R-GMA Server Installation

Tony Calanducci
INFN Catania - Italy

First Latin American Workshop for Grid Administrators
21-25 November 2005

Outline

- Introduction to R-GMA
- How to install it
- How to configure it
- Troubleshooting

Introduction to R-GMA

- **Relational Grid Monitoring Architecture (R-GMA)**
  - Provides Information (which resources are available on the Grid) and Monitoring Services
  - Developed as part of the EuropeanDataGrid Project (EDG)
  - Now as part of the EGEE project.
  - Implementation of the Grid Monitoring Architecture (GMA) from the Global Grid Forum (GGF).

- **Uses a relational data model.**
  - Data are viewed as tables.
  - Data structure defined by the columns.
  - Each entry is a row (tuple).
  - Queried using Structured Query Language (SQL).

Grid Monitoring Architecture (GMA)

- The Producer stores its location (URL) in the Registry.
- The Consumer looks up producer URLs in the Registry.
- The Consumer contacts the Producer to get all the data or the Consumer can listen to the Producer for new data.

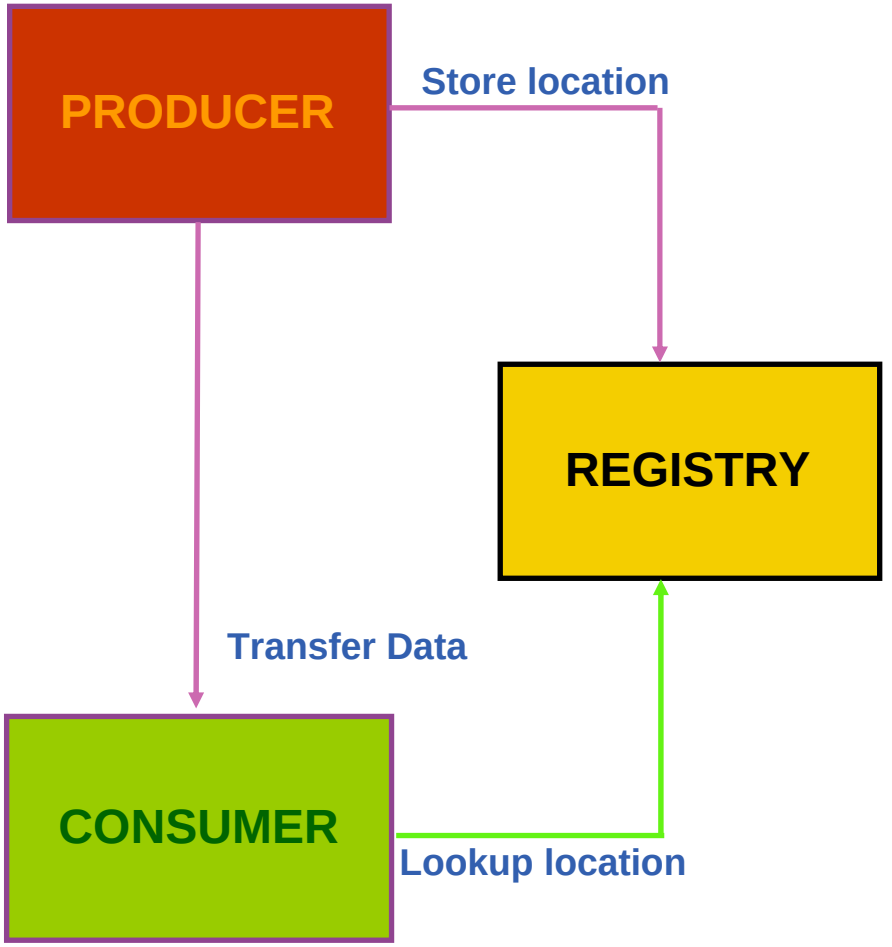

The Relational Model

name	ID	birth	Group
Tom	4	1977-08-20	HR

SELECT * FROM people WHERE group='HR'

R-GMA

- The data model is relational.
- The table definition is globally unique and is stored in the Schema.
- The Registry stores the Producers table name as well as the URL.
- The data is inserted in the form of a tuple.
- The Consumer gets the tuple from Producer.
- **Producers**
  - publish: SQL “INSERT”
- **Consumers**
  - collect: SQL “SELECT”

