



NORDUGRID-TECH-4

19/4/2013

THE NORDUGRID-ARC INFORMATION SYSTEM

Technical Description and Usage Manual

Balázs Kónya*

Daniel Johansson†

ARC version 1.0.0

*Comments to: balazs.konya@hep.lu.se

†Comments to: daniel@ndgf.org

Contents

| | | |
|----------|--|----------|
| 1 | Introduction | 6 |
| 2 | Overview | 6 |
| 2.1 | Operational overview of ARC LDAP Infosys components | 7 |
| 2.2 | Overview of ARC LDAP Infosys schemas | 8 |
| 3 | ARIS | 8 |
| 3.1 | Security considerations | 9 |
| 4 | The ARC information model | 9 |
| 4.1 | LDAP technicalities: namespace, OID, objectclasses, attributes | 9 |
| 4.2 | The structure of the local LDAP tree: Arc DIT | 10 |
| 4.3 | Globus Mds | 11 |
| | mds-validfrom | 11 |
| | mds-validto | 11 |
| | mds-keepsto | 11 |
| | mds-vo-name | 11 |
| | mds-vo-op-name | 11 |
| | mds-service-type | 12 |
| | mds-service-protocol | 12 |
| | mds-service-port | 12 |
| | mds-service-hn | 12 |
| | mds-service-url | 12 |
| | mds-service-ldap-suffix | 12 |
| | mds-service-ldap-timeout | 13 |
| | mds-service-ldap-sizelimit | 13 |
| | mds-service-ldap-cachettl | 13 |
| | mds-service-ldap-ttl | 13 |
| | mds-reg-status | 13 |
| | mds-bind-method-servers | 13 |
| 4.4 | Grouping authuser and job entries: nordugrid-info-group objectclass | 13 |
| | nordugrid-info-group-name | 13 |
| 4.5 | Computing resources: nordugrid-cluster and nordugrid-queue objectclass | 14 |
| | nordugrid-cluster-name | 14 |
| | nordugrid-cluster-aliasname | 14 |
| | nordugrid-cluster-contactstring | 15 |
| | nordugrid-cluster-support | 15 |
| | nordugrid-cluster-location | 15 |
| | nordugrid-cluster-owner | 15 |

| | |
|---|----|
| nordugrid-cluster-acl | 15 |
| nordugrid-cluster-issuerca | 16 |
| nordugrid-cluster-issuerca-hash | 16 |
| nordugrid-cluster-trustedca | 16 |
| nordugrid-cluster-credentialexpirationtime | 16 |
| nordugrid-cluster-lrms-type | 17 |
| nordugrid-cluster-lrms-version | 17 |
| nordugrid-cluster-lrms-config | 17 |
| nordugrid-cluster-homogeneity | 17 |
| nordugrid-cluster-architecture | 17 |
| nordugrid-cluster-opsys | 18 |
| nordugrid-cluster-nodecpu | 18 |
| nordugrid-cluster-benchmark | 18 |
| nordugrid-cluster-nodememory | 18 |
| nordugrid-cluster-totalcpus | 19 |
| nordugrid-cluster-cpudistribution | 19 |
| nordugrid-cluster-sessiondir-free | 19 |
| nordugrid-cluster-sessiondir-total | 19 |
| nordugrid-cluster-sessiondir-lifetime | 20 |
| nordugrid-cluster-cache-free | 20 |
| nordugrid-cluster-cache-total | 20 |
| nordugrid-cluster-runtimeenvironment | 20 |
| nordugrid-cluster-locale | 20 |
| nordugrid-cluster-middleware | 21 |
| nordugrid-cluster-totaljobs | 21 |
| nordugrid-cluster-usedcpus | 21 |
| nordugrid-cluster-queuedjobs | 21 |
| nordugrid-cluster-prelrmsqueued | 22 |
| nordugrid-cluster-nodeaccess | 22 |
| nordugrid-cluster-comment | 22 |
| nordugrid-cluster-interactive-contactstring | 22 |
| nordugrid-queue-name | 22 |
| nordugrid-queue-status | 23 |
| nordugrid-queue-comment | 23 |
| nordugrid-queue-schedulingpolicy | 23 |
| nordugrid-queue-homogeneity | 23 |
| nordugrid-queue-nodecpu | 24 |
| nordugrid-queue-nodememory | 24 |
| nordugrid-queue-architecture | 24 |
| nordugrid-queue-opsys | 24 |
| nordugrid-queue-benchmark | 25 |

| | |
|--|----|
| nordugrid-queue-maxrunning | 25 |
| nordugrid-queue-maxqueueable | 25 |
| nordugrid-queue-maxuserrun | 25 |
| nordugrid-queue-maxtotalcputime | 25 |
| nordugrid-queue-maxcputime | 26 |
| nordugrid-queue-mincputime | 26 |
| nordugrid-queue-defaultcputime | 26 |
| nordugrid-queue-maxwalltime | 26 |
| nordugrid-queue-minwalltime | 27 |
| nordugrid-queue-defaultwalltime | 27 |
| nordugrid-queue-running | 27 |
| nordugrid-queue-gridrunning | 27 |
| nordugrid-queue-queued | 27 |
| nordugrid-queue-gridqueued | 28 |
| nordugrid-queue-localqueued | 28 |
| nordugrid-queue-prelrmsqueued | 28 |
| nordugrid-queue-totalcpus | 28 |
| 4.6 Grid jobs: nordugrid-job objectclass | 29 |
| nordugrid-job-globalid | 29 |
| nordugrid-job-globalowner | 29 |
| nordugrid-job-jobname | 29 |
| nordugrid-job-execcluster | 29 |
| nordugrid-job-execqueue | 30 |
| nordugrid-job-executionnodes | 30 |
| nordugrid-job-submissionui | 30 |
| nordugrid-job-submissiontime | 30 |
| nordugrid-job-sessiondirerasetime | 31 |
| nordugrid-job-proxyexpirationtime | 31 |
| nordugrid-job-completiontime | 31 |
| nordugrid-job-runtimeenvironment | 31 |
| nordugrid-job-gmlog | 32 |
| nordugrid-job-clientsoftware | 32 |
| nordugrid-job-stdout | 32 |
| nordugrid-job-stderr | 32 |
| nordugrid-job-stdin | 32 |
| nordugrid-job-cpucount | 33 |
| nordugrid-job-reqcputime | 33 |
| nordugrid-job-reqwalltime | 33 |
| nordugrid-job-queuerank | 33 |
| nordugrid-job-lrmscomment | 33 |
| nordugrid-job-comment | 34 |

| | |
|--|-----------|
| nordugrid-job-usedcputime | 34 |
| nordugrid-job-usedwalltime | 34 |
| nordugrid-job-usedmem | 34 |
| nordugrid-job-exitcode | 35 |
| nordugrid-job-errors | 35 |
| nordugrid-job-status | 35 |
| nordugrid-job-rerunable | 37 |
| 4.7 Grid users: nordugrid-authuser objectclass | 37 |
| nordugrid-authuser-name | 37 |
| nordugrid-authuser-sn | 37 |
| nordugrid-authuser-freecpus | 38 |
| nordugrid-authuser-diskspace | 38 |
| nordugrid-authuser-queuelength | 38 |
| 4.8 Storage Resources: the nordugrid-se objectclass | 38 |
| nordugrid-se-name | 39 |
| nordugrid-se-aliasname | 39 |
| nordugrid-se-type | 39 |
| nordugrid-se-freespace | 39 |
| nordugrid-se-totalspace | 40 |
| nordugrid-se-url | 40 |
| nordugrid-se-location | 40 |
| nordugrid-se-owner | 40 |
| nordugrid-se-acl | 40 |
| nordugrid-se-issuerca | 41 |
| nordugrid-se-issuerca-hash | 41 |
| nordugrid-se-trustedca | 41 |
| nordugrid-se-middleware | 41 |
| nordugrid-se-comment | 41 |
| nordugrid-se-authuser | 42 |
| nordugrid-se-accesscontrol | 42 |
| 4.9 Other services: nordugrid-rc objectclass for data catalogues | 42 |
| 4.10 Mapping to the GLUE Information model | 42 |
| 4.10.1 The ComputingElement entity of the LCG Glue model | 43 |
| 4.10.2 The Cluster and SubCluster entity of the LCG Glue model | 47 |
| 4.10.3 The Site entity of the LCG Glue model | 50 |
| 5 Registration Processes, EGIIS: Topology | 52 |
| 5.1 Resource discovery | 54 |
| A Clients of the ARC information system | 55 |
| B Glue 1.2 with ARC, gLite interoperability | 56 |

1 Introduction

A stable, robust, scalable, dynamic and reliable information system is a cornerstone of any kind of Grid system. Without a properly working information system it is not possible to construct a production quality Grid. A scalable Grid information system is inherently distributed, a centralized system is not able to cope with the dynamism of the Grid.

The information system acts as a nervous system of the Grid and its main tasks consist of

- **Resource Description:** Characterization of Grid resources by specifying static, semi-static and dynamic properties (e.g. information about Grid jobs and user quotas are presented as dynamic local resource properties).

Grid clients are relying on the resource description functionality of the information system during their matchmaking and brokering process. Grid monitoring and job status queries also rely on resource description.

- **Resource aggregation:** Individual resources are connected to an "information mesh" by dynamically registering to some of the information index services. The information index services are responsible for the resource aggregation, they maintain a dynamic list of available Grid resources. Furthermore, the index services are connected to each other following a specific topological order. The resulting structure is the "information mesh" of the Grid.

Grid clients performing resource discovery scan the "information mesh" utilizing its topological structure in order to find available Grid resources. Therefore resource discovery is delicately coupled to the topological structure of the "information mesh" that is to the resource aggregation process.

The ARC middleware implements a scalable, production quality dynamic distributed information system. The ARC information system has been deployed and being used in a large scale production environment since May 2002. As of writing, the information system aggregates 70 resources providing 40 thousand CPUs and serves 400 thousand jobs per month.

LDAP-based ARC information system relies on OpenLDAP [2] and is derived from the Globus Monitoring and Discovery Services framework [3]. It uses BDII [17] together with a updated MDS LDAP Schema.

This document presents a technical overview of ARC's LDAP information system. It describes the architecture, implementation of the main components and the NorduGrid-ARC information model. The document is also intended to serve as reference manual by giving a detailed description of the available Grid information.

2 Overview

The ARC middleware implements a dynamic LDAP-based distributed information system via a set of coupled resource lists (index services) and local LDAP databases. The system consists of three main components:

1. the ARC Resource Information Service (ARIS),
2. the Enhanced Grid Information Indexing Service (EGIIS),
3. and the Registration Processes (RP)

ARIS instances are responsible for resource (computing or storage) description and characterization. The local information is generated on the resource, and it can be cached. Upon client requests it is presented via LDAP interface.

The main task of EGIIS services is to maintain a dynamic list of resources (LDAP URLs of the ARISes) and index services. The index services are further coupled together implementing a specific topology.

ARIS services make use of registration processes running locally on the resources in order to list themselves in some of the resource lists maintained by EGIIS services. Registrations are always initiated by the registrants (bottom-up model).

ARIS has to be present at each Grid resource, and is therefore essentially an integral part of a Grid service, like for example a Compute Element. EGIIS is a stand-alone service and is not coupled to any resource. While EGIIS is not needed for any particular service to function, there is no Grid without EGIIS: EGIIS instances aggregate information and thus create a coupled infrastructure from individual resources.

Grid clients such as monitors, Web portals or user interfaces perform two type of queries:

1. During the resource discovery process clients query EGIIS services in order to collect list of LDAP contact URLs of ARIS services describing Grid-connected resources.
2. During a direct resource query the clients directly contact each ARIS by making use of the obtained LDAP contact URLs.

Both type of queries are carried out and served via LDAP protocol.

Figure 1 presents an overview of the ARC information system components.



Figure 1: Overview of ARC information system components.

2.1 Operational overview of ARC LDAP Infosys components

ARC information system consists of a few different components that are to be split over a few different machines:

1. Client – is typically installed on the user's machine by users themselves, and can either be an OpenLDAP LDAP-client for presenting information in a human readable form, or an ARC client that is responsible for job submission or status querying. The standard ARC client does not output LDAP information directly, instead it tailors it to the task that was asked of it. For example, if a job status was queried, only information about that job will be shown to the user.

2. ARIS – is installed by system administrators that are responsible for the clusters that are connected to the Grid. Most importantly, it runs on the ARC Computing Element and is responsible for publishing information about the resource (cluster). More information about ARIS is available in Section 3. This service has to be configured to register to one or more EGIIS instances.
3. EGIIS – is normally installed as a standalone service by a system administrator. It can co-exist with any other service; when co-deployed with ARIS, it will make use of the same LDAP database instances. This service is responsible for aggregating information from multiple ARISes and other EGIISes, and can register this information to a higher level EGIIS. EGIISes can create a hierarchical topology of any depth. More information about EGIIS can be found in Section 5.

2.2 Overview of ARC LDAP Infosys schemas

ARC information system currently can present information in three different formats, or schemas. These can be enabled simultaneously. The schemas are:

1. NorduGrid-ARC schema – this is the NorduGrid default schema, described in detail in this document. It was inspired by Globus MDS, but has been improved a lot over the years and due to incompatible changes was moved into the NorduGrid LDAP namespace. If you want standard NorduGrid clients to submit jobs to your resource, you want to publish this schema.
2. Glue 1.2 – This is the schema that is used by gLite [6]. Currently, gLite support Glue 1.3 schema, but Glue 1.2 is sufficient to be compatible. If you configure ARC to publish information in the Glue 1.2 format, you will first produce data in NorduGrid-ARC schema which will then be translated to Glue 1.2 (see Section 4.10 for mapping details). If you want to allow gLite clients to submit to your resource, you want to publish this schema. Please note, that you will also need to hook in your ARC cluster into the gLite information system in order to get this interoperability to work. Some information about this is available in Appendix B.
3. Glue 2.0 – This is the schema that will become the common schema for the EMI [7]. This schema can be published both through LDAP and XML interfaces of ARC Compute Element.

3 ARIS

ARIS is the information service that is installed on the ARC Compute Element. It contains information about the local computing cluster, like: operating system, amount of main memory, computer architecture, information about running and finished jobs, users allowed to run and trusted certificate authorities. The information can be published in either NorduGrid-ARC schema, Glue 1.2 schema or Glue 2.0 schema.

The NorduGrid-ARC schema is the main ARC schema, which needs to be published in order for an ARC client to submit jobs to the resource.

The Glue 1.2 schema is the main schema of gLite, if you want to make your cluster compatible with gLite clients, then you will want to enable this schema. Please take a look at Appendix B.

The ARIS component of the information system is responsible for generating the dynamic state information, implementing the first-level caching of the local information and providing the requested Grid information to the clients through the LDAP protocol. ARIS is basically nothing more but a specially populated and customized OpenLDAP database.

The dynamic resource state information is generated on the resource. Small and efficient programs, called information providers, are used to collect local state information from the batch system, from the local Grid layer (e.g. A-REX [5], Grid Manager or GridFTP server [4]) or from the local operating system (e.g. information available in the `/proc` area). Currently, ARC is capable interfacing to the following batch systems (or local resource management system LRMS in the ARC terminology): UNIX fork, the PBS-family (OpenPBS, PBS-Pro, Torque), Condor, Sun Grid Engine, IBM LoadLeveler and SLURM.

The output of the information providers (generated in LDIF format) is used to populate the local LDAP tree. This OpenLDAP back-end implements two things: it is capable caching the providers output and upon client query request it triggers the information providers unless the data is already available in its cache.

The caching feature of the OpenLDAP back-end provides protection against overloading the local resource by continuously triggering the information providers.

The default information stored in ARIS follows the NorduGrid-ARC information model. The section 4 gives a detailed technical account of the ARC information model.

3.1 Security considerations

ARIS is implemented via an LDAP database which implies the security and confidentiality capabilities of the system.

OpenLDAP [2] contains two methods for specifying access control. The first is static, i.e. you define the rights in configuration files. From an operational point of view, the problem of this method is that needs a server restart at every security configuration change. The second method for access control, called as ACI (Access Control Information), inserts access control information inside the directory itself by augmenting every LDAP entry with a dynamically modifiable ACL. Unfortunately the ACI method is still considered to be experimental.

The current ARC setup makes use of the static LDAP access control, the trees are configured to be fully readable by anybody: ARC provides anonymous read access to every information stored in the local trees.

There are considerations to experiment with the ACI access control method or to modify the static configuration and require authentication from the clients.

4 The ARC information model

A Grid information model should be a result of a delicate design process how to represent the resources and what is the best way to structure this information.

ARC implements an LDAP-based information system. In an LDAP-based system the information is being stored as attribute-value pairs grouped together in entries which are organized into a hierarchical tree. Therefore an LDAP-based information model is technically specified via an LDAP schema AND the structure of the LDAP-tree (DIT).

The ARC information model naturally describes the main Grid components:

- computing resources with Grid jobs and Grid users,
- storage elements,
- and metadata catalogues

though the latter two are treated in a rather simplistic manner.

4.1 LDAP technicalities: namespace, OID, objectclasses, attributes

The NorduGrid-ARC LDAP schema (available in appendix D) makes use of the *nordugrid*- namespace, the objectclass and attribute names starts with the *nordugrid*- prefix.

NorduGrid is assigned to the 1.3.6.1.4.1.11604 Private Enterprise Number which is utilized according to the Table1.

The OID's used in the LDAP schema are shown in Table2 and are taken from the range 1.3.6.1.4.1.11604.2.* Table2 also serves as a list of the NorduGrid objectclasses. The ARC implementation follows a "one LDAP entry = one objectclass" approach. The ARC information system objects such as Grid-enabled clusters, queues, storages, Grid users and Grid jobs are described by specific LDAP entry which utilizes a single objectclass. As a result a one-to-one correspondence exists between ARC LDAP entries and ARC objectclasses.

The detailed description of the objectclasses and attributes are given in the following subsections. First the main purpose behind the objectclass is outlined followed by the one-by-one description of the attributes. The attribute descriptions also contain information about the attributes role played in the brokering[8], the

| Object Identifier | Service area |
|---------------------|--------------------|
| 1.3.6.1.4.1.11604.1 | security |
| 1.3.6.1.4.1.11604.2 | information system |
| 1.3.6.1.4.1.11604.3 | data management |
| 1.3.6.1.4.1.11604.4 | user management |

Table 1: The OID space utilization within ARC

| | |
|---------------------------|------------------------|
| 1.3.6.1.4.1.11604.2.1.1 | cluster objectclass |
| 1.3.6.1.4.1.11604.2.1.1.x | cluster attributes |
| 1.3.6.1.4.1.11604.2.1.2 | info-group objectclass |
| 1.3.6.1.4.1.11604.2.1.2.x | info-group attributes |
| 1.3.6.1.4.1.11604.2.1.3 | queue objectclass |
| 1.3.6.1.4.1.11604.2.1.3.x | queue attributes |
| 1.3.6.1.4.1.11604.2.1.4 | job objectclass |
| 1.3.6.1.4.1.11604.2.1.4.x | job attributes |
| 1.3.6.1.4.1.11604.2.1.5 | authuser objectclass |
| 1.3.6.1.4.1.11604.2.1.5.x | authuser attributes |
| 1.3.6.1.4.1.11604.2.1.6 | se objectclass |
| 1.3.6.1.4.1.11604.2.1.6.x | se attributes |
| 1.3.6.1.4.1.11604.2.1.7 | rc objectclass |
| 1.3.6.1.4.1.11604.2.1.7.x | rc attributes |
| 1.3.6.1.4.1.11604.2.1.8 | Mds objectclass |
| 1.3.6.1.4.1.11604.2.1.8.x | Mds attributes |

Table 2: The OIDs from the NorduGrid-ARC schema

job submission or the monitoring process. If applicable, the corresponding xRSL attribute[9] is displayed. Please notice that the most of the attribute values documented below are not enforced, misconfigured or rough sites can publish incorrect information.

