

EGEE

R-GMA User Guide for Python Programmers

Document identifier:	EGEE-JRA1-TEC-503614
Date:	September 4, 2006
Activity:	JRA1: Middleware Engineering and Integration (UK Cluster)
Document status:	FINAL
Document link:	https://edms.cern.ch/document/503614/

Abstract: This document provides the Python programmer with the information necessary to get started with R-GMA.

Document Change Log

Issue	Date	Comment	Author
1.0	31 March 2005	First release	JRA1-UK
1.1	18 April 2006	First release, first revision	JRA1-UK

Document Change Record

Item	Reason for Change
Updated consumer example	To avoid hanging on isExecuting()
Added explanation to allowed SQL on duplicate columns	Needed to support chunking.
Added resilient consumer and producer examples	To aid users

Copyright ©Members of the EGEE Collaboration. 2004. See <http://eu-egEE.org/partners> for details on the copyright holders.

EGEE (“Enabling Grids for E-science in Europe”) is a project funded by the European Union. For more information on the project, its partners and contributors please see <http://www.eu-egEE.org>.

You are permitted to copy and distribute verbatim copies of this document containing this copyright notice, but modifying this document is not allowed. You are permitted to copy this document in whole or in part into other documents if you attach the following reference to the copied elements: “Copyright ©2004. Members of the EGEE Collaboration. <http://www.eu-egEE.org>”

The information contained in this document represents the views of EGEE as of the date they are published. EGEE does not guarantee that any information contained herein is error-free, or up to date.

EGEE MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, BY PUBLISHING THIS DOCUMENT.

CONTENTS

1	INTRODUCTION	6
1.1	PURPOSE AND STRUCTURE OF THIS DOCUMENT	6
1.2	R-GMA ARCHITECTURE	6
1.2.1	VIRTUAL DATABASE	6
1.2.2	WEB SERVICES	7
1.2.3	PRODUCERS	7
1.2.4	CONSUMERS	7
1.2.5	RETENTION PERIODS	8
1.2.6	RESOURCE FRAMEWORK AND THE TERMINATION INTERVAL	8
2	GETTING STARTED WITH R-GMA	9
2.1	PREREQUISITES	9
2.2	SETTING UP FOR USAGE	9
2.3	SIMPLE INTERACTION USING THE COMMAND LINE TOOL	10
2.4	SIMPLE PRODUCING AND CONSUMING INFO WITH THE COMMAND LINE TOOL	11
3	R-GMA INSTALLATION	12
3.1	INSTALLED COMPONENTS	12
3.2	CHECK YOUR INSTALLATION	12
3.3	SECURITY	13
4	PRIMARY PRODUCERS	13
4.1	TERMINATION INTERVAL	13
4.2	PRODUCER PROPERTIES	13
4.3	PRIMARY PRODUCER EXAMPLES	14
4.3.1	SIMPLE PRIMARY PRODUCER EXAMPLE	14
4.3.2	RUNNING THE EXAMPLE	15
4.3.3	RESILIENT PRIMARY PRODUCER EXAMPLE	15
5	CONSUMING INFORMATION	17
5.1	TYPES OF QUERY	17
5.2	CONSUMER EXAMPLES	17
5.2.1	SIMPLE CONSUMER EXAMPLE	17
5.2.2	CONSUMING CONTINUOUS PLUS OLD INFORMATION	18
5.2.3	ONE-OFF QUERIES	18
5.2.4	CONSUMER EXTRACTING INFORMATION FROM RESULT SET	19
5.2.5	RESILIENT CONSUMER EXAMPLE	19

6	REUBLISHING VIA SECONDARY PRODUCERS	21
6.1	SECONDARY PRODUCER EXAMPLES	21
6.1.1	SIMPLE SECONDARY PRODUCER EXAMPLE	21
6.1.2	AVOIDING A PERMANENT CONNECTION TO A SECONDARY PRODUCER	22
7	THE RGMA COMMAND LINE TOOL	23
7.1	INTRODUCTION	23
7.1.1	STARTING THE R-GMA COMMAND LINE TOOL	23
7.1.2	ENTERING COMMANDS	23
7.2	COMMANDS	24
7.2.1	GENERAL COMMANDS	24
7.2.2	QUERYING DATA	24
7.2.3	INSERTING DATA	24
7.2.4	SECONDARY PRODUCERS	25
7.2.5	INFORMATION COMMANDS	25
7.2.6	DIRECTED QUERIES	26
8	USING THE WEB TO BROWSE R-GMA INFORMATION	26
8.1	SECURITY	26
9	ADMINISTRATION	26
9.1	TABLE CREATION	26
9.2	RECOVERY FOLLOWING RESTART	26
10	SQL	27
10.1	CREATING A DATABASE AS STORAGE FOR A PRODUCER	27
10.2	EXAMPLES OF SQL QUERIES	27
10.3	SUPPORTED SQL	28
11	ADVICE ON USING R-GMA	28
11.1	GENERAL ADVICE	28
11.2	PRIMARY PRODUCERS	29
11.3	SECONDARY PRODUCERS	29
11.4	CONSUMERS	29
12	RELEASE NOTES FROM A USER PERSPECTIVE	30
12.1	OVERVIEW	30
12.2	NEW FEATURES	30
12.3	DEPRECATED FEATURES	31
12.4	FEATURES WITHDRAWN	32

13 KNOWN PROBLEMS AND CAVEATS	32
13.1 FUNCTIONALITY NOT YET IMPLEMENTED	32
13.2 OTHER KNOWN ISSUES	32
13.3 REPORTING BUGS AND GETTING HELP	33

1 INTRODUCTION

1.1 PURPOSE AND STRUCTURE OF THIS DOCUMENT

This document is intended to get people started with R-GMA. It is one of a set, with each member customised for a different programming language.

After this introduction there are sections explaining what should be done to ensure that R-GMA is correctly installed, how to publish information via a “Primary Producer”, how to get information back via a “Consumer”, how to set-up a “Secondary Producer” and how to use the command line and web based tools.

The APIs (in C, C++, Java and Python) are all described in detail in the documentation linked from <http://hepunix.rl.ac.uk/egEE/jra1-uk/glite-r1.5/>. In addition the documentation is all distributed with the software and may be found as `$RGMA_HOME/share/doc/<module>/manual.pdf`, where “module” identifies the document. Look at the directory `$RGMA_HOME/share/doc` to see the naming scheme.

Some brief release notes from a user’s perspective may be found in section 12 which is useful to anyone who is familiar with the previous version of R-GMA.


To understand more detail of what R-GMA is meant to do, you may choose to read the specification[1]. However we do not expect the average user to need the specification document.

This document contains a number of code examples. You can find a copy of these in the directory: `$RGMA_HOME/share/doc/rgma-base/examples` wherever R-GMA has been installed.

1.2 R-GMA ARCHITECTURE

1.2.1 VIRTUAL DATABASE

R-GMA is an implementation of the Grid Monitoring Architecture (GMA) proposed by the Global Grid Forum (GGF), which models the information infrastructure of a Grid as a set of *Consumers* (who request information), *Producers* (who provide information) and a single *Registry* (which mediates the communication between Producers and Consumers). R-GMA imposes a standard query language (a subset of SQL) on this model – so Producers publish *tuples* (database rows) with an SQL insert statement and Consumers query them using SQL select statements. R-GMA also ensures that all tuples carry a *time-stamp*, so that monitoring systems (which require time-sequenced data) are inherently supported.



R-GMA presents the information resources of a Virtual Organisation (VO)¹ as a single *virtual database* containing a set of *virtual tables*. As the picture above shows, a single² schema contains the name and structure (column names, types and settings) of each virtual table in the system. A single registry contains a list, for each table, of Producers who have offered to *publish* (provide data for) rows for the table. A Consumer runs an SQL query against a table, and the registry selects the best Producers to answer the query in a process called *mediation*. The Consumer then contacts each Producer directly, combines the information, and returns a set of tuples. The mediation process is hidden from the user. Note that there is no central repository holding the contents of the virtual table; it is in this sense, that the database is virtual.

