

EUROPEAN MIDDLEWARE INITIATIVE

BDII DOCUMENTATION

Document Version:	1.2.0
Date:	30.01.2012
Author:	Maria Alandes Laurence Field

Document Change Record

Issue	Item	Reason for Change
1.0.1	Description of BDII_IPV6_SUPPORT variable.	A new variable has been defined in glite-yaim-bdii 4.3.11-1 to enable IPv6 support in the BDII. The Sys Admin Guide has been updated to include the description of the variable.
1.0.2	BDII core details.	Installation and Configuration details specific for resource BDII were missing.
	YAIM variables documented in Appendix.	A new appendix containing tables describing the YAIM variables needed to configure a resource, site and top BDII have been created. Formatting has been improved. BDII_REGIONS example has been fixed.
	Updated contents of /etc/glite/glite-info-update-endpoints.conf	Contents of the file were not up to date.
	Improved configuration details.	Included a more detailed description of the configuration files used in both top and site BDII, including usage of the different commands. Added a section about FCR.
	Reference to EGI BDII site documentation.	Added a link to the BDII site documentation created by EGI.
	Several minor fixes. Section about Glue 1.3/Glue 2.0 publication. Section about caching mode.	Review by Stephen Burke as reported in https://savannah.cern.ch/bugs/?97380
1.1.0	Information System overview diagram	New diagram representing also the interactions with EGI and OSG GOC DBs, and SAM.
	Nagios probes documentation EMIR integration Openldap versions LDAP logging Information provider order LDAP optimisations	A series of improvements have been requested as tracked in https://savannah.cern.ch/bugs/index.php?98391
1.2.0	ARC integration Support for GLUE 2.0 ldap URLs Sysconfig documentation More details on Nagios probe Improve troubleshooting with ldapsearch Update upgrade table to include EMI 3	A series of improvements have been requested as tracked in https://savannah.cern.ch/bugs/index.php?99551

Contents

1. ABOUT THIS DOCUMENT	5
2. FUNCTIONAL DESCRIPTION.....	6
2.1. INTRODUCTION	6
2.2. GLUE SCHEMA VERSIONS	7
2.3. CACHING MODE.....	7
3. SYSTEM ADMINISTRATOR GUIDE	8
3.1. PREREQUISITES	8
3.1.1 <i>BDII core</i>	8
3.1.2 <i>BDII top</i>	8
3.1.3 <i>BDII site</i>	9
3.2. INSTALLATION	9
3.2.1 <i>Software Repositories</i>	9
3.2.2 <i>Clean Installation</i>	10
3.2.3 <i>Upgrade Installation</i>	10
3.3. CONFIGURATION	11
3.3.1 <i>YAIM configuration</i>	11
3.3.2 <i>Advanced configuration</i>	11
3.4. SERVICE OPERATION	16
3.5. SERVICE MONITORING	16
3.5.1 <i>Status Information</i>	16
3.5.2 <i>Nagios Probes</i>	18
3.6. TROUBLESHOOTING	18
3.7. USER SUPPORT	19
4. SERVICE REFERENCE CARD	20
4.1. DAEMONS RUNNING	20
4.2. INIT SCRIPTS AND OPTIONS.....	20
4.3. CONFIGURATION FILE LOCATION	20
4.4. LOGFILE LOCATION	20
4.5. OPEN PORTS	20
4.6. UNIT TESTS	20
4.7. SERVICE STATE	20
4.8. CRON JOBS	21
4.9. SECURITY INFORMATION.....	21
4.9.1 <i>Access control mechanisms</i>	21
4.9.2 <i>How to block/ban a user</i>	21
4.9.3 <i>Network usage</i>	21
4.9.4 <i>Firewall configuration</i>	21
4.9.5 <i>Security recommendations</i>	21
4.9.6 <i>Security incompatibilities</i>	21
4.9.7 <i>List of external packages that are not maintained by the supported OS</i>	21
5. GLUE-VALIDATOR GUIDE	22
6. APPENDIX A: LDAP OPTIMISATIONS	23



EUROPEAN MIDDLEWARE INITIATIVE

7.	APPENDIX B: EMIR INTEGRATION	25
8.	APPENDIX C: ARC INTEGRATION	26
9.	APPENDIX D: YAIM VARIABLES	27
	9.1. BDII BASICS	27
	9.2. BDII TOP	28
	9.3. BDII SITE	29

1. ABOUT THIS DOCUMENT

This document is for System Administrators who want to install, configure and manage a BDII. It contains the following information:

- **Functional Description:** a summary of the BDII structure and functionalities.
- **System Administrator Guide:** Detailed document covering all the aspects of installing, configuring, managing and monitoring the resource, site and top BDII's.
- **Service Reference Card:** Useful summary for top and site BDII operation. Aimed at more advanced system administrators with hands-on experience.
- **glue-validator guide:** guide to use the glue-validator command. Glue-validator checks whether the information published by a service is compliant with Glue 1.3 or Glue 2.0. Aimed at any system administrator managing middleware services that are published in the information system.

The document also contains a series of appendixes for more detailed information in the following areas:

- **LDAP optimisations:** details on LDAP configuration options used to implement a series of optimisations for better performance in the different BDII flavours.
- **EMIR integration:** guide on how to publish site and top BDII in EMIR.
- **ARC integration:** details on the changes required in the BDII to be able to publish ARC resources.
- **YAIM variables:** detailed documentation of YAIM variables involved in the configuration of the different BDII flavours.

2. FUNCTIONAL DESCRIPTION

2.1. INTRODUCTION

Grid information systems are mission-critical components in today's production grid infrastructures. They provide detailed information about grid services which is needed for various different tasks. The EMI information system has a hierarchical structure of three levels. The fundamental building block used in this hierarchy is the Berkley Database Information Index (BDII). Although the BDII has additional complexity, it can be visualized as an LDAP database. The resource level or core BDII is usually co-located with the grid service and provides information about that service. Each grid site runs a site level BDII. This aggregates the information from all the resource level BDIIs running at that site. The top level BDII aggregates all the information from all the site level BDIIs and hence contains information about all grid services. There are multiple instances of the top level BDII in order to provide a fault tolerant, load balanced service. The information system clients query a top level BDII to find the information that they require.

