

GLUE Schema v1.2 / v1.3 Mapping to Old ClassAd
Format (v1.1)
June 20, 2008

For v1.1:
Gabriele Garzoglio, Tanya Levshina

For v1.0:
Sergio Andreozzi (*INFN*)
Gabriele Garzoglio, Tanya Levshina, Sudhamsh Reddy (*Fermilab*)
Marco Mambelli (*University of Chicago*)
Alain Roy (*University of Wisconsin*),
Shaowen Wang (*University of Iowa*),
Torre Wenaus (*BNL*)

July 14, 2008

Abstract

This document describes a mapping between the GLUE schema v1.2 / v1.3 and old ClassAd format. The old ClassAd representation of resource characteristics can be used in out-of-the-box resource selection technologies, such as the Condor Matchmaking Service. This mapping has been written focussing on the problem of resource selection and is presented as a series of rules with examples.

Contents

1	Introduction	3
2	Limitations	3
3	Mapping Rules	3
3.1	Implementations	3
3.2	Basic Rules	4
3.2.1	Definitions	4
3.2.2	Considerations	6
3.2.3	Example	6
3.3	Extended Rules v1.0	6
3.3.1	Definitions	7
3.3.2	Example	8
3.4	Extended Rules v1.1	8
3.4.1	Definitions	8
	References	10

1 Introduction

The GLUE schema [1] provides a model to describe Grid resources and their relationship. Previous works [2] [3] present a mapping between the GLUE schema and different popular data formats, such as LDIF and XML. This document presents a mapping between the GLUE schema and the old ClassAd format. Such representation can be immediately used by out-of-the-box software, such as the Condor Matchmaking Service, to implement resource selection.

2 Limitations

In this document, the main motivation for the mapping of the GLUE schema to old ClassAd format is providing a resource description that can be easily used in resource selection and discovery services. We do not focus on other potential uses of such mappings, such as resource monitoring or job environment enhancing. With “job environment enhancing” we intend the process of adding to the job environment by dereferencing abstract attribute values to concrete values, after resource selection is completed¹. For these reasons, some entities of the GLUE schema, such as “Job” or “Computing Element”, or the “Path” attribute of the “Location” entity², are explicitly excluded from the mapping.

3 Mapping Rules

This section presents a set of rules to translate the GLUE schema into old ClassAd format. The section describes three sets of rules: (1) basic rules, (2) extended rules v1.0, (3) extended rules v1.1. There are different set of rules because new rules have been introduced at different times. These rules apply to the GLUE schema v1.2 and v1.3, unless otherwise noted.

3.1 Implementations

In March 2006, the basic rules have been implemented as the “old_ClassAd” dialect (OSG_CE sensor plug-in) of the CEMon service (v1.7.2) from gLite [4].

The mapping uses as input the GLUE representation of the resources in LDIF format [2, 6] from the Generic Information Providers [5]. As such, it should be noted that it is not mandatory for all the GLUE attributes to have defined values at all sites. Developers should keep this in mind when implementing code that works with this mapping e.g. Condor matchmaking functions.

The extended rules v1.0 are an extension/modification to the basic rules and focus on the “VOView” and “Location” entities. The extended rules v1.0 for the VOView entity has been implemented in CEMon v1.7.2 as of May 2006.

¹In condor ClassAd, attribute dereferencing is expressed with the `$(AttributeName)` syntax

²The “Location” entity has been deprecated in Glue Schema v1.3 and has been replaced by the “Software” entity.

The extended rules v1.1 are an extension to the basic rules and extended rules v1.0. These new rules provide mappings for the “SE” and “Software” entities. These rules have been implemented in CEMon v1.9.1-0 as of June 2008.

3.2 Basic Rules

In this section, we describe the basic rules of the mapping of LDIF representation of Grid resources based on the GLUE Schema into the old ClassAd format.

