Querying the Information System

Last review date	Reviewer
2009-09-15	Marco Bencivenni
	Enrico Fattibene

Table of Contents

Querying the Information System

This page explains how to query the Information System in order to get an up- to- date view of the resources available on the Grid. It also explains the topology of the Information System, and gives a brief introduction to the GLUE schema.

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.

lcg- info- sites

The lcg- infosites command can be used to obtain VO- specific information on existing grid resources. The syntax is the following:

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

Use `lcg- infosites – h` or `-- help` for a description of the various flags.

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

```
lcg-infosites --vo lhcb -f cern-prod se

Avail Space(Kb) Used Space(Kb)  Type    SEs
-------------------------------------------------------
300000000000    160000000000    n.a    srm-lhcb.cern.ch
1000000000000   500000000000    n.a    srm-durable-lhcb.cern.ch
1000000000000   500000000000    n.a    castorsrm.cern.ch
[...]
```

lcg- info

The `lcg- info` command can be used to list either CEs or SEs and their attributes. The general format of the command for listing CE or SE information is:

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

Use `lcg- info – h` or `-- help` for a description of the various flags.

An example of the use of this command would be to find out which CEs have installed a particular version of an experiment's software. For example:

```
lcg-info --vo cms --list-ce --attrs Tag --query 'tag=*ORCA_8_7_1*'
```

The `-- list- attrs` option can be used to get a list of the supported attributes. This is especially helpful if one wishes to construct more complicated LDAP queries.

```
lcg-info --list-attrs

Attribute name        Glue object class      Glue attribute name

WorstRespTime         GlueCE                 GlueCEStateWorstResponseTime
CEAppDir              GlueCE                 GlueCEInfoApplicationDir
TotalCPUs             GlueCE                 GlueCEInfoTotalCPUs
MaxRunningJobs        GlueCE                 GlueCEPolicyMaxRunningJobs
CE                    GlueCE                 GlueCEUniqueID
WaitingJobs           GlueCE                 GlueCEStateWaitingJobs
[...]
```

Information System Topology

Collecting information on the Grid is done in a hierarchical manner. The following diagram illustrates the hierarchy. At the lowest level, resource- level BDIIs collect information on the state of resources (using scripts called "information providers"). Site- level BDIIs aggregate that information, and make it available to the top- level BDIIS. Periodically, the higher level servers make LDAP queries to the lower- level ones. There are multiple instances of the top level BDII in order to provide fault tolerance.
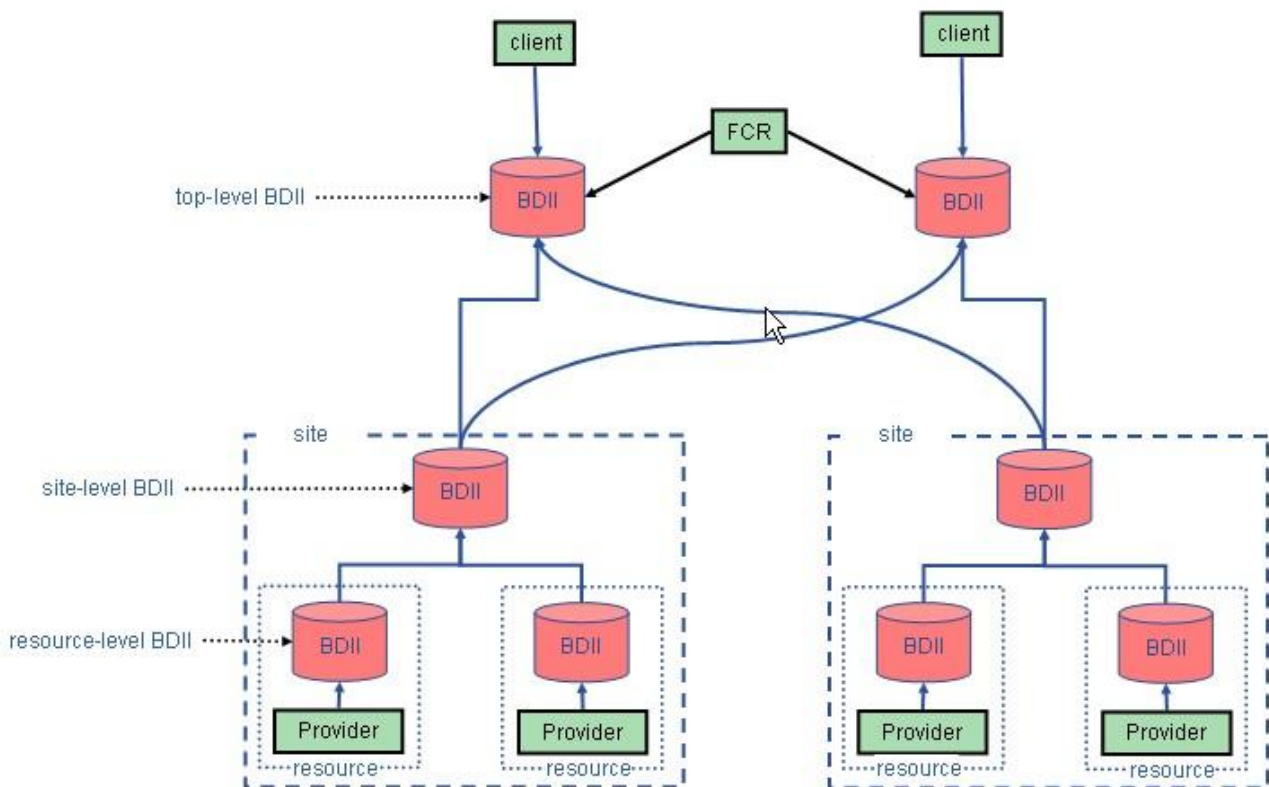
Top- level BDIIs are therefore the best places to look for comprehensive information on the resources available. As an example, Resource Brokers (RBs) query a top- level BDII to get the information used during the match- making process.