4.2 The structure of the local LDAP tree: Arc DIT

The ARC local LDAP tree gives a natural representation of a Grid-enabled resource. The *mds-vo-name=local,o=grid* top level entry of the tree carries no information, its role is purely structural. This entry accommodates the subtrees of the different Grid services offered by the local machine. A Grid resource in ARC can host a computing service and several storage or data indexing services. A computing service is described by the cluster subtree, while the storage and data indexing services are characterised by the *se* and *rc* single-entry subtrees, respectively.

Figure 2 shows the local LDAP tree of two Grid-enabled resources. The first machine *bambi.hep.lu.se* offers both a computing service, a storage service and a data indexing service, therefore the ARIS of *bambi.hep.lu.se* contains a cluster subtree under the *nordugrid-cluster-name=bambi.hep.lu.se* entry a storage *nordugrid-se-name=.* and a data indexing *nordugrid-rc-name=.* entry. The second resource *hathi.hep.lu.se* serves as a dedicated storage hosting two storage elements, therefore the ARIS of *hathi.hep.lu.se* consists of the two storage entries.

The schematic structure of the cluster subtree is shown enlarged in Fig. 3. The *cluster* top entry of the subtree describes the hardware, software and middleware properties of a cluster. Grid-enabled queues are represented by their *queue* entries. Active Grid jobs and authorized Grid users are described by their corresponding *job* and *authuser* entries which are located under their hosting queues. The *job* and *authuser* entries belonging to the same queue are grouped in two distinct subtrees, the branching is accomplished by

structural *nordugrid-info-group=job* and *nordugrid-info-group=user* entries.

The storage and data indexing services are represented by their corresponding single LDAP entries, currently no LDAP subtree is associated to them.

4.3 Globus Mds

These are the globus-mds attributes that we have incorporated into arc. The schema contains special objectclasses: *Mds*, *MdsVo*, *MdsVoName*, *MdsServiceLdap* and *MdsService* whose role it is to create structural entries in the LDAP tree.

mds-validfrom

Attribute value: {time}
 Example: 20050307103026Z
 Related xRSL: none

mds-validto

Attribute value: {time}
 Example: 20050307103026Z
 Related xRSL: none

mds-keepsto

Attribute value: {time}
 Example: 20050307103026Z
 Related xRSL: none

mds-vo-name

Attribute value: {Locally unique VO name}
 Example: local
 Related xRSL: none



Figure 2: The local information tree on two resources. The first machine bambi.hep.lu.se provides both computing, storage and data indexing services while the second resource hathi.hep.lu.se hosts two storage elements.

mds-vo-op-name

Attribute value: {Locally unique Op name}
Related xRSL: none

mds-service-type

Attribute value: {Mds service type}
Example: ldap
Related xRSL: none

mds-service-protocol

Attribute value: {Service protocol OID}
Related xRSL: none

mds-service-port

Attribute value: {Service TCP port}
Related xRSL: none

mds-service-hn

Attribute value: {Service FQDN hostname}
Related xRSL: none

mds-service-url

Attribute value: {Service URL}
Related xRSL: none

mds-service-ldap-suffix



Figure 3: The schematic picture of an LDAP subtree representing a computing resource. The cluster subtree is part of ARIS shown in Fig: 2

Attribute value: {DN suffix of service}
Example: Mds-Vo-Name=local, o=grid
Related xRSL: none

mds-service-ldap-timeout

Attribute value: {time in minutes}
Related xRSL: none

mds-service-ldap-sizelimit

Attribute value: {size}
Related xRSL: none

mds-service-ldap-cachettl

Attribute value: {time in minutes}
Related xRSL: none

mds-service-ldap-ttl

Attribute value: {time in minutes}
Related xRSL: none

mds-reg-status

Attribute value: {status}
Example: VALID
Related xRSL: none

mds-bind-method-servers

Attribute value: {type of server}
Example: ANONYM-ONLY
Related xRSL: none

4.4 Grouping authuser and job entries: nordugrid-info-group objectclass

The schema contains a special objectclass the *nordugrid-info-group* whose role is to create structural entries in the LDAP tree. The entries *nordugrid-info-group=jobs* and *nordugrid-info-group=users* of ARIS are separating nordugrid-job and nordugrid-authuser entries of a grid queue by grouping them in two separate LDAP branches under the queue entry (see Fig. 2).

The objectclass comes with a single attribute.

nordugrid-info-group-name

Attribute value: {users,jobs}

Related xRSL: none

The *nordugrid-info-group-name* attribute is used to distinguish between jobs or users grouping: nordugrid-job entries are grouped under the structural entry *nordugrid-info-group-name=jobs* while nordugrid-authuser entries are grouped under the *nordugrid-info-group-name=users* entry.

4.5 Computing resources: nordugrid-cluster and nordugrid-queue objectclass

The nordugrid schema provides two objectclasses for the description of a computing resource. The *nordugrid-cluster* is used to describe general properties such as ownership, name, location, contact URL, pre-installed software environments, Grid scratch space, batch system, node properties such as architecture, CPUs, network connectivity. Dynamic cluster load information, such as number of queued/total jobs, is also part of the objectclass information.

The generalized concept of a computing queue plays a central role in ARC: queues are the job submission targets in an ARC-based Grid, during the brokering process clients select a Grid-enabled queue on a computing resource. An ARC queue represents either a traditional batch queue of a local resource management system (LRMS) such as the PBS or describes an entire LRMS when the LRMS does not support conventional queues (Condor and SGE is handled this way). The very special LRMS, the UNIX fork is also described as queue. The *nordugrid-queue* objectclass is designed to describe the generalized concept of a computing queue. Besides the usual queue-specific information (queue status and limits, number of running/queueing jobs) queue-level node attributes are also introduced to describe hardware/software characteristics of computing nodes assigned to a certain queue. Also notice that the schema makes possible the distinction of Grid and non-grid jobs being managed by the queue.

The ARC schema introduces cluster- and queue-level node attributes together with two homogeneity attributes in order to handle possible inhomogeneity within a computing resource. The schema is designed to be capable describing inhomogeneous resources with a queue-level homogeneity assumption but higher level inhomogeneity can also be treated with less accuracy.

In case of homogeneous nodes the *nordugrid-cluster-homogeneity=true* is set and the cluster-level node attributes carry the relevant information. If the nodes are inhomogeneous the *nordugrid-cluster-homogeneity=no* is set and the cluster-level node attributes are either not set or their value refers to the smallest/slowest/least powerful node. Suppose the nodes can be organized into homogeneous subgroups, this case the queue-level node attributes are used to describe the properties of the homogeneous nodes assigned to the same queue. Clients should always treat the queue-level node attributes with higher priority than the cluster-level ones. The *nordugrid-queue-homogeneity=true* attribute value is used to specify the node homogeneity within a queue. The *nordugrid-queue-homogeneity=no* means that if a given queue-level node attribute is set it refers to the smallest/slowest/least powerful node.

nordugrid-cluster-name

Attribute value: FQDN

Example: nordugrid-cluster-name: gate1.monstercluster.nordugrid.org

Related xRSL: cluster

UI role: used in matchmaking, monitoring, job manipulation

Description: The fully qualified domain name of the front-end machine of the cluster. This attribute is used in the Distinguished Name of a cluster LDAP entry.

nordugrid-cluster-aliasname

Attribute value: free form text
Example: `nordugrid-cluster-aliasname: Grid Monster`
Related xRSL: none
UI role: ?

Description: A free form text attribute displaying the alias name of the computing resource.

nordugrid-cluster-contactstring

Attribute value: URL
Example: `nordugrid-cluster-contactstring: gsiftp://bambi.hep.lu.se:2811/jobs`
Related xRSL: none
UI role: used during the job submission process

Description: The URL of the job submission service running on the cluster front-end. Clients must use this attribute to determine the URL of the job submission gateway available on the cluster.

nordugrid-cluster-support

Attribute value: RFC822 email address
Example: `nordugrid-cluster-support: help@gridcluster.gridcenter.org`
Related xRSL: none
UI role: none

Description: The support email address of the Grid-enabled computing resource, users are suggested to use this address in case they need to contact the site.

nordugrid-cluster-location

Attribute value: Postal ZIP code with two letter country prefix
Example: `nordugrid-cluster-location: SE-22100`
Related xRSL: none
UI role: none

Description: The geographical location of the cluster, preferably specified as a postal code with a two letter country prefix.

nordugrid-cluster-owner

Attribute value: free form text
Example: `nordugrid-cluster-owner: Danish Center for Grid Computing`
Example: `nordugrid-cluster-owner: Copenhagen University`
Related xRSL: none
UI role: none

Description: The multivalued attribute is used to display the owner of the resource.

nordugrid-cluster-acl

Attribute value: fixed form syntax
Example: `nordugrid-cluster-acl: VO:ATLAS`
Example: `nordugrid-cluster-acl: VO:developers.nordugrid.org`
Related xRSL: none
UI role: none

Description: The multivalued attribute is used to display authorization rules applied on a cluster. The attribute value follows a *fixed form syntax*. Current syntax is rather coarse-grained and primitive: a "VO:" prefix followed by a VO name means that the given VO is authorized on the cluster. Note that there are no standards for VO names.

nordugrid-cluster-issuerca

Attribute value: LDAP Distinguished Name
Example: `nordugrid-cluster-issuerca: /C=DE/O=GermanGrid/CN=GridKa-CA`
Related xRSL: none
UI role: used during the job submission, matchmaking

Description: The certificate issuer of the cluster, the DN of the CA which issued the host certificate is shown by the attribute.

nordugrid-cluster-issuerca-hash

Attribute value: Hash value of a CA certificate
Example: `nordugrid-cluster-issuerca-hash: 1f0e8352`
Related xRSL: none
UI role: ???

Description: The hash value of the certificate of the issuer CA of the cluster, the hash of the certificate of the CA which issued the host certificate used by the cluster is shown by the attribute.

nordugrid-cluster-trustedca

Attribute value: LDAP Distinguished Name
Example: `nordugrid-cluster-trustedca: /C=DE/O=GermanGrid/CN=GridKa-CA`
Example: `nordugrid-cluster-trustedca: /DC=ORG/DC=SEE-GRID/CN=SEE-GRID CA`
Related xRSL: none
UI role: ???

Description: The DN's trusted by the cluster are shown by this multivalued attribute.

nordugrid-cluster-credentialexpirationtime

Attribute value: GMT formatted time stamp
Example: `nordugrid-job-credentialexpirationtime: 20050222120449Z`
Related xRSL: none
UI role: ?

Description: The expiration date of the shortest living credential affecting the cluster's x509 environment in GMT time format.

nordugrid-cluster-lrms-type

Attribute value: {PBSPro, OpenPBS, torque, SGE,SGEEE, fork, Condor,ll,SLURM}

Example: nordugrid-cluster-lrms-type: torque

Related xRSL: none

UI role: it will be used in the brokering (not implemented yet)

Description: The type of the Local Resource Management System (LRMS) running on the cluster. ARC currently supports the PBS family, the SGE family, the Condor, IBM LoadLeveler, SLURM and the UNIX fork batch systems.

nordugrid-cluster-lrms-version

Attribute value: version string

Example: nordugrid-cluster-lrms-version: 1.0.1p5

Related xRSL: none

UI role: none

Description: The vendor specific version string of the Local Resource Management System. The original vendor-provided LRMS version string is displayed without any modification.

nordugrid-cluster-lrms-config

Attribute value: free form text

Example: Short parallel jobs are prioritised"

Related xRSL: none

UI role: none

Description: A free form text attribute for additional remarks on the LRMS setup of the cluster. The attribute is purely for 'human consumption'.

nordugrid-cluster-homogeneity

Attribute value: {True, False}

Example: nordugrid-cluster-homogeneity: False

Related xRSL: none

UI role: ?

Description: A logical flag indicating the homogeneity of the cluster nodes. The front-end is not needed to be homogeneous with the nodes. If the nodes are declared inhomogeneous on the cluster-level, then the the cluster-level node attributes are referring to the properties of the slowest/smallest/least powerful node.

nordugrid-cluster-architecture

Attribute value: {i686, x86_64, alpha, sun4u}
 Example: `nordugrid-cluster-architecture: i686`
 Related xRSL: `architecture`
 UI role: used in matchmaking/brokering

Description: This is a cluster-level node attribute describing the 'hardware type' of the nodes of the cluster. The 'hardware type' is defined as the output of the `uname -m` UNIX command.

nordugrid-cluster-opsys

Attribute value: fixed format string
 Example: `nordugrid-cluster-opsys: Redhat-7.2`
 Example: `nordugrid-cluster-opsys: Linux-2.4.21-mypatch`
 Example: `nordugrid-cluster-opsys: glibc-2.3.1`
 Related xRSL: `none`
 UI role: not yet used in the brokering

Description: The multivalued cluster-level node attribute is meant to describe the operating system of the computing nodes. The attribute describes the operating system via the specification of the software distribution. The same multivalued attribute is also used to specify the kernel or libc version in case those differ from the originally shipped ones. The attribute value follows a *fixed form syntax*: the distribution name is given as `distroname-version.number` where spaces are not allowed. *Kernel* and *libc* versions are specified according to a fixed form: `kernelname-version.number`, `libcname-version.number`.

nordugrid-cluster-nodecpu

Attribute value: fixed format string
 Example: `nordugrid-cluster-nodecpu: Dual AMD Athlon(tm) MP Processor 1800+ @ 1500 MHz`
 Related xRSL: `gridTime`
 UI role: used in the brokering process

Description: This cluster-level node attribute gives the CPU type information of the cluster nodes in a fixed format. The string is constructed as `CPU-model-name CPU-frequency MHZ`, where CPU-model-name and CPU-frequency are vendor specified values (on Linux systems the data is taken from the `/proc/cpuinfo`).

nordugrid-cluster-benchmark ***

Attribute value: fixed format string
 Example: `nordugrid-cluster-benchmark: SPECINT2000 @ 222`
 Example: `nordugrid-cluster-benchmark: SPECFP2000 @ 333`
 Related xRSL: `benchmark`
 UI role: used in brokering

Description[‡]: The multivalued cluster-level node attribute shows the performance of the computing nodes with respect to specified benchmarks. The attribute value follows a fixed syntax: the benchmark name and value is separated by "@".

[‡]not in a stable release yet

nordugrid-cluster-nodememory

Attribute value: a number showing the amount in MBs
Example: **nordugrid-cluster-nodememory: 900**
Related xRSL: **memory**
UI role: used in matchmaking/brokering

Description: The amount of node memory in MBs which can be guaranteed to be available for the application running on the node. Please note in most cases it is less than the physical memory installed in the nodes.

nordugrid-cluster-totalcpus

Attribute value: number
Example: **nordugrid-cluster-totalcpus: 60**
Related xRSL: **count**
UI role: used in matchmaking/brokering

Description: The total number of CPUs of the computing resource being controlled by the LRMS. It is possible that not all of them are available for Grid jobs (e.g. the cluster has a non-grid queue with dedicated nodes).

nordugrid-cluster-cpudistribution

Attribute value: fixed format string
Example: **nordugrid-cluster-cpudistribution: 1cpu:36 2cpu:7**
Related xRSL: **none**
UI role: **none**

Description: The CPU distribution over the nodes given in the form of *ncpu:m* where *n* is the number of CPUs per machine and *m* is the number of such computers, an example: *1cpu:3,2cpu:4,4cpu:1* represents a cluster with 3 single CPU machines, 4 dual machines and one computer with 4 CPUs.

nordugrid-cluster-sessiondir-free

Attribute value: number showing the amount in MBs
Example: **nordugrid-cluster-sessiondir-free: 447870**
Related xRSL: **disk**
UI role: used in matchmaking/brokering

Description: Each Grid job has a dedicated Grid scratch area called the session directory. This attribute shows the available free disk space in MBs for the session directories. As a minimum protection the broker compares the available disk space to the size of the uploadable input data and rejects the clusters with insufficient free space.

nordugrid-cluster-sessiondir-total

Attribute value: number showing the amount in MBs
Example: **nordugrid-cluster-sessiondir-total: 1447870**
Related xRSL: **none**
UI role: **?**

Description: The total disk space in MB allocated on the cluster to host the grid job's session directories.

nordugrid-cluster-sessiondir-lifetime

Attribute value: time interval specified in minutes
Example: **nordugrid-cluster-sessiondir-lifetime: 10080**
Related xRSL: none
UI role: ?

Description: The lifetime of the job's session directory after the job has completed. The job's session directory together with all the user's data is erased when the *nordugrid-cluster-sessiondir-lifetime* has expired counted from the completion of the job.

nordugrid-cluster-cache-free

Attribute value: number showing the amount in MBs
Example: **nordugrid-cluster-cache-free: 2048**
Related xRSL: **disk**
UI role: used in matchmaking/brokering process

Description: ARC clusters can provide a cache area to store frequently used input data. Upon user request the input data is placed into the cache instead of the session directory of the job (input data in a session directory is not accessible by a consequent jobs). This attribute shows the available space in the cache in MBs.

nordugrid-cluster-cache-total

Attribute value: number showing the amount in MBs
Example: **nordugrid-cluster-cache-total: 8048**
Related xRSL: none
UI role: none

Description: The total space in MBs allocated for the cache service.

nordugrid-cluster-runtimeenvironment

Attribute value: Runtime Environment string[10]
Example: **nordugrid-cluster-runtimeenvironment: APPS/MODELCHECK/DUPPAAL**
Related xRSL: **runtimeenvironment**
UI role: used in matchmaking

Description: Runtime Environments are special pre-installed and pre-configured software packages provided in a standard way by the computing resources. A Runtime Environment Registry [10] maintains a list of available REs with pointers to RE descriptions. The multivalued attribute is used to display the REs available and supported on the cluster.

nordugrid-cluster-locale

Attribute value: URL
 Example: `nordugrid-cluster-locale: gsiftp://hypatia.uio.no/scratch/`
 Related xRSL: none
 UI role: used in brokering

Description: This multivalued parameter tells the broker that certain storage URLs should be considered "locally" available on the cluster. The attribute gives the URL of storage elements considered to be "local" to the cluster.

nordugrid-cluster-middleware

Attribute value: free form string to represent a software package
 Example: `nordugrid-cluster-middleware: nordugrid-0.4.4`
 Example: `nordugrid-cluster-middleware: globus-2.4.3-15ng`
 Related xRSL: `middleware`
 UI role: used in matchmaking

Description: This multivalued attribute specifies the middleware packages installed on the cluster.

nordugrid-cluster-totaljobs

Attribute value: number
 Example: `nordugrid-cluster-totaljobs: 580`
 Related xRSL: none
 UI role: ?

Description: The total number of non-completed jobs in the cluster. Totaljobs includes both Grid and non-grid jobs, non-grid jobs are those batch jobs which are directly submitted to the LRMS by a local user. Grid jobs with *FINISHING*, *FINISHED*, *DELETED* status are discarded.

nordugrid-cluster-usedcpus

Attribute value: number
 Example: `nordugrid-cluster-usedcpus: 12`
 Related xRSL: none
 UI role: ?

Description: The total number of occupied CPUs in the cluster. The attribute displays the number of busy/used CPUs regardless if the CPU is occupied by a Grid or a non-grid job.

nordugrid-cluster-queuedjobs

Attribute value: number
 Example: `nordugrid-cluster-queuedjobs: 812`
 Related xRSL: none
 UI role: ?

Description: The total number of jobs (grid and non-grid) not-yet running: preparing (e.g. Grid stage-in process) or waiting to run on a cluster. A Grid job submitted to the cluster needs to complete several stages

before it arrives to the LRMS. All these 'pre-LRMS Grid jobs' plus the LRMS queuing jobs are taken into account in the *nordugrid-cluster-queuedjobs* attribute. WARNING: The attribute is DEPRECATED in the 0.6 release!

nordugrid-cluster-prelrmsqueued

Attribute value: number
Example: **nordugrid-cluster-prelrmsqueued: 423**
Related xRSL: none
UI role: ?

Description: The total number of grid jobs not-yet reached the LRMS. These jobs are being processed or put on hold by the grid layer of the cluster. A Grid job submitted to the cluster needs to complete several stages before it arrives to the LRMS. All these 'pre-LRMS Grid jobs' are taken into account in the *nordugrid-cluster-prelrmsqueued* attribute.

nordugrid-cluster-nodeaccess

Attribute value: {inbound,outbound}
Example: **nordugrid-cluster-nodeaccess: inbound**
Example: **nordugrid-cluster-nodeaccess: outbound**
Related xRSL: **nodeaccess**
UI role: used in matchmaking

Description: The inbound/outbound network accessibility of the nodes determines how the nodes can connect to the Internet: *outbound* access means the nodes can connect to the outside world while *inbound* access means the nodes can be connected from outside. Specifying both *inbound*, *outbound* means the nodes are sitting on an open network. If a cluster has not set this attribute then the nodes are assumed to be sitting on a private isolated network.

nordugrid-cluster-comment

Attribute value: free form text
Example: **nordugrid-cluster-comment: This cluster is dedicated for CMS calculations**
Related xRSL: none
UI role: none

Description: The free form attribute displays some additional information about the cluster. Sometimes it contains an URL where more information can be read about the cluster.

nordugrid-cluster-interactive-contactstring

Attribute value: URL
Example: **nordugrid-cluster-interactive-contactstring: gsissh://atlas.hpc.unimelb.edu.au:2200**
Related xRSL: none
UI role: ?

Description: The URL for interactive login to the cluster. Some clusters offer GSI-enabled ssh services, this attribute presents the URL of that service.

nordugrid-queue-name

Attribute value: string representing a queue name
Example: `nordugrid-queue-name: longqueue`
Related xRSL: `queue`
UI role: used during job submission

Description: The name of the Grid-enabled batch queue. The special value *fork* is used for the 'UNIX fork' system. This attribute constitutes the Distinguished Name of a queue LDAP entry.

nordugrid-queue-status

Attribute value: `{active
inactive
inactive, grid-manager does not accept new jobs
inactive, grid-manager is down
inactive, gridftp is down}`
Example: `nordugrid-queue-status: inactive, grid-manager is down`
Related xRSL: `none`
UI role: used in brokering

Description: The generalized status of the queue. Besides the usual batch system queue status the attribute also takes into account the status of the Grid services such as the *grid-manager* and the *gridftp server* serving the queue. Grid jobs are only submitted to queues with *active* status.

nordugrid-queue-comment

Attribute value: free form text
Example: `nordugrid-queue-comment: Special queue dedicated to BIO Apps`
Related xRSL: `none`
UI role: `none`

Description: A free form attribute containing additional information about the queue.

nordugrid-queue-schedulingpolicy

Attribute value: free form text
Example: `nordugrid-queue-schedulingpolicy: SIMPLE FIFO`
Related xRSL: `none`
UI role: `none`

Description: The attribute is used to describe the implied scheduling policy of the queue (i.e. FIFO).

nordugrid-queue-homogeneity

Attribute value: `{True, False}`
Example: `nordugrid-queue-homogeneity: False`
Related xRSL: `none`
UI role: `?`

Description: A logical flag indicating the homogeneity of the queue nodes. If the nodes are declared inhomogeneous on the queue-level, then the queue-level node attributes are referring to the properties of the slowest/smallest/least powerful node within the queue.

nordugrid-queue-nodecpu

Attribute value: fixed format string

Example: `nordugrid-queue-nodecpu: Dual AMD Athlon(tm) MP Processor 1800+ @ 1500 MHz`

Related xRSL: `gridTime`

UI role: used in brokering

Description: This queue-level node attribute gives the CPU type information of the queue nodes in a fixed format. The string is constructed as `CPU-model-name CPU-frequency MHZ`, where CPU-model-name and CPU-frequency are vendor specified values (on Linux systems the data is taken from the `/proc/cpuinfo`).

nordugrid-queue-nodememory

Attribute value: a number showing the amount in MBs

Example: `nordugrid-queue-nodememory: 600`

Related xRSL: `memory`

UI role: used in matchmaking/brokering

Description: The queue-level node attribute shows the amount of node memory in MBs which can be guaranteed to be available for the application running on the node. Please note in most cases it is less than the physical memory installed in the nodes.

nordugrid-queue-architecture

Attribute value: `{i686, x86_64, alpha, sun4u}`

Example: `nordugrid-queue-architecture: x86_64`

Related xRSL: `architecture`

UI role: used in matchmaking

Description: This is a queue-level node attribute describing the 'hardware type' of the nodes of the queue. The 'hardware type' is defined as the output of the `uname -m` unix command.

nordugrid-queue-opsys

Attribute value: fixed format string

Example: `nordugrid-queue-opsys: Redhat-7.2`

Example: `nordugrid-queue-opsys: Linux-2.4.21-mypatch`

Example: `nordugrid-queue-opsys: glibc-2.3.1`

Related xRSL: `none`

UI role: not yet used in brokering

Description: The multivalued queue-level node attribute is meant to describe the operating system of the computing nodes. The attribute describes the operating system via the specification of the software distribution. The same multivalued attribute is also used to specify the kernel or libc version in case those differ from the originally shipped ones. The attribute value follows a *fixed form syntax*: the distribution name

is given as `distroname-version.number` where spaces are not allowed. *Kernel* and *libc* versions are specified according to a fixed form: `kernelname-version.number`, `libcname-version.number`.

nordugrid-queue-benchmark ***

Attribute value: fixed format string
Example: `nordugrid-queue-benchmark: SPECINT2000 @ 111`
Example: `nordugrid-queue-benchmark: SPECFP2000 @ 555`
Related xRSL: `benchmark`
UI role: used in brokering

Description[§]: The multivalued queue-level node attribute shows the performance of the computing nodes with respect to specified benchmarks. The attribute value follows a fixed syntax: the benchmark name and value is separated by "@".

nordugrid-queue-maxrunning

Attribute value: number
Example: `nordugrid-queue-maxrunning: 120`
Related xRSL: `none`
UI role: ?

Description: The batch queue limit indicating the maximum number of jobs allowed to run from this queue.

nordugrid-queue-maxqueueable

Attribute value: number
Example: `nordugrid-queue-maxqueueable: 500`
Related xRSL: `none`
UI role: ?

Description: The batch queue limit indicating the maximum number of jobs allowed to reside in the queue (both queuing and running).

nordugrid-queue-maxuserrun

Attribute value: number
Example: `nordugrid-queue-maxuserrun: 12`
Related xRSL: `none`
UI role: ?

Description: The batch queue limit indicating the maximum number of jobs a user can run at the same time in the queue.

nordugrid-queue-maxtotalcputime

[§]not in a stable release yet

Attribute value: number showing the time interval in minutes
Example: `nordugrid-queue-maxtotalcputime: 120`
Related xRSL: `cpuTime`
UI role: used in matchmaking, relevant for parallel jobs

Description: The batch queue limit giving the maximum total CPU time (in minutes) a job can use/request within this queue. The total is calculated over all processes belonging to the job (on all nodes, in case of parallel jobs). Only published on clusters that support such a limit.

nordugrid-queue-maxcputime

Attribute value: number showing the time interval in minutes
Example: `nordugrid-queue-maxcputime: 120`
Related xRSL: `cpuTime`
UI role: used in matchmaking

Description: The batch queue limit giving the maximum CPU time per CPU (in minutes) a job can use/request within this queue. If `nordugrid-queue-maxtotalcputime` is also published, clients should ignore `nordugrid-queue-maxcputime`.

nordugrid-queue-mincputime

Attribute value: number showing the time interval in minutes
Example: `nordugrid-queue-mincputime: 10`
Related xRSL: `cpuTime`
UI role: used in matchmaking

Description: The queue limit giving the lower value of job CPU time requests in minutes allowed in the queue.

nordugrid-queue-defaultcputime

Attribute value: number showing the time interval in minutes
Example: `nordugrid-queue-defaultcputime: 70`
Related xRSL: `cpuTime`
UI role: ?

Description: The default CPU time assigned to this queue in minutes. Jobs not specifying their CPU time requests are set to this default CPU time value by the LRMS.

nordugrid-queue-maxwalltime

Attribute value: number showing the time interval in minutes
Example: `nordugrid-queue-maxwalltime: 140`
Related xRSL: `cpuTime`
UI role: ???

Description: The batch queue limit gives the maximum walltime (in minutes) a job can use/request within this queue.

Comment: The *nordugrid-queue-maxwalltime* attribute is introduced in the 0.6.1 ARC release.

nordugrid-queue-minwalltime

Attribute value: number showing the time interval in minutes

Example: **nordugrid-queue-minwalltime: 30**

Related xRSL: **cpuTime**

UI role: ?

Description: The queue limit giving the lower value of job walltime requests in minutes allowed in the queue.

Comment: The *nordugrid-queue-minwalltime* attribute is introduced in the 0.6.1 ARC release.

nordugrid-queue-defaultwalltime

Attribute value: number showing the time interval in minutes

Example: **nordugrid-queue-defaultwalltime: 90**

Related xRSL: **cpuTime**

UI role: ?

Description: The default walltime assigned to this queue in minutes. Jobs not specifying their CPU time requests are set to this default CPU time value by the LRMS.

Comment: The *nordugrid-queue-defaultwalltime* attribute is introduced in the 0.6.1 ARC release.

nordugrid-queue-running

Attribute value: number

Example: **nordugrid-queue-running: 14**

Related xRSL: **none**

UI role: ?

Description: The attribute gives the number of CPUs being occupied by running jobs in the queue including both the Grid and non-Grid jobs. Multi-node jobs are counted with their multiplicity: a four-node running job increases the value of *nordugrid-queue-running* by four.

nordugrid-queue-gridrunning

Attribute value: number

Example: **nordugrid-queue-gridrunning: 6**

Related xRSL: **none**

UI role: ?

Description: The attribute gives the number of CPUs currently being occupied by running Grid jobs in the queue. Multi-node Grid jobs are counted with their multiplicity: a four-node running job increases the value of *nordugrid-queue-running* by four.

nordugrid-queue-queued

Attribute value: number
Example: **nordugrid-queue-queued: 23**
Related xRSL: none
UI role: ?

Description: The attribute gives the number of jobs, including both Grid and non-Grid, waiting in the queue. Each queuing job counts as one regardless their multiplicity. **WARNING: The attribute is DEPRECATED in the 0.6 release!**

nordugrid-queue-gridqueued

Attribute value: number
Example: **nordugrid-queue-gridqueued: 11**
Related xRSL: none
UI role: ?

Description: The attribute gives the number of waiting Grid jobs in the batch queue. Each queuing job counts as one regardless their multiplicity.

nordugrid-queue-localqueued

Attribute value: number
Example: **nordugrid-queue-localqueued: 24**
Related xRSL: none
UI role: ?

Description: The attribute gives the number of locally submitted non-Grid jobs waiting in the queue. Each queuing job counts as one regardless their multiplicity.

nordugrid-queue-prelrmsqueued

Attribute value: number
Example: **nordugrid-queue-prelrmsqueued: 25**
Related xRSL: none
UI role: ?

Description: The attribute gives the number of Grid jobs belonging to this queue being processed or waiting in the Grid-layer before the LRMS submission. Each queuing job counts as one regardless their multiplicity.

nordugrid-queue-totalcpus

Attribute value: number
Example: **nordugrid-queue-totalcpus: 11**
Related xRSL: ?
UI role: ?

Description: Some of the batch systems provides the possibility of assigning nodes to queues. This attribute shows the total number of CPUs exclusively dedicated to the queue within such batch system.

4.6 Grid jobs: nordugrid-job objectclass

In the NorduGrid-ARC information system every Grid job submitted to a Grid-enabled resource is represented by a *job* entry. Job entries are generated and optionally cached in the local LDAP tree of the hosting resource. This implies that job information within ARC is coupled to the execution Grid resource, namely for job status query or job monitoring the LDAP server of the hosting resource has to be contacted, this way ARC implements a fully distributed job status monitoring system: no central database or service is used for job status query/monitoring.

A job entry is generated and stored in the local LDAP tree for every existing Grid job on a resource. The job entry is kept in the local LDAP tree as long as the job is handled by the resource, when a job is removed from a resource the corresponding job entry is also deleted from the local LDAP tree. This implies that the ARC information system contains no information about non-existing deleted Grid jobs, another ARC service, the logging service is designed to store historical job information [11].

Job monitoring and status query of active Grid jobs is entirely based upon the LDAP job entries stored in the local information trees. Job entries carry information collected from the grid layer running on the resource (read from the job control files of the ARC Grid manager) and from the LRMS system. The attributes of the *nordugrid-job* objectclass are designed to provide all the necessary information.

nordugrid-job-globalid

Attribute value: URL
Example: `gsiftp://farm.hep.lu.se:2811/jobs/243361109008699845213642`
Related xRSL: none
UI role: used as a job handle in job management and as an URL in data movement

Description: ARC uses a GridFTP URL as a unique global jobID. The globally unique GridFTP URL is used as a handle in job manipulations such as rerun, kill or output retrieval. The GridFTP URL can also be used to access the session directory of the Grid job during the job's existence on the resource. The *nordugrid-job-globalid* attribute constitutes to the DN of the job entry.

nordugrid-job-globalowner

Attribute value: LDAP Distinguished Name
Example: `/O=Grid/O=NorduGrid/OU=nordugrid.org/CN=Lars_Jenssen`
Related xRSL: none
UI role: used during the job discovery process of *ngsync*

Description: The LDAP Subject Name of the job owner as specified in his/her Grid credentials. A Grid user or a client can easily find his/her own jobs on the Grid-enabled resource by issuing an LDAP search with a filter of *nordugrid-job-globalowner=his/her SN*.

nordugrid-job-jobname

Attribute value: free form text
Example: `nordugrid-job-jobname: ngtest-job-80`
Related xRSL: `jobname`
UI role: *ngget* optionally makes use of it

Description: The job name specified by the user with the *jobname* xRSL attribute. The client tools optionally can use the user-specified job name as the name of the local copy of the session directory of the job.

nordugrid-job-execcluster

Attribute value: FQDN
Example: `nordugrid-job-execcluster: farm.hep.lu.se`
Related xRSL: `cluster`
UI role: ?

Description: The name of the execution cluster specified as the fully qualified domain name of the front-end machine.

nordugrid-job-execqueue

Attribute value: string representing a queue name
Example: `nordugrid-job-execqueue: fastq`
Related xRSL: `queue`
UI role: ?

Description: The name of the execution queue hosting the Grid job. Within ARC the queues are coupled to clusters and used together as submission targets. Therefore the execution queue is selected together with the executing cluster during the brokering process, which means that the value of the *nordugrid-job-execqueue* is known for all the accepted Grid jobs even if they are not yet handed over to the local batch system. Also recall that Grid job entries are linked under their hosting queue entries in the local LDAP tree.

nordugrid-job-executionnodes

Attribute value: string representing a node name
Example: `nordugrid-job-executionnodes: n3`
Example: `nordugrid-job-executionnodes: n4`
Example: `nordugrid-job-executionnodes: n5`
Related xRSL: `none`
UI role: `none`

Description: The multivalued attribute presents the local node names of the cluster nodes which are occupied by the running Grid job. Every node being used by the job is listed with an attribute value pair. The shown example corresponds to a 3-node-job running on the nodes n3,n4,n5. Obviously, the *nordugrid-job-executionnodes* attribute is only available for jobs being run or already completed in the local batch system.

nordugrid-job-submissionui

Attribute value: fixed format string
Example: `nordugrid-job-submissionui: 130.235.91.118:45447;guest4.hep.lu.se`
Related xRSL: `none`
UI role: ?

Description: The attribute specifies client machine from where the job was submitted in a fixed format string. The string contains the submission host's IP, the port and the host name.

nordugrid-job-submissiontime

Attribute value: GMT formatted time stamp
Example: nordugrid-job-submissiontime: 20050220155311Z
Related xRSL: none
UI role: ?

Description: The time stamp of the submission of the job specified in Globus MDS time format (GMT). Job submission is the process when the client handles over the job request to the selected resource and the resource returns a job handle (the globally unique job ID).

nordugrid-job-sessiondirerasetime

Attribute value: GMT formatted time stamp
Example: nordugrid-job-sessiondirerasetime: 20050220165311Z
Related xRSL: lifeTime
UI role: none

Description: Within an ARC Grid every Grid job is confined to a dedicated area on the execution cluster which is called the *session directory*. After job completion the *session directory* of the grid job contains all the job and debugging output which was not requested to be uploaded to a storage element. The *session directory* can be accessed and the output data within the directory be downloaded for a limited time after the job completion. The date when the *session directory* is removed from the cluster is given in GMT time format by the *nordugrid-job-sessiondirerasetime* attribute.

nordugrid-job-proxyexpirationtime

Attribute value: GMT formatted time stamp
Example: nordugrid-job-proxyexpirationtime: 20050222120449Z
Related xRSL: none
UI role: ?

Description: The expiration time of the proxy assigned to the job displayed in GMT time format. A valid proxy is required for the stage-out phase of the Grid job if the stage out target makes use of GSI-based authentication and authorization. Intelligent clients can use this attribute to check if the job possesses a valid proxy and automatically initiate proxy renewal in case a proxy expiration.

nordugrid-job-completiontime

Attribute value: GMT formatted time stamp
Example: nordugrid-job-completiontime: 20050222120449Z
Related xRSL: none
UI role: ?

Description: The completion time of the Grid job expressed in GMT time format. Job completion refers to the *FINISHED* job state when the job completed all the requested operations including both job execution and stage out.

nordugrid-job-runtimeenvironment

Attribute value: string representing a valid RuntimeEnvironment
Example: `nordugrid-job-runtimeenvironment: APPS/CHEM/DALTON-1.2.1-1.0`
Related xRSL: `runtimeenvironment`
UI role: none

Description: The multivalued attribute lists the RuntimeEnvironments requested by the job.

nordugrid-job-gmlog

Attribute value: string representing a directory name
Example: `nordugrid-job-gmlog: Grid_manager_logdir`
Related xRSL: `gmlog`
UI role: optionally used for status monitoring

Description: The name of the directory which contains the Grid session related logs within the *session directory* of the job. The *gmlog* directory contains plenty of useful information for tracking or debugging the Grid job being processed on the execution site.

nordugrid-job-clientsoftware

Attribute value: string
Example: `nordugrid-job-clientsoftware: nordugrid-0.5.21`
Related xRSL: none
UI role: none

Description: The client software which was used to submit the job. The client software needs to be able to communicate its version to the Grid layer of the resource in order to have this attribute set.

nordugrid-job-stdout

Attribute value: string representing a file name
Example: `nordugrid-job-stdout: JG.out`
Related xRSL: `stdout`
UI role: ?

Description: The name of the file which contains the standard output of the job.

nordugrid-job-stderr

Attribute value: string representing a file name
Example: `nordugrid-job-stderr: JG.err`
Related xRSL: `stderr`
UI role: ?

Description: The name of the file which contains the standard error of the job.

nordugrid-job-stdin

Attribute value: string representing a file name
Example: `nordugrid-job-stdin: my.job_input`
Related xRSL: `stdin`
UI role: ?

Description: The name of the file which is used as the standard input of the job.

nordugrid-job-cpucount

Attribute value: number
Example: `nordugrid-job-cpucount: 7`
Related xRSL: `count`
UI role: none

Description: The number of CPUs requested by the job.

nordugrid-job-reqcputime

Attribute value: number showing the time interval in minutes
Example: `nordugrid-job-reqcputime: 146`
Related xRSL: `cpuTime`
UI role: none

Description: The CPU time request by the job specified in minutes.

nordugrid-job-reqwalltime ***

Attribute value: number showing the time interval in minutes
Example: `nordugrid-job-reqwalltime: 146`
Related xRSL: `wallTime`
UI role: none

Description[¶]: The wallclock time request of the job specified in minutes.

nordugrid-job-queuerank

Attribute value: number
Example: `nordugrid-job-queuerank: 13`
Related xRSL: none
UI role: the information can be used to initiate resubmission

Description: The attribute displays the queue position of the Grid job being idle in a batch queue. Most of the cases the given value is rather approximate since the majority of schedulers are not able to provide accurate information.

[¶]not in a stable release yet

nordugrid-job-lrmscomment

Attribute value: free form text

Example: `nordugrid-job-lrmscomment: Job is not running no available resources`

Related xRSL: none

UI role: none

Description: The optional comment provided by the Local Resource Management System. The attribute is only available in version 0.4.x, it was replaced by the more general *nordugrid-job-comment*.

nordugrid-job-comment

Attribute value: free form text

Example: `nordugrid-job-comment: LRMS: Job is not running no available resources`

Example: `nordugrid-job-comment: GM: The grid-manager is down`

Related XRSL: none

UI role: none

Description: The multivalued attribute contains the optional job comments provided by either the Grid Layer (e.g. Grid Manager) or the Local Resource Management System. The attribute has been introduced as a more general replacement of the *nordugrid-job-lrmscomment* and available with the ARC 0.6 version.

nordugrid-job-usedcputime

Attribute value: number showing the time interval in minutes

Example: `nordugrid-job-usedcputime: 144`

Related xRSL: none

UI role: none

Description: The consumed CPU time of the job in minutes as it was reported by the local batch system.

nordugrid-job-usedwalltime ***

Attribute value: number showing the time interval in minutes

Example: `nordugrid-job-usedwalltime: 166`

Related xRSL: none

UI role: none

Description^{||}: The consumed wall clock time of the job in minutes as it was reported by the local batch system.

nordugrid-job-usedmem

Attribute value: number representing memory consumption in KBs

Example: `nordugrid-job-usedmem: 4376`

Related xRSL: none

UI role: none

^{||}not in a stable release yet

Description: The memory usage of the job reported in KBs.

nordugrid-job-exitcode ***

Attribute value: number
Example: nordugrid-job-exitcode: 127
Related xRSL: none
UI role: used in job status monitoring

Description**: The exit code of the executable of the job obtained from the Local Resource Management System.

nordugrid-job-errors

Attribute value: free form text
Example v0.4.x: nordugrid-job-errors: JOB FAILURE: Failed extracting LRMS ID due to internal error
Example v0.6: nordugrid-job-errors: Failed extracting LRMS ID due to internal error
Related xRSL: none
UI role: used in job status monitoring

Description: Textual explanation of the job's failure, error message provided by the Grid layer running on the resource.

0.4.x release implementation: This attribute was/is used to determine the failure of the job, the presence of the attribute indicates a failure in the Grid job execution. The attribute text starts with the *JOB FAILURE:* prefix.

0.6 release implementation: The *JOB FAILURE:* prefix is dropped from the attribute value, new *FINISHED*, *KILLED*, *FAILED* final states were introduced, clients don't have to rely on the *nordugrid-job-errors* any longer to determine job failure.

nordugrid-job-status

Attribute value v0.4.x: {ACCEPTED,PREPARING,SUBMITTING,INLRMS: X,FINISHING,FINISHED,CANCELLING,PENDING:ACCEPTED,PENDING:PREPARING,PENDING:INLRMS,DELETED}
Attribute value v0.6: {ACCEPTING,ACCEPTED,PREPARING,PREPARED,SUBMITTING,INLRMS:X KILLING,EXECUTED FINISHING,FINISHED,FAILED,KILLED,DELETED}
Example v0.4.x: nordugrid-job-status: FINISHED at: 20020402161013Z
Example v0.6: nordugrid-job-status: FINISHED
Related xRSL: none
UI role: used in job status monitoring

Description: The status of the Grid job. The job state representation is undergoing a major change with the upcoming 0.6 release.

0.4.x release implementation: The attribute fully exposes the internal Grid Manager job states, for the explanation of the states *ACCEPTED*, *PREPARING*, *SUBMITTING*, *INLRMS*, *FINISHING*, *CANCELLING*, *FINISHED* consult the Grid manager manual[4]. The internal *INLRMS* state, meaning the job is under the control of the Local Resource Management System, is expanded by the information system to display the batch system status of the job as well: *INLRMS: R* or *INLRMS: Q* states are used to represent Grid jobs running or queuing in the local batch system.

** not in a stable release yet

Completed jobs are labeled by the *FINISHED* state regardless their success or failure. The *nordugrid-job-errors* attribute is used to distinguish between failed and successfully completed jobs. The *FINISHED* job state also carries information about the completion time of the job expressed in the GMT time format (see the example v0.4.x above)

0.6 release implementation: The Grid Manager internals are not fully exposed to the clients, new job states are introduced. The terminal job state is separated into three new states *FINISHED*, *KILLED*, *FAILED*. Furthermore, the *nordugrid-job-completiontime* attribute was introduced to separate the completion time from the *FINISHED* state.

ACCEPTING This is the initial job state The job has reached the cluster, a session directory was created, the UI has optionally started to upload some files to the sessiondir, the job waits to be detected by the Grid manager (GM).

internal state: ACCEPTED

ACCEPTED The job has been detected by the GM but can't go to the next state due to an internal GM-limit. The job may also be in the ACCEPTED state if the grid manager died during the stagein process (while the job was in the PREPARING state)

internal state: PENDING:ACCEPTED

internal state: PREPARING (in case of dead grid-manager)

PREPARING The input data is being gathered into the session directory (or to the cache), the GM downloads files specified in the user's xRSL. This is the Grid-stage-in process to the cluster. This is the latest state when the upload from the UI finishes. During this state, the UI can still upload files to the session directory.

internal state: PREPARING

PREPARED The stage-in process has successfully completed, the job is held waiting in the GM's internal queue due to an exceeded internal GM limit. The job may also be in the PREPARED state if the grid manager died during the LRMS job submission process (while the job was in the SUBMITTING state)

internal state: PENDING:PREPARING

internal state: SUBMIT (in case of a dead grid-manager)

SUBMITTING The GM prepares the LRMS job submission script and submits the job to the LRMS.

internal state: SUBMIT

INLRMS:X The job is in the local batch system, the job is controlled, managed by the LRMS. This state has several sub-states which are general mappings of native batch system states. Currently implemented sub-states:

INLRMS:R The job is running in the LRMS, executing on a node controlled by the batch system.

internal state: INLRMS

INLRMS:Q The job is queuing in the LRMS, waiting for a node, being put on hold, for some reason the job is in a 'pending state' of the LRMS.

internal state: INLRMS

INLRMS:S An already running job is in a suspended state.

internal state: INLRMS

INLRMS:E The job is finishing in the LRMS, it is 'exiting' from the batch system either because the job is completed or because it was cancelled.

internal state: INLRMS

INLRMS:O Any other native LRMS state which can not be mapped to the above general states must be labeled as 'O', meaning "other"

internal state: INLRMS

KILLING The job was requested to be killed and it is being killed by the GM, the GM interacts with the LRMS by running the job-cancel script

internal state: CANCELING

EXECUTED The job has completed in the batch system. There are two internal states corresponding to this state:

The job left the LRMS but the GM has not yet recognized this fact. The information system can't find the job in the batch system any longer but the GM still thinks the job is in the batch system (due to its latency).

internal state: INLRMS

The job has completed in the batch system AND the GM scanning process has recognized the job left the batch system BUT the job is held waiting in the GM internal queue due to an exceeded GM limit. The job may also be in the EXECUTED state if the grid manager died during the stageout process (while the job was in the FINISHING state)

internal state: PENDING:INLRMS

internal state: FINISHING (in case of dead grid-manager)

FINISHING The GM is doing the Grid stage-out process, specified output files are moved to their locations, GM is cleaning up the session directory removing everything which was not requested to be kept.

internal state: FINISHING

FINISHED The job has finished ALL its activity on the cluster AND no errors occurred during the Grid job's lifetime (no job.xx.failed file was created in the control-directory). The state *FINISHED* corresponds to the successful Grid job completion.

internal state: FINISHED

FAILED The job has finished ALL its activity on the cluster AND there occurred some problem during the lifetime of the Grid job. The *nordugrid-job-errors* and *nordugrid-job-exitcode* attributes contain more information about the job failure.

internal state: FINISHED

KILLED The job has finished ALL its activity on the cluster as a result of being killed by a client.

internal state: FINISHED

DELETED The job's session directory has removed from the cluster due to the expired session-directory-lifetime, only minimal set of info is kept about such a job.

internal state: DELETED

nordugrid-job-rerunable

Attribute value: {none, PREPARING, INLRMS, FINISHING}

Example: nordugrid-job-rerunable: PREPARING

Related xRSL: rerun

UI role: ngrerun utility makes use of it

Description: The attribute is only set for FAILED jobs and its value is either *none* or the name of the Grid job state from which the job can be rerun following a client request.

4.7 Grid users: nordugrid-authuser objectclass

Within the ARC information model every authorized Grid user of a resource is described by an *authuser* entry in the local tree. The user entries are used to present user-specific view of the resource, information such as free CPUs and available disk space are shown for every authorized Grid user. The existence of an *nordugrid-authuser* entry implies the granted access to the queue of the resource for the corresponding Grid identity.

nordugrid-authuser-name

Attribute value: string

Example: nordugrid-authuser-name: Lars Jenssen...8

Related xRSL: none

UI role: ?

Description: The Common Name of the authorized user appended by a local unique number. The Common Name is determined from the Certificate of the user. This *nordugrid-authuser-name* attribute constitutes to the DN of the user entry.

nordugrid-authuser-sn

Attribute value: LDAP Distinguished Name
Example: **nordugrid-authuser-sn:** /O=Grid/O=NorduGrid/OU=nordugrid.org/CN=Lars Jenssen
Related xRSL: none
UI role: used while searching for available resources of a user

Description: The LDAP Subject Name of the authorized Grid user as specified in his/her Grid credentials. A Grid user or a client can easily find the resources where he/she is authorized by issuing an LDAP search with a filter of *nordugrid-authuser-sn=his/her SN*.

nordugrid-authuser-freecpus

Attribute value: fixed format string
Example: **nordugrid-authuser-freecpus:** 2 4:25 5:180
Related xRSL: **count**
UI role: used in brokering

Description: The number of freely available CPUs with their time limits for the specific Grid identity is given by this attribute according to the following syntax:
ncpus[:min] [ncpus:min] ... where the pair *ncpus:min* gives the number of free CPUs with their time limit in minutes. The time limit information is optional. When there are blocked grid jobs in the Grid layer on a cluster this attribute is set to zero regardless of available free slots on the cluster.

nordugrid-authuser-diskspace

Attribute value: number showing the amount in MBs
Example: **nordugrid-authuser-diskspace:** 13964
Related xRSL: **disk**
UI role: used in the matchmaking

Description: The free disk space available for the session directory of the user's Grid job given in MBs.

nordugrid-authuser-queuelength

Attribute value: number
Example: **nordugrid-authuser-queuelength:** 0
Related xRSL: none
UI role: used in the brokering

Description: The number of queuing jobs of a particular user due to the Grid layer and batch system queue. The attribute takes into account the user's jobs accumulated both in the Grid layer and in the LRMS and shows the sum of the user's "queuing" jobs, gives the "length" of the user's personal queue.

4.8 Storage Resources: the nordugrid-se objectclass

The *nordugrid-se* objectclass is used to describe storage resources within the NorduGrid-ARC model. A storage resource consists of physical data source (the storage space itself) plus the protocols, policies, services

and interfaces which make the storage space available to the clients. The attributes of the objectclass are designed to describe all of these layers.

Some of the current attributes are not yet supported in the information system implementation. Furthermore the following attributes are being discussed to be added to the schema (see <http://bugzilla.nordugrid.org>, bug #181):

- **se-iospeed**: The average IO capability of the storage (MB/s)
- **se-networkspeed**: The network capability of the storage (MB/s)
- **se-architecture**: The hardware architecture of the Storage
Enumeration values: disk, raid of disks, memory, raid of memory, Tape, hierarchical storage, network storage
- **se-status**: The status of the storage service.
Enumeration values following glue status info: OK, Warning, Critical, Unknown, other. Most probable the OK and Critical is enough
- **se-load**: a numeric value representing the load on the storage element. load values from the UNIX top?
- **se-accessprotocol**: The protocol supported to access/transform files, e.g. the SRM flavour.
- **se-backupfrequency**: The backup service provided by the storage.
Enumeration values: never, occasionally, monthly, weekly, nightly.

nordugrid-se-name

Attribute value: fixed format string

Example: **nordugrid-se-name**: HEP:hathi.hep.lu.se

Related xRSL: none

Description: The globally unique name of the Storage Element composed as *local-name colon FQDN*. The *nordugrid-se-name* attribute constitutes to the DN of the se entry.

nordugrid-se-aliasname

Attribute value: free form text

Example: **nordugrid-se-aliasname**: Lund HEP SE

Related xRSL: none

Description: A free form text attribute displaying the alias name of the storage resource.

nordugrid-se-type

Attribute value: {gridftp, SSE, other}

Example: **nordugrid-se-type**: gridftp

Related xRSL: none

Description: The type of the storage element. ARC currently comes with two class of storage elements, the conventional GridFTP-based and the Smart Storage Element (SSE)[15]

nordugrid-se-freespace

Attribute value: number showing the amount in MBs

Example: **nordugrid-se-freespace: 253870**

Related xRSL: none

Description: The total amount of free space available on the SE in MBs. Not all of this space may be available for every Grid user.

nordugrid-se-totalspace

Attribute value: number showing the amount in MBs

Example: **nordugrid-se-totalspace: 1531381**

Related xRSL: none

Description: The total capacity of the storage resource displayed in MBs.

nordugrid-se-url

Attribute value: URL

Example: **nordugrid-se-url: gsiftp://hathi.hep.lu.se:2811/hep**

Related xRSL: none

Description: The URL to contact the Storage Element. Multivalued attribute, an SE can be accessed via several URLs.

nordugrid-se-location

Attribute value: Postal ZIP code with two letter country prefix

Example: **nordugrid-se-location: SE-22100**

Related xRSL: none

Description: The geographical location of the storage resource, preferably specified as a postal code with a two letter country prefix. Not yet supported.

nordugrid-se-owner

Attribute value: free form text

Example: **nordugrid-se-owner: Danish Center for Grid Computing**

Related xRSL: none

Description: The multivalued attribute is used to display the owner of the resource. Not yet supported.

nordugrid-se-acl

Attribute value: fixed form syntax

Example: **nordugrid-se-acl: VO:ATLAS**

Example: **nordugrid-se-acl: VO:developers.nordugrid.org**

Related xRSL: none

UI role: none

Description: The multivalued attribute is used to display authorization rules applied on an SE. The attribute value follows a *fixed form syntax*. Current syntax is rather coarse-grained and primitive: a "VO:" prefix followed by a VO name means that the given VO is authorized on the SE. Note that there are no standards for VO names.

nordugrid-se-issuerca

Attribute value: LDAP Distinguished Name
Example: `nordugrid-se-issuerca: /C=DE/O=GermanGrid/CN=GridKa-CA`
Related xRSL: none
UI role: ?

Description: The certificate issuer of the storage resource. The DN of the CA which issued the host certificate is shown by the attribute.

nordugrid-se-issuerca-hash

Attribute value: Hash value of a CA certificate
Example: `nordugrid-cluster-se-hash: 1f0e8352`
Related xRSL: none
UI role: ???

Description: The hash value of the certificate of the issuer CA of the SE, the hash of the certificate of the CA which issued the host certificate used by the SE is shown by the attribute.

nordugrid-se-trustedca

Attribute value: LDAP Distinguished Name
Example: `nordugrid-se-trustedca: /C=DE/O=GermanGrid/CN=GridKa-CA`
Example: `nordugrid-se-trustedca: /DC=ORG/DC=SEE-GRID/CN=SEE-GRID CA`
Related xRSL: none
UI role: ???

Description: The DNS trusted by the SE are shown by this multivalued attribute.

nordugrid-se-middleware

Attribute value: free form string representing a software package
Example: `nordugrid-se-middleware: nordugrid-200501280505cvs`
Example: `nordugrid-cluster-middleware: globus-2.4.3-15ng`
Related xRSL: none

Description: The multivalued attribute specifies the middleware packages installed on the storage resource.

nordugrid-se-comment

Attribute value: free form text
Example: `nordugrid-se-comment: Dedicated HEP storage`
Related xRSL: none

Description: The free form attribute displays additional information about the storage. Sometimes it contains an URL where more information can be read about the resource.

nordugrid-se-authuser

Attribute value: LDAP Distinguished Name
Example: /O=Grid/O=NorduGrid/OU=nordugrid.org/CN=Lars Jenssen
Example: /O=Grid/O=NorduGrid/OU=nordugrid.org/CN=Leif Jenssen
Related xRSL: none
UI role: ?

Description: The multivalued attribute lists the DN's of the authorized users.

nordugrid-se-accesscontrol ***

Attribute value: {trivial, gac1, other}
Example: nordugrid-se-accesscontrol: gac1
Related xRSL: none
UI role: ?

Description^{††}: The access control framework provided by the storage element and can be utilized by the users. Currently ARC implements two type of access control frameworks, the *GACL* [16] and a trivial one [4].

4.9 Other services: nordugrid-rc objectclass for data catalogues

Originally, the *nordugrid-rc* objectclass was introduced to describe the Replica Catalogue services. The objectclass was rarely used and by now it has become obsolete. The *nordugrid-rc* objectclass is to be deprecated and replaced by a more general concept of Grid service.

For the sake of completeness the attributes of the *nordugrid-rc* are listed without proper description, please refer to the Storage Resource section or to the Appendix in case more information is needed.

The list of the *nordugrid-rc* attributes: nordugrid-rc-name, nordugrid-rc-aliasname, nordugrid-rc-baseurl, nordugrid-rc-authuser, nordugrid-rc-location, nordugrid-rc-owner, nordugrid-rc-issuerca.

4.10 Mapping to the GLUE Information model

The GLUE Information model, or the "GLUE Schema" represents an alternative information model [13] developed parallel to the ARC schema. The version 1.2 of the GLUE Schema is widely used in large Grid deployments such as LCG/EGEE or the Open Science Grid. Below a mapping between the ARC and the GLUE information model is presented, in particular the most commonly used GLUE 1.2 entities are populated using information stored in ARC attributes and objectclasses. Unfortunately, the GLUE specification [13] is incomplete, the relation between the GLUE entities, the DIT is not fully specified within the model. Therefore in our mapping we had to rely on a particular implementation, the mapping presented below follows the GLUE implementation of the LCG/EGEE project.

The GLUE model aims to describe Grid entities such as *Site*, *Service*, *Cluster*, *SubCluster*, *Host*, *ComputingElement*, *Job*, *StorageElement*, however not all of the above entities are used in production deployments. The ARC - GLUE-1.2 mapping presented below deals with the most relevant *ComputingElement*, *Cluster/SubCluster* and *Site* entities.

^{††}not in a stable release yet

4.10.1 The ComputingElement entity of the LCG Glue model

The LCG implementation of the Glue-1.2 ComputingElement entity is populated with information stored in ARC schema: attributes from the *nordugrid-cluster* and *nordugrid-queue* entries (objectclass) are used to fill an LCG-Glue ComputingElement. The ARC schema is much richer than the Glue one, almost every Glue CE attribute can be filled from ARC info while there are lots of ARC information which has no place within the Glue-1.2 representation.

GlueCEUniqueID

ARC mapping: `nordugrid-cluster-contactstring/nordugrid-queue-name`

Example: `GlueCEUniqueID: gsiftp://gridbox.nordugrid.org:2811/gridqueue`

Comment: ARC computing elements are defined by cluster contact information and queue information. Notice the slash / introduced to extend the *nordugrid-cluster-contactstring* with the *nordugrid-queue-name*.

GlueCENAME

ARC mapping: `nordugrid-queue-name`

Example: `GlueCENAME: gridqueue`

Comment: -

GlueCEInformationServiceURL

ARC mapping: an infrastructure specific LDAP URL

Example: `GlueCEInformationServiceURL: ldap://arc-bdii.cern.ch:2170/o=grid`

Comment: The attribute should contain the LDAP URL of the BDII service providing the ARC-Glue mapping.

GlueCEInfoLRMSType

ARC mapping: `nordugrid-cluster-lrms-type`

Example: `GlueCE: SGE`

Comment: -

GlueCEInfoLRMSVersion

ARC mapping: `nordugrid-cluster-lrms-version`

Example: `GlueCEInfoLRMSVersion: 5.1`

Comment: -

GlueCEInfoGRAMVersion

ARC mapping: this attribute is irrelevant for ARC

Comment: ARC does not use GRAM for job submission. GlueCEInfoGRAMVersion Glue attribute is too Globus specific.

GlueCEInfoHostName

ARC mapping: `nordugrid-cluster-name`

Example: `GlueCEInfoHostName: gridbox.nordugrid.org`

Comment: -

GlueCEInfoGatekeeperPort

ARC mapping: this attribute is irrelevant for ARC

Comment: The Glue attribute is too much Globus specific, again. ARC runs no Globus gatekeeper.

GlueCEInfoJobManager

ARC mapping: this attribute is irrelevant for ARC

Comment: Too much Globus specific, again. No Globus jobmanager on ARC.

GlueCEInfoContactString

ARC mapping: `nordugrid-cluster-contactstring?queue=nordugrid-queue-name`

Example: `GlueCEInfoContactString: gsiftp://gridbox.nordugrid.org:2811/jobs?queue=gridqueue`

Comment: The attribute is constructed from the *nordugrid-cluster-contactstring* and the *nordugrid-queue-name*. In order to access a CE in ARC you need to specify the queue name together with the `gsiftp://` cluster contact string. `GlueCEInfoContactString` attribute obsoletes the `GlueCEInfo.HostName`, `GlueCEInfo.GatekeeperPort` and `GlueCEInfo.Jobmanager` attributes. The attribute name is rather misleading since it is not the contact string for the Information service but the contact info for the job submission service of the CE.

GlueCEInfoTotalCPUs

ARC mapping: if defined use the *nordugrid-queue-totalcpus* otherwise the *nordugrid-cluster-totalcpus*

Example: `GlueCEInfoTotalCPUs: 66`

Comment: This Glue attribute is DEPRECATED in Glue version 1.2, better not to set.

GlueCEInfoApplicationDir

ARC mapping: *ApplicationDir* does not exist on ARC computing resource

Comment: none-existing attributes shouldn't be set.

GlueCEInfoDataDir

ARC mapping: *DataDir* does not exist on ARC computing resource

Comment: none-existing attributes shouldn't be set.

GlueCEDefaultSE

ARC mapping: This concept does not exist on ARC computing resource
Comment: none-existing attributes shouldn't be set.

GlueCEStateStatus

ARC mapping: perform mapping on the *nordugrid-queue-status* values
Example: `GlueCEstateStatus:production`
Comment: Use the following mapping of the *nordugrid-queue-status* values: active (ARC) = Production (Glue); inactive (ARC) = Closed (Glue).

GlueCEStateRunningJobs

ARC mapping: *nordugrid-queue-running*
Example: `GlueCEStateRunningJobs: 12`
Comment: -

GlueCEStateWaitingJobs

ARC mapping: sum of the *nordugrid-queue-gridqueued*, *nordugrid-queue-localqueued*, *nordugrid-queue-prelrmsqueued*
Example: `GlueCEStateWaitingJobs: 12`
Comment: -

GlueCEStateTotalJobs

ARC mapping: sum of the *nordugrid-queue-running*, *nordugrid-queue-gridqueued*, *nordugrid-queue-localqueued*, *nordugrid-queue-prelrmsqueued*
Example: `GlueCE: 54`
Comment: -

GlueCEStateEstimatedResponseTime

ARC mapping: ARC does not use this attribute and recommends to set it to a fix value
Example: `GlueCEStateEstimatedResponseTime: 1000`
Comment: Without advanced reservation it is almost impossible to come up with a reasonable value for queue "EStateEstimatedResponseTime". NorduGrid does not believe in the usefulness of this attribute. Best is to set it for the same value for all CEs and let the broker decide using other attributes.

GlueCEStateWorstResponseTime

ARC mapping: ARC does not use this attribute and recommends to set it to a fix value

Example: `GlueCEStateWorstResponseTime: 2000`

Comment: NorduGrid does not believe in the usefulness of this attribute. Best is to set it for the same value for all CEs and let the broker decide using other attributes.

GlueCEStateFreeJobSlots

ARC mapping: calculate it from the *nordugrid-queue-totalcpus* / *nordugrid-cluster-totalcpus* and the *nordugrid-queue-running*

Example: `GlueCEStateFreeJobSlots: 3`

Comment: Subtract the *nordugrid-queue-running* from the *nordugrid-queue-totalcpus* or from the *nordugrid-cluster-totalcpus* in case the *nordugrid-queue-totalcpus* attribute is not set.

GlueCEStateFreeCPUs

ARC mapping: - same as the `GlueCEState.FreeJobSlots`

Comment: This attribute is DEPRECATED in glue-1.2, better not to set.

GlueCEPolicyMaxWallTime

ARC mapping: use the value of *nordugrid-queue-maxcputime*

Example: `GlueCEPolicyMaxWallTime: 120`

Comment: *queue-maxwalltime* attribute currently does not exist in ARC schema, the value of the *queue-maxcputime* can be used. *nordugrid-queue-maxwalltime* will be added to the ARC schema.

GlueCEPolicyMaxCPUTime

ARC mapping: *nordugrid-queue-maxcputime*

Example: `GlueCEPolicyMaxCPUTime: 160`

Comment: -

GlueCEPolicyMaxTotalJobs

ARC mapping: *nordugrid-queue-maxqueueable*

Example: `GlueCEPolicyMaxTotalJobs: 120`

Comment: -

GlueCEPolicyMaxRunningJobs

ARC mapping: `nordugrid-queue-maxrunning`
Example: `GlueCEPolicyMaxRunningJobs: 16`
Comment: -

GlueCEPolicyPriority

ARC mapping: does not exist in ARC, set it to fix value
Example: `GlueCE: 1`
Comment: Set it to 1 for every CE, it is a rather useless internal LRMS value.

GlueCEPolicyAssignedJobSlots

ARC mapping: use one of the *nordugrid-queue-totalcpus*, *nordugrid-queue-maxrunning*, *nordugrid-cluster-totalcpus* attributes
Example: `GlueCEPolicyAssignedJobSlots: 30`
Comment: One of the above listed three nordugrid attributes in the given priority order can be used to set this strange Glue attribute: there is large overlap with the `GlueCEPolicyMaxRunningjobs` attribute, one of them should be enough.

GlueCEAccessControlBaseRule

ARC mapping: `nordugrid-cluster-acl`
Example: `GlueCEAccessControlBaseRule: V0:atlas`
Comment: The *nordugrid-cluster-acl* attribute is introduced in the 0.6 ARC release. In case *nordugrid-cluster-acl* is not defined, *GlueCEAccessControlBaseRule* will contain the default value `V0:ops`.

4.10.2 The Cluster and SubCluster entity of the LCG Glue model

Originally, the Glue model offers the *Cluster*, *Subcluster* and *Host* entities for describing physical (hardware) properties of computing services. Heterogeneity was supposed to be addressed by the *SubCluster* and *Host* entities: clusters can be heterogeneous while subclusters are assumed to be homogeneous. The *Cluster* entity was designed to link the hardware description (*SubCluster*, *Host*) to the computing service view of the resource: the *ComputingElement* entity.

Unfortunately, in the actual LCG/EGEE deployment the Glue *Cluster*, *SubCluster*, *Host* entities got mixed up. The *Cluster* entry is rather empty, it carries almost zero information. Clusters have only one *SubCluster* subentry (no heterogeneity is treated), *SubCluster* entries got merged with the *Host* entities, thus *SubClusters* carry all the *Host* properties. Furthermore, due to the flat LDAP tree the *ComputingElement* entries are not linked to the *Cluster/SubCluster* entries.

Since the Glue model comes with no DIT, it does not specify the relation of particular entities. In the LCG/EGEE deployment it lead to an almost completely flat structure where *ComputingElement* and *Cluster* entities (and even *Site* entities !) are placed on the same level of the LDAP tree.

The ARC infosys representation has cluster and queue entries. More importantly, it contains the notion of cluster- and queue-level attributes and a cluster-homogeneity flag. Populating the Glue-LCG *Cluster/SubCluster* entities using the ARC schema is rather strateforward since ARC offers cluster and queue-level view of computing resources to treat heterogeneity (see Section 4.5). Queue-level view is used for heterogeneous subset of computing resurces. For homogeneous clusters (where the *nordugrid-cluster-homogeneity=true* is set) create only one Glue *SubCluster* and use the cluster-level attributes from the ARC information model. For inhogeneous clusters (*nordugrid-cluster-homogeneity=false*) create a Glue *SubCluster* for every queue and populate the Glue *SubCluster* with the corresponding ARC queue-level attributes.

Final remark: as of writing, on the whole LCG/EGEE Grid there is not a single Cluster with more than one SubCluster.

GlueClusterUniqueID

ARC mapping: `nordugrid-cluster-name`

Example: `GlueClusterUniqueID: gridbox.nordugrid.org`

Comment: -

GlueClusterName

ARC mapping: `nordugrid-cluster-aliasname`

Example: `GlueClusterName: Gridbox cluster`

Comment: -

GlueClusterTmpDir

ARC mapping: does not exist in ARC

Comment: The location of a temporary directory on the frontend is an internal cluster property, shouldn't be directly advertised on the Grid. Don't set this attribute.

GlueClusterWNTmpDir

ARC mapping: does not exist in ARC

Comment: The location of a temporary directory on the worker node is an internal cluster property, shouldn't be directly advertised on the Grid. Don't set this attribute.

GlueSubClusterUniqueID

ARC mapping: `nordugrid-cluster-name/nordugrid-queue-name`

Example: `GlueSubClusterUniqueID: gridbox.nordugrid.org/gridqueue`

Comment: notice the "Glue SubCluster" = "ARC queue" assumption.

GlueSubClusterName

ARC mapping: `nordugrid-cluster-aliasname/nordugrid-queue-name`

Example: `GlueSubClusterName: Gridbox cluster/gridqueue`

Comment: notice the "Glue SubCluster" = "ARC queue" assumption.

GlueSubClusterPhysicalCPUs

ARC mapping: `nordugrid-queue-totalcpus` or `nordugrid-cluster-totalcpus`

Comment: The queue or cluster level attribute should be used depending on the cluster-homogeneity flag.

GlueSubClusterLogicalCPUs

ARC mapping: not available in ARC

Comment: No need to set this attribute.

GlueSubClusterLocationName/Version/Path

ARC mapping: does not exist in ARC

Example: `GlueSubCluster:`

Comment: The name/version/path of an installed software is an internal cluster property, shouldn't be advertised on the grid. The 'Location' Glue entity which defines the above three attributes is a bad idea. Don't set these attributes.

GlueHostOperatingSystemName/OperatingSystemRelease/OperatingSystemVersion

ARC mapping: `nordugrid-queue-opsys` or `nordugrid-cluster-opsys`

Example: `GlueHostOperatingSystemName: Redhat`

Example: `GlueHostOperatingVersion: 7.3`

Comment: The ARC schema has only two multivalued (!) attributes for describing the operatingsystem. Use the *nordugrid-queue-opsys* or the *nordugrid-cluster-opsys* multivalued (!) attribute, depending on the cluster-homogeneity flag, to set some of the above Glue attributes.

GlueHostProcessorModel/ProcessorVersion/ProcessorVendor/ProcessorClockSpeed/ ProcessorInstructionSet/ProcessorOtherDescription

ARC mapping: `nordugrid-queue-nodecpu` or `nordugrid-cluster-nodecpu`

Comment: ARC has only two attributes for describing the node CPU, the *nordugrid-queue-nodecpu* or the *nordugrid-cluster-nodecpu*. Use one of them depending on the cluster-homogeneity flag.

GlueHostRAMSize

ARC mapping: `nordugrid-queue-nodememory` or `nordugrid-cluster-nodememory`

Example: `GlueHostRAMSize: 1024`

Comment: The queue or cluster level attribute should be used depending on the cluster-homogeneity flag. Notice the different semantics of the LCG-Glue and ARC memory attributes: ARC memory is 'amount of memory guaranteed for the application' while LCG-Glue memory is 'amount of RAM'.

GlueHostVirtualSize

ARC mapping: not available in ARC

Comment: Not available in the ARC model, don't set this attribute.

GlueHostNetworkAdapterOutboundIP

ARC mapping: nordugrid-cluster-nodeaccess

Example: GlueHostNetworkAdapterOutboundIP: FALSE

Comment: Set *TRUE* if *nordugrid-cluster-nodeaccess=outbound*, otherwise *FALSE*. There is no queue-level nodeaccess property in ARC, it is defined on the cluster-level.

GlueHostNetworkAdapterInboundIP

ARC mapping: nordugrid-cluster-nodeaccess

Example: GlueHostNetworkAdapterInboundIP: TRUE

Comment: Set *TRUE* if *nordugrid-cluster-nodeaccess=inbound*, otherwise *FALSE*. There is no queue-level nodeaccess property in ARC, it is defined on the cluster-level.

GlueHostArchitecturePlatformType

ARC mapping: nordugrid-queue-architecture or nordugrid-cluster-architecture

Example: GlueHostArchitecturePlatformType: x86_64

Comment: The queue or cluster level attribute should be used depending on the cluster-homogeneity flag.

GlueHostBenchmarkSI00/BenchmarkSF00

ARC mapping: nordugrid-queue-benchmark or nordugrid-cluster-benchmark

Example: GlueHostBenchmarkSI00: 100

Comment: Take the benchmark value from the *nordugrid-queue-benchmark: SPECINT2000 @ 111* or use the nordugrid-cluster-benchmark for homogeneous clusters. The way Glue addresses (hardcodes) benchmark names are really bad.

GlueHostApplicationSoftwareRunTimeEnvironment

ARC mapping: nordugrid-cluster-runtimeenvironment

Example: GlueHost: APPS/CHEM/DALTON-1.2.1-1.0

Comment: Runtimeenvironments are defined on cluster-level in ARC.

4.10.3 The Site entity of the LCG Glue model

The LCG-Glue Site entity is meant to be an administrative and structuring entry in the information system representation. Its main purpose is to collect resources belonging to the same "site" and their administrative description in a single entry. Unfortunately, in the LCG Glue deployment the Site entry is NOT used in the LDAP tree to group the resources/entries belonging to the same site. The introduction of the Site concept in LCG-Glue has a consequence that Glue-LCG resource entries (Cluster, ComputingElement, StorageElement) have no administrative attributes.

The ARC schema follows a different approach. The ARC infosys representation has no Site object, resources (clusters, SEs) are not grouped on a site level, administrative information (owner, location, comments, etc) are stored within the resource entries themselves.

In order to create an LCG-Glue Site entry from the ARC information tree the best is to follow the "one resource = one site" approach and take the administrative info present e.g. in the cluster ARC entry and use that to create a Glue Site entry. In the proposed mapping below every ARC cluster corresponds to a Glue-LCG site. Optionally, resources having the same *owner* attribute can be grouped under a common Site entry. Unfortunately there is no easy and reliable method to find out which resources would belong to the same administrative unit.

GlueSiteUniqueID

ARC mapping: `nordugrid-cluster-name`

Example: `GlueSiteUniqueID: gridbox.nordugrid.org`

Comment: The proposed mapping expresses the "one cluster = one site" mapping which provides real uniqueness. btw, the currently used LCG Site names such as "Budapest" don't guarantee uniqueness at all.

GlueSiteName

ARC mapping: `nordugrid-cluster-aliasname`

Example: `GlueSiteName: Gridbox cluster`

Comment: Use the cluster alias or let site admins manually fill the site names.

GlueSiteDescription

ARC mapping: `nordugrid-cluster-comment`

Example: `GlueSiteDescription: General purpose Grid box`

Comment: -

GlueSiteUserSupportContact/SysAdminContact/SecurityContact

ARC mapping: `mailto: nordugrid-cluster-support`

Example: `GlueSiteSysAdminContact: mailto: contact@gridbox.nordugrid.org`

Comment: ARC schema has only one "contact" email address, that can be used to set all the three Glue contact attributes. Notice the *mailto:* Glue prefix in front of the *nordugrid-cluster-support*.

GlueSiteLocation

ARC mapping: `nordugrid-cluster-location`

Comment: Observe the different semantics! ARC location is a postal ZIP code while LCG-location is 'city,state,country'

GlueSiteLatitude/Longitude

ARC mapping: these don't exist in ARC

Comment: This information is not available in ARC schema.

GlueSiteWeb

ARC mapping: does not exist in ARC

Comment: This info (website) sometimes is available in the *nordugrid-cluster-comment* attribute.

GlueSiteSponsor

ARC mapping: nordugrid-cluster-owner

Example: `GlueSiteSponsor: SweGrid`

Comment: Note that the *nordugrid-cluster-owner* attribute is multivalued.

GlueSiteOtherInfo

ARC mapping: add the fixed text shown in the example

Example: `GlueSiteOtherInfo: This site is running on the ARC middleware`

Comment: Let us have this small marketing. Btw, many LCG deployments use the *GlueSiteOtherInfo* to advertise their webpage or a hostname, though there are attributes for those.

5 Registration Processes, EGIIS: Topology

The individual ARIS instances need to be connected and organized into some sort of topological structure in order to create a coherent Grid system. NorduGrid-ARC utilizes registration processes and EGIIS servers to build a distributed information system out of the individual local information trees.

Connecting information sources together is usually referred to as the resource information aggregation. ARC implements a minimalistic aggregation process: an EGIIS collects only the contact information of the information resources and no information is gathered or cached from ARISes.

The ARIS and EGIIS services are linked together via the *registration processes*. During a registration process the registrant (lower level) sends its registration packet to an EGIIS. A registration packet contains information about the host (the registrant) initiating the registration process and about the information service running on the registrant (either ARIS or EGIIS): the registration message is basically the LDAP contact URL of an information service running on the registrant (additionally some timing parameter is also transferred indicating how long the registration information is expected to be kept in EGIIS). The target EGIIS can filter out registration processes coming from registrants, unfortunately the filtering capability of the current EGIIS is rather limited, it is based only on the FQDN and NOT on the LDAP URL. Registrations are sent periodically to the target EGIISes, thus the registration mechanism follows a periodic push model. Technically, the registrations are implemented as periodic *ldapadd* operations.

EGIISSes are used to maintain dynamic lists of available resources, containing the LDAP contact information of the registrants. A registrant can either be an ARIS or another EGIIS. The content of an EGIIS, that is the information on the registrants, is periodically purged, this way maintaining a dynamic registrant list. EGIIS is implemented as a special purpose LDAP database: the registration entries are stored as LDAP entries in the EGIIS LDAP back-end. The periodic purging of the registrant entries are done by NOT removing the obsoleted entry BUT setting the *Mds-Reg-status: PURGED* attribute. ARC EGIISes are purely used to

```
ldapsearch -x
  -h host
  -p port
  -b basedn
  -s base
```

Figure 4: The LDAP query to obtain the content of an EGIIS.

```
#grid.tsl.uu.se, Sweden, grid
dn: nordugrid-cluster-name=grid.tsl.uu.se, Mds-Vo-name=Sweden,o=grid
objectClass: Mds
objectClass: MdsVoOp
objectClass: MdsService
objectClass: MdsServiceLdap
Mds-Service-type: ldap
Mds-Service-hn: grid.tsl.uu.se
Mds-Service-port: 2135
Mds-Service-Ldap-suffix: nordugrid-cluster-name=grid.tsl.uu.se,
                        Mds-Vo-name=local, o=grid
Mds-Service-Ldap-sizelimit: 0
Mds-Service-Ldap-timeout: 45
Mds-Service-Ldap-cachettl: 15
Mds-Bind-Method-servers: ANONYM-ONLY
Mds-Reg-status: VALID
```

Figure 5: A LDAP registration entry obtained from the Sweden EGIIS describing the valid registration of a computing resource

maintain a set of LDAP registration entries shown in Figures 5,6. No any other information is stored in the EGIISes.

The LDAP registration entries stored in an EGIIS running on *host*, *port* under the LDAP *basedn* can be obtained by an LDAP query executed against the EGIIS:

That is an anonymous (-x) ldapsearch against the LDAP server of the Index Service with search scope *base*. Examples for other EGIIS queries can be found in the the Appendix C.

A valid LDAP registration entry, stored in an EGIIS and obtained by the above described LDAP query is shown in Figure 5. The LDAP contact URL of the ARIS running on the *grid.tsl.uu.se* machine is given by the attributes *Mds-Service-hn*, *Mds-Service-port*, *Mds-Service-Ldap-suffix*. The *dn: nordugrid-cluster-name=grid.tsl.uu.se, Mds-Vo-name=Sweden,o=grid* of the registration entry indicates that the registration corresponds to computing resource registering to an EGIIS called *Mds-Vo-name=Sweden,o=grid*. The *Mds-Reg-status* attribute shows the validity of the registration. Figure 6 shows another LDAP registration entry corresponding to a PURGED state registration of an (lower level) EGIIS to another (higher level) EGIIS.

The ARISes and the EGIISes of the NorduGrid-ARC Grid are organized into a multi-level tree hierarchy. The ARISes describing the actual Storage or Computing resources represent the lowest level of the tree-like topology. Resources are registering to first level EGIISes which are registering to Second level services, and so forth. The registration chain ends at the Top Level Indices which represent the root of the tree hierarchy. The structure is built from bottom to top: always the lower level registers to the higher one. The tree-like hierarchical structure is motivated by the natural geographical organization where resources belonging to the same region register under a region index, region indices are registering to the appropriate country index while country indices are grouped together and register to the top level Grid index services. In order to avoid any single point of failure, NorduGrid-ARC operates a multi-rooted tree with several top-level Indices (Table 3 lists the LDAP contact URL of the Top Level EGIISes). Figure 7 shows simplified schematic view of the multi-rooted tree topology of ARC-connected resources and EGIISes. Besides the geographical structuring there are some EGIISes which group resources by specific application area or organization. These application/organization EGIISes either link themselves to a country EGIIS or register directly to a Top Level EGIIS.

```

# SweGrid, Sweden, grid
dn: Mds-Vo-name=SweGrid, Mds-Vo-name=Sweden,o=grid
objectClass: Mds
objectClass: MdsVoOp
objectClass: MdsService
objectClass: MdsServiceLdap
Mds-Service-type: ldap
Mds-Service-hn: hagrid.it.uu.se
Mds-Service-port: 2135
Mds-Service-Ldap-suffix: Mds-Vo-name=SweGrid, o=grid
Mds-Service-Ldap-sizelimit: 0
Mds-Service-Ldap-timeout: 120
Mds-Service-Ldap-cachettl: 30
Mds-Bind-Method-servers: ANONYM-ONLY
Mds-Reg-status: PURGED

```

Figure 6: A LDAP registration entry obtained from the Sweden EGIIS describing the PURGED registration of the SWEGRID EGIIS: The SWEGRID EGIIS is running on the *hagrid.it.uu.se* machine on port *2135* and with LDAP base *Mds-Vo-name=SweGrid, o=grid*.

| host | port | LDAP baseDN |
|----------------------|------|------------------------------|
| index1.nordugrid.org | 2135 | mds-vo-name=nordugrid,o=grid |
| index2.nordugrid.org | 2135 | mds-vo-name=nordugrid,o=grid |
| index3.nordugrid.org | 2135 | mds-vo-name=nordugrid,o=grid |
| index4.nordugrid.org | 2135 | mds-vo-name=nordugrid,o=grid |

Table 3: LDAP URL of the TOP Level ARC Index Services

5.1 Resource discovery

Resource discovery is the process when clients walk through the EGIISes organized in a tree topology (see Figure 7) and collect LDAP contact URLs of the Computing and Storage resources. The discovery process usually starts at the top of the tree by querying some of the Top Level EGIISes (Table 3 lists the Top EGIISes). A Top Level EGIIS is queried by a LDAP search (Figure 4) which returns all the registration entries stored in the EGIIS. EGIISes contain registration information of both resources and other EGIISes. In the second step the client has to separate the entries corresponding to EGIISes from entries describing registrations of actual resources (ARISes). Figures 5 and 6 show a registration table of a Computing resource and an EGIIS, respectively. The *Ldap-suffix* of a resource always contains the special string *"mds-vo-name=local,o=grid"* referring to the fact that the resource runs an ARIS. Then, the client contact the newly discovered EGIISes and obtains the registration tables. The tables are separated into Indices and Resources, again. The process is repeated until all the EGIISes are queried and the full list of LDAP Contact URL of ARISes are collected.

Once the client has collected a list of ARIS LDAP contact URLs from the EGIISes, the second phase of the information collection begins: the client directly contacts every Resource and initiates an LDAP query against the ARIS. This is the real information gathering process in contrast to the first phase in which only the LDAP URLs were collected. Remember, unlike other systems (Globus GIIS, GT4 aggregator, R-GMA, LCG-BDII) ARC has no service which caches or aggregates Resource specific information on a higher level, ARC EGIISes are not used to store local information, Indices maintain only LDAP URLs.

Acknowledgements

This work was supported in parts by: the Nordunet 2 programme, the Nordic DataGrid Facility, the EU KnowARC project (Contract nr. 032691) and the EU EMI project (Grant agreement nr. 261611).

A Clients of the ARC information system

The entire content of the Information System including both the ARIS and EGIIS are presented via an LDAP interface. LDAP is a very well supported protocol, therefore it is very easy to construct clients making use of the ARC information system.

The Grid Monitor [12] is a simple yet powerful graphical interface to the Information System written in PHP. The monitor can be set up at any Web server and it provides a browsable hierarchial representation of the information content of the Grid. The monitor is localised in many human languages; still, it does not pre-process or modify information, and displays contents of ARIS as is.

The standard ARC command line interface [14] or Graphical UIs, as well as ARClib being developed provide another examples of straightforward interface-ability to the LDAP-based Information System.

B Glue 1.2 with ARC, gLite interoperability

The main reason for publishing ARC information in the Glue schema is for interoperability with gLite. In order to do this, knowledge about both the ARC information system and the gLite information system is required.

In ARC, you have a local (on ARC CE) information system called ARIS. In gLite, this is translated to a Resource-BDII.

In ARC, the ARIS registers to one or more EGIIS, while in gLite, the Resource-BDII registers to a Site-BDII. ARC can be configured to act as a Site-BDII as well by configuring it in `arc.conf`.

In ARC, the topmost level is an EGIIS. In gLite, it is a Top-BDII. You can not set up a Top-BDII with ARC software.

What you are likely to want to do is to set up ARC to act as a Resource-BDII and as a Site-BDII. And then register your site-BDII to a gLite top-bdii.

C LDAP Examples

- LDAP query against a local tree with a filter for nordugrid-cluster objectclass:

```
ldapsearch -h bambi.hep.lu.se -p 2135 -x -b 'mds-vo-name=local,o=grid' 'objectclass=nordugrid-cluster'
version: 2

#
# filter: objectclass=nordugrid-cluster
# requesting: ALL
#

# bambi.hep.lu.se, local, grid
dn: nordugrid-cluster-name=bambi.hep.lu.se,Mds-Vo-name=local,o=grid
objectClass: Mds
objectClass: nordugrid-cluster
nordugrid-cluster-name: bambi.hep.lu.se
nordugrid-cluster-aliasname: Bambi Test Cluster
nordugrid-cluster-owner: EHEP
nordugrid-cluster-location: SE-221 00
nordugrid-cluster-issuerca: /O=Grid/O=NorduGrid/CN=NorduGrid Certification Authority
nordugrid-cluster-contactstring: gsiftp://bambi.hep.lu.se:2811/jobs
nordugrid-cluster-support: grid.support@mysite.org
nordugrid-cluster-support: grid.support@myproject.org
nordugrid-cluster-lrms-type: fork
nordugrid-cluster-lrms-config: single job per processor
nordugrid-cluster-architecture: i686
nordugrid-cluster-opsys: Mandrake-10.0
nordugrid-cluster-opsys: glibc-2.3.3-10mdk
nordugrid-cluster-benchmark: SPECINT2000 @ 222
nordugrid-cluster-benchmark: SPECFP2000 @ 333
nordugrid-cluster-homogeneity: True
nordugrid-cluster-nodecpu: Intel(R) Pentium(R) 4 CPU 3.00GHz @ 2993.100 Mhz
nordugrid-cluster-nodeaccess: inbound
nordugrid-cluster-nodeaccess: outbound
nordugrid-cluster-totalcpus: 2
nordugrid-cluster-usedcpus: 0
nordugrid-cluster-cpudistribution: 2cpu:1
nordugrid-cluster-queuedjobs: 0
nordugrid-cluster-totaljobs: 0
nordugrid-cluster-sessiondir-free: 13958
nordugrid-cluster-sessiondir-total: 26403
nordugrid-cluster-cache-free: 9536
nordugrid-cluster-cache-total: 9536
nordugrid-cluster-middleware: nordugrid-0.5.20
nordugrid-cluster-middleware: globus-2.4.3-12ng
nordugrid-cluster-middleware: my grid software
nordugrid-cluster-runtimeenvironment: tt
Mds-validfrom: 20050307103026Z
Mds-validto: 20050307103029Z

# search result
search: 2
result: 0 Success

# numResponses: 2
# numEntries: 1
```

- Query for active Grid jobs stored in the local tree describing a computing resource:

```
ldapsearch -h quark.hep.lu.se -p 2135 -x -b 'mds-vo-name=local,o=grid' 'objectclass=nordugrid-job'
version: 2

#
# filter: objectclass=nordugrid-job
# requesting: ALL
#

# gsiftp://quark.hep.lu.se:2811/jobs/131601109950874935622127, jobs, pc, quark.hep.lu.se, local, grid
dn: nordugrid-job-globalid=gsiftp://quark.hep.lu.se:2811/jobs/131601109950874935622127,nordugrid-info-group-name=jobs,nordugrid-queue-name=pc,nordugrid-cluster-name=quark.hep.lu.se,Mds-Vo-name=local,o=grid
objectClass: Mds
objectClass: nordugrid-job
nordugrid-job-globalid: gsiftp://quark.hep.lu.se:2811/jobs/131601109950874935622127
nordugrid-job-globalowner: /O=Grid/O=NorduGrid/OU=cmn.ki.se/CN=Roxana Merino
nordugrid-job-jobname: TFB2
nordugrid-job-submissiontime: 20050304154114Z
nordugrid-job-execcluster: quark.hep.lu.se
nordugrid-job-execqueue: pc
nordugrid-job-cpucount: 1
nordugrid-job-sessiondirresettime: 20050305154232Z
nordugrid-job-stdout: tfbs.out
nordugrid-job-stderr: tfbs.out
nordugrid-job-gmlog: gmlog
nordugrid-job-runtimeenvironment: BIO-GEIJER-0.0.2
nordugrid-job-submissionui: 217.208.119.237:18236;10.0.0.1
nordugrid-job-clientsoftware: nordugrid-0.4.4
nordugrid-job-proxyexpirationtime: 20050304212024Z
nordugrid-job-status: DELETED
nordugrid-job-reqcputime: 2
nordugrid-job-executionnodes: node2
Mds-validfrom: 20050307110037Z
Mds-validto: 20050307110107Z

# gsiftp://quark.hep.lu.se:2811/jobs/5381111006551269125810, jobs, pclong, quark.hep.lu.se, local, grid
dn: nordugrid-job-globalid=gsiftp://quark.hep.lu.se:2811/jobs/5381111006551269125810,nordugrid-info-group-name=jobs,nordugrid-queue-name=pclong,nordugrid-cluster-name=quark.hep.lu.se,Mds-Vo-name=local,o=grid
objectClass: Mds
objectClass: nordugrid-job
nordugrid-job-globalid: gsiftp://quark.hep.lu.se:2811/jobs/5381111006551269125810
nordugrid-job-globalowner: /O=Grid/O=NorduGrid/OU=fys.ku.dk/CN=Brian Moller Andersen
```

```

nordugrid-job-jobname: selfconsistentMDSSAFNAF100slabonemu
nordugrid-job-submissiontime: 20050305233152Z
nordugrid-job-execcluster: quark.hep.lu.se
nordugrid-job-execqueue: pclong
nordugrid-job-cpucount: 1
nordugrid-job-stdout: output
nordugrid-job-stderr: output
nordugrid-job-submissionui: 130.225.102.149:33853;johansen.fys.ku.dk
nordugrid-job-clientsoftware: nordugrid-0.3.39
nordugrid-job-proxyexpirationtime: 20050306100826Z
nordugrid-job-status: INLRMS: R
nordugrid-job-usedmem: 30736
nordugrid-job-usedwalltime: 2127
nordugrid-job-usedcputime: 2125
nordugrid-job-reqcputime: 10080
nordugrid-job-executionnodes: node1/0
nordugrid-job-lrmscomment: Job started on Sun Mar 06 at 00:32
Mds-validfrom: 20050307110038Z
Mds-validto: 20050307110108Z

# gsiftp://quark.hep.lu.se:2811/jobs/52301110065467612188831, jobs, pclong, q
uark.hep.lu.se, local, grid
dn: nordugrid-job-globalid=gsiftp://quark.hep.lu.se:2811/jobs/5230111006546761
2188831, nordugrid-info-group-name=jobs, nordugrid-queue-name=pclong,nordugri
d-cluster-name=quark.hep.lu.se,Mds-Vo-name=local,o=grid
objectClass: Mds
objectClass: nordugrid-job
nordugrid-job-globalid: gsiftp://quark.hep.lu.se:2811/jobs/5230111006546761218
8831
nordugrid-job-globalowner: /O=Grid/O=NorduGrid/OU=fys.ku.dk/CN=Brian Moller An
dersen
nordugrid-job-jobname: selfconsistentMDSSAFNAF100slabonemu
nordugrid-job-submissiontime: 20050305233107Z
nordugrid-job-execcluster: quark.hep.lu.se
nordugrid-job-execqueue: pclong
nordugrid-job-cpucount: 1
nordugrid-job-stdout: output
nordugrid-job-stderr: output
nordugrid-job-submissionui: 130.225.102.149:33846;johansen.fys.ku.dk
nordugrid-job-clientsoftware: nordugrid-0.3.39
nordugrid-job-proxyexpirationtime: 20050306100826Z
nordugrid-job-status: INLRMS: R
nordugrid-job-usedmem: 30736
nordugrid-job-usedwalltime: 2126
nordugrid-job-usedcputime: 2126
nordugrid-job-reqcputime: 10080
nordugrid-job-executionnodes: node3/0
nordugrid-job-lrmscomment: Job started on Sun Mar 06 at 00:32
Mds-validfrom: 20050307110038Z
Mds-validto: 20050307110108Z

# gsiftp://quark.hep.lu.se:2811/jobs/49831110065360908977417, jobs, pclong, q
uark.hep.lu.se, local, grid
dn: nordugrid-job-globalid=gsiftp://quark.hep.lu.se:2811/jobs/4983111006536090
8977417, nordugrid-info-group-name=jobs, nordugrid-queue-name=pclong,nordugri
d-cluster-name=quark.hep.lu.se,Mds-Vo-name=local,o=grid
objectClass: Mds
objectClass: nordugrid-job
nordugrid-job-globalid: gsiftp://quark.hep.lu.se:2811/jobs/4983111006536090897
7417
nordugrid-job-globalowner: /O=Grid/O=NorduGrid/OU=fys.ku.dk/CN=Brian Moller An
dersen
nordugrid-job-jobname: selfconsistentMDSSAFNAF100slabonemu
nordugrid-job-submissiontime: 20050305232920Z
nordugrid-job-execcluster: quark.hep.lu.se
nordugrid-job-execqueue: pclong
nordugrid-job-cpucount: 1
nordugrid-job-stdout: output
nordugrid-job-stderr: output
nordugrid-job-submissionui: 130.225.102.149:33836;johansen.fys.ku.dk
nordugrid-job-clientsoftware: nordugrid-0.3.39
nordugrid-job-proxyexpirationtime: 20050306100826Z
nordugrid-job-status: INLRMS: R
nordugrid-job-usedmem: 30736
nordugrid-job-usedwalltime: 2128
nordugrid-job-usedcputime: 2128
nordugrid-job-reqcputime: 8640
nordugrid-job-executionnodes: node2/0
nordugrid-job-lrmscomment: Job started on Sun Mar 06 at 00:30
Mds-validfrom: 20050307110038Z
Mds-validto: 20050307110108Z

# search result
search: 2
result: 0 Success

# numResponses: 5
# numEntries: 4

```

- LDAP query to obtain the registration entries stored in an Index Service:

```

ldapsearch -h quark.hep.lu.se -p 2135 -x -b 'mds-vo-name=Sweden,o=Grid' -s base giisregistrationstatus
version: 2

#
# filter: (objectclass=*)
# requesting: giisregistrationstatus
#

# sel:sel.hpc2n.umu.se, Sweden, grid
dn: nordugrid-se-name=sel:sel.hpc2n.umu.se, Mds-Vo-name=Sweden,o=grid
objectClass: Mds
objectClass: MdsVoOp
objectClass: MdsService
objectClass: MdsServiceLdap
Mds-Service-type: ldap
Mds-Service-hn: ido-i.hpc2n.umu.se
Mds-Service-port: 2135
Mds-Service-Ldap=suffix: nordugrid-se-name=sel:sel.hpc2n.umu.se, Mds-Vo-name=l
ocal, o=grid
Mds-Service-Ldap-sizelimit: 0
Mds-Service-Ldap-timeout: 45
Mds-Service-Ldap-cachettl: 15

```

```

Mds-Bind-Method-servers: ANONYM-ONLY
Mds-Reg-status: VALID

# bphysics:quark.hep.lu.se, Sweden, grid
dn: nordugrid-se-name=bphysics:quark.hep.lu.se, Mds-Vo-name=Sweden,o=grid
objectClass: Mds
objectClass: MdsVoOp
objectClass: MdsService
objectClass: MdsServiceLdap
Mds-Service-type: ldap
Mds-Service-hn: quark.hep.lu.se
Mds-Service-port: 2135
Mds-Service-Ldap-suffix: nordugrid-se-name=bphysics:quark.hep.lu.se, Mds-Vo-na
me=local, o=grid
Mds-Service-Ldap-sizelimit: 0
Mds-Service-Ldap-timeout: 45
Mds-Service-Ldap-cachettl: 15
Mds-Bind-Method-servers: ANONYM-ONLY
Mds-Reg-status: VALID

# Quark:quark.hep.lu.se, Sweden, grid
dn: nordugrid-se-name=Quark:quark.hep.lu.se, Mds-Vo-name=Sweden,o=grid
objectClass: Mds
objectClass: MdsVoOp
objectClass: MdsService
objectClass: MdsServiceLdap
Mds-Service-type: ldap
Mds-Service-hn: quark.hep.lu.se
Mds-Service-port: 2135
Mds-Service-Ldap-suffix: nordugrid-se-name=Quark:quark.hep.lu.se, Mds-Vo-name=
local, o=grid
Mds-Service-Ldap-sizelimit: 0
Mds-Service-Ldap-timeout: 45
Mds-Service-Ldap-cachettl: 15
Mds-Bind-Method-servers: ANONYM-ONLY
Mds-Reg-status: PURGED

# quark.hep.lu.se, Sweden, grid
dn: nordugrid-cluster-name=quark.hep.lu.se, Mds-Vo-name=Sweden,o=grid
objectClass: Mds
objectClass: MdsVoOp
objectClass: MdsService
objectClass: MdsServiceLdap
Mds-Service-type: ldap
Mds-Service-hn: quark.hep.lu.se
Mds-Service-port: 2135
Mds-Service-Ldap-suffix: nordugrid-cluster-name=quark.hep.lu.se, Mds-Vo-name=l
ocal, o=grid
Mds-Service-Ldap-sizelimit: 0
Mds-Service-Ldap-timeout: 45
Mds-Service-Ldap-cachettl: 0
Mds-Bind-Method-servers: ANONYM-ONLY
Mds-Reg-status: VALID

# files:grid.tsl.uu.se, Sweden, grid
dn: nordugrid-se-name=files:grid.tsl.uu.se, Mds-Vo-name=Sweden,o=grid
objectClass: Mds
objectClass: MdsVoOp
objectClass: MdsService
objectClass: MdsServiceLdap
Mds-Service-type: ldap
Mds-Service-hn: grid.tsl.uu.se
Mds-Service-port: 2135
Mds-Service-Ldap-suffix: nordugrid-se-name=files:grid.tsl.uu.se, Mds-Vo-name=l
ocal, o=grid
Mds-Service-Ldap-sizelimit: 0
Mds-Service-Ldap-timeout: 45
Mds-Service-Ldap-cachettl: 15
Mds-Bind-Method-servers: ANONYM-ONLY
Mds-Reg-status: PURGED

# toto7.lunarc.lu.se, Sweden, grid
dn: nordugrid-cluster-name=toto7.lunarc.lu.se, Mds-Vo-name=Sweden,o=grid
objectClass: Mds
objectClass: MdsVoOp
objectClass: MdsService
objectClass: MdsServiceLdap
Mds-Service-type: ldap
Mds-Service-hn: toto7.lunarc.lu.se
Mds-Service-port: 2135
Mds-Service-Ldap-suffix: nordugrid-cluster-name=toto7.lunarc.lu.se, Mds-Vo-nam
e=local, o=grid
Mds-Service-Ldap-sizelimit: 0
Mds-Service-Ldap-timeout: 45
Mds-Service-Ldap-cachettl: 0
Mds-Bind-Method-servers: ANONYM-ONLY
Mds-Reg-status: VALID

# Svegrid, Sweden, grid
dn: Mds-Vo-name=Svegrid, Mds-Vo-name=Sweden,o=grid
objectClass: Mds
objectClass: MdsVoOp
objectClass: MdsService
objectClass: MdsServiceLdap
Mds-Service-type: ldap
Mds-Service-hn: nexus.svegrid.se
Mds-Service-port: 2135
Mds-Service-Ldap-suffix: Mds-Vo-name=Svegrid, o=grid
Mds-Service-Ldap-sizelimit: 0
Mds-Service-Ldap-timeout: 120
Mds-Service-Ldap-cachettl: 30
Mds-Bind-Method-servers: ANONYM-ONLY
Mds-Reg-status: VALID

# search result
search: 2
result: 0 Success

# numResponses: 8
# numEntries: 7

```

D NorduGrid LDAP schema file

```
#-----
# These classes and attributes are imported from globus mds
# slightly to be proper LDAP schemas.

attributetype ( 1.3.6.1.4.1.11604.2.1.8.0.1
  NAME 'Mds-validfrom'
  DESC 'Object creation time'
  EQUALITY generalizedTimeMatch
  ORDERING generalizedTimeOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.24
  SINGLE-VALUE
)

attributetype ( 1.3.6.1.4.1.11604.2.1.8.0.2
  NAME 'Mds-validto'
  DESC 'Object expiration time'
  EQUALITY generalizedTimeMatch
  ORDERING generalizedTimeOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.24
  SINGLE-VALUE
)

attributetype ( 1.3.6.1.4.1.11604.2.1.8.0.3
  NAME 'Mds-keepsto'
  DESC 'Object purge time'
  EQUALITY generalizedTimeMatch
  ORDERING generalizedTimeOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.24
  SINGLE-VALUE
)

objectclass ( 1.3.6.1.4.1.11604.2.1.8
  NAME 'Mds'
  ABSTRACT
  MUST ( Mds-validfrom $ Mds-validto )
  MAY Mds-keepsto
)

attributetype ( 1.3.6.1.4.1.11604.2.1.8.1.4.0.1
  NAME 'Mds-Vo-name'
  DESC 'Locally unique VO name'
  EQUALITY caseIgnoreMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE
)

objectclass ( 1.3.6.1.4.1.11604.2.1.8.1.4
  NAME 'MdsVo'
  SUP 'Mds'
  STRUCTURAL
  MUST Mds-Vo-name
)

attributetype ( 1.3.6.1.4.1.11604.2.1.8.1.4.1.0.1
  NAME 'Mds-Vo-Op-name'
  DESC 'Locally unique Op name'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE
)

objectclass ( 1.3.6.1.4.1.11604.2.1.8.1.4.1
  NAME 'MdsVoOp'
  SUP 'Mds'
  STRUCTURAL
  MUST Mds-Vo-Op-name
)

attributetype ( 1.3.6.1.4.1.11604.2.1.8.2.7.1.0.1
  NAME 'Mds-Service-type'
  DESC 'Service protocol'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE
)

attributetype ( 1.3.6.1.4.1.11604.2.1.8.2.7.1.0.2
  NAME 'Mds-Service-protocol'
  DESC 'Service protocol UID'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
)

attributetype ( 1.3.6.1.4.1.11604.2.1.8.2.7.1.0.3
  NAME 'Mds-Service-port'
  DESC 'Service TCP port'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE
)

attributetype ( 1.3.6.1.4.1.11604.2.1.8.2.7.1.0.4
  NAME 'Mds-Service-hn'
  DESC 'Service FQDN hostname'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE
)
```

```

)

attributetype ( 1.3.6.1.4.1.11604.2.1.8.2.7.1.0.5
    NAME 'Mds-Service-url'
    DESC 'Service URL'
    EQUALITY caseIgnoreMatch
    ORDERING caseIgnoreOrderingMatch
    SUBSTR caseIgnoreSubstringsMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
    SINGLE-VALUE
)

objectclass ( 1.3.6.1.4.1.11604.2.1.8.2.7.1
    NAME 'MdsService'
    SUP 'Mds'
    AUXILIARY
    MUST ( Mds-Service-type $ Mds-Service-protocol $ Mds-Service-port $ Mds-Service-hn )
    MAY Mds-Service-url
)

attributetype ( 1.3.6.1.4.1.11604.2.1.8.2.7.1.1.0.1
    NAME 'Mds-Service-Ldap-suffix'
    DESC 'DN suffix of service'
    EQUALITY caseIgnoreMatch
    ORDERING caseIgnoreOrderingMatch
    SUBSTR caseIgnoreSubstringsMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
)

attributetype ( 1.3.6.1.4.1.11604.2.1.8.2.7.1.1.0.2
    NAME 'Mds-Service-Ldap-timeout'
    DESC 'suggested timeout'
    EQUALITY caseIgnoreMatch
    ORDERING caseIgnoreOrderingMatch
    SUBSTR caseIgnoreSubstringsMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
)

attributetype ( 1.3.6.1.4.1.11604.2.1.8.2.7.1.1.0.3
    NAME 'Mds-Service-Ldap-sizelimit'
    DESC 'suggested sizelimit'
    EQUALITY caseIgnoreMatch
    ORDERING caseIgnoreOrderingMatch
    SUBSTR caseIgnoreSubstringsMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
)

attributetype ( 1.3.6.1.4.1.11604.2.1.8.2.7.1.1.0.4
    NAME 'Mds-Service-Ldap-cachettl'
    DESC 'suggested cacheability'
    EQUALITY caseIgnoreMatch
    ORDERING caseIgnoreOrderingMatch
    SUBSTR caseIgnoreSubstringsMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
)

attributetype ( 1.3.6.1.4.1.11604.2.1.8.2.7.1.1.0.5
    NAME 'Mds-Service-Ldap-ttl'
    DESC 'suggested ttl'
    EQUALITY caseIgnoreMatch
    ORDERING caseIgnoreOrderingMatch
    SUBSTR caseIgnoreSubstringsMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
)

attributetype ( 1.3.6.1.4.1.11604.2.1.8.2.7.1.1.0.10
    NAME 'Mds-Reg-status'
    DESC 'VALID/INVALID/PURGED'
    EQUALITY caseIgnoreMatch
    ORDERING caseIgnoreOrderingMatch
    SUBSTR caseIgnoreSubstringsMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
)

attributetype ( 1.3.6.1.4.1.11604.2.1.8.2.7.1.1.0.11
    NAME 'Mds-Bind-Method-servers'
    DESC 'AUTHC-ONLY/AUTHC-FIRST/ANONYM-ONLY'
    EQUALITY caseIgnoreMatch
    ORDERING caseIgnoreOrderingMatch
    SUBSTR caseIgnoreSubstringsMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
)

objectclass ( 1.3.6.1.4.1.11604.2.1.8.2.7.1.1
    NAME 'MdsServiceLdap'
    SUP 'MdsService'
    AUXILIARY
    MUST Mds-Service-Ldap-suffix
    MAY ( Mds-Service-Ldap-timeout $ Mds-Service-Ldap-sizelimit $ Mds-Service-Ldap-cachettl $ Mds-Service-Ldap-ttl $ Mds-Reg-status $ Mds-Bind-Method-servers )
)

# attributes for the nordugrid-cluster objectclass
#
attributetype ( 1.3.6.1.4.1.11604.2.1.1.1
    NAME 'nordugrid-cluster-name'
    DESC 'The name of the cluster specified as the domain name of the frontend'
    EQUALITY caseIgnoreMatch
    ORDERING caseIgnoreOrderingMatch
    SUBSTR caseIgnoreSubstringsMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
    SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.2
    NAME 'nordugrid-cluster-aliasname'
    DESC 'The alias name of the cluster'
    EQUALITY caseIgnoreMatch
    ORDERING caseIgnoreOrderingMatch
    SUBSTR caseIgnoreSubstringsMatch
    SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
    SINGLE-VALUE)

attributetype ( 1.3.6.1.4.1.11604.2.1.1.3
    NAME 'nordugrid-cluster-contactstring'
    DESC 'The URL of the job submission service running on the cluster frontend'
    EQUALITY caseExactIA5Match
    SUBSTR caseExactIA5SubstringsMatch

```

```

SYNTAX 1.3.6.1.4.1.1466.115.121.1.26
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.4
  NAME 'nordugrid-cluster-support'
  DESC 'RFC822 email address of support'
  EQUALITY caseIgnoreIA5Match
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26{256})

attributetype ( 1.3.6.1.4.1.11604.2.1.1.5
  NAME 'nordugrid-cluster-lrms-type'
  DESC 'The type of the Local Resource Management System'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.6
  NAME 'nordugrid-cluster-lrms-version'
  DESC 'The version of the Local Resource Management System'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.7
  NAME 'nordugrid-cluster-lrms-config'
  DESC 'Additional remark on the LRMS config'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.8
  NAME 'nordugrid-cluster-architecture'
  DESC 'The architecture of the cluster nodes'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.9
  NAME 'nordugrid-cluster-opsys'
  DESC 'The operating system of the machines of the cluster'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15)

attributetype ( 1.3.6.1.4.1.11604.2.1.1.10
  NAME 'nordugrid-cluster-homogeneity'
  DESC 'A logical flag indicating the homogeneity of the cluster nodes'
  EQUALITY booleanMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.7
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.11
  NAME 'nordugrid-cluster-nodectpu'
  DESC 'The cpu type of the nodes expressed in a fixed form (model name + MHz)'
  EQUALITY caseIgnoreMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.12
  NAME 'nordugrid-cluster-nodememory'
  DESC 'The amount of memory which can be guaranteed to be available on the node in MB'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.13
  NAME 'nordugrid-cluster-totalcpus'
  DESC 'The total number of cpus in the cluster'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.14
  NAME 'nordugrid-cluster-cpudistribution'
  DESC 'The cpu distribution of the nodes given in the form of 1cpu:3 2cpu:4 4cpu:1'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.15
  NAME 'nordugrid-cluster-sessiondir-free'
  DESC 'Free disk space in MB of the sessiondirectory on the cluster'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.16
  NAME 'nordugrid-cluster-sessiondir-total'
  DESC 'Total disk space in MB of the sessiondirectory on the cluster'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.17
  NAME 'nordugrid-cluster-cache-free'
  DESC 'Free disk space in MB of the cache area on the cluster'
  EQUALITY integerMatch

```

```

ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.18
  NAME 'nordugrid-cluster-cache-total'
  DESC 'Total disk space in MB of the cache area on the cluster'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.19
  NAME 'nordugrid-cluster-runtimeenvironment'
  DESC 'preinstalled software packages of the cluster'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.20
  NAME 'nordugrid-cluster-locale'
  DESC 'The URL of a storage element considered to be local to the cluster'
  EQUALITY caseExactMatch
  ORDERING caseExactOrderingMatch
  SUBSTR caseExactSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.21
  NAME 'nordugrid-cluster-middleware'
  DESC 'The middleware packages on the cluster'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.22
  NAME 'nordugrid-cluster-totaljobs'
  DESC 'The total number of jobs (Grid + non-Grid) in the cluster'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.23
  NAME 'nordugrid-cluster-usedcpus'
  DESC 'The total number of occupied cpus in the cluster'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.24
  NAME 'nordugrid-cluster-queuedjobs'
  DESC 'The total number of jobs (grid and non-grid) not-yet running: preparing or waiting to run
  on the cluster, either in the grid-manager or in the LRMS.
  The attribute is TO BE DEPRECATED'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.25
  NAME 'nordugrid-cluster-location'
  DESC 'The geographical location of the cluster expressed in terms of a Postal ZIP code'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.26
  NAME 'nordugrid-cluster-owner'
  DESC 'The owner of the resource'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.27
  NAME 'nordugrid-cluster-issuerca'
  DESC 'The DN of the Certificate Authority which issued the certificate of the cluster'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.28
  NAME 'nordugrid-cluster-nodeaccess'
  DESC 'The inbound/outbound network accessibility of the nodes'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44 )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.29
  NAME 'nordugrid-cluster-comment'
  DESC 'Free form comment'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.30
  NAME 'nordugrid-cluster-interactive-contactstring'
  DESC 'The URL for interactive login'
  EQUALITY caseExactIA5Match
  SUBSTR caseExactIA5SubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26 )

```



```

attributetype ( 1.3.6.1.4.1.11604.2.1.1.31
  NAME 'nordugrid-cluster-benchmark'
  DESC '@ separated benchmark_name, benchmark_value pair characterizing the cluster nodes'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.32
  NAME 'nordugrid-cluster-sessiondir-lifetime'
  DESC 'The lifetime of the sessiondir after the job has completed (in minutes)'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.33
  NAME 'nordugrid-cluster-prelrmsqueued'
  DESC 'The total number of grid jobs not-yet reached the LRMS: preparing or queuing in the grid-layer'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.34
  NAME 'nordugrid-cluster-issuerca-hash'
  DESC 'The HASH of the Certificate Authority which issued the certificate for the cluster'
  EQUALITY caseExactMatch
  ORDERING caseExactOrderingMatch
  SUBSTR caseExactSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.35
  NAME 'nordugrid-cluster-trustedca'
  DESC 'The DN of a Certificate Authority trusted by the cluster'
  EQUALITY caseIgnoreMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.36
  NAME 'nordugrid-cluster-acl'
  DESC 'Cluster authorization information'
  EQUALITY caseExactMatch
  ORDERING caseExactOrderingMatch
  SUBSTR caseExactSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44 )

attributetype ( 1.3.6.1.4.1.11604.2.1.1.37
  NAME 'nordugrid-cluster-credentialexpirationtime'
  DESC 'The expiration date of the shortest living credential affecting the cluster's x509 environment in GMT'
  EQUALITY generalizedTimeMatch
  ORDERING generalizedTimeOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.24
  SINGLE-VALUE )

objectclass ( 1.3.6.1.4.1.11604.2.1.1
  NAME 'nordugrid-cluster'
  DESC 'Description of a Nordugrid cluster'
  SUP 'Mds'
  STRUCTURAL
  MUST ( nordugrid-cluster-name $ nordugrid-cluster-contactstring )
  MAY ( nordugrid-cluster-aliasname $ nordugrid-cluster-support $
    nordugrid-cluster-lrms-type $ nordugrid-cluster-lrms-version $
    nordugrid-cluster-lrms-config $ nordugrid-cluster-architecture $
    nordugrid-cluster-opsys $ nordugrid-cluster-homogeneity $
    nordugrid-cluster-nodescpu $ nordugrid-cluster-nodememory $
    nordugrid-cluster-totalcpus $ nordugrid-cluster-cpudistribution $
    nordugrid-cluster-sessiondir-free $ nordugrid-cluster-sessiondir-total $
    nordugrid-cluster-cache-free $ nordugrid-cluster-cache-total $
    nordugrid-cluster-runtimeenvironment $ nordugrid-cluster-locale $
    nordugrid-cluster-middleware $ nordugrid-cluster-totaljobs $
    nordugrid-cluster-usedcpus $ nordugrid-cluster-queuedjobs $
    nordugrid-cluster-location $ nordugrid-cluster-owner $
    nordugrid-cluster-issuerca $ nordugrid-cluster-nodeaccess $
    nordugrid-cluster-comment $ nordugrid-cluster-interactive-contactstring $
    nordugrid-cluster-benchmark $ nordugrid-cluster-sessiondir-lifetime $
    nordugrid-cluster-prelrmsqueued $ nordugrid-cluster-issuerca-hash $
    nordugrid-cluster-trustedca $ nordugrid-cluster-acl $
    nordugrid-cluster-credentialexpirationtime ) )

#-----
# attributes for the nordugrid-info-group objectclass
#
attributetype ( 1.3.6.1.4.1.11604.2.1.2.1
  NAME 'nordugrid-info-group-name'
  DESC 'Locally unique info group name'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE )

objectclass ( 1.3.6.1.4.1.11604.2.1.2
  NAME 'nordugrid-info-group'
  DESC 'A general entry for grouping together MDS entries'
  SUP 'Mds'
  STRUCTURAL
  MUST ( nordugrid-info-group-name ) )

#-----
# attributes for the nordugrid-queue objectclass
#
attributetype ( 1.3.6.1.4.1.11604.2.1.3.1
  NAME 'nordugrid-queue-name'
  DESC 'The queue name'
  EQUALITY caseExactMatch
  ORDERING caseExactOrderingMatch
  SUBSTR caseExactSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.2

