.1.x client libraries, instead it requires MySQL 4.0.x client libraries. Also, we have noticed that psqlodbc-07.03.x will build against postgres-7.4.x but when using the driver it fails with little or no information to indicate why it failed. We have found that psqlodbc-08.00.x works well with postgres-7.4.6. In all cases, consult the vendor's documentation.

Support RLS development by sending your working RLS + ODBC Manager + ODBC Driver + RDBMS configuration to the rls-user@globus.org list! We cannot test on all platforms with all combinations so we rely on our user community to help advise new users. Your support in this matter is appreciated.

We provide the following tips from the troubleshooting section of this guide:

- [Section 16, “Which libraries do you recommend to new users?”](#)

2. Installation of RLS

The following section describes various options available to install RLS given that you have satisfied the prerequisites. The options include:

- [Section 2.1, “Installing RLS using Globus Toolkit Binary Bundle”](#)
- [Section 2.2, “Installing RLS using Globus Toolkit Source Bundle”](#)
- [Section 2.3, “Updating RLS using Globus Toolkit Source Bundle”](#)

Note

If you set the ODBCINI environment variable to the path of the `odbc.ini` file prior to installation, the `globus-rls-server.conf` will be seeded with the corresponding value. If you are not already familiar with `odbc.ini` it will be explained in [Section 3.5, “Specify Data Source Names \(DSNs\)”](#). The setting is optional, since the `globus-rls-server.conf` can be changed at any time.

2.1. Installing RLS using Globus Toolkit Binary Bundle

Follow instructions provided by [Installing GT 4.2.0](#) and [Installing GT](#) in order to install the RLS from one of the available binary bundles of the Globus Toolkit. As noted in [Section 1.2, “ODBC Manager”](#), the `globus-rls-server` found in the binary bundles is linked to the default ODBC Manager.

Once you have unpackaged the binary bundle, the following commands may be used to install RLS (the `--enable-rls` option is not required):

```
% ./configure --prefix=$GLOBUS_LOCATION --enable-rls
% make
% make install
```

2.2. Installing RLS using Globus Toolkit Source Bundle

Follow instructions provided by [Installing GT 4.2.0](#) and [Installing GT](#) in order to install the RLS from the available source bundle of the Globus Toolkit. In addition to the general requirements, you may optionally set environment variables pertaining to the include and lib path for your ODBC Manager installation. The following environment variables are not required to install RLS.

If you wish to use `unixODBC`, set the following environment variables.

```
% setenv GLOBUS_UNIXODBC_INCLUDES /path/to/unixODBC/include
% setenv GLOBUS_UNIXODBC_LIBS /path/to/unixODBC/lib
```

If you wish to use `iODBC`, set the following environment variables.

```
% setenv GLOBUS_IODBC_INCLUDES /path/to/libiodbc/include
```

```
% setenv GLOBUS_IODBC_LIBS /path/to/libiodbc/lib
```



Note

Your shell format for setting environment variables may be `export VAR=VAL` instead of `setenv VAR VAL`.

Alternatively, instead of setting the above environment variables you may choose to specify the includes and libs path for your ODBC Manager using the Globus Toolkit `configure` script's options, which include `--with-unixodbc-includes=DIR` and `--with-unixodbc-libs=DIR` or `--with-iodbc-includes=DIR` and `--with-iodbc-libs=DIR`.

Once you have unpackaged the source bundle, the following commands may be used to install RLS (the `--enable-rls` option is not required):

```
% ./configure --prefix=$GLOBUS_LOCATION --enable-rls
% make
% make install
```

2.3. Updating RLS using Globus Toolkit Source Bundle

In some situations, you may wish to update an RLS installation using RLS source from the Globus Toolkit Source Bundle. For instance, the RLS found in the Globus Toolkit binary bundles is linked against an ODBC Manager. If you would like to link to a different ODBC Manager, it may be easier to install the toolkit from the binary bundle and just update the RLS server package from source. This may save considerable time compared to building the entire Globus Toolkit from source.

To do this, you must first follow the complete instructions found in [Section 2.1, “Installing RLS using Globus Toolkit Binary Bundle”](#) then set the environment according to the instructions found in [Section 2.2, “Installing RLS using Globus Toolkit Source Bundle”](#). If you have not already set the `GPT_LOCATION` then set it to the `GLOBUS_LOCATION` (i.e., `setenv GPT_LOCATION $GLOBUS_LOCATION`).

Once you have unpackaged the source bundle, the following commands may be used to install RLS:

```
% cd source-trees-thr/replica/rls/server/
% $GPT_LOCATION/sbin/gpt-build -verbose -force gcc32dbgpthr
...
% $GPT_LOCATION/sbin/gpt-postinstall
...
```

You may select a flavor other than `gcc32dbgpthr` that is suitable to your platform and existing installation environment. You may omit the `-verbose` flag.

3. Database Setup

By default, the RLS embedded database and related settings are automatically created and configured. Unless you chose to setup your own database management system for use with the RLS, you may skip this section and proceed to [Chapter 2. Configuring RLS](#).

The following section describes the procedures necessary to complete the database setup for your RLS installation. Topics addressed include:

- [Section 3.1, “Choose databases for RLS Server”](#)
- [Section 3.2, “Configure database related RLS settings”](#)
- [Section 3.3, “Create a database user account”](#)
- [Section 3.4, “Create databases”](#)
- [Section 3.5, “Specify Data Source Names \(DSNs\)”](#)



Note

The rest of this section assumes that the DBMS has been started.

3.1. Choose databases for RLS Server

Decide which database(s) the RLS server will use:

- If the RLS server is a Local Replica Catalog (LRC) server you, will need to create the LRC database.
- If the server is a Replica Location Index (RLI) server, you may need to create a RLI database.

An RLI server can receive updates from LRC servers in one of two forms, as LFN lists (in which case the RLI database must be created) or as highly compressed Bloom filters. Since Bloom filters are relatively small, they are kept in memory and no database is required (though if you plan to use the hierarchical RLI updates feature, you will still need the RLI database). An RLS server can be configured as both an LRC and RLI server.

3.2. Configure database related RLS settings

The RLS Server configuration file (`$GLOBUS_LOCATION/etc/globus-rls-server.conf`) contains a few settings related to your database setup. Throughout this guide we refer to the username as *dbuser* and the password as *dbpassword* and to the LRC database as *lrc1000* and the RLI database as *rli1000*. Substitute your chosen username, password, and database names where applicable. Also, you may change the `odbcini` setting to reflect the location of your ODBC Manager's `odbc.ini` file. You should be familiar with the location of this file from the installation of your ODBC Manager. As of the GT 4.1.x release, the RLS setup includes a default `odbc.ini` file installed to `$GLOBUS_LOCATION/var/odbc.ini` which is the initial default value of the `odbcini` configuration parameter. We discuss more about `odbc.ini` in [Section 3.5, “Specify Data Source Names \(DSNs\)”](#).

Edit the following fields in `$GLOBUS_LOCATION/etc/globus-rls-server.conf`:

```
# Database connection options
db_user          dbuser
db_pwd           dbpasswordgoeshere
odbcini          /path/to/var/odbc.ini
...
# LRC options
lrc_server       true
lrc_dbname       lrc1000 # optional (if LRC server)
...
#RLI options
```

```
rli_server          true
rli_dbname          rli1000 # optional (if RLI server)
...
rli_bloomfilter     false
...
```

3.3. Create a database user account

You or your system administrator must create a DBMS user account to be used by the RLS Server to log on to the DBMS. SQLite users may safely skip this step.

