



The EU DataGrid - Information and Monitoring Services



The European DataGrid Project Team

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



Aim

- ◆ The aim of the Information and Monitoring Service is to deliver a flexible infrastructure that provides information on both
 1. The EU DataGrid itself
 - Mainly for the middleware packages
 2. Grid applications
 - For users
- ◆ The current system is based on MDS which only offers the first service - to the middleware.
 - The user may however query it to understand the status of the DataGrid



Overview of Talk

- ◆ Requirements of an information and monitoring service
- ◆ Globus MDS (Metacomputing Directory Service or Monitoring and Discovery Service as it is now called)
 - OpenLDAP, a hierarchical database
 - EDG GRIS/GIIS hierarchy
 - EDG information providers
 - Searching the information system
- ◆ R-GMA (Relational Grid Monitoring Architecture)
 - A relational implementation of the Global Grid Forums GMA



Some Requirements of a Grid Information & Monitoring Service

- ◆ The system must be able to cope with nodes in a distributed environment
- ◆ Dynamic addition and deletion of information providers 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 to be defined



Globus MDS

Monitoring and Discovery Services





LDAP attributes

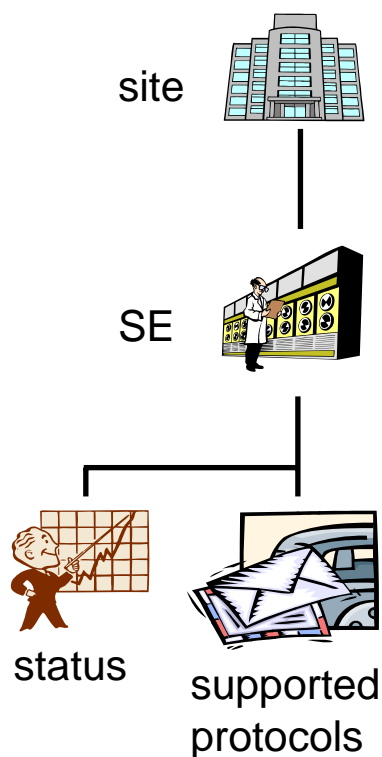
- ◆ EDG currently uses Globus MDS which is built on OpenLDAP
- ◆ A schema describes the attributes and the types of the attributes associated with data objects
- ◆ Example - some attributes of SiteInfo:
 - siteName: RALDEV
 - sysAdminContact: grid.sysadmin@rl.ac.uk
 - userSupportContact: grid.support@rl.ac.uk
 - siteSecurityContact: grid.security@rl.ac.uk
 - dataGridVersion: 1.2
 - InstallationDate: 20020704142800Z



LDAP hierarchy

- ◆ Lightweight Directory Assess Protocol (LDAP) offers a hierarchical view of information
- ◆ The objects are arranged in a Directory Information Tree (DIT)
- ◆ One or more attributes represent the Relative Distinguished Name (RDN)
- ◆ An object is identified by its Distinguished name
 - This is its RDN with the Distinguished name of its parent