```

```

NAME 'nordugrid-queue-status'
DESC 'The queue status'
EQUALITY caseIgnoreMatch
ORDERING caseIgnoreOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.3
NAME 'nordugrid-queue-running'
DESC 'Number of running jobs (Grid + non-Grid) in the queue with multi-node jobs multiplicity'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.4
NAME 'nordugrid-queue-queued'
DESC 'The number of jobs (Grid + non-Grid) waiting in the queue.
The attribute is TO BE DEPRECATED'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.5
NAME 'nordugrid-queue-maxrunning'
DESC 'The maximum number of jobs allowed to run from this queue'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.6
NAME 'nordugrid-queue-maxqueueable'
DESC 'The maximum number of jobs allowed to reside in the queue'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.7
NAME 'nordugrid-queue-maxuserun'
DESC 'Maximum number of jobs a user can run at the same time'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.8
NAME 'nordugrid-queue-maxcputime'
DESC 'The maximum cputime allowed in this queue (in minutes)'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.9
NAME 'nordugrid-queue-mincputime'
DESC 'The minimum possible cputime of this queue (in minutes)'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.10
NAME 'nordugrid-queue-defaultcputime'
DESC 'The default cputime assigned to this queue (in minutes)'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.11
NAME 'nordugrid-queue-schedulingpolicy'
DESC 'The scheduling policy of the queue (i.e. FIFO)'
EQUALITY caseIgnoreMatch
ORDERING caseIgnoreOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.12
NAME 'nordugrid-queue-totalcpus'
DESC 'Total number of cpus assigned to the queue'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.13
NAME 'nordugrid-queue-nodectpu'
DESC 'The cpu type of the nodes assigned to the queue (model name + MHz)'
EQUALITY caseIgnoreMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.14
NAME 'nordugrid-queue-nodememory'
DESC 'The installed memory of a node assigned to the queue in MB'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.15

```

```

NAME 'nordugrid-queue-architecture'
DESC 'The architecture of the machines in the queue'
EQUALITY caseIgnoreMatch
ORDERING caseIgnoreOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.16
NAME 'nordugrid-queue-opsys'
DESC 'The operating system of the nodes in the queue'
EQUALITY caseIgnoreMatch
ORDERING caseIgnoreOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.15)

attributetype ( 1.3.6.1.4.1.11604.2.1.3.17
NAME 'nordugrid-queue-gridrunning'
DESC 'Number of running Grid jobs in the queue with multi-node jobs multiplicity'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.18
NAME 'nordugrid-queue-gridqueued'
DESC 'The number of Grid jobs waiting in the LRMS queue'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.19
NAME 'nordugrid-queue-comment'
DESC 'Free form comment'
EQUALITY caseIgnoreMatch
ORDERING caseIgnoreOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.20
NAME 'nordugrid-queue-benchmark'
DESC 'Colon separated benchmark_name, benchmark_value pair characterizing the queue'
EQUALITY caseIgnoreMatch
ORDERING caseIgnoreOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.21
NAME 'nordugrid-queue-homogeneity'
DESC 'A logical flag indicating the homogeneity of the queue nodes'
EQUALITY booleanMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.7
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.22
NAME 'nordugrid-queue-prelrmsqueued'
DESC 'The number of Grid jobs belonging to this queue being processed
or waiting in the Grid-layer before the LRMS submission.'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.23
NAME 'nordugrid-queue-localqueued'
DESC 'The number of non-Grid jobs waiting in the LRMS queue.'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.24
NAME 'nordugrid-queue-maxwalltime'
DESC 'The maximum walltime allowed in this queue (in minutes)'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.25
NAME 'nordugrid-queue-minwalltime'
DESC 'The minimum possible walltime of this queue (in minutes)'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.26
NAME 'nordugrid-queue-defaultwalltime'
DESC 'The default walltime assigned to this queue (in minutes)'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.3.27
NAME 'nordugrid-queue-maxtotalcputime'
DESC 'The maximum total cputime allowed in this queue (in minutes)'
EQUALITY integerMatch
ORDERING integerOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
SINGLE-VALUE )

objectclass ( 1.3.6.1.4.1.11604.2.1.3
NAME 'nordugrid-queue'
DESC 'An LRMS queue'

```

```

SUP 'Mds'
STRUCTURAL
MUST ( nordugrid-queue-name $ nordugrid-queue-status)
MAY ( nordugrid-queue-running $ nordugrid-queue-queued $
      nordugrid-queue-maxrunning $ nordugrid-queue-maxqueueable$
      nordugrid-queue-maxuserun $ nordugrid-queue-maxcputime $
      nordugrid-queue-mincputime $ nordugrid-queue-defaultcputime $
      nordugrid-queue-schedulingpolicy $ nordugrid-queue-totalcpus $
      nordugrid-queue-nodescpu $ nordugrid-queue-nodememory $
      nordugrid-queue-opsys $ nordugrid-queue-architecture $
      nordugrid-queue-gridrunning $ nordugrid-queue-gridqueued $
      nordugrid-queue-comment $ nordugrid-queue-benchmark $
      nordugrid-queue-homogeneity $ nordugrid-queue-prelmsqueued $
      nordugrid-queue-localqueued $ nordugrid-queue-maxwalltime $
      nordugrid-queue-minwalltime $ nordugrid-queue-defaultwalltime $
      nordugrid-queue-maxtotalcputime ))

#-----
#Attributes for the nordugrid-job objectclass
#
attributetype ( 1.3.6.1.4.1.11604.2.1.4.1
  NAME 'nordugrid-job-globalid'
  DESC 'The global job identifier string'
  EQUALITY caseExactIA5Match
  SUBSTR caseExactIA5SubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.26
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.2
  NAME 'nordugrid-job-globalowner'
  DESC 'The SubjectName of the job owner'
  EQUALITY caseExactMatch
  ORDERING caseExactOrderingMatch
  SUBSTR caseExactSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.3
  NAME 'nordugrid-job-execcluster'
  DESC 'The name of the execution cluster'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.4
  NAME 'nordugrid-job-execqueue'
  DESC 'The name of the execution queue'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.5
  NAME 'nordugrid-job-stdout'
  DESC 'The name of the file which contains the stdout'
  EQUALITY caseExactMatch
  ORDERING caseExactOrderingMatch
  SUBSTR caseExactSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.6
  NAME 'nordugrid-job-stderr'
  DESC 'The name of the file which contains the stderr'
  EQUALITY caseExactMatch
  ORDERING caseExactOrderingMatch
  SUBSTR caseExactSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.7
  NAME 'nordugrid-job-stdin'
  DESC 'The name of the file which contains the stdin'
  EQUALITY caseExactMatch
  ORDERING caseExactOrderingMatch
  SUBSTR caseExactSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.8
  NAME 'nordugrid-job-reqcputime'
  DESC 'The cputime request by the job in minutes'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.9
  NAME 'nordugrid-job-status'
  DESC 'The status of the grid job'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.10
  NAME 'nordugrid-job-queuerank'
  DESC 'The queue position of the job'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.11
  NAME 'nordugrid-job-comment'
  DESC 'Free form comment about the job'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

```

```

attributetype ( 1.3.6.1.4.1.11604.2.1.4.12
  NAME 'nordugrid-job-submissionui'
  DESC 'The name of the UI from where the job was submitted'
  EQUALITY caseIgnoreMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.13
  NAME 'nordugrid-job-submissiontime'
  DESC 'The submission time of the job in GMT'
  EQUALITY generalizedTimeMatch
  ORDERING generalizedTimeOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.24
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.14
  NAME 'nordugrid-job-usedcputime'
  DESC 'The consumed cputime of the job in minutes'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.15
  NAME 'nordugrid-job-usedwalltime'
  DESC 'The consumed walltime of the job in minutes'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.16
  NAME 'nordugrid-job-sessiondirerasetime'
  DESC 'The date when the session dir will be deleted in GMT'
  EQUALITY generalizedTimeMatch
  ORDERING generalizedTimeOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.24
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.17
  NAME 'nordugrid-job-usedmem'
  DESC 'The memory usage of the job (in KB)'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.18
  NAME 'nordugrid-job-errors'
  DESC 'Error messages from the cluster'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.19
  NAME 'nordugrid-job-jobname'
  DESC 'The jobname specified by the user with the jobname RSL attribute'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.20
  NAME 'nordugrid-job-runtimeenvironment'
  DESC 'The runtimeenvironment requested by the job'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.21
  NAME 'nordugrid-job-cpucount'
  DESC 'The number of CPUs requested by the job'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.22
  NAME 'nordugrid-job-executionnodes'
  DESC 'The list of nodenames where the job is running'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.23
  NAME 'nordugrid-job-gmlog'
  DESC 'The name of the directory which contains the grid session related logs'
  EQUALITY caseExactMatch
  ORDERING caseExactOrderingMatch
  SUBSTR caseExactSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.24
  NAME 'nordugrid-job-clientsoftware'
  DESC 'The client software which submitted the job'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15)

attributetype ( 1.3.6.1.4.1.11604.2.1.4.25
  NAME 'nordugrid-job-proxyexpirationtime'
  DESC 'The expiration time of the proxy of the job in GMT'

```

```

EQUALITY generalizedTimeMatch
ORDERING generalizedTimeOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.24
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.26
  NAME 'nordugrid-job-completiontime'
  DESC 'The completion time of the grid job in GMT'
  EQUALITY generalizedTimeMatch
  ORDERING generalizedTimeOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.24
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.27
  NAME 'nordugrid-job-exitcode'
  DESC 'The exit code of the executable of the job obtained from the LRMS'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.28
  NAME 'nordugrid-job-rerunable'
  DESC 'Rerunability of the FAILED grid jobs'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.4.29
  NAME 'nordugrid-job-reqwalltime'
  DESC 'The request wallclock time of the job in minutes'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

objectclass ( 1.3.6.1.4.1.11604.2.1.4
  NAME 'nordugrid-job'
  DESC 'A Grid job'
  SUP 'Mds'
  STRUCTURAL
  MUST ( nordugrid-job-globalid $ nordugrid-job-globalowner $
    nordugrid-job-status )
  MAY ( nordugrid-job-queuerank $ nordugrid-job-submissionui $
    nordugrid-job-submissiontime $
    nordugrid-job-usedcputime $ nordugrid-job-usedwalltime $
    nordugrid-job-usedmem $ nordugrid-job-comment $
    nordugrid-job-execcluster $ nordugrid-job-execqueue $
    nordugrid-job-stdout $ nordugrid-job-stderr $
    nordugrid-job-stdin $
    nordugrid-job-sessiondirerasetime $ nordugrid-job-reqcputime $
    nordugrid-job-errors $ nordugrid-job-jobname $
    nordugrid-job-runtimeenvironment $ nordugrid-job-cpucount $
    nordugrid-job-executionnodes $ nordugrid-job-gmlog $
    nordugrid-job-clientsoftware $ nordugrid-job-proxyexpirationtime $
    nordugrid-job-completiontime $ nordugrid-job-exitcode $
    nordugrid-job-rerunable $ nordugrid-job-reqwalltime))

#-----
# attributes for the nordugrid-authuser objectclass
#

attributetype ( 1.3.6.1.4.1.11604.2.1.5.1
  NAME 'nordugrid-authuser-name'
  DESC 'The Common Name of the authorized user plus a local unique number'
  EQUALITY caseIgnoreMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.5.2
  NAME 'nordugrid-authuser-sn'
  DESC 'The SubjectName of the authorized user'
  EQUALITY caseIgnoreMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.5.3
  NAME 'nordugrid-authuser-freecpus'
  DESC 'The number of freely available cpus with their timelimits in minutes
    for a user in the queue. Given in the form ncpus:min,
    min is optional (example: 2 4:25 5:180)'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.5.4
  NAME 'nordugrid-authuser-diskspace'
  DESC 'The free diskpace available for the job (in MB)'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.5.5
  NAME 'nordugrid-authuser-queuelength'
  DESC 'The number of queuing jobs of a particular user,
    both queuing in the LRMS and in the Grid-layer'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

objectclass ( 1.3.6.1.4.1.11604.2.1.5
  NAME 'nordugrid-authuser'
  DESC 'An authorised user of a NorduGrid cluster'

```

```

SUP 'Mds'
STRUCTURAL
MUST ( nordugrid-authuser-name $ nordugrid-authuser-sn )
MAY ( nordugrid-authuser-queue length $ nordugrid-authuser-diskspace $
      nordugrid-authuser-freecpu )

#-----
#
# nordugrid-se

attributetype ( 1.3.6.1.4.1.11604.2.1.6.1
  NAME 'nordugrid-se-name'
  DESC 'The name of the Storage Element'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.6.2
  NAME 'nordugrid-se-aliasname'
  DESC 'The alias name of the SE'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.6.3
  NAME 'nordugrid-se-type'
  DESC 'The type of the SE'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.6.4
  NAME 'nordugrid-se-freespace'
  DESC 'The free space available in the SE (in MB)'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.6.5
  NAME 'nordugrid-se-url'
  DESC 'The URL to contact the Storage Element'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.6.6
  NAME 'nordugrid-se-authuser'
  DESC 'The DN of an authorized user'
  EQUALITY caseExactMatch
  ORDERING caseExactOrderingMatch
  SUBSTR caseExactSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.6.7
  NAME 'nordugrid-se-location'
  DESC 'The geographical location of the SE expressed in terms of a Postal ZIP code: SE-22363'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44 )

attributetype ( 1.3.6.1.4.1.11604.2.1.6.8
  NAME 'nordugrid-se-owner'
  DESC 'The owner of the resource'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.6.9
  NAME 'nordugrid-se-issuerca'
  DESC 'The DN of the Certificate Authority which issued the certificate of the SE'
  EQUALITY caseExactMatch
  ORDERING caseExactOrderingMatch
  SUBSTR caseExactSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.6.10
  NAME 'nordugrid-se-totalspace'
  DESC 'The total capacity of the SE (in MB)'
  EQUALITY integerMatch
  ORDERING integerOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.27
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.6.11
  NAME 'nordugrid-se-middleware'
  DESC 'The middleware packages on the SE'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.6.12
  NAME 'nordugrid-se-comment'
  DESC 'Free form comment'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.6.13
  NAME 'nordugrid-se-accesscontrol'
  DESC 'The access control framework of the SE'

```

```

EQUALITY caseIgnoreMatch
ORDERING caseIgnoreOrderingMatch
SUBSTR caseIgnoreSubstringsMatch
SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.6.14
  NAME 'nordugrid-se-issuerca-hash'
  DESC 'The HASH of the Certificate Authority which issued the certificate for the SE'
  EQUALITY caseExactMatch
  ORDERING caseExactOrderingMatch
  SUBSTR caseExactSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.6.15
  NAME 'nordugrid-se-trustedca'
  DESC 'The DN of a Certificate Authority trusted by the SE'
  EQUALITY caseExactMatch
  ORDERING caseExactOrderingMatch
  SUBSTR caseExactSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.6.16
  NAME 'nordugrid-se-acl'
  DESC 'Storage Element authorization information'
  EQUALITY caseExactMatch
  ORDERING caseExactOrderingMatch
  SUBSTR caseExactSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

objectclass ( 1.3.6.1.4.1.11604.2.1.6
  NAME 'nordugrid-se'
  DESC 'A storage element in the Nordugrid'
  SUP 'Mds'
  STRUCTURAL
  MUST ( nordugrid-se-name $ nordugrid-se-url )
  MAY ( nordugrid-se-aliasname $ nordugrid-se-type $
        nordugrid-se-freespace $ nordugrid-se-authuser $
        nordugrid-se-location $ nordugrid-se-owner $
        nordugrid-se-issuerca $ nordugrid-se-totalspace $
        nordugrid-se-middleware $ nordugrid-se-comment $
        nordugrid-se-accesscontrol $ nordugrid-se-issuerca-hash $
        nordugrid-se-trustedca $ nordugrid-se-acl )

#-----
# nordugrid-rc
#
attributetype ( 1.3.6.1.4.1.11604.2.1.7.1
  NAME 'nordugrid-rc-name'
  DESC 'The domain name of the machine hosting the Replica Catalog'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.7.2
  NAME 'nordugrid-rc-aliasname'
  DESC 'The alias name of the rc'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.7.3
  NAME 'nordugrid-rc-baseurl'
  DESC 'The URL of the Replica Catalog'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15
  SINGLE-VALUE )

attributetype ( 1.3.6.1.4.1.11604.2.1.7.4
  NAME 'nordugrid-rc-authuser'
  DESC 'An authorized user of the replica catalog'
  EQUALITY caseExactMatch
  ORDERING caseExactOrderingMatch
  SUBSTR caseExactSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.7.5
  NAME 'nordugrid-rc-location'
  DESC 'The geographical location of the RC expressed in terms of a Postal ZIP code'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.44 )

attributetype ( 1.3.6.1.4.1.11604.2.1.7.6
  NAME 'nordugrid-rc-owner'
  DESC 'The owner of the resource'
  EQUALITY caseIgnoreMatch
  ORDERING caseIgnoreOrderingMatch
  SUBSTR caseIgnoreSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

attributetype ( 1.3.6.1.4.1.11604.2.1.7.7
  NAME 'nordugrid-rc-issuerca'
  DESC 'The DN of the Certificate Authority which issued the certificate of the RC'
  EQUALITY caseExactMatch
  ORDERING caseExactOrderingMatch
  SUBSTR caseExactSubstringsMatch
  SYNTAX 1.3.6.1.4.1.1466.115.121.1.15 )

objectclass ( 1.3.6.1.4.1.11604.2.1.7
  NAME 'nordugrid-rc'
  DESC 'A replica catalogue in the Nordugrid'
  SUP 'Mds'
  STRUCTURAL
  MUST ( nordugrid-rc-name $ nordugrid-rc-baseurl )
  MAY ( nordugrid-rc-aliasname $ nordugrid-rc-authuser $
        nordugrid-rc-location $ nordugrid-rc-owner $
        nordugrid-rc-issuerca )

```


References

- [1] NorduGrid Web site. <http://www.nordugrid.org>
- [2] Openldap. <http://www.openldap.org>
- [3] Monitoring and Discovery Services. <http://www.globus.org/mds/mds2/>
- [4] The NorduGrid Grid Manager And GridFTP Server: Description And Administrator's Manual. NORDUGRID-TECH-2
- [5] The ARC Computational Job Management Module - A-REX. NORDUGRID-TECH-14
- [6] Lightweight middleware for Grid Computing. <http://glite.cern.ch/>
- [7] European Middleware Initiative, Web site. <http://www.eu-emi.eu>
- [8] The NorduGrid Brokering Algorithm, M.Ellert, <http://www.nordugrid.org/papers.html>
- [9] xRSL (Extended Resource Specification Language), O.Smirnova. NORDUGRID-MANUAL-4
- [10] Runtime Environment Registry, <http://gridrer.csc.fi/>
- [11] Usage statistics and usage patterns on the NorduGrid, K.Pajchel, <http://www.nordugrid.org/papers.html>
- [12] The Grid Monitor: Usage Manual, <http://www.nordugrid.org/documents/monitor.pdf>
- [13] The GLUE Information model versin 1.2 <http://infnforgue.cnaf.infn.it/glueinfomodel/>
- [14] The NorduGrid Toolkit User Interface: User's Manual <http://www.nordugrid.org/documents/ui.pdf>
- [15] The NorduGrid "Smart" Storage Element, A.Konstantinov, <http://www.nordugrid.org/papers.html>
- [16] The Gridsite, <http://www.gridsite.org>
- [17] BDII homepage, <https://twiki.cern.ch/twiki/bin/view/EGEE/BDIIv4>



Figure 7: Resources (hosts) and EGIISes are linked via the registration process creating a multi-rooted tree topology