For MySQL users, the command may look something like this:

```
% /path/to/mysql/bin/mysql -u root -p
Enter password:
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 71 to server version: 4.0.18-standard-log

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql> use mysql;
Database changed
mysql> grant all on lrc1000.* to dbuser@localhost identified by 'dbpassword';
mysql> grant all on rli1000.* to dbuser@localhost identified by 'dbpassword';
```

For PostgreSQL user, the command may look something like this:

```
% /path/to/pgsql/bin/createuser -P dbuser -W
Enter password for new user:
Enter it again:
Shall the new user be allowed to create databases? (y/n) y
Shall the new user be allowed to create more new users? (y/n) n
Password:
CREATE USER
```

3.4. Create databases

We provide SQL Data Definition (DDL) scripts to create the databases required by the RLS Server. Configure the script(s) for the database(s) you will create. The SQL scripts for the LRC and RLI databases are installed in `$GLOBUS_LOCATION/setup/globus`.

For PostgreSQL, use:

- `globus-rls-lrc-postgres.sql`
- `globus-rls-rli-postgres.sql`

For MySQL, use:

- `globus-rls-lrc-mysql.sql`
- `globus-rls-rli-mysql.sql`

For SQLite, use:

- `globus-rls-lrc-sqlite.sql`
- `globus-rls-rli-sqlite.sql`

For Oracle DBMS, use:

- `globus-rls-lrc-oracle.sql`
- `globus-rls-rli-oracle.sql`

Edit these files to set the name of the database user you created for RLS and the names of the databases configured in `$GLOBUS_LOCATION/etc/globus-rls-server.conf`.

Create the database(s) using the appropriate tool for your chosen DBMS.

For PostgreSQL, the commands should look like this:

```
% createdb -O dbuser -U dbuser -W lrc1000
% createdb -O dbuser -U dbuser -W rli1000
% psql -W -U dbuser -d lrc1000 -f $GLOBUS_LOCATION/setup/globus/globus-rls-lrc-postgres.
% psql -W -U dbuser -d rli1000 -f $GLOBUS_LOCATION/setup/globus/globus-rls-rli-postgres.
```

For MySQL, the commands should look like this:

```
% mysql -p -u dbuser < $GLOBUS_LOCATION/setup/globus/globus-rls-lrc-mysql.sql
% mysql -p -u dbuser < $GLOBUS_LOCATION/setup/globus/globus-rls-rli-mysql.sql
```

For SQLite, the commands should look like this:

```
% sqlite3 $GLOBUS_LOCATION/var/globus-rls-lrc1000.db < $GLOBUS_LOCATION/setup/globus/glob
% sqlite3 $GLOBUS_LOCATION/var/globus-rls-lrc1000.db < $GLOBUS_LOCATION/setup/globus/glob
```

3.5. Specify Data Source Names (DSNs)

Data Source Names (DSNs) are used by ODBC software (and the applications that depend on ODBC software) to define connection settings between the database client and the DBMS. The RLS depends on the data source names configured in [Section 3.2, “Configure database related RLS settings”](#) for the `lrc_dbname` and `rli_dbname` settings.

DSNs are configured in a file named `odbc.ini`. Most ODBC Managers find this file in the following search order.

- The `ODBCINI` environment variable with value of `/path/to/odbc.ini`
- The user `odbc.ini` found at `~/odbc.ini`

- The system `odbc.ini` found at `/etc/odbc.ini` or similar location

DSNs may be configured in a variety of ways, but they essentially contain the same information which includes:

- *DSN* such as `lrc1000` or `rli1000`
- *Driver*: the path to the ODBC driver (such as `psqlodbc.so`, `libmyodbc3.so`, or `libsqlite3odbc.so`)
- *Database*: the name of the database (in the case of SQLite, this parameter's value will indicate the path to the database file)
- *Server name*: the hostname of the DBMS server
- *Port*: the port number of the DBMS server
- *Socket*: in some cases a local socket file is used if the DBMS and RLS are installed on the same server

As of the Globus Toolkit 4.1.x release, the RLS setup includes a default `odbc.ini` file installed to `$GLOBUS_LOCATION/var/odbc.ini`. The settings of the included `odbc.ini` are appropriate for usage with the `psqlODBC` driver and reference `psqlODBC` libraries installed to `$GLOBUS_LOCATION/lib/psqlodbc.so`, which is the location of the libraries as installed by the Globus Toolkit binary installer and the source installer by using the `make globus_database_psqlodbc` target. When using the installed `odbc.ini` file, you may still need to set your `ODBCINI` environment variable to the `odbc.ini` file's location.

We provide the following tips from the troubleshooting section of this guide:

- [Section 17, "I need help configuring my DSNs"](#)
- [Section 18, "How can I tell which odbc.ini is being used?"](#)
- In particular, please see [Section 19, "Testing your ODBC configuration"](#)

Chapter 2. Configuring RLS

1. Configuration overview

RLS configuration involves statically-defined, system settings as defined in the RLS configuration file (see `$GLOBUS_LOCATION/etc/globus-rls-server.conf`), settings changed temporarily at run-time using the RLS Admin tool (see `globus-rls-admin(1) -C option value` command), and finally LRC-to-RLI and RLI-to-RLI updates configured using the RLS Admin tool (see `globus-rls-admin(1) -a, -A, -d` commands).

2. Server configuration file (`globus-rls-server.conf`)

Configuration settings for the RLS are specified in the `globus-rls-server.conf` file. If the configuration file is not specified on the command line (see the `-c` option) then it is looked for in both:

- `$GLOBUS_LOCATION/etc/globus-rls-server.conf`
- `/usr/local/etc/globus-rls-server.conf` if `GLOBUS_LOCATION` is not set

Note

Command line options always override items found in the configuration file.

The configuration file is a sequence of lines consisting of a keyword, whitespace, and a value. Comments begin with `#` and end with a newline.

3. Basic configuration

Review the server configuration file `$GLOBUS_LOCATION/etc/globus-rls-server.conf` and change any options you want. The server man page `globus-rls-server(8)` has complete details on all options. The complete details are also provided later in this section.

A minimal configuration file for both an LRC and RLI server would be:

```
# Configure the database connection info
db_user      dbuser
db_pwd       dbpassword

# If the server is an LRC server
lrc_server    true
lrc_dbname    lrc1000

# If the server is an RLI server
rli_server    true
rli_dbname    rli1000 # Not needed if updated by Bloom filters

# Configure who can make requests of the server
acl .*: all
```

```
# RE matching grid-mapfile users or DNs from x509 certs
...
```

4. Host key and certificate configuration

The server uses a host certificate to identify itself to clients. By default this certificate is located in the files `/etc/grid-security/hostcert.pem` and `/etc/grid-security/hostkey.pem`. Host certificates have a distinguished name of the form `/CN=host/FQDN`. If the host you plan to run the RLS server on does not have a host certificate, you must obtain one from your Certificate Authority. The RLS server must be run as the same user who owns the host certificate files (typically root). The location of the host certificate files may be specified in `$GLOBUS_LOCATION/etc/globus-rls-server.conf`:

```
rlscertfile      path-to-cert-file    # default /etc/grid-security/hostcert.pem
rlskeyfile       path-to-key-file   # default /etc/grid-security/hostkey.pem
```

It is possible to run the RLS server without authentication, by starting it with the `-N` option, and using URL's of the form `rlsn://server` to connect to it. Notice that the URL scheme is `rlsn` as opposed to `rls`.

It is generally recommended to run the server with a user account other than root for added security. In order to do so, you will need to create complimentary key and certificate files owned by a designated user account, `globus` for instance.