RDNs and DN



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



MDS GRISs & GIISs

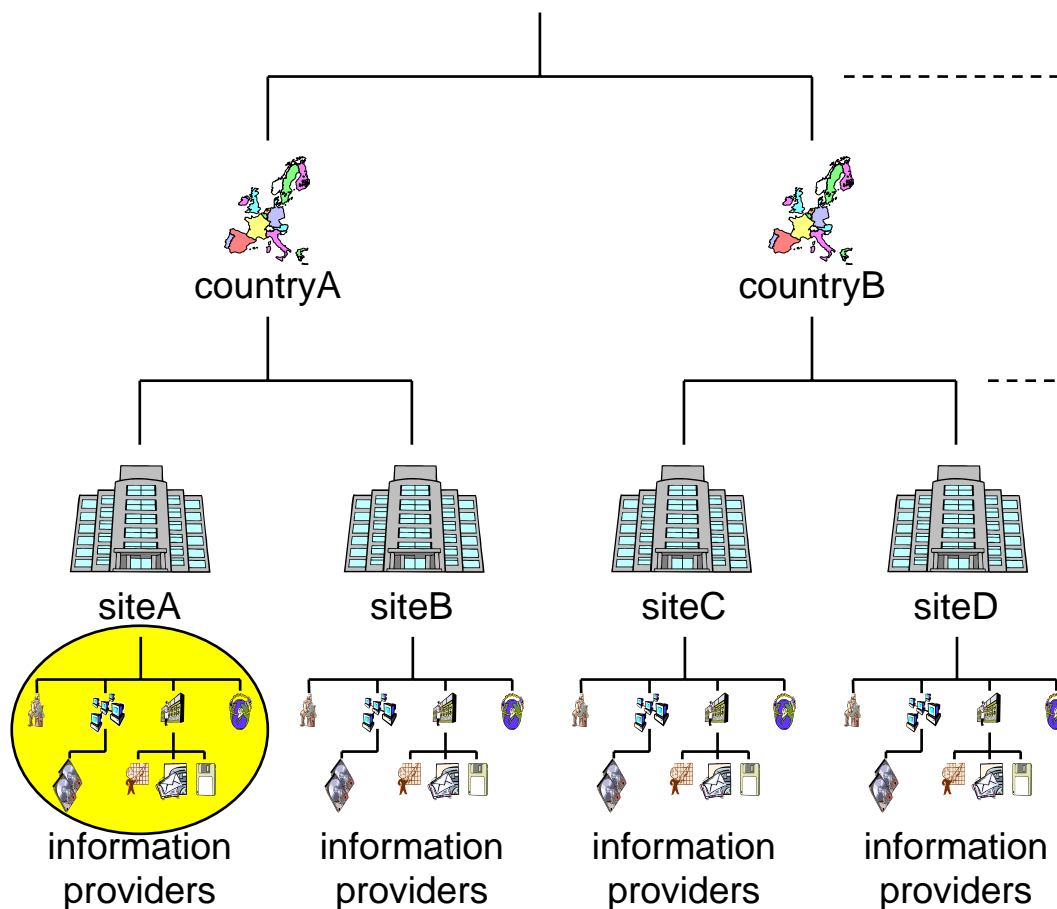
- ◆ Information providers are scripts which when invoked by the LDAP server make available the desired information
 - Information is cached by the server to improve performance
 - LDAP was designed for use with relatively static data, not rapidly changing data
- ◆ Within MDS the EDG information providers 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



EDG GRIS/GIIS Hierarchy



datagrid



- ◆ There is a top level datagrid GIIS to which all of the country GIISs register

- ◆ Each country has a GIIS to which all of the site GIISs register

- ◆ Each Site has a Grid Information Index Server (GIIS) which acts as a single point of contact for all of the sites resources. The GRISs register with their site GIIS

- ◆ Information providers publish information to a local LDAP server known as a Grid Resource Information Server (GRIS)

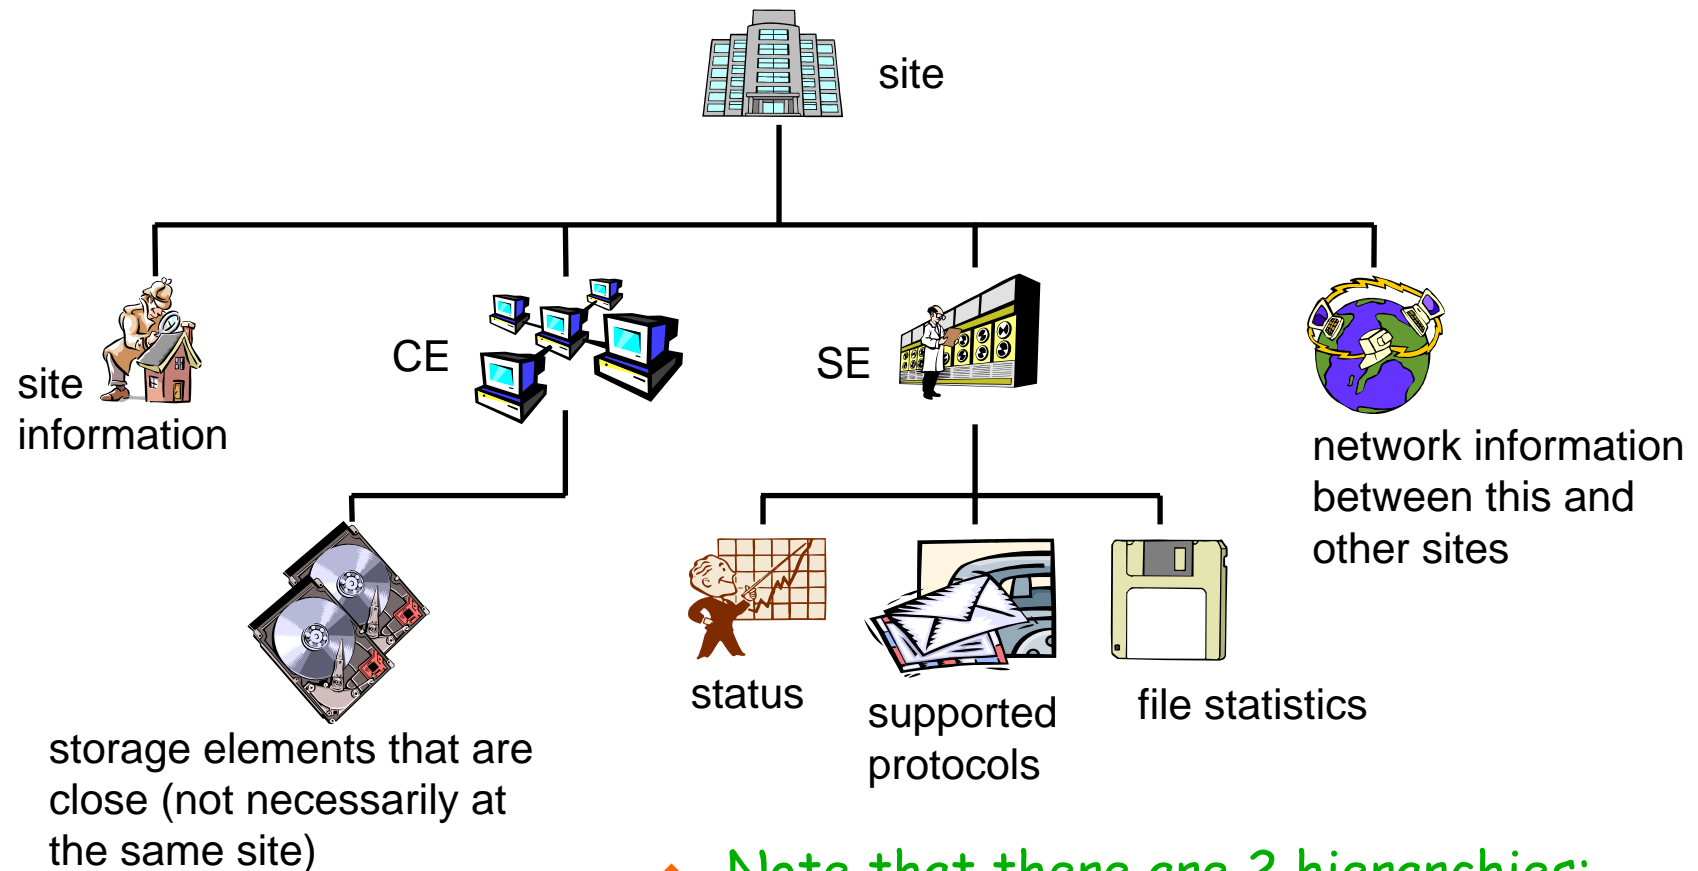


EDG Information Providers

- ◆ The EDG have produced information providers:
 - Site information
 - The Computing Element
 - The Storage Element
 - Network Monitoring
- ◆ All of the EDG data objects are dynamic, they have a time stamp and a time to live (used by the cache mechanism) associated with them



EDG Information Providers & the Directory Information Tree



- ◆ Note that there are 2 hierarchies:
 - The GIIS/GRIS structure
 - The DIT



Siteinfo



in=siteinfo,Mds-Vo-name=ral-dev,Mds-Vo-name=uk,o=Grid

objectClass: SiteInfo

objectClass: DataGridTop

objectClass: DynamicObject

siteName: RALDEV

sysAdminContact: grid.sysadmin@rl.ac.uk

userSupportContact: grid.support@rl.ac.uk

siteSecurityContact: grid.security@rl.ac.uk

dataGridVersion: 1.2

installationDate: 20020704142800Z



Storage Element



seId=dev02.hepgrid.clrc.ac.uk,Mds-Vo-name=ral-dev,Mds-Vo-name=uk,o=Grid

objectClass: StorageElement

objectClass: DataGridTop

objectClass: DynamicObject

SEId: dev02.hepgrid.clrc.ac.uk

CloseCE: dev01.hepgrid.clrc.ac.uk:2119/jobmanager-pbs-M

SEtypearchitecture: disk

SEsize: 13177

SEResourceContactString: grid.support@rl.ac.uk

SEvo: wpsix, :/flatfiles/05/wpsix



Storage Element Protocols



seProtocol=gridftp, seId=dev02.hepgrid.clrc.ac.uk,Mds-Vo-name=ral-dev,Mds-Vo-name=uk,o=Grid
objectClass: StorageElementProtocol
objectClass: DataGridTop
objectClass: DynamicObject
SEId: dev02.hepgrid.clrc.ac.uk
SEProtocol: gridftp
Port: 2811

seProtocol=rfio, seId=dev02.hepgrid.clrc.ac.uk,Mds-Vo-name=ral-dev,Mds-Vo-name=uk,o=Grid
objectClass: StorageElementProtocol
objectClass: DataGridTop
objectClass: DynamicObject
SEId: dev02.hepgrid.clrc.ac.uk
SEProtocol: rfio
Port: 3147

seProtocol=file, seId=dev02.hepgrid.clrc.ac.uk,Mds-Vo-name=ral-dev,Mds-Vo-name=uk,o=Grid
objectClass: StorageElementProtocol
objectClass: DataGridTop
objectClass: DynamicObject
SEId: dev02.hepgrid.clrc.ac.uk
SEProtocol: file



Storage Element Status



in=status,seId=dev02.hepgrid.clrc.ac.uk,Mds-Vo-name=ral-dev,Mds-Vo-name=uk,o=Grid

