

Information and Monitoring

The European DataGrid Project Team

<http://www.eu-datagrid.org>





Contents


- ◆ Grid Information Systems
- ◆ GMA and R-GMA
- ◆ Topologies of components
- ◆ APIs and Tools
- ◆ Info-Providers
- ◆ Deployment
- ◆ GLUE Schema

Features of a grid information system



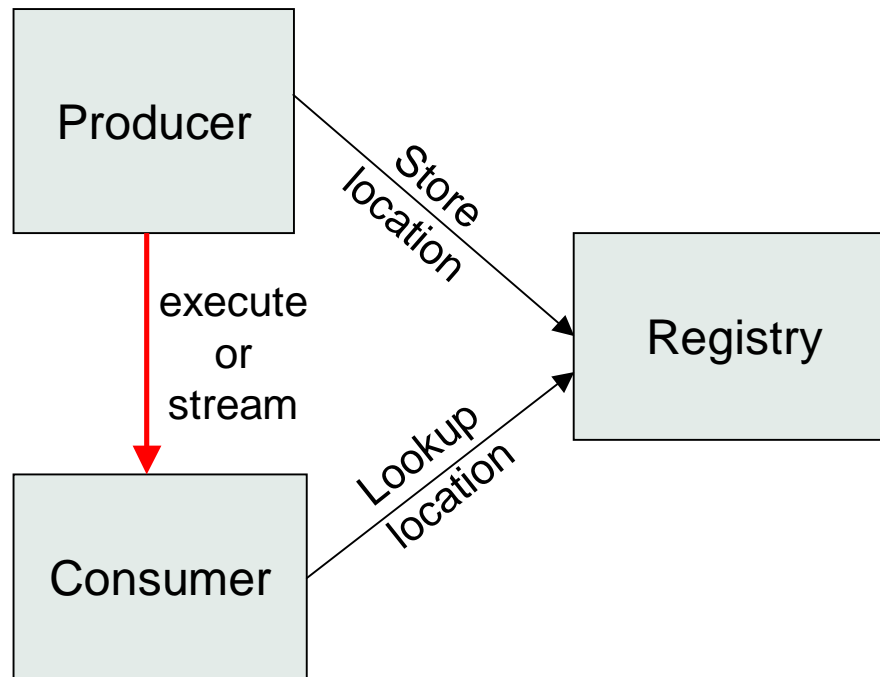
- ◆ Provides information on both:
 - The Grid itself
 - Mainly for the middleware packages
 - The user may query it to understand the status of the Grid
 - Grid applications
 - For users
- ◆ Flexible infrastructure
 - Able to cope with nodes in a distributed environment with an unreliable network
 - Dynamic addition and deletion of information producers
 - Security system able to address the access to information at a fine level of granularity
 - Allow new data types to be defined
 - Scalable
 - Good performance
 - Standards based

GMA



- ◆ From GGF
- ◆ Very simple model
- ◆ Does not define:
 - Data model
 - Data transfer mechanism
 - Registry implementation

R-GMA



- ◆ Use the GMA from GGF
- ◆ A relational implementation
 - Powerful data model and query language
 - All data modelled as tables
 - SQL can express most queries in one expression
- ◆ Applied to both information and monitoring
- ◆ **Creates impression that you have one RDBMS per VO**