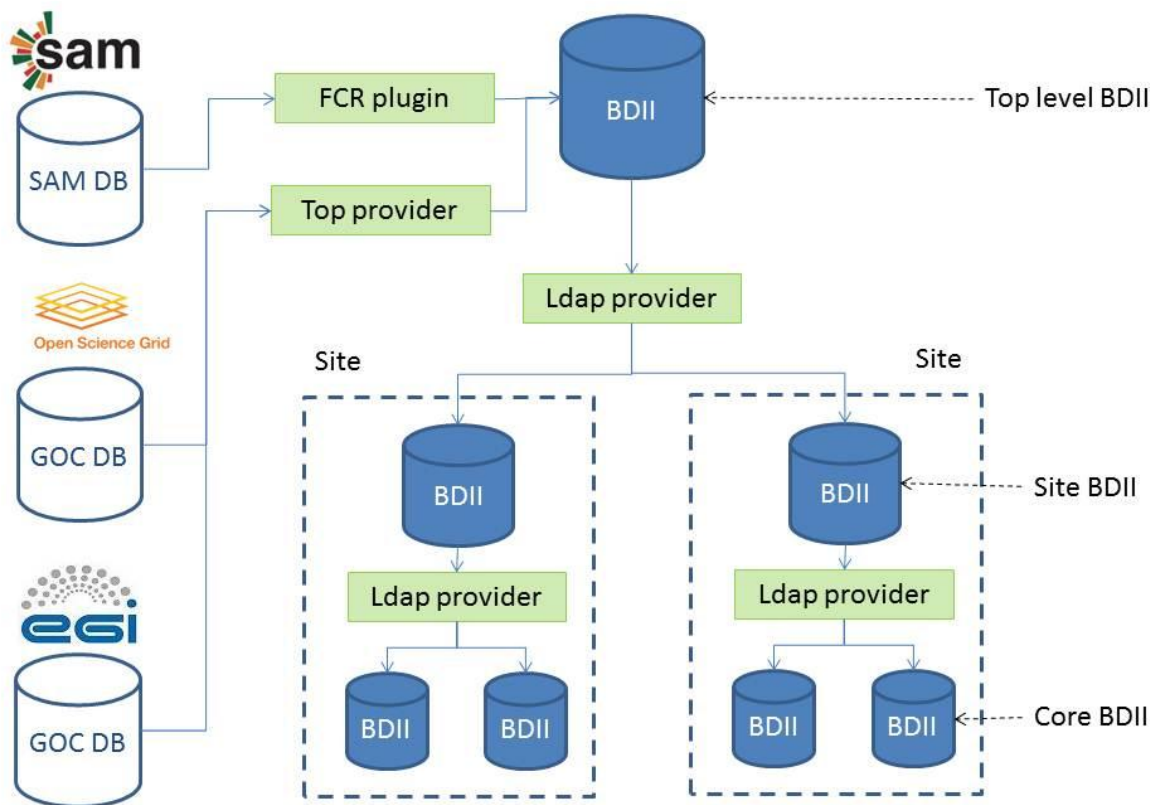


Figure 1 - Information System Overview

The BDIIs are populated with information by running information providers. These are scripts which obtain information, format it as LDIF and print the result to standard out. These information providers can also be used to query other BDIIs which is how the hierarchy is built. The order in which these information providers are run is random.

The information in the information systems conforms to a schema called the GLUE schema. The GLUE schema started as collaboration effort between European and US grid projects to facilitate interoperation between them. The Open Grid Forum (OGF) is now responsible for the GLUE schema.

The Freedom of Choice for Resources mechanism (FCR) is used within a top level BDII to enable the Virtual Organization (VO) to influence their usage of specific services. The FCR portal generates a list of services available to a VO. The portal can be used by the VO manager to either white list or black list sites. This information is then downloaded by the top level BDII and it deletes the ACL for the VO on that specific service from the database. This results in the service not being matched in a query from that VO.

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.

2.2. GLUE SCHEMA VERSIONS

The Information System publishes information in both Glue 1.3 and Glue 2.0 versions. This is done by using two different LDAP Data Information Trees (DIT) for each Glue Schema version. The Glue 1.3 DIT root is o=grid, whereas the Glue 2.0 DIT root is o=glue.

LDAP URLs are stored in GOC DB using the Glue 1.3 binding. For example:

```
ldap://prod-bdii.cern.ch:2170/mds-vo-name=CERN-PROD,o=grid
```

In order to get Glue 2.0 information from the grid resources, the Glue 1.3 LDAP URL is transformed into a Glue 2.0 LDAP URL internally by the BDII, using the glite-info-provider-ldap information provider. As of version 1.4.4-1, GLUE 2.0 LDAP URLs can be directly defined in the GOCDB.

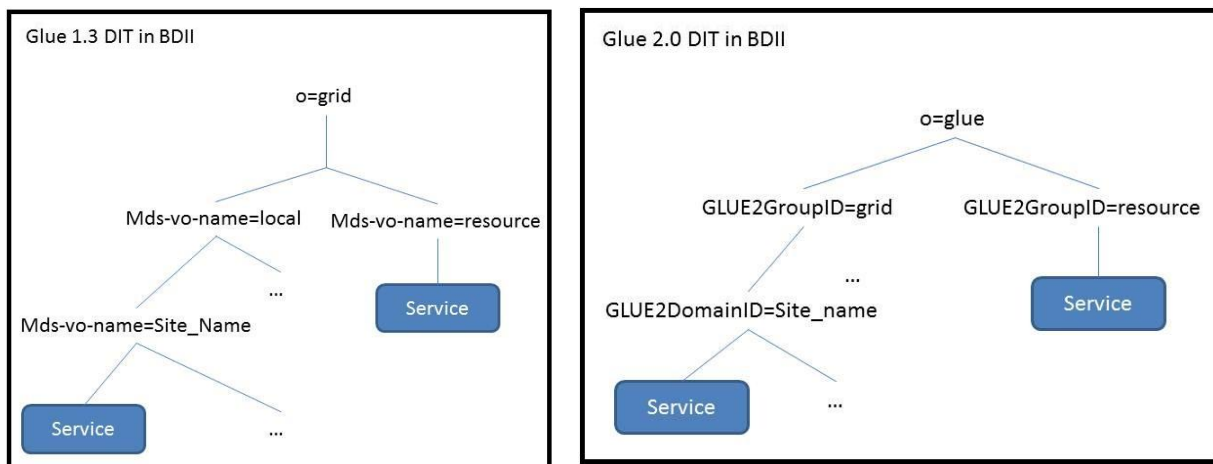


Figure 2 - Glue 1.3 and Glue 2.0 DIT structure in a top level BDII

2.3. CACHING MODE

The BDII uses a snapshot approach to publish information about grid resources. Snapshots are produced in a regular way by configuring the frequency of the BDII updates. Services that disappear between subsequent snapshots because they are temporarily offline can give the appearance of information volatility. This is a disadvantage to discover services using the information system. The top level BDII offers the possibility of enabling a caching mode whereby information is cached for a certain amount of time.