1. Begin by copying the `/etc/grid-security/hostcert.pem` and `/etc/grid-security/hostkey.pem` to `/etc/grid-security/containercert.pem` and `/etc/grid-security/containerkey.pem`. Note that we use the prefix "container" to conform with the recommended naming scheme for other services distributed with the Globus Toolkit.

```
% cp /etc/grid-security/hostcert.pem /etc/grid-security/containercert.pem
% cp /etc/grid-security/hostkey.pem /etc/grid-security/containerkey.pem
```

2. Then change ownership of the files to the designated user account, `globus` in our example.

```
% chown globus /etc/grid-security/containercert.pem
% chown globus /etc/grid-security/containerkey.pem
```

3. Change the `rlskeyfile` and `rlscertfile` settings in the RLS configuration file (`$GLOBUS_LOCATION/etc/globus-rls-server.conf`) to reflect the appropriate filenames.

```
rlscertfile      /etc/grid-security/containercert.pem
rlskeyfile       /etc/grid-security/containerkey.pem
```

4. Finally, bear in mind that your certificate and key files must always have file permissions 644 and 400 respectively.

```
% ls -l /etc/grid-security/*.pem
-rw-r--r--  1 globus  gridstaff    818 Dec  8  2005 /etc/grid-security/containercer
-r-----   1 globus  gridstaff    887 Dec  8  2005 /etc/grid-security/containerkey
-rw-r--r--  1 root    root        818 Dec  8  2005 /etc/grid-security/hostcert.pem
-r-----   1 root    root        887 Dec  8  2005 /etc/grid-security/hostkey.pem
```

If authentication is enabled, RLI servers must include `acl` configuration options that match the identities of LRC servers that update it and that grant the `rli_update` permission to the LRCs.

5. Configuring LRC to RLI updates

One of the key benefits to using the RLS for managing replica location information is its distributed architecture. In a distributed deployment, one or more Local Replica Catalog (LRC) services will send updates of its contents to one or more Replica Location Index (RLI) services.

By default the installed LRC is not configured to send updates to any RLI, even the local RLI co-located with the local LRC. Use the `globus-rls-admin(1)` tool to configure the LRC to send updates to one or more RLI services.

- To configure the LRC to send uncompressed lists of its logical names to a RLI, use the following command:

```
% $GLOBUS_LOCATION/bin/globus-rls-admin -a rls://rli_host rls://lrc_host
```

- To configure the LRC to send compressed bitmaps (using Bloom filters) of its logical names to a RLI, use the following command:

```
% $GLOBUS_LOCATION/bin/globus-rls-admin -A rls://rli_host rls://lrc_host
```

- To configure the LRC to stop sending updates to a RLI, use the following command:

```
% $GLOBUS_LOCATION/bin/globus-rls-admin -d rls://rli_host rls://lrc_host
```

Note

While any given LRC is capable of sending uncompressed or compressed updates to any RLI. The RLI service must be configured to accept either uncompressed or compressed updates but not both. See the `rli_bloom-filter` setting of the RLS configuration file for more details.

There are tradeoffs between using uncompressed and compressed updates in your configuration. The advantage of using compressed updates, not surprisingly, is a significant reduction in network overhead and memory usage. As replica location mappings grow into the 10's of millions or more, the savings of using compressed updates becomes important. On the other hand, due to the compressed nature of the Bloom filter bitmap used to represent the logical names in the LRC, the wildcard query at the RLI cannot be supported when update compression is used.

6. Configuring the RLS Server for the WS MDS Index Service

The server package includes a script `$GLOBUS_LOCATION/libexec/aggrexec/globus-rls-aggregator-source.pl` that may be used as an Execution Aggregator Source by WS MDS. See [GT 4.2.0 Index Services](#) for more information on setting up and using the Execution Aggregator Source scripts in WS MDS. The script may be invoked as follows and will generate output in the format as depicted.

```
% $GLOBUS_LOCATION/libexec/aggrexec/globus-rls-aggregator-source.pl rls://mysite
<?xml version="1.0" encoding="UTF-8"?>
<rlsStats>
  <site>rls://mysite</site>
  <version>4.0</version>
  <uptime>03:08:15</uptime>
  <serviceList>
    <service>lrc</service>
    <service>rli</service>
  </serviceList>
  <lrc>
    <updateMethodList>
      <updateMethod>lfnlist</updateMethod>
      <updateMethod>bloomfilter</updateMethod>
    </updateMethodList>
    <updatesList>
      <updates>
        <site>rls://myothersite:39281</site>
        <method>bloomfilter</method>
        <date>08/01/05</date>
        <time>16:16:38</time>
      </updates>
    </updatesList>
    <numlfn>283902</numlfn>
    <numpfn>593022</numpfn>
    <nummap>593022</nummap>
  </lrc>
  <rli>
    <updatedViaList>
      <updatedVia>bloomfilters</updatedVia>
    </updatedViaList>
    <updatedByList>
      <updatedBy>
        <site>rls://myothersite:39281</site>
        <date>08/01/05</date>
        <time>10:03:21</time>
      </updatedBy>
    </updatedByList>
  </rli>
</rlsStats>
```

Important

Be sure to configure the security context of the container running the MDS, and be sure that the security configuration on the RLS host recognizes the MDS security context.

When following the instructions provided by the [GT 4.2.0 Index Services](#), you will need to consider the security context used by the MDS to invoke the Execution Aggregator Source script provided by RLS. Most deployments of RLS run the service with security enabled. Therefore any client connections, including administrative status operations, require authentication and authorization. In order for MDS to use the provided script to check RLS status, it must invoke the script with a valid user proxy or user certificate and key. The RLS must recognize the DN from the user certificate (i.e., the DN should be in the gridmap file).

One way to configure the MDS security context for use with RLS monitoring is to set the environment variables `X509_USER_CERT` and `X509_USER_KEY` to point to the container certificate and key. Run the MDS with these environment settings. Also, add the DN from the container certificate to the gridmap file on the host running the RLS.

Alternatively, you could modify the provided script so that it sets the environment variables to another user certificate and key (or proxy) as desired before calling the RLS.

7. Configuring the RLS Server for the MDS2 GRIS

The server package includes a program called `globus-rls-reporter` that will report information about an RLS server to the MDS2 GRIS. Use this procedure to enable this program:

1. To enable Index Service reporting, add the contents of the file `$GLOBUS_LOCATION/setup/globus/rls-ldif.conf` to the MDS2 GRIS configuration file `$GLOBUS_LOCATION/etc/grid-info-resource-ldif.conf`.
2. If necessary, set your virtual organization (VO) name in `$GLOBUS_LOCATION/setup/globus/rls-ldif.conf`. The default value is `local`. The VO name is referenced twice, on the lines beginning `dn:` and `args:`.
3. You must restart your MDS (GRIS) server after modifying `$GLOBUS_LOCATION/etc/grid-info-resource-ldif.conf`. You can use the following commands to do so:

```
$GLOBUS_LOCATION/sbin/SXXgris stop
$GLOBUS_LOCATION/sbin/SXXgris start
```

8. Complete RLS Server settings (globus-rls-server.conf)

This section describes the complete details of the RLS Server configuration settings.

Table 2.1. Complete RLS Server settings (globus-rls-server.conf)

| | |
|--|---|
| <pre>acl user: permission [permission]</pre> | <p>acl entries may be a combination of DNs and local usernames. If a DN is not found in the gridmap file then the file is used to search the acl list.</p> <p>A gridmap file may also be used to map DNs to local usernames, which in turn are matched against the regular expressions in the acl list to determine the user's permissions.</p> <p>user is a regular expression matching distinguished names (or local usernames if a gridmap file is used) of users allowed to make calls to the server.</p> <p>There may be multiple acl entries, with the first match found used to determine a user's privileges.</p> <p>[permission] is one or more of the following values:</p> <ul style="list-style-type: none"> • lrc_read Allows client to read an <i>LRC</i>. • lrc_update Allows client to update an LRC. • rli_read Allows client to read an <i>RLI</i>. • rli_update Allows client to update an RLI. • admin Allows client to update an LRC's list of RLIs to send updates to. • stats Allows client to read performance statistics. • all Allows client to do all of the above. |
| <pre>authentication true false</pre> | <p>Enable or disable GSI authentication.</p> <p>The default value is true.</p> <p>If authentication is enabled (true), clients should use the URL schema rls: to connect to the server.</p> <p>If authentication is <i>not</i> enabled (false), clients should use the URL schema rlsn:.</p> |
| <pre>db_pwd password</pre> | <p>Password to use to connect to the database server.</p> <p>The default value is changethis.</p> |
| <pre>db_user data-baseuser</pre> | <p>Username to use to connect to database server.</p> <p>The default value is dbperson.</p> |
| <pre>idletimeout seconds</pre> | <p>Seconds after which idle connections close.</p> <p>The default value is 900.</p> |
| <pre>loglevel N</pre> | <p>Sets loglevel to N (default is 0). Higher levels mean more verbosity.</p> |