3.2.1 Definitions

- Attribute names are built using the same rules as for the LDIF mapping [2] and are not described here.
- The mapping is built by considering all the possible combinations of inter-related ComputingElement, Cluster, and Subcluster entities. In other words, each combination contains a single ComputingElement, Cluster, and Subcluster entity. Attributes in each combination are then mapped to a single old ClassAd. In general, the full representation of the GLUE schema for a site in old ClassAd format consists of a set of old ClassAds. Each ClassAd represents a sort of virtual homogeneous cluster with a single entry point and waiting queue.
- Attributes of the CE, Cluster, and Subcluster entities that, according to the schema, do not appear with more than one value (e.g. GlueCEInfoHostName, GlueCEInfoTotalCPUs) are translated in old ClassAd format as `AttributeName = Value`
- Attributes that can appear with more than one value (indicated in the UML representation of the GLUE schema with an asterisk [*] next to the attribute type or without the keyword `SINGLE-VALUE` in the LDAP attribute definition), must be treated differently. In an old ClassAd, in fact, attributes must be unique, i.e. they appear only once in the ClassAd. Every GLUE attribute that can be present multiple times with different values (e.g. GlueHostApplicationSoftwareRunTimeEnvironment) appears only once in the old ClassAd mapping. The multiple values are represented as a single comma-separated string of the original values. For the purpose of resource selection using the Condor Matchmaking service, the values can be individually used in expressions using custom (call-out) functions. For example, to select a ClassAd that contains a certain value in one of these attributes, the Matchmaking Service could provide the users a syntax like

```
MultiValueAttribute ==  
    matchString("myValue",MultiValueAttributeValues)
```

- In some General Information Provider implementations (GIP), the ContactString Info attribute of the ComputingElement entity is not necessarily a GRAM URL and may have the form `<GRAM-URL>-<Batch-System-Queue>`. This form cannot be used directly as a contact string by most Grid clients. In LCG, the Resource Broker processes

this value so that it can be used to submit jobs. In the mapping to old ClassAd, instead, the value is published in the form <GRAM-URL>, irrespectively of the GIP implementation. The CEMon implementation uses the value of the attribute GlueCEName as <Batch-System-Queue>, in order to match and strip it off from the ContactString. This is important in order to use the Condor Matchmaking Service out-of-the-box for Resource Selection.

- The “Job” and “VOView” entities of “Computing Element” AND the “Location” entity of “SubCluster” are not considered for the mapping to old ClassAd in the basic mapping rules. “VOView” and “Location” are used in the extended mapping rules v1.0.
- Mappings for “SE” and “Software” entities are provided by the extended mapping rules v1.1.
- The following rules are followed to quote string values in old ClassAd

ClassAd Keywords

Special keyword are not quoted.

Special keywords: undefined, true, t, false, f, error

Note: int, float, and the like are not keywords in the ClassAd language and have no special meaning

Variables

If an attribute refers to another attribute, the attribute name is a variable, and is not quoted.

For instance, if you have:

```
X = 3
```

```
Y = X
```

then if you evaluate Y, you get 3. In this case, no quote marks are used.

Variables can be combined into expressions, like:

```
Y = X + 3
```

```
Y = X > 3
```

Numbers

Numbers are not quoted, unless you want them to be interpreted as strings. There are floating point numbers and integers. If the C function strtod() recognizes it as a floating point number, then it is, otherwise it isn't.

Strings

The only other element that you can have is a string, and it must be quoted.

```
X = "blah"
```

```
Y = X
```

```
Z = "foo"
```

A backslash (\) can also be used for quoting the following strings:

```
\ (\<blank>)  
\"  
\n  
\r
```

3.2.2 Considerations

These rules are useful when using the Condor Matchmaking Service to implement Resource Selection. These rules are implementation-specific, thus are kept outside of the “generic” mapping rules.

- When using the Condor Matchmaking Service for Resource Selection, every generated ClassAd should have the following three attributes added to it: “MyType”, “TargetType”, and “Name”. These attribute are used internally by Condor for Matchmaking and have the following constant values:

```
MyType = "Machine"
```

```
TargetType = "Job"
```

```
Name = <GlueCeName>.<GlueClusterName>.<GlueSubclusterName>
```

where the notation <GlueCeName> indicates the value of the GlueCeName attribute and similarly for the other attributes.

It should be noted that the attribute “Name” must be unique among all ClassAds considered for resource selection.

3.2.3 Example

Figure 6 and 7 show an example of a set of two old ClassAds. The set is derived applying the “basic rules” (section 3.2) to the LDIF representation of the example GLUE schema in figures 1, 2, 3, 4, and 5.

The attributes in the examples are based on the Glue Schema v1.2. Even if the attributes introduced in v1.3 are not present, the mapping rules apply nevertheless.