3. SYSTEM ADMINISTRATOR GUIDE

3.1. PREREQUISITES

3.1.1 BDII core

Please, check the following hardware and software requirements before installing a resource BDII.

<i>Hardware</i>	
CPU	No particular requirements.
Memory	
Disk	
Network	Port 2170 needs to be open.
<i>Software</i>	
OS version	Scientific Linux 5 (EMI 1 and 2) and Scientific Linux 6 (EMI 2)
Python version	Python 2.5 in SL5 and python 2.6 in SL6
OpenLDAP software	openldap 2.4 is required. This is the default version in SL6 repositories. In the case of SL5, note that this version is not available. In this case, the version is distributed as third party software in the EMI repository.
Host Certificates	No host certificate is required to be installed.
Recommended deployment scenarios	Install BDII core in the same host where the service that is going to be published in the Information System is running.

3.1.2 BDII top

Please, check the following hardware and software requirements before installing a top BDII.

<i>Hardware</i>	
CPU	Dual core CPU
Memory	2-3 GB RAM. If you decide to set <i>BDII_RAM_DISK=yes</i> in your YAIM configuration, it's advisable to have 4GB of RAM.
Disk	10 GB of hard disk space.
Network	Port 2170 needs to be open.
<i>Software</i>	
OS version	Scientific Linux 5 (EMI 1 and 2) and Scientific Linux 6 (EMI 2)
Python version	Python 2.5 in SL5 and python 2.6 in SL6
OpenLDAP software	openldap 2.4 is required. This is the default version in SL6 repositories. In the case of SL5, note that this version is not available. In this case, the version is distributed as third party software in the EMI repository.
Host Certificates	No host certificate is required to be installed.
Recommended	Due to the critical nature of the information system with respect to the

deployment scenarios	operation of the grid, the Top BDII should be installed as a stand-alone service to ensure that problems with other services do not affect the Top BDII. In no circumstances should the Top BDII be co-hosted with a service which has the potential to generate a high load.
High Availability	Check EGI documentation for more details on how to provide high availability for a top level BDII: https://wiki.egi.eu/wiki/MAN05

3.1.3 BDII site

Please, check the following hardware and software requirements before installing a site BDII.

<i>Hardware</i>	
CPU	Dual core CPU.
Memory	2-3 GB RAM.
Disk	10 GB of hard disk space.
Network	Port 2170 needs to be open.
<i>Software</i>	
OS version	Scientific Linux 5 (EMI 1 and 2) and Scientific Linux 6 (EMI 2)
Python version	Python 2.5 in SL5 and python 2.6 in SL6
OpenLDAP software	openldap 2.4 is required. This is the default version in SL6 repositories. In the case of SL5, note that this version is not available. In this case, the version is distributed as third party software in the EMI repository.
Host Certificates	No host certificate is required to be installed.
Recommended deployment scenarios	Due to the critical nature of the information system with respect to the operation of the grid, the site BDII should be installed as a stand-alone service to ensure that problems with other services do not affect the site BDII. In no circumstances should the site BDII be co-hosted with a service which has the potential to generate a high load.

3.2. INSTALLATION

3.2.1 Software Repositories

In order to install a core, site or top BDII, the EMI repositories have to be used. Please, refer to:

- EMI 1 Generic Installation and Configuration guide:
<https://twiki.cern.ch/twiki/bin/view/EMI/GenericInstallationConfigurationEMI1>
- EMI 2 Generic Installation and Configuration guide:
<https://twiki.cern.ch/twiki/bin/view/EMI/GenericInstallationConfigurationEMI2>
- EMI 3 Generic Installation and Configuration guide:
<https://twiki.cern.ch/twiki/bin/view/EMI/GenericInstallationConfigurationEMI3>

3.2.2 Clean Installation

The following command has to be run to install a resource BDII:

```
yum install emi-resource-information-service
```

The following EMI packages are installed:

- bdi
- glue-schema
- glite-info-provider-service
- glue-validator
- glue-validator-cron
- emi-version

The following command has to be run to install a top BDII:

```
yum install emi-bdi-top
```

The following EMI packages are installed:

- emi-resource-information-service
- glite-info-provider-ldap
- glite-yaim-bdi
- glite-yaim-core
- bdi-config-top
- glite-info-plugin-fcr
- glite-info-update-endpoints

The following command has to be run to install a site BDII:

```
yum install emi-bdi-site
```

The following EMI packages are installed:

- emi-resource-information-service
- glite-info-provider-ldap
- glite-yaim-bdi
- glite-yaim-core
- bdi-config-site
- glite-info-site
- glite-info-static
-

3.2.3 Upgrade Installation

Running a *yum upgrade* is sufficient to update a BDII installation.

The following table explains which upgrade paths are supported:

BDII version from -> to	gLite 3.2	EMI 1 SL5	EMI 2 SL5	EMI 2 SL6	EMI 3 SL5	EMI 3 SL6
gLite 3.2	YES (*)	NO	NO	NO	NO	NO
EMI 1 SL5	NO	YES	YES (**)	NO	YES(**)	NO
EMI 2 SL5	NO	NO	YES	NO	YES	NO
EMI 2 SL6	NO	NO	NO	YES	NO	YES

(*) gLite 3.2 is no longer supported since 30.04.2012.

(**) This upgrade path hasn't been tested.

3.3. CONFIGURATION

3.3.1 YAIM configuration

YAIM is a tool for simple configuration of middleware services. For more information, please check <https://twiki.cern.ch/twiki/bin/view/EGEE/YAIM>. YAIM can be used by system administrators who want to abstract from the configuration details of each service. For a more advanced configuration of the BDII, please check the next section.

The YAIM variables defined to configure the basics of a BDII are listed in Appendix 9.1. They are defined for resource, site and top BDII.

In order to configure a resource BDII with yaim, the function *config_bdii_5.2* has to be included in the node-info.d file of the service.

The YAIM variables defined to configure a top BDII are listed in Appendix 9.2.

The following YAIM command has to be run to configure a top BDII:

```
yaim -c -s site-info.def -n BDII_top
```

The YAIM variables defined to configure a site BDII are listed in Appendix 9.3.

The following YAIM command has to be run to configure a site BDII:

```
yaim -c -s site-info.def -n BDII_site
```

3.3.2 Advanced configuration

YAIM enables the system administrator to abstract from the specific configuration details of each service. In this section, the detailed configuration files of the BDII are presented in case a system administrator does not want to use YAIM and prefers to directly configure the BDII.