Relational Data Model in R-GMA



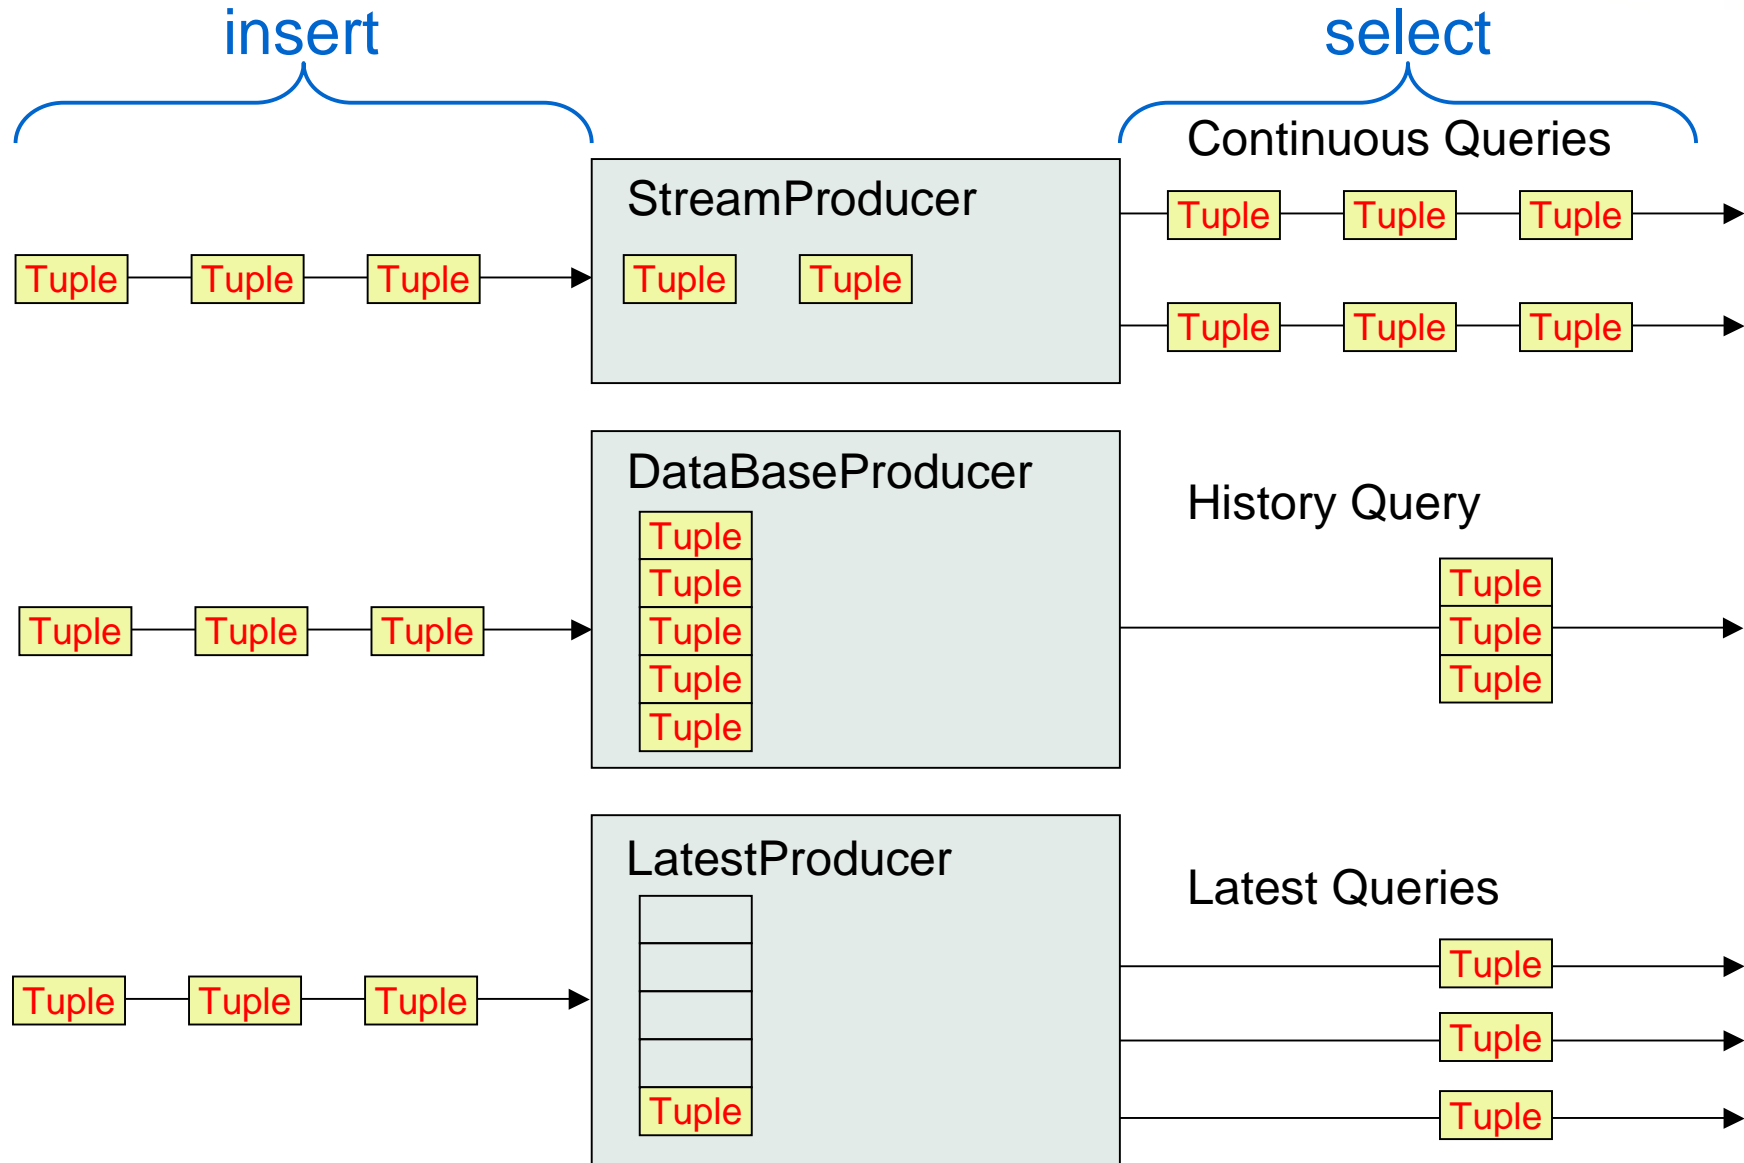
- ◆ **Not** a general distributed RDBMS system, but a way to use the relational model in a distributed environment **where global consistency is not important**
- ◆ **Producers** announce: SQL "CREATE TABLE"
publish: SQL "INSERT"
- ◆ **Consumers** collect: SQL "SELECT"
- ◆ Some producers, the Registry and Schema make use of RDBMS as appropriate – but what is central is the relational **model**



Data Transfer: Producer ➔ Consumer

- ◆ Consumer can issue one-off queries
 - Similar to normal database query
- ◆ Consumer can also start a continuous query
 - Requests all data published which matches the query
 - As data matching the query is produced it is streamed to the Consumer
 - Can be seen as an alert mechanism

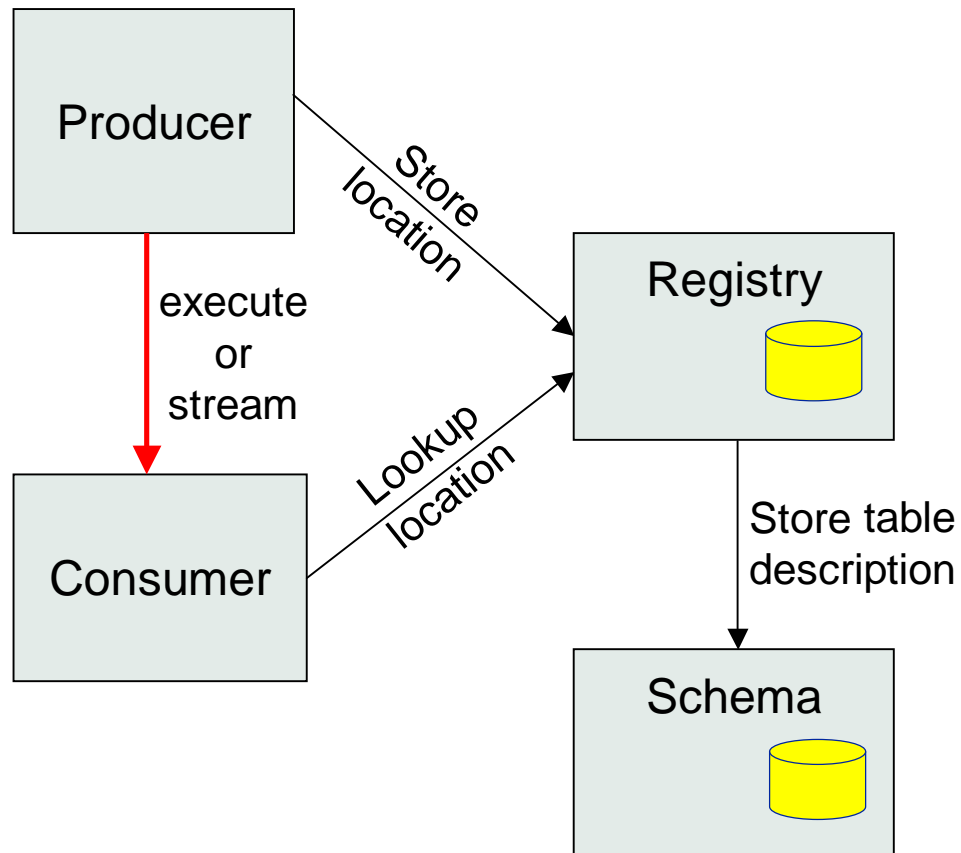
3 Kinds of Query



Producers

- ◆ StreamProducer – Supports **Continuous** Queries
 - In memory data structure
 - Can define minimum retention period
- ◆ ResilientStreamProducer – Supports **Continuous** Queries
 - Like the StreamProducer but won't lose data if system crashes
 - So slightly slower
- ◆ DataBaseProducer – Supports **History** Queries
 - Information not lost
 - Supports joins
 - Clean up strategy
- ◆ LatestProducer – Supports **Latest** Queries
 - Just holds the latest information for any "primaryish" key
 - Supports joins
- ◆ CanonicalProducer – Supports anything
 - Offers "anything" as relations
 - User has to write code to handle SQL etc.