It should be also noted that the LDIF example is derived from a real configuration of resources and does not include the GLUE attributes “GlueClusterTmpDir”, “GlueClusterWNTmpDir”, “GlueHostArchitecturePlatformType”, “GlueHostProcessorVersion”, “GlueHostProcessorInstructionSet”, “GlueHostProcessorOtherDescription”, and “GlueCEInfoGRAMVersion”. These attributes are included for completeness in the old ClassAd set example.

3.3 Extended Rules v1.0

In this section, we describe the extended rules of the mapping of LDIF representation of Grid resources based on the GLUE Schema into the old ClassAd format. The extended rules

provide a mapping to old ClassAd for the GLUE schema entities “VOView” and “Location”. It should be noted that the Glue Schema v1.3 deprecates the use of the “Location” entity and introduces the “Software” entity, with a similar semantic. In this mapping we do not define rules for the “Software” entity.

These extended rules must be used in addition to the basic rules of section 3.2.

3.3.1 Definitions

The extended rules v1.0 are the following:

- The mapping is built by considering all the possible combinations of inter-related ComputingElement, Cluster, Subcluster, and VOView entities. In other words, each combination contains a single ComputingElement, Cluster, Subcluster, and VOView entity. Attributes in each combination are then mapped to a single old ClassAd. VOView attributes that have the same names as the ones for CE overwrite the CE values (e.g. GlueCEStateFreeJobSlots from the linked VOView, if present, overwrites GlueCEStateFreeJobSlots from CE). The original values appear in the classad as the value of new attribute names, derived from the original attribute names, appending the string “Original” to each. For example, the original value of the attribute GlueCEStateRunningJobs appears in the classad as the value of the new attribute GlueCEStateRunningJobsOriginal. This makes it easier for users to express requirements on resource characteristics that affect the VO/group/role. In fact, the requirement expression does not change depending on whether VOView is present or not.
- After building all ClassAds, it should be checked whether the value of the attribute GlueCEAccessControlBaseRule of the CE entity is the same as one of the GlueCEAccessControlBaseRule attributes in the VOView entities. If this is NOT the case, the CE entity refers to an Access Control Base Rule that is not qualified using a VOView entity. Therefore, in order to advertise this CE-level Access Rule, an extra set of ClassAds should be generated by building all possible combinations of the inter-related ComputingElement, Cluster, Subcluster only (excluding the VOView entities). If this were not done, the GlueCEAccessControlBaseRule attribute of the CE would go lost, as the GlueCEAccessControlBaseRule attributes of VOView always overwrite the one for CE.
- This rule applies for Glue Schema v1.2 only. The location entity appears in old ClassAd as a multivalued attribute. The mapping is built as follows. For a given subcluster entity, we consider all its related Location entities. For a given attribute in the Location entity, we build the list of values in all the related Location entities as a comma-separated string of values. The order of the values must be the same across different attributes, i.e. values from the same entity have the same position on the list. This way, pairs of values from two different Location attributes can be selected using custom (call-out) functions in the Condor Matchmaking Service. Example signatures for such functions could be

```
matchOrderedPair(AttributeName1, AttributeName2,
```

Value1, Value2)

or

```
matchOrderedNTuple(NumberOfAttributes,  
                   AttributeName1, AttributeName2, ... ,  
                   Value1, Value2, ...)
```

In summary, in old ClassAd format, for every subcluster, every attribute in the Location entity appears once with the list of values expressed as a list.

- The attribute Path from the Location entity is not advertised, as it is deemed to be of marginal importance in the resource selection process.
- At least these two attributes from the Site entity are provided: GlueSiteUniqueID and GlueSiteName. In particular, the GlueSiteName is of interest in the selection of resources.

3.3.2 Example

Figure 8, 9, 10, and 11 show an example of a set of two old ClassAds (each ClassAd is split in two figures to fit on the page). The set is derived applying the “extended rules” (section 3.3) to the LDIF representation of the example GLUE schema in figures 1, 2, 3, 4, and 5.

The ClassAd are identical to the ones generated with the basic rules, with the addition of Location attributes. VOView attributes have overwritten CE attributes in the case of the cdf VO. The GlueLocationPath attribute of the Location entity is not included in the mapping. The examples are based on the Glue Schema v1.2, but the extended rules apply also to v1.3.