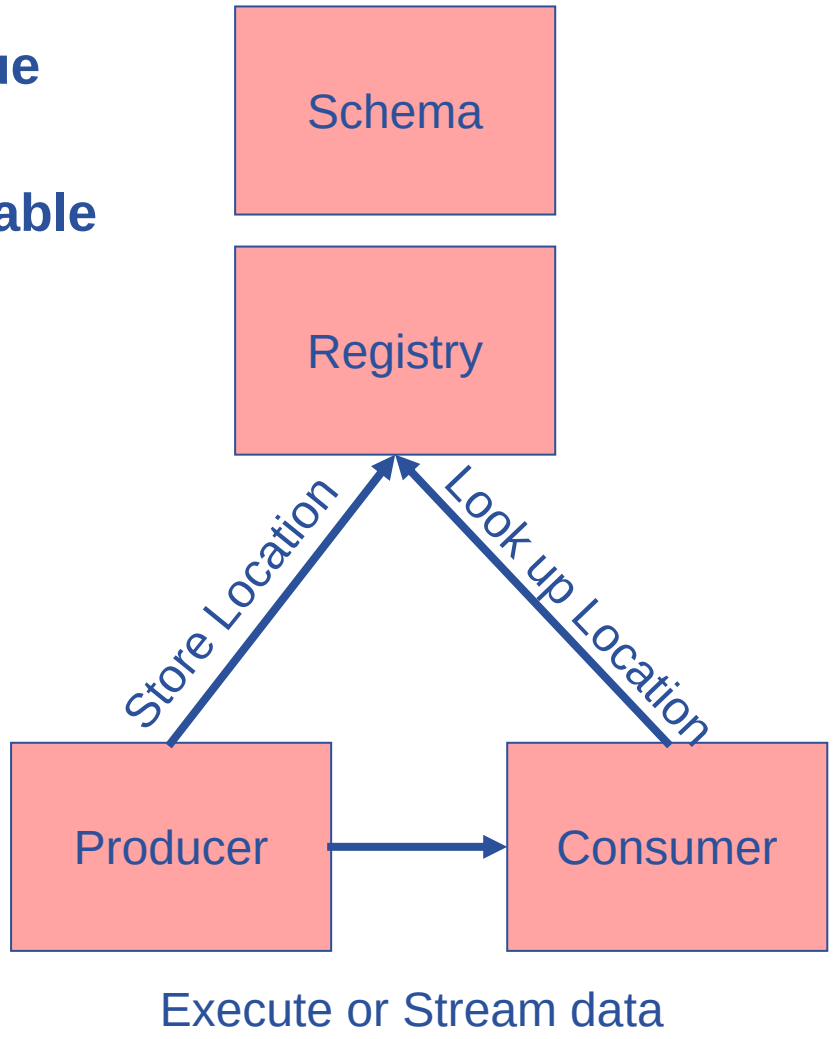

Multiple Producers

- The Consumer will get all the URLs that could satisfy the query.
- The Consumer will connect to all the Producers.
- Producers that can satisfy the query will send the tuples to the Consumer.
- The Consumer will merge these tuples to form one result set.

