Grid Information System

Paolo Veronesi (paolo.veronesi@cnaf.infn.it)

IV Scuola utenti INFN della Grid,

CNAF, Bologna

21-24/02/2011

Outline

- What is the Information System
  - BDII & LDAP
- Model: the GLUE Schema
  - Overview
  - Core entities examples
- Information System Architecture
  - Top BDII and Site BDII
  - GIP
- GOC-DB
- HANDS ON

Grid Middleware

- When using a PC or workstation you
  - Login with a username and password (“Authentication”)
  - Use rights given to you (“Authorisation”)
  - Run jobs
  - Manage files: create them, read/write, list directories
- Components are linked by a bus
- Operating system
- One admin domain

- **When using a Grid you**
  - **Login with digital credentials (“Authentication”)**
  - **Use rights given you (“Authorisation”)**
  - **Run jobs**
  - **Manage files: create them, read/write, list directories**
- **Services are linked by the Internet**
- **Middleware**
- **Many admin domains**

The GRID architecture: general view

- The GRID is a collection of geographically distributed resources
- GRID users:
  - Organized in Virtual Organizations
  - Need to run programs without the need to know
    - Where to run a job
    - Where to get the input data from
    - Where to store the output data to
- The GRID consists of
  - An Authorization and Authentication System
  - An Information System
  - A Workload Management System
  - A Data Management System
  - An Accounting System
  - Various monitoring services
  - Various installation services

gLite Components

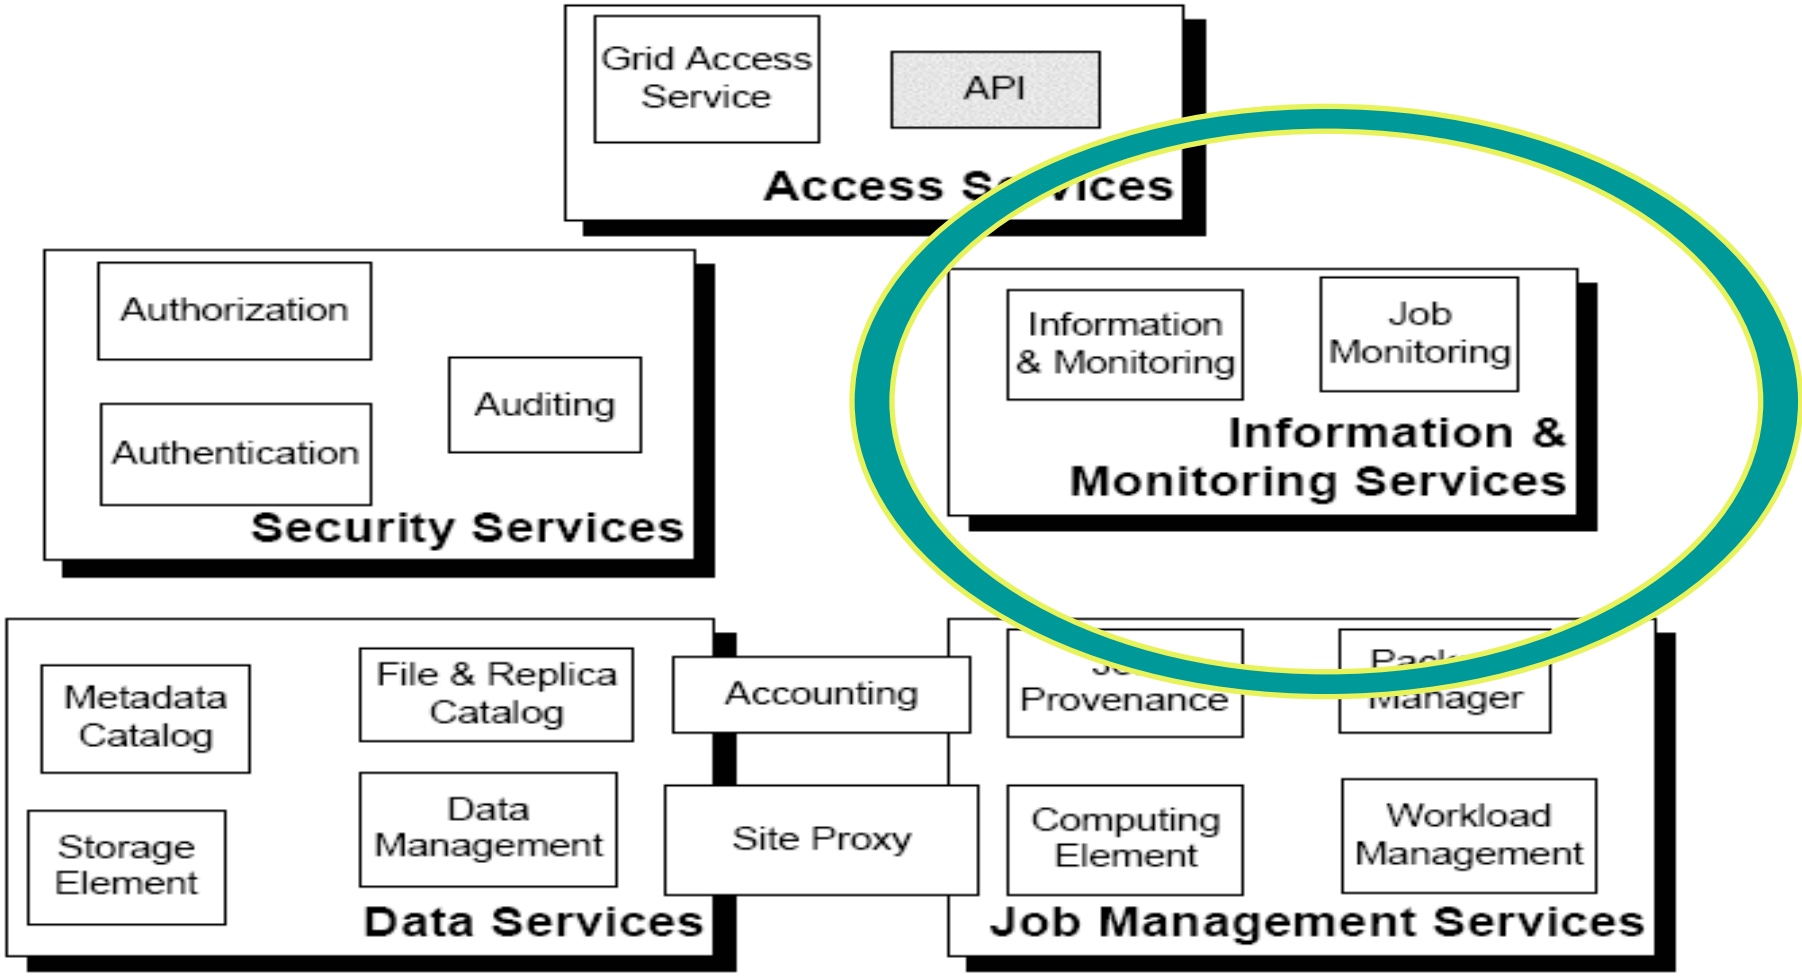

gLite Components

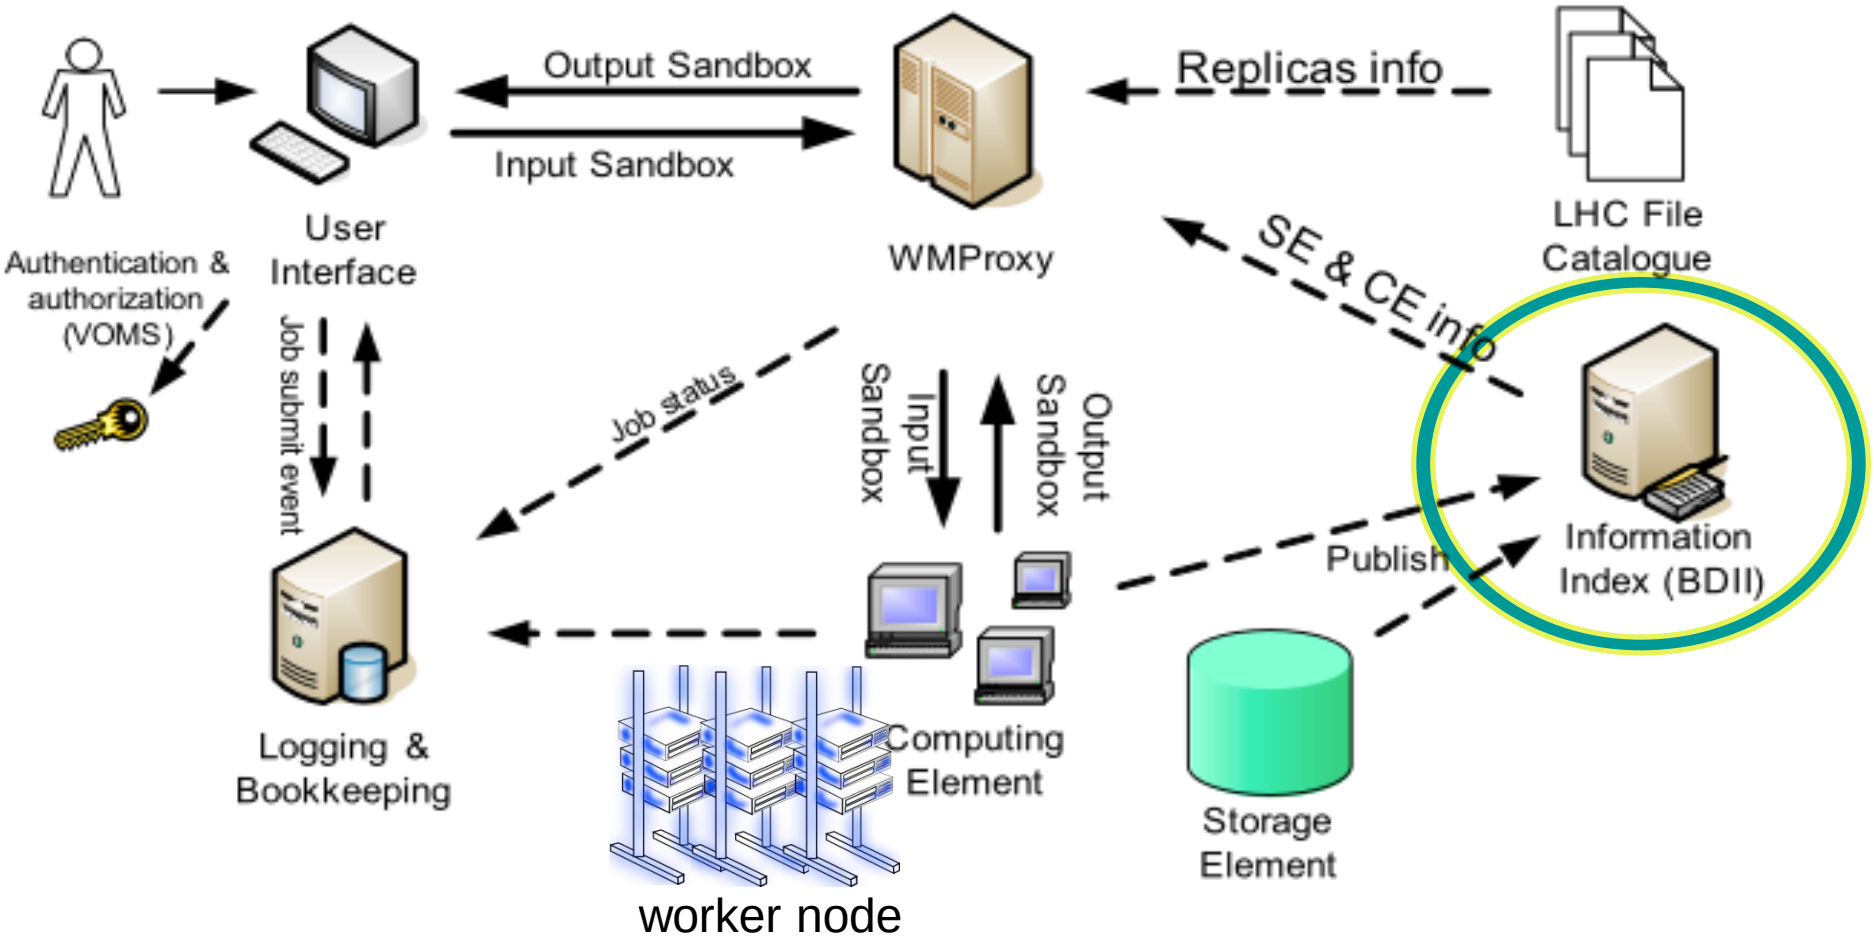

Information System

- **What is?**
  - System to collect information on the state of resources
- **Why?**
  - To discover resources of the grid and their nature
  - To have useful data that helps who is in charge of managing the workload to do it more efficiently.
  - To check for health status of resources.
- **How?**
  - Monitoring state of resources locally and publishing right information on the information system.
  - Adopting a data model that MUST be well known to all components that want to access monitored information
  - Using different approaches that we are going to investigate in next slides

Design of Information Systems

- **About Measures**
  - Measures SHOULD be sensitive to the aim the users want to achieve.
  - Measures SHOULD be enough accurate to be considered valid.
  - Rate of taking measures MUST be adequate to be used.
- **About the gathering of Information**
  - How and when collected info should be published?
  - Where should collected info be stored?
  - How long should this info be maintained in the storage?
- **Querying the Information System**
  - Where should queries be sent to have a response?
  - What syntax and protocols have to be adopted to make queries?
  - What is the adopted data model to describe resources?
- **Security**
  - Who is allowed to execute queries against the IS and what type of queries is he allowed to do?
  - Management of user rights and credentials.