3.4 Extended Rules v1.1

The extended rules v1.1 provide LDIF to old ClassAd mappings for the “SE” (Storage Element) and “Software” entities. These rules apply for Glue Schema v1.3 only.

These extended rules must be used in addition to the basic rules and the extended rules v1.0 (sections 3.2 and 3.3).

3.4.1 Definitions

The extended rules v1.1 define a mapping for SE and Software entities into a set of old ClassAds. They then provide rules to append such ClassAds to the ClassAds created for Cluster, CE, SubCluster, and VO View entities (basic and extended v1.0 rules).

- The mapping for Storage Elements is built by considering all the possible combinations of inter-related StorageElement entities (SE), StorageAreas (SA) entities, and VO supported by each SA (as listed in the GlueSAAccessControlBaseRule attribute for each SA). In other words, each combination contains a single StorageElement, StorageArea, and VO. Attributes in each combination are then mapped to a single old ClassAd.

- According to the Glue Schema v1.3, for each SA, one can define a VOInfo entity, to redefine particular attributes (e.g. Path) for a restricted list of VOs. The list of VOs is defined in the GlueVOInfoAccessControlBaseRule attribute list. If VOInfo is defined, each classad created for the SA / VO combination is augmented with the attributes in the VOInfo for those VOs. For example, if VOInfo is defined for VO X, the classad for the corresponding SA and VO X will have 2 Path attributes (GlueSAPath and GlueVOInfoPath).
- If multiple ControlProtocol and AccessProtocol entities are present, their attributes are expressed in each classad as the comma-separated list of their values.
- Information from the classads generated for the CEs (from Cluster, Computing Element, SubCluster, VOView entities, using basic rules and extended rules v1.0) and for the SEs (from Storage Elements, Storage Areas, Control Protocol, Access Protocol, VOInfo entities, using extended rules v1.1 above) are then combined into single classads. Classads are linked via the CESEBind entity. Attributes in the CESEBind entity are added to the final classad. If no binding is defined for an SE, only CE classads are generated.
- For every VO (from GlueCEAccessControlBaseRule) in the classads generated for the CEs, find the classad for the corresponding VO (from GlueSAAccessControlBaseRule) in the set of classads generated for the SEs. These two classads are appended to one another. Note: if CE supports a VO (from GlueCEAccessControlBaseRule) AND SE does not support it for any SA (from the list in GlueSAAccessControlBaseRule), no classad is generated i.e. the information for that VO and CE is NOT propagated. Also, if SE supports a VO AND no CE supports it, no classad is generated.
- Software entities are mapped as follows: if more than one Software entity is linked to a SubCluster entity, each attribute of the Software entity is mapped as a comma-separated list of the values of that attribute from all the Software entities. The order of the list is undefined, but values from the same software entity will have the same index in the respective attribute lists. For example, 2 Software entities are defined with LocalID=SW1 and LocalID=SW2; the classad will have attributes like GlueSoftwareVersion = "SW2.Version,SW1.Version", GlueSoftwareInstalledRoot = "SW2.InstalledRoot,SW1.InstalledRoot", and so on.
- In the Glue Schema v1.3, the SoftwareData entity can extend the Software entity with additional attribute names (i.e. the value of the Key attribute) and values (i.e. the value of the Value attribute). The information in the SoftwareData entities linked to a Software entity (e.g. with localID=SWx) will be represented as a string as "SW1.Key1-value=SW1.Value1-value; SW1.Key2-value=SW1.Value2-value; SW1.Key3-value=SW1.Value3-value;...", where SW1.Key1-value is the value of the Key attribute of the first SoftwareData entity linked to the SW1 Software entity, SW1.Value1-value is the value of the Value attribute of the first SoftwareData entity linked to the SW1 Software entity, and so on. This string is propagated in the

classad as the value of the attribute GlueSoftwareData. Since multiple Software entities can be defined for a given SubCluster, GlueSoftwareData is the comma-separated list of all the SoftwareData information for all Software Entities, as discussed in the mapping rule above.