The basic configuration file needed for resource, site and top BDII is **/etc/bdii/bdii.conf**. This configuration file is used to configure the basic parameters of the BDII. The format is key-value pairs with a “=” sign as separator. A default configuration file comes with any BDII installation. It requires editing for the BDII to work in a different way. The table below describes the key-value pairs found in the configuration file:

Key	Default Value	Description
BDII_LOG_FILE	/var/log/bdii/bdii-update.log	The log file for the BDII update process
BDII_PID_FILE	/var/run/bdii/bdii-update.pid	PID file for the bdii-update daemon
BDII_LOG_LEVEL	ERROR	The log level for the update process [ERROR, WARNING, INFO, DEBUG]
BDII_LDIF_DIR	/var/lib/bdii/gip/ldif	The directory containing LDIF files
BDII_PROVIDER_DIR	/var/lib/bdii/gip/provider	The directory containing information providers
BDII_PLUGIN_DIR	/var/lib/bdii/gip/plugin	The directory containing plugins
BDII_PORT	2170	The port which is used for the LDAP server
BDII_BREATHE_TIME	120	Time in seconds to wait before updating the next database
BDII_READ_TIMEOUT	300	Time to wait for LDAP sources to return
BDII_ARCHIVE_SIZE	0	The BDII stores a compressed snapshot of the whole information tree in BDII_VAR_DIR/archive every BDII_BREATHE_TIME. This variable defines the total number of snapshots to be kept in the archive directory. Older ones are deleted.
BDII_DELETE_DELAY	43200	Time to wait in seconds before deleting removed entries. Default is 12 hours. This variable activates the caching mode to be used only in the top level BDII.
BDII_USER	ldap	The user running the update process and the slapd databases
BDII_VAR_DIR	/var/lib/bdii	The directory used by the BDII for writing data.

BDII top specific configuration files:

- **/etc/sysconfig/bdii:** it contains a set of basic BDII configuration variables. If YAIM is used to configure the top BDII, these attributes are taken from the YAIM configuration. Otherwise, this file has to be manually updated with the proper values:

Variable and default value	Description
SLAPD_CONF=/etc/bdii/bdii-slapd.conf	Location of the LDAP configuration file. This has to be changed to /etc/bdii/bdii-top-slapd.conf for a top level BDII.
SLAPD=/usr/sbin/slapd	LDAP daemon. Default is LDAP version 2.3. Change this to /usr/bin/slapd2.4 .
BDII_RAM_DISK=no	This variable defines whether to use a RAM disk to store the LDAP tree. Set this to yes for the top level BDII.
BDII_IPV6_SUPPORT=yes	This variable is not present in the file but needs to be manually defined if IPv6 support is needed. The default behaviour of the BDII is not supporting IPv6.

- **/etc/bdii/bdii-top-slapd.conf**: slapd configuration files. For more details on OpenLDAP configuration, please check www.openldap.org.
- **/etc/glite/glite-info-update-endpoints.conf**: top BDII configuration file to specify which site BDIIs are going to be queried. By default the file uses EGI and OSG list of site BDIIs in status *Certified*, taken from their corresponding GOC DB URLs.

It is possible to manually specify a different list of URLs by setting *OSG* and *EGI* to False (case sensitive), *manual* to True (case sensitive), and defining a manual file in *manual_file*. Default values are:

```

EGI = True
OSG = True
certification_status = Certified
manual = False
manual_file =
output_file = /var/cache/glite/top-urls.conf
cache_dir = /var/cache/glite/glite-info-update-endpoints

```

The OSG GOC DB URL is:

```

http://myosg.grid.iu.edu/rgsummary/xml?datasource=summary&summary_attrs_showwlcg=on&all_resources=on&gridtype=on&gridtype_1=on&active=on&active_value=1&disable_value=1

```

The EGI GOC DB URL is:

```

http://goc.egi.eu/gocdbpi/public/?method=get_site_list&certification_status=Certified&production_status=Production

```

In case a manual file is specified, the following syntax is needed:

```

SITE_NAME SITE_BDII_LDAP_URL
Example:
CERN-PROD ldap://prod-bdii.cern.ch:2170/mds-vo-name=CERN-PROD,o=grid

```

The */var/cache/glite/top-urls.conf* contains the list of site BDIIs and is used by the *glite-info-provider-top* information provider, which in turn uses the *glite-info-provider-ldap* command to query the site BDIIs and publish their resources.

```
Usage: /opt/glite/libexec/glite-info-provider-ldap -c <config file>
[-m <mds-vo-name>] [-v validity] [-d dynamic] [-t timeout] [-s size]
[-g site/region]

This information provider takes a list of LDAP urls from the
configuration file and queries the LDAP sources. It then merges the
results and modifies the dn to append the mds-vo-name as specified on
the command line.

The configuration file has lines typically with the following format:
entity ldap://host.domain:2170/mds-vo-name=something,o=grid
Comments start with "#" and are ignored.

Options:

-c The configuration file listing all the LDAP URLs.
-m The mds-vo-name which should be used.
-h Displays this helpful message.
-d This option will change the search filter to only
  retrieve the dynamic information. (currently not supported)
-g This option will modify the search endpoint to obtain glue
  2.0 information and dynamically modify the dn if value is not
  "none".
-t The timeout of the ldapsearch in seconds.
-v The validity of the cache files in seconds.
-s Maximum file size in megabytes for a single source.
```

- **FCR mechanism:** The Freedom of Choices mechanism is currently hardcoded in the BDII top code and it is only active for the CMS VO. The *generate-fcr-exclude-file* cron job executed every hour contacts the SAM server to obtain the list of blacklisted resources. The used URL is: http://grid-monitoring.cern.ch/myegi/sam-pi/status_of_service_in_profile?vo_name=cms&profile_name=CMS_BLACKLIST&output=xml

BDII site specific configuration files:

- **/etc/sysconfig/bdii:** it contains a set of basic BDII configuration variables. If YAIM is used to configure the site BDII, these attributes are taken from the YAIM configuration. Otherwise, this file has to be manually updated with the proper values:

Variable and default value	Description
SLAPD_CONF=/etc/bdii/bdii-slapd.conf	Location of the LDAP configuration file. The default value is the correct one for a site BDII.
SLAPD=/usr/sbin/slapd	LDAP daemon. Default is LDAP version 2.3. Change this to /usr/bin/slapd2.4 .
BDII_RAM_DISK=no	This variable defines whether to use a RAM disk to store the LDAP tree. The site BDII LDAP tree does not need to be in RAM disk.
BDII_IPV6_SUPPORT=yes	This variable is not present in the file but needs to be manually defined if IPv6 support is needed. The default behaviour of the BDII is not supporting IPv6.