Information System

- **Information System** (IS) provides information about the Grid resources and their status.
- Much of the data published to the IS conforms to the **GLUE Schema**.
- The most common implementation is called **BDII**
- The BDII implements the GLUE Schema using **LDAP**

BDII

- **Berkley Database Information Index (BDII)**
- It has been adopted in the gLite middleware as the Information System technology.
- The BDII consists of a standard **LDAP** database which is updated by an external process
- Within the BDII, one finds elements that have attributes and links to other elements.

LDAP

- **Lightweight Directory Access Protocol**
- The protocol used to **query** the information system.
- It is a **lightweight** protocol for **accessing directory** services optimized for reading, browsing and searching.
- It is the internal protocol used by the EGEE/LCG services to share information
- Various graphical LDAP browsers available.

LDAP Data Model (1/2)

- The LDAP information model is based on **entries**
- Entries are collections of **attributes** → each entry's attributes have a type and one or more values
- Each entry has **Distinguished Names** (DN) → a sequence of attribute/value pairs
  - DN uniquely identifies an entry
- Based on their DNs, entries can be arranged into a hierarchical tree-like structure, called a **Directory Information Tree** (DIT).

LDAP Data Model (2/2)

The information is imported and exported by **LDIF files** (LDAP Data Interchange Format) which follow such structure

```
dn: <distinguished name>
objectclass:<objectclassname>
<attributetype>:<attributevalue>
<attributetype>:<attributevalue>

dn: <distinguished name>
objectclass:<objectclassname>
<attributetype>:<attributevalue>
<attributetype>:<attributevalue>
```