1.2.2 WEB SERVICES

R-GMA will conform to the Web Services Architecture [4]. The version currently released uses servlet technology to run the services rather than WSDL-defined Web Services. When web services are introduced the APIs will be unchanged. Each service has a well defined set of *operations* that it can carry out, as specified in a machine-readable XML document conforming to the Web Services Description Language (WSDL [5]). All operations are requested by applications through an exchange of messages with the service. We provide APIs to offer a convenient way to interact with the services.

1.2.3 PRODUCERS

There are three classes of Producer: *Primary*, *Secondary* and *On-demand*. Each is created by a user application and returns tuples in response to queries from other user applications. The main difference is in where the tuples come from.

For a Primary Producer, user's code periodically inserts tuples which are then stored internally by the Producer. The Producer answers Consumer queries from this storage. The Secondary Producer service also answers queries from its internal storage, but it populates this storage itself by running its own query against the virtual table: the user code only sets the process running; the tuples come from other Producers. In the On-demand Producer, there is no internal storage; data is provided by the user code in direct response to a query forwarded on to it by the Producer service. No examples of the use of the On-demand Producer are included in this document.

Each virtual table has a *key* column (or group of columns) declared in the schema. Each tuple also carries a *time-stamp*, added by the Primary Producer when the tuple is first published into the system, which, together with the key columns, is similar to a primary key for the table. Tuples with the same key, but different values for the time-stamp, can also be thought of as different versions of the same tuple. R-GMA works consistently in UTC³.

1.2.4 CONSUMERS

Each Consumer represents a single SQL SELECT query on the virtual database. The request is initiated by user code, but the *Consumer service* carries out all of the work on its behalf. The query is first passed to the Registry to identify which Producers, for each virtual table in the query, must be contacted to answer it. This process is called *mediation*. The query is then passed by the Consumer service to each relevant Producer, to obtain the answer tuples directly.

¹Currently all VOs share the one namespace.

²Although there is only one logical schema and registry pair per VO, identical replicas will be made for scalability and robustness. This is discussed in detail later. Currently this is not implemented.

³UTC refers to Coordinated Universal Time, which used to be known as GMT.

There are four types of query: *continuous*, *latest*, *history* and *static*. The first three types (continuous, latest and history) can optionally take a *TimeInterval* parameter.

A continuous query (which can only act on a single table) causes all new tuples matching the query to be automatically streamed to the Consumer when they are inserted to the virtual table by a Producer. If a *TimeInterval* parameter is specified, all existing tuples newer than (*now - TimeInterval*) will additionally be returned when the query is first started.

A latest query is evaluated on that set of tuples, which for each table and key have the greatest time-stamp value and that have not exceeded their Latest Retention Period. In addition, if a *TimeInterval* is specified, only tuples that are newer than (*now - TimeInterval*) will be used. Whether or not you specify a time interval, the query will never use tuples that have exceeded their Latest Retention Period. The LatestRetention Period is described in 1.2.5.

A history query is evaluated over all available versions of tuples. The set can also be restricted by specifying a *TimeInterval*.

A static query is handled by an On-demand Producer like a normal one-off database query. There are no time-stamps or retention periods associated with static queries.

1.2.5 RETENTION PERIODS

To allow Primary and Secondary Producers to periodically purge “old” tuples, and to give a precise meaning to the “current state” for a latest query, *retention periods* are used.

The *LatestRetentionPeriod* defines how old a tuple can be before it should no longer be considered to be the latest. This time interval is added to the *time-stamp* and inserted into each tuple published by a Primary Producer, and remains there when a tuple is re-published by a Secondary Producer. In addition, Primary and Secondary Producers declare a *HistoryRetentionPeriod* for each table to which they are publishing tuples.

Primary and Secondary Producers therefore have two logical tuple-stores, one supporting latest-queries and the other supporting continuous and history queries. Producers undertake to retain the *most recent* version of any tuple which has not exceeded its LatestRetentionPeriod, and *all* versions of any tuple which have not exceeded the HistoryRetentionPeriod.

Primary and Secondary Producers all support continuous queries, however they may not support history and latest queries.

The HistoryRetentionPeriod may be longer or shorter than the LatestRetentionPeriod. The HistoryRetentionPeriod is a (per-table) property of the producer, whereas the LatestRetentionPeriod is a property of the tuple.

1.2.6 RESOURCE FRAMEWORK AND THE TERMINATION INTERVAL

Each instance of a Producer or Consumer in a running R-GMA system exists as a *resource* on a server. A resource contains the private data or threads associated with that particular instance (such as the tuple-storage and tuple-streamer in a Primary Producer), and it is created by an R-GMA service, when a user sends a “create” request. It resides on the server with the service, and is given an identifier which is passed back to the client. The client then includes the resource identifier with all subsequent requests relating to that instance (the API takes care of this). A resource is normally destroyed at the explicit request of the user, but in order to protect itself from an accumulation of redundant resources, an R-GMA service requires the user to specify a termination interval (in a restricted range) when it creates the resource. We plan to remove this from user control. To allow for easy migration the termination interval should be set to one hour. If the service doesn’t hear from the user for any period exceeding the

termination interval, the resource is destroyed. This is the concept of *soft-state registration*. It puts the onus on the user to keep the resource alive, by making periodic contact with the service, but any contact will do. A *showSignOfLife* call is provided which does nothing but establish this periodic contact. The registry protects itself in the same way, against Producers and Consumers which register then disappear, so a periodic keep-registered message has also to be sent to the registry, by the Consumer and Producer services (this is hidden from the user).

2 GETTING STARTED WITH R-GMA

Here we present one of the simplest forms of interaction with R-GMA, where no coding is required, looking at information already in the tables and adding some very simple information using the command line tool. More information on the command line tool is in section 7 of this document, and in the command line tool document, available from the JRA1-UK web page <http://hepunix.rl.ac.uk/egee/jra1-uk/glite-r1.5/> To carry out the simplest form of interaction, publishing and consuming information using the chosen language, see sections 4 and 5.

2.1 PREREQUISITES

In order to do anything in the EGEE grid, a user will need to

1. Obtain a Certificate from one of the Certificate Authorities in the EGEE PMA (see <http://www.eugridpma.org/>)
2. Register with a Virtual Organisation.
3. Obtain an account on a user interface machine.

Note that at present, it is possible to interact with R-GMA without being a member of a virtual organisation, as there is no authorization in R-GMA, but you will still need a certificate for authentication purposes.

2.2 SETTING UP FOR USAGE

You need to set up a few environment variables in order to use R-GMA client applications. On a gLite User Interface machine, these are defined for you by the `glite_setenv.sh` bash script. If you have not already done so, run this script by typing:

```
bash
```

followed by:

```
. /etc/glite/profile.d/glite_setenv.sh
```

at the shell's command prompt. If this is not found, try executing the `glite-env.sh` script, whose location depends on the installation. For example, enter:

```
. /afs/cern.ch/project/egee/jra1/prototype/ui/glite-env.sh
```

If this doesn't work either, see your system administrator.

You can now generate a grid proxy (for authentication) by typing:

```
grid-proxy-init
```

Note that `voms-proxy-init` can be used instead: `grid-proxy-init` will work until authorization is added to R-GMA after which `voms-proxy-init` must be used. Either way, this will generate a proxy file based on your uid and write it to `/tmp`. Check this file is present, owned by you and read/writable by you, by typing:

```
ls -l /tmp/x509*
```

You also need to tell R-GMA where to find the proxy file by defining the `X509_USER_PROXY` environment variable. The `glite_setenv.sh` script may have done this for you, but if not, type:

```
export X509_USER_PROXY=/tmp/x509up_uxxxxx
```

where `xxxxx` is replaced by your Unix uid.