References

- [1] <http://glueschema.forge.cnaf.infn.it/>
- [2] <http://glueschema.forge.cnaf.infn.it/Mapping/LDAP>
- [3] <http://glueschema.forge.cnaf.infn.it/Mapping/XMLSchema>
- [4] <http://grid.pd.infn.it/cemon/>
- [5] L. Field. Generic Information Provider. <http://lfield.home.cern.ch/lfield/gip>
- [6] G. Good. The LDAP Data Interchange Format (LDIF) - Technical Specification. IETF RFC 2849. June 2000.

```

# Computing Element Entity (1)
dn: GlueCEUniqueID=rsgri3.its.uiowa.edu:2119/jobmanager-condor-cdf,mds-vo-name=local,o=grid
objectClass: GlueCETop
objectClass: GlueCE
objectClass: GlueSchemaVersion
objectClass: GlueCEAccessControlBase
objectClass: GlueCEInfo
objectClass: GlueCEPolicy
objectClass: GlueCEState
objectClass: GlueInformationService
objectClass: GlueKey
GlueCEHostingCluster: rsgri3.its.uiowa.edu
GlueCEName: cdf
GlueCEUniqueID: rsgri3.its.uiowa.edu:2119/jobmanager-condor-cdf
GlueCEInfoGatekeeperPort: 2119
GlueCEInfoHostName: rsgri3.its.uiowa.edu
GlueCEInfoLRMSType: condor
GlueCEInfoLRMSVersion: 6.7.13
GlueCEInfoTotalCPUs: 1
GlueCEInfoJobManager: condor
GlueCEInfoContactString: rsgri3.its.uiowa.edu:2119/jobmanager-condor
GlueCEInfoApplicationDir: /export/grid/app
GlueCEInfoDataDir: /export/grid/data
GlueCEInfoDefaultSE: rsgri3.its.uiowa.edu
GlueCEStateEstimatedResponseTime: 0
GlueCEStateFreeCPUs: 0
GlueCEStateRunningJobs: 0
GlueCEStateStatus: Production
GlueCEStateTotalJobs: 0
GlueCEStateWaitingJobs: 0
GlueCEStateWorstResponseTime: 0
GlueCEStateFreeJobSlots: 1
GlueCEPolicyMaxCPUTime: 0
GlueCEPolicyMaxRunningJobs: 0
GlueCEPolicyMaxTotalJobs: 0
GlueCEPolicyMaxWallClockTime: 0
GlueCEPolicyPriority: 0
GlueCEPolicyAssignedJobSlots: 1
GlueCEAccessControlBaseRule: VO:cdf
GlueForeignKey: GlueClusterUniqueID=rsgri3.its.uiowa.edu
GlueInformationServiceURL: ldap://rsgri3.its.uiowa.edu:2135/mds-vo-name=local,o=grid
GlueSchemaVersionMajor: 1
GlueSchemaVersionMinor: 2

```

Figure 1: An LDIF representation of a Computing Element entity of the GLUE schema v1.2

```

# Computing Element Entity (2)
dn: GlueCEUniqueID=rsgri3.its.uiowa.edu:2119/jobmanager-condor-fermilab,mds-vo-name=local,o=grid
objectClass: GlueCETop
objectClass: GlueCE
objectClass: GlueSchemaVersion
objectClass: GlueCEAccessControlBase
objectClass: GlueCEInfo
objectClass: GlueCEPolicy
objectClass: GlueCEState
objectClass: GlueInformationService
objectClass: GlueKey
GlueCEHostingCluster: rsgri3.its.uiowa.edu
GlueCEName: fermilab
GlueCEUniqueID: rsgri3.its.uiowa.edu:2119/jobmanager-condor-fermilab
GlueCEInfoGatekeeperPort: 2119
GlueCEInfoHostName: rsgri3.its.uiowa.edu
GlueCEInfoLRMSType: condor
GlueCEInfoLRMSVersion: 6.7.13
GlueCEInfoTotalCPUs: 1
GlueCEInfoJobManager: condor
GlueCEInfoContactString: rsgri3.its.uiowa.edu:2119/jobmanager
GlueCEInfoApplicationDir: /export/grid/app
GlueCEInfoDataDir: /export/grid/data
GlueCEInfoDefaultSE: rsgri3.its.uiowa.edu
GlueCEStateEstimatedResponseTime: 0
GlueCEStateFreeCPUs: 0
GlueCEStateRunningJobs: 0
GlueCEStateStatus: Production
GlueCