

# Information Systems



**Heinz Stockinger**  
**CERN & INFN**

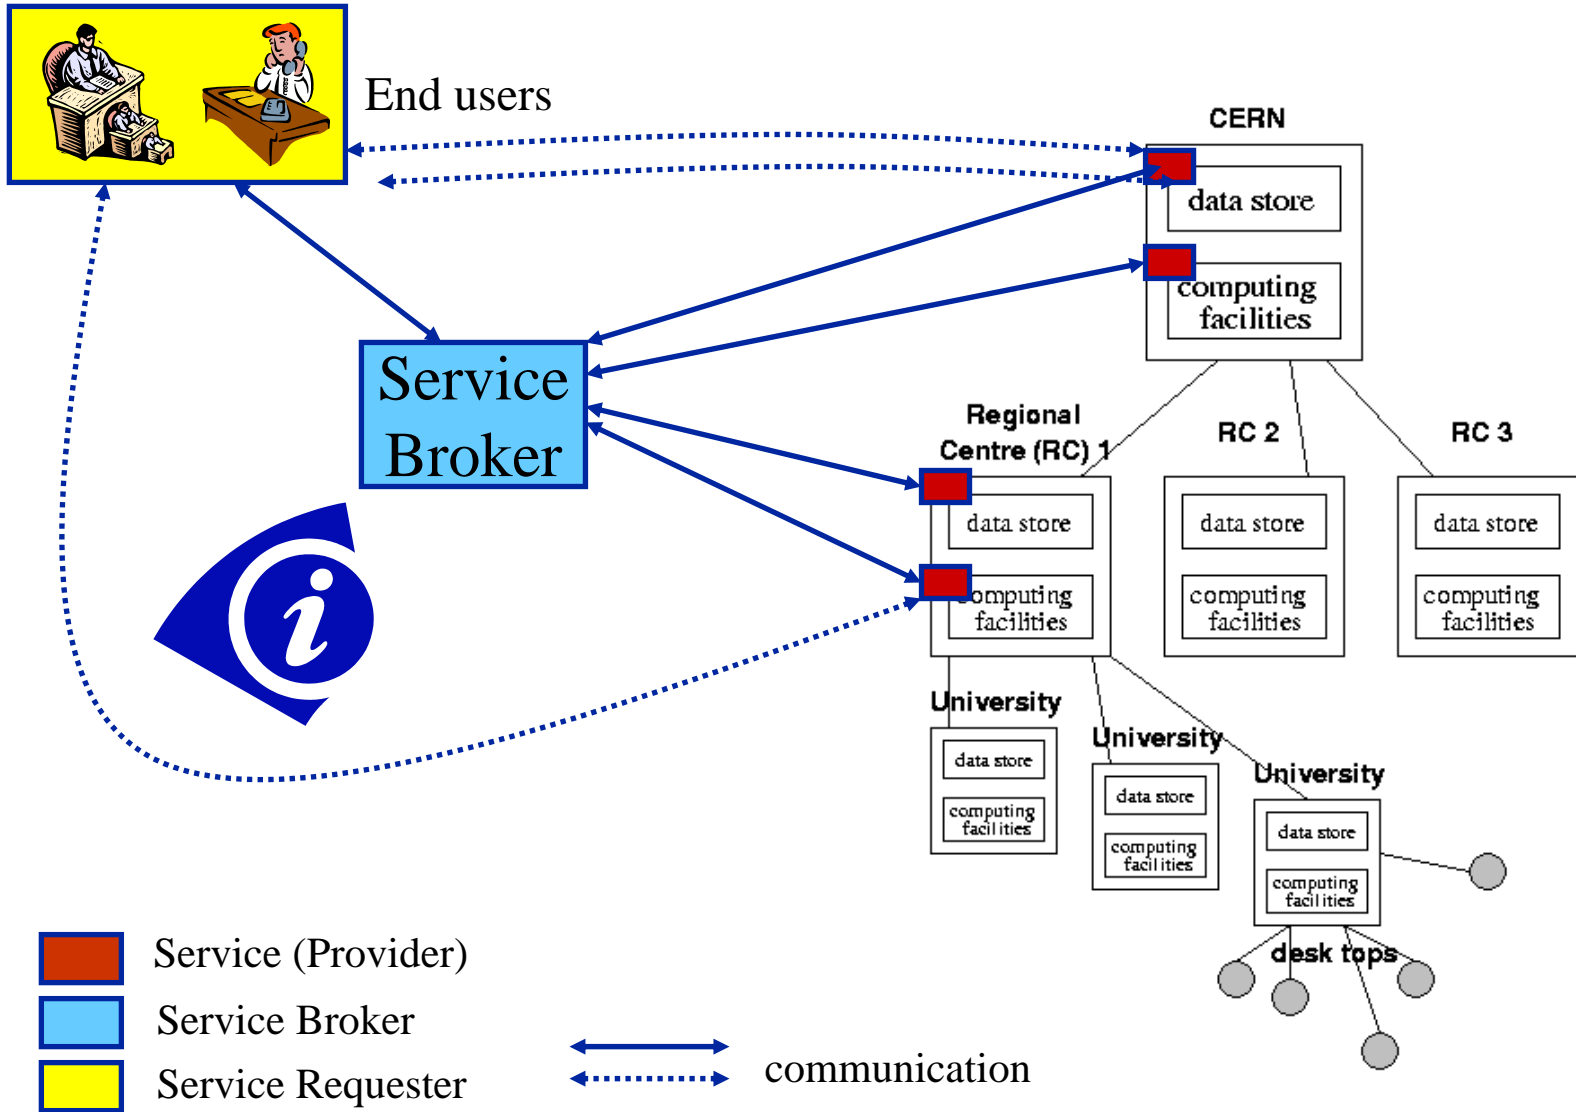
Information Systems in General

Specific Implementations

R-GMA – MDS2 – MDS3

Schema

# A Reference Grid



# Some General Definitions

## ◆ Computing Element (CE)

- Gateway to computing resources
- Grid interface to Local Resource Management (LRMS) systems (PBS, LSF, Condor,..)

Can have 1 or many Worker Nodes

## ◆ Worker Node (WN)

- Resource to execute application

## ◆ Storage Element (SE)

- Basic Storage Resource
- Grid interface to storage



Can have several storage types like disk, tape etc.

## ◆ Site

- Set of (Grid) resources

Can have 0 or many CEs


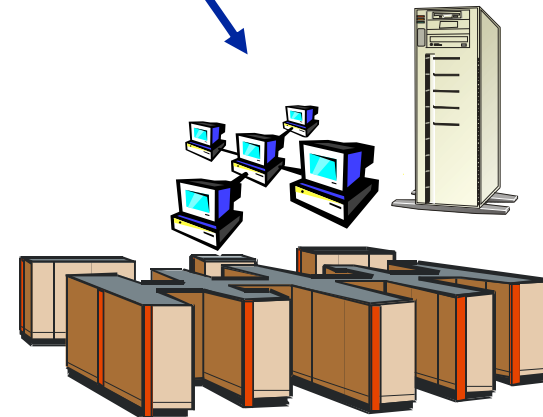

Can have 0 or many SEs



# Information Systems

## ◆ Some Requirements

- The system must be able to cope with nodes in a **distributed environment**
  - Dynamic **addition and deletion information providers** (producers) should be supported
  - It must have a **security system** able to address the access to information at a fine level of granularity
  - It must be able to work well on an unreliable **Wide Area Network (WAN)**
  - The system must allow new **data types (schema)** to be defined
  - Scalability, performance, reliability
- ◆ Often, information service and monitoring are separated: here we discuss them as a unit
- ◆ For simplicity: we also consider any cataloging system as an “information service”
- E.g. Replica Catalogue that provides information about replica locations (details in Data Management lecture)



# Features of an Information System

## ◆ Provides **information on** both:

- The **Grid** itself (i.e. the basic resources and services)
  - Mainly for the middleware services
  - The user may query it to understand the status of the Grid
- Grid applications
  - For users

## ◆ **Schema** required

- To describe all resources in the Grid
- Example: Computing resource and available CPU

## ◆ **Query "language"** required

Features are similar to a Database Management System

## ◆ Several architectural **approaches**

- No common agreement yet but several standardisation efforts
- Standards are required to make Grids interoperable
  - Analogy: SQL/ODBC can be used to query database implementations of several different vendors

# Information Services in Distributed Systems

- ◆ DNS (Domain Name Service (DNS))
  - Name/location information about hosts (resources)
- ◆ UDDI (Universal Description, Discovery & Integration)
  - For web services (**registry**)
- ◆ Other **resource and service registries** like
  - ... ? *You name it*
  - Often different semantics are used
- ◆ Different **catagorisations**
  - White pages
  - Yellow pages
- ◆ Information can be **static** or **dynamic** ...

# Grid Approaches - 1


- ◆ Many Grid projects provide their own implementations:
  - Globus: MDS (Monitoring and Discovery Service)
  - EGEE (EDG): R-GMA (Relational Grid Monitoring Architecture)
  - GridLab: extensions to MDS
  - UNICORE: Incarnation Database (IDB)
  - etc.
- ◆ Need for standardisation since a **service registry is required by all Grid systems that use service oriented architectures**
- ◆ Grid Monitoring Architecture (GMA)
  - For information storage (**registry**) and general monitoring of services and status
  - Standardisation in GGF
  - One of the early standard proposals in GGF



# Grid Approaches - 2

- ◆ Several groups in Global Grid Forum work in the area of **Information Systems and Performance**
  - CIM based Grid Schema
  - Discovery and Monitoring Event Description
  - Network Measurement
  - Grid Information Retrieval
  - Grid Benchmarking
  - Relational Grid Information Services
- ◆ Here we only concentrate on three representative implementations that cover many (all?) of the above fields
  - R-GMA
  - MDS 2 (early Grid solution), MDS 3 (web service approach)
- ◆ The three approaches use **different protocols** but **same schema**

# Grid Monitoring Architecture (GMA)



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

Defined in GGF Document: GFD-I.7

<http://forge.gridforum.org/projects/ggf-editor/document/GFD-I.7>

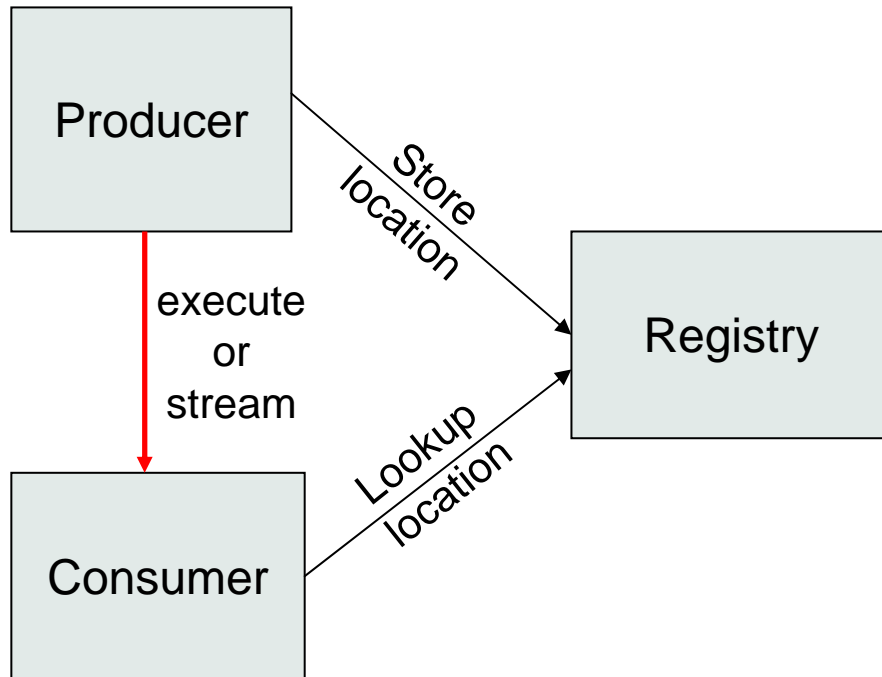
Information Systems in General

## Specific Implementations

R-GMA – MDS2 – MDS3

Schema

# Relational Grid Monitoring Architecture (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

## R-GMA (2)

- ◆ **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.**
- ◆ All R-GMA tuples are time-stamped
- ◆ Uses web service technology:
  - Tomcat web server – client server interaction via HTTP
  - SOAP is not yet used
  - Plans for "OGSIfying" R-GMA, i.e. use OGSA approach
- ◆ Replaces MDS in EU DataGrid project

# 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**
- ◆ All R-GMA tuples are time-stamped

# Example: 2 tables

## ◆ Service

|              |              |   |
|--------------|--------------|---|
| <b>URI</b>   | VARCHAR(255) | URI to contact the service  |
| <b>VO</b>    | VARCHAR(50)  | Where info should be published – or an empty string to indicate all |
| type         | VARCHAR(50)  | Type of service   |
| emailContact | VARCHAR(50)  | The e-mail of a human being to complain to                          |
| site         | VARCHAR(50)  | Domain name of site hosting the service                             |
| secure       | VARCHAR(1)   | 'y' or 'n' - indicates whether or not this is a secure service      |
| majorVersion | INT          | Version of protocol not implementation                              |
| minorVersion | INT          | Version of protocol not implementation                              |
| patchVersion | INT          | Version of protocol not implementation                              |

## ◆ ServiceStatus

|            |              |   |
|------------|--------------|---|
| <b>URI</b> | VARCHAR(255) | URI to contact the service              |
| status     | INT          | status code. 0 means the service is up. |
| message    | VARCHAR(255) | Message corresponding to status code    |

# SQL example 1

◆ SELECT DISTINCT type  
FROM Service

```
+-----+  
| type |  
+-----+  
| GridFTP |  
| GRIS |  
| RFIO |  
+-----+
```

---

```
| R-GMA.ResilientStreamProducerService |  
| R-GMA.ArchiverService |  
| R-GMA.StreamProducerService |  
| R-GMA.CanonicalProducerService |  
| R-GMA.DBProducerService |  
| R-GMA.LatestProducerService |  
| GIN |  
| R-GMA.RegistryService |  
| R-GMA.SchemaService |  
| R-GMA.BrowserService |  
| GOUT |  
| edg-netmon |  
| edg-iperf |  
| edg-udpmon |  
| myproxy |  
| edg-pinger |  
+-----+
```

25 Rows in set

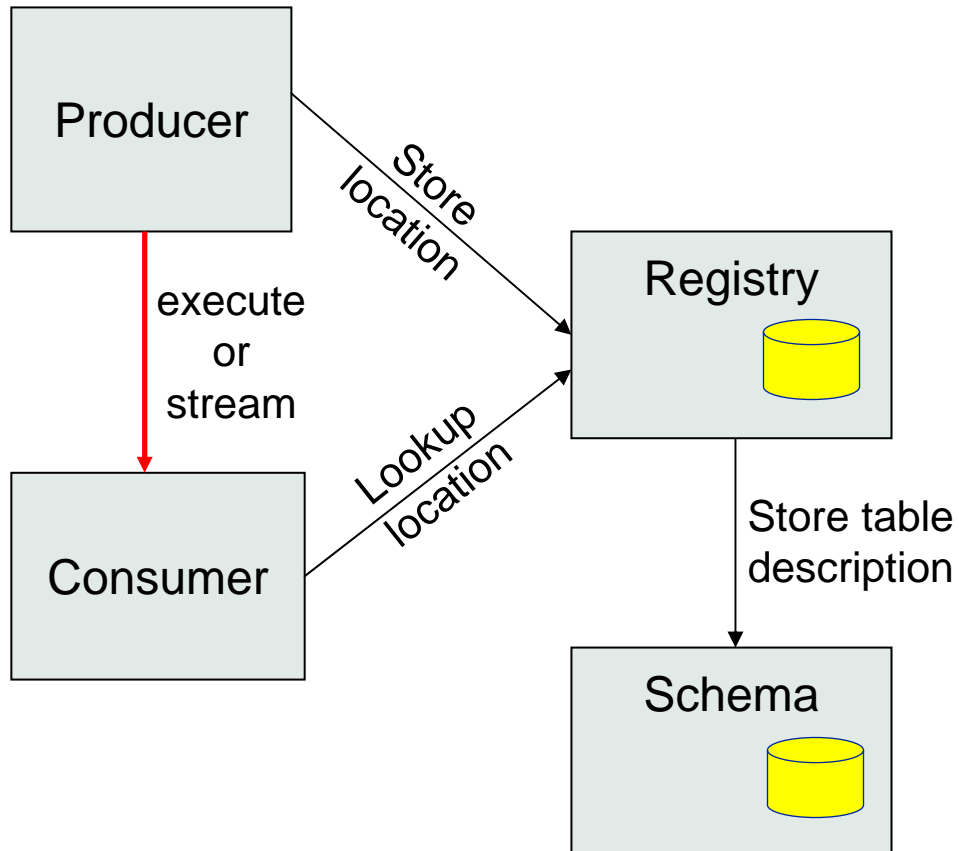


## SQL Example 2

```
SELECT Service.site, ServiceStatus.status,  
ServiceStatus.message, Service.URI  
FROM Service,ServiceStatus  
WHERE Service.URI = ServiceStatus.URI  
AND ServiceStatus.status <> 0  
AND Service.Type = 'GIN'
```

```
+-----+-----+-----+-----+  
| site      | status | message          | URI          |  
+-----+-----+-----+-----+  
| nikhef.nl | 2      | Gin is stopped  | http://tbn03.nikhef.nl/GIN |  
| nikhef.nl | 2      | Gin is stopped  | http://tbn09.nikhef.nl/GIN |  
| nikhef.nl | 2      | Gin is stopped  | http://tbn16.nikhef.nl/GIN |  
+-----+-----+-----+-----+  
3 Rows in set
```

# 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 (Aggregation)

| 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 |

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

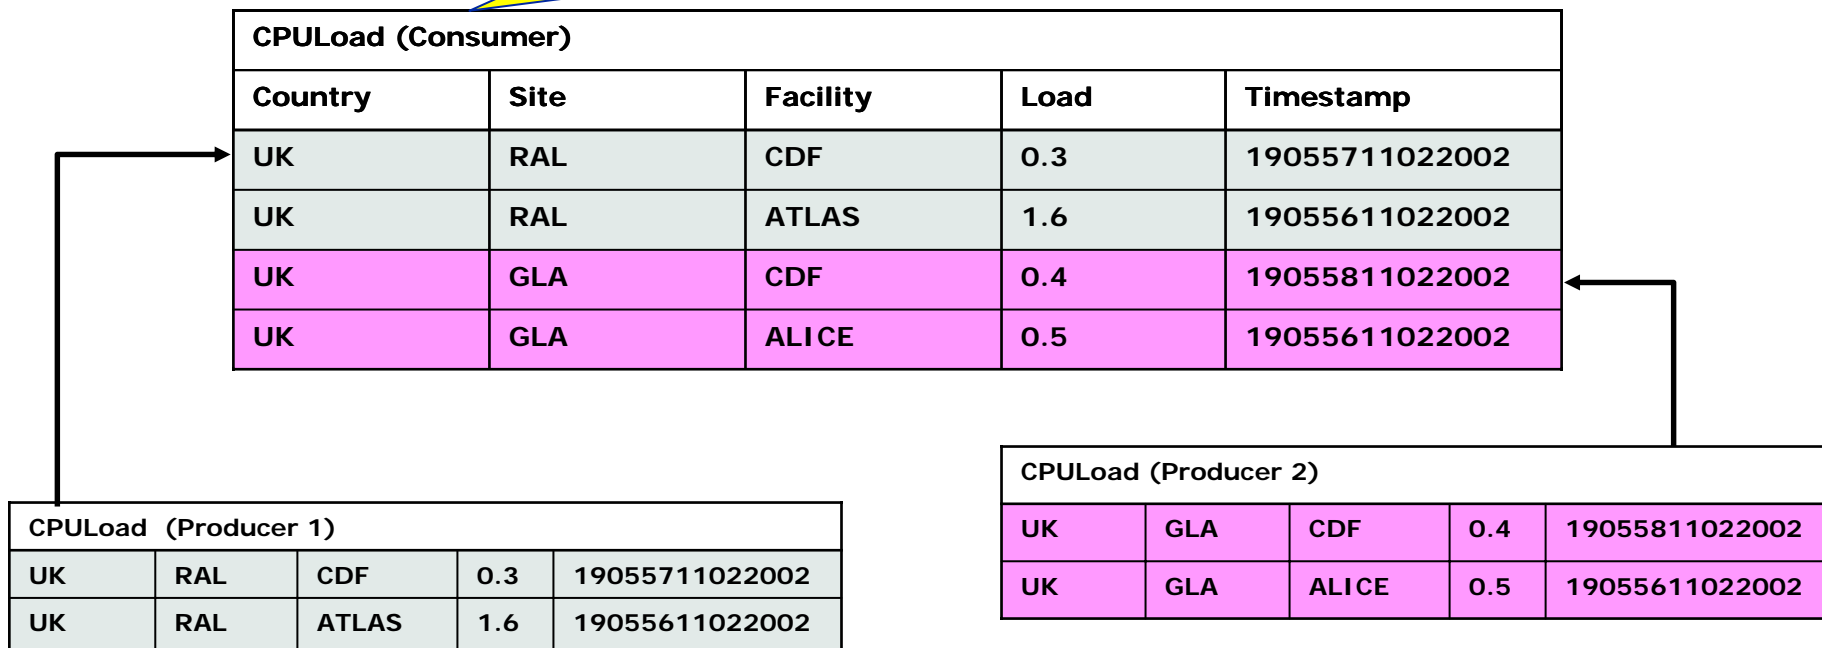
| CPULoad (Producer 3) |      |       |     |                |
|----------------------|------|-------|-----|----------------|
| CH                   | CERN | ATLAS | 1.6 | 19055611022002 |
| CH                   | CERN | CDF   | 0.6 | 19055511022002 |

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

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

# Queries over "global" table – merging streams

SELECT \* from CPUload WHERE country = 'UK'



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

| CPUload (Producer 3) |      |          |      |                |
|----------------------|------|----------|------|----------------|
| Country              | Site | Facility | Load | Timestamp      |
| 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<br>(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           |  | up | message       |
|   |       |      |                |      |        |              |              | gppse01       |  | y  | SE is running |
|   |       |      |                |      |        |              |              | gppse02       |  | n  | SE ERROR 101  |
|   |       |      |                |      |        |              |              | lxshare0404   |  | y  | SE is running |

# 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
- ◆ R-GMA packaged SQL (edg-rgma-util)
  - e.g. edg-rgma-util contacts:
    - Command: `SELECT siteName, sysAdminContact, userSupportContact, siteSecurityContact FROM SiteInfo`



## R-GMA Browser

[All tables](#)[EDG Info Providers](#)[Network Monitoring](#)[CMS](#)[Home](#)[Predefined Queries](#)[Service Status](#)[Site Info](#)[Table Sets](#)

## EDG Info Providers

[GlueCE](#)[GlueCEAccessControlBaseRule](#)[GlueCESEBind](#)[GlueCluster](#)[GlueHostRemoteFileSystem](#)[GlueSA](#)[GlueSAAccessControlBaseRule](#)[GlueSE](#)[GlueSEAccessProtocol](#)[GlueSEAccessProtocolSupportedSecur](#)[GlueSL](#)[GlueSubCluster](#)[GlueSubClusterSoftwareRunTimeEnvirc](#)[SiteInfo](#)

SELECT

FROM **GlueCE**WHERE 

Query

[Description of table](#)

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://gpprg06.gridpp.rl.ac.uk:8080/R-GMA/LatestProducerServlet ConnectionId:301164355

Continuous Producer

 producerServlet:http://gpprg06.gridpp.rl.ac.uk:8080/R-GMA/StreamProducerServlet ConnectionId:291549138

 producerServlet:http://gpprg06.gridpp.rl.ac.uk:8080/R-GMA/StreamProducerServlet ConnectionId:291549226

Query



# edg-rgma – Example

```
$> edg-rgma
```

```
rgma> stream declare userTable
```

```
rgma> stream minret 0.2
```

```
rgma> stream INSERT into userTable (userId, aString,  
    aReal, anInt) values ('fisher', 'hello', 3.162, 21)
```

```
rgma> timeout 0.3
```

```
rgma> old continuous SELECT * from userTable
```

```
+-----+-----+-----+-----+-----+-----+  
| userId | aString | aReal | anInt | MeasurementDate | MeasurementTime |  
+-----+-----+-----+-----+-----+-----+  
| fisher | hello   | 3.162 | 21    | 2003-11-11      | 11:06:01        |  
+-----+-----+-----+-----+-----+-----+  
1 Rows in set
```



# 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 = edg_rgma_perl::StreamProducer_new();
```

## Python

```
myProducer = edg_rgma_python.StreamProducer_new()     or  
myProducer = rgma.StreamProducer()
```

Information Systems in General

## Specific Implementations

R-GMA – MDS2 – MDS3

Schema

# Information Systems in Globus

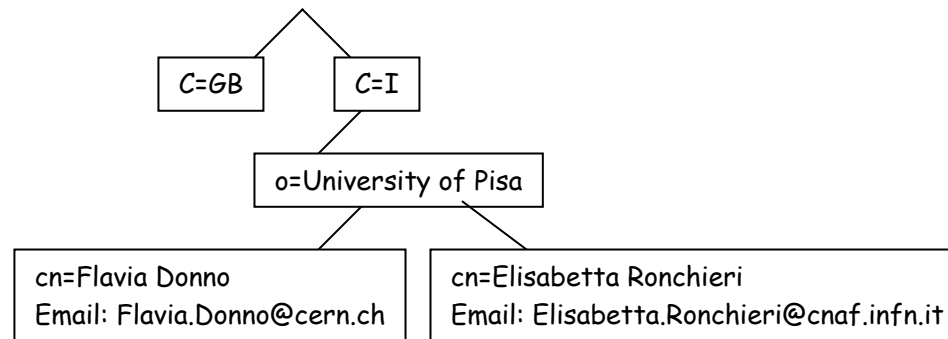
- ◆ Currently, two versions of MDS (Monitoring & Discovery Service):
  - MDS 2 in GT2: based on LDAP protocol
  - MDS 3 in GT3: based on web service approach
    - Provides the same features as the earlier version MDS 2
- ◆ MDS2 has been in used at several Grid projects for a few year
- ◆ Uses the “early” set of protocols in GT2
  - LDAP (Lightweight Directory Assess Protocol) – network protocol
    - Was originally used in other services too like the “early” Replica Catalogue provided by Globus

# LDAP-based Services

- ◆ LDAP (Lightweight Directory Assess Protocol)
  - See RFC 1777 for details: <ftp://ftp.rfc-editor.org/in-notes/rfc1777.txt>
- ◆ Many “early” services provided Globus or the EU DataGrid project use the **LDAP protocol** for:
  - Information service (Globus’ **MDS** (Monitoring and Discovery Service))
  - Replica Catalogue (Globus): to locate replicas
- ◆ Several services are/will be replaced by newer services that use different technologies:
  - Globus’ MDS vs. EDG’s R-GMA vs. MDS3
  - LDAP based replica catalogue is replaced by SQL/web service based service
- ◆ General trend in the Grid community:
  - Web service approach (HTTP, SOAP, WSDL)
  - Open Grid Service Architecture (OGSA) – based on web services
  - SQL based services rather than LDAP based for storing/retrieving information

# LDAP Hierarchy

- ◆ Lightweight Directory Access Protocol (LDAP) offers a hierarchical view of information
- ◆ The entries are arranged in a Directory Information Tree (DIT)

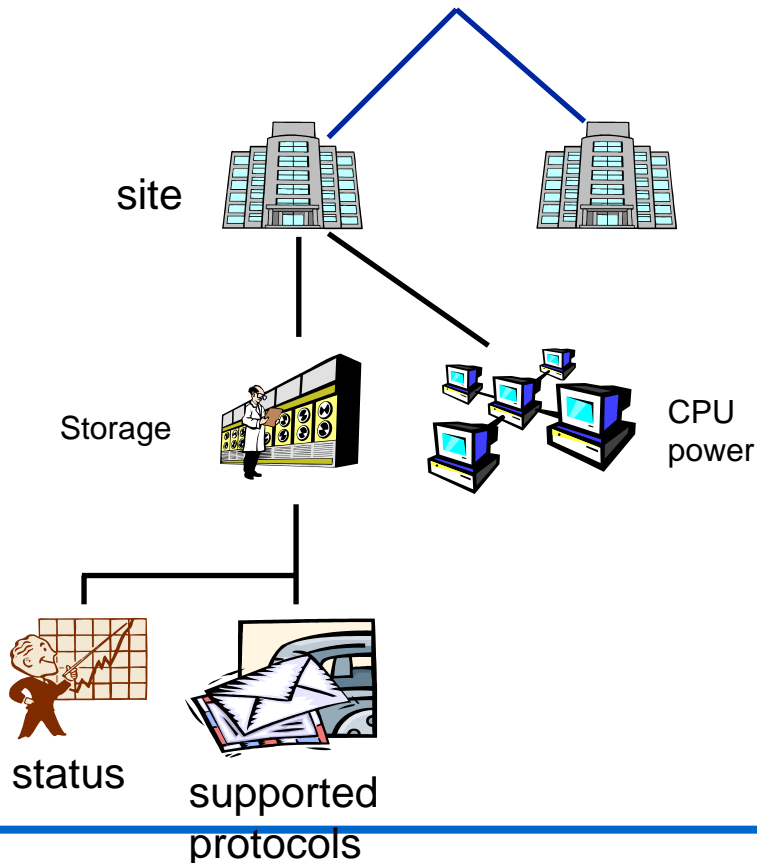


- ◆ The entry name represents the Relative Distinguished Name (RDN). *example: "cn=Flavia Donno"*
- ◆ An entry is identified by its Distinguished Name
  - This is its RDN with the Distinguished name of its parent.
  - *example: "C=I, o=University of Pisa, cn=Flavia Donno"*

# Usage of LDAP-based Services

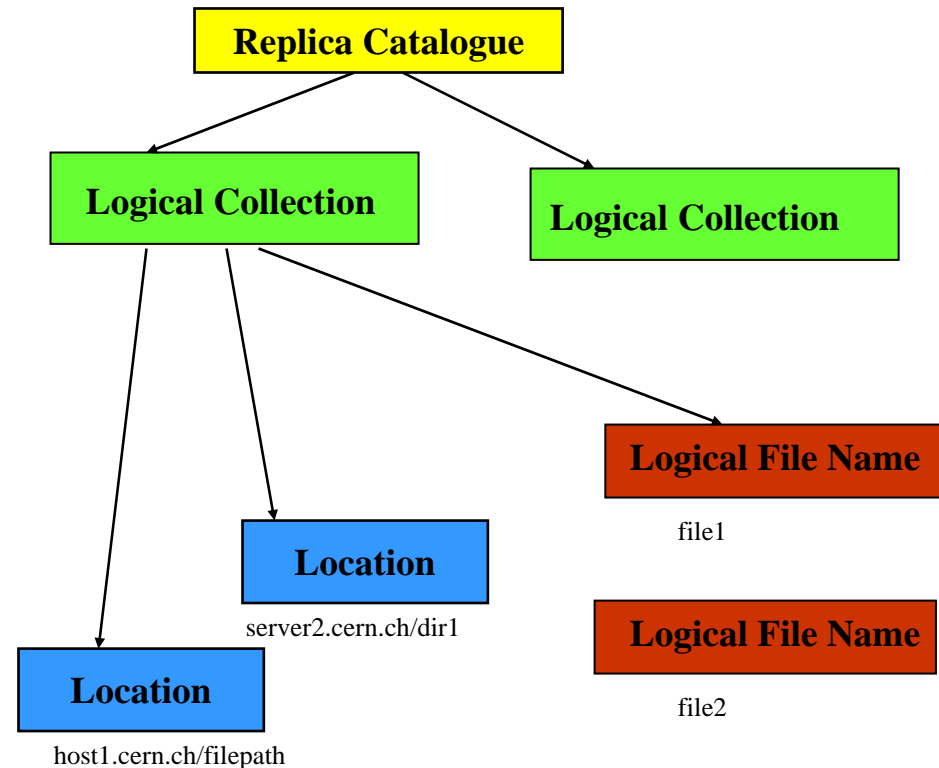
## Information Service

- ◆ Used for **identifying and locating** all **resources** in the Grid




## Replica Catalogue

- ◆ Used for **locating identical replicas** of files
  - Logical and physical file names are used



# RDNs and DNs in Globus Information Service (MDS)

## ◆ RDN



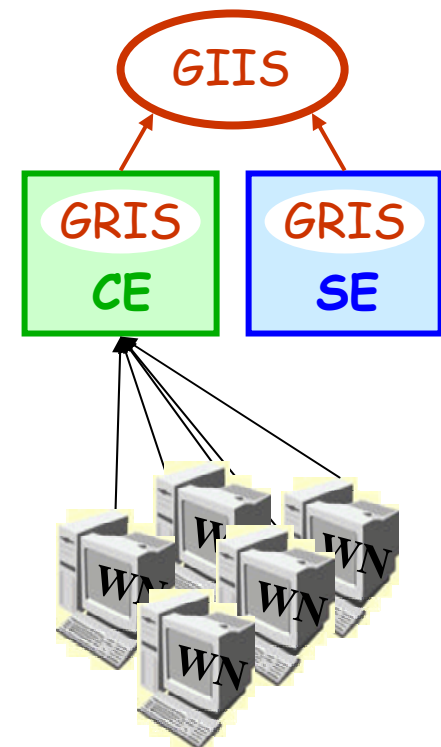
- SE
  - `seId=dev02.hepgrid.clrc.ac.uk`
- Status
  - `in=status`
- Protocols
  - `seProtocol=gridftp`
  - `seProtocol=rfio`
  - `seProtocol=file`

## ◆ DN

- Site
  - `Mds-Vo-name=ral-dev, o=Grid`
- SE
  - `seId=dev02.hepgrid.clrc.ac.uk, Mds-Vo-name=ral-dev, o=Grid`
- Status
  - `in=status, seId=dev02.hepgrid.clrc.ac.uk, Mds-Vo-name=ral-dev, o=Grid`
- Protocols
  - `seProtocol=gridftp, seId=dev02.hepgrid.clrc.ac.uk, Mds-Vo-name=ral-dev, o=Grid`
  - `seProtocol=rfio, seId=dev02.hepgrid.clrc.ac.uk, Mds-Vo-name=ral-dev, o=Grid`
  - `seProtocol=file, seId=dev02.hepgrid.clrc.ac.uk, Mds-Vo-name=ral-dev, o=Grid`

# MDS Hierarchy: GRISs & GIISs

- ◆ **Information providers** are scripts which when invoked by the LDAP server **make available the desired information**
  - Information is cached by the server to improve performance
  - LDAP was designed for use with relatively static data, not rapidly changing data
- ◆ Within MDS the EDG **information providers** are **invoked by a local LDAP server**, the Grid Resource Information Server (GRIS)
- ◆ **"Aggregate directories"**, Grid Information Index Servers (GIIS), are used to group resources
- ◆ The GRISs use **soft state registration** to register with one or more GIISs
- ◆ The GIIS can then act as a single point of contact for a number of resources
  - A GIIS may represent a site, country, virtual organization, etc.
- ◆ In turn a GIIS may register with another GIIS





# Querying the Information & Monitoring Service

- ◆ Queries can be posed to the current Information and Monitoring Service using LDAP search commands:

```
$ldapsearch\  
-x\  
-H ldap://lxshare0225.cern.ch:2135\  
-b 'Mds-Vo-name=datagrid,o=grid'\  
'objectclass=StorageElement'\  
seId SEsize \  
-s base|one|sub
```

- ◆ “simple” authentication
- ◆ uniform resource identifier
- ◆ base distinguished name for search filter
- ◆ attributes to be returned
- ◆ scope of the search specifying just the base object, one-level or the complete subtree

# Querying the GRIS/GIIS Hierarchy



Mds-Vo-name  
=datagrid



Mds-Vo-name  
=countryA



Mds-Vo-name  
=countryB



Mds-Vo-name  
=siteA

Mds-Vo-name  
=siteB

Mds-Vo-name  
=siteC

Mds-Vo-name  
=siteD

- ◆ Mds-Vo-name=datagrid,o=grid
  - This will look at all the data
- ◆ Mds-Vo-name=siteB, Mds-Vo-name=countryA,Mds-Vo-name=datagrid,o=grid
  - This will look at all the data from siteB
- ◆ Mds-Vo-name=countryA,o=grid
  - This will look at all the data from countryA
- ◆ Mds-Vo-name=siteB,Mds-Vo-name=countryA,o=grid
  - This will look at all the data from siteB
- ◆ Mds-Vo-name=siteB,o=grid
  - This will look at all the data from siteB

Information Systems in General

## Specific Implementations

R-GMA – MDS2 – MDS3

Schema

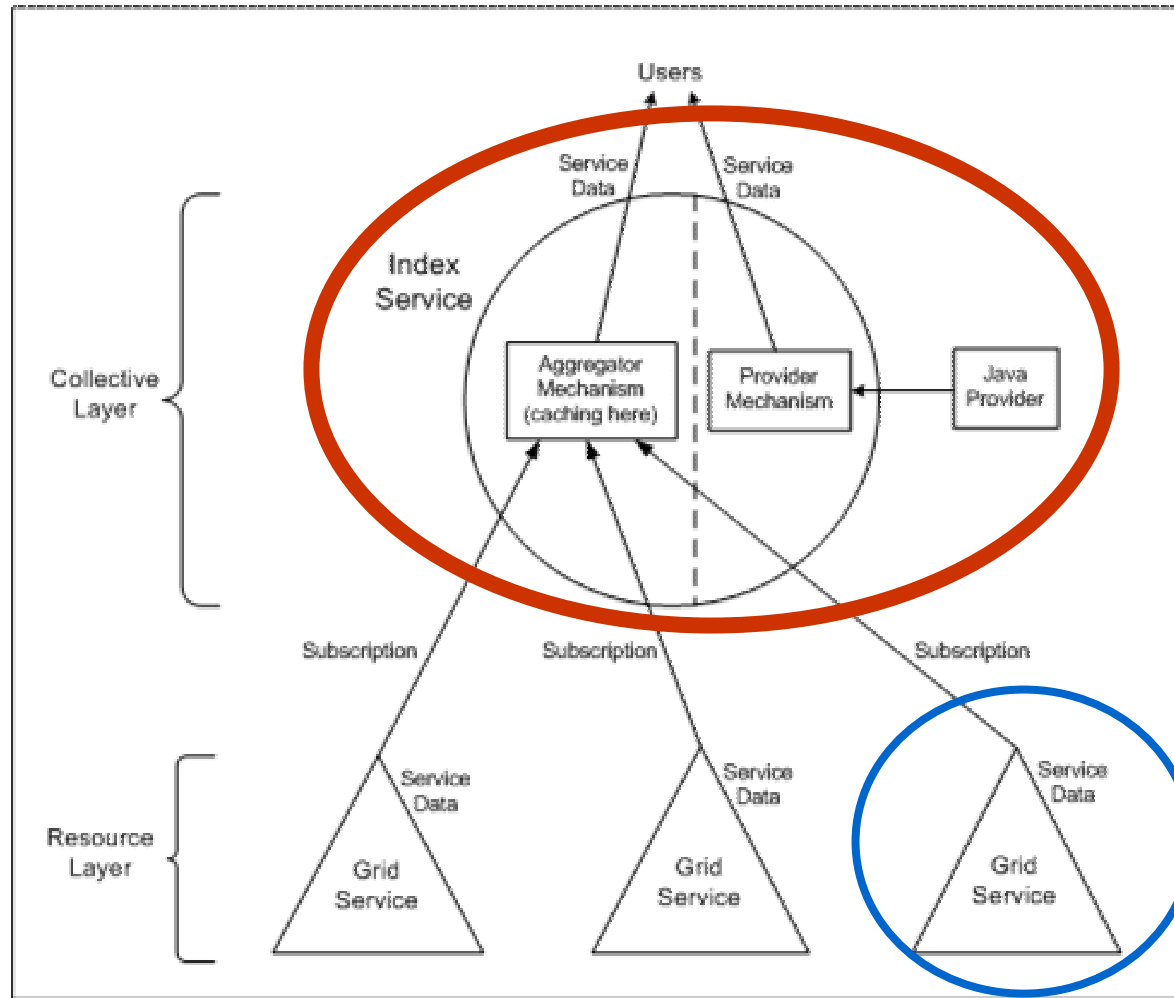
# GT3 Index Service - MDS 3

- ◆ Key capabilities of Index Service:
  - interface for connecting external Service Data Provider programs to service instances
  - Generic framework for aggregation of service data
  - Registry of Grid services
- ◆ Information Service: *generates, registers, indexes, aggregates, subscribes, monitors, queries, or displays Service Data*
- ◆ The GT3 **Index Service** provides an **information aggregation service** that is more extensible than the GT2 GIIS
  - GRIS is not needed since each Grid service has its **Service Data Element** that contains the service status

# GT3 Index Service cont.

- ◆ MDS 2 and MDS 3 are not interoperable:
  - MDS2: Idif format
  - MDS3: XML format
  - But: GT2 information providers can be ported into GT3
- ◆ Host information is represented by default in the **GLUE schema**
- ◆ Queries:
  - **FindServiceData** operation from the GridService interface is used
  - standard, extensible query operation against a service's Service Data Elements may use either the simple default "by name" query or a more complex language like **XPath** or **XQuery**

# GT3 Index Service Overview



**Standard  
Grid Service**

Information Systems in General

## Specific Implementations

R-GMA – MDS2 – MDS3

Schema

# MDS vs. R-GMA

|                     | MDS 2   | R-GMA  | MDS 3   |
|---------------------|---|--|---|
| communicaton        | LDAP  | HTTP   | SOAP  |
| Information storage | LDAP based backends   | Relational Data Base                                   | Any web-service compliant backend   |
| queries             | LDAP queries<br><pre>Ldapsearch -x -H ldap://host.cern.ch:2135\ -b 'Mds-Vo-name=datagrid,o=grid'\ 'objectclass=StorageElement'\ seld SEdize</pre> | SQL queries<br><pre>Select * from StorageElement</pre> | default "by name" or<br>Xpath/Xquery<br><pre>ogsi-find-service-data -service http://128.9.72.46:9103/ogsa/services/core /registry/ContainerRegistryService -sde gsdl:entry -output registry</pre> |
| Components          |   |  |   |



Information Systems in General

Specific Implementations

R-GMA – MDS2 – MDS3

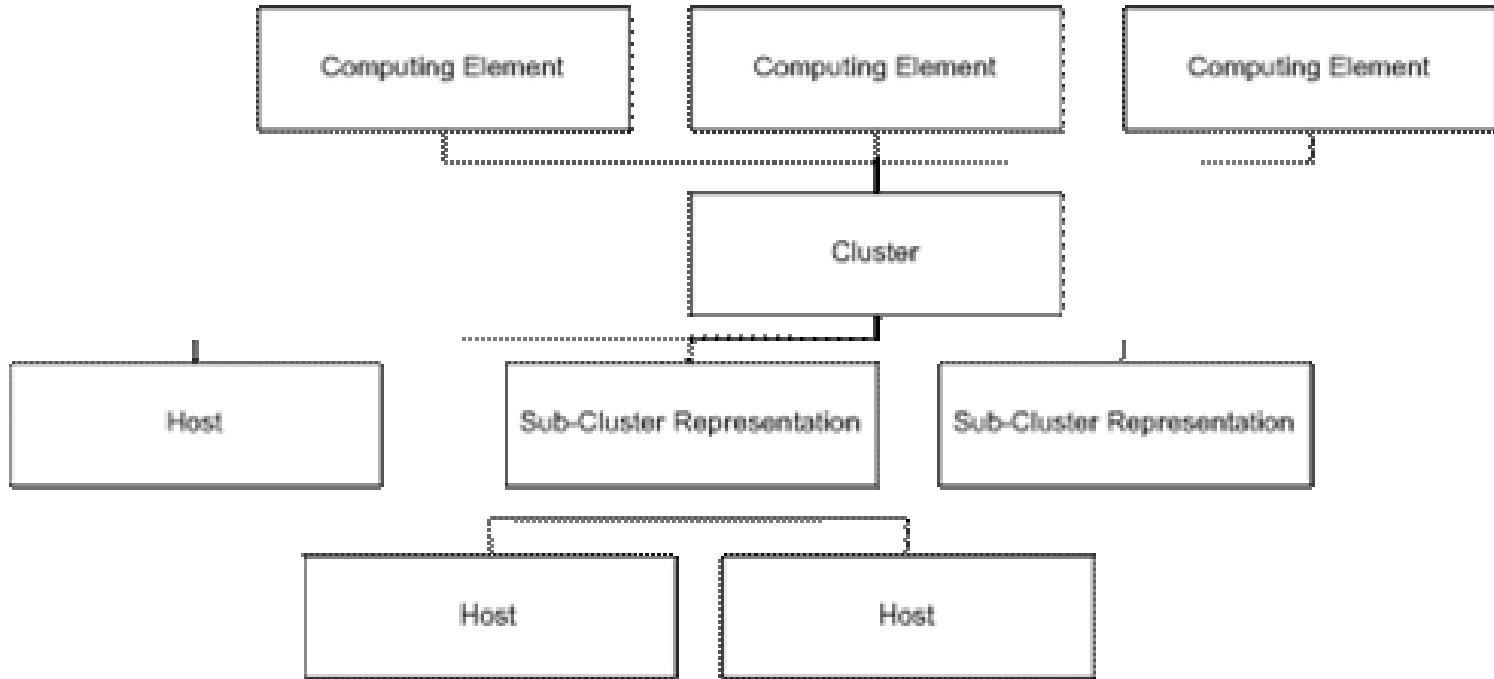
**Schema**

# The GLUE schema

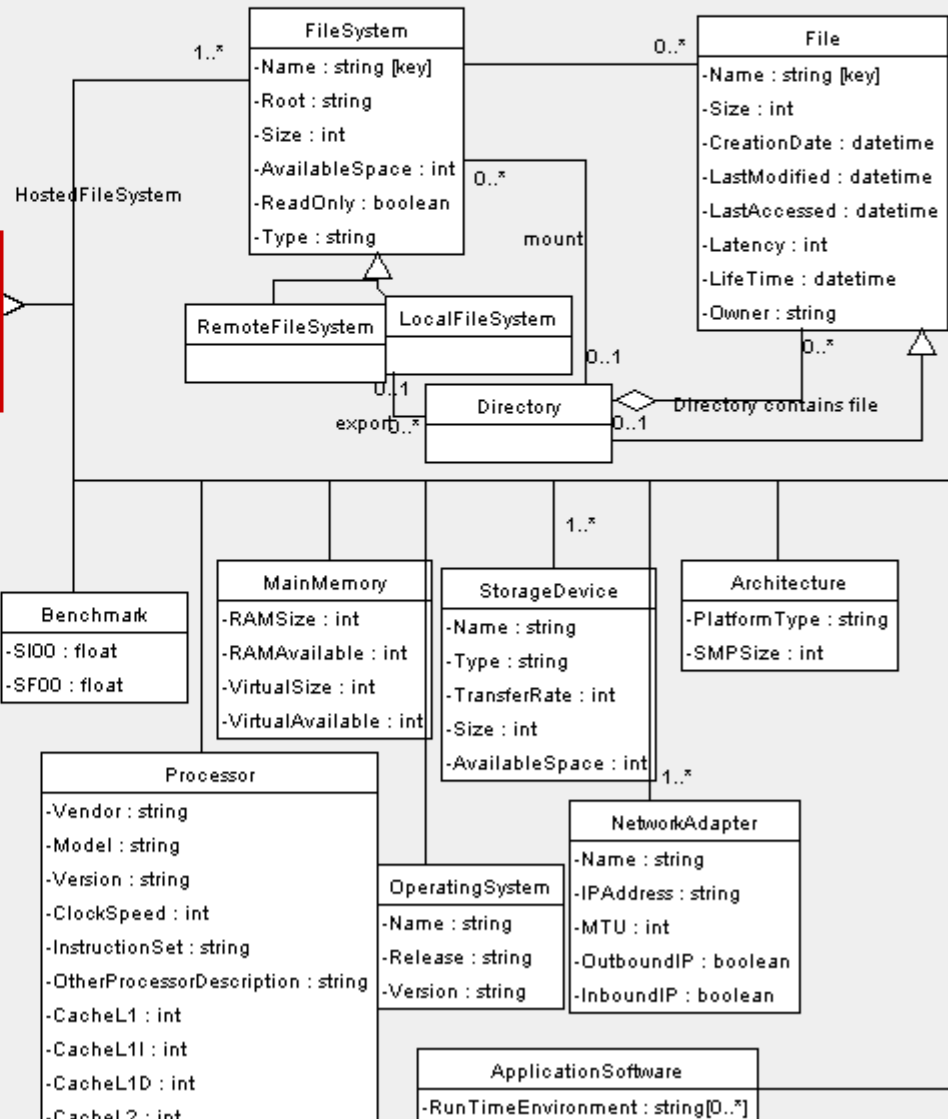






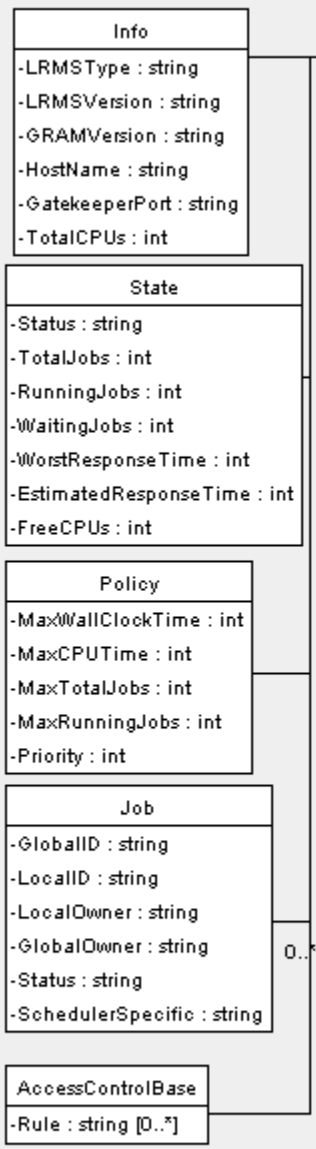
- ◆ Developed within High Energy Physics (HEP) community
  - DataGrid
  - DataTAG
  - Globus
- ◆ Currently defines CEs and SEs
- ◆ Coming soon...
  - Network monitoring
  - Services

# Computing Element Hierarchy

- ◆ GLUE Schema v1.1



# Computing Element




# 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      |
| <b>UniqueID</b>     |
| Name                |
| GlueClusterUniqueID |
| RAMSize             |
| RAMAvailable        |
| VirtualSize         |
| VirtualAvailable    |
| PlatformType        |
| SMPSize             |
| OSName              |
| OSRelease           |
| OSVersion           |
| Vendor              |
| .....               |

|                     |
|---------------------|
| GlueCluster         |
| UniqueID            |
| Name                |
| InformationServiceU |

|                               |
|-------------------------------|
| GlueSubCluster                |
| SoftwareRunTime               |
| Environment                   |
| <b>GlueSubClusterUniqueID</b> |
| 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 |

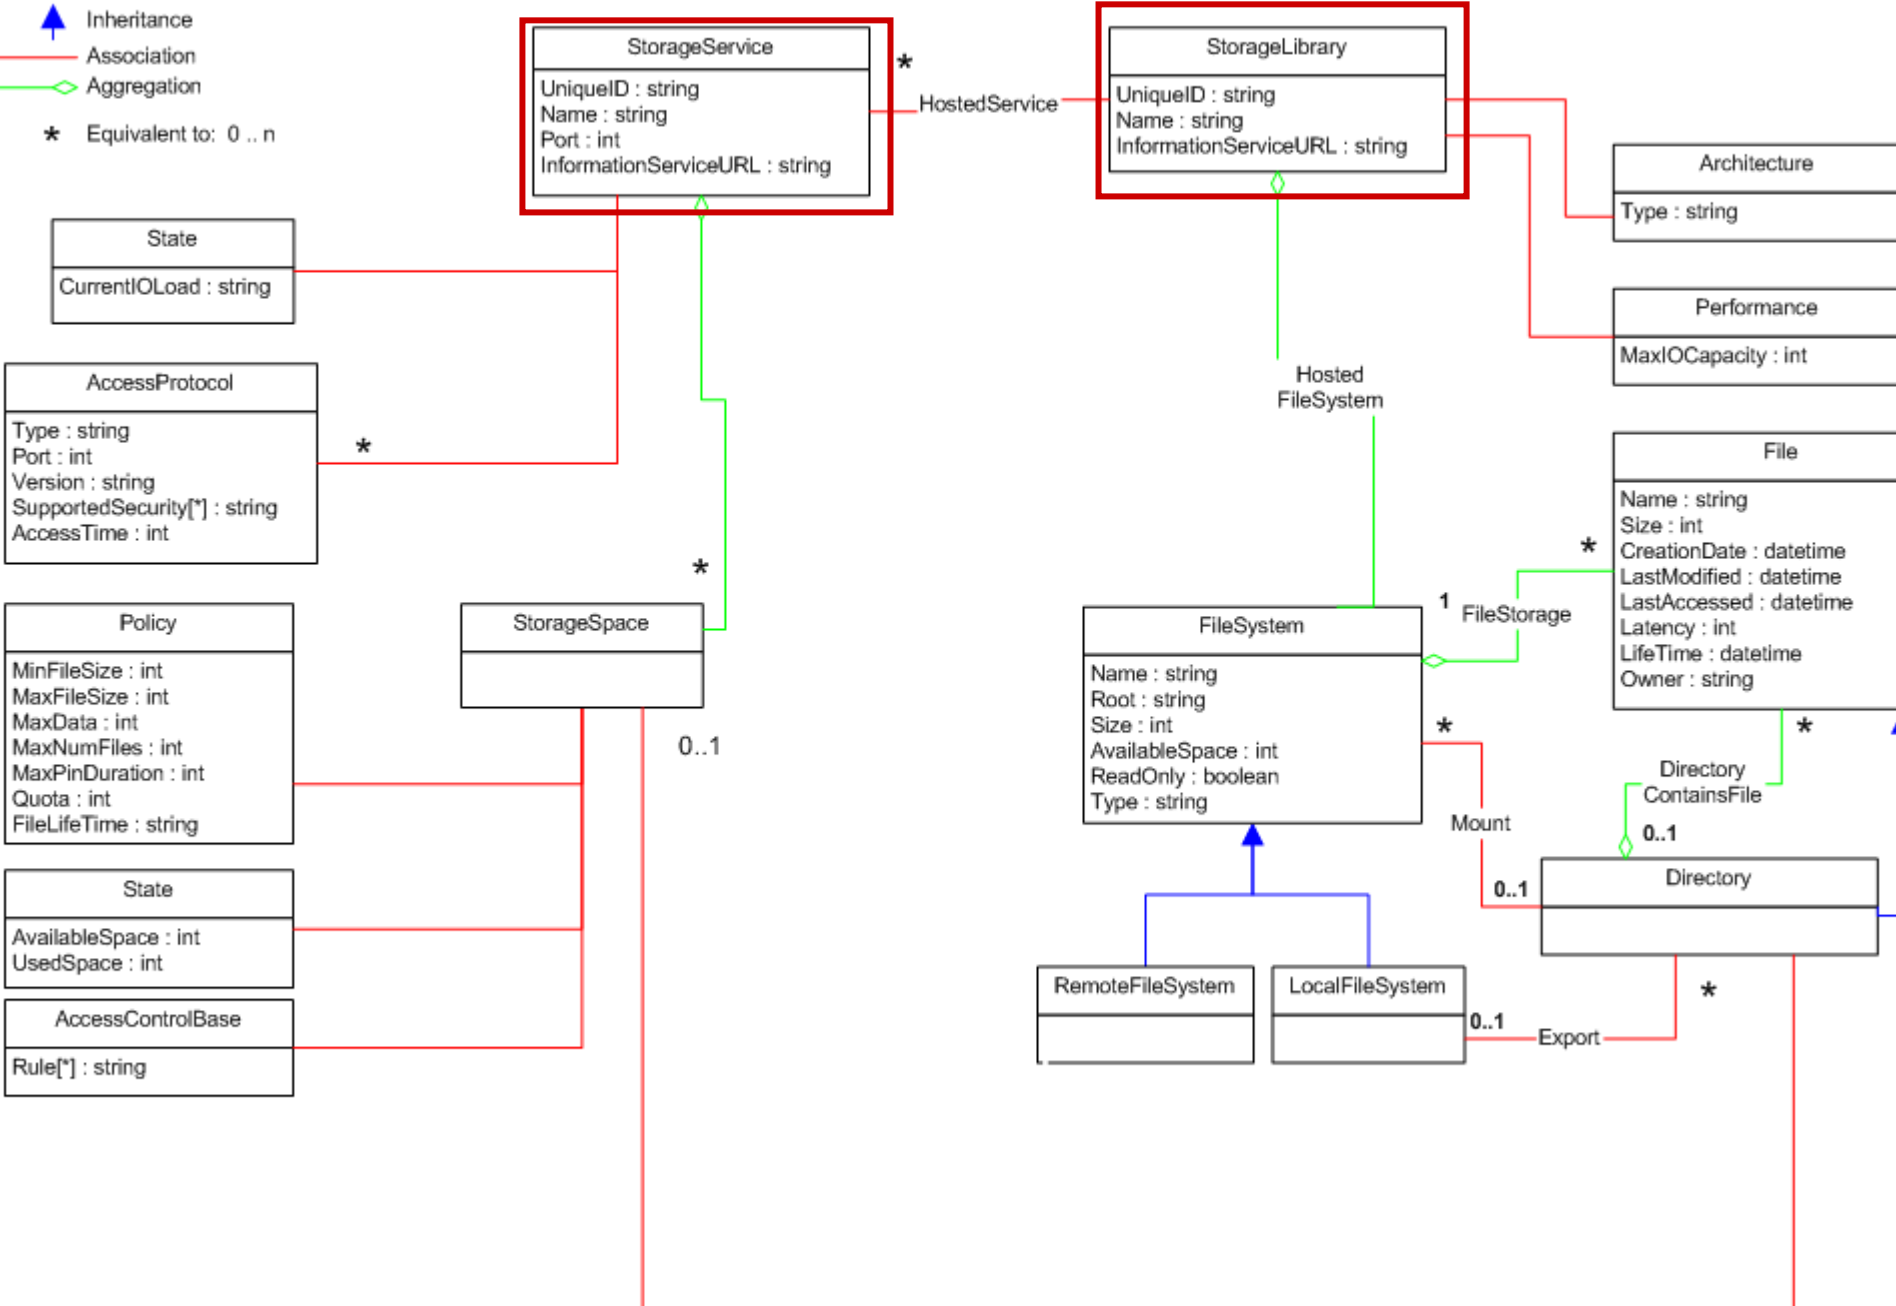
|                                 |
|---------------------------------|
| GlueCEAccess<br>ControlBaseRule |
| GlueCEUniqueID                  |
| Value                           |

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



# Storage Element

- Inheritance
- Association
- Aggregation
- Equivalent to: 0..n



# 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

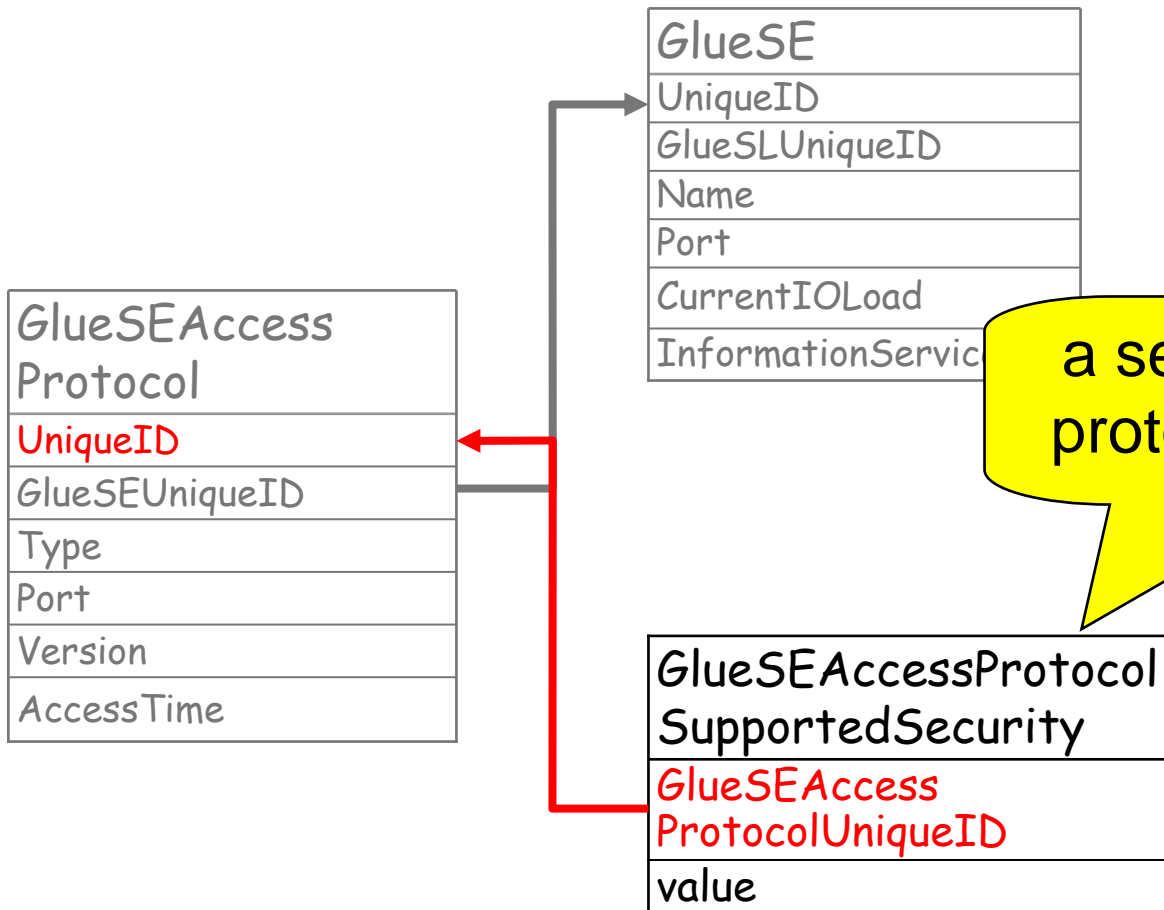


|                |
|----------------|
| GlueSEAccess   |
| Protocol       |
| UniqueID       |
| GlueSEUniqueID |
| Type           |
| Port           |
| Version        |
| AccessTime     |

|                       |
|-----------------------|
| GlueSE                |
| UniqueID              |
| GlueSLUniqueID        |
| Name                  |
| Port                  |
| CurrentIOLoad         |
| InformationServiceURL |

a protocol supported  
by the storage service

# R-GMA Storage Service Tables



a security feature the protocol can deal with

# 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      |

storage space

# 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<br>ControlBaseRule |
| GlueSAUniqueID                  |
| Value                           |





# R-GMA Binding Table



|                       |
|-----------------------|
| GlueSE                |
| UniqueID              |
| GlueSLUniqueID        |
| Name                  |
| Port                  |
| CurrentIOLoad         |
| InformationServiceURL |

related GlueCE  
and GluseSE

|                |
|----------------|
| GlueCESEBind   |
| GlueCEUniqueID |
| GlueSEUniqueID |
| Accesspoint    |

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

# Conclusion

- ◆ Several information systems have been discussed
- ◆ More standardisation is going on
- ◆ Further information
  - <http://www.r-gma.org>
  - <http://www.globus.org/mds>
  - <https://forge.gridforum.org/projects/isp>
  - <http://www.cnaf.infn.it/~sergio/datatag/glue/v11/SE/index.htm>

# Abbreviations

- ◆ CE – Computing Element
- ◆ GMA – Grid Monitoring Architecture
- ◆ MDS – Metadata and Directory Service
- ◆ LDAP – Light-weight Directory Access Protocol
- ◆ LRMS – Local Resource Management System
- ◆ R-GMA – Relational Grid Monitoring Architecture
- ◆ SE – Storage Element
- ◆ WN – Worker Node