Registry and Schema



- ◆ Registry has two main tables:
 - Producer
 - Table name
 - Predicate
 - Location
 - Consumer
 - Query
 - Location
- ◆ Schema holds description of tables
 - Column names and types of each table
- ◆ Registry predicate defines subset of "global" table

Contributions to the "global" table



CPUload (Global Schema)				
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

WHERE
country = 'UK'
AND site =
'RAL'

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

WHERE
country = 'CH'
AND site =
'CERN'

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



Mediator

- ◆ Queries posed against a virtual data base
- ◆ The Mediator must:
 - find the right Producers
 - combine information from them
- ◆ Hidden component – but vital to R-GMA
- ◆ Will eventually support full distributed queries but for now will only merge information from multiple producers for queries on one table or over multiple tables from one producer

Queries over "global" table – merging streams

```
SELECT * from CPUload 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

Mediator handles merging information from multiple producers for queries on one table

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

Queries over "global" table – joining tables

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

Service/ServiceStatus (Consumer)	
URI	emailContact
gppse02	sysad@rl.ac.uk


Service/ServiceStatus (Latest Producer)													
Service													
URI	VO	type	emailContact	site	secure	majorVersion	minorVersion	patchVersion					
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



Archiver (Re-publisher)

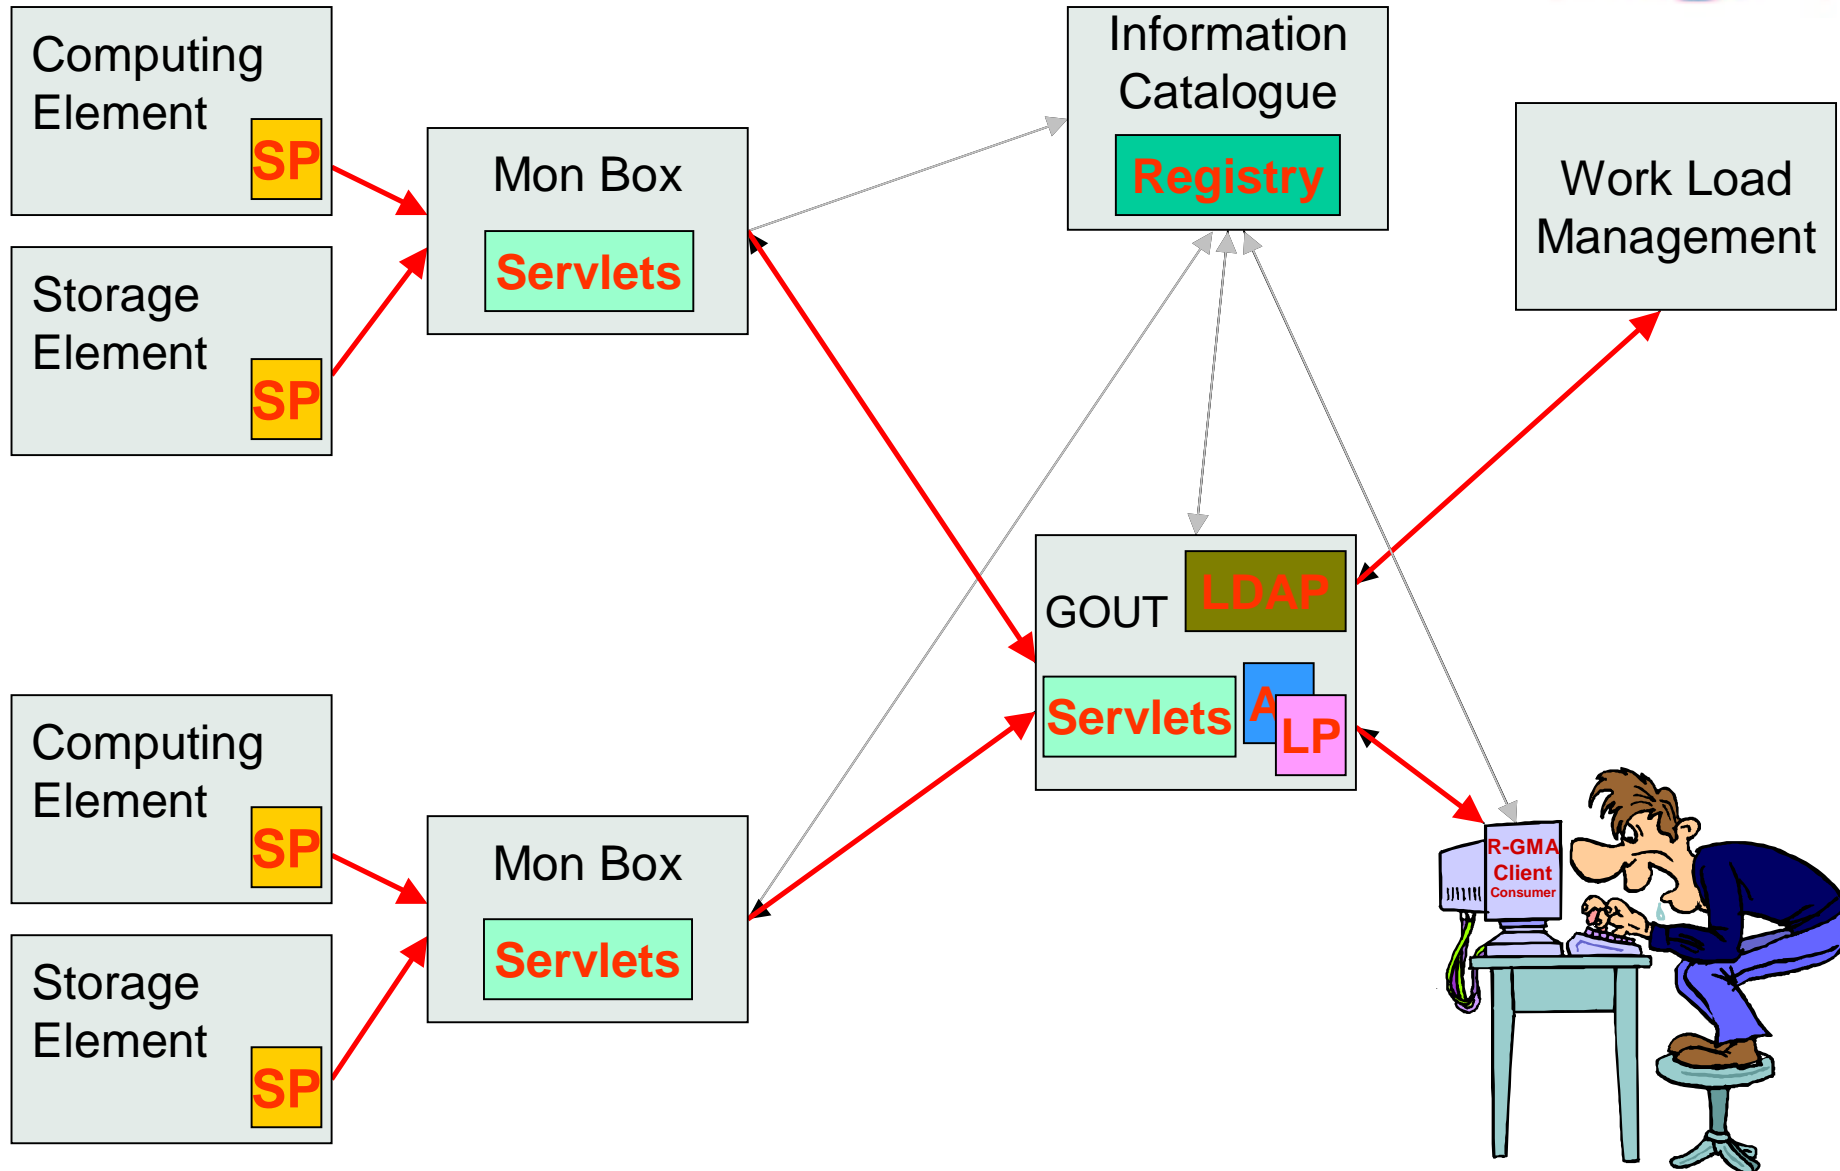
- ◆ It is a combined Consumer-Producer
 - Follows the GMA concept but packaged for ease of use
- ◆ You just have to tell it what to collect and it does so on your behalf
- ◆ Re-publishes to **any** kind of "Insertable" (i.e. not to the CanonicalProducer)
 - Can support joins if archiving to a DataBaseProducer or a LatestProducer

Topologies



- ◆ Normally publish via a StreamProducer **SP**
- ◆ Archivers **A** instantiated with a Producer and a Predicate.
 - May re-publish via:
 - StreamProducer
 - LatestProducer **LP**
 - DataBaseProducer **DBP**
- ◆ Must avoid cycles in the connections – i.e. must be a DAG.

Topology for EDG



APIs

- ◆ Exist in Java, C++, C, Python and Perl
- ◆ C, Python and Perl follow an object based style reflecting the Java and C++ APIs

Java