objectClass: StorageElementStatus

objectClass: DataGridTop

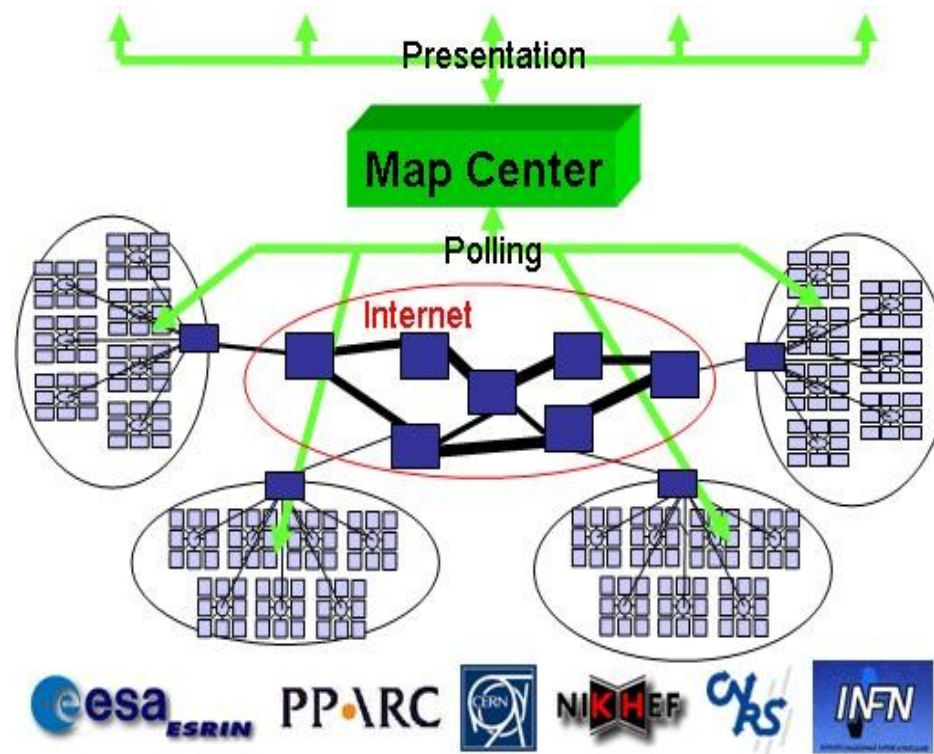
objectClass: DynamicObject

SEfreespace: 12031

SEId: dev02.hepgrid.clrc.ac.uk

Map Centre

- ◆ An easy way to browse the data available in the Information and Monitoring Service is to use Map Centre
- ◆ <http://ccwp7.in2p3.fr/mapcenter/datagrid-d/>





Querying the Information & Monitoring Service

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

```
$ldapsearch\
```

```
-x\
```

"simple" authentication

```
-H ldap://lxshare0225.cern.ch:2135\
```

uniform resource identifier

```
-b 'Mds-Vo-name=datagrid,o=grid\
```

base distinguished name for search

```
'objectclass=StorageElement'\
```

filter

```
seId SEsize \
```

attributes to be returned

```
-s base|one|sub
```

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



R-GMA

Relational - Grid Monitoring Architecture



- ◆ LDAP does not allow queries over different objects
 - I.e. you can only query based on attributes of an object (no "Joins")
- ◆ MDS is not designed for applications to publish their own data
 - It has relatively static descriptions of the data being published - the schema.
- ◆ R-GMA is a relational implementation of the Grid Monitoring Architecture (GMA) of the GGF
 - The relational model is very flexible and allows complex queries which make use of information in multiple objects
 - R-GMA provides a means for anyone to publish any information on the Grid - can also do the job of the current MDS
 - It is highly dynamic - with new Producers of information being noticed by existing Consumers



R-GMA

The Consumer Producer Model



Producer

Registry



Consumer

- ◆ Use the Grid Monitoring Architecture from Global Grid Forum
- ◆ A relational implementation
- ◆ Applied to both information and monitoring
- ◆ Creates impression that you have one RDBMS per Virtual Organization

→ Command flow
→ Information flow



Relational Approach

- ◆ **Not** a general distributed RDBMS system, but a way to use the relational model in a distributed environment **where ACID properties are not generally important.**
- ◆ **Producers**
 - announce: SQL "CREATE TABLE"
 - publish: SQL "INSERT"
- ◆ **Consumers**
 - collect: SQL "SELECT"
- ◆ The mediator is a component within the Consumer which locates one or more Producers and combines the information as necessary



Summary

- ◆ The current information system is *Globus MDS*
 - This uses LDAP as its underlying data interface
 - There are a number of EDG information providers
- ◆ The new Information and Monitoring Service is *R-GMA*
 - Early version in 1.2
 - It can support complex SQL queries
 - It can return a single result set or stream data
 - Grid applications can publish data



The End

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