| | |
|--|--|
| <code>lrc_bloomfilter_numhash N</code> | <p>Number of hash functions to use in <i>Bloom filters</i>.</p> <p>The default value is 3.</p> <p>Possible values are 1 through 8.</p> <p>This value, in conjunction with <code>lrc_bloomfilter_ratio</code>, will determine the number of false positives that may be expected when querying an RLI that is updated via Bloom filters.</p> <p><i>Note:</i> The default values of 3 and 10 give a false positive rate of approximately 1%.</p> |
| <code>lrc_bloomfilter_ratio N</code> | <p>Sets ratio of bloom filter size (in bits) to number of <i>LFNs</i> in the LRC catalog (in other words, size of the Bloom filter as a multiple of the number of LFNs in the LRC database.) This is only meaningful if Bloom filters are used to update an RLI. Too small a value will generate too many false positives, while too large a value wastes memory and network bandwidth.</p> <p>The default value is 10.</p> <p><i>Note:</i> The default values of 3 and 10 give a false positive rate of approximately 1%.</p> |
| <code>lrc_buffer_time N</code> | <p>LRC to RLI updates are buffered until either the buffer is full or this much time in seconds has elapsed since the last update.</p> <p>The default value is 30.</p> |
| <code>lrc_dbname</code> | <p>Name of LRC database.</p> <p>The default value is <code>lrcdb</code>.</p> |
| <code>lrc_server true false</code> | <p>If LRC server, the value should be <code>true</code>.</p> <p>The default value is <code>false</code>.</p> |
| <code>lrc_update_bf seconds</code> | <p>Interval in seconds between LRC to RLI updates when the RLI is updated by Bloom filters. In other words, how often an LRC server does a Bloom filter soft state update.</p> <p>This can be much smaller than the interval between updates without using Bloom filters (<code>lrc_update_ll</code>).</p> <p>The default value is 300.</p> |
| <code>lrc_update_factor N</code> | <p>If <code>lrc_update_immediate</code> mode is on, and the LRC server is in sync with an RLI server (an LRC and RLI are synced if there have been no failed updates since the last full soft state update), then the interval between RLI updates for this server (<code>lrc_update_ll</code>) is multiplied by the value of this option.</p> |
| <code>lrc_update_immediate true false</code> | <p>Turns LRC to RLI immediate mode updates on (<code>true</code>) or off (<code>false</code>).</p> <p>The default value is <code>false</code>.</p> |
| <code>lrc_update_ll seconds</code> | <p>Number of seconds before an LRC server does an LFN list soft state update.</p> <p>The default value is 86400.</p> |

| | |
|---|---|
| <code>lrc_update_retry seconds</code> | Seconds to wait before an LRC server will retry to connect to an RLI server that it needs to update. The default value is 300. |
| <code>maxbackoff seconds</code> | Maximum seconds to wait before re-trying listen in the event of an I/O error. The default value is 300 . |
| <code>maxfreethreads N</code> | Maximum number of idle threads. Excess threads are killed. The default value is 5 . |
| <code>maxconnections N</code> | Maximum number of simultaneous connections. The default value is 100. |
| <code>maxthreads N</code> | Maximum number of threads running at one time. The default value is 30. |
| <code>myurl URL</code> | URL of server. The default value is <code>rls://<hostname>:port</code> |
| <code>odbcini filename</code> | Sets environment variable ODBCINI. If not specified, and ODBCINI is not already set, then the default value is <code>\$GLOBUS_LOCATION/var/odbc.ini</code> . |
| <code>pidfile filename</code> | Filename where pid file should be written. The default value is <code>\$GLOBUS_LOCATION/var/<programname>.pid</code> . |
| <code>port N</code> | Port the server listens on. The default value is 39281 . |
| <code>result_limit limit</code> | Sets the maximum number of results returned by a query. The default value is 0 (zero), which means no limit. If a query request includes a limit greater than this value, an error (GLOBUS_RLS_BADARG) is returned. If the query request has no limit specified, then at most <code>result_limit</code> records are returned by a query. |
| <code>rli_bloomfilter true false</code> | RLI servers must have this set to accept Bloom filter updates. If <code>true</code> , then only Bloom filter updates are accepted from LRCs. If <code>false</code> , full LFN lists are accepted. <i>Note:</i> If Bloom filters are enabled, then the RLI does <i>not</i> support wildcarded queries. |

| | |
|--|---|
| <code>rli_bloomfilter_dir</code> <code>none default path-</code> <code>name</code> | <p>If an RLI is configured to accept bloom filters (<code>rli_bloomfilter true</code>), then Bloom filters may be saved to this directory after updates.</p> <p>This directory is scanned when an RLI server starts up and is used to initialize Bloom filters for each LRC that updated the RLI.</p> <p>This option is useful when you want the RLI to recover its data immediately after a restart rather than wait for LRCs to send another update.</p> <p>If the LRCs are updating frequently, this option is unnecessary and may be wasteful in that each Bloom filter is written to disk after each update.</p> <ul style="list-style-type: none"> • <code>none</code> <p>Bloom filters are not saved to disk.</p> <p>This is the default.</p> • <code>default</code> <p>Bloom filters are saved to the default directory:</p> <ul style="list-style-type: none"> • <code>\$GLOBUS_LOCATION/var/rls-bloomfilters</code> if <code>GLOBUS_LOCATION</code> is set • <code>else, /tmp/rls-bloomfilters</code> • <code>pathname</code> <p>Bloom filters are saved to the named directory.</p> <p>Any other string is used as the directory name unchanged.</p> <p>The Bloom filter files in this directory have the name of the URL of the LRC that sent the Bloom filter, with slashes(/) changed to percent signs (%) and ".bf" appended.</p> |
| <code>rli_dbname</code> <code>database</code> | <p>Name of the RLI database.</p> <p>The default value is <code>rlidb</code>.</p> |
| <code>rli_expire_int</code> <code>seconds</code> | <p>Interval (in seconds) between RLI expirations of stale entries. In other words, how often an RLI server will check for stale entries in its database.</p> <p>The default value is <code>28800</code>.</p> |
| <code>rli_expire_stale</code> <code>seconds</code> | <p>Interval (in seconds) after which entries in the RLI database are considered stale (presumably because they were deleted in the LRC).</p> <p>The default value is <code>86400</code>.</p> <p>This value should be no smaller than <code>lrc_update_ll</code>.</p> <p>Stale RLI entries are not returned in queries.</p> <p><i>Note:</i> If the LRC server is responding, this value is not used. Instead the value of <code>lrc_update_ll</code> or <code>lrc_update_bf</code> is retrieved from the LRC server, multiplied by 1.2, and used as the value for this option.</p> |

| | |
|--|---|
| <code>rli_server</code> <code>true false</code> | If an RLI server, the value should be <code>true</code> . The default value is <code>false</code> . |
| <code>rlscertfile filename</code> | Name of the X.509 certificate file identifying the server. This value is set by setting environment variable <code>X509_USER_CERT</code> . |
| <code>rlskeyfile filename</code> | Name of the X.509 key file for the server. This value is set by setting environment variable <code>X509_USER_KEY</code> . |
| <code>startthreads N</code> | Number of threads to start initially. The default value is 3. |
| <code>timeout seconds</code> | Timeout (in seconds) for calls to other RLS servers (e.g., for LRC calls to send an update to an RLI). |

Chapter 3. Starting the RLS Server

The RLS Server may be started by directly invoking the executable and passing it command-line arguments or by using the startup script. For complete details on starting the server directly, refer to [globus-rls-server\(1\)](#). We provide a startup script found at `$GLOBUS_LOCATION/sbin/SXXrls`, which you may use in places like `/etc/init.d` for automated startup and shutdown of the server. You may start the RLS Server by using this script as follows:

```
% $GLOBUS_LOCATION/sbin/SXXrls start
```

In addition to the startup script, the RLS Server may be started by using the executable as follows:

```
% $GLOBUS_LOCATION/bin/globus-rls-server [-d] [-L log-level]
```

We provide the following tips from the troubleshooting section of this guide:

- [Section 20, “Important notes on RLS initialization”](#)

Chapter 4. Stopping the RLS Server

The RLS Server may be stopped by directly issuing a KILL signal, by using the RLS Admin client, or by using the startup script.

If you have started the RLS Server directly, using `globus-rls-server(1)`, and you have not detached the process, you may issue a CTRL-C in the terminal window. If the server is running detached or was started using the startup script, you may issue a KILL signal to the `globus-rls-server` process. The RLS Server will catch the KILL signal and safely shutdown.

Using the RLS Admin client (see `globus-rls-admin(1)` for details), you may shutdown the RLS Server using the following command (this is also effective for shutting down remote RLS Servers):

```
% $GLOBUS_LOCATION/bin/globus-rls-admin -q rls[n]://hostname:port
```

Using the RLS Server startup script found at `$GLOBUS_LOCATION/sbin/SXXrls`, you may shutdown the RLS Server using the following command:

```
% $GLOBUS_LOCATION/sbin/SXXrls stop
```

Chapter 5. Testing

You can use the programs `globus-rls-admin` and `globus-rls-cli` to test functionality. See their respective man pages for details on their use.

1. Start the server in debug mode with the command:

```
$GLOBUS_LOCATION/bin/globus-rls-server [-N] -d -L 3
```

The `-N` option is helpful: if you do not have a host certificate for the server host, or a user certificate for yourself, it disables authentication.

2. Ping the server using `globus-rls-admin`:

```
$GLOBUS_LOCATION/bin/globus-rls-admin -p rls://serverhost
```

If you disabled authentication (by starting the server with the `-N` option), then use this command:

```
$GLOBUS_LOCATION/bin/globus-rls-admin -p rlsn://serverhost
```

Chapter 6. Security Considerations

1. Replica Location Service (RLS) Security Considerations

Security recommendations include:

- *Dedicated User Account:* It is recommended that users create a dedicated user account for installing and running the RLS service (e.g., `globus` as recommended in the general GT installation instructions). This account may be used to install and run other services from the Globus Toolkit.
- *Key and Certificate:* It is recommended that users do not use their `hostkey` and `hostcert` for use by the RLS service. Create a `containerkey` and `containercert` with permissions 400 and 644 respectively and owned by the `globus` user. Change the `rlskeyfile` and `rlscertfile` settings in the RLS configuration file (`$GLOBUS_LOCATION/etc/globus-rls-server.conf`) to reflect the appropriate filenames.
- *LRC and RLI Databases:* Users must ensure security of the RLS data as maintained by their chosen database management system. Appropriate precautions should be made to protect the data and access to the database. Such precautions may include creating a user account specifically for RLS usage, encrypting database users' passwords, etc.
- *RLS Configuration:* It is recommended that the RLS configuration file (`$GLOBUS_LOCATION/etc/globus-rls-server.conf`) be owned by and accessible only by the dedicated user account for RLS (e.g., `globus` account per above recommendations). The file contains the database user account and password used to access the LRC and RLI databases along with important settings which, if tampered with, could adversely affect the RLS service.

Chapter 7. Troubleshooting

General information on troubleshooting can be found in the [FAQ](#)¹. For a list of common errors in GT, see [Error Codes](#).

1. Errors

Table 7.1. Replica Locator Service (RLS) Errors

| Error Code | Definition | Possible Solutions |
|---|--|--|
| Error with credential: The proxy credential: <credential> with subject: <subject> expired <minutes> minutes ago | Expired proxy credential | Create a new proxy with grid-proxy-init . |
| Unable to connect to localhost:xxxx | Unable to connect to the local host. This can be due to a variety of reasons, including a wrong address or port number in the RLS connection URL or an issue with a firewall configuration. | <ul style="list-style-type: none"> • Double-check the address and port number in the RLS connection URL. parameters are correct. • If a firewall configuration is preventing connections to the target host for a particular port, you may need to consult the system administrator. |
| "connection timeout" | <p>At times, a client may experience a connection timeout when interacting with the RLS server due to a variety of reasons:</p> <ul style="list-style-type: none"> • One reason could simply be due to wide-area network latency or congestion. • Another situation that users eventually encounter is due to scaling of the system. As the RLS server's database of replica location mappings grows in size, some query operations, such as bulk queries involving large quantities of mappings or wildcard queries that result in a large subset of mappings, will begin to take more time both to process the query and to return the large results set to the client over the network. | <p>If timeouts are experienced with increasing frequency, increase the RLS server's timeout configuration parameter found in the <code>\$GLOBUS_LOCATION/var/globus-rls-server.conf</code> file. You may also use the <code>-t</code> timeout option of the globus-rls-cli tool.</p> |

¹ http://www.globus.org/toolkit/data/rls/rls_faq.html

2. I need help installing MySQL DBMS

There is simply too much information on DBMS installation for us to provide useful tips. In general, a binary package with a native installer (e.g., RPM, PKG, etc.) is the best option if you have root privileges. A binary package without a native installer is an excellent option if you do not have root privileges. And if all else fails, you can always try source installers. We have found most source installers of the DBMS to be reliable.

We have used RLS with MySQL versions 4.0.x, 4.1.x, and 5.0.x.

3. How can I check if MySQL is installed correctly?

There are probably numerous things you could do, but one simple check is to see if you can log on to the database. Of course, this requires that you have a user account. Assuming you installed the DBMS yourself, you should be able to log on as root:

```
% /path/to/mysql/bin/mysql -u root -p
Enter password:
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 64 to server version: 4.0.18-standard-log

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.
```

```
mysql> show databases;
+-----+
| Database          |
+-----+
| mysql             |
| test              |
+-----+
2 rows in set (0.03 sec)

mysql> quit
Bye
```

If you did not install the DBMS yourself but were given an account on a remote server, you should be able to log on using your user account:

```
% /path/to/mysql/bin/mysql -h hostname -P 3306 -u dbuser -p
Enter password:
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 66 to server version: 4.0.18-standard-log

Type 'help;' or '\h' for help. Type '\c' to clear the buffer.

mysql>
```

4. I need help installing PostgreSQL DBMS

There is simply too much information on DBMS installation for us to provide useful tips. In general, a binary package with a native installer (e.g., RPM, PKG, etc.) is the best option if you have root privileges. A binary package without a native installer is an excellent option if you do not have root privileges. And if all else fails, you can always try source installers. We have found most source installers of the DBMS to be reliable.

We have used RLS with PostgreSQL versions 7.4.x and it seems to work well with 8.x also.

5. How can I check if PostgreSQL is installed correctly?

There are probably numerous things you could do, but one simple check is to see if you can log on to the DBMS. Of course, this requires that you have a user account. Assuming you installed the DBMS yourself, you should be able to log on as root:

```
% /path/to/pgsql/bin/psql
Welcome to psql 7.4.6, the PostgreSQL interactive terminal.
```

```
Type:  \copyright for distribution terms
        \h for help with SQL commands
        \? for help on internal slash commands
        \g or terminate with semicolon to execute query
        \q to quit
```

```
user=#
```

If you did not install the DBMS yourself but were given an account on a remote server, you should be able to log on using your user account:

```
% /path/to/pgsql/bin/psql -h hostname -p 5432 -U dbuser -W dbname
Password:
Welcome to psql 7.4.6, the PostgreSQL interactive terminal.
```

```
Type:  \copyright for distribution terms
        \h for help with SQL commands
        \? for help on internal slash commands
        \g or terminate with semicolon to execute query
        \q to quit
```

```
dbname=>
```

6. I cannot connect to PostgreSQL

If you cannot connect to PostgreSQL but you know that the process is alive, you may have forgotten to use the `-i` when starting the DBMS. This option instructs `postmaster` to listen on TCP/IP socket. Without it, `postmaster` will not accept socket connections from clients.

7. Performance of RLS with PostgreSQL has degraded over time

Be sure to do periodic `vacuum` and `analyze` commands on all your PostgreSQL databases. The PostgreSQL documentation recommends doing this daily from `cron`. Failure to do this can seriously degrade performance, to the point where routine RLS operations (such as LRC to RLI soft state updates) timeout and fail. Please see the PostgreSQL documentation for further details.

8. I need help installing SQLite embedded database

Please be aware that we have relatively limited exposure to using RLS with SQLite than other database management systems. As of the GT 4.2 development branch, we have added SQLite setup scripts to create the RLS schema in a sqlite embedded database. Our testing has been limited to various Linux flavors on 32-bit platforms.

You are encouraged to review the SQLite site for authoritative details on installation of `sqlite3`. The following commands have been used by us to install `sqlite-3.3.6`:

```
% tar zxvf sqlite-3.3.6.tar.gz
% mkdir bld
% cd bld
% ../sqlite-3.3.6/configure --prefix=/path/to/install --enable-threadsafe
% make
% make install
```

9. How can I check if SQLite is installed correctly?

There are probably numerous things you could do, but one simple check is to see if you can open a test database. You should be able to open a test database and perform various SQL operations:

```
% /path/to/sqlite/bin/sqlite3 test.db
SQLite version 3.3.6
Enter ".help" for instructions
sqlite>
```

10. I need help installing iODBC

The `iodbc.org` website contains information on installation of the iODBC (or `libiodbc`) library. We have used the following commands to install iODBC version 3.52.4 in our development and test sites. From the `libiodbc` build directory:

```
% ./configure --prefix=/path/to/libiodbc-3.52.4 \
--with-iodbc-inidir=/path/to/etc --disable-gui --disable-gtktest
```

```
% make
% make install
```

The configure parameters are optional. The defaults enable the features required by RLS, namely pthreads. We prefer to install to a specified directory due to testing against multiple ODBC managers, however, this would not apply to most end users. We also prefer not to install the GUI features.

Keep in mind, if you install the ODBC Manager library to a non-standard location, you will need to update your library path (e.g., `LD_LIBRARY_PATH`) to reflect your non-standard installation.

11. I need help installing unixODBC

The `unixodbc.org` website contains information on installation of the unixODBC library. We have used the following commands to install unixODBC version 2.2.11 in our development and test sites. From the unixODBC build directory:

```
% ./configure --prefix=/path/to/unixODBC --disable-gui \
  --disable-drivers
% make
% make install
```

The configure parameters are optional. The defaults enable the features required by RLS, namely threads. We prefer to install to a specified directory due to testing against multiple ODBC managers, however, this would not apply to most end users. We also prefer not to install the GUI features or the supplied drivers.

Keep in mind, if you install the ODBC Manager library to a non-standard location, you will need to update your library path (e.g., `LD_LIBRARY_PATH`) to reflect your non-standard installation.

12. Building unixODBC failed on OS X

Building unixODBC-2.2.11 on OS X failed for us too. Our unixODBC tarball came from the `unixodbc.org` project site. It may be that other sources of this tarball may have been ported properly. What we found by doing a brief search was that making a small change to one of the files corrected the problem. In the file `Drivers/txt/SQLTables.c`, we made the following change from:

```
#ifdef OSXHEADER
#include <i386/types.h>
#endif
```

to:

```
/*
#ifdef OSXHEADER
#include <i386/types.h>
#endif
*/
```

13. I need help installing MySQL Connector/ODBC

The MySQL project/company provide several binary packages for several major platforms using a variety of native installers. We recommend using this approach if possible, though it may require root privileges and help from your system administrator. There are also binary packages available that do not use native installers, which can be used without root privileges and provide a better option than source installation.

From what we can tell, the 3.51.12 version of MySQL Connector/ODBC from the binary bundles provided by MySQL will work with unixODBC-2.2.11 but not with libiodbc-3.52.4 or earlier. The last known combinations of drivers from MySQL that worked with iODBC that we know of are:

- MyODBC-3.51.06 and libiodbc-3.51.1
- MyODBC-3.51.06 and libiodbc-3.51.2

Unfortunately, MySQL does not provide links to packages of MyODBC-3.51.06 as far as we can tell. You may find mirrors of the MySQL site that still have this driver available by searching the web for "myodbc 3.51.06". We found this one.²

13.1. MySQL Connector/ODBC 3.51.12 (binary)

With `mysql-connector-odbc-3.51.12-linux-i686`, we have used the following commands to install the driver.

```
% tar -zxvf /path/tomysql-connector-odbc-3.51.12-linux-i686.tar.gz
```

We have successfully used this driver with unixODBC-2.2.11. However, when testing against libiodbc-3.52.4 using the `iodbctest` utility, the driver failed with the following error.

```
% ./bin/iodbctest "DSN=lrc1000_mysql_3_51_12;UID=dbpassword;PWD:dbpassword"
iODBC Demonstration program
This program shows an interactive SQL processor
Driver Manager: 03.52.0406.0126
1: SQLDriverConnect = [MySQL][ODBC 3.51 Driver]Invalid window handle for
connection completion argument. (0) SQLSTATE=IM008
1: ODBC_Connect = [MySQL][ODBC 3.51 Driver]Invalid window handle for connection
completion argument. (0) SQLSTATE=IM008
```

Have a nice day.

We did not "Have a nice day" after this. There is limited information on this error, though it did appear that the driver loaded and attempted to make a connection.

² <http://sunsite.mff.cuni.cz/MIRRORS/ftp.mysql.com/downloads/api-myodbc-3.51.html>

13.2. MySQL Connector/ODBC 3.51.12 (source)

We attempted building and installing the mysql-connector-odbc-3.51.12 driver from the source package. We attempted building against unixODBC and libiodbc libraries and mysql-4.0.18 and mysql-5.0.21. We used a variety of configure options. In the end, we encountered configure or compiler errors that prevented us from using the source package.

13.3. MyODBC 3.51.06 (binary)

With MyODBC-3.51.06-pc-linux-gnu-i686, we have used the following commands to install the driver.

```
% tar -zxvf /path/to/MyODBC-3.51.06-pc-linux-gnu-i686
```

We have successfully used this driver with unixODBC-2.2.11 and libiodbc-3.52.4.

13.4. MyODBC 3.51.06 (source) and unixODBC

With MyODBC-3.51.06, we have used the following commands to install myodbc in our development and test sites. From the MyODBC-3.51.06 build directory:

```
% ./configure --prefix=/path/to/odbc/Drivers/MyODBC-3.51.06-unixODBC \  
--with-unixODBC --with-unixODBC-includes=/path/to/odbc/unixODBC-2.2.11/include \  
--with-unixODBC-libs=/path/to/odbc/unixODBC-2.2.11/lib \  
--with-odbc-ini=/path/to/odbc/etc/odbc.ini \  
--with-mysql-includes=/path/to/mysql/include --with-mysql-libs=/path/to/mysql/lib \  
--enable-thread-safe --without-samples \  
% gmake \  
% gmake install
```

Note: when building against unixODBC you must use the `--with-unixODBC` flag even if you use the `--with-unixODBC-includes=DIR` and `--with-unixODBC-libs=DIR` parameters.

You must use the `--enable-thread-safe` flag.

Also, be prepared to see a large number of compiler warnings when building this library.

13.5. MyODBC 3.51.06 (source) and iODBC

With MyODBC-3.51.06, we have used the following commands to install myodbc in our development and test sites. From the MyODBC-3.51.06 build directory:

```
% ./configure --prefix=/path/to/odbc/Drivers/MyODBC-3.51.06-iODBC \  
--with-iodbc=/path/to/libiodbc-3.51.2 --with-odbc-ini=/path/to/odbc/etc/odbc.ini \  
--with-mysql-includes=/path/to/mysql/include --with-mysql-libs=/path/to/mysql/lib \  
--enable-thread-safe --without-samples \  
% gmake \  
% gmake install
```

Note: when building against iODBC you must use the `--with-iodbc=DIR` parameter. The `--with-iodbc-
includes=DIR` and `--with-iodbc-libs=DIR` parameters DO NOT work.

You must use the `--enable-thread-safe` flag.

Also, be prepared to see a large number of compiler warnings when building this library.

14. I need help installing psqLODBC

The README file or docs directory included with the psqLODBC package contains information on installation of the psqlodbc library. We have noticed some differences between building and installing psqlodbc-7.2.5, psqlodbc-08.00.0102, and psqlodbc-08.01.0200 against iODBC and unixODBC.

Note: In earlier versions of this guide, we warned that the `--with-iodbc` configure option resulted in an unstable psqlodbc library that caused memory corruption and failure (e.g., core dump) when an application (RLS Server) attempted to open the database connection. We have not experienced this problem with the more recent versions of psqlodbc that we have tested and which are documented below.

14.1. psqlodbc-7.2.5 (source) and unixODBC-2.2.11

With psqlodbc-7.2.5, we have used the following commands to install psqlodbc in our development and test sites. From the psqlodbc build directory:

```
% ./configure --prefix=/path/to/odbc/Drivers/psqlodbc-7.2.5-unixODBC \  
--enable-pthreads --with-unixodbc  
% gmake  
% gmake install
```

You may also want to consider setting CPPFLAGS and LDFLAGS if you install your ODBC Manager in a non-standard location.

14.2. psqlodbc-7.2.5 (source) and libiodbc-3.52.4

With psqlodbc-7.2.5, we have used the following commands to install psqlodbc in our development and test sites. From the psqlodbc build directory:

```
% ./configure --prefix=/path/to/odbc/Drivers/psqlodbc-7.2.5-iODBC \  
--enable-pthreads --with-iodbc  
% gmake  
% gmake install
```

You may also want to consider setting CPPFLAGS and LDFLAGS if you install your ODBC Manager in a non-standard location.

14.3. psqlodbc-08.00.0102 (source) and unixODBC-2.2.11

With psqlodbc-08.00.0102, we have used the following commands install psqLODBC with a unixODBC-2.2.11 installation. From the psqlodbc build directory:

```
% ./configure --prefix=/path/to/odbc/Drivers/psqlodbc-08.00.0102-unixODBC \  
  --enable-pthreads --with-unixodbc LDFLAGS="-L/path/to/odbc/unixODBC-2.2.11/lib" \  
  CPPFLAGS="-I/path/to/odbc/unixODBC-2.2.11/include"  
% gmake  
% gmake install
```

14.4. psqlodbc-08.00.0102 (source) and libiodbc-3.52.4

With psqlodbc-08.00.0102, we have used the following commands to configure psqlodbc but it failed to build. We have no advice on resolving this problem other than to suggest using a different combination of ODBC Manager and ODBC Driver.

```
% ./configure --prefix=/path/to/odbc/Drivers/psqlodbc-08.00.0102-iODBC \  
  --enable-pthreads --with-iodbc LDFLAGS=-L/path/to/odbc/libiodbc-3.52.4/lib \  
  CPPFLAGS=-I/path/to/odbc/libiodbc-3.52.4/include  
% gmake  
...ends in errors...
```

14.5. psqlodbc-08.00.0102 (source) and libiodbc-3.51.2

With psqlodbc-08.00.0102, we have used the following commands to configure psqlodbc and to build it successfully. In this case, we used psqlodbc along with libiodbc-3.51.2.

```
% ./configure --prefix=/path/to/odbc/Drivers/psqlodbc-08.00.0102-iODBC \  
  --enable-pthreads --with-iodbc LDFLAGS=-L/path/to/odbc/libiodbc-3.51.2/lib \  
  CPPFLAGS=-I/path/to/odbc/libiodbc-3.51.2/include  
% gmake  
% gmake install
```

14.6. psqlodbc-08.01.0200 (source)

We have NOT succeeded installing psqlodbc-08.01.0200 from source with either iODBC or unixODBC.

15. I need help installing SQLite ODBC Driver

The README file included with the SQLite ODBC Driver package contains information on installation of the sqliteodbc library. Please be aware that the author of the SQLite ODBC Driver states that the package is of experimental quality.

15.1. sqliteodbc-0.69 (source) and iODBC-3.51.2

With sqliteodbc-0.69, we have used the following commands to install sqliteodbc in our development and test sites. From the sqliteodbc build directory:

```
% ./configure --prefix=/path/to/install --disable-winterface --with-sqlite3=/path/to/sqlite3
```

```
% make
% make install
```

Note the distinction between using `--with-sqlite=DIR` and `--with-sqlite3=DIR`. Use the latter for the sqlite3 database library (recommended).

16. Which libraries do you recommend to new users?

The short answer, consider the following options:

- PostgreSQL (recent stable) with `psqlODBC (08.00.0102)` with `unixODBC (2.2.11)`
- MySQL (recent stable) with MySQL Connector/ODBC (3.51.12, binary package) with `unixODBC (2.2.11)`

The long answer follows. First of all, if you are a new user and not confident installing system level software, you will save yourself a lot of time and trouble by consulting with your system administrator. If you are an administrator and you still find this troubling, then you are in good company.

As of the writing of this guide, we have experienced a higher success rate using `unixODBC` (current stable release at this time is version 2.2.11) than we have with `iODBC` (current stable release is version 3.52.4).

As of the writing of this guide, we have found PostgreSQL and `psqlODBC` installation from source to be reliable. We would recommend the use of `psqlODBC` version 08.00.0102 with `unixODBC`. If you are an `iODBC` user, you will need to use `psqlODBC 7.2.5`.

As of the writing of this guide, we have found MySQL and MySQL Connector/ODBC installation from binary to be reliable. We would recommend the use of MySQL Connector/ODBC version 3.51.12 with `unixODBC`. If you are a `iODBC` user, you will need to find a binary of `MyODBC 3.51.06`.

17. I need help configuring my DSNs

There are a variety of ways that you can specify the Data Source Names (DSNs) in the `odbc.ini` file. You may use a GUI or command-line tool provided with your ODBC Manager (`unixODBC` or `iODBC`). We tend to edit the `odbc.ini` file directly when working in our development and test sites, however, you should avoid this practice unless you are experienced with using ODBC tools (a small typo such as an unwanted space character can prevent your configuration from working -- with little in the way of useful error messages).

The following is an example of a typical setup when using PostgreSQL:

```
[lrc1000]
Driver      = /path/to/lib/psqlodbc.so
Database    = lrc1000
Servername  = hostname
Port        = 5432
```

```
[rli1000]
Driver      = /path/to/lib/psqlodbc.so
Database    = rli1000
Servername  = hostname
```

Port = 5432

The following is an example of a typical setup when using MySQL:

```
[lrc1000]
Driver      = /path/to/lib/libmyodbc3.so
Database    = lrc1000
SERVER      = hostname
PORT        = 3306
SOCKET      = /path/to/mysql/mysql.sock
```

```
[rli1000]
Driver      = /path/to/lib/libmyodbc3.so
Database    = rli1000
SERVER      = hostname
PORT        = 3306
SOCKET      = /path/to/mysql/mysql.sock
```

The following is an example of a typical setup when using SQLite:

```
[lrc1000]
Driver      = /path/to/lib/libsqlite3odbc.so
Database    = /path/to/globus-rls-lrc1000.db

[rli1000]
Driver      = /path/to/lib/libsqlite3odbc.so
Database    = /path/to/globus-rls-rli1000.db
```

18. How can I tell which odbc.ini is being used?

The ODBC libraries look for the odbc.ini file along a search path. We are not experts in this matter and it seems that it varies from vendor to vendor over time. But to the best of our knowledge the general search path is in the following order:

- According to environment variable `$ODBCINI`
- In the user's home at `~/odbc.ini`
- In the system default `/etc/odbc.ini`

When we are troubleshooting a system, we tend to use the `strace` utility to see exactly which files it is reading.