```
myProducer = new StreamProducer();
```

C++

```
myProducer= new edg::info::StreamProducer();
```

C

```
myProducer = StreamProducer_new();
```

Perl

```
$myProducer = rgmainfo::StreamProducer_new();
```

Python

```
myProducer = rgmainfo.StreamProducer_new()
```

R-GMA Tools

- ◆ R-GMA Browser
 - Application dynamically generating web pages
 - Supports pre-defined and user-defined queries
- ◆ R-GMA CLI (edg-rgma)
 - Command Line Interface (similar to MySQL)
 - Supports single query and interactive modes
 - Can perform simple operations with Consumers, Producers and Archivers
- ◆ Pulse
 - R-GMA Java client-based GUI
 - Supports streaming and simple graphical displays

R-GMA Browser

Home
Predefined Queries
Service Status
Site Info
Table Sets



- All tables
[EDG Info Providers](#)
[Network Monitoring](#)
[CMS](#)

- EDG Info Providers**
- [GlueCE](#)
[GlueCEAccessControlBaseRu](#)
[GlueCESEBind](#)
[GlueCluster](#)
[GlueHostRemoteFileSystem](#)
[GlueSA](#)
[GlueSAAccessControlBaseRu](#)
[GlueSE](#)
[GlueSEAccessProtocol](#)
[GlueSEAccessProtocolSupport](#)
[GlueSL](#)
[GlueSubCluster](#)
[GlueSubClusterSoftwareRunT](#)
[SiteInfo](#)

SELECT
 Name
 GlueClusterUniqueID
 TotalCPUs
 LRMSType

FROM **GlueCE**

WHERE

Query

Type of query:
 History Latest Continuous Cont.+Old
 Queries wait for seconds

Use Mediator
 Select Producers you want to query:
 There are no available History producers for table GlueCE

Latest Producer

producerServlet:http://lxshare0382.cern.ch:8080/R-GMA/LatestProducerServlet ConnectionId:2

Continuous Producer

producerServlet:http://lxshare0382.cern.ch:8080/R-GMA/StreamProducerServlet ConnectionId:365

producerServlet:http://gprrg06.gridpp.rl.ac.uk:8080/R-GMA/StreamProducerServlet ConnectionId:5

producerServlet:http://testbed012.cnaf.infn.it:8080/R-GMA/StreamProducerServlet ConnectionId:3

Query



edg-rgma

- ◆ show tables
- ◆ describe Service
- ◆ show producers of Service
- ◆ latest select * from Service
- ◆ old continuous select * from Service

edg-rgma – Example

```
./edg-rgma \
  -c "timeout 0.1" \
  -c timeout \
  -c "str decl Service" \
  -c "str minr .2" \
  -c "str minr" \
  -c "stream INSERT into Service (URI, VO, type, secure,
    emailContact, site, majorVersion, minorVersion,
    patchVersion) values ('a','b','c','y', 'd','e',1,2,3)" \
  -c "old continuous select * from Service
```

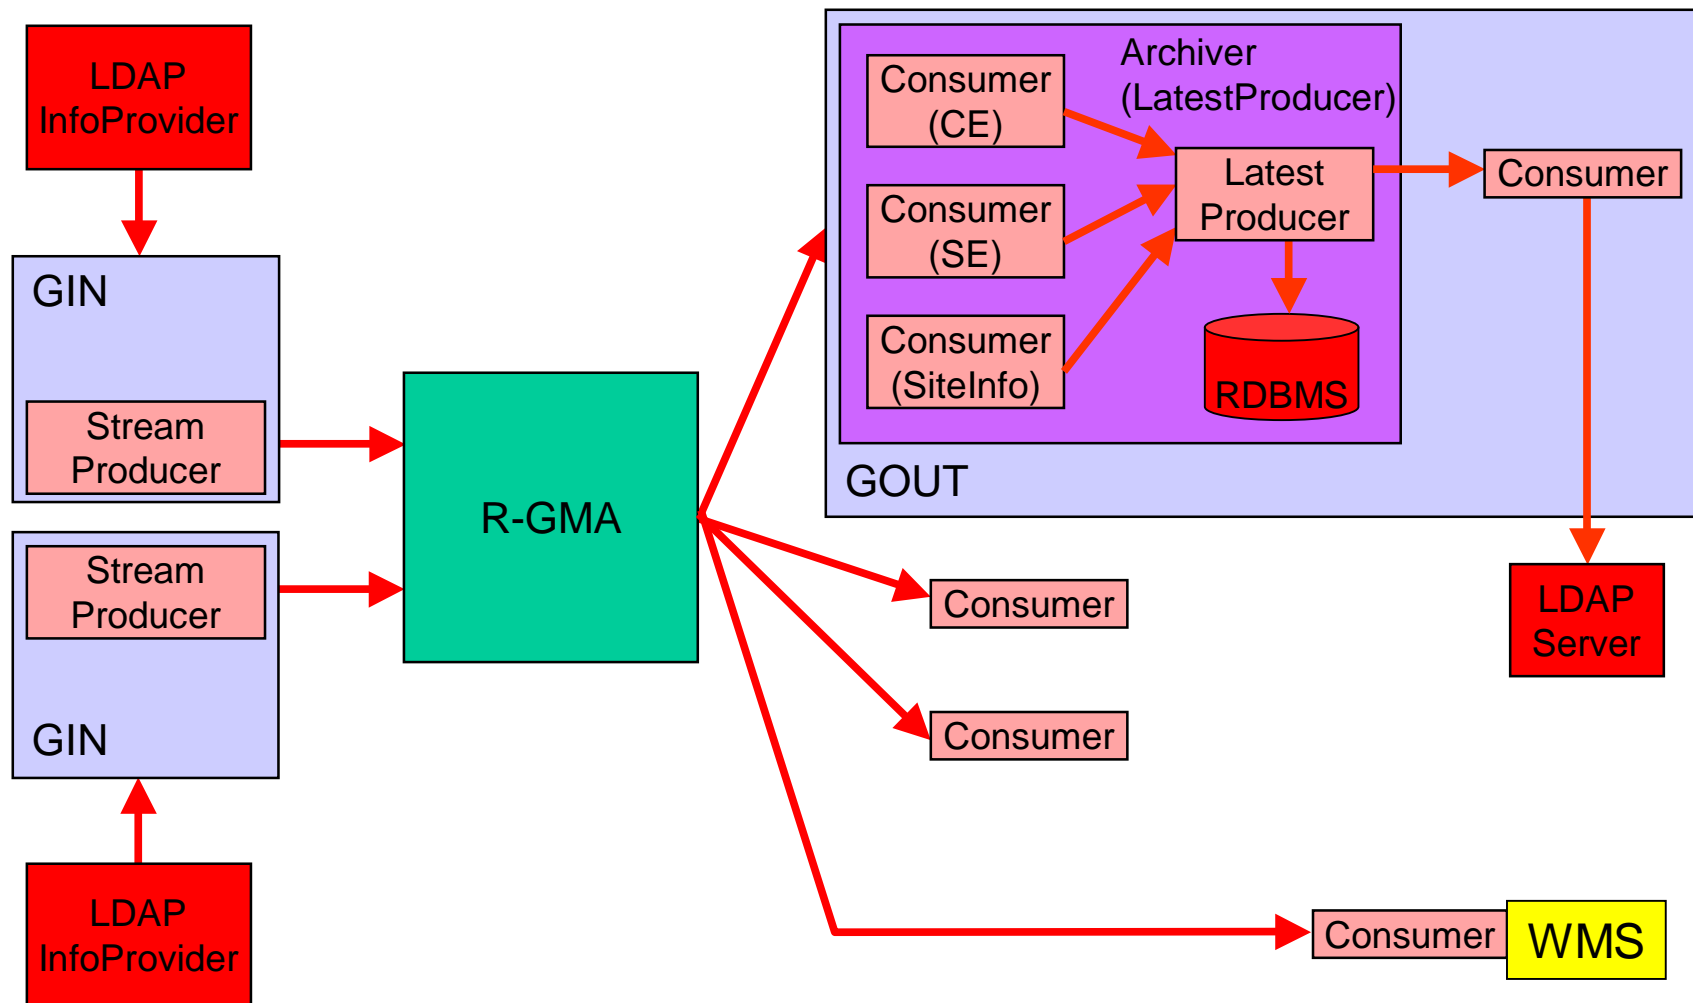
```
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
-+-----+
| URI | VO | type | emailContact | site | secure | majorVersion | minorVersion | patchVersion | MeasurementDate
| MeasurementTime |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
-+-----+
| a | b | c | d | e | y | 1 | 2 | 3 | 2003-07-08
| 10:26:58 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
-+-----+
1 Rows in set
```



EDG Information Providers

- ◆ EDG information providers
 - Software that provides information about resources and infrastructure
 - Provided by the work packages responsible for the resource
- ◆ The information providers produce data in LDIF format
 - This is a legacy from when Globus MDS was the primary information system
- ◆ R-GMA publishes the data
 - Gin (gadget in) is used to invoke the information provider scripts and publish via StreamProducers
 - Gout (gadget out) republishes the data via a LatestProducer and then to an OpenLDAP database
 - This is to provide backwards compatibility during the transition from MDS to R-GMA

Current deployment: GIN and GOUT





The GLUE schema


- ◆ Developed within HEP community
 - DataGrid
 - DataTAG
 - Globus
- ◆ It is a compromise - so not perfect for DataGrid
- ◆ Currently defines CEs and SEs
- ◆ Coming soon...
 - Network monitoring
 - Services

R-GMA Computing Service Tables



GlueCluster
UniqueID
Name
InformationServiceURL

R-GMA Computing Service Tables




GlueSubCluster
UniqueID
Name
GlueClusterUniqueID
RAMSize
RAMAvailable
VirtualSize
VirtualAvailable
PlatformType
SMPSize
OSName
OSRelease
OSVersion
Vendor
.....

GlueCluster
UniqueID
Name
InformationServiceURL

a sub cluster is a group of homogeneous machines

R-GMA Computing Service Tables




GlueSubCluster
UniqueID
Name
GlueClusterUniqueID
RAMSize
RAMAvailable
VirtualSize
VirtualAvailable
PlatformType
SMPSize
OSName
OSRelease
OSVersion
Vendor
.....

GlueCluster
UniqueID