- **/etc/bdii/bdii-slapped.conf**: slapd configuration files. For more details on OpenLDAP configuration, please check www.openldap.org.
- **/etc/glite-info-static/site/site.cfg**: BDII site configuration file to specify the site attributes. Default values are:

```
# Site name, a domain name.
SITE_NAME = MySite
# Optional description of the site
SITE_DESC = Testing
# Optional web address in full URL format. E.g.:
# http://www.cern.ch/gridinfo
SITE_WEB = http://www.cern.ch/gridinfo
# New optional GLUE 2.0 parameter to identify if the site is
# geographically distributed
# true or false
SITE_DISTRIBUTED = no
# Location of the site
SITE_LOC = Geneva, Switzerland
SITE_COUNTRY = Switzerland
# Latitude
SITE_LAT = 0.0
# Longitude
SITE_LONG = 0.0
# Contact email
SITE_EMAIL = admin@domain.invalid
# Security email
SITE_SECURITY_EMAIL = admin@domain.invalid
# User support email
SITE_SUPPORT_EMAIL = admin@domain.invalid
# Other info
OTHERINFO = GRID=EGEE
```

If YAIM is used to configure the site BDII, these attributes are taken from the YAIM configuration. Otherwise, this file has to be manually updated with the proper site values.

The `/etc/glite-info-static/site/site.cfg` file is used by the `glite-info-provider-site-entry` which in turns uses the `glite-info-static` command to publish the site information.

```
Usage: /usr/sbin/glite-info-static -m <module> [OPTIONS]
       glite-info-create.sh -m <module> [-i <ifaces>] [-t <templates>] [-c
<configs>] [-p <path>] [-o <outpath>] [-d debug]
```

Parameters:

```
-m <module>      The module you are using. E.g.: site
-i <ifaces>      The interface you want to use. E.g.: glue, wlcg
                 (default)
-t <templates>  The template you want to use. E.g.: glue1, glue2
                 (all default)
-c <config>     The config file location if outside from the module
                 directory
-p <path>       The path for the module directory.
                 Default: /etc/glite-info-static
-d <debug>     Debug level: 0:ERROR, 1:WARNING, 2:INFO, 3:DEBUG.
                 Default: 0
```

Examples:

```
glite-info-create.sh -m site
glite-info-create.sh -m site -i 'glue wlcg' -t glue2 -c /etc/site.cfg
```

- **/etc/bdii/gip/site-urls.conf** : BDII site configuration file to specify the list of resource BDIIs that are going to be published in the site BDII. The syntax is a list of service name identifier and URL pairs:

```
SERVICE_NAME_ID RESOURCE_BDII_LDAP_URL
```

Example:

```
CE1 ldap://myce.cern.ch:2170/mds-vo-name=resource,o=grid
```

The */etc/bdii/gip/site-urls.conf* file is used by the *glite-info-provider-site* information provider, which in turn uses the *glite-info-provider-ldap* command (check the description in the previous BDII top configuration section).

For further details on BDII site configuration, check the EGI documentation:

https://wiki.egi.eu/wiki/MAN01#Configuring_a_site_BDII

3.4. SERVICE OPERATION

The BDII is started and stopped by the daemon script */etc/init.d/bdii*.

When a BDII is started, the following processes run:

- One multithreaded slapd process. The number of (active) threads may depend on the query load and/or the */etc/bdii/bdii-slapd.conf* file or */etc/bdii/bdii-top-slapd.conf*.
- 1 *bdii-update* process.
- Periodically, one *ldapadd*, *ldapdelete* or *ldapmodify* process.

3.5. SERVICE MONITORING

3.5.1 Status Information

Status information about the BDII is available by querying the **o=Infosys** root for the *UpdateStats* object.

```
ldapsearch -LLL -x -h localhost -p 2170 -b o=Infosys
```


This entry contains a number of metrics relating to the latest update such as the time to update the database and the total number of entries. An example of such entry is shown below.

```
dn: Hostname=localhost,o=infosys
objectClass: UpdateStats
Hostname: lxbra2510.cern.ch
FailedDeletes: 0
ModifiedEntries: 4950
DeletedEntries: 1318
UpdateTime: 150
FailedAdds: 603
FailedModifies: 0
TotalEntries: 52702
QueryTime: 8
NewEntries: 603
DBUpdateTime: 11
ReadTime: 0
PluginsTime: 4
ProvidersTime: 113
```

The following table shows the meaning of these metrics:

Metric	Description
Hostname	The hostname of the BDII instance. This is hardcoded to localhost since the statistics are a mechanism for sys admin diagnosis and can only be accessed per BDII instance.
FailedDeletes	The number of delete statements which failed.
ModifiedEntries	The number of objects to modify.
DeletedEntries	The number of objects to delete.
UpdateTime	To total update time in seconds.
FailedAdds	The number of add statements which failed.
FailedModifies	The number of modify statements which failed.
TotalEntries	The total number of entries in the database.
QueryTime	The time taken to query the database.
NewEntries	The number of new objects.
DBUpdateTime	The time taken to update the database in seconds.
ReadTime	The time taken to read the LDIF sources in seconds.
PluginsTime	The time taken to run the plugins in seconds.
ProvidersTime	The time taken to run the information providers in seconds.

3.5.2 Nagios Probes

Nagios probes can also be used to monitor the state of the BDII. The interface has been designed to be used with Nagios, but the probe can be used independently.

```
Usage: check_bdii_entries -H host [OPTION]
Options:
  -p      The port of the BDII
  -b      The base of the BDII
  -f      The filter to use for the search
         Note: Shortcuts 'site', 'service' and 'stats' can be used
  -w      Warning threshold <time/freshness warning>:<entry warning>
  -c      Critical threshold <time/freshness critical>:<entry critical>
  -t      Timeout in seconds (not implemented)
```

The output of the command has the following syntax:

- Freshness of the information (Only available for the *stats* filter) :
`Exit_code: freshness=num_sec, entries=num |freshness=num_sec;entries=num`
- Number of entries (default filter is *service* if nothing is specified):
`Exit_code: time=num_sec, entries=num |time=num_sec;entries=num`

Where *Exit_code* is OK, WARNING or CRITICAL; *num_sec* is a time in seconds; and *num* is the total number of entries. The exit code is calculated based on the thresholds specified when launching the probe.

The following examples show different uses cases calling the probe.

Command to check the number of sites in a top BDII:

```
check_bdii_entries -H host.invalid -p 2170 -b mds-vo-name=local,o=grid -f site, -w 10:300 -c 15:200
```

Command to check the number of services in a site BDII:

```
check_bdii_entries -H host.invalid -p 2170 -b mds-vo-name=SiteName,o=grid -f service, -w 5:4 -c 10:1
```

Command to check the freshness of information:

```
check_bdii_entries -H host.invalid -p 2170 -b o=stats -f stats, -w 600:1 -c 1000:1
```

3.6. TROUBLESHOOTING

When experiencing problems with the BDII, the steps below must be followed to be able to identify the nature of the problem:

- Identify in which level of the BDII the problem appears: top, site or resource. Some useful ldapsearch commands for this are:

- GLUE 1:
 - Resource BDII: `ldapsearch -LLL -x -h localhost -p 2170 -b "mds-vo-name=resource,o=grid"`
 - Site BDII: `ldapsearch -LLL -x -h localhost -p 2170 -b "mds-vo-name=SITE_NAME,o=grid"`
 - Top BDII: `ldapsearch -LLL -x -h localhost -p 2170 -b "o=grid"`
- GLUE 2:
 - Resource BDII: `ldapsearch -LLL -x -h localhost -p 2170 -b GLUE2GroupID=resource,o=glue`
 - Site BDII: `ldapsearch -LLL -x -h localhost -p 2170 -b GLUE2DomainID=SITE_NAME,o=glue`
 - Top BDII: `ldapsearch -LLL -x -h localhost -p 2170 -b o=glue`
- Identify whether the problem is in the BDII itself or in the data sources: ldif, provider or plugins. Some log files that could be useful to check:
 - Site BDII glite-info-provider-ldap: check `BDII_VAR_DIR/gip/tmp/gip/log/site-urls.conf(-glue2)`
 - top level BDII glite-info-update-endpoints: check `/var/log/glite/glite-info-update-endpoints.log`.
 - top level BDII glite-info-provider-ldap: check `BDII_VAR_DIR/gip/tmp/gip/log/top-urls.conf(-glue2)`
 - For all BDII flavours: `BDII_LOG_FILE` for general error messages. Change the `BDII_LOG_LEVEL` variable to `DEBUG` for more details and restart the BDII.
- LDAP logs are disabled by default. In order to enable logging for the LDAP server, please follow the steps below:
 - Define `loglevel` in the `/etc/bdii/bdii-(top-)slapd.conf` file. Possible values for `loglevel` are documented here: <http://www.openldap.org/doc/admin24/slapdconfig.html>
 - Add the following line in `/etc/syslog`:

```
Local4.*      /var/log/slapd.log
```
 - Restart the syslog daemon and the BDII.

For a list of common problems, please check <https://twiki.cern.ch/twiki/bin/view/EGEE/InfoTrouble>.

3.7. USER SUPPORT

Incidents with the BDII should be reported using GGUS: <http://ggus.org>.

4. SERVICE REFERENCE CARD

4.1. DAEMONS RUNNING

```
/usr/sbin/slapd -f /etc/bdii/bdii-(top-)slapd.conf -h ldap://0.0.0.0:2170 -u ldap
```

```
/usr/bin/python /usr/sbin/bdii-update -c /etc/bdii/bdii.conf -d
```

4.2. INIT SCRIPTS AND OPTIONS

```
/etc/init.d/bdii (start|stop|restart|condrestart|status)
```

4.3. CONFIGURATION FILE LOCATION

/etc/sysconfig/bdii: startup configuration file for the bdii.

/etc/bdii/bdii.conf: general bdii configuration file.

BDII top

- /etc/bdii/bdii-top-slapd.conf: LDAP configuration file.
- /etc/glite/glite-info-update-endpoints.conf: configuration file defining the endpoints to download the list of site BDII's to be published in the top level BDII.

BDII site

- /etc/bdii/bdii-slapd.conf: LDAP configuration file.
- /etc/glite-info-static/site/site.cfg: configuration file defining the site attributes.
- /etc/bdii/gip/site-urls.conf: configuration file containing the list of Resource level BDII's of the services published by the site BDII. The syntax is a list of service name-URL pairs.

4.4. LOGFILE LOCATION

- /var/log/bdii/bdii-update.log: general BDII log file.
- /var/log/slapd.log: LDAP daemon log file if logging is activated (See Troubleshooting section for more details).

BDII top

- BDII_VAR_DIR/gip/tmp/gip/log/top-urls.conf(-glue2): glite-info-provider-ldap log file for a top BDII.
- /var/log/glite/glite-info-update-endpoints.log: glite-info-update-endpoints log file.

BDII site

- BDII_VAR_DIR/gip/tmp/gip/log/site-urls.conf(-glue2): glite-info-provider-ldap log file for a site BDII.

4.5. OPEN PORTS

2170

4.6. UNIT TESTS

<https://tomtools.cern.ch/confluence/display/IS/Tests>

4.7. SERVICE STATE

```
/etc/init.d/bdii status
```

For more detailed metrics on the status of the BDII:

```
ldapsearch -LLL -x -h localhost -p 2170 -b o=Infosys
```

4.8. CRON JOBS

BDII top

- generate-fcr-exclude-file: cron job that generates the file for the FCR mechanism.
- glite-info-update-endpoints: cron job that updates the URLs of the site BDII's that are being published by the top level BDII.

4.9. SECURITY INFORMATION

4.9.1 Access control mechanisms

The database is world readable but only the root user can make modifications.

4.9.2 How to block/ban a user

The database is world readable. All users can be blocked/banned by stopping the service.

4.9.3 Network usage

The slapd service running requires network access. The service can have a high network usage, due to the amount of information about the infrastructure it keeps and provides.

4.9.4 Firewall configuration

Port 2170/TCP needs to be open. The firewall configuration should allow to access this port from everywhere in the Internet. An example for IPTABLES configuration is:

```
-A RH-Firewall-1-INPUT -m state --state NEW -m tcp -p tcp --dport 2170 -j ACCEPT
```

4.9.5 Security recommendations

The service is based on OpenLDAP, as such, the security recommendations from OpenLDAP should be followed www.openldap.org.

4.9.6 Security incompatibilities

None

4.9.7 List of external packages that are not maintained by the supported OS

None

5. GLUE-VALIDATOR GUIDE

The glue-validator command is a very useful command for system administrators who want to validate whether the information published by the service they are managing is compliant with Glue 1.3 and Glue 2.0.