This is an **entry**; collection of **attributes**. It's defined by a unique DN

Objectclass: special attributes can be defined for each entry

a) Defines the tree structure of a certain entry

b) Can be used to filter the entries containing this objectclass

White space to separate entries from each other

The types and objectclass names follow the **Glue Schema**

Directory Information Tree

- LDAP 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=Gene
   va,c=Switzerland,o=grid"
```

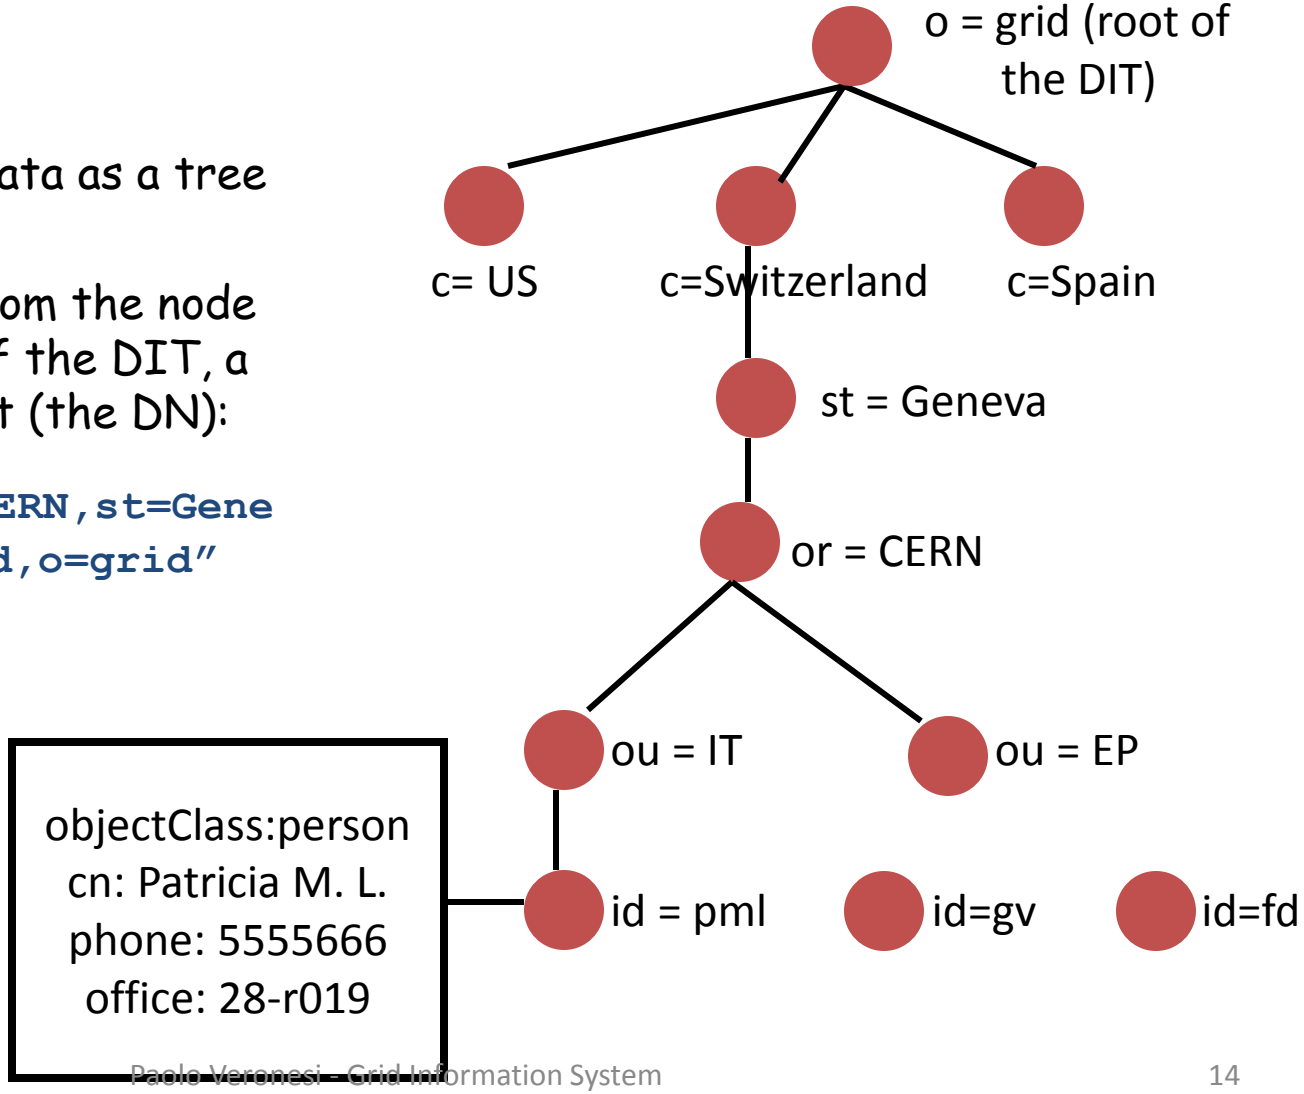

Why a schema?

- A Grid consists of many sites with a wide variety of resources
- Users, applications and middleware need to know what resources are available and what their properties are:
  - What Workload Managers are available to CMS?
  - Find a Computing Service running SL5 with > 3 Gb memory
  - Find a Storage Service with 20 TB of free space
- Grid and VO management and operations staff need an overview of the state of the Grid:
  - How many jobs are running in the UK?
  - How much disk space has ATLAS used?
  - What is the total installed CPU power available to LCG?

GLUE Schema

- **Grid Laboratory Uniform Environment schema**.
- The GLUE information schema provides a standardized description of a Grid computing system, to enable resources and services to be presented to users and external services in a uniform way
- The information is transported via an information system, but the schema is logically independent of it.
- The entities of the GLUE Schema are organised hierarchically
- The schema is split into a number of high-level pieces.
  - **Site** defines some information about an entire Grid site.
  - **Service** provides a general abstraction for any Grid service.
  - CE and SE provide detailed information about **Computing** and **Storage Elements**, as these are the most important components of the Grid.
    - the schema also describes a Cluster which represents the hardware (Worker Nodes) which can be accessed via a CE.

GLUE Schema (1.3) overview

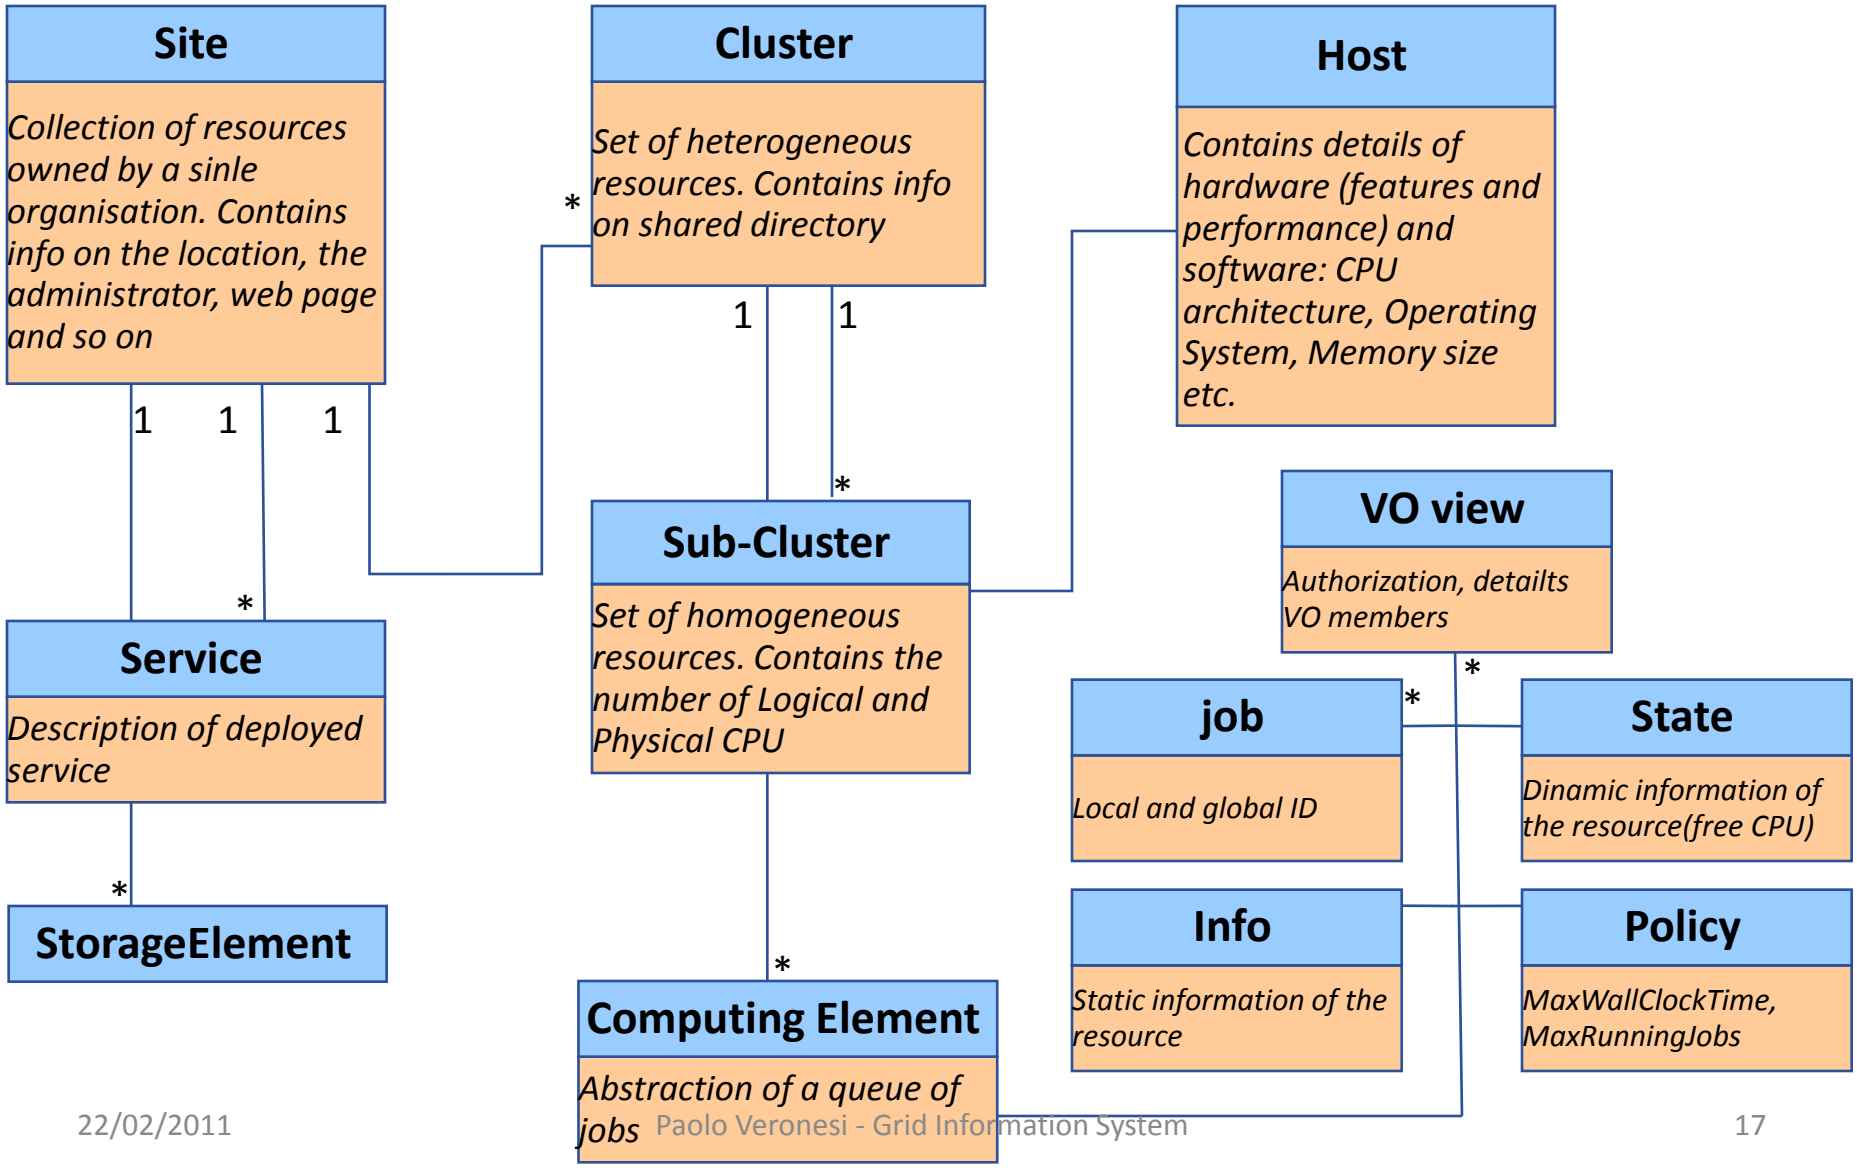

GLUE Schema example: Site information

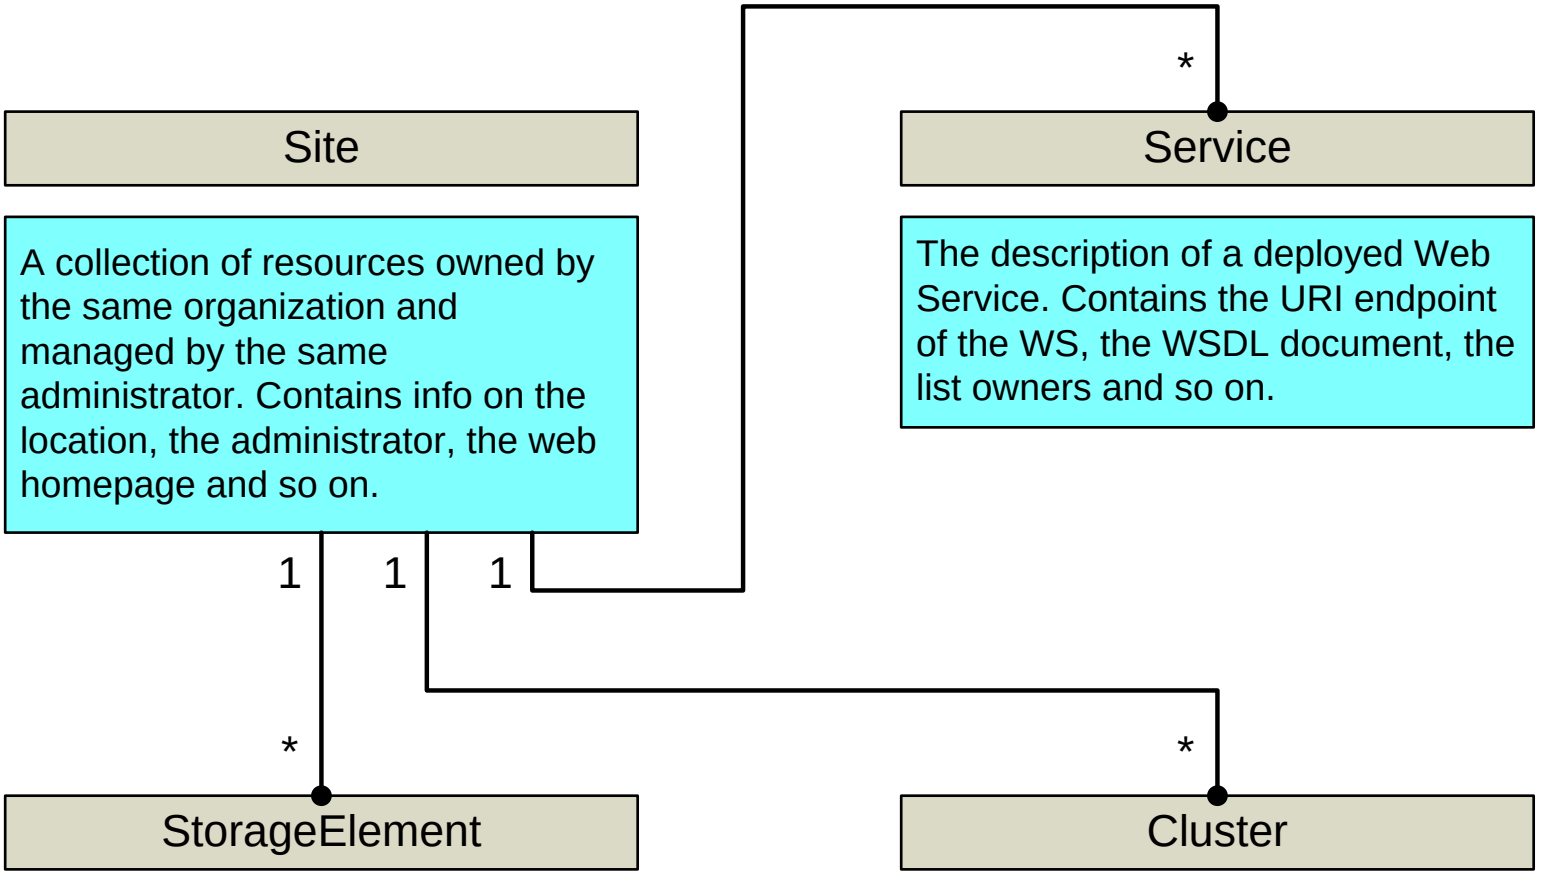

GLUE Schema example: Site information

- **This provides information about a grid site as a whole. GlueSite object:**
- **GlueSiteUniqueID**: This is the unique name for the site, as defined in the GOC database.
- **GlueSiteDescription**: A general description of the site.
- **GlueSiteEmailContact**: A mailto: URL defining a general email contact address for the site.
- **GlueSiteLocation**: The geographical location of the site as a string, normally in the form City, State, Country.
- **GlueSiteLatitude, GlueSiteLongitude**: The map reference for the site, in degrees. The resolution usually locates the site to within 100m.
- **GlueSiteWeb**: A URL pointing to a web page relating to the site.
- **GlueSiteSponsor**: The organization(s) providing funding for the site.
- **GlueSiteOtherInfo**: A multivalued string which may contain any further information the site considers useful.

GLUE Schema example: Service information

- **This provides a general abstraction of a Grid service. GlueService object:**
- **GlueServiceType:** The service type, taken from a defined list which can be found on the Glue website.
- **GlueServiceVersion:** The version of the service, in the form major.minor.patch.
- **GlueServiceEndpoint:** The network endpoint for the service.
- **GlueServiceStatus:** The status of the service (one of OK, Warning, Critical, Unknown, Other).
- **GlueServiceStatusInfo:** A textual explanation of the Status.
- **GlueServiceWSDL:** For web services this is a URL pointing to the WSDL definition of the service.
- **GlueServiceSemantics:** This is a URL which would typically point to a web page explaining how to use the service.
- **GlueServiceOwner:** The service owner, if any; typically a VO name.
- **AccessControlBaseRule:** A set of ACLs defining who is allowed access to the service.

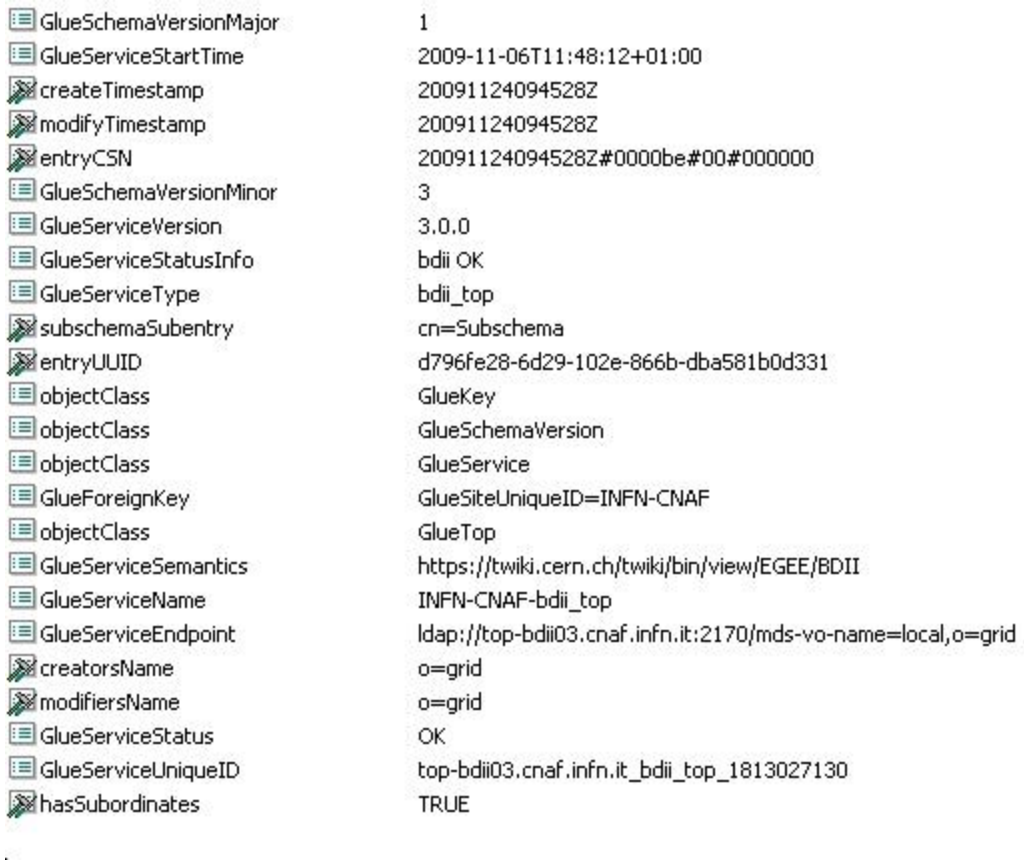

Cluster Element

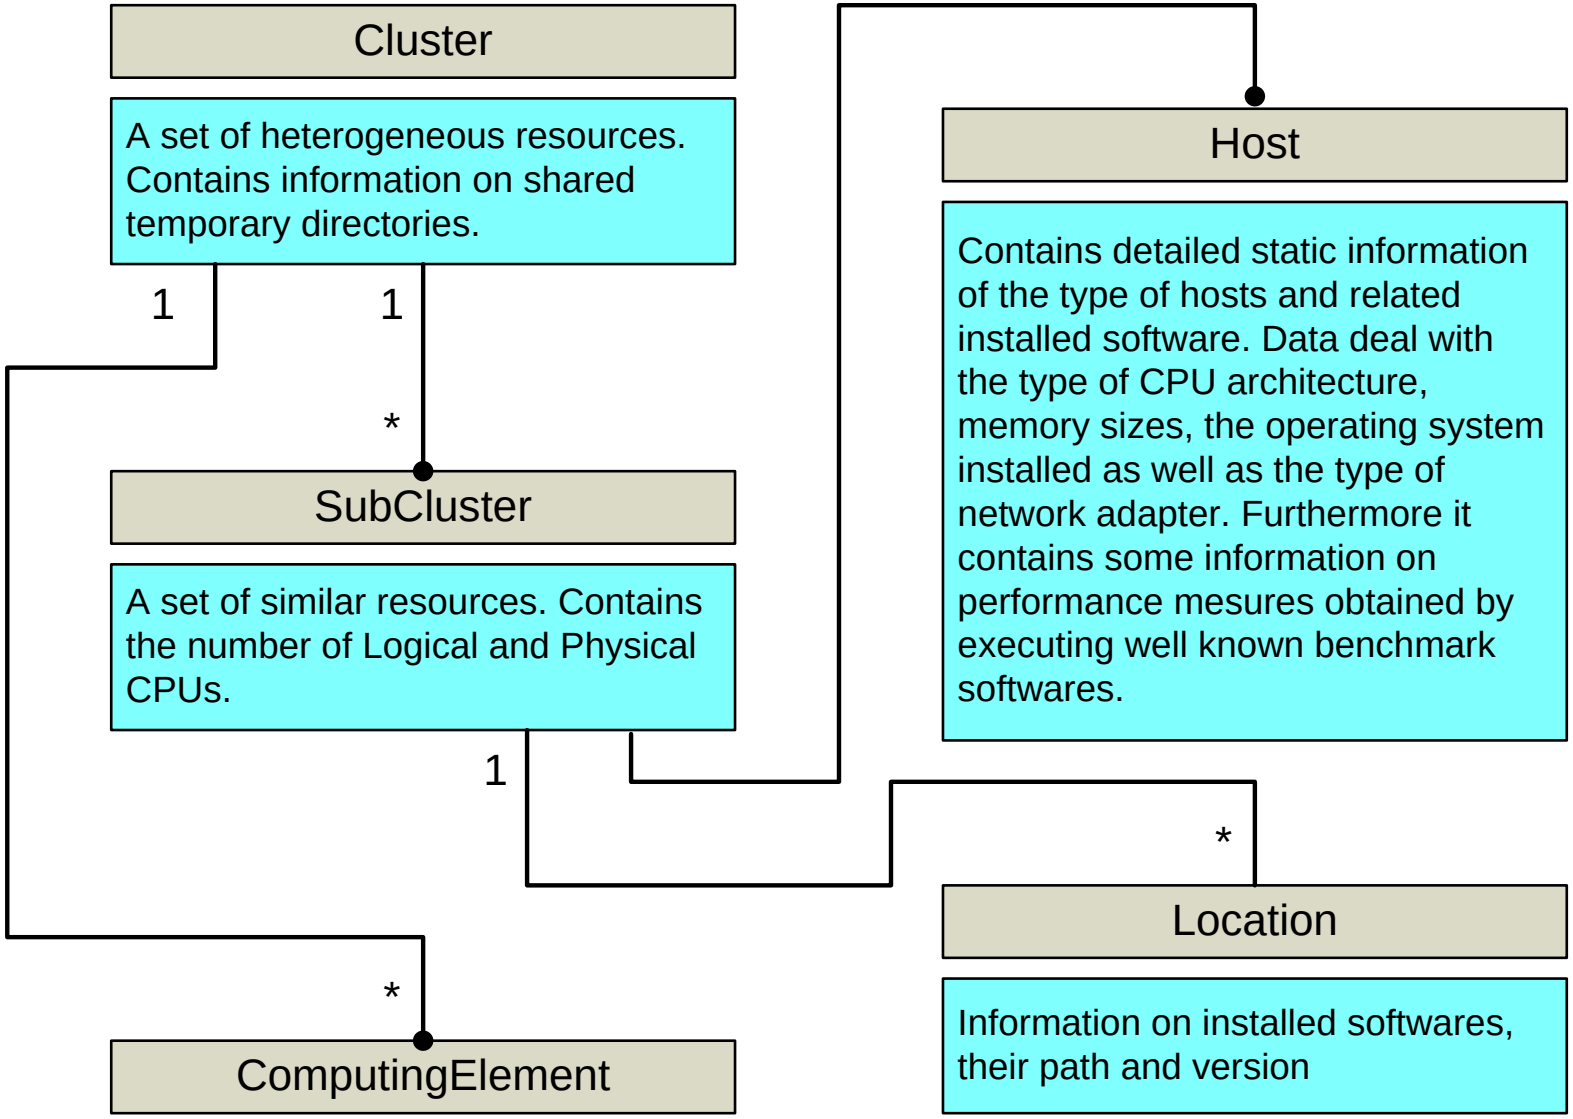

GLUE: cluster and subcluster

GlueClusterName: infn-ce-01.ct.trigrid.it
GlueClusterService: infn-ce-01.ct.trigrid.it:2119/jobmanager-lcglsf-short
GlueClusterService: infn-ce-01.ct.trigrid.it:2119/jobmanager-lcglsf-long
GlueClusterService: infn-ce-01.ct.trigrid.it:2119/jobmanager-lcglsf-infinite
GlueClusterService: infn-ce-01.ct.trigrid.it:2119/jobmanager-lcglsf-cert
GlueClusterService: infn-ce-01.ct.trigrid.it:2119/jobmanager-lcglsf-cometa