To check the setup, type:

```
rgma-client-check
```

This should produce something like:

```
*** Running R-GMA client tests on lxplus063.cern.ch ***

Checking C API: Success
Checking C++ API: Success
Checking CommandLine API: Success
Checking Python API: Success
Checking Java API: Success

*** R-GMA client test successful ***
```

If not, see section 3 on installation or ask your system administrator for help.

2.3 SIMPLE INTERACTION USING THE COMMAND LINE TOOL

Enter:

```
rgma
```

This starts a mini shell and subsequent commands are typed into this. (Note that when using some commands table and column names are case sensitive, so it is best to assume everything is case sensitive.)

```
show tables
```

This should list the tables currently available. It is simple to show the column definitions in a table, try for example:

```
describe GlueSite
```

To see the sites enter:

```
select * from GlueSite
```

To see the services available enter:

```
select * from GlueService
```

To see the status of the various services enter:

```
select * from GlueServiceStatus
```

To see the values for specific fields enter

```
select Name,Description from GlueSite
```

To get a list of valid commands enter

```
help
```

To find out about the select command enter

```
help select
```

To exit the rgma command line interface enter:

```
exit
```

Note that you may also use 'q' or CTRL D to exit

2.4 SIMPLE PRODUCING AND CONSUMING INFO WITH THE COMMAND LINE TOOL

This shows how to place very simple information into one table, making it remain for at least 10 minutes and examining it. We use the pre-defined `userTable` which is available for people to use for tests. Enter the following to place the information into the table:

```
rgma
show tables
describe userTable
set producer latest
set producer LRP 10 minutes
insert into userTable values ('JoeBloggs', 'Hello', 1.23, 4)
```

Note that the values should be of the same type and column order as those in the table. Enter the following to read the data back from the table:

```
select * from userTable
```

You might see entries from other users who are also using this table. To select only your own entries you could change the query to:

```
select * from userTable where userId='JoeBloggs'
```

To exit the rgma command line interface enter:

```
exit
```

Some more examples are available within the command line tool by entering:

```
rgma  
help examples
```

For more information on the R-GMA command line tool see section 7.

3 R-GMA INSTALLATION

Though this is a user guide, it is worth understanding some aspects of the installation.

3.1 INSTALLED COMPONENTS

Clients communicate with an R-GMA server. For operations using a producer or a consumer this is the local R-GMA server on the site, however some operations of the command line tool contact a registry and/or the schema. A client installation needs to be configured to know its local R-GMA site server, the set of servers running a registry and the server running a schema. Each R-GMA server has the ability to run all services, but most servers only run the producer and consumer services. An R-GMA installation is defined by its schema and registries and sites should *not* configure an additional registry except with the agreement of those having planning authority for that grid. A number of secondary producers, republishing to a database, will normally be set up to answer latest and history queries over aggregated information. An R-GMA server can also be configured to run the browser service to offer web access to R-GMA.

3.2 CHECK YOUR INSTALLATION

You should check your installation by typing “rgma-client-check”. If it returns with an OK message then the system is probably installed correctly, if not you have some kind of problem and you should either check the installation as described in the installation guide linked from <http://hepunix.rl.ac.uk/egEE/jra1-uk/glite-r1.5/>, or ask your sysadmin to install the system correctly.⁴

The file `$RGMA_HOME/etc/rgma/rgma.conf` must contain the URLs of the R-GMA services you want the API to connect to: the standard R-GMA installation creates this for you. If you need to use an HTTP proxy server, set the environment variable `http_proxy` to contain the URL (including port number) of the proxy server.

⁴This test checks that all the four APIs are installed. If you have not installed all of them then you should expect it to report a different number.

3.3 SECURITY

R-GMA client applications require a suitable grid certificate or proxy to access secure R-GMA services (on <https://...:8443/R-GMA>). The location of the security credentials can be specified in one of two ways:

1. Setting environment variable `X509_USER_PROXY` to the location of a grid proxy certificate file (e.g. `/tmp/x509up_u00000`).
2. Setting environment variable `TRUSTFILE` to the location of a security properties file. An example is installed at `$RGMA_HOME/etc/rgma/ClientAuthentication.props`. This should contain a setting for either a `gridProxyFile` (as above) or `sslCertFile/sslKey/sslKeyPasswd` combination. You should make your own copy of this file.

Note that Java clients will need to pass in one of these environment variables to the JVM explicitly using the `-D` option.

Since R-GMA authentication is mutual, your system administrator will also need to install the Certificate Authority files for the CA that signs the server certificate of any server you wish to use, on your machine. These are normally located in `/etc/grid-security/certificates`.

4 PRIMARY PRODUCERS

4.1 TERMINATION INTERVAL

Each Primary Producer resource has a Termination Interval as described in 1.2.6, that is a time interval within which the user must make contact with the producer service, in order to keep the resource alive and maintain its table entries in the registry. The Termination Interval is set by the user when the resource is created, however this option will probably be removed from control in the future. For now please set it to one hour. If the producer's publication interval (the interval between calls to `insert`) is greater than the Termination Interval, then the user should call `showSignOfLife` periodically to keep it alive.

If the user sends a `close` request the Primary Producer will continue to be available to Consumers until all tuples are older than the *HistoryRetentionPeriod* but will no longer be contactable from the user API. If the termination interval is exceeded, the resource behaves as though it had received a `close` request.

The resource is destroyed immediately when the user issues an explicit `destroy` request.

The user should make allowance for the variable time it will take for any messages to reach the service and be processed. This extra time should be short as the system has been designed with the expectation that clients will contact local services. Even then, the user should always be aware that a resource he was using may have timed out when he next tries to contact it. Various recommendations on timing are provided in the "Advice on using R-GMA" section 11 of this guide.

4.2 PRODUCER PROPERTIES

All producers support continuous queries, but you may also specify in the producer properties that you want the producer to also support history and/or latest queries.

The tuple-storage maintained by Primary and Secondary Producers can either be in memory or in a real database table. You should choose, as part of the producer properties, whichever is the most appropriate. Memory gives the best performance, whereas an RDBMS gives the best performance for complex queries

- especially those requiring joins⁵. A common pattern is to use a memory based producer as the primary one and then to use a more resilient secondary producer by storing in an RDBMS. This is the example we consider here.

4.3 PRIMARY PRODUCER EXAMPLES

Here we show an example of a piece of code and how to run it.

4.3.1 SIMPLE PRIMARY PRODUCER EXAMPLE

```

0  #!/usr/bin/env python

1  import sys
2  import rgma

3  if len(sys.argv) != 2:
4      print("Exactly one argument must be specified")
5      sys.exit(1)

6  try:
7      storage = rgma.Storage(rgma.StorageType.MEMORY)
8      properties = rgma.ProducerProperties(storage)
9      terminationInterval = rgma.TimeInterval(60, rgma.Units.MINUTES)
10     producer = rgma.PrimaryProducer(terminationInterval, properties)

11     predicate = "WHERE (userId = '%s')" % sys.argv[1]
12     latestRetentionPeriod = rgma.TimeInterval(60, rgma.Units.MINUTES)
13     historyRetentionPeriod = rgma.TimeInterval(60, rgma.Units.MINUTES)
14     producer.declareTable("userTable",
15                           predicate,
16                           latestRetentionPeriod,
17                           historyRetentionPeriod)

18     insert = "INSERT INTO userTable (userId, aString, aReal, anInt) \
19             VALUES ('%s', 'Python producer', 3.1415962, 42)" % sys.argv[1]
20     producer.insert(insert)

21     producer.close()

22 except (rgma.RGMAException, rgma.RemoteException, rgma.UnknownResourceException), e:
23     print "RGMA Error", e
24     sys.exit(1)
25 except:
26     print "Unexpected error:", sys.exc_info()[0]
27     sys.exit(1)

```

Line 2 imports the rgma library. It must be imported to make any use of the R-GMA facilities.

Line 6 is the start of a try block. All the effective code is within this try block.

Line 7 defines the storage for the producer to be of type memory

Line 8 defines the properties. ProducerProperties can take a second argument, for example to specify that the producer should also support history or latest queries.

⁵Currently only RDBMS storage supports join operations

Line 9 defines a termination interval of 60 minutes. This is in line with the recommendations in section 11 giving advice on how to use R-GMA.

Lines 10 creates the primary producer with the specified termination interval and properties.

Lines 11–17 define the predicate and a pair of retention periods (60 minutes in each case) and then declare the intention of publishing to a table called `userTable`. A producer is able to publish to more than one table.

Lines 18–20 build an SQL insert statement and then call the insert method of the primary producer with that SQL statement as an argument.

Line 21 closes the producer. This is not essential as the producers destructor will tidy up for you if you forget.

Finally lines 22–27 handle any exceptions which are thrown. The `RemoteException` is thrown when a connection to a remote service fails, the `UnknownResourceException` is thrown when a remote resource cannot be located (maybe because it has timed out) and `RGMAExceptions` are thrown by RGMA itself.

4.3.2 RUNNING THE EXAMPLE

The environment variable `PYTHONPATH` should include `$RGMA_HOME/lib/python`⁶.

4.3.3 RESILIENT PRIMARY PRODUCER EXAMPLE

This example illustrates how to write a resilient producer according to the recommendations in section 11. The Primary Producer publishes information periodically every 30 seconds. If there is a network problem or the buffer is full it retries after one minute, if the producer resource no longer exists a new one is created.

```
0  #!/usr/bin/env python
1  import sys, time
2  import rgma
3  if len(sys.argv) != 2:
4      sys.stderr.write("Exactly one argument must be specified\n")
5      sys.exit(1)
6  data = 0
7  predicate = "WHERE userId = '%s'" % sys.argv[1]
8  storage = rgma.Storage(rgma.StorageType.MEMORY)
9  properties = rgma.ProducerProperties(storage)
10 terminationInterval = rgma.TimeInterval(60, rgma.Units.MINUTES)
11 historyRP = rgma.TimeInterval(50, rgma.Units.MINUTES)
12 latestRP = rgma.TimeInterval(25, rgma.Units.MINUTES)
13 producerCreated = False
14 tableDeclared = False
15 while True:
16     try:
17         while not producerCreated:
18             producer = rgma.PrimaryProducer(terminationInterval, properties)
```

⁶Currently, you also need to ensure that `LD_LIBRARY_PATH` includes `$RGMA_HOME/lib`.

```

19         producerCreated = True

20     while not tableDeclared:
21         producer.declareTable("userTable", predicate, latestRP, historyRP)
22         tableDeclared = True

23     while True:
24         insert = "INSERT INTO userTable (userId, aString, aReal, anInt) \
25             VALUES ('%s', '', 0.0, %d)" % (sys.argv[1], data)
26         producer.insert(insert)
27         print insert
28         data += 1
29         time.sleep(30)

30     except (rgma.UnknownResourceException), e:
31         sys.stderr.write("Producer has died - will try to create a new one\n");
32         producerCreated = False
33         tableDeclared = False

34     except (rgma.RemoteException), e:
35         sys.stderr.write("Failed to contact R-GMA server: %s - will retry in 60s\n" % e);
36         time.sleep(60)

37     except (rgma.RGMABufferFullException), e:
38         sys.stderr.write("Producer buffer is full - will retry insert in 60s\n");
39         time.sleep(60)

40     except (rgma.RGMAUserException), e:
41         sys.stderr.write("WARNING: Invalid tuple: %s (%s)\n" % (insert, e.getMessage()))

42     except (rgma.RGMAException), e:
43         sys.stderr.write("Unexpected R-GMA error - exiting\n")
44         if producer:
45             try:
46                 producer.close()
47             except:
48                 pass
49         raise e

```

Line 6 initialises the data value that is going to be inserted into each tuple.

Lines 7–12 set up the producer properties.

Lines 13–14 set up initial states for the producer and table

Line 15 loops forever.

Lines 17–19 create a new Primary Producer. This is only done on the first time through the loop or if the producer needs to be recreated following a failure.

Lines 20–22 declare the userTable. This is only done on the first time through the loop or if the table needs to be declared following a failure.

Lines 23–29 loops forever - each iteration will insert a new tuple. Our simple example tuples just contain the userId and the current value of the data variable that gets incremented on each successful iteration of the loop.

Lines 30–33 an UnknownResourceException is a result of the client contacting the server and finding that the producer resource no longer exists. Hence an UnknownResourceException will result in the producer being recreated.

Lines 34–36 a `RemoteException` is a result of the client being unable to contact the server, possibly due to network problems. Wait for a minute before retrying the operation.

Lines 37–39 a `RGMABufferFullException` is caused by the tuple buffer on the server temporarily being full. Wait for a minute before retrying the operation.

Lines 40–41 a `RGMAUserException` is caused when there is an error with the inserted tuple. A warning is printed and the code continues.

Lines 42–49 other R-GMA exceptions will cause the program to exit. Before exiting an attempt is made to close the producer.

5 CONSUMING INFORMATION

5.1 TYPES OF QUERY

There are four types of query: *continuous*, *latest*, *history* and *static* (the latter are only supported by On-demand Producers). The set of queries that a particular producer supports is recorded in the registry. All query types except static can take an optional time interval parameter.

A continuous query causes all new tuples that match the query, to be streamed into the consumer's tuple-storage, as soon as they are inserted into the virtual table by the producers. Streaming continues until the consumer requests it to stop. If a time interval is specified, the consumer will additionally receive any tuples which are already in the virtual table when the query starts, and which are no older than the time interval. There is no guarantee that tuples are time-ordered. All Primary and Secondary producers support continuous queries. On-demand producers do not.

Latest and history queries are *one-time* queries: they execute on the current contents of the virtual table, then terminate. In a history-query, all versions of any matching tuples are returned; in a latest-query, only those representing the “current state” (see 1.2.4) are returned. In both cases, a time interval may be specified with the query, to limit the age of the tuples returned. Primary and Secondary Producers may optionally support one-time queries. On-demand producer do not.

5.2 CONSUMER EXAMPLES

This provides examples of code you might use or adapt.

5.2.1 SIMPLE CONSUMER EXAMPLE

```
0  #!/usr/bin/env python
1  import sys, time
2  import rgma
3
4  try:
5      terminationInterval = rgma.TimeInterval(60, rgma.Units.MINUTES)
6      queryProperties = rgma.QueryProperties(rgma.QueryType.CONTINUOUS)
7      select = "SELECT * FROM userTable"
8      consumer = rgma.Consumer(terminationInterval, queryProperties, select)
9
10     timeout = rgma.TimeInterval(300, rgma.Units.SECONDS)
11     consumer.start(timeout)
```

```

10     endOfResults = False
11     while not endOfResults:
12         time.sleep(5)
13         results = consumer.pop(2000)
14         for row in results:
15             print row
16         endOfResults = results.endOfResults

17     consumer.close()

18 except (rgma.RGMAException, rgma.RemoteException, rgma.UnknownResourceException), e:
19     print "RGMA Error: ", e
20     sys.exit(1)
21 except:
22     print "Unexpected error:", sys.exc_info()[0]
23     sys.exit(1)

```

Line 2 imports the library to allow R-GMA to be used.