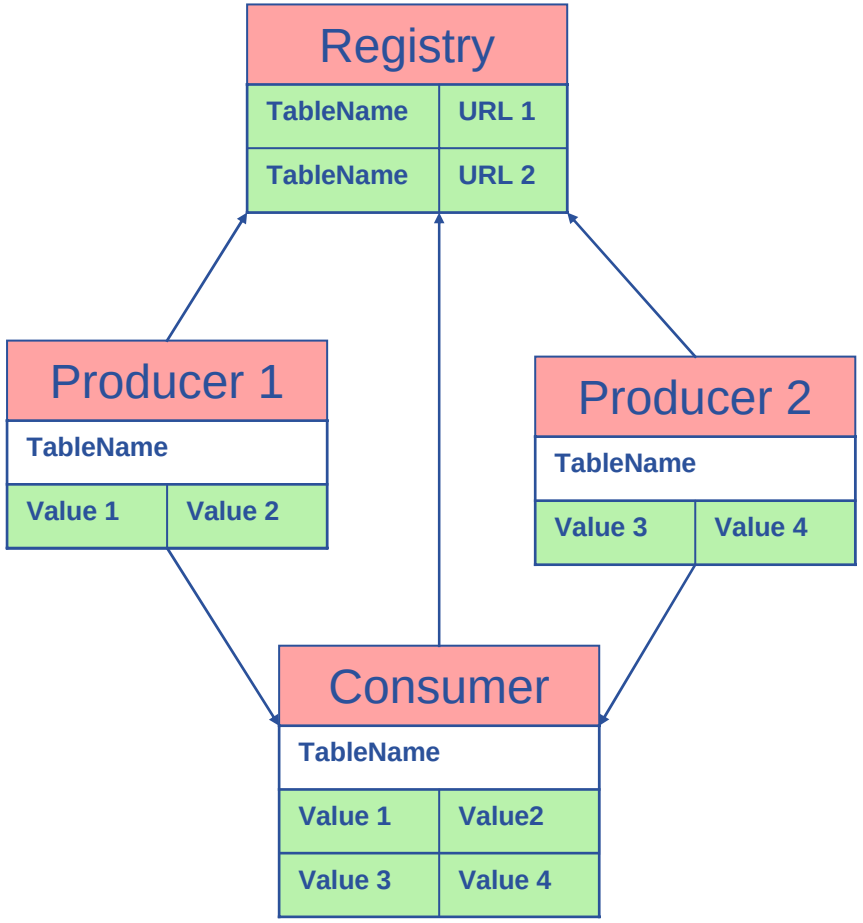

Select * from CPULoad

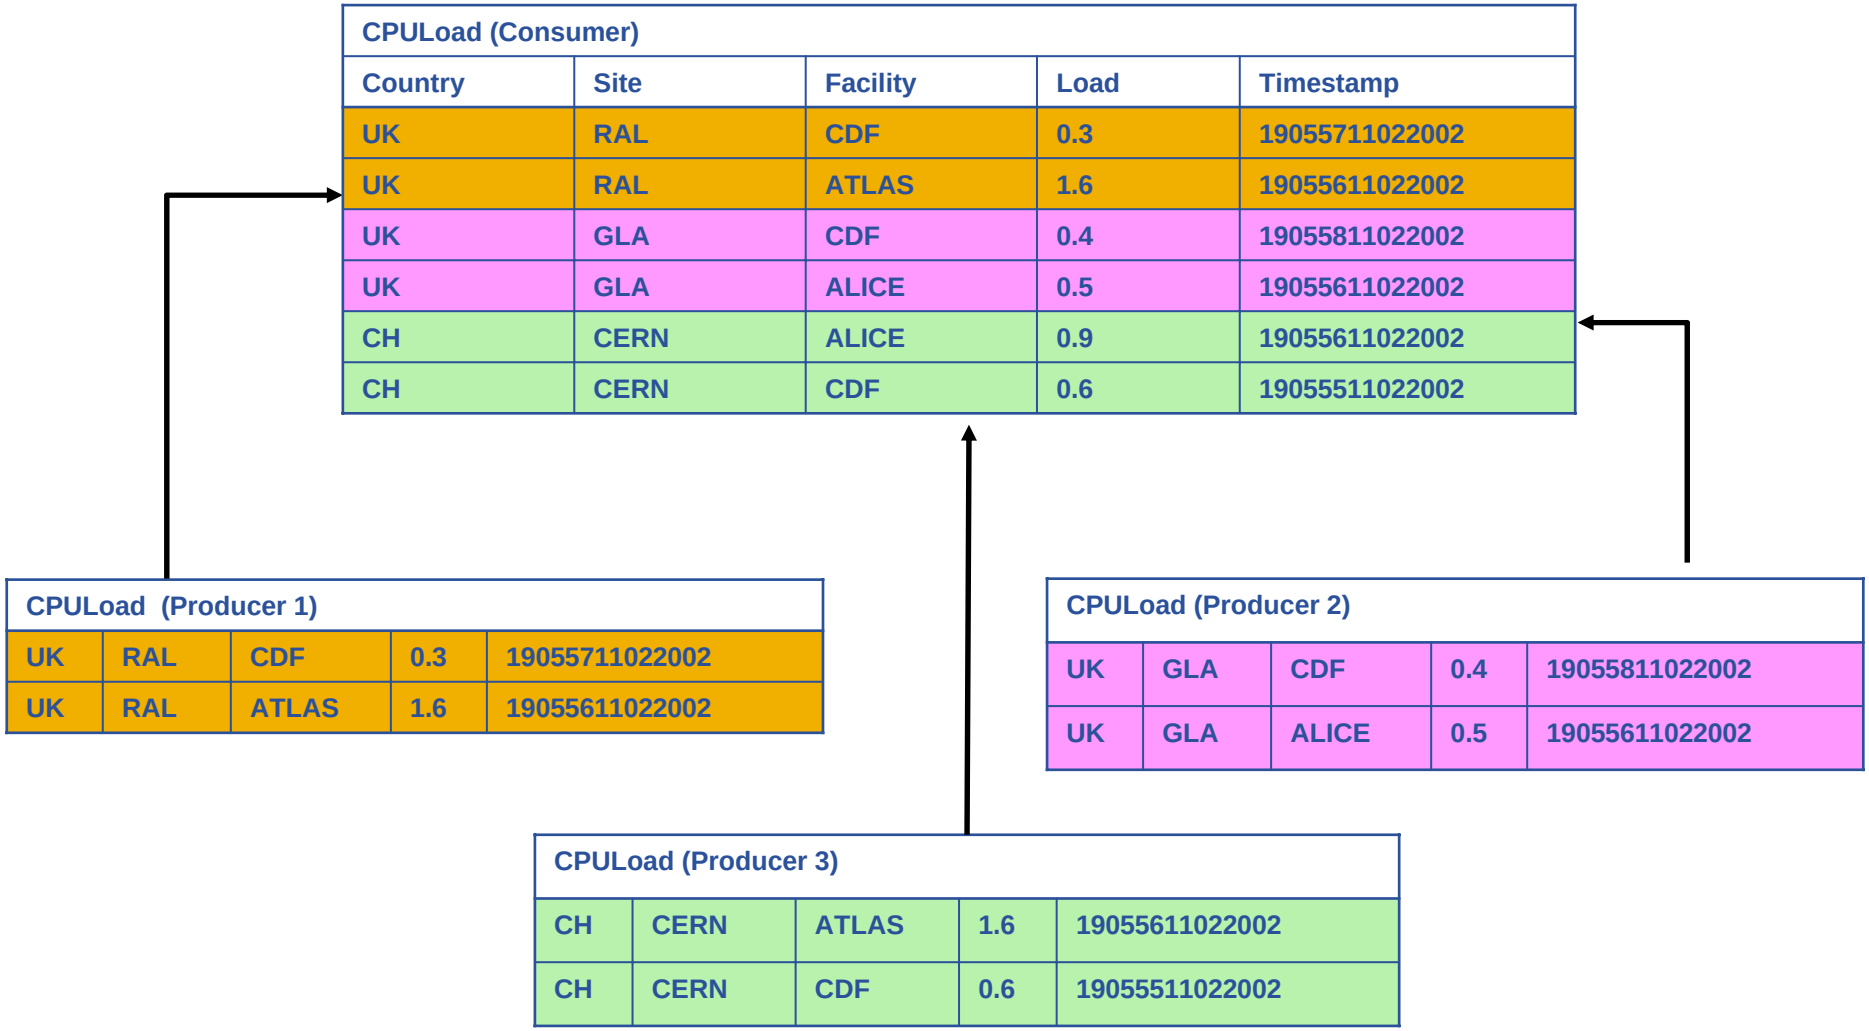

CPULoad (Consumer)

Country	Site	Facility	Load	Timestamp
UK	RAL	CDF	0.3	19055711022002
UK	RAL	ATLAS	1.6	19055611022002
UK	GLA	CDF	0.4	19055811022002
UK	GLA	ALICE	0.5	19055611022002
CH	CERN	ALICE	0.9	19055611022002
CH	CERN	CDF	0.6	19055511022002

CPULoad (Producer 1)

UK	RAL	CDF	0.3	19055711022002
UK	RAL	ATLAS	1.6	19055611022002

CPULoad (Producer 2)

UK	GLA	CDF	0.4	19055811022002
UK	GLA	ALICE	0.5	19055611022002

CPULoad (Producer 3)

CH	CERN	ATLAS	1.6	19055611022002
CH	CERN	CDF	0.6	19055511022002

The Mediator

- **The Mediator is the intelligence of R-GMA**
  - Not a single component, but distributed.
  - Enables queries to be accurately and efficiently returned.
- **The table name is stored next to the URL in the Registry.**
  - For simple queries, only the URLs that can answer query are passed to the Consumer.
  - If the query has a predicate, only the URLs that could satisfy the query will be passed to the Consumer.
- **The Mediator will also try to do joins.**
  - For complex queries the query must use a Producer with a database backend (secondary producer).
  - Merges and produces the resulting result set.
- **The Consumers URL and query is also stored in the Registry.**
  - Enables the Registry to notify listening Consumers about new Producers.

Where country='UK'

CPULoad (Consumer)				
Country	Site	Facility	Load	Timestamp
UK	RAL	CDF	0.3	19055711022002
UK	RAL	ATLAS	1.6	19055611022002
UK	GLA	CDF	0.4	19055811022002
UK	GLA	ALICE	0.5	19055611022002

CPULoad (Producer 1)				
UK	RAL	CDF	0.3	19055711022002
UK	RAL	ATLAS	1.6	19055611022002

CPULoad (Producer 2)				
UK	GLA	CDF	0.4	19055811022002
UK	GLA	ALICE	0.5	19055611022002

CPULoad (Producer 3)				
CH	CERN	ATLAS	1.6	19055611022002
CH	CERN	CDF	0.6	19055511022002

Producer and Query Types

- **There are two different types of producers.**
  - Primary Producer.
    - In memory or database.
    - Publish data.
  - Secondary Producers.
    - Republish data.
- **Producers can have different properties providing three different query type.**
  - Stream (Continuous)
    - Sends new tuples directly to the Consumer.
  - Latest.
    - Only stores that latest tuple for the primary key.
    - Latest Retention Period, property of tuple.
  - History.
    - Stores all tuples.
    - History Retention Period, property of the producer.

History or Latest Query

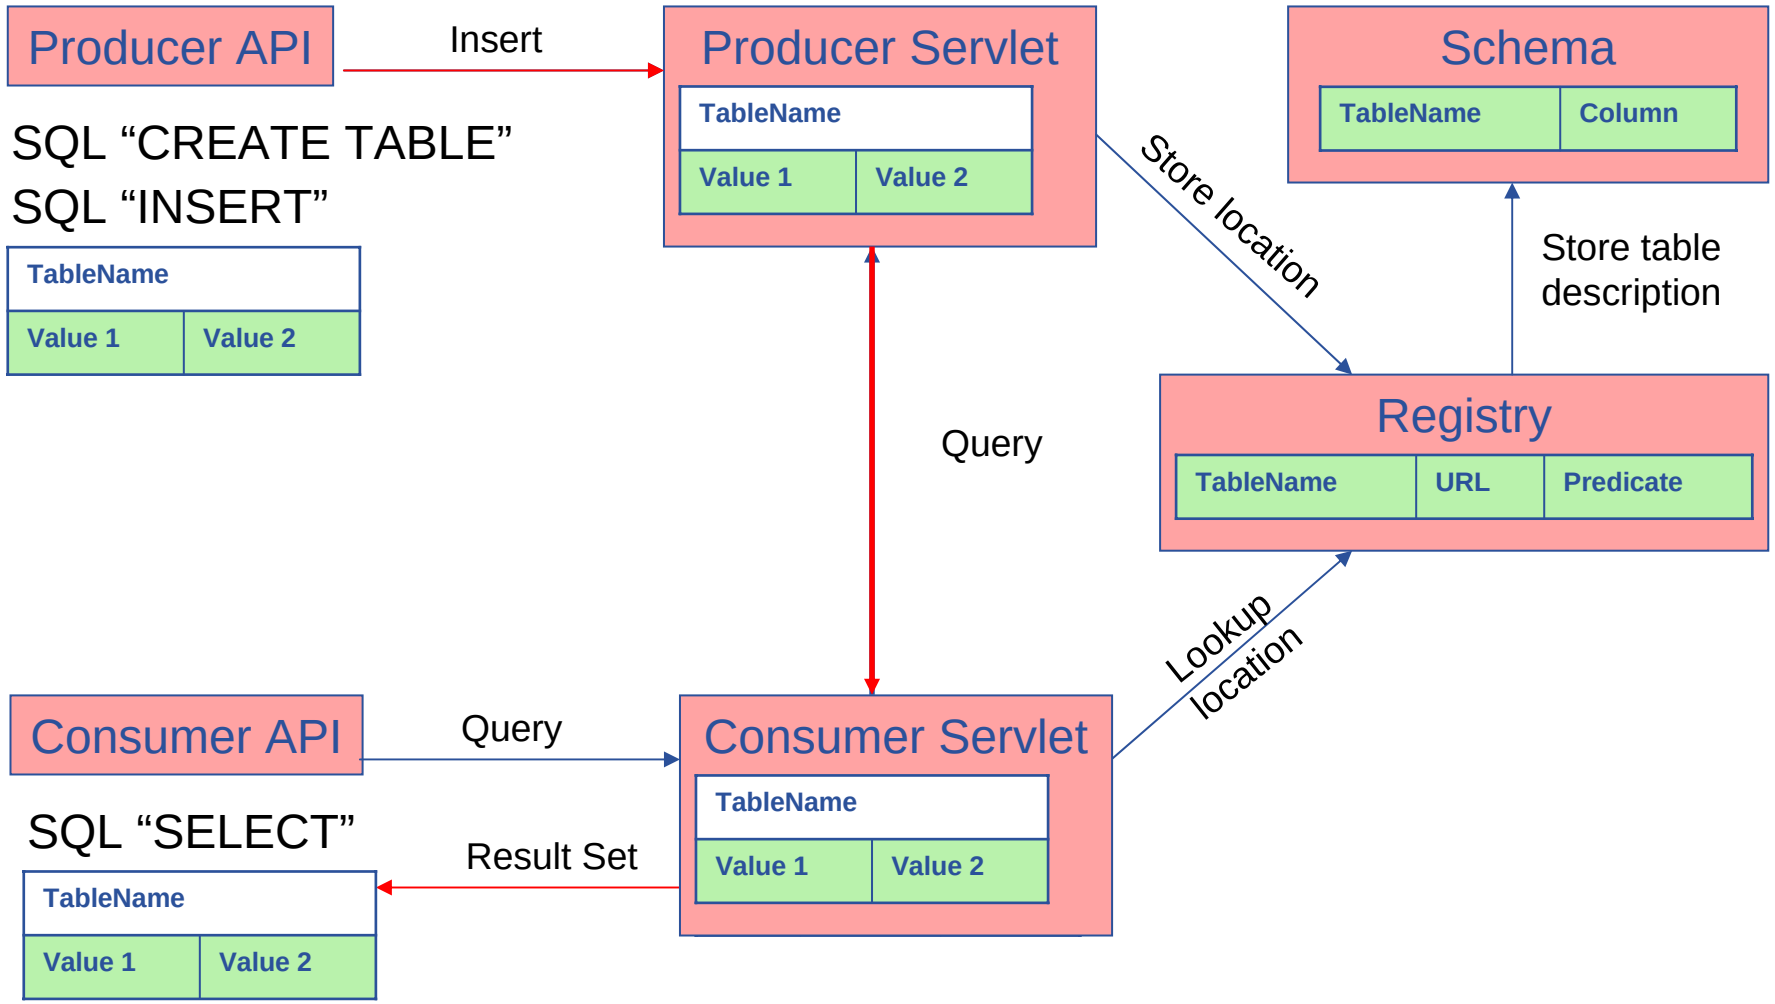

R-GMA summary

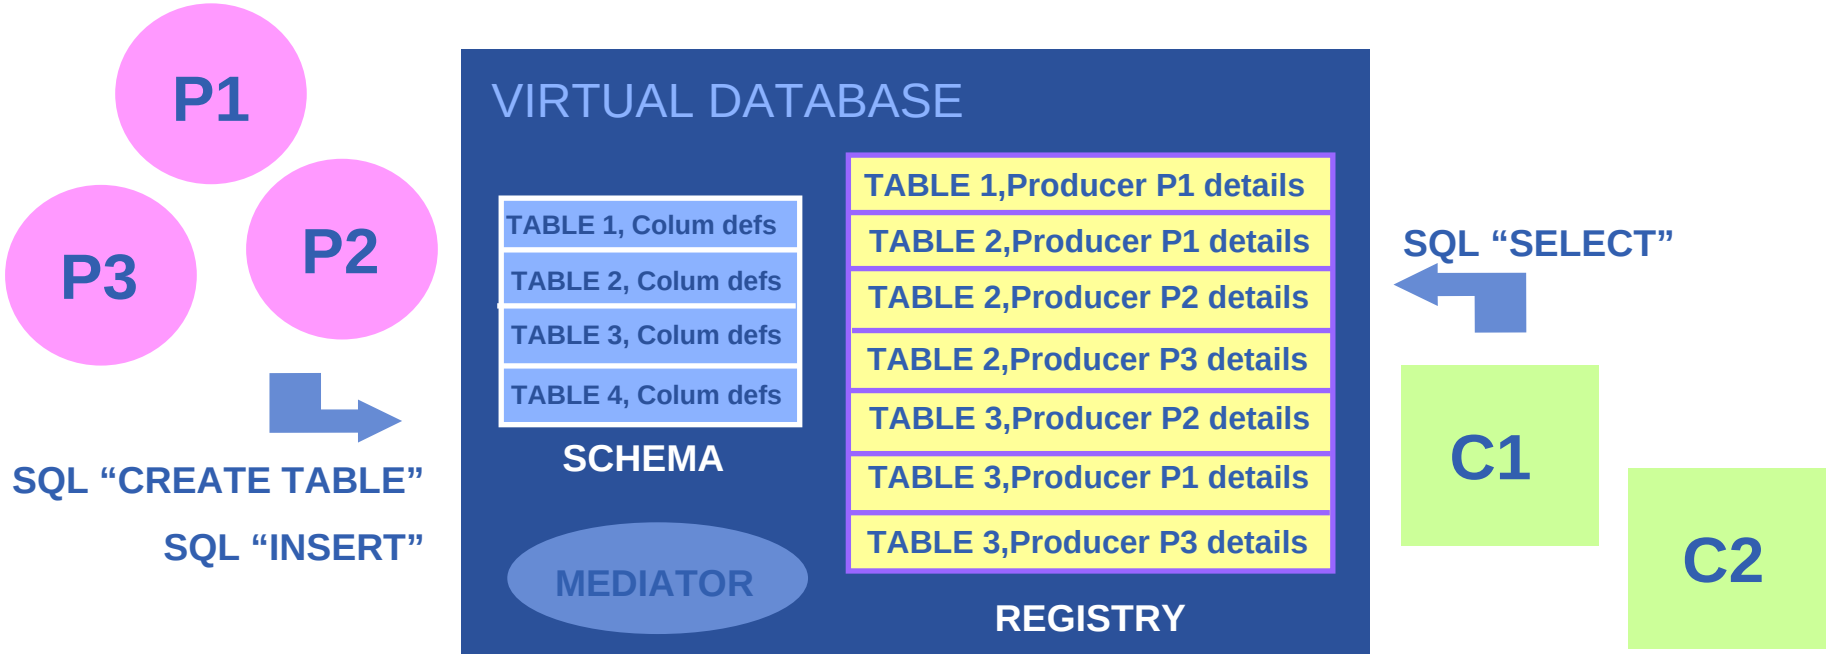

There is no central repository!!! There is only a "*Virtual Database*".

Schema is a list of table definitions.

Registry is a list of data producers with all its details.

Producers publish data.

Consumer read data published.

R-GMA Server Installation

Installation Pre-requisites

- Start from a fresh install of SLC 3.0.x

- Java JRE or JDK are required to run R-GMA Server. Due to license issues, it is not distributed by gLite. So please download and install a release of Java JRE/JDK >= 1.4.2_08 from http://java.sun.com/j2se/1.4.2/download.html

```
chmod +x j2sdk-1_4_2_10-linux-i586-rpm.bin
./j2sdk-1_4_2_10-linux-i586-rpm.bin
rpm -ivh j2sdk-1_4_2_10-linux-i586.rpm
Preparing...                ########################################### [100%]
   1:j2sdk                  ########################################### [100%]
```

Security pre-requisites

- Request host certificates for R-GMA Server from your CA. For example, to get host certs from GILDA CA go to:
  - https://gilda.ct.infn.it/CA/mgt/restricted/srvreq.php
- Install host certificates (hostcert.pem and hostkey.pem) in **/etc/grid-security**.
  - *chmod 644 hostcert.pem*
  - *chmod 400 hostkey.pem*

- If you plan to use certificates released by unsupported EGEE CA's, be sure that their public key, signing policy and CRLs (usually distributed with an rpm) are installed.
  - For the VO GILDA, the RPM is available from

  https://gilda.ct.infn.it/RPMS/ca_GILDA-0.28.1.i386.rpm

Installing R-GMA Server via apt

1. **Verify if apt is present in your installation:**
   - **rpm -qa | grep apt**
   - **Install apt if necessary:**
     - **rpm -ivh** http://linuxsoft.cern.ch/cern/slc30X/i386/SL/RPMS/apt-0.5.15cnc6-8.SL.cern.i386.rpm

2. **Add gLite apt repository:**
   - Put this line in a new file (e.g. glite.list) inside the /etc/apt/sources.list.d directory (R 1.4):

```
rpm http://glitesoft.cern.ch/EGEE/gLite/APT/R1.4/ rhel30
   externals Release1.4 updates
```

   - `apt-get update`
   - `apt-get upgrade`

3. **Install R-GMA Server:**
   - **`apt-get install glite-rgma-server-config`**

More info at http://glite.web.cern.ch/glite/packages/APT.asp

Installing R-GMA Server via apt (cont.)

- This will install the following deployment modules:
  - R-GMA server
  - R-GMA servicetools
  - Security Utils
- If the installation is performed successfully, the following components are installed:
  - *gLite in /opt/glite*
  - *gLite-essentials-java in /opt/glite/externals/share*
  - *MySQL-server e MySQL-client in /usr*
  - *Tomcat in /var/lib/tomcat5*
- The gLite R-GMA server configuration script is installed in
  - $GLITE_LOCATION/etc/config/scripts/glite-rgma-server-config.py
- Templates configurations files are located in
  - $GLITE_LOCATION/etc/config/templates

R-GMA Server configuration

- **Copy the configuration template files from**
  - $GLITE_LOCATION/etc/config/templates
- **To**
  - $GLITE_LOCATION/etc/config

- **P.S: You can delete the *glite-rgma-servicetool-serviceName.cfg.xml* template file because it will not be used.**

- **Now we will start to customize the just copied configuration files by replacing the “changeme” value in all user-defined parameters with the proper value.**

R-GMA Server Configuration (cont.)

- **List of XML files to customize:**

glite-global.cfg.xml
glite-rgma-common.cfg.xml
glite-rgma-server.cfg.xml
glite-rgma-servicetool.cfg.xml
glite-security-utils.cfg.xml

glite-global.cfg.xml

<JAVA_HOME description="Environment variable pointing to the SUN Java JRE or J2SE package for example '/usr/java/j2re1.4.2_08/' or '$JAVA_HOME' (if it is defined as an environment variable)" value="/usr/java/j2sdk-1.4.2_10"/>

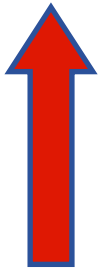

Check the correct path of your JAVA JRE/JDK.

glite-rgma-common.cfg.xml

<rgma.server.hostname
description="Host name of the R-GMA server.
[Example: lxb1420.cern.ch] [Type: 'string']"
value="**rgma-test.trigrid.it**"/>

<rgma.schema.hostname
description="Host name of the R-GMA schema service.
(See also configuration parameter 'rgma.server.run_schema_service'
in the R-GMA server configuration file in case you install a server).
[Example: lxb1420.cern.ch] [Type: 'string']"
value="**rgma-test.trigrid.it**"/>

<rgma.registry.hostname
description="Host name of the R-GMA registry service.
You must specify at least one hostname and you can specify several
if you want to use several registries.
(See also configuration parameter 'rgma.server.run_registry_service'
in the R-GMA server configuration file in case you install a server).
[Example: lxb2029.cern.ch] [Type: 'string']">
<value>**rgma-test.trigrid.it**</value>
</rgma.registry.hostname>

glite-rgma-server.cfg.xml

```
<rgma.server.run_schema_service
        description="Run a schema service for the R-GMA server on your machine yes|no).
        value="yes"/>

<rgma.server.run_registry_service
        description="Run a registry service for the R-GMA server on your machine (yes|no)
        value="yes"/>

<rgma.server.run_browser
        description="Run a browser (yes|no). Running a browser is optional but useful.
        value="yes"/>

<rgma.server.run_archiver
        description="Run the R-GMA data archiver (yes|no).
        value="yes"/>

<rgma.server.run_site-publisher
        description="Run the R-GMA site-publisher (yes|no).
        value="yes"/>

<rgma.server.mysql_root_password
        description="MySQL root password.
        value="HolaVenezuela"/>
```

glite-rgma-server.cfg.xml (cont.)

```
<!-- site-publisher specific configuration parameters -->

        <rgma.site-publisher.contact.system_administrator
              description="Contact email address of the site system administrator.
        value="tony.calanducci@ct.infn.it"/>

        <rgma.site-publisher.contact.user_support
              description="Contact email address of the user support.
        value="usersupport@ct.infn.it"/>

        <rgma.site-publisher.contact.site_security
              description="Contact email address of the site security responsible.
        value="tony.calanducci@ct.infn.it"/>

        <rgma.site-publisher.location.latitude
              description="Latitude of your site. Please go to
                    'http://www.multimap.com/' to find the correct value for your site.
        value="8.6015"/>

        <rgma.site-publisher.location.longitude
              description="Longitude of your site.
        value="-71.1487"/>
```

glite-rgma-servicetool.cfg.xml

- <rgma.servicetool.sitename

  description="DNS name of the site publisher node. This parameter must have the same value as the rgma.site- publisher.sitename parameter in the R-GMA Server configuration.

  value="rgma-test.trigrid.it"/>

- <rgma.servicetool.vo

  description="List of VOs that this service is considered part of.>

  <value>gilda</value>

  </rgma.servicetool.vo>

glite-security-utils.cfg.xml

```
<cron.mailto
        description="E-mail address for sending cron job notifications"
        value="root@localhost"/>
```

Configure MySQL

- Make sure that the MySQL root password that you have specified in the configuration file (**glite-rgma-server.cfg.xml**, *rgma.server.mysql_root_password* **paramenter**) matches the password that is set in the MySQL database.
- To set the MySQL root password you should issue the following commands as root:

/usr/bin/mysqladmin –u root password 'SaluTony'

/usr/bin/mysqladmin –u root –h rgma-test.trigrid.it password 'SaluTony'

Configure and start the R-GMA Server

- Change to the script directory:
  - cd /opt/glite/etc/config/scripts
- Configure the R-GMA server by executing the R-GMA Server configuration script:
  - ./glite-rgma-server-config.py --configure
- If no error message is reported by the previous step, start the R-GMA server by running
  - ./glite-rgma-server-config.py --start

Verify the installation

- Check that all the R-GMA services are running properly by running:
  - ./glite-rgma-server-config.py --status
- or by connecting to the R-GMA Browser with your Internet Browser at the following address:
  - https://yourRGMAserverHostname:8443/R-GMA/
- If you want to stop the R-GMA Server, you can do it by issuing
  - ./glite-rgma-server-config.py --stop

R-GMA Browser screenshoot

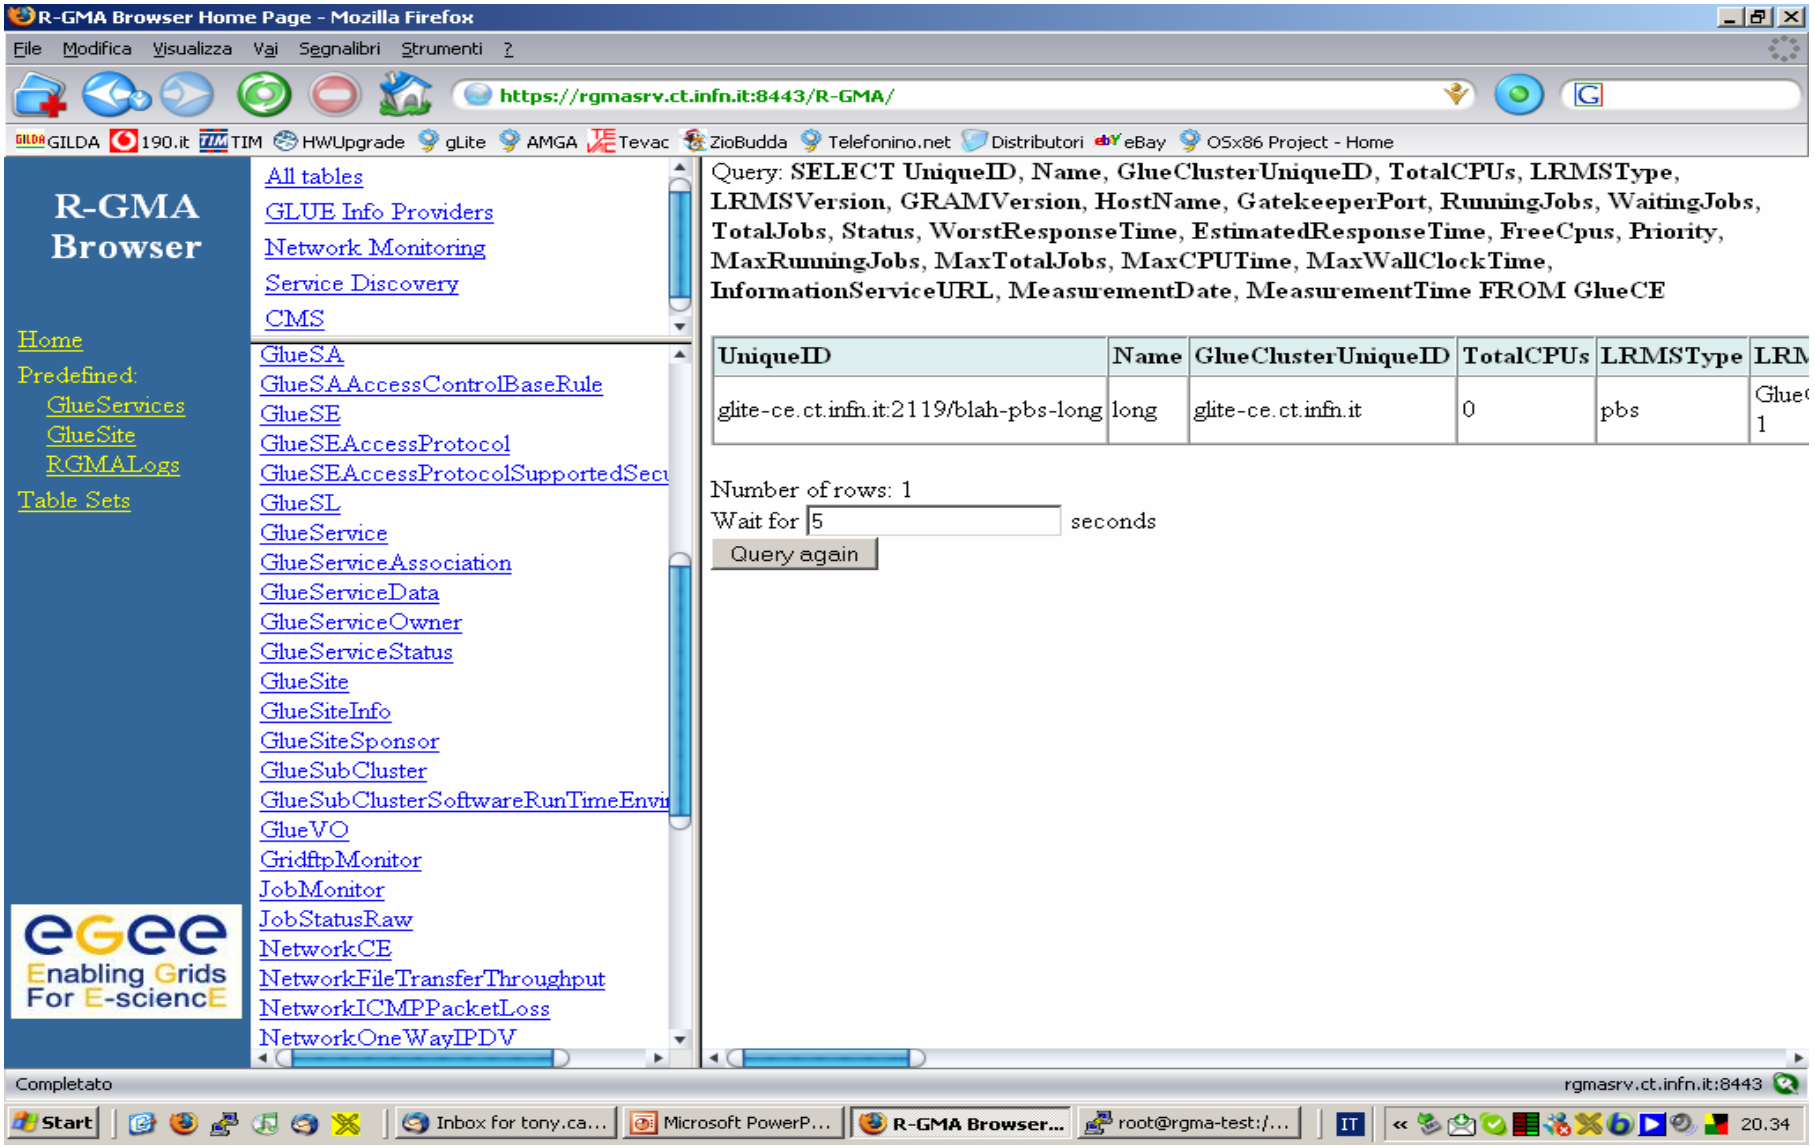

Troubleshooting

- **If you get the following error during service startup:**
  - `Starting R-GMA glue-archiver. Please be patient this may take several minutes ...`
    `Starting R-GMA glue-archiver [FAILED]`
    `An error occurred while starting the gLite R-GMA Server service`

- **Check you host certificates and you (above all) have installed you Certification Autority certificates!!!**