The glue-validator is installed by default with the BDII core product (package emi-resource-information-service). If it is not installed by default when installing a service, please run the following command:

```
yum install glue-validator
```

Check section 3.2.1 for more information on the software repositories that need to be used.

```
Usage: /usr/bin/glue-validator -t <test class> [OPTIONS]
-t --test The test class [glue1|glue2].
```

```
Server Mode: Obtains LDIF from an OpenLDAP server.
-h --host      Hostname of the LDAP server.
-p --port      Port for the LDAP server.
-b --bind      The bind point for the LDAP server.
```

```
File Mode: Obtains LDIF directly from a file.
-f --file      An LDIF file
```

```
Options:
-d --debug     Debug level 0-3, default 0
```

To test GLUE 1.3 information, the following command must be run, replacing the *site-bdii.host* and *SITE-NAME* with the hostname of the site BDII and the name of the site respectively.

```
glue-validator -h site-bdii.host -p 2170 -b mds-vo-name=SITE_NAME,o=grid -t glue1
```

To test GLUE 2.0 information, the following command must be run, replacing the *site-bdii.host* and *SITE-NAME* with the hostname of the site BDII and the name of the site respectively.

```
glue-validator -h site-bdii.host -p 2170 -b GLUE2DomainID=SITE_NAME,o=glue -t glue2
```

Both commands show a summary with the number of test run along with the number of errors found. The output includes an *AssertionError* for each Error found which should provide more details on the cause of the problem.

6. APPENDIX A: LDAP OPTIMISATIONS

A series of optimisations have been done in the LDAP configuration of all BDII flavours, both for the GLUE 1.3 and GLUE 2.0 databases. The optimisations are implemented using a series of directives in /etc/bdii/bdii-(top)-slapd.conf file or flags in /etc/bdii/DB_CONFIG. They are described below.

cachsize

The cachsize directive defines the number of entries that the LDAP backend will maintain in memory.

This directive has been defined to 300000 for a top BDII, and to 30000 for a site/resource BDII. In the case of a top BDII, this number is several orders of magnitude higher than the existing entries in the EGI and OSG infrastructures together. In the case of the site BDII, it is also several orders of magnitude higher than the existing entries in big sites like CERN.

checkpoint

The checkpoint directive defines the time between checkpoint operations in the database. A checkpoint operation flushes the database buffers to disk and writes a checkpoint record in the log. The frequency of checkpointing determines the time during which data may be unrecoverable by the database in the event of a system failure. The checkpoint will occur if either certain amount of data has been written or certain amount of minutes have passed since the last checkpoint.

For all BDII flavours and databases, this directive is defined as 1024 0, which means check point whenever 1024 KB are written. The *minutes* argument is ignored since it is zero.

dbnosync

The dbnosync directive specifies that the on-disk database contents do not need to be updated immediately with any in-memory records. This option will increase performance on writes but has the disadvantage that if there is a system failure before the disk and memory are re-synchronised then data may be lost. For this reason, this directive is used together with the checkpoint directive explained in the previous section.

For all BDII flavours and databases, this directive is defined. This directive is the same as the DB_TXN_NOSYNC flag, which is also present in all DB_CONFIG files.

Database Index

Database indexes improve the speed of data retrieval operations on a database at the cost of slower writes and increased storage space. Indexes consume memory (more indexes, more memory) and write or modify operations take longer due to index updates, therefore, only the attributes that are believed to be searched more often have been indexed. All indexed attributes have been declared as eq (EQUALITY rule, no wildcards are included), some as sub (SUBSTR rule, wildcards are included) and pres in the case of the *objectClass*.

The following table summarises the defined indexes for all flavours of BDII in GLUE 1.3 and GLUE 2.0 databases.

GLUE 1.3	GLUE 2.0
GlueCEAccessControlBaseRule eq	GLUE2GroupID eq
GlueCESEBindCEUniqueID eq	GLUE2ExtensionLocalID eq
GlueCESEBindSEUniqueID eq	GLUE2LocationID eq
GlueCEUniqueID eq	GLUE2ContactID eq
GlueChunkKey eq	GLUE2DomainID eq
GlueClusterUniqueID eq	GLUE2ServiceID eq
GlueSAAccessControlBaseRule eq	GLUE2EndpointID eq

GlueSALocalID eq GlueSEAccessProtocolType pres GlueSEUniqueID eq GlueServiceAccessControlRule eq GlueServiceAccessControlBaseRule eq GlueServiceType eq,sub GlueServiceEndpoint eq,sub GlueServiceURI eq,sub GlueServiceDataKey eq GlueSubClusterUniqueID eq GlueVOInfoAccessControlBaseRule eq objectClass eq,pres	GLUE2ShareID eq GLUE2ManagerID eq GLUE2ResourceID eq GLUE2ActivityID eq GLUE2PolicyID eq GLUE2BenchmarkID eq GLUE2ApplicationEnvironmentID eq GLUE2ApplicationHandleID eq GLUE2ToStorageServiceID eq GLUE2StorageServiceCapacityID eq GLUE2StorageAccessProtocolID eq GLUE2StorageShareSharingID eq GLUE2StorageShareCapacityID eq GLUE2EndpointInterfaceName eq GLUE2PolicyRule eq objectClass eq,pres
--	--

Relays and overlays

The purpose of the relay backend to slapd is to map a naming context defined in a database running in the same slapd instance into a virtual naming context with attributeType and objectClass manipulation using the rwm overlay. The rwm overlay performs basic DN/data rewrite and objectClass/attributeType mapping to provide virtual views of existing data.

This has been used in the GLUE 1.3 database of the top level BDII to define a *o=shadow* root on top of *o=grid*. This is indeed transparent because a relay has been used where *o=grid* actually messages *o=grid,o=shadow*. This has been implemented in this way after carrying out some performance testing where it has been observed that performance increases when the actual root of the DIT tree is defined in the second level of the tree.

7. APPENDIX B: EMIR INTEGRATION

Note: EMIR integration will be available in the EMI 3 release. Refer to the EMI web pages for more information: <http://www.eu-emi.eu/releases>.

In order to publish site and top BDII in EMIR, the following steps need to be executed:

1. Install EMIR service publisher in the top or site BDII host. This component is available in the EMI repository:

```
yum install emir-serp
```

2. In the YAIM site-info.def in the top or site BDII host, define the following variables:

```
BDII_EMIR_ENABLE=yes
```

```
BDII_EMIR_HOST=http://emir-hostname:port
```

3. Make sure the BDII is started.
4. Run the following YAIM command:

```
/opt/glite/yaim/bin/yaim -r -s site-info.def -n node_type -f config_emir_serp
```

After all these steps have been executed, the top or site BDII service will be published in the EMIR server. This can be checked by browsing <http://emir-hostname:port/services>.