Line 3 is the start of the *try* block.

Lines 4–7 define the termination interval for the consumer to be 60 minutes, and that we require a continuous query with the query specified on line 6. The query is then created on line 7. Note that it is recommended to set the termination interval to 60 minutes (see section 11) It is also possible to specify how far back in time to start, see section 5.2.2

Lines 8–9 start the query running with a timeout of 5 minutes. With a latest or history query we can expect the query to complete within the timeout. However the continuous query can only be terminated by an explicit abort call or by the timeout being exceeded. So we expect this query to take exactly 5 minutes.

Lines 10–16 loop until the all results have been retrieved, sleeping for five seconds between each call to `pop`.

Line 17 closes the consumer

Lines 18–23 show the exception handling.

5.2.2 CONSUMING CONTINUOUS PLUS OLD INFORMATION

If you want a continuous query starting from 10 minutes ago, you need to say it is continuous with a time interval. Replace line 5 by: `queryProperties = rgma.QueryProperties(rgma.QueryType.CONTINUOUS, rgma.TimeInterval(10, rgma.Units.MINUTES))`. This is what you would need to pick up the tuple already published by the Primary Producer example.

5.2.3 ONE-OFF QUERIES

For one-off queries, either history or latest, it is preferable to check to see if the query aborted. This can be achieved by looking to see if `consumer.hasAborted()` is true after either kind of consumer loop.

This can be the result of hitting the timeout (300 seconds in this case) or making an explicit `consumer.abort()` call. Note that continuous queries *only* stop by one of these means so there is no point in checking in that case.

5.2.4 CONSUMER EXTRACTING INFORMATION FROM RESULT SET

If you wish to extract the specific fields from the ResultSet, and handle them separately this can be done by replacing lines 13–14 and 16–17 in the above example with the following loop

```
0   for row in rows:

1       user = row['userId']
2       aString = row['aString']
3       aReal = row['aReal']
4       anInt = row['anInt']
5       tupleTime = row['MeasurementTime']

6       print "Read: userId = %s, aString = %s, aReal = %f, anInt = %i, time = %s" %\
7           (user, aString, float(aReal), int(anInt), tupleTime)
```

5.2.5 RESILIENT CONSUMER EXAMPLE

This example illustrates how to write a resilient consumer according to the recommendations in section 11. The consumer retrieves information periodically every five seconds. If there is a network problem it retries after one minutes, if the consumer resource no longer exists a new one is created.

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

1   import sys, time
2   import rgma

3   terminationInterval = rgma.TimeInterval(60, rgma.Units.MINUTES)
4   oldData = rgma.TimeInterval(30, rgma.Units.SECONDS)
5   queryProperties = rgma.QueryProperties(rgma.QueryType.CONTINUOUS, oldData)
6   select = "SELECT * FROM userTable"

7   consumerCreated = False
8   consumerStarted = False
9   endOfResults = False;

10  while True:

11      try:

12          while not consumerCreated:
13              consumer = rgma.Consumer(terminationInterval, queryProperties, select)
14              consumerCreated = True

15          while not consumerStarted:
16              consumer.start(rgma.TimeInterval(365, rgma.Units.DAYS))
17              consumerStarted = True

18          while not endOfResults:
19              time.sleep(5)
20              resultSet = consumer.pop(2000)
21              for row in resultSet:
22                  print row
23              if resultSet.warning:
24                  print "WARNING: %s" % resultSet.warning.getMessage()
25              endOfResults = resultSet.endOfResults
```

```

26         try:
27             consumer.close()
28         except Exception, e:
29             pass

30         consumerCreated = False;
31         consumerStarted = False;
32         endOfResults = False;

33     except (rgma.UnknownResourceException), e:
34         sys.stderr.write("Consumer has died - will try to create a new one\n")
35         consumerCreated = False
36         consumerStarted = False
37         endOfResults = False

38     except (rgma.RemoteException), e:
39         sys.stderr.write("Failed to contact R-GMA server: %s - will retry in 60s\n" % e)
40         time.sleep(60)

41     except (rgma.RGMAException), e:
42         sys.stderr.write("Unexpected R-GMA error - exiting")
43         if consumer:
44             try:
45                 consumer.close()
46             except:
47                 pass
48         raise e

```

Lines 3–6 set up the consumer properties. This uses a continuous plus old information query (5.2.2) so that if the consumer is restarted data should not be lost. N.B. It is possible that this may cause the some tuples to be retrieved a second time if the consumer restarts.

Lines 7–9 set up initial states for the consumer and resultset

Line 10 loops forever.

Lines 12–14 create a new consumer. This is only done on the first time through the loop or if the consumer needs to be recreated following a failure.

Lines 15–17 start the consumer. This is only done on the first time through the loop or if the consumer needs restarting following a failure.

Lines 18–25 loops for a year - each iteration will retrieve data and check for warnings. See the Consumers section in 11 for further explanation about the pop and sleep.

Lines 26–29 after a year of looping close the consumer.

Lines 30–32 reset the states for the consumer and resultset so that a new consumer will be recreated by the main loop.

Lines 33–37 an UnknownResourceException is a result of the client contacting the server and finding that the consumer resource no longer exists. Hence an UnknownResourceException will result in the consumer being recreated.

Lines 38–40 a RemoteException is a result of the client being unable to contact the server, possibly due to network problems. Wait for a minute before retrying the operation.

Lines 41–48 other R-GMA exceptions will cause the program to exit. Before exiting an attempt is made to close the consumer.

6 REPUBLISHING VIA SECONDARY PRODUCERS

As explained in section 1.2.3, a Secondary Producer populates its internal storage by running a query and the user code only sets the process running. This is demonstrated by the code in the examples below. Note that as this example stores information in the database, unlike the primary producer and simple consumer above, this example cannot be used without editing to specify your own local database.

6.1 SECONDARY PRODUCER EXAMPLES

6.1.1 SIMPLE SECONDARY PRODUCER EXAMPLE

```

0  #!/usr/bin/env python

1  import sys, time
2  import rgma

3  try:

4      storage = rgma.Storage(rgma.StorageType.DATABASE, \
5                             "jdbc:mysql://localhost:3306/somedatabase", \
6                             "fred", \
7                             "bloggs")
8      properties = rgma.ProducerProperties(storage, rgma.QueryType.LATEST)
9      terminationInterval = rgma.TimeInterval(60, rgma.Units.MINUTES)
10     sp = rgma.SecondaryProducer(terminationInterval, properties)

11     predicate = ""
12     historyRetentionPeriod = rgma.TimeInterval(2, rgma.Units.HOURS)
13     sp.declareTable("userTable",
14                    predicate,
15                    historyRetentionPeriod)

16     while True:
17         sp.showSignOfLife()
18         time.sleep(55*60)

19     except (rgma.RGMAException, rgma.RemoteException, rgma.UnknownResourceException), e:
20         print "RGMA Error", e
21         sys.exit(1)
22     except:
23         print "Unexpected error:", sys.exc_info()[0]

```

Line 2 is the usual R-GMA import statement for python

Line 3 is the start of the try block

Lines 4–10 create the secondary producer. Lines 4–7 define the storage location parameters and line 10 actually does the creation. In this case we have specified a termination interval of 65 minutes with a database providing support for latest queries. It is currently the user's responsibility to ensure that the specified database exists, as explained in 10.1.

Lines 11–15 declare a table that the secondary producer will deal with. The predicate (line 11) is an empty string meaning that this secondary producer will collect and republish the whole table. The history retention period is set to 2 hours. This means that tuples will be available until they are 2 hours old - these tuples will be made available to both continuous and history queries. In addition tuples will be available for latest queries. In this case however the latest retention period is a property of the individual tuple as defined at the primary producer.

Lines 16–18 keep the secondary producer alive. On line 9 it was given a termination interval of 60 minutes so here we must issue an `showSignOfLife` call slightly more often. In this case the sleep is for 55*60 seconds (55 minutes). It is possible that the `showSignOfLife` call may fail if it is unable to contact the service, or the service is unable to locate the remote resource. In this case the loop will be exited.

Lines 19–24 are the exception handling.

Notice that if the secondary producer does fail it could be an hour before the program detects this and exits, so you need to set the retention period and sleep parameter according to your needs.

6.1.2 AVOIDING A PERMANENT CONNECTION TO A SECONDARY PRODUCER

If you don't want to keep the process running continuously you could leave the secondary producer running after disconnecting from the API and then periodically run a job which reconnects. To do this replace lines 16–18 to the end by:

```
0      rep = sp.getEndpoint()
1      print rep.getUrl()
2      print rep.getResourceId()
3      sys.exit(0)
```

Lines 0 obtains the `ResourceEndpoint`. This holds everything which is needed to reestablish a connection to a remote resource.

Lines 1–3 write the URL and the `resourceId` to stdout. It is assumed that you will capture this output. You then need to run every hour the following code:

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

1  import sys
2  import rgma

3  if len(sys.argv) != 3:
4      print("Usage: python SecondaryProducerReconnect.py <url> <resourceId>");
5      sys.exit(1);

6  try:

7      url = sys.argv[1]
8      resourceId = int(sys.argv[2])
9      rep = rgma.ResourceEndpoint(url, resourceId)

10     sp = rgma.reconnect(rgma.SecondaryProducer, rep)
11     sp.showSignOfLife()

12 except (rgma.RGMAException, rgma.RemoteException, rgma.UnknownResourceException), e:
13     print "RGMA Error", e
14     sys.exit(1)
15 except:
16     print "Unexpected error:", sys.exc_info()[0]
```

Line 2 is the normal R-GMA import statement.

Lines 3–5 check that two arguments have been specified. These are the URL and `resourceId` output by the first program.

Lines 7–9 create the `ResourceEndpoint`

Line 10–11 reconnect using the ResourceEndpoint. Note that the first argument to `reconnect` is the class. The `showSignOfLife` call tells the service to keep the remote resource alive.

Lines 12–16 are the exception handling.

7 THE RGMA COMMAND LINE TOOL

7.1 INTRODUCTION

The R-GMA command line tool offers simple command-based access to the R-GMA virtual database. The interface is intended to be similar to the command-line tools supplied with databases, e.g. MySQL.

A complete description of the command line tool is available in the users manual linked from the R-GMA documentation web page: <http://hepunix.rl.ac.uk/egee/jra1-uk/glite-r1.5/>.

7.1.1 STARTING THE R-GMA COMMAND LINE TOOL

To start the R-GMA command line tool, run the command `rgma`.

On startup you should receive the following message:

```
Welcome to the R-GMA virtual database for Virtual Organisations.  
=====
```

```
Your local R-GMA server is:
```

```
https://yourlocal.server.machine:8443/R-GMA
```

```
You are connected to the following R-GMA Registry services:
```

```
https://amachine.somewhere:8443/R-GMA/RegistryServlet
```

```
You are connected to the following R-GMA Schema service:
```

```
https://amachine.somewhere:8443/R-GMA/SchemaServlet
```

```
Type "help" for a list of commands.
```

```
rgma>
```

7.1.2 ENTERING COMMANDS

Commands are entered by typing at the `rgma>` prompt and hitting “enter” to execute the command. A history of commands executed can be accessed using the Up and Down arrow keys. Commands can be entered in lower or upper case (but not a mixture of both).

Command autocompletion is supported – hit the “Tab” key when you have partly entered a command and it will either be completed automatically or a list of matching alternatives will be displayed.

7.2 COMMANDS

7.2.1 GENERAL COMMANDS

<code>help</code>	Displays general help information
<code>help <command></code>	Displays help for a specific command
<code>exit</code>	Exit the R-GMA command line
<code>quit</code>	Same as “exit”

7.2.2 QUERYING DATA

Querying data uses the standard SQL `SELECT` statement, e.g.

```
rgma> SELECT * FROM ServiceStatus
```

The type of query can be changed using the `SET QUERY` command:

```
rgma> SET QUERY LATEST
rgma> SET QUERY CONTINUOUS
```

The maximum age of tuples to return can also be controlled using the `SET MAXAGE` command which takes a value and a time unit (seconds, minutes, hours, days - default is seconds).

```
rgma> SET MAXAGE 2 minutes
rgma> SET MAXAGE 120
```

If a maximum age is specified for a continuous query, the query will initially return a history of matching tuples up to the specified maximum age. It will then return new tuples as they are inserted.⁷

To disable the maximum age, set it to `none`:

```
rgma> SET MAXAGE none
```

The query timeout controls how long the query will execute for before exiting automatically.

```
rgma> SET TIMEOUT 3 minutes
rgma> SET TIMEOUT 180
```

7.2.3 INSERTING DATA

The SQL `INSERT` statement may be used to add data to the system:

```
rgma> INSERT INTO Table (col1, col2, col3, col4) VALUES ('a', 'b', 'c', 'd')
```

Data is inserted into the system using a Producer component. All producers can answer continuous queries, but ability to answer latest and history queries is optional and is changed using the `SET PRODUCER` command⁸:

⁷Currently the use of maxage with a continuous query is not fully implemented. If any value is specified for maxage, the query will initially return a history of all matching tuples that have not expired.

⁸Primary and Secondary producers able to answer latest and history queries have not yet been implemented


```
rgma> SET PRODUCER latest
rgma> SET PRODUCER latest history
rgma> SET PRODUCER continuous
```

The default setting is for a producer that only answers continuous queries.

A producer may have a predicate associated with it describing the subset of a table it provides. For example, if a table `userTable` has the column `userId` which for your producer will always have the value `me`, you can express this restriction using:

```
rgma> SET PRODUCER PREDICATE userTable WHERE userId = 'me'
```

To remove the predicate use:

```
rgma> SET PRODUCER PREDICATE <table name> none
```

For a producer that can answer latest and/or history queries, the latest and history retention periods can be controlled using:

```
rgma> SET PRODUCER latestretentionperiod 30 minutes
rgma> SET PRODUCER historyretentionperiod 2 hours
```

7.2.4 SECONDARY PRODUCERS

A Secondary producer does not insert new data to the system, but collects data from individual Producers and makes it available via its own Producer component.

To instruct the secondary producer to consume from the table `MyTable`, use the following command:

```
rgma> SECONDARYPRODUCER MyTable
```

Like the producer, the secondary producer may be configured to answer latest and/or history queries, e.g.

```
rgma> SET SECONDARYPRODUCER latest
```

If the secondary producer can answer history queries, it has an associated history retention period, as for a producer. This is controlled in the same way:

```
rgma> SET SECONDARYPRODUCER historyretentionperiod 1 day
```

7.2.5 INFORMATION COMMANDS

To show a list of all R-GMA producers that produce the table `MyTable`:

```
rgma> SHOW PRODUCERS OF MyTable
```

To show a list of all table names:

```
rgma> SHOW TABLES
```

To show information about a table `MyTable`:

```
rgma> DESCRIBE MyTable
```

7.2.6 DIRECTED QUERIES

Normally a component of R-GMA called the *mediator* selects which Producers are contacted to answer a query. For debugging purposes it may be useful to specify a particular Producer to use instead. This is called a *directed query* and can be specified with the `USE PRODUCER` command:

```
rgma> USE PRODUCER <url> <resource id>
```

All future `SELECT` queries will be directed to this Producer. Only one producer may be specified. The `<url>` and `<resource id>` should correspond to a valid Producer that can answer the type of queries you put to it or no results will be returned. The `SHOW PRODUCERS` command displays urls and resource IDs of registered Producers.

To revert back to using the Mediator to select producers, use the command:

```
rgma> USE MEDIATOR
```

8 USING THE WEB TO BROWSE R-GMA INFORMATION

The R-GMA browser can be used to:

- View definitions of available tables in the schema.
- View Producers that are publishing to a table.
- Pose a mediated query on a table.
- Pose a query to specific Producers of a table.

8.1 SECURITY

To access the R-GMA Browser when it is hosted on a secure (HTTPS) server, a suitable client certificate needs to be imported into the Web Browser.

9 ADMINISTRATION

9.1 TABLE CREATION

The new `createTable` API method should now be used for adding tables to an R-GMA schema.

9.2 RECOVERY FOLLOWING RESTART

R-GMA instances (Web Service Resources) are always destroyed when the corresponding service restarts. All calls to services in R-GMA (either from the user or another service) must therefore be prepared to discover that a resource no longer exists, or is no longer registered in the registry, and to handle the error gracefully. The use of time-outs (soft-state registration) on the registry and in the producer and consumer services, ensures entries to non-existent resources are automatically removed within a reasonable length of time.

Those R-GMA services which use permanent storage (registry service, schema service and primary and secondary producers with user-specified tuple stores) do have some degree of resilience, because new

resources can connect to existing storage, provided the storage itself is recoverable. In the case of a producer's file or database tuple-store, existing tuples will be automatically available, and will remain in the store until they exceed their retention period, in the usual way.

10 SQL

10.1 CREATING A DATABASE AS STORAGE FOR A PRODUCER

To create the database in the examples (somedatabase with user/password of fred/bloggs) you will need sufficient privileges to be able to create a database. You (or your sysadmin) should connect to the DBMS with the command `mysql` e.g. by entering:-

```
mysql -u root -p
```

and enter:

```
CREATE DATABASE somedatabase;  
GRANT ALL PRIVILEGES ON somedatabase.* TO fred@localhost IDENTIFIED BY 'bloggs';  
GRANT ALL PRIVILEGES ON somedatabase.* TO fred%"%" IDENTIFIED BY 'bloggs';
```

The first line creates the database. The other two lines create user/password combinations if they do not already exist and give "fred", logged in from any machine, full privileges on "somedatabase"

10.2 EXAMPLES OF SQL QUERIES

There are many good SQL tutorials on the web[3, 2]. Here are listed a few simple examples of working queries.

To list all rows in a table called `GlueSite`:

```
SELECT * FROM GlueSite
```

To list selected information about all sites in this table:

```
SELECT SysAdminContact, UserSupportContact FROM GlueSite
```

To list the names of all sites in this table in the Northern Hemisphere:

```
SELECT Name FROM GlueSite WHERE Latitude > 0
```

To list the endpoints and types of services in this table hosted by sites in the Northern Hemisphere:

```
SELECT Endpoint, Type FROM GlueSite, GlueService WHERE GlueSite.UniqueId = GlueService.GlueSite
```

10.3 SUPPORTED SQL

SQL SELECT statements are restricted by 3 components: the R-GMA SQL Parser, limitations on continuous queries, and limitations imposed by external components.

Currently the only external component which limits the SQL is the MySQL RDBMS when it is used by a producer for managing storage. As R-GMA is ported to other RDBMS systems in the future, users will be obliged to use only that dialect of SQL which is common to those RDBMSs able to participate in answering a query.

Continuous queries must be of a form which can be evaluated on each tuple in isolation.

Finally the R-GMA SQL parser accepts SQL92 entry level SELECT statements except as listed below:

- Trailing decimal points in a number – e.g. 123.
- The keywords AS and ESCAPE.
- Nested SELECT statement after keyword SOME.
- The keyword HAVING after a table name or WHERE clause.
- “,” used after COUNT (*) – e.g. SELECT COUNT (*) , SUM(NUMKEY) FROM UPUNIQ
- Lack of space between elements in expression.
- Column names in quotes.
- Expressions with 2 column name elements inside MAX (), SUM (), and AVG () .
- Every column name in a SELECT must be unique. This also implies that queries of the form SELECT * FROM A,B are also not allowed as the * would include the R-GMA system columns twice in the join. In the unlikely event that you need two columns in a join with the same name but with a different meaning, the R-GMA team can propose a work-around.

11 ADVICE ON USING R-GMA

This section contains various recommendations, many of which relate to timing, to help the user make the best use of R-GMA. The aim is to ensure the most reliable throughput of tuples and reduce unnecessary load on the servers

11.1 GENERAL ADVICE

- All R-GMA calls raise exceptions - they must all be caught and handled for reliable use.
- If you try to contact a producer or consumer and the server returns a `UnkownResourceException` exception the resource no longer exists so you must recreate the producer or consumer.
- If the server cannot be contacted a `RemoteException` exception is returned. Wait for a period of one minute and retry the operation and repeat every minute.
- The termination interval should be set to 60 minutes for all producers and consumers. In future, it will probably not be possible for the user to set this parameter.
- Remember to close consumers and producers when you have finished with them. Otherwise they will only be closed, at best, an hour later when the termination interval has passed.

- Always list the columns in INSERTs and SELECTs to provide some protection from schema changes.

11.2 PRIMARY PRODUCERS

- Set the Latest Retention Period for tuples (in `declareTable`) to match the life-time for which you think the tuple should be considered to be "latest" information. Typically this will be a little greater than the publication interval.
- The primary producer service has a buffer which can fill up. If this happens, the server returns an `RMGABufferFull` exception. On receipt of this exception, the user should wait for one minute before trying to insert the tuple(s) again.
- If data are not being published at least every half an hour create a Primary Producer when it is needed rather than keep it alive. If data are published at less than half hour intervals it is better to keep the producer running.
- In general, it is not necessary to send `showSignOfLife` messages, and we recommend that you do not send them, because `insert` messages will keep a producer resource alive. If a problem should occur you can always create a new producer.
- Set the History Retention Period (HRP) to hold at least the last measurement so that a new Secondary Producer can pick up the latest data. It should not normally be significantly longer than this. If you do make it long you may run out of memory on the service (if using memory storage). In this case you have a problem - if you just leave it, or close it, then it will continue to occupy memory until the HRP expires. The only way to clean it up is to `destroy` it. For this reason HRP for primary producers should normally be short: between 10 and 60 minutes. If the rate of publishing data is very high it may be necessary to have a shorter HRP, however this will lead to some tuples being lost if a Secondary Producer fails.

11.3 SECONDARY PRODUCERS

- Call `showSignOfLife` for a Secondary Producer to keep it alive. The interval should be 5 minutes less than the shortest History Retention Period (HRP) of the contributing Primary Producers. This way if a Secondary Producer dies it can be restarted without loss of data. If the shortest HRP is more than 60 minutes then the `showSignOfLife` should be sent every 55 minutes (i.e 5 minutes less than the termination interval).
- The HRP can be set to be large with database storage, however with memory consider how much space will be taken up by the tuples.
- gLite 1.4 includes a flexy archiver which will do most of the work for you.

11.4 CONSUMERS

- You may get better performance and will reduce the load on the server by introducing a delay between successive calls to `pop`. This allows time for the consumer's buffer to fill up a bit rather than getting back very few tuples (or none at all) each time. Smart code would adjust the polling rate to be as low as possible to keep up with the data. A simpler algorithm would be to wait for 5 seconds after any `pop` which returns no tuples; otherwise not include a delay.
- Do not use `showSignOfLife` but rely on the `pop` calls to keep it alive.

12 RELEASE NOTES FROM A USER PERSPECTIVE

This section describes the changes that a user will see when going from gLite 1.4 to 1.5.