GlueClusterService: infn-ce-01.ct.trigrid.it:2119/jobmanager-lcglsf-inaf
GlueClusterService: infn-ce-01.ct.trigrid.it:2119/jobmanager-lcglsf-alice
GlueClusterService: infn-ce-01.ct.trigrid.it:2119/jobmanager-lcglsf-cometa
[..]
GlueSubClusterPhysicalCPUs: 4
GlueSubClusterLogicalCPUs: 4
GlueSubClusterTmpDir: /tmp
GlueSubClusterWNTmpDir: /tmp

GLUE: Host

GlueHostApplicationSoftwareRunTimeEnvironment: GLITE-3_0_0
GlueHostApplicationSoftwareRunTimeEnvironment: INFN-CATANIA
GlueHostApplicationSoftwareRunTimeEnvironment: MPICH
[..]
GlueHostArchitectureSMPSize: 4
GlueHostBenchmarkSF00: 1937
GlueHostBenchmarkSI00: 1483
GlueHostMainMemoryRAMSize: 4096
GlueHostMainMemoryVirtualSize: 8192
GlueHostNetworkAdapterInboundIP: TRUE
GlueHostNetworkAdapterOutboundIP: TRUE
GlueHostOperatingSystemName: Scientific Linux CERN
GlueHostOperatingSystemRelease: 3.0.6
GlueHostOperatingSystemVersion: SLC
GlueHostProcessorClockSpeed: 2392
GlueHostProcessorModel: Dual Core Opteron 280
GlueHostProcessorVendor: AMD

Computing Element

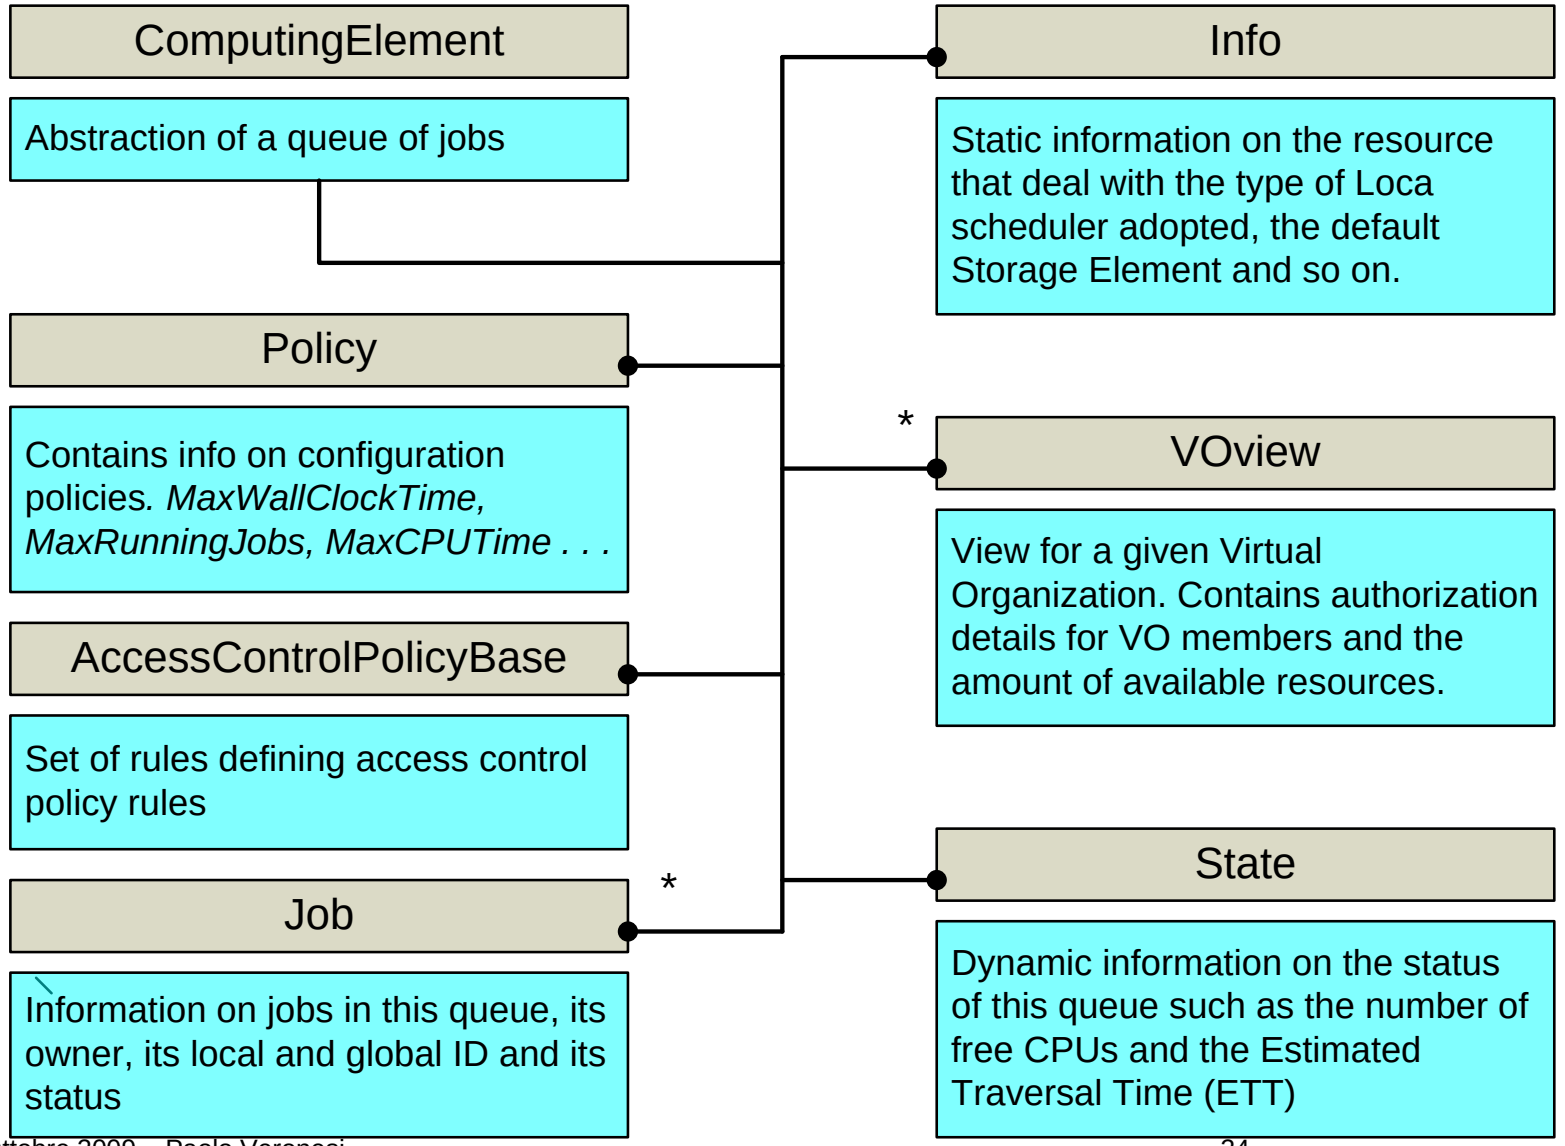

GLUE: Host

GlueCEName: cometa
GlueCEUniqueID: infn-ce-01.ct.trigrid.it:2119/jobmanager-lcglsf-cometa
GlueCEInfoGatekeeperPort: 2119
GlueCEInfoHostName: infn-ce-01.ct.trigrid.it
GlueCEInfoLRMSType: lsf
GlueCEInfoLRMSVersion: 6.1
GlueCEInfoTotalCPUs: 98
GlueCEInfoJobManager: lcglsf
GlueCEInfoContactString: infn-ce-01.ct.trigrid.it:2119/jobmanager-lcglsf-cometa
GlueCEInfoApplicationDir: /opt/exp_soft
GlueCEInfoDataDir: unset
GlueCEInfoDefaultSE: infn-se-01.ct.trigrid.it
GlueCEStateEstimatedResponseTime: 61713
GlueCEStateFreeCPUs: 26
GlueCEStateRunningJobs: 70
GlueCEStateStatus: Production
GlueCEStateTotalJobs: 70
GlueCEStateWaitingJobs: 0
GlueCEStateWorstResponseTime: 123427
GlueCEStateFreeJobSlots: 26
GlueCEPolicyMaxCPUTime: 2880
GlueCEPolicyMaxRunningJobs: 98
GlueCEPolicyMaxTotalJobs: 0
GlueCEPolicyMaxWallClockTime: 2880
GlueCEPolicyPriority: -10
GlueCEPolicyAssignedJobSlots: 98
GlueCEAccessControlBaseRule: VO:cometa

GLUE Schema example: storage element

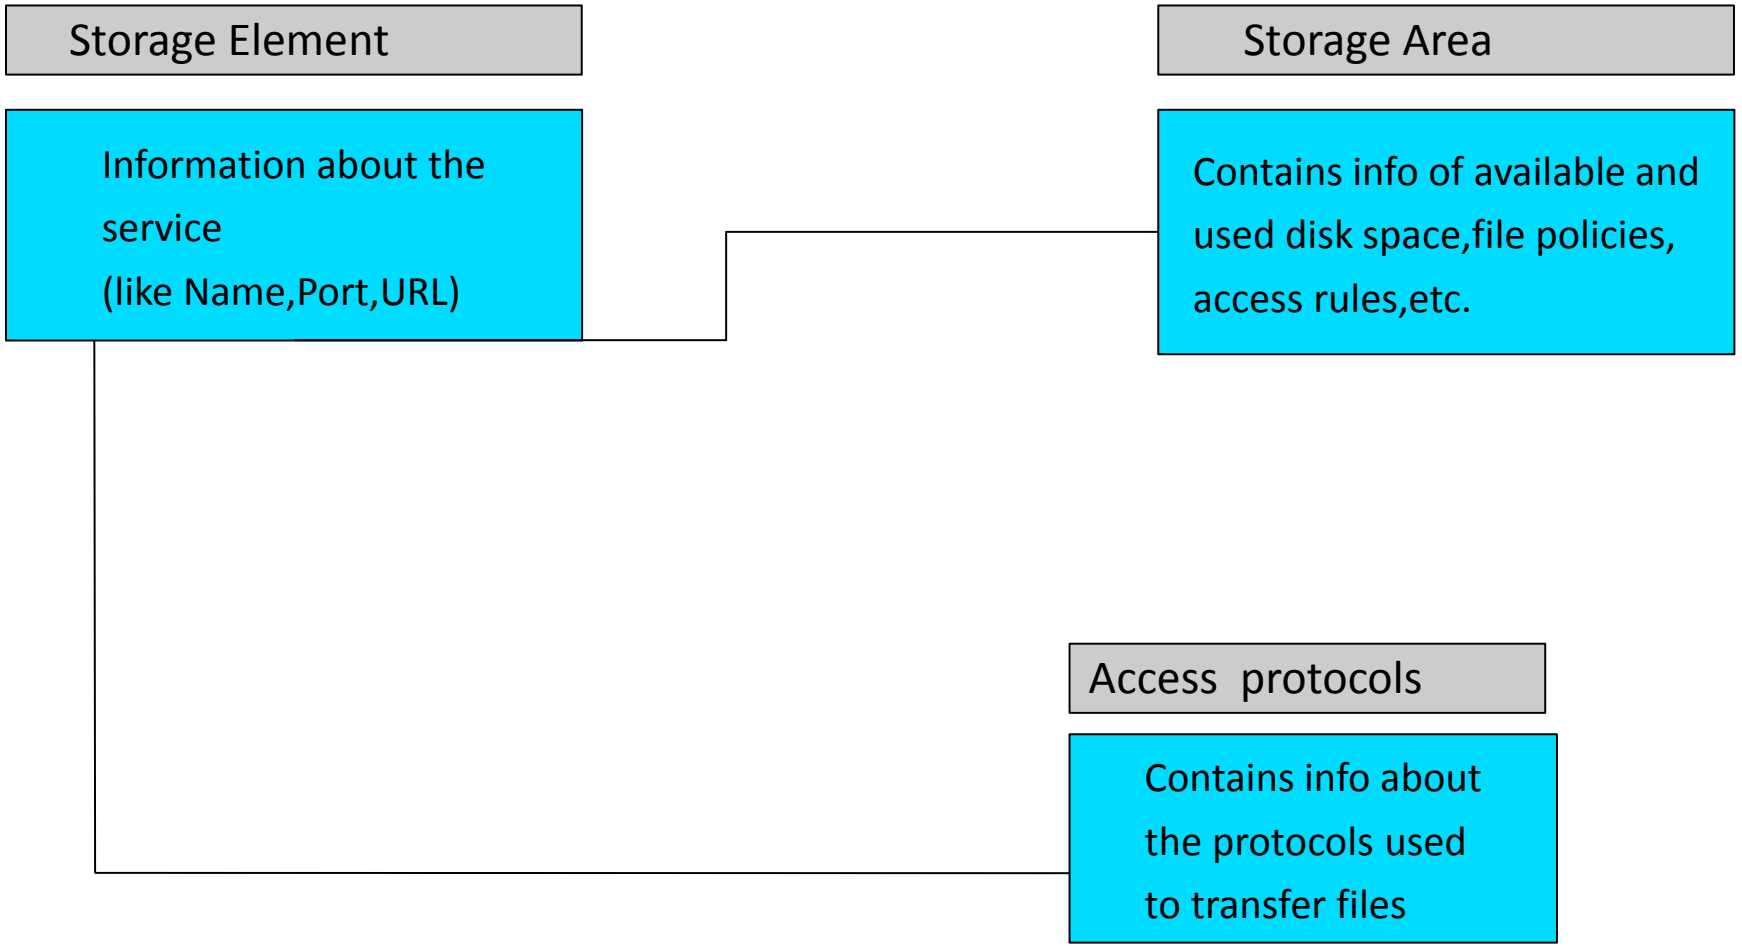

GLUE Schema example: storage element

GlueSE object

- **GlueSEUniqueID:** A globally unique identifier for the SE. See note on unique identifiers
- **GlueSEName:** A human readable name for the cluster. Typically the host name is used
- **GlueSEType** :The SE Type. Deprecated
- **GlueSEArchitecture:** Underlying architectural. *disk*, *tape*, *multi-disk*, *other*
- **GlueSEImplementationName** :The name of the storage system implementation
- **GlueSEImplementationVersion:** The version of the storage system implementation
- **GlueSEStatus** :The status of the SE. *Queueing*, *Production*, *Closed* or *Draining*
- **GlueSETotalOnlineSize** :Total size of online (disk) storage space in GB
- **GlueSETotalNearlineSize** :Total size of nearline (tape) storage space in GB
- **GlueSEUsedOnlineSize:** Used size of online storage space in GB
- **GlueSEUsedNearlineSize** :Used size of nearline storage space in GB

GLUE Schema example: storage element

- Auxiliary concepts are the **Control Protocol** and **Access Protocol**.
  - **Control Protocol** relates to the protocol used to manage the SE; in WLCG/EGEE this currently means some version of the SRM protocol.
  - **Access Protocol** specifies the protocols used for data transfer; a typical SE will support several of these (e.g. gsiftp or rfio).

GlueSEAccessProtocol object

- **GlueSEAccessProtocolType**: The protocol type, e.g. gsiftp or rfio.
- **GlueSEAccessProtocolVersion**: The protocol version.
- **GlueSEAccessProtocolEndpoint**: A URL specifying the endpoint for this protocol. Note that with an SRM the endpoint is normally obtained dynamically.
- **GlueSEAccessProtocolCapability**: A multivalued string allowing arbitrary capabilities to be advertised.
- **GlueSEAccessProtocolMaxStreams**: The maximum number of data streams allowed for a single transfer using this protocol.

Information System Architecture (1/2)

- Information is collected on the leaves of a hierarchical tree and travels towards the root
- Clients can query the hierarchical tree at every level
- The higher the level against which queries are made, the older is the obtained information
- FCR "Freedom of Choice for Resources" is a mechanism to allow Virtual Organizations to mask sites or services from their users if they are known not to be working correctly.

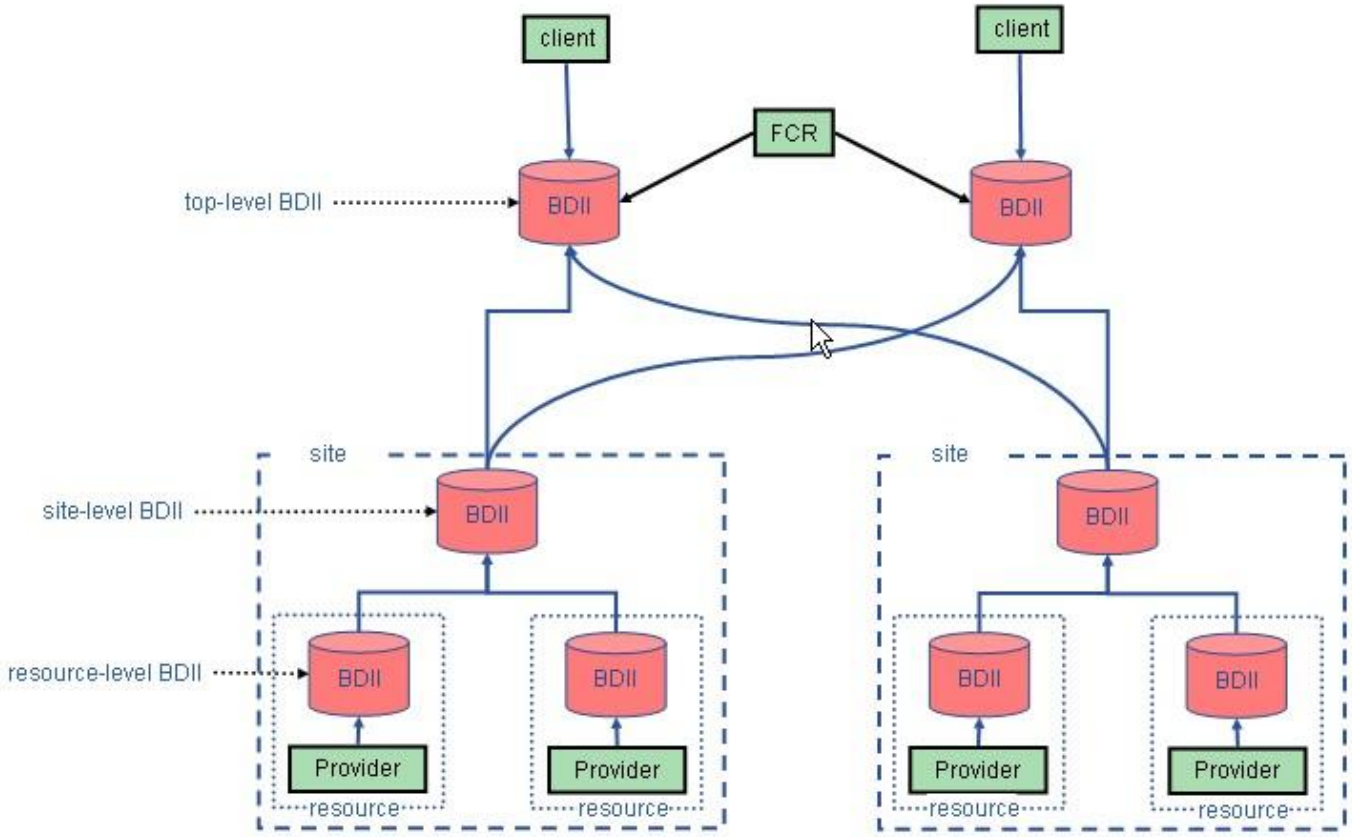

Information System Architecture (2/2)

- Every node (except UI and WNs) has a BDII service in order to publish its information.
- A node in every site collects all site BDIIs and publishes them using a **site BDII** → The **top BDII** collects all site BDIIs.
- User can run a **set of commands** to query the top BDII.
- The difference between resource level , site level and top level is just information content and scope.

IS way of working

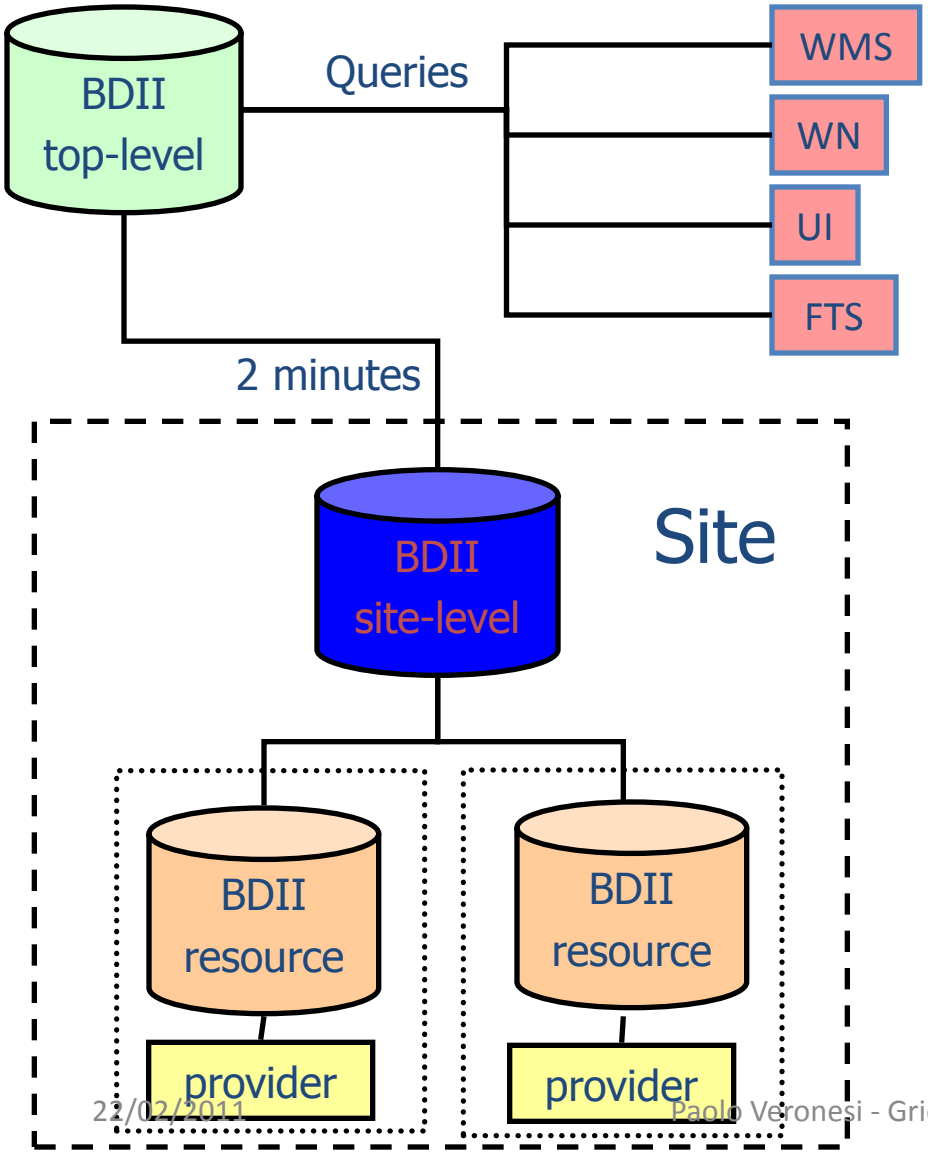

- One BDII for each resource
- One BDII for each site collecting info from below BDII systems
- Every two minutes a cron-job runs a script and collects info from a list of BDII sites
- The list of site BDII is placed in the configuration file of the top BDII

Way of collecting info

- Pull model (higher level servers periodically query lower level servers)
- LDAP query model

Resource-level BDIIs

- Grid Site Level: Lower level
- Computing and storage resources at a site run a piece of software called an Information Provider.
- Information Provider generates the relevant information about the resource (both static, like the type of SE, and dynamic, like the used space in an SE).
- This information is published via a server called a resource-level BDII, which normally runs on the resource itself.

The Site BDII

- Grid Site Level: Medium level
- The site-level BDII is used to store and publish data from all the resource-level BDIIs at a site.

The top-level BDII

- Grid Core Service level: Higher level
- top-level BDII is used as the top of the hierarchy.
- BDIIs at this level are configured to read from a specific set of sites, which effectively defines a view of the overall Grid resources.
- These BDIIs act as a cache by storing information about the Grid status in their database.
- The BDIIs therefore contain all the available information about the Grid sites they look at.
- Nevertheless, it is always possible to get information about specific resources by directly contacting the site- or resource-level BDIIs.

Information Providers

- The information presented in the schema is produced by programs known as **Information Providers**.
- These are structured as a framework, **Generic Information Provider (GIP)**, with a set of plugins for specific parts of the schema (e.g. to cope with different batch systems or SE technologies).
- These divide into **static and dynamic providers**.
  - **Dynamic information** (e.g. the number of running jobs) changes on a short timescale and therefore has to be collected from some other system (e.g. a batch system) every time the provider is run.
  - **Static information** (e.g. the maximum CPU time for a queue) changes infrequently, and the providers therefore read it from a configuration file which is updated only when the system is reconfigured.

GOC Database

- For a list of all sites and all resources present, refer to the GOC database
- The information system is *bootstrapped* from the information in the **Grid Operations Center Database** (GOC DB).
- When a site registers, it enters the URL for the *site level* BDII into the GOC DB.
- The GOC DB generates a list of LDAP URLs for all the sites in the grid and this is downloaded by the information provider running on the *top level* BDII.
- These URLs are then used to query all the *site level* BDII and the result is used to populate the *top level* BDII.
- Site managers can insert in GOC database the contact address of their BDII as well as other useful information about the site.

INFN –CNAF site:
https://goc.gridops.org/portal/index.php?Page_Type=View_Object&object_id=72&grid_id=0

Monitoring

- The ability to monitor resource related parameters in such a heterogeneous and complex system as the Grid becomes fundamental.
- A monitoring system implies the existence of a central repository of operational information (in WLCG/EGEE, the GOCDB).
- The monitoring system should be able to collect data from the resources in the system, in order to analyze the usage, behavior and performance of the Grid, detect and notify fault conditions, contract violations and userdefined events.
- Several different monitoring tools are in use, including general-purpose monitoring tools and Grid specific systems

Recommended Query Tools

- User can run a set of commands to query the top BDII.
- Two utilities are provided to allow users to query top-level BDIIs *without having to know the details of LDAP syntax*
- They are simply wrappers for the corresponding LDAP queries:
  - **lcg-infosites**
  - **lcg-info**

lcg-infosites

- The lcg-infosites command can be used to obtain VO-specific information on existing grid resources.

```
lcg-infosites --vo voname -[v] -f [site name]
   [option(s)] [-h| --help] [--is BDII]
```

- For example, to list the Storage Elements (SEs) available to the lhcb VO at the UNINA-EGEE site, one could issue the following command:

```
lcg-infosites --vo lhcb -f UNINA-EGEE se

Avail Space(Kb) Used Space(Kb)  Type    SEs
----------------------------------------------------------
16804291653     236974264       n.a     se.scope.unina.it
```

lcg-info

- The lcg-info command can be used to list either CEs or SEs satisfying a given set of conditions on their attributes, and to print, for each of them, the values of a given set of attributes.
- The information is taken from the BDII specified by the LCG GFAL INFOSYS environment variable or in the command line.

```
lcg-info [--list-ce | --list-se] [--query
  <query>] [--attrs <attrs>]
```

ldapsearch vs lcg-infosites

- LDAP queries can be quite complex, and can potentially put a heavy load on a BDII and return lots of data, so caution should be used.

```
ldapsearch -x -H ldap://gridit-bdii-
  01.cnaf.infn.it:2170 -b
  "GlueSEUniqueID=se.scope.unina.it,mds-vo-
  name=UNINA-EGEE,mds-vo-name=local,o=grid"
  GlueSALocalID=lhcb GlueSAFreeOnlineSize
  GlueSAUsedOnlineSize
```

Conclusion: Who and how use the IS

- Users
  - Retrieve information about resources:
    - Where can I run my job?
    - Where can I store my files?
    - Which software packages are available on a give CE?
- Site Managers
  - Publish information about the resources and services they provide
- Middleware
  - WMS: matching job requirements and allocating the resources
  - Monitoring Services: retrieving information about the status and availability of resources

More information

- gLite User Guide

  https://edms.cern.ch/file/722398/1.2/gLite-3-UserGuide.pdf

- Use cases

  http://www.eu-egee.org/fileadmin/documents/UseCases/Index.html

- gLite Services

  http://glite.web.cern.ch/glite/documentation/

- GLUE Schema Documentation

  https://twiki.cern.ch/twiki//bin/view/EGEE/GlueUse

The End?

Questions ?


hands
on

Hands on

- Advanced Information System Queries:
  - ldapsearch
  - lcg-infosites
  - lcg-info
- LDAP browser tool
  - Web tool
  - LDAP Browser 2.6 (Windows)
  - Luma (part of KDE)

Advanced Information System Queries: ldapsearch

Getting the Unique ID for all Ces

This is an elementary query. It would work without any filter as the attribute name only appears in the GlueCE object, but would be somewhat less efficient. Similarly the attribute name could be omitted as it gets selected anyway with grep - this is done to remove DNs and blank lines from the LDIF output. An extra step could be to use cut -d: -f2- to strip off the attribute name.

```
ldapsearch -x -LLL -H ldap://egee-bdii.cnaf.infn.it:2170/ -b mds-vo-name=local,o=grid \
objectclass=GlueCE GlueCEUniqueID | grep GlueCEUniqueID:
```

Advanced Information System Queries: ldapsearch

Finding out which CE job managers are in use at INFN sites

The query here is slightly more complex, in that it uses a wildcard filter to select only CEs with hostnames in INFN domains. There is also a bit more post-processing to present the result in a useful form. This also illustrates the lack of case-sensitivity in the filter. However, remember that grep is case-sensitive by default.

```
ldapsearch -x -LLL -H ldap://egee-bdii.cnaf.infn.it:2170/ -b mds-vo-name=local,o=grid \
'(&(objectclass=gluece)(glueceinfohostname=*.infn.it))' \
glueceinfojobmanager | grep GlueCEInfoJobManager: \
| cut -d: -f2 | sort | uniq -c
```

Advanced Information System Queries: ldapsearch

How many SEs support my VO?

This uses the fact that the **GlueSA object**, which contains the VO assignment information, has an attribute **GlueChunkKey** to provide a link to the parent SE. There could potentially be several matching SAs per SE, but the "sort | uniq" ensures that only one is counted. The access control format is expanding, but this gets all possible variants currently in use.

For variety, this query uses a different BDII, the alternative format to specify the host and port, and queries from o=grid. It also doesn't explicitly select an attribute name or use the -LLL option as the output gets filtered by the grep anyway. The objectclass is omitted as the attribute filter implies it.

```
# export MYVO=put_here_your_vo_name
export MYVO=gridit

ldapsearch -x -h egee-bdii.cnaf.infn.it -p 2170 -b o=grid \
'(|(GlueSAAccessControlBaseRule=$MYVO)(GlueSAAccessControlBaseRule=VO:$MYVO)(GlueSAAccessControlBaseRule=VOMS:/$MYVO/*))' \
| grep 'GlueChunkKey: GlueSEUniqueID=' | sort | uniq | wc -l
```

Advanced Information System Queries: ldapsearch

Which sites have SEs which support the gsiftp protocol?

This example illustrates a **double query**: first find the GlueSEAccessProtocol objects which match a condition, extract the GlueChunkKey which contains a reference to the parent GlueSE object, and then query that for the GlueForeignKeywhich contains the site name. The post-processing here just strips out the attribute name and removes any duplication. **Efficiency is more of an issue here** because the second query is executed multiple times (once per matching SE). Adding an explicit restriction objectclass=GlueSE reduces the time for the query by around a factor 2, from 20 seconds to 10.

```
for i in `ldapsearch -x -H ldap://egee-bdii.cnaf.infn.it:2170 -b o=grid \
GlueSEAccessProtocolType=gsiftp | grep GlueChunkKey: | cut -d= -f2 \
| sort | uniq`; do ldapsearch -x -H ldap://egee-bdii.cnaf.infn.it:2170 \
 -b o=grid "GlueSEUniqueID=$i*" GlueForeignKey; done \
 | grep GlueSiteUniqueID | cut -d= -f2 | sort | uniq
```

Advanced Information System Queries: ldapsearch

Which sites have SEs which support the gsiftp protocol?

Here is another way to get the same result, using perl to join any split lines and sed to process the DN and use the result in the second query. This is somewhat more efficient (it takes about 7 seconds), **but harder to write**.

```
for dn in `ldapsearch -h egee-bdii.cnaf.infn.it -p 2170 -x -b o=grid \
GlueSEAccessProtocolType=gsiftp | perl -p00e 's/\r?\n //g' \
| sed -ne 's/^dn: [^,]*,[^,]*,\(.*\)/\1/p' | sort -u`; do \
ldapsearch -h egee-bdii.cnaf.infn.it -p 2170 -x -b $dn \
'(objectClass=GlueSite)' GlueSiteUniqueID \
| perl -p00e 's/\r?\n //g' | sed -ne 's/^GlueSiteUniqueID: \(.*\)/\1/p'; done
```

Advanced Information System Queries: lcg-infosites and lcg-info

Recommended Query Tools

For convenience, two utilities are provided to allow users to query top-level BDIIs without having to know the details of LDAP syntax and the GLUE Schema.

They are, however, simply wrappers for the corresponding LDAP queries. The Glue Schema describes the information available from the above tools.

In both tools, if the BDII to be queried is not explicitly specified on the command line, it defaults to the one defined by the **LCG_GFAL_INFOSYS** environment variable.

```
env |grep LCG_GFAL_INFOSYS
```

Advanced Information System Queries:

lcg-infosites

The **lcg-infosites** command can be used to obtain VO-specific information on existing grid resources.

Getting information on computing elements (CE) and storage element (SE)

```
lcg-infosites --vo gridit ce –v
lcg-infosites --vo gridit se –v
```

Getting binding information between CE and SE

It is also possible to obtain a list of storage elements (SE) closest to each computing element (CE). The term "**close**" in grids can have several meanings including physical distance and speed of data access (available bandwidth). The closest SE to a CE is defined by the manager of the CE.

```
lcg-infosites --vo gridit closeSE
```

Advanced Information System Queries: lcg-info

The **lcg-info** command can be used to list either CEs or SEs and their attributes.

Finding out which software tag is supported on which CE

One of the most common use of the "lcg-info" is to check which software packages are available on CE's. The software packages are stored in the "Tag" attribute

```
lcg-info --vo gridit --list-ce --attrs Tag
```

Finding out which CEs have more than 10 cpus

List all the Computing Elements in the BDII satisfying the given conditions (TotalCPUs=10) and print the desired attributes (FreeCPUs,!TotalJobs)