Name
InformationServiceU

GlueSubCluster
SoftwareRunTime
Environment
GlueSubClusterUniqueID
Value

software accessible by the nodes, e.g. ALICE-3.09.06, ALIEN-1.29.9

R-GMA Computing Service Tables




GlueCluster
UniqueID
Name
InformationServiceURL

GlueCE
UniqueID
Name
GlueClusterUniqueID
TotalCPUs
LRMSType
LRMSVersion
GRAMVersion
HostName
GatekeeperPort
RunningJobs
WaitingJobs
TotalJobs
Status
.....

a computing element relates to a queue on batch system

R-GMA Computing Service Tables




authorised user

GlueCluster
UniqueID
Name
InformationServiceURL

GlueCEAccess ControlBaseRule
GlueCEUniqueID
Value

GlueCE
UniqueID
Name
GlueClusterUniqueID
TotalCPUs
LRMSType
LRMSVersion
GRAMVersion
HostName
GatekeeperPort
RunningJobs
WaitingJobs
TotalJobs
Status
.....



R-GMA Storage Service Tables



GlueSL
UniqueID
GlueSEUniqueID
Name
ArchitectureType
MaxIOCapacity
InformationServiceURL


storage library, the machine providing the storage service

R-GMA Storage Service Tables



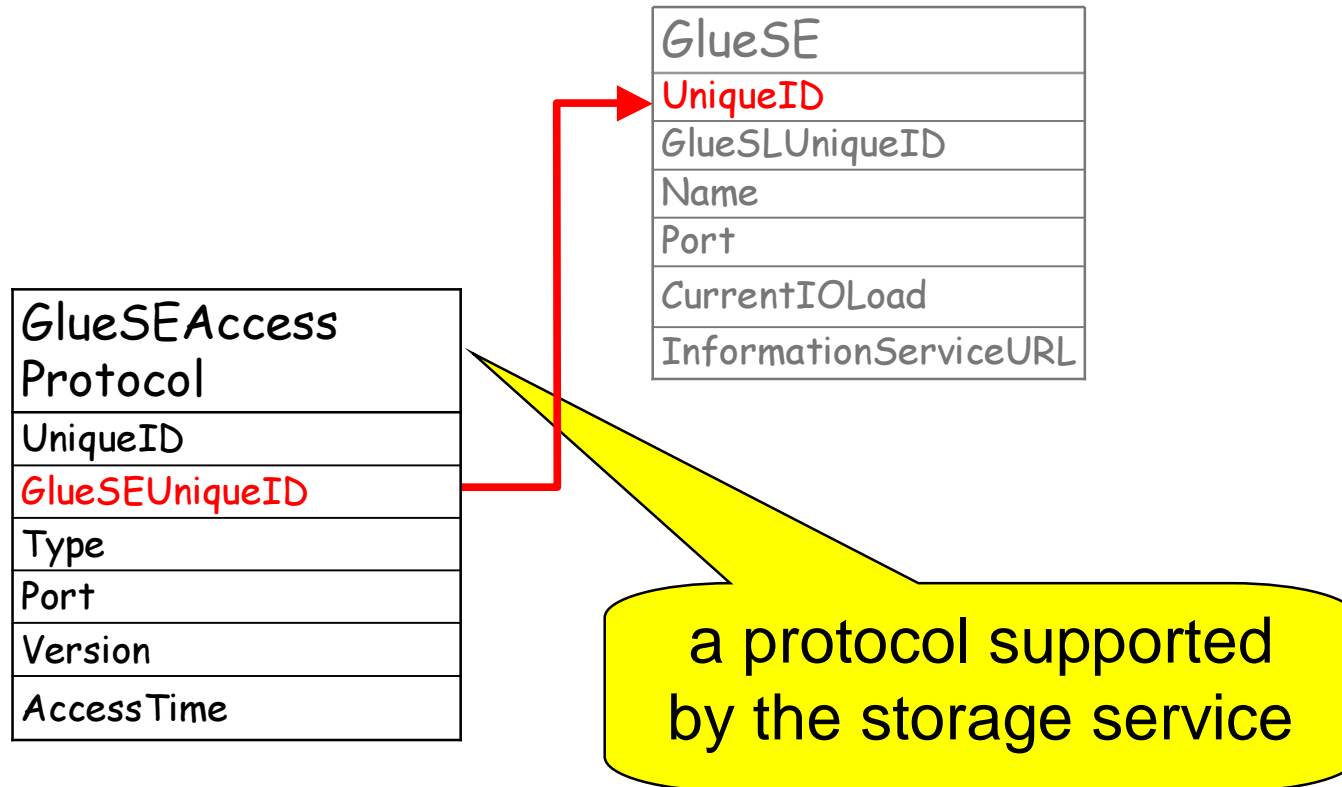
GlueSL
UniqueID
GlueSEUniqueID
Name
ArchitectureType
MaxIOCapacity
InformationServiceURL

GlueSE
UniqueID
GlueSLUniqueID