12.1 OVERVIEW

This release of R-GMA in gLite 1.5 fixes more than 40 bugs. Our most serious bug number 8099: "Archivers are Inconsistent" has not been fixed completely but this release should show much better behaviour.

The version is finally free of the old EDG API layer. The APIs communicate directly with the services and are no longer interdependent. One consequence of this is much improved error messages.

APIs and services must be matched for the transition from 1.4 to 1.5 - so a site should upgrade in one go. For C we provide a 1.4 compatibility library so that code linked against a shareable library will still work. This was not feasible for the C++, so C++ application code must be recompiled. An extra compiler switch is needed for both C and C++ compilation.

Error handling is much improved. You can now expect to receive a clear error message if there is a problem. If you don't - please submit a bug. There are three new exception types:

- `RGMASecurityException`
- `RGMAUserException`
- `RGMABufferFullException`

We are now checking that published data are consistent with the schema. Previously a primary producer would accept almost anything, but then a secondary producer might modify the fields - e.g. truncating strings or it could just reject the tuple. Now we make the check when you try to insert the data into R-GMA this allows us to reject the tuple at source and inform the guilty party.

In future R-GMA will support multiple Virtual Databases (VDBs). We imagine that a VO will "own" and administer more than one VDB. We had begun to introduce the framework for this but realised that we had not done it in the best way. In future the VDB will simply be a prefix to the table name, currently however, the prefix may not be specified. Users should start think about how they wish to use the VDB when it is introduced.

The `reconnect` operation no longer sends a `showSignOfLife`. This is an important change of semantics. Typically if you reconnect, you will immediately use the service so you will be unaware of the change.

Finally, there used to be some non-advertised facilities for determining service status. We now offer `getProperty` operations to find out what is happening. These are of limited interest to the normal user. A `setProperty` call is also provided.

12.2 NEW FEATURES

- (LRP on insert) You can now specify a latest retention period for each tuple in the Primary Producer's `insert` call, overriding the one specified in `declareTable`. This is useful, for example, for the R-GMA Service Tool as different LRPs are appropriate for data from different services about which information is being published.
- (Logically named tuple stores) You can now access the database used by a Primary or Secondary producer just by giving it a logical name: R-GMA will map it to a physical database for you. Two new API calls `listTupleStores` and `dropTupleStores` are provided to manage them.

- (Tuple checking) Attempts to insert tuples that are inconsistent with the table definition in the schema will now be rejected.
- (numSuccessfulOps) The Primary Producer's `insertList` call may fail after inserting a number of tuples. You can now determine how far it got with the new `numSuccessfulOps` field in the `RGMAException`. It will contain the number of successfully inserted tuples.
- (endOfResults) A new method called `endOfResults` has been added to the result set returned by the Consumer's `pop` method to allow a much cleaner consumer loop to be constructed. A (Python) example looks like this:

```
consumer.start(timeout)
while 1:
    results = consumer.pop(1000)
    for result in results:
        print result

    if results.endOfResults:
        break

    time.sleep(5)
```

The calls for the old looping constructs `isExecuting` and `count` are deprecated.

- (Registry and Schema calls) Calls to access the Registry and Schema have been added to the APIs. The most important new calls are `getAllProducersForTable` on the Registry and `createTable`, `dropTable`, `getAllTables` and `getTableDefinition` on the Schema. The temporary `create-table` command line utility has been withdrawn. Note that users should be very wary of calling `dropTable` because table definitions are shared by everyone using the Schema.

See the next section for known limitations of the current R-GMA services.

12.3 DEPRECATED FEATURES

The following features are now deprecated:

- VO-names list in `createPrimaryProducer`, `createSecondaryProducer` and `createOnDemandProducer`.
- `IgnoreSlowConsumers` property in `createPrimaryProducer` and `createSecondaryProducer`: the flag will be ignored and tuples dropped when the History Retention Period (HRP) is exceeded.
- `FILE` storage type in `createPrimaryProducer` and `createSecondaryProducer`: it was never implemented so nobody should miss it.
- `HTTP` protocol in URI in `createOnDemandProducer`: this was unused.
- `isExecuting` and `count` Consumer calls: use the new `endOfResults` call as described above.

12.4 FEATURES WITHDRAWN

- `setTerminationInterval` for all producers and consumers: this was deprecated in previous releases.
- `popAll` on a consumer: this was deprecated in previous releases.
- `setMetaData`, `setEndOfResults`, `addRow` and `setWarning` in the C++ API's `ResultSet` class: these were not deprecated previously, but should never have been made public.

13 KNOWN PROBLEMS AND CAVEATS

13.1 FUNCTIONALITY NOT YET IMPLEMENTED

Some API calls and options are not yet implemented by the R-GMA services. These are:

- LATEST or HISTORY producers with MEMORY storage.
- CONTINUOUS Secondary Producers.
- Combined LATEST and HISTORY producers.
- HISTORY and LATEST retention periods (functionality is approximated).
- Primary Producer `getHistoryRetentionPeriod()` method (functionality is approximated).
- Primary Producer `getLatestRetentionPeriod()` method (functionality is approximated).
- Secondary Producer `getHistoryRetentionPeriod()` method (functionality is approximated).
- The `getVersion()` method on all services.
- The following registry methods: `createRegistry()`, `destroyRegistry()`.
- The following schema methods: `createSchema()`, `destroySchema()`, `getTableIndex()`, `setAuthorizationRules()`, `getAuthorizationRules()`, `createIndex()`, `dropIndex()`, `createView()`, `dropView()`.
- The VDB name parameter on all registry and schema calls is currently ignored.

Schema replication is also not yet implemented.

13.2 OTHER KNOWN ISSUES

There are a number of other known bugs and limitations in the current release.

- Table and column names are currently case sensitive in R-GMA. To get consistent results always respect case.
- Continuous queries are case sensitive but latest and history queries ignore case with a default installation of MySQL.
- Only a subset of SQL92 is currently supported by R-GMA, and continuous queries are even more restricted. Time functions in SQL should not be used.

- Some operations of the command line tool try to access the schema and registry directly rather than just the local R-GMA server. This access may be blocked by a firewall.
- Use Java version 1.4.2.08 or later. There is a known issue with Java 1.4.2 running with the Scientific Linux 3 SMP kernel on dual processor machines, where the Java Virtual Machine can crash without warning. The problem does not seem to occur in RH7.3 or RHEL4.
- The Python API and the command line tool (that uses it) make a new connection to the server for each API operation. When connecting to a secure server this adds a substantial overhead, so multiple API calls (e.g. lots of inserts) will run slowly. The other three APIs are not affected.
- All VDBs share the same name space and information. Though the VDBName parameter is ignored, users are recommended to set it to an empty (zero length) string. This will facilitate the introduction of VDBs. Table names must not contain "." This will avoid confusion with VDB-Names when they become active.
- If authentication is enabled, to use the R-GMA Browser you must load your certificate into the web browser.
- When creating a table, integer columns must be specified as type INTEGER; the abbreviation INT does not work.

13.3 REPORTING BUGS AND GETTING HELP

If you think you have found a bug, please use Savannah at: <https://savannah.cern.ch/projects/jra1mdw/>. Check first to see if a bug report has already been submitted.

If you want to consult a human first, or if you have suggestions or need help using R-GMA please feel free to send an email to: <mailto:jra1-uk@physics.gla.ac.uk>.

REFERENCES

- [1] JRA1-UK. Information and Monitoring Service (R-GMA) System Specification. Technical Report EDMS 490223, EGEE, 2004.
- [2] Sql tutorial. <http://www.w3schools.com/sql/default.asp>.
- [3] Sql92 tutorial. <http://www.firstsql.com/tutor.htm>.
- [4] Web services architecture working group note. <http://www.w3.org/TR/ws-arch/>.
- [5] Web services description language. <http://www.w3.org/TR/wsdl/>.