19. Testing your ODBC configuration

The commands below demonstrate the `show tables` command with a MySQL database. Use the equivalent of this command if you are using a different DBMS. Alternately, you may simply perform a few `SELECT` queries against the listed tables.

Testing with unixODBC, run:

```
% /path/to/unixODBC-2.2.11/bin/isql lrc1000 dbuser dbpassword
```

```
+-----+
| Connected!
|
| sql-statement
| help [tablename]
| quit
|
+-----+
```

```
SQL> show tables;
```

```
+-----+
| Tables_in_lrc1000|
+-----+
| t_attribute
| t_date_attr
| tflt_attr
| t_int_attr
| t_lfn
| t_map
| t_pfn
| t_rli
| t_rlipartition
| t_str_attr
|
+-----+
```

```
SQLRowCount returns 10
```

```
10 rows fetched
```

```
SQL> quit
```

Testing with iODBC, run:

```
% /path/to/bin/iodbctest "DSN=lrc1000;UID=dbuser;PWD=dbpassword"
```

```
iODBC Demonstration program
```

```
This program shows an interactive SQL processor
```

```
Driver Manager: 03.51.0002.0224
```

```
Driver: 03.51.06
```

```
SQL>show tables;
```

```
Tables_in_lrc1000
```

```
-----
t_attribute
t_date_attr
tflt_attr
t_int_attr
t_lfn
t_map
t_pfn
t_rli
t_rlipartition
t_str_attr
```

```
result set 1 returned 10 rows.
```

```
SQL>quit
```

```
Have a nice day.
```

20. Important notes on RLS initialization

Please be advised (and advise other users responsible for bringing up the RLS) that the startup initialization may take a few minutes before the RLS may be accessible. The initialization involves two key operations that may consume significant resources causing the server to appear temporarily unresponsive. Users of RLS may mistakenly assume that RLS failed to startup and may kill the server and start over. Some users may fall into this in a repeated cycle, believing that the RLS is unable to startup properly.

If the RLS is configured to send compressed updates (Bloom filters) to other RLIs, the RLS startup will involve initialization of the Bloom filter representing the current contents of the local replica catalog (LRC). This step is a prerequisite before any additional operations may be allowed, therefore no client connections are permitted until the initialization is complete. In our test environment, we have seen over 30 seconds delay due to creation of the Bloom filter corresponding to 1 million LFN names on a system with Dual 1 GHz CPU and 1.5 GB RAM. You may experience greater delays at larger scales and/or when running RLS with more limited system resources.

If the RLS is configured to send uncompressed updates (LFN lists) to other RLIs, the RLS startup will not involve any additional initialization delay. However, the RLS will spawn an initial full catalog update to all RLIs it updates. Though these updates will take place on separate threads of execution after the initialization of the system, they will consume a great amount of processor activity. Depending on the volume of the local replica catalog (LRC), this processor activity may initially interfere with a client operation. In our test environment, we have seen our initial "globus-rls-admin ping..." operation may suffer a delay and timeout in 30 seconds, the second "ping" may delay for a few seconds but will successfully return, and the third and every subsequent "ping" operation will successfully return immediately throughout the duration of the update. The system exhibits the same behavior for any other client operation, such as a `globus-rls-cli query...` operation.

21. Linux kernel and glibc incompatibility

This note originated based on observations of RLS Server running on RedHat 9 but could also apply to other Linux distributions. There have been occurrences of RLS servers hanging on RedHat 9 systems. The external symptoms are:

- The server does not accept new connections from clients, with an error message similar to:

```
connect(rls://XXXXX): globus_rls_client: IO timeout:
globus_io_tcp_register_connect() timed out after 30 seconds
```

- Often, the server continues to receive and send updates as configured and respond to signals. You can check this by querying other servers that interact with the one that's hung. Under `gdb`: All the server threads are waiting to be signaled on a condition variable. Sometimes, this is in `globus_io` functions, particularly in `globus_io_cancel()`.

21.1. Probable cause

This seems to be due to a problem in the new kernel and thread libraries of RedHat 9. A problem in `pthread_cond_wait()` causes threads not to wake up correctly. This problem has been seen with the following kernels and glibc packages:

- Kernels: 2.4.20-30.9 and 2.4.20-8
- glibc: glibc-2.3.2-27.9.7

21.2. Suggested workaround

The problems don't seem to arise when RLS is linked with older pthread libraries. This can be done as by adding a couple of lines to the RLS startup script in `$GLOBUS_LOCATION/sbin/SXXrls`, as shown:

```
#!/bin/sh

GLOBUS_LOCATION=/opt/gt3.2
MYSQL=/opt/mysql
IODBC=/opt/iodbc

export GLOBUS_LOCATION

#RedHat 9 workaround
LD_ASSUME_KERNEL=2.4.1
export LD_ASSUME_KERNEL
...
```

On i586 systems, set:

```
...
LD_ASSUME_KERNEL=2.2.5
...
```

Chapter 8. Usage statistics collection by the Globus Alliance

1. RLS-specific usage statistics

The following usage statistics are sent by RLS Server by default in a UDP packet:

- Component identifier
- Usage data format identifier
- Time stamp
- Source IP address
- Source hostname (to differentiate between hosts with identical private IP addresses)
- Version number
- Uptime
- *LRC* service indicator
- *RLI* service indicator
- Number of *LFNs*
- Number of *PFNs*
- Number of Mappings
- Number of RLI LFNs
- Number of RLI LRCs
- Number of RLI Senders
- Number of RLI Mappings
- Number of threads
- Number of connections

The RLS sends the usage statistics at server startup, server shutdown, and once every 24 hours when the service is running.

If you wish to disable this feature, you can set the following environment variable before running the RLS:

```
export GLOBUS_USAGE_OPTOUT=1
```

By default, these usage statistics UDP packets are sent to `usage-stats.globus.org:4180` but can be redirected to another host/port or multiple host/ports with the following environment variable:

```
export GLOBUS_USAGE_TARGETS="myhost.mydomain:12345 myhost2.mydomain:54321"
```

You can also dump the usage stats packets to stderr as they are sent (although most of the content is non-ascii). Use the following environment variable for that:

```
export GLOBUS_USAGE_DEBUG=MESSAGES
```

Also, please see our [policy statement](#)¹ on the collection of usage statistics.

¹ ../../Usage_Stats.html

Glossary

B

Bloom filter Compression scheme used by the Replica Location Service (RLS) that is intended to reduce the size of soft state updates between Local Replica Catalogs (LRCs) and Replica Location Index (RLI) servers. A Bloom filter is a bit map that summarizes the contents of a Local Replica Catalog (LRC). An LRC constructs the bit map by applying a series of hash functions to each logical name registered in the LRC and setting the corresponding bits.

L

Local Replica Catalog (LRC) Stores mappings between logical names for data items and the target names (often the physical locations) of replicas of those items. Clients query the LRC to discover replicas associated with a logical name. Also may associate attributes with logical or target names. Each LRC periodically sends information about its logical name mappings to one or more RLIs.

See also [RLI](#)⁶.

logical file name A unique identifier for the contents of a file.

P

physical file name The address or the location of a copy of a file on a storage system.

R

Replica Location Index (RLI) Collects information about the logical name mappings stored in one or more Local Replica Catalogs (LRCs) and answers queries about those mappings. Each RLI periodically receives updates from one or more LRCs that summarize their contents.

⁶ #rli

Index

E

errors, 25

I

installing

from binary, 5

from source, 5

prerequisites, 1

database, 1

ODBC driver, 2

ODBC manager, 2

recommended software, 3

setting up the database, 6

configure settings, 7

create databases, 8

create user account, 8

DSNs, 9

which database?, 7

S

server operations

starting RLS server, 21

stopping RLS server, 22

T

testing

your installation, 23

troubleshooting

for admin, 25

U

updating

from source, 6

usage statistics, 39