8. APPENDIX C: ARC INTEGRATION

In order to publish ARC resources into the BDII tree, the following changes in the LDAP configuration file have been defined for all flavours of the BDII, that is /etc/bdii/bdii-slapd.conf and /etc/bdii/bdii-top-slapd.conf:

```
#####
# Relay DB to address DIT changes requested by ARC
#####

database      relay
suffix        "GLUE2GroupName=services,o=glue"
overlay       rwm
suffixmessage "GLUE2GroupID=resource,o=glue"

database      relay
suffix        "GLUE2GroupName=services,GLUE2DomainID=*,o=glue"
overlay       rwm
suffixmessage "GLUE2GroupID=resource,GLUE2DomainID=*,o=glue"

database      relay
suffix        "GLUE2GroupName=services,GLUE2DomainID=*,GLUE2GroupName=grid,o=glue"
overlay       rwm
suffixmessage "GLUE2GroupID=resource,GLUE2DomainID=*,GLUE2GroupID=grid,o=glue"
```

These changes define an alias for *GLUE2GroupID* and *resource*, so that *GLUE2GroupName* and *services* can be also used. For example, the following four bindings can now be used and are identical:

```
ldapsearch -LLL -x -h localhost -p 2170 -b "GLUE2GroupID=resource,o=glue"
ldapsearch -LLL -x -h localhost -p 2170 -b "GLUE2GroupName=resource,o=glue"
ldapsearch -LLL -x -h localhost -p 2170 -b "GLUE2GroupName=services,o=glue"
ldapsearch -LLL -x -h localhost -p 2170 -b "GLUE2GroupID=services,o=glue"
```

The *glite-info-provide-ldap* script has been also modified, to allow ARC resources to decouple the *GLUE2Domain* object from the published services, and at the same time keeping the existing BDII DIT having the ARC services published under a domain.

Note that these changes are only performed for GLUE 2.0 and they are backwards compatible with previous versions of the BDII.

9. APPENDIX D: YAIM VARIABLES

9.1. BDII BASICS

The following YAIM variables are defined for resource, site and top BDIIs.

Variable Name	Description	Example	Default Value
SITE_NAME	It defines the <code>GlueSiteUniqueID</code> attribute. It's the human-readable name of a site. It should be the same as the site name in the GOCDB.	<code>SITE_NAME=yaim-testbed</code>	To be defined by the sys admin.
BDII_BREATHE_TIME	Time in seconds between updates of the bdi.	<code>BDII_BREATHE_TIME=120</code>	120
BDII_READ_TIMEOUT	Amount of time in seconds to wait until an information is assumed to have timed out.	<code>BDII_READ_TIMEOUT=300</code>	300
BDII_ARCHIVE_SIZE	Number of dumps of the database to keep for debugging purposes.	<code>BDII_ARCHIVE_SIZE=0</code>	0
BDII_IPV6_SUPPORT	Variable that enables IPv6 support. If set to 'yes' it works with both IPv4 and IPv6. Possible values are 'yes' or 'no' (case sensitive).	<code>BDII_IPV6_SUPPORT=yes</code>	no

9.2. BDII TOP

The following YAIM variables are defined in a top BDII:

Variable Name	Description	Example	Default Value
BDII_HOST	Top BDII hostname. It should be a DNS alias if defined.	BDII_HOST=yaim-bdii.cern.ch	To be defined by the sys admin.
BDII_RAM_DISK	It defines whether a RAM disk is used to improve performance.	BDII_RAM_DISK=yes	Yes
SLAPD	Stand-alone LDAP daemon location.	SLAPD=/usr/sbin/slapd	/usr/sbin/slapd
SLAPD_CONF	Configuration file location for slapd.	SLAPD_CONF=/etc/bdii/bdii-top-slapd.conf	/etc/bdii/bdii-top-slapd.conf
BDII_DELETE_DELAY	Time to wait in seconds before deleting removed entries. Default is 12 hours. This variable activates the caching mode.	BDII_DELETE_DELAY=43200	43200

9.3. BDII SITE

The following YAAM variables are defined in a site BDII:

Variable Name	Description	Example	Default Value
SITE_BDII_HOST	Site BDII hostname. It should be a DNS alias if defined.	SITE_BDII_HOST=yaim-bdii.cern.ch	To be defined by the sys admin.
BDII_RAM_DISK	It defines whether a RAM disk is used to improve performance. Possible values are 'yes' or 'no' (case sensitive).	BDII_RAM_DISK= no	no
SLAPD	Stand-alone LDAP daemon location.	SLAPD=/usr/sbin/slapd	/usr/sbin/slapd
SLAPD_CONF	Configuration file location for slapd.	SLAPD_CONF=/etc/bdii/bdii-slapd.conf	/etc/bdii/bdii-slapd.conf
BDII_DELETE_DELAY	Time to wait in seconds before deleting removed entries. Default is 0.	BDII_DELETE_DELAY=0	0
BDII_REGIONS	List of host identifiers publishing information to the BDII. It should be one per host which runs a resource BDII. For each item listed in the BDII_REGIONS variable you need to create a BDII_<host-id>_URL variable.	BDII_REGIONS="CE1 CE2"	To be defined by the sys admin.
BDII_<host-id>_URL	URL of the information producer (e.g. BDII_host1_URL="ldap://host1_hostname:2170/mds-vo-name=resource,o=grid". Where host1 is a host where several node types may be installed, for example a lcg CE and a site BDII. It's therefore not necessary to create one variable per node type, but per host). It should be a real hostname, not an alias.	BDII_CE1_URL="ldap://information_producer_host_name:2170/mds-vo-name=resource,o=grid"	
SITE_DESC	Long format name of a site.	SITE_DESC="A long format description"	

Variable Name	Description	Example	Default Value
SITE_EMAIL	To set the GlueSiteEmailContact attribute. It's the main email contact for the site. The syntax is a mailto:URL.	<code>SITE_EMAIL="mailto:yaim-contact@cern.ch,admin-yaim@cern.ch"</code>	To be defined by the sys admin.
SITE_LOC	Location of the site BDII.	<code>SITE_LOC="City, Country"</code>	
SITE_LAT	To set the GlueSiteLatitude attribute. It's the position of the site north or south of the equator measured from -90° to 90° with positive values going north and negative values going south.	<code>SITE_LAT=46.20</code>	
SITE_LONG	To set the GlueSiteLongitude attribute. It's the position of the site east or west of Greenwich, England measured from -180° to 180° with positive values going east and negative values going west.	<code>SITE_LONG=6.1</code>	
SITE_WEB	Site web site.	<code>SITE_WEB="site-web-url"</code>	
SITE_SECURITY_EMAIL	Contact e-mail for security.	<code>SITE_SECURITY_EMAIL="sec@site.mail"</code>	
SITE_SUPPORT_EMAIL	Contact e-mail for support.	<code>SITE_SUPPORT_EMAIL="support@site.mail"</code>	
SITE_OTHER_*	Check the YAIM instructions in https://wiki.egi.eu/wiki/MAN01#Configuring_a_site_BDII to know how to configure these variables.		