The Information System

Mike Mineter, NeSC
Ruediger Berlich, FZK

2 February 2006, Brisbane

Acknowledgement

- **Most slides are from EGEE presentations prepared by INFN / University of Catania**

Please download this file from the agendamaker page:

- Go to http://egee.nesc.ac.uk
- Choose "**events and registration**" from menu
- Select "**jump to now**" for latest events
- Select **Brisbane** event
- Download **Information Systems** talk

Overview

- Information System concepts
- Information Systems in the current production system
- Information Systems for gLite

Current production middleware

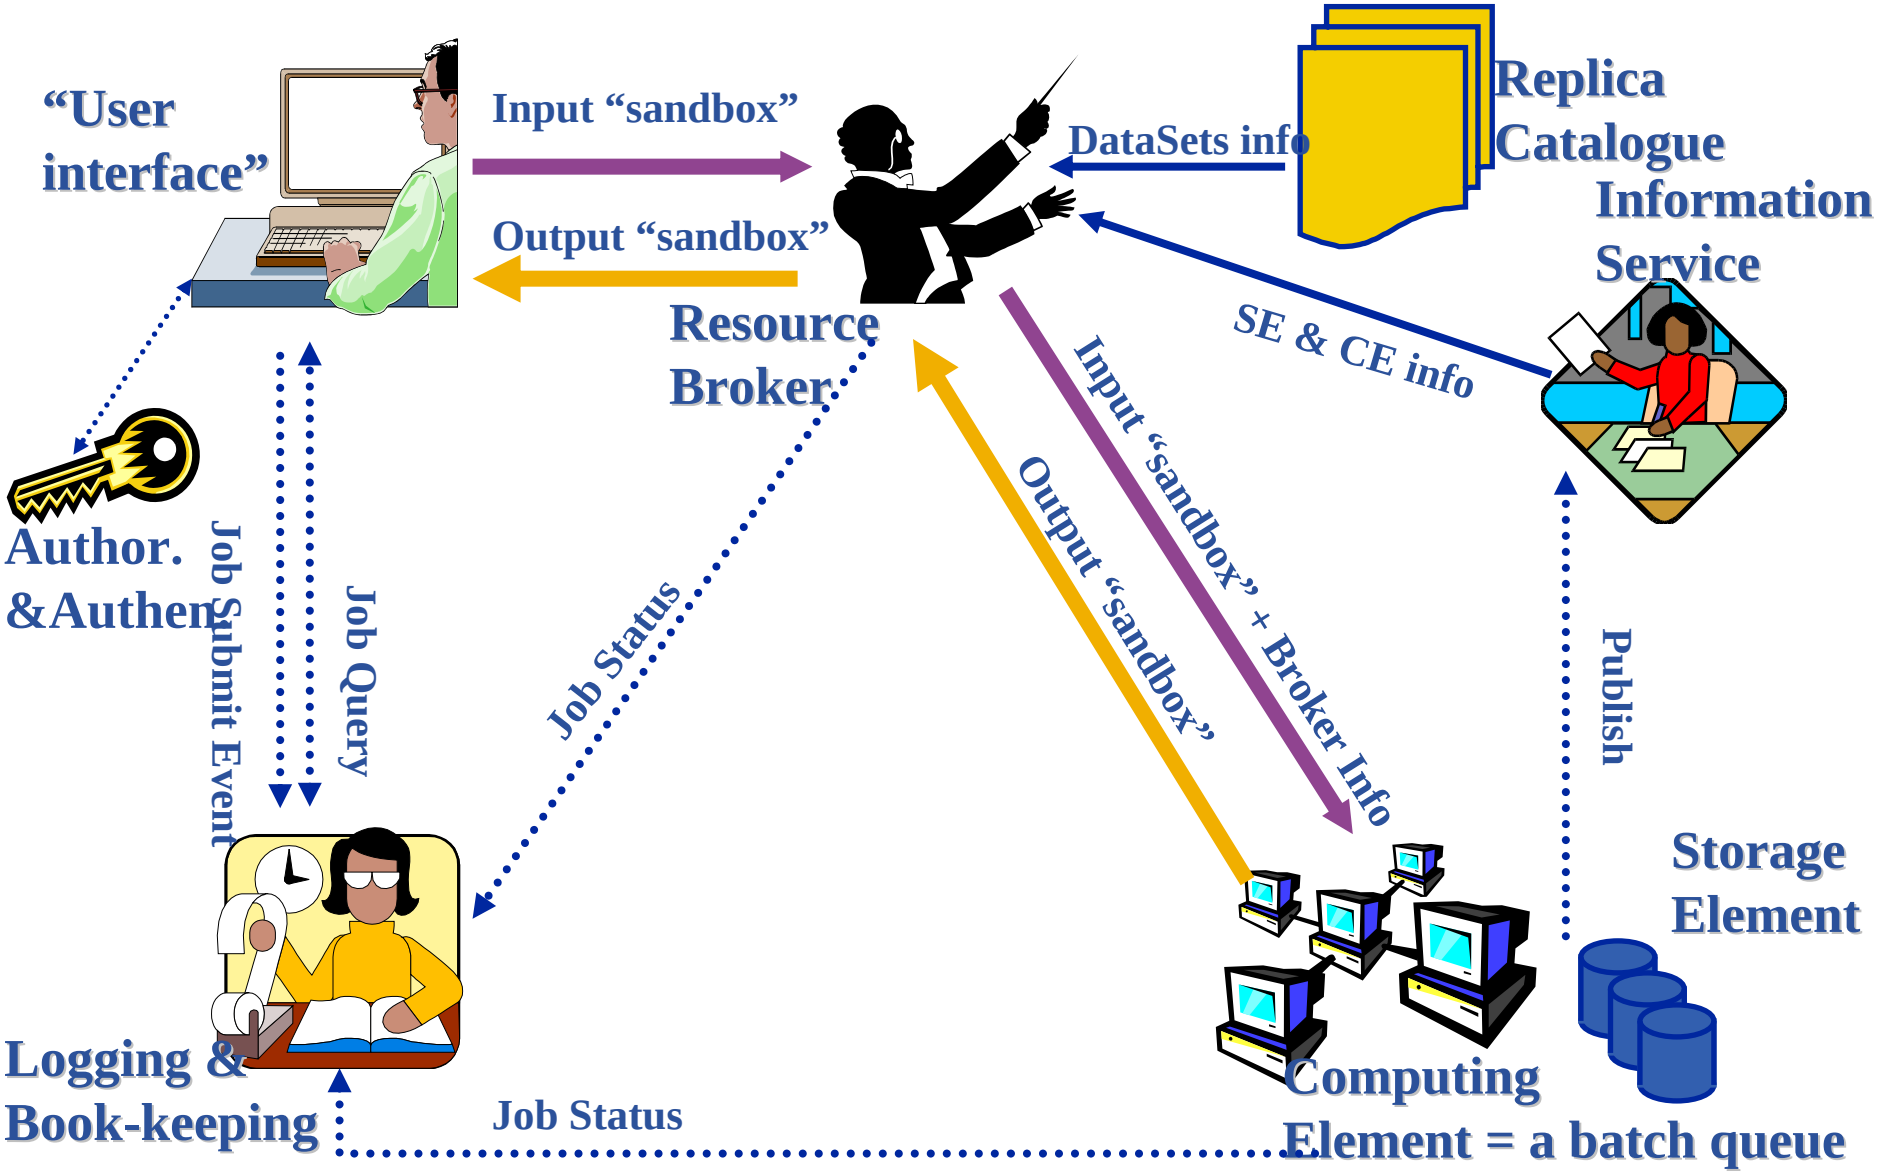

Uses of the Information System

If you are a user

Retrieve information about
- Grid resources and status
- Resources that can run your job
- Status of your jobs

If you are a middleware developer

Workload Management System:
Matching job requirements and Grid resources

Monitoring Services:
Retrieving information of Grid Resources status and availability

If you are site manager or service

You “generate” the information for example relative to your site or to a given service

Evolution

- The data published in the IS conforms to the GLUE (Grid Laboratory for a Uniform Environment) Schema. The GLUE Schema aims to define a common conceptual data model to be used for Grid resources. http://infnforge.cnaf.infn.it/glueinfomodel/

- In LCG-2, the BDII (Berkeley DB Information Index), based on an updated version of the Monitoring and Discovery Service (MDS), from Globus, was adopted as main provider of the Information Service.
- R-GMA (Relational Grid Monitoring Architecture) is now adopted as IS in both the EGEE production grid (mainly “LCG-2”) and in the pre-production grid (with new gLite services)

The LDAP Protocol

- **Lightweight Directory Access Protocol:**
  structures data as a tree

- Following a path from the node back to the root of the DIT, a unique name is built (the DN):

```
"id=pml,ou=IT,or=CERN,st=Geneva, \
c=Switzerland,o=grid"
```

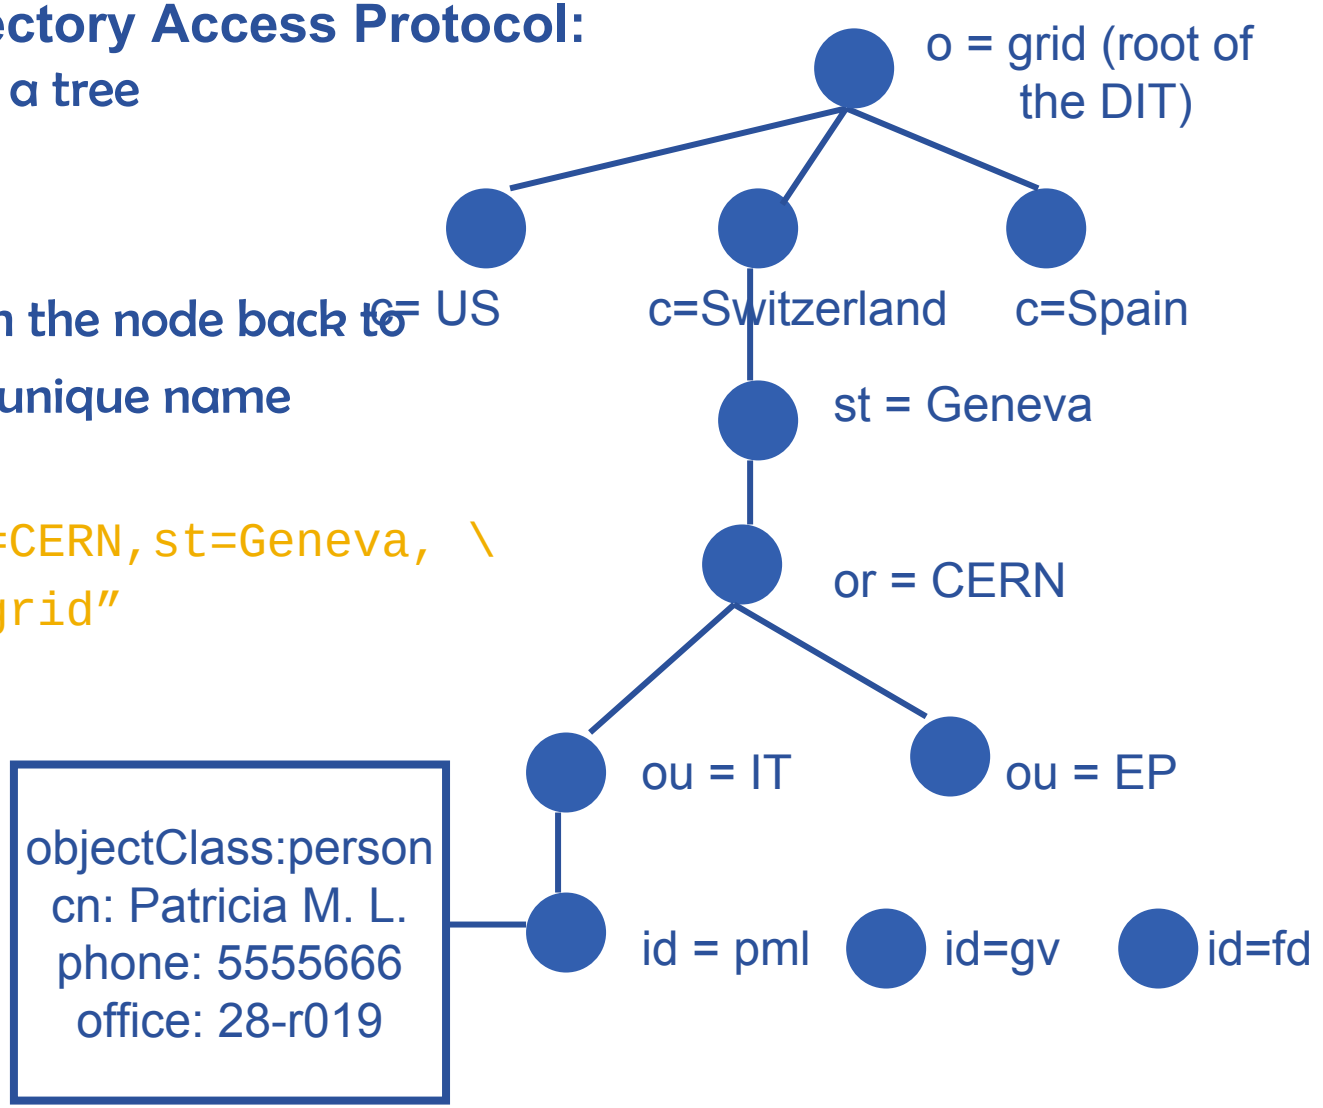

lcg-infosites

LCG Information Service

- a user or a service can query
  - the BDII (usual mode)
  - LDAP servers on each site

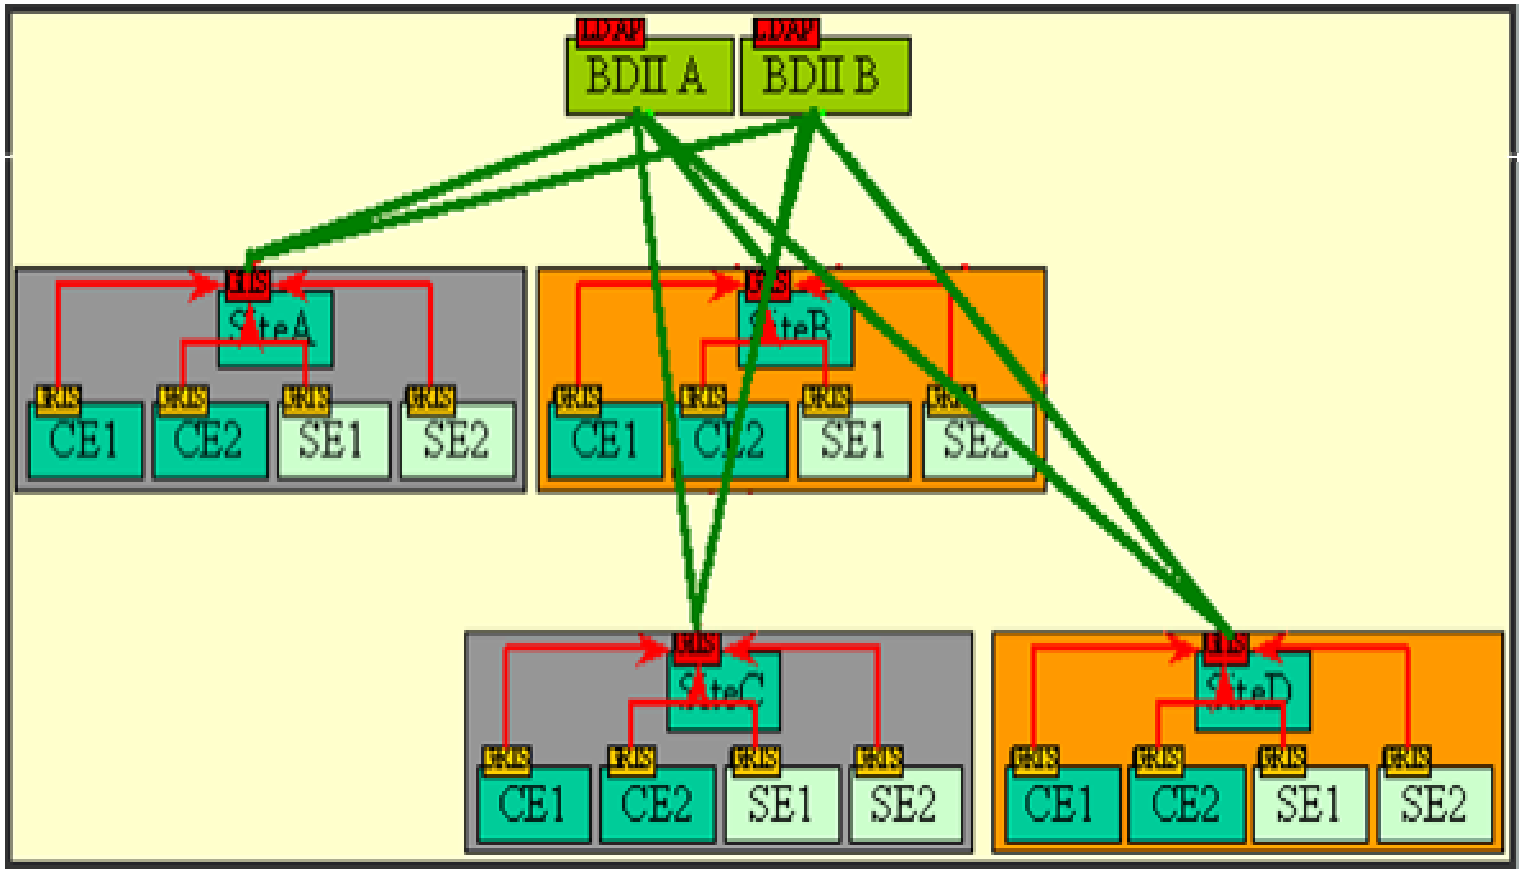

lcg-infosites

- The lcg-infosites command can be used as an easy way to retrieve information on Grid resources for most use cases. (For others, use lcg-info)

USAGE: lcg-infosites --vo <vo name> options -v <verbose level> --is <BDII to query>

lcg-infosites options

ce	The information related to number of CPUs, running jobs, waiting jobs and names of the CEs are provided. All these data group all VOs together. With "-v 1" only the names of the queues will be printed while with "-v 2" The RAM Memory together with the operating system and its version and the processor included in each CE are printed.
se	The names of the SEs supported by the user's VO together with the kind of Storage System,the used and available space will be printed. With "-v 1" only the names of the SEs will be printed.
closeSE	The names of the CEs where the user's VO is allowed to run together with their corresponding closest SEs are provided.
lfc	Name of the lfc Catalog for the user's VO.
tag	The names of the tags relative to the software installed in site is printed together with the corresponding CE.
all	It groups together the information provided by ce, se, lrc and rmc.
is	If not specified the BDII defined in default by the variable LCG GFAL INFOSYS will be queries. However the user may want to query any other BDII without redefining this environment variable. This is possible specifying this argument followed by the name of the BDII which the user wants to query. All options admits this argument.

PRACTICAL: lcg-infosites

- In the next 15 minutes, run the commands shown in following slides to explore GILDA using lcg-infosites

Obtaining information about CE

$ lcg-infosites --vo gilda ce

```
*************************************************************
These are the related data for gilda: (in terms of queues and CPUs)
*************************************************************
```

#CPU	Free	Total Jobs	Running	Waiting	ComputingElement
4	3	0	0	0	cn01.be.itu.edu.tr:2119/jobmanager-lcglsf-long
4	3	0	0	0	cn01.be.itu.edu.tr:2119/jobmanager-lcglsf-short
34	33	0	0	0	grid010.ct.infn.it:2119/jobmanager-lcgpbs-long
16	16	0	0	0	grid011f.cnaf.infn.it:2119/jobmanager-lcgpbs-long
1	1	0	0	0	grid006.cecalc.ula.ve:2119/jobmanager-lcgpbs-log
2	1	1	0	1	gildace.oact.inaf.it:2119/jobmanager-lcgpbs-short

[..]

$ lcg-infosites --vo gilda ce --v 2

RAMMemory	Operating System	System Version	Processor	CE Name
1024	SLC	3	P4	ced-ce0.datagrid.cnr.it
4096	SLC	3	Xeon	cn01.be.itu.edu.tr
1024	SLC	3	PIII	cna02.cna.unicamp.br
917	SLC	3	PIII	gilda-ce-01.pd.infn.it
1024	SLC	3	Athlon	gildace.oact.inaf.it
1024	SLC	3	Xeon	grid-ce.bio.dist.unige.it

[..]

Obtaining information about SE

```
$ lcg-infosites --vo gilda se

****************************************************************
These are the related data for gilda: (in terms of SE)
****************************************************************

Avail Space(Kb) Used Space(Kb)    Type            SEs
----------------------------------------------------------------------------
143547680       2472756           disk            cn02.be.itu.edu.tr
168727984       118549624         disk            grid009.ct.infn.it
13908644        2819288           disk            grid003.cecalc.ula.ve
108741124       2442872           disk            gildase.oact.inaf.it
28211488        2948292           disk            testbed005.cnaf.infn.it
349001680       33028             disk            gilda-se-01.pd.infn.it
31724384        2819596           disk            cna03.cna.unicamp.br
387834656       629136            disk            grid-se.bio.dist.unige.it
```

Listing the close Storage Elements

$ lcg-infosites --vo gilda closeSE

Name of the CE: cn01.be.itu.edu.tr:2119/jobmanager-lcglsf-long
Name of the close SE: cn02.be.itu.edu.tr

Name of the CE: cn01.be.itu.edu.tr:2119/jobmanager-lcglsf-short
Name of the close SE: cn02.be.itu.edu.tr

Name of the CE: grid010.ct.infn.it:2119/jobmanager-lcgpbs-long
Name of the close SE: grid009.ct.infn.it

Name of the CE: grid011f.cnaf.infn.it:2119/jobmanager-lcgpbs-long
Name of the close SE: testbed005.cnaf.infn.it

Listing tags of installed software

$ lcg-infosites --vo gilda tag

****
Information for gilda relative to their software tags included in each CE
****

Name of the TAG: VO-gilda-GEANT
Name of the TAG: VO-gilda-GKS05
Name of the CE:cn01.be.itu.edu.tr

Name of the TAG: VO-gilda-slc3_ia32_gcc323
Name of the TAG: VO-gilda-CMKIN_5_1_1
Name of the TAG: VO-gilda-GEANT
Name of the TAG: VO-gilda-GKS05
Name of the CE:grid010.ct.infn.it

[..]

lcg-info intro

- This command can be used to list either CEs or the SEs that satisfy a given set of conditions, and to print the values of a given set of attributes.
- The information is taken from the BDII specified by the LCG_GFAL_INFOSYS environment variable.

- The query syntax is like this:

```
attr1 op1 valueN, ... attrN opN valueN
```

where *attrN* is an attribute name

op is =, >= or <=, and the cuts are ANDed.

The cuts are comma-separated and spaces are not allowed.

lcg-info usage

USAGE

lcg-info --list-ce [--bdii bdii] [--vo vo] [--sed] [--query query] [--attrs list]

lcg-info --list-se [--bdii bdii] [--vo vo] [--sed] [--query query] [--attrs list]

lcg-info --list-attrs

lcg-info --help

lcg-info options

--list-attrs	Prints a list of the attributes that can be queried.
--list-ce	Lists the CEs which satisfy a query, or all the CEs if no query is given.
--list-se	Lists the SEs which satisfy a query, or all the SEs if no query is given.
--query	Restricts the output to the CEs (SEs) which satisfy the given query.
--bdii	Allows to specify a BDII in the form :. If not given, the value of the environmental variable LCG_GFAL_INFOSYS is used. If that is not defined, the command returns an error.
--sed	Print the output in a "sed-friendly" format.
--attrs	Specifies the attributes whose values should be printed.
--vo	Restricts the output to CEs or SEs where the given VO is authorized. Mandatory when VO-dependent attributes are queried upon.

Get the list of supported attributes

```
$ lcg-info --list-attrs
```

Attribute name	Glue object class	Glue attribute name
MaxTime	GlueCE	GlueCEPolicyMaxWallClockTime
CEStatus	GlueCE	GlueCEStateStatus
TotalJobs	GlueCE	GlueCEStateTotalJobs
CEVOs	GlueCE	GlueCEAccessControlBaseRule
TotalCPUs	GlueCE	GlueCEInfoTotalCPUs
FreeCPUs	GlueCE	GlueCEStateFreeCPUs
CE	GlueCE	GlueCEUniqueID
WaitingJobs	GlueCE	GlueCEStateWaitingJobs
RunningJobs	GlueCE	GlueCEStateRunningJobs
CloseCE	GlueCESEBindGroup	GlueCESEBindGroupCEUniqueID
CloseSE	GlueCESEBindGroup	GlueCESEBindGroupSEUniqueID
SEVOs	GlueSA	GlueSAAccessControlBaseRule
UsedSpace	GlueSA	GlueSAStateUsedSpace
AvailableSpace	GlueSA	GlueSAStateAvailableSpace
Type	GlueSE	GlueSEType
SE	GlueSE	GlueSEUniqueID
Protocol	GlueSEAccessProtocol	GlueSEAccessProtocolType
ArchType	GlueSL	GlueSLArchitectureType
Processor	GlueSubCluster	GlueHostProcessorModel
OS	GlueSubCluster	GlueHostOperatingSystemName
Cluster	GlueSubCluster	GlueSubClusterUniqueID
Tag	GlueSubCluster	GlueHostApplicationSoftwareRunTimeEnvironment
Memory	GlueSubCluster	GlueHostMainMemoryRAMSize

lcg-info examples

List all the CE(s) in the BDII satisfying given conditions

$ lcg-info --list-ce --query 'TotalCPUs>=30,OS=SL*' --attrs 'RunningJobs,FreeCPUs'

```
- CE: grid010.ct.infn.it:2119/jobmanager-lcgpbs-long
  - RunningJobs        0
  - FreeCPUs          33

- CE: grid010.ct.infn.it:2119/jobmanager-lcgpbs-short
  - RunningJobs        0
  - FreeCPUs          33

- CE: grid010.ct.infn.it:2119/jobmanager-lcgpbs-infinite
  - RunningJobs        1
  - FreeCPUs          33

- CE: skurut1.cesnet.cz:2119/jobmanager-lcgpbs-long
  - RunningJobs        0
  - FreeCPUs          26

- CE: skurut1.cesnet.cz:2119/jobmanager-lcgpbs-gilda
  - RunningJobs        0
  - FreeCPUs          26
[..]
```

lcg-info examples (cont.)

List all the CE(s) which satisfying the condition FreeCPU >=30

```
$ lcg-info --list-ce --query 'FreeCPUs >= 30‘--attrs ‘FreeCPUs’

- CE: grid010.ct.infn.it:2119/jobmanager-lcgpbs-long
  - FreeCPUs          33

- CE: grid010.ct.infn.it:2119/jobmanager-lcgpbs-short
  - FreeCPUs          33

- CE: grid010.ct.infn.it:2119/jobmanager-lcgpbs-infinite
  - FreeCPUs          33

[..]
```

Print all the tags published by a specific query

$ lcg-info --list-ce --query 'CE=*grid010.ct.infn.it:2119*' --attrs 'Tag'

PBS
INFN
CATANIA
LCG-2
LCG-2_1_0
LCG-2_1_1
LCG-2_2_0
LCG-2_3_0
LCG-2_3_1
LCG-2_4_0
R-GMA
AFS
CMS-1.1.0
ATLAS-6.0.4
GATE-1.0.0-3
LHCb-1.1.1
IDL-5.4
CMSIM-125
ALICE-4.01.00
ALIEN-1.32.14
POVRAY-3.5
DEMTOOLS-1.0
CMKIN-VALID
CMKIN-1.1.0
CMSIM-VALID
CSOUND-4.13
MPICH
VIRGO-1.0
CMS-OSCAR-2.4.5
LHCb_dbase_common-v3r1
GEANT4-6
VLC-0.7.2
EGEODE-1.0
RASTER3D
SCILAB-2.6
G95-3.5.0
MAGIC-6.19
CODESA3D-1.0
VO-gilda-slc3_ia32_gcc323
VO-gilda-CMKIN_5_1_1
VO-gilda-GEANT
VO-gilda-GKS05

List the CEs with a particular SW

$ lcg-info –vo gilda --list-ce --query 'Tag=*MPICH*' --attrs 'CE'

- CE: cn01.be.itu.edu.tr:2119/jobmanager-lcglsf-long
- CE cn01.be.itu.edu.tr:2119/jobmanager-lcglsf-long

- CE: cn01.be.itu.edu.tr:2119/jobmanager-lcglsf-short
- CE cn01.be.itu.edu.tr:2119/jobmanager-lcglsf-short

- CE: grid010.ct.infn.it:2119/jobmanager-lcgpbs-long
- CE grid010.ct.infn.it:2119/jobmanager-lcgpbs-long

- CE: grid011f.cnaf.infn.it:2119/jobmanager-lcgpbs-long
- CE grid011f.cnaf.infn.it:2119/jobmanager-lcgpbs-long

- CE: ced-ce0.datagrid.cnr.it:2119/jobmanager-lcgpbs-long
- CE ced-ce0.datagrid.cnr.it:2119/jobmanager-lcgpbs-long

[..]

List the SEs satisfying given query

$ lcg-info –vo gilda --list-se --query 'AvailableSpace>=100000' --attrs 'CloseCE'

- SE: cn02.be.itu.edu.tr
- CloseCE cn01.be.itu.edu.tr:2119/jobmanager-lcglsf-long
cn01.be.itu.edu.tr:2119/jobmanager-lcglsf-short
cn01.be.itu.edu.tr:2119/jobmanager-lcglsf-infinite

- SE: grid009.ct.infn.it
- CloseCE grid010.ct.infn.it:2119/jobmanager-lcgpbs-long
grid010.ct.infn.it:2119/jobmanager-lcgpbs-short
grid010.ct.infn.it:2119/jobmanager-lcgpbs-infinite

- SE: ced-se0.datagrid.cnr.it
- CloseCE ced-ce0.datagrid.cnr.it:2119/jobmanager-lcgpbs-long
ced-ce0.datagrid.cnr.it:2119/jobmanager-lcgpbs-short
ced-ce0.datagrid.cnr.it:2119/jobmanager-lcgpbs-infinite

- SE: grid003.cecablc.ula.ve
- CloseCE grid006.cecalc.ula.ve:2119/jobmanager-lcgpbs-cert
grid006.cecalc.ula.ve:2119/jobmanager-lcgpbs-long
grid006.cecalc.ula.ve:2119/jobmanager-lcgpbs-short
grid006.cecalc.ula.ve:2119/jobmanager-lcgpbs-infinite

[..]

R-GMA

Introduction to R-GMA

- **Relational Grid Monitoring Architecture (R-GMA)**
  - Developed as part of the EuropeanDataGrid Project (EDG)
  - Now as part of the EGEE project.
  - Based the Grid Monitoring Architecture (GMA)

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

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

SELECT * FROM people WHERE group='HR'

Service orientation

- 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.

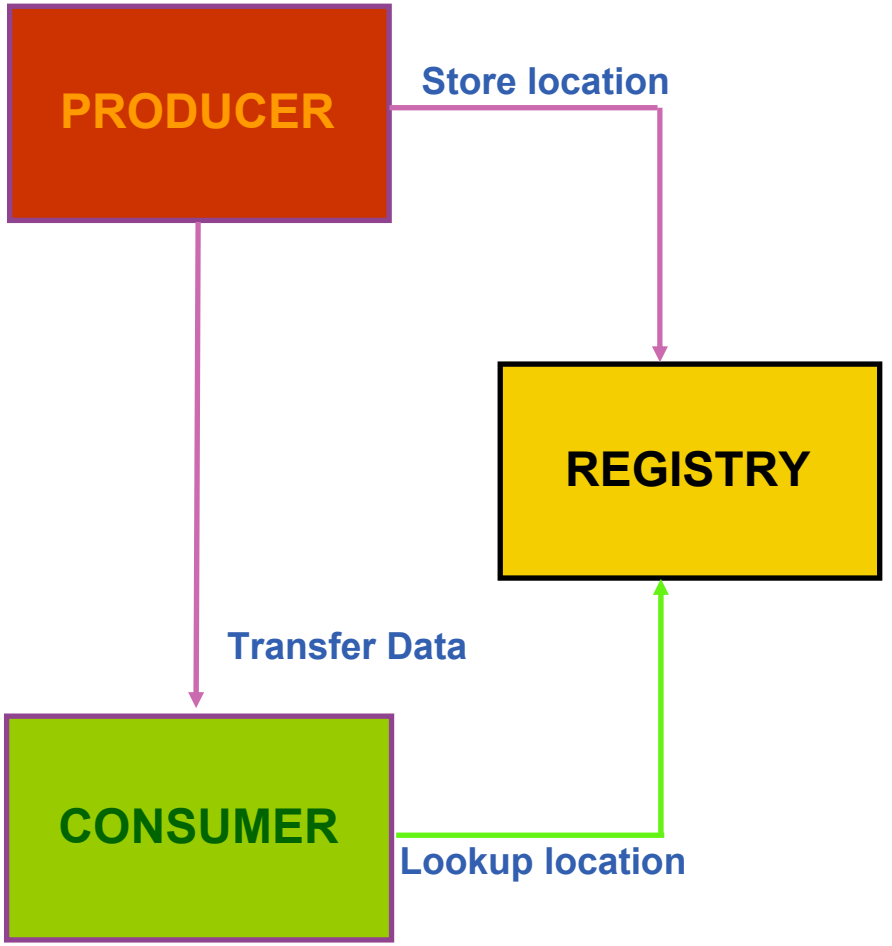

Multiple Producers

- 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.