In the above diagram, FCR refers to "Freedom of Choice for Resources", which is a mechanism to allow Virtual Organizations to mask sites or services from their users if they are known not to be working correctly.

BDII

The BDII (Berkley Database Information Index) has been adopted in the gLite middleware as the Information System technology. It is based on Lightweight Directory Access Protocol (LDAP) servers. Good documentation on the BDII in the context of EGEE can be found at https:// twiki.cern.ch/ twiki// bin/ view/ EGEE/ BDII.

Within the BDII, one finds elements that have attributes and links to other elements. For example, a site element defines a Grid site in terms of its location, homepage, contact names etc., and also in terms of the Storage Elements, services and computing clusters that it offers. A cluster can comprise several Computing Elements (CEs), which are abstractions for a queue of jobs. A CE can have several attributes, such as configuration policies (e.g. MaxRunningJobs), access control policies, information on the jobs in the queue, and status information.

GLUE Schema

To query the BDII, one needs to understand the layout of the information therein - which is specified by the GLUE schema. The Grid Laboratory Uniform Environment (GLUE) schema is a data model to describe, in a precise and systematic way, information on static and dynamic Grid resources (including state and VO- specific views). Grid resources are geographically dispersed, span multiple trust domains and are heterogenous - hence the necessity to have a common method for discovering objects and their respective attributes. The GLUE schema started as collaboration effort between European and US grid projects to facilitate interoperation between them.

Documentation on GLUE Schema usage within EGEE can be found at https:// twiki.cern.ch/ twiki// bin/ view/ EGEE/ GlueUse.

LDAP

The protocol used to query the information system (BDII) is LDAP, an open standard. LDAP is a lightweight protocol for accessing directory services optimised for reading, browsing and searching. The LDAP information model is based on entries, which are collections of attributes that have globally- unique Distinguished Names (DN). Each entry's attributes have a type and one or more values, and entries are arranged in a hierarchical tree- like structure. A good overview of LDAP can be found here.

The port on which a BDII server responds to queries is typically 2170. 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. A trivial example of using LDAP would be to find the list of Virtual Organizations (VOs) supported at a particular site with a query such as the following:

```
ldapsearch -x -H ldap://lnx112.eela.if.ufrj.br:2170 \
      -b mds-vo-name=EELA-UNAM,mds-vo-name=local,o=grid \
        | grep ControlBaseRule | sort -u | awk '{print $2 }'| sed '/^VO/d'
```

In this example, the site name is EELA- UNAM and the output would be similar to:

```
alice
dteam
edteam
eela
```

More details about LDAP and the use of the `ldapsearch` command can be found in chapter "Advanced Information System Queries: ldapsearch".

There are also various graphical LDAP browsers available. One which is often included as standard in Linux distributions is gq. A java-based browser which can work on any patform is no longer being developed but is still useful. Other browsers can be found by using a search engine.

GOC Database

The Grid Operations Center Database (GOCDB) at https:// goc.gridops.org/ provides the authoritative list of sites in EGEE. This is where valid values for the `lcg- info- sites – f` parameter can be found. It can also be used to find the network names of BDIIs which can be queried. To access the GOCDB website, you will need a valid certificate.

The Information System is bootstrapped from the information in the GOCDB. When a site registers, it enters the URL for the site level BDII into the GOCDB. The GOCDB 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 BDIIs and the result is used to populate the top level BDII.

Further Information

For a more detailed description of the Information System, and further examples of querying it, please refer to the gLite User Guide (https:// edms.cern.ch/ file/722398// gLite-3- UserGuide.pdf).