Name
Port
CurrentIOLoad
InformationServiceURL

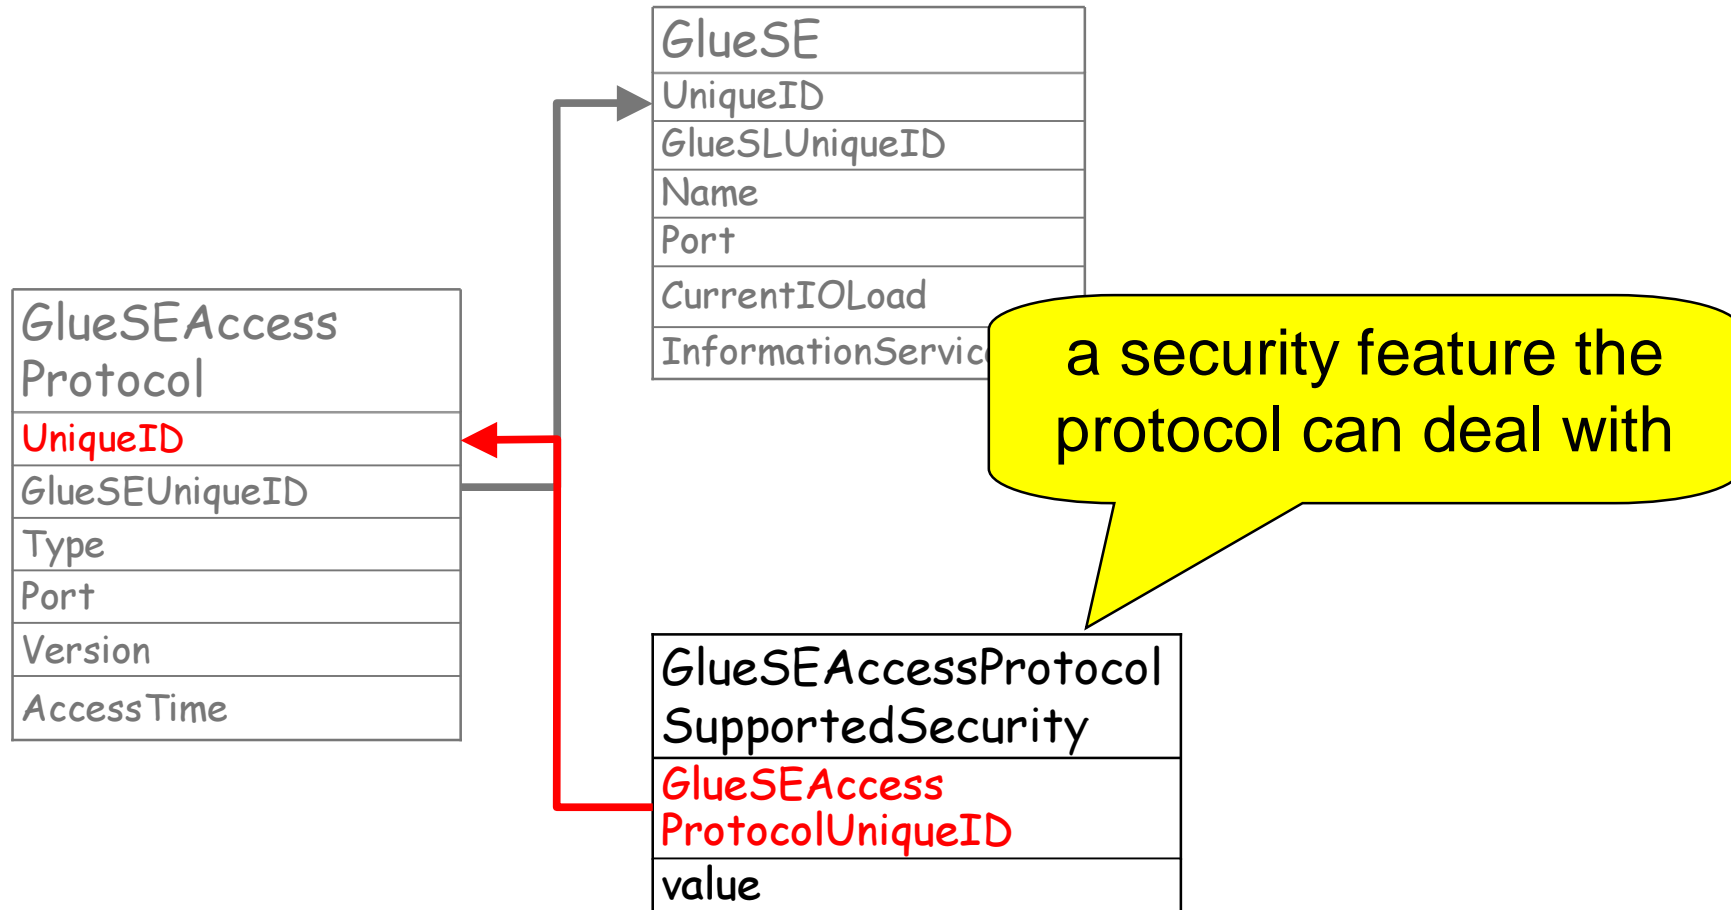


storage service

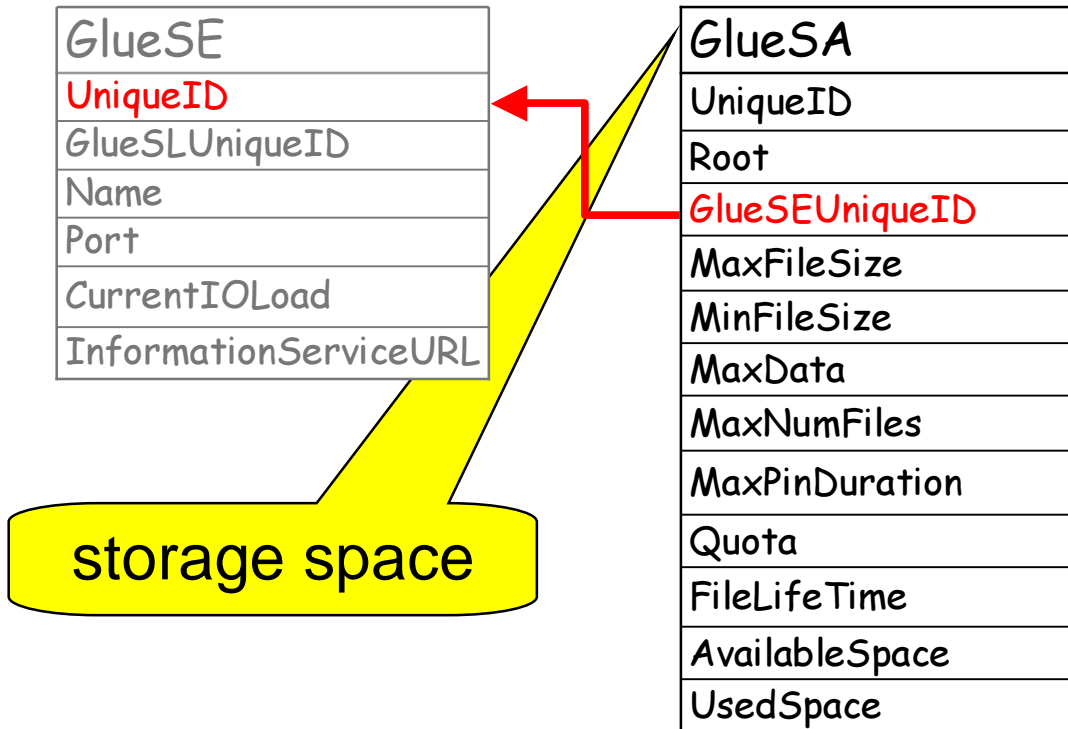
R-GMA Storage Service Tables



R-GMA Storage Service Tables



R-GMA Storage Service Tables



R-GMA Storage Service Tables




GlueSE
UniqueID
GlueSLUniqueID
Name
Port
CurrentIOLoad
InformationServiceURL

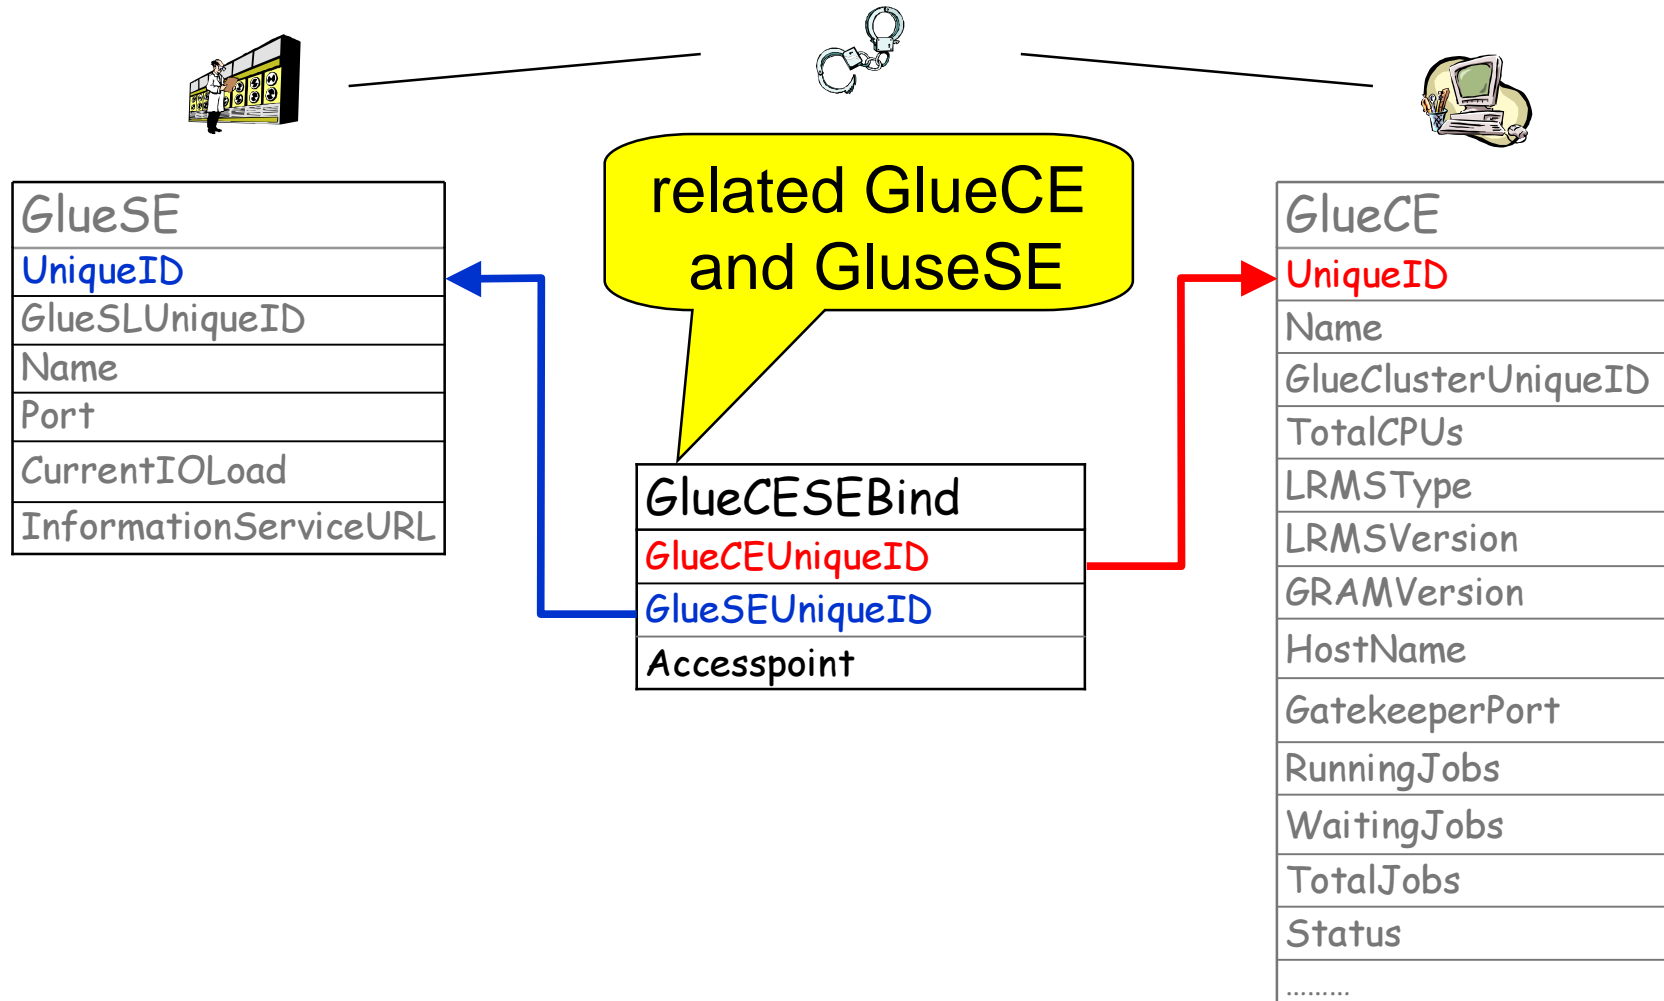
GlueSA
UniqueID
Root
GlueSEUniqueID
MaxFileSize
MinFileSize
MaxData
MaxNumFiles
MaxPinDuration
Quota
FileLifeTime
AvailableSpace
UsedSpace

authorised user

GlueSAAccess ControlBaseRule
GlueSAUniqueID
Value



R-GMA Binding Table



Summary

- ◆ R-GMA
 - is a relational implementation of the GGF's GMA
 - has different Producer types
 - has an API available in multiple languages
 - components can be deployed in various topologies
 - mediator creates the impression of a single RDBMS
- ◆ The current EDG deployment with GIN and GOUT is backwards compatible with MDS2
- ◆ The GLUE Schema gives a data model independent view of information



Further Information

- ◆ Information and Monitoring Services
 - <http://hepunx.rl.ac.uk/edg/wp3/>