```
lcg-info --vo gridit --list-ce --query 'TotalCPUs = 10' --attrs 'RunningJobs,FreeCPUs'
```

LDAP browser tool

- Web Tool
  - GSTAT: http://gstat-prod.cern.ch/gstat/ldap
- LDAP Browser (Windows)
  - LDAP Browser 2.6: http://www.ldapadministrator.com/download.htm
- LDAP Browser (Linux)
  - Luma (part of KDE): http://kde-apps.org/content/show.php?content=9771

GSTAT

http://gstat-prod.cern.ch/gstat/ldap

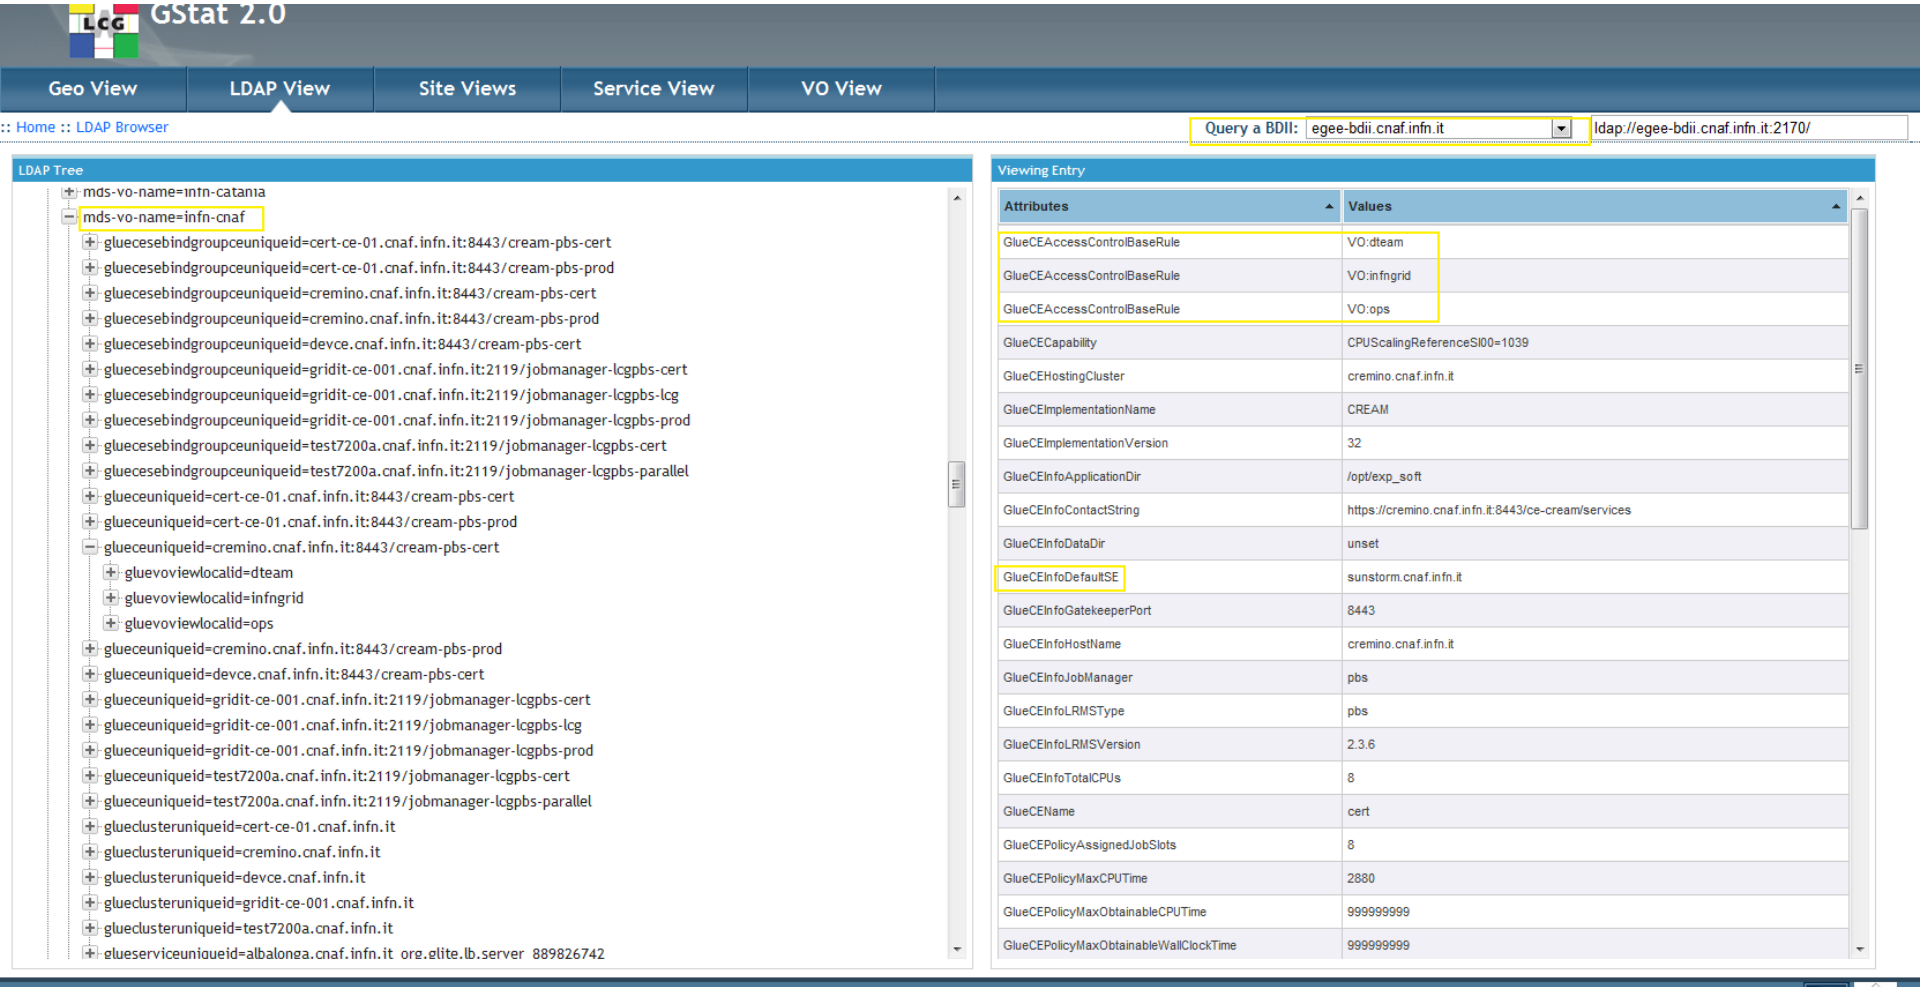

LDAP Browser 2.6

http://www.ldapadministrator.com/download.htm

FILE->New Profile

LABEL: ex. TOP-BDI

Host: egee.bdii.cnaf.infn.it

Port: 2170

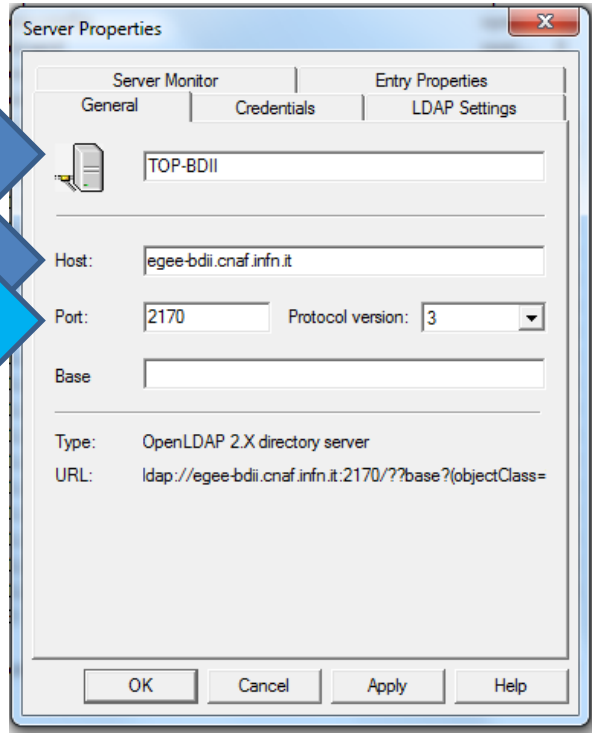

LDAP Browser 2.6

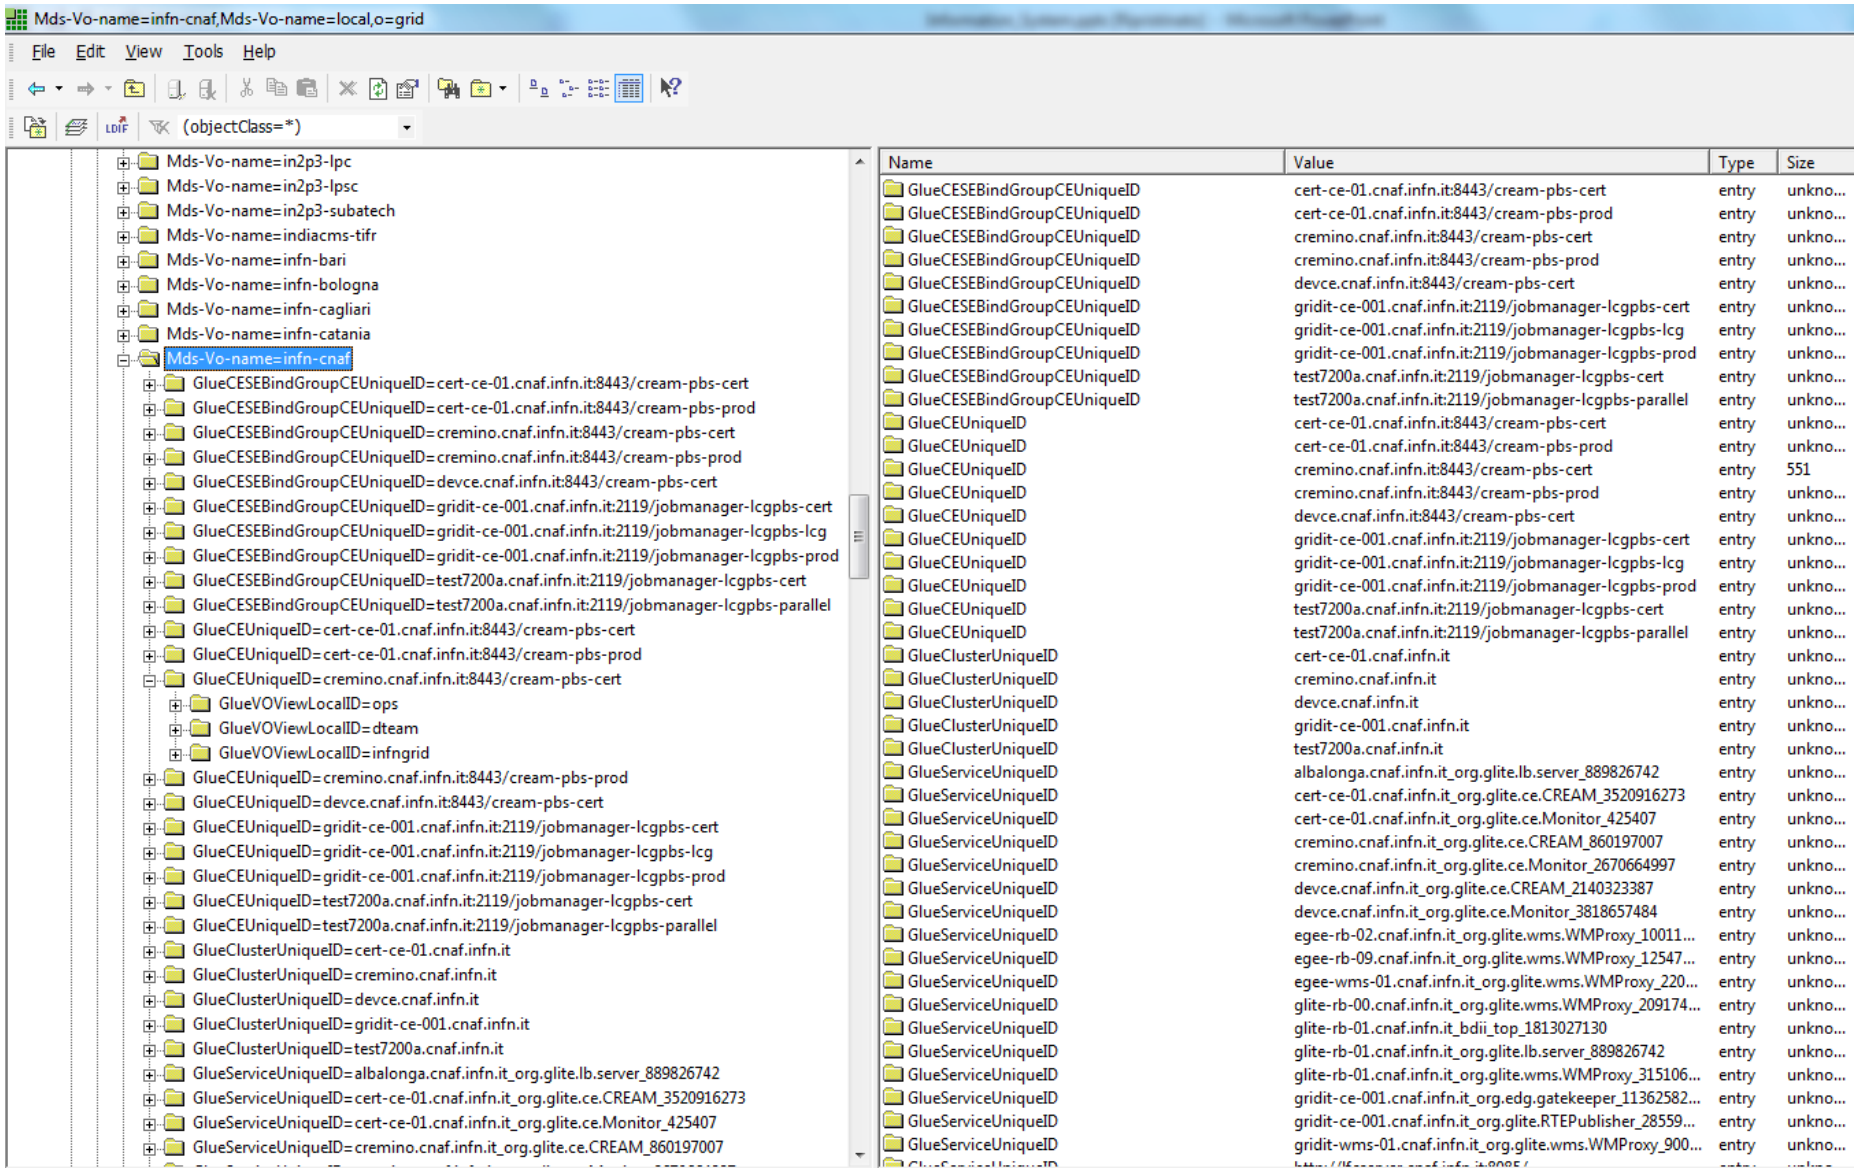

Luma (part of KDE)

- http://kde-apps.org/content/show.php?content=9771

- Apt-get install luma
- Yum install luma

- Settings -> Edit server list

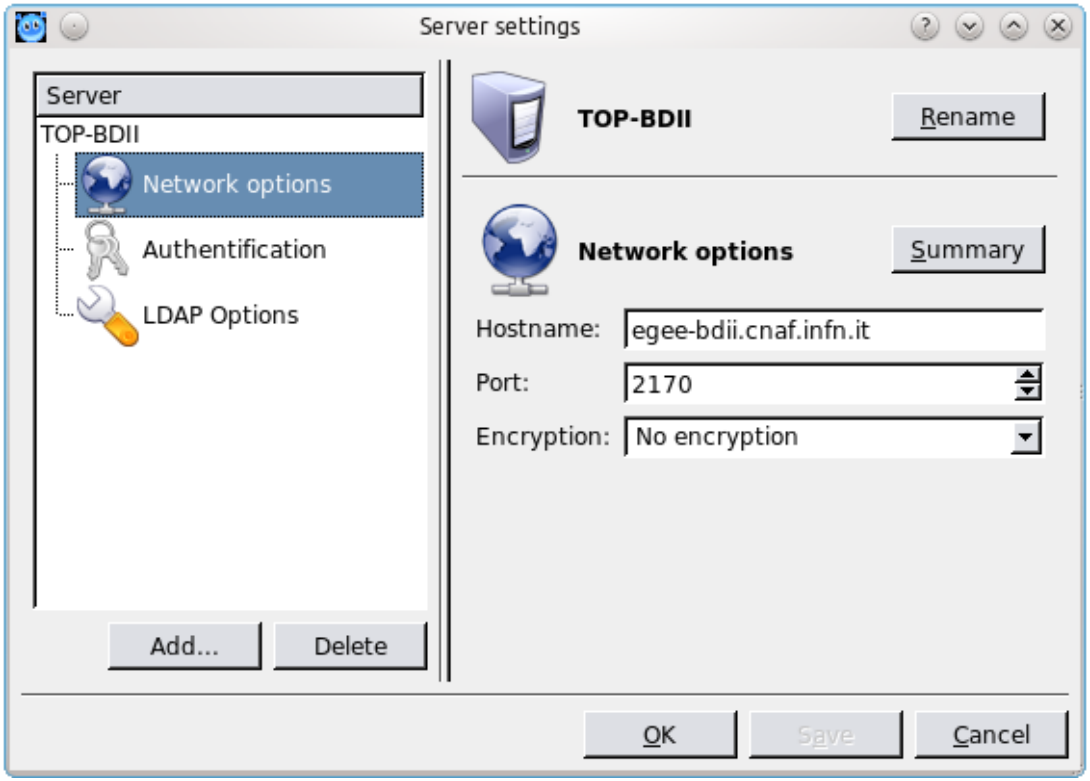

Luma (part of KDE)

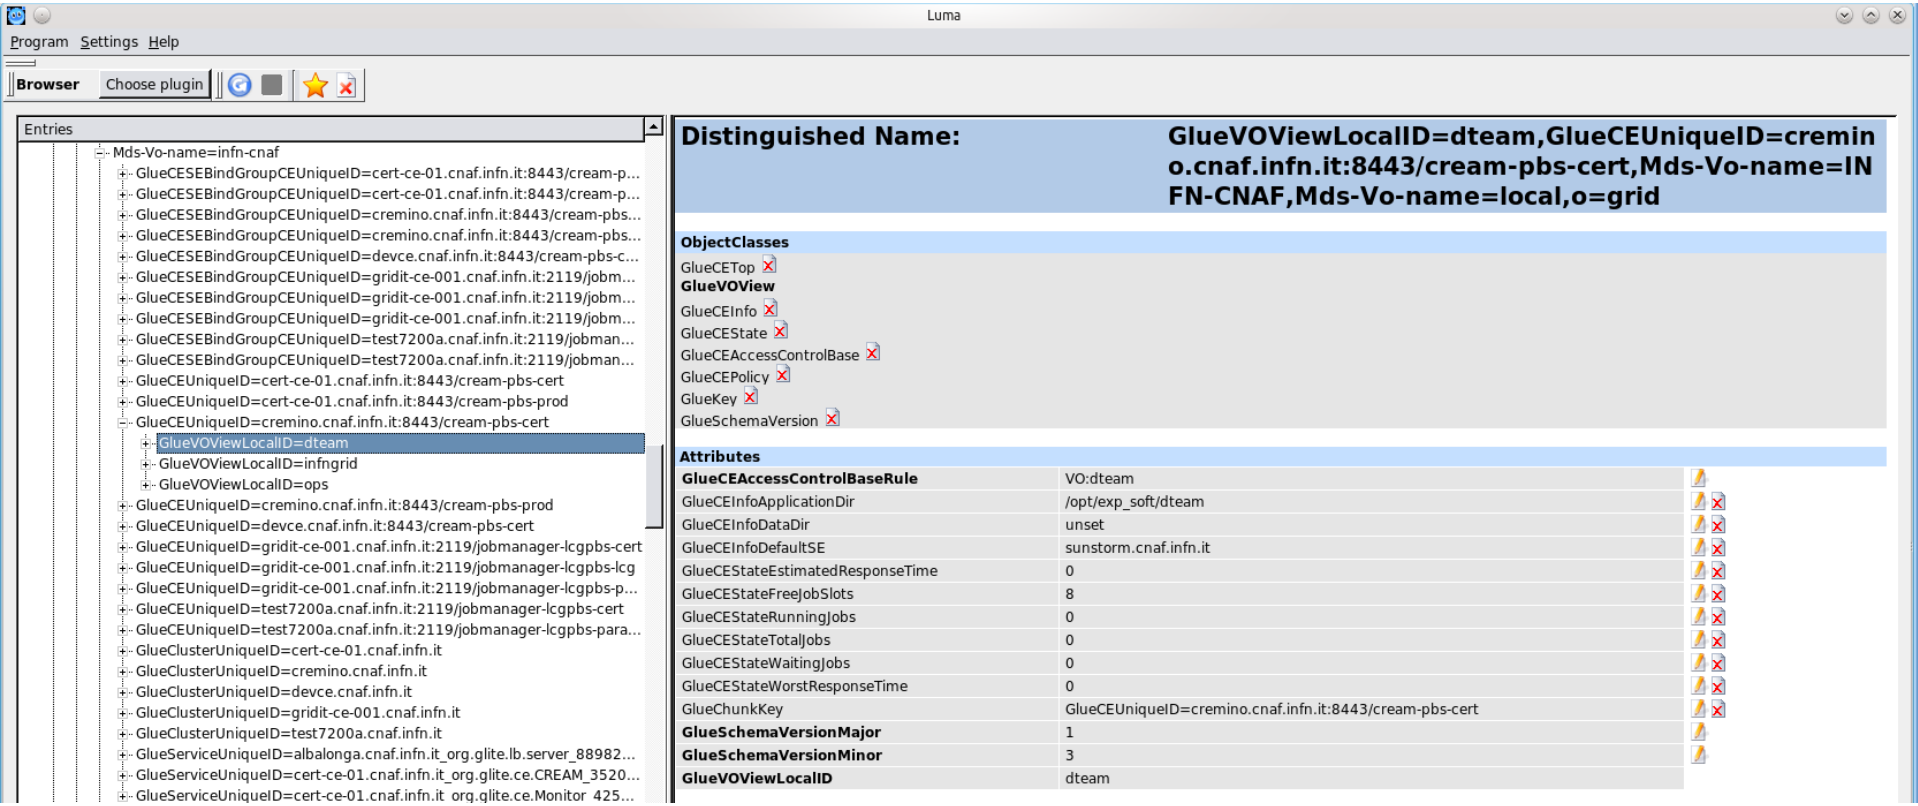

Conclusion

Although the use of the **ldapsearch** tool can be really flexible it needs a good knowledge of its usage and also the user should know exactly the names of all the needed attributes of the *glue schema*. In order to find out more information about the glue schema please refer to the following web address: http://glueschema.forge.cnaf.infn.it/

On the other hand the use of the **lcg-infosites** and **lcg-info** tools is simple and both commands are enough powerful to obtain information about Grid resources.

The End!

Questions ?