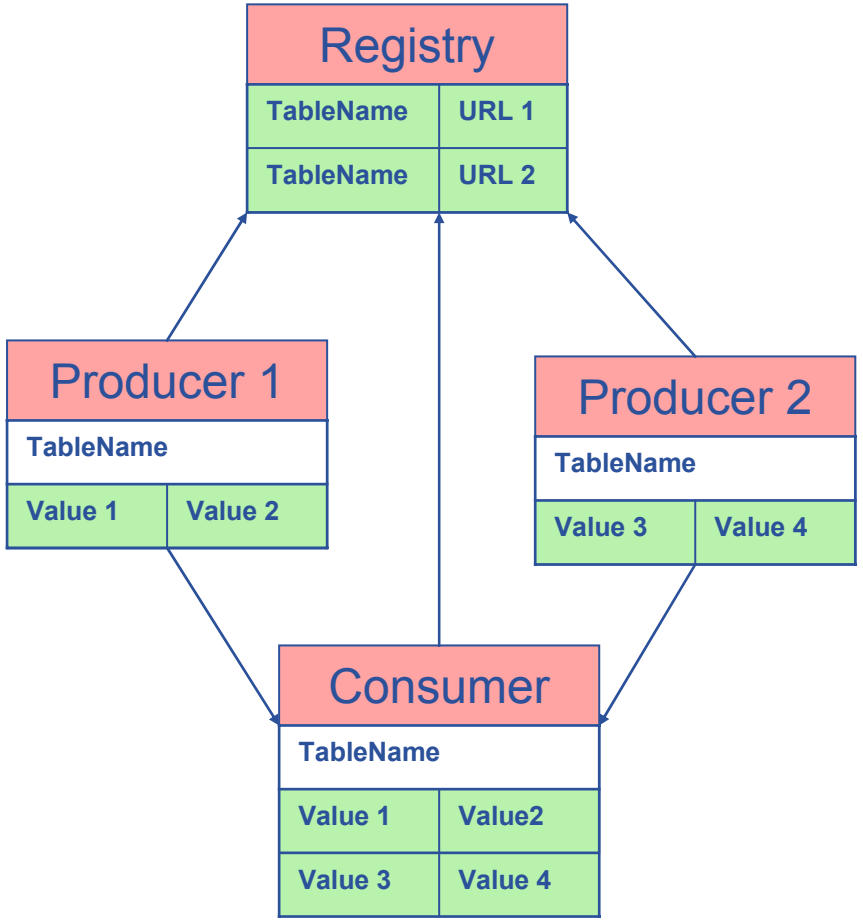

Select * from CPULoad

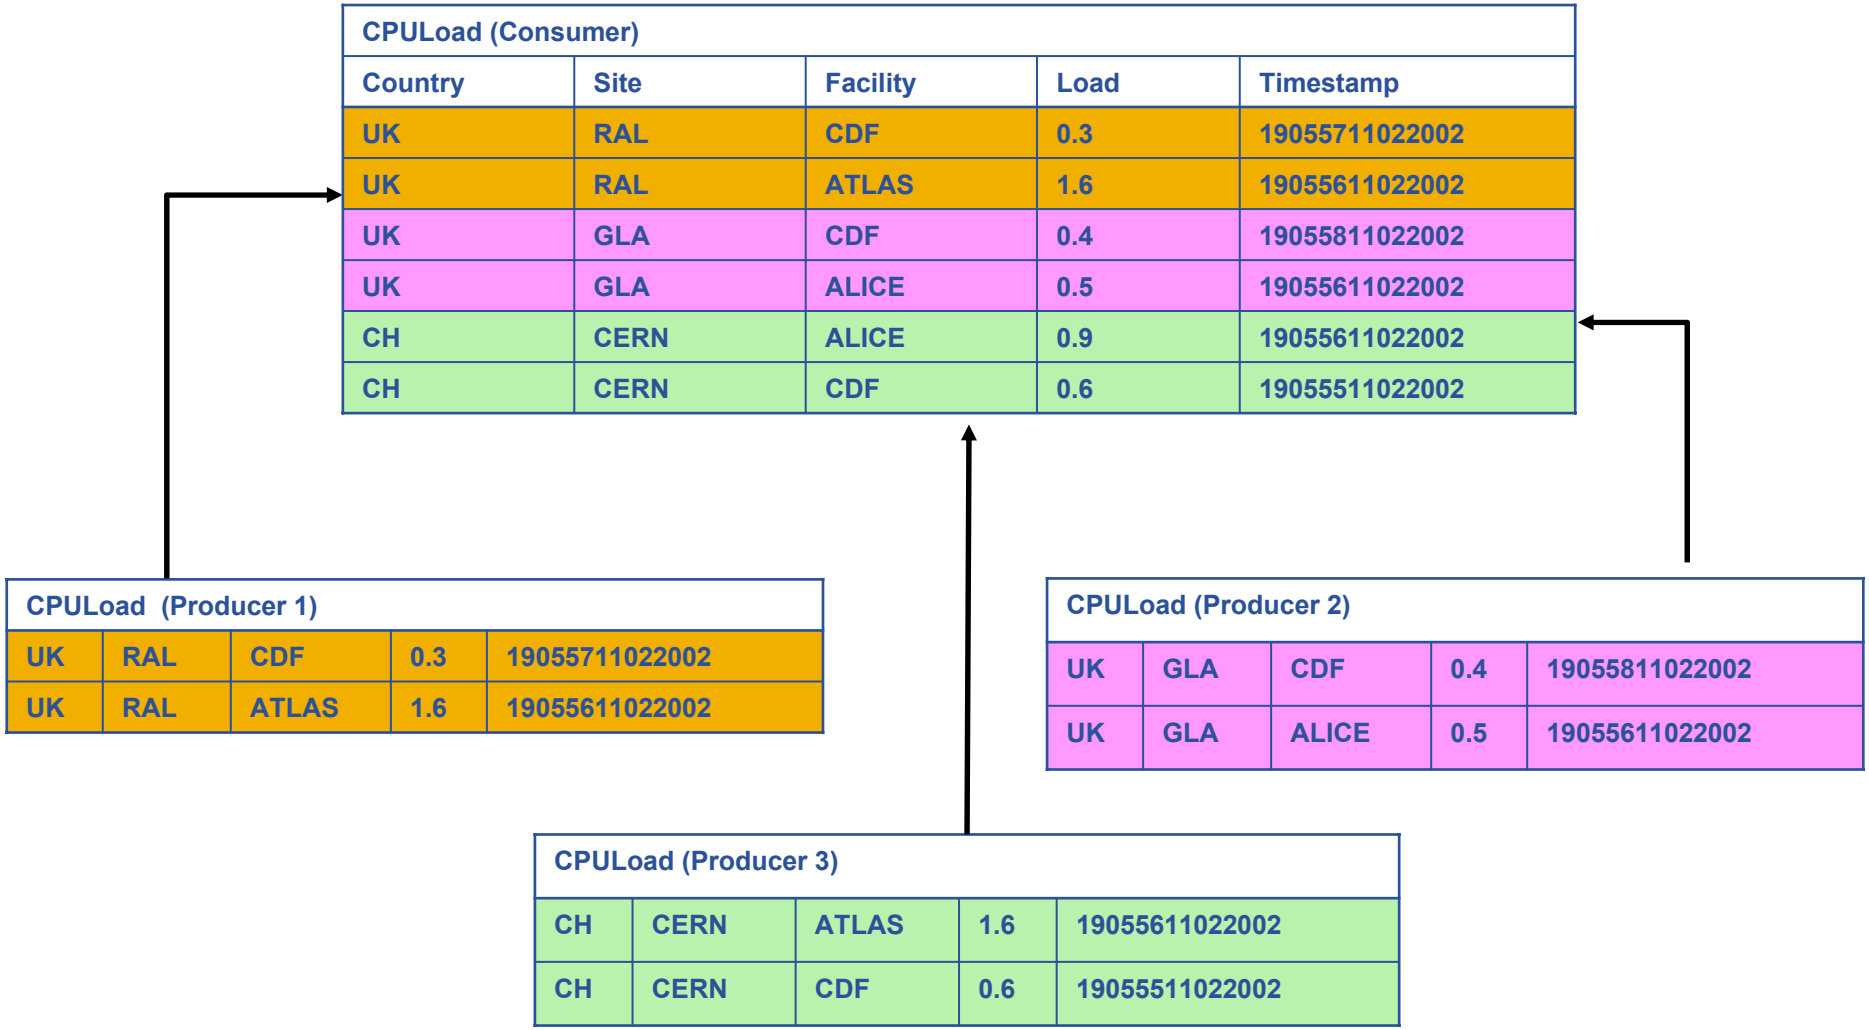

CPULoad (Consumer)

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

Joins

Service				
URI	VO	type	emailContact	site
gppse01	alice	SE	sysad@rl.ac.uk	RAL
gppse01	atlas	SE	sysad@rl.ac.uk	RAL
gppse02	cms	SE	sysad@rl.ac.uk	RAL
lxshare0404	alice	SE	sysad@cern.ch	CERN
lxshare0404	atlas	SE	sysad@cern.ch	CERN

ServiceStatus				
URI	VO	type	up	status
gppse01	alice	SE	y	SE is running
gppse01	atlas	SE	y	SE is running
gppse02	cms	SE	n	SE ERROR 101
lxshare0404	alice	SE	y	SE is running
lxshare0404	atlas	SE	y	SE is running

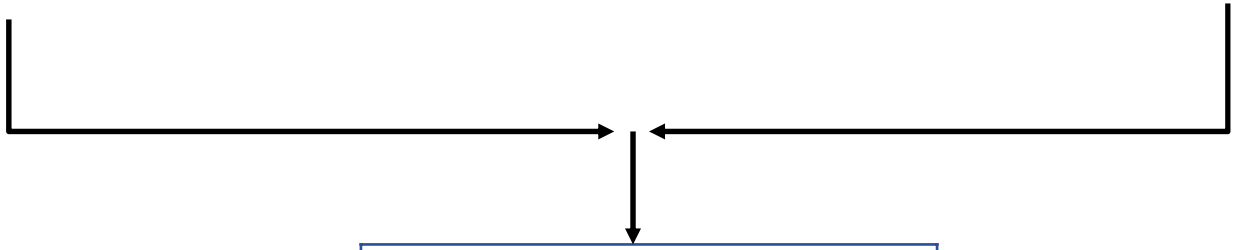

Result Set (Consumer)	
URI	emailContact
gppse02	sysad@rl.ac.uk

SELECT Service.URI Service.emailContact FROM Service S, ServiceStatus SS
WHERE (S.URI= SS.URI and SS.up='n')

Roles

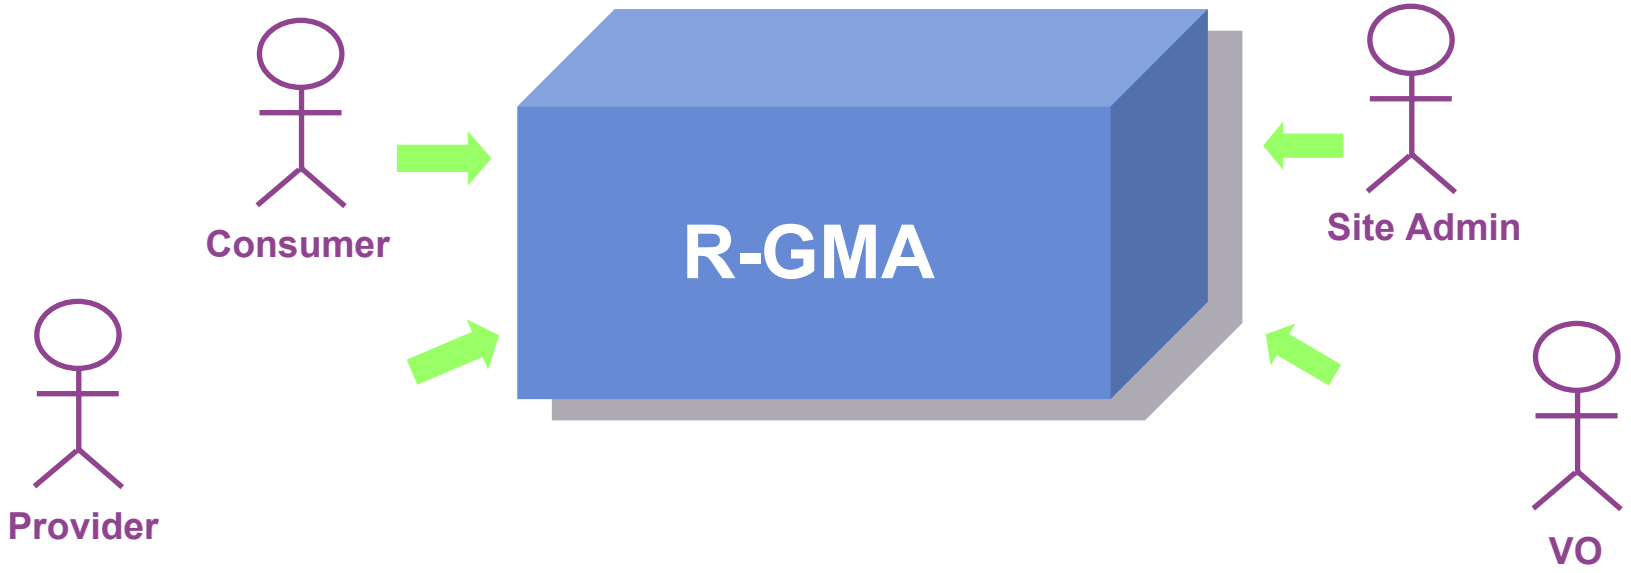

- Consumer users: who requests information.
- Producer users: who provides information.
- Site administrators: who runs R-GMA services.
- Virtual Organizations: who “own” the schema and registry.

Security

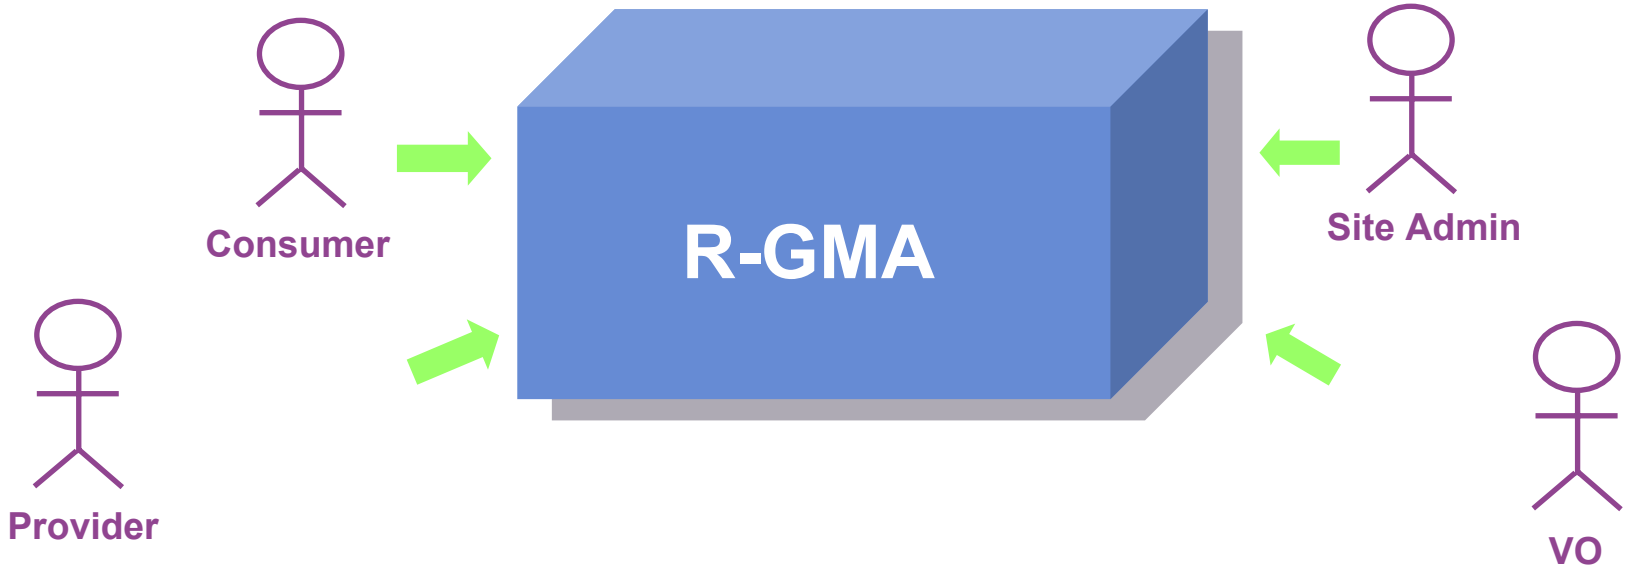

- Mutual Autentication: guaranteeing who is at each end of an exchange of messages.
- Encryption: using an encrypted transport protocol (HTTPS).
- Authorization: implicit or explicit.

R-GMA

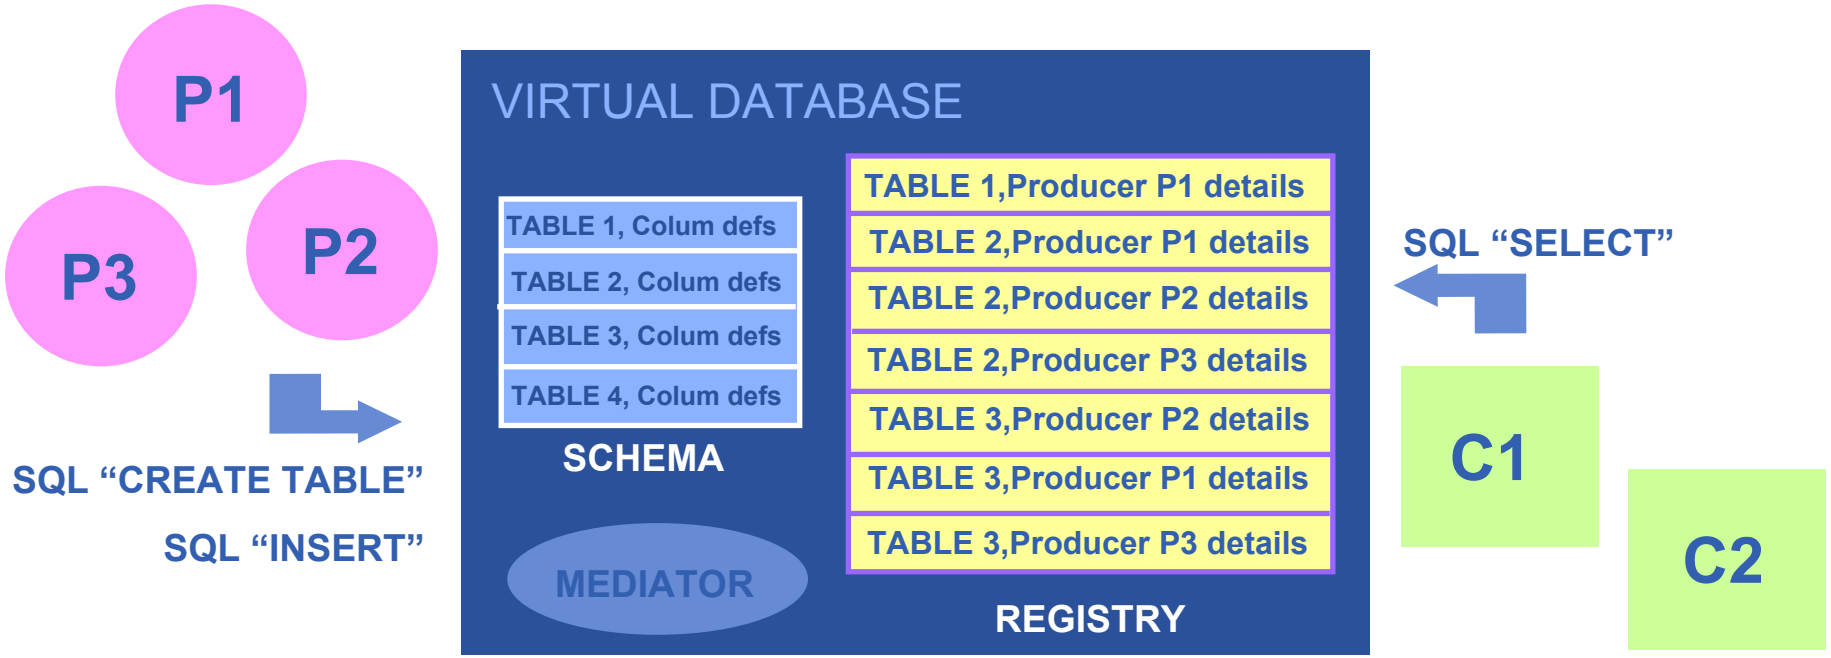

There is no central repository!!! There is only a "*Virtual Database*".

Schema is a list of table definitions: additional tables/schema can be defined by applications

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

Producers publish data.

Consumer read data published.

Deployment

- **Producer and Consumer Services are typically on a one per site basis**
- **Centralized Registry and Schema.**
- **The Registry and Schema may be replicated, to avoid a single point of failure**
  - … when you use RGMA CLI you will see which are being used

Producer Types

- **Primary Producer**

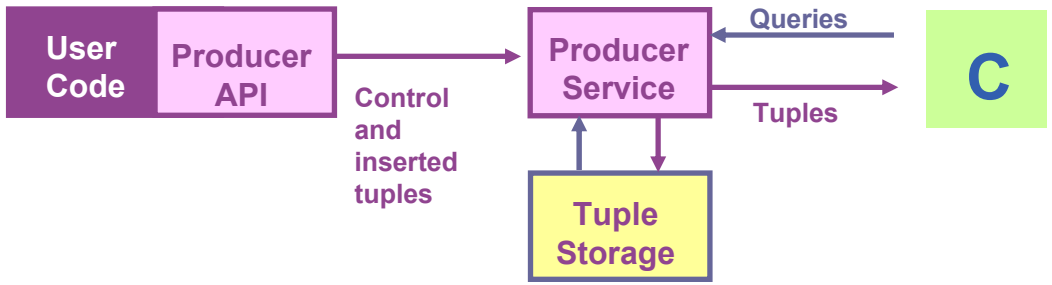

- **Secondary Producer**

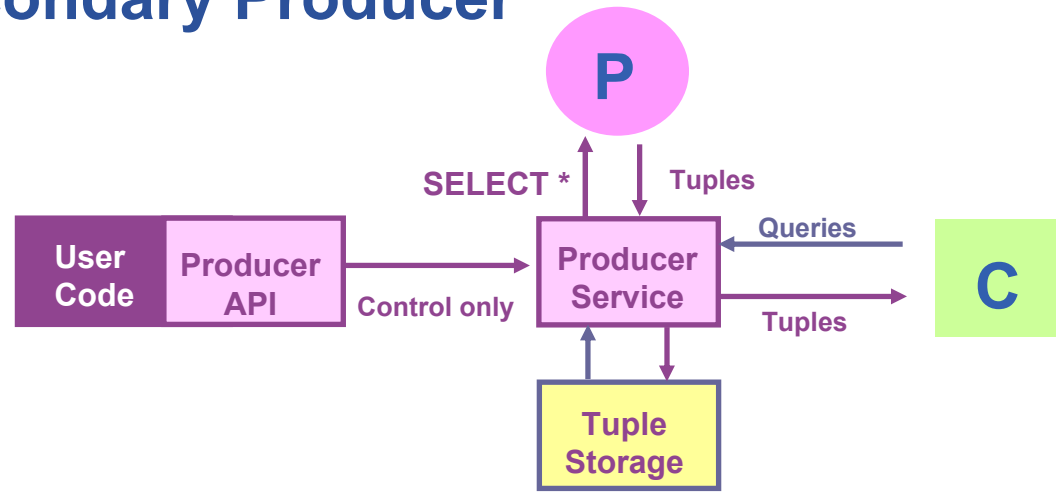

- **On-Demand Producer**
  - No internal storage
  - Queries passed to user code

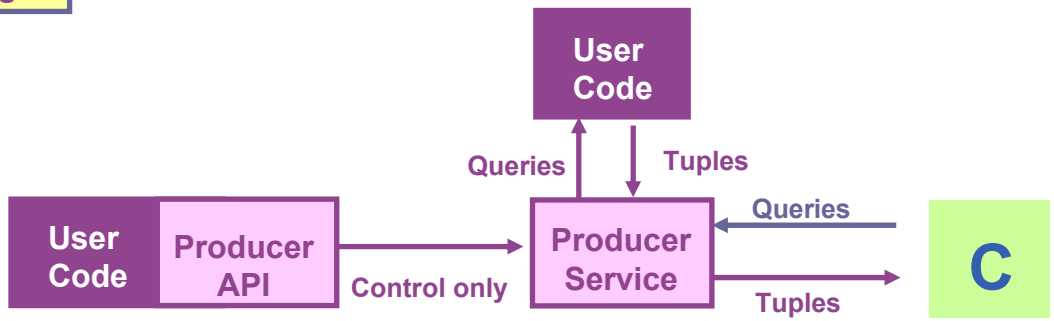

Query Types

Continuous

- **Latest**
- **History**
- **Static**

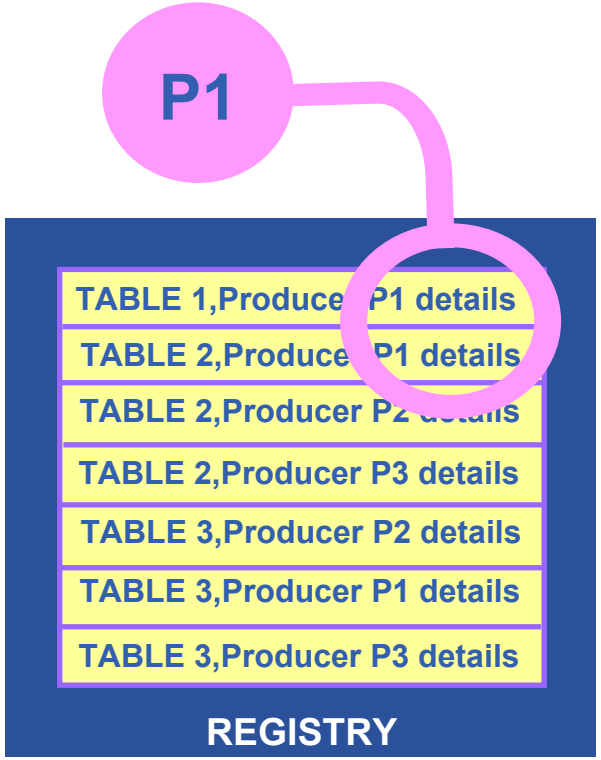

Continuous

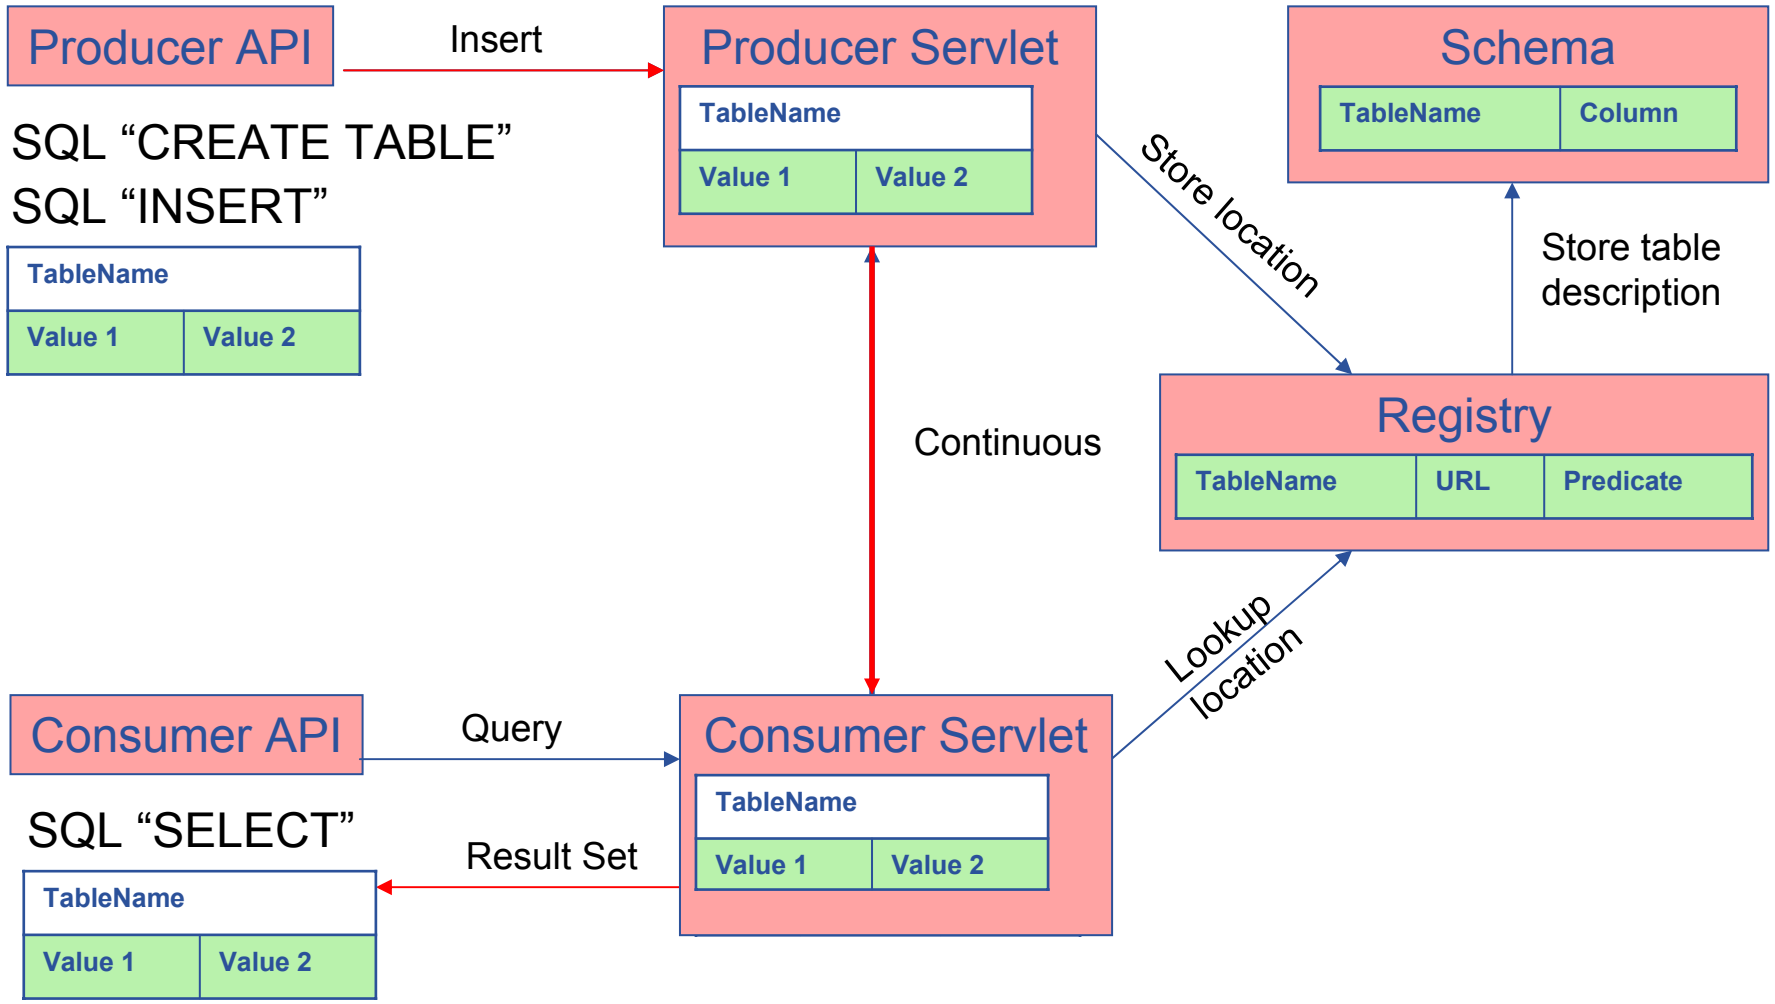
Producer API
Insert
Producer Servlet
TableName
Value 1
Value 2
Schema
TableName
Column
SQL "CREATE TABLE"
SQL "INSERT"
TableName
Value 1
Value 2
Store location
Store table description
Continuous
Registry
TableName
URL
Predicate
Lookup location
Consumer API
Query
Consumer Servlet
TableName
Value 1
Value 2
SQL "SELECT"
Result Set
TableName
Value 1
Value 2

Query Types

- Continuous
- Latest
- History
- Static

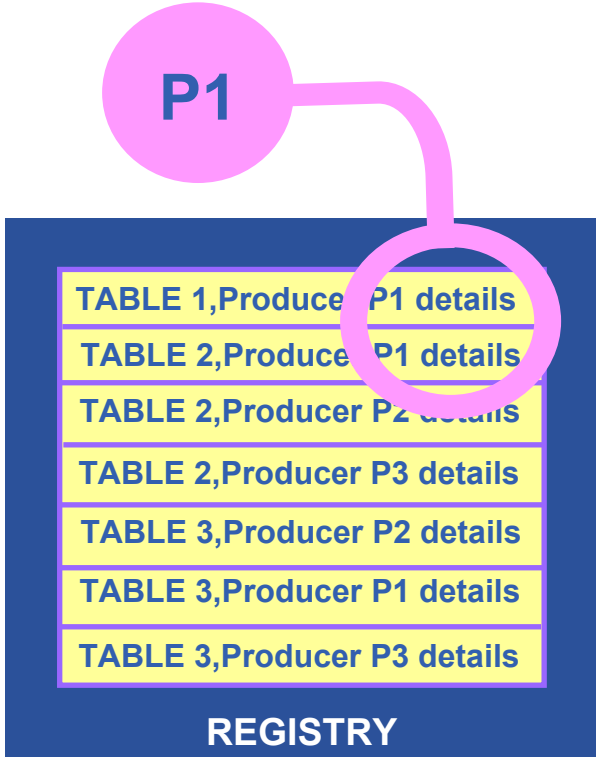

History or Latest

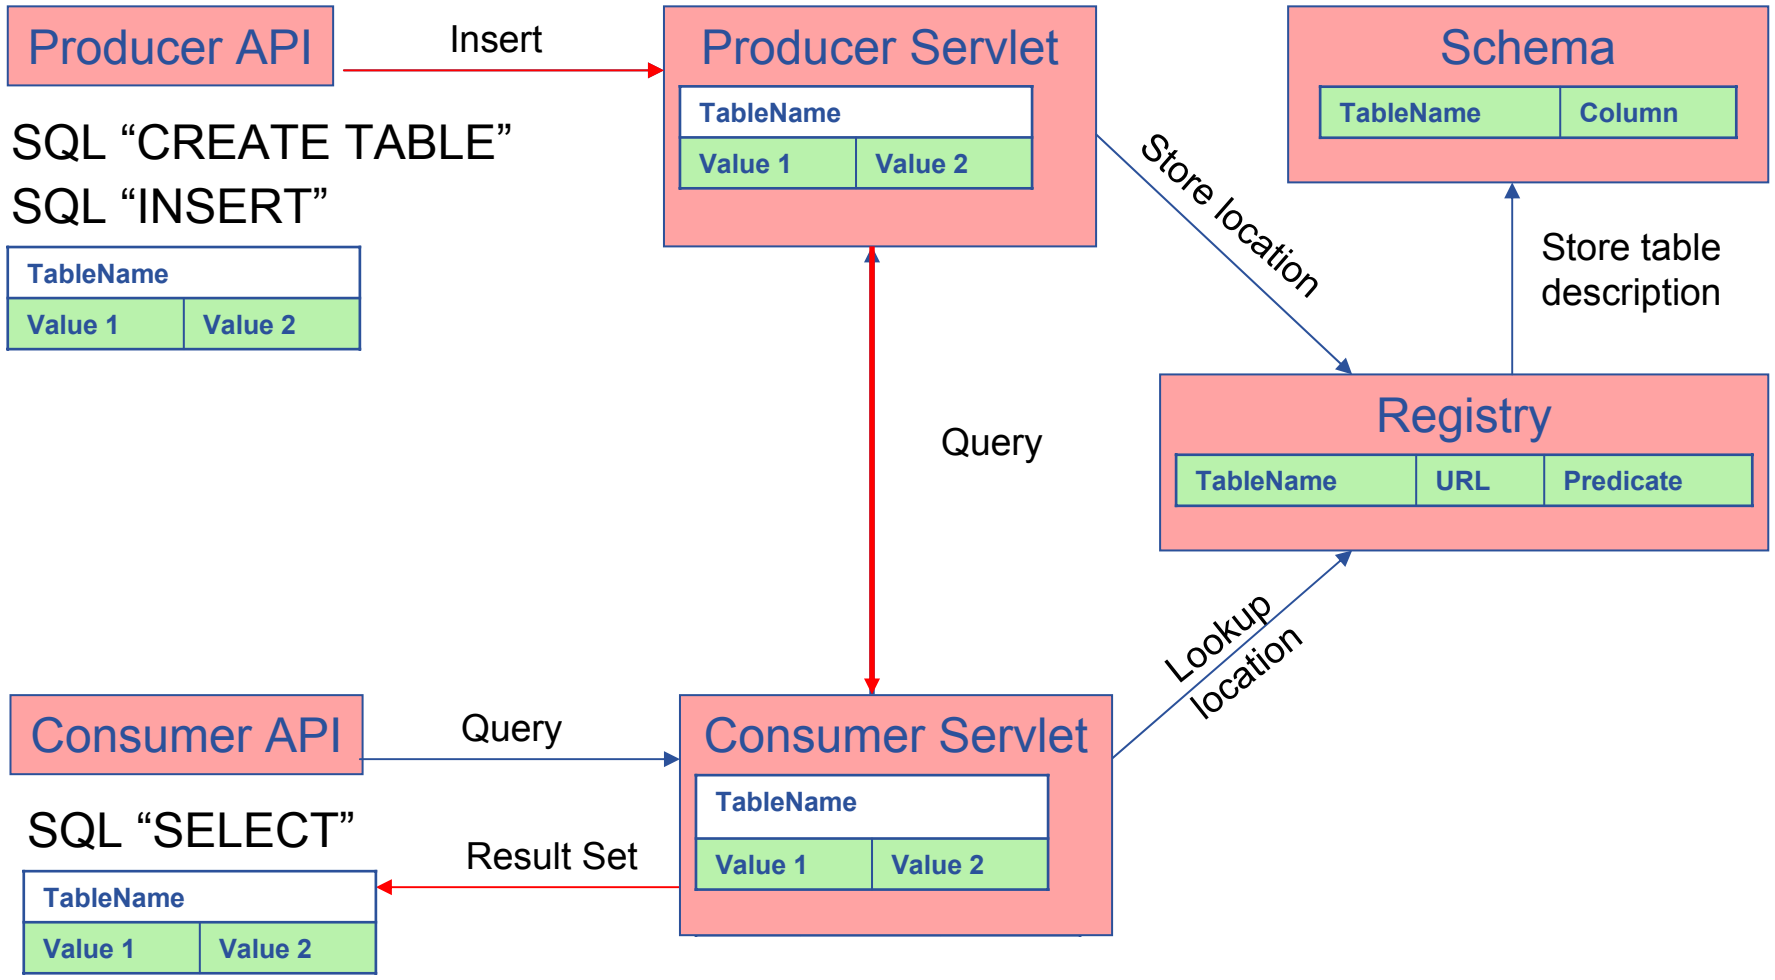

Query Types

- Continuous
- Latest
- History
- Static

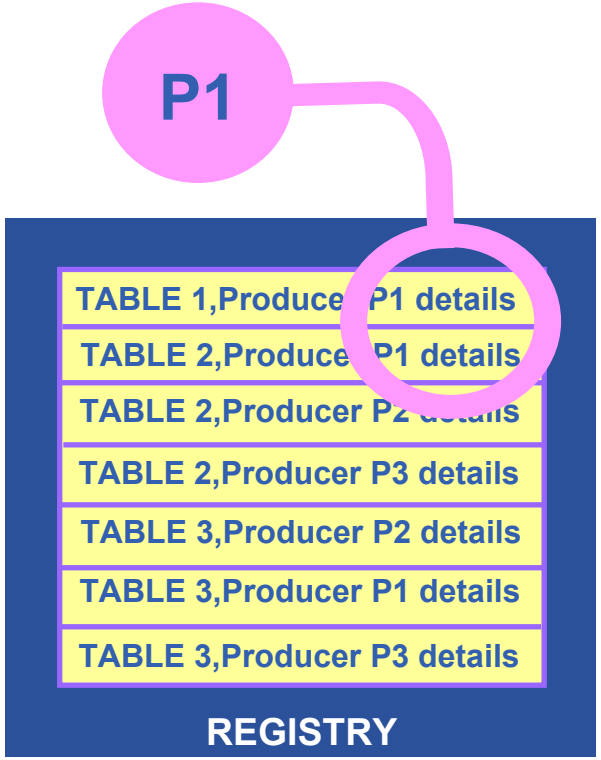

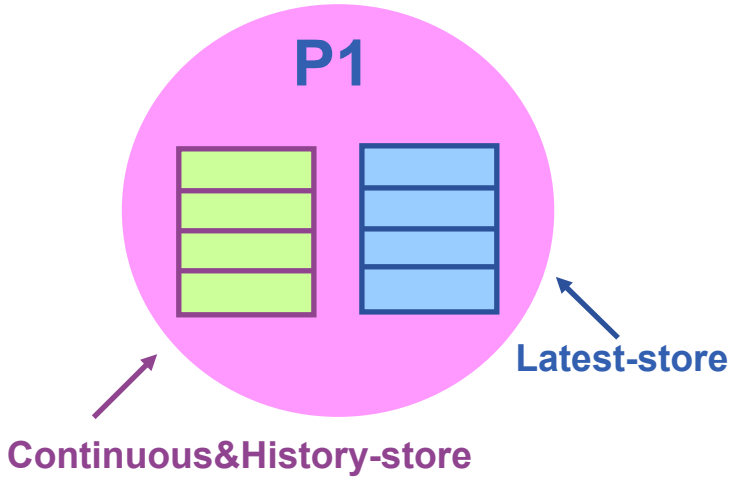

Latest Retention Period

History Retention Period

R-GMA APIs

- **APIs exist in Java, C, C++, Python.**
  - For clients (servlets contacted behind the scenes)
- **They include methods for…**
  - Creating consumers
  - Creating primary and secondary producers
  - Setting type of queries, type of produces, retention periods, time outs…
  - Retrieving tuples, inserting data
  - …
- **You can create your own Producer or Consumer.**

More information

- **R-GMA overview page.**
  - http://www.r-gma.org/
- **R-GMA in EGEE**
  - http://hepunx.rl.ac.uk/egee/jra1-uk/
- **R-GMA Documentation**
  - http://hepunx.rl.ac.uk/egee/jra1-uk/LCG/doc/

R-GMA practical

R-GMA Command Line Tool (1)

- **CHECK YOU HAVE A VOMS PROXY CERTIFICATE**
- To Start the R-GMA command line tool run the following command:

>rgma

- On startup you should receive the following message:

```
Welcome to the R-GMA virtual database for Virtual Organisations.
You are connected to the R-GMA registry service at

http://<registry-host>:8080/R-GMA/RegistryServlet

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

Entering Command

- Commands are entered by typing at the rgma> prompt and hitting ‘enter’ to execute the command.
- A history of the commands executed can be accessed using the Up and Down arrow keys.
- To search a command from history use CTRL-R and type the first few letters of the command to recall.
- Command autocompletion is supported (use Tab when you have partly entered a command).

General Commands

General Commands

- **help**

  Display general help information.

- **help <command>**

  Display help for a specific command.

- **exit or quit**

  Exit from R-GMA command line interface.

- **Show tables**

  Display the name of all tables existing in the Schema

- **Describe <tablename>**

  Show all information about the structure of a table

Querying Data (1)

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

  rgma> SELECT * FROM GlueService

  The behaviour of SELECT varies according to the type of query being executed. In R-GMA there are three basic types of query:
- LATEST Queries only the most recent tuple for each primary key
- HISTORY Queries all historical tuples for each primary key
- CONTINUOUS Queries returns tuples continuously as they are inserted.

Querying Data (2)

- The type of query can be changed using the SET QUERY command as follow:

```
rgma> SET QUERY LATEST
```

or

```
rgma> SET QUERY CONTINUOUS
```

- The current query type can be displayed using

```
rgma> SHOW QUERY
```

Exercises

1. Display all the table of the Schema

   rgma>show tables

2. Display information about GlueSite table

   rgma>describe GlueSite

3. Basic select query on the table named GlueSite

   rgma>set query latest

   rgma>show query

   rgma> select Name,Latitude,Longitude from GlueSite

Maximum AGE of tuples

- The maximum age of tuples to return can also be controlled. To limit the age of latest or historical tuples use the SET MAXAGE command. The following are equivalent:

  rgma> SET MAXAGE 2 minutes

  rgma> SET MAXAGE 120

- The current maximum tuple age can be displayed using rgma> SHOW MAXAGE
- To disable the maximum age, set it to none:

  rgma> SET MAXAGE none

Query Timeout

- The final property affecting queries is timeout.
  - For a latest or history query the timeout exists to prevent a problem (e.g. network failure) from stopping the query from completing.
  - For a continuous query, timeout indicates how long the query will continue to return new tuples. Default timeout is 1 minute and it can be changed using

rgma>SET TIMEOUT 3 minutes or SET TIMEOUT 180

- The current timeout can be displayed using

rgma>SHOW TIMEOUT

Producer & Inserting Data

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

  rgma> INSERT INTO userTable VALUES ('a', 'b', 'c', 'd')

- **In R-GMA, data is inserted into the system using a Producer component which handles the INSERT statement.**
- **Using the command line tool you may work with one producer at a time.**
- The current producer type can be displayed using:

  rgma>show producer
- The producer type can be set using:

  rgma>set producer latest

Exercise

Choose a role for the exercise as consumer or as producer (alernate if you wish)

PRODUCERS

rgma> set producer continuous

rgma> set maxage 3 minutes

rgma> insert into userTable values('brisbanexx','any string',1.4,66)

CONSUMERS

rgma> set query continuous OR set query history

rgma> set timeout 5 seconds

rgma> select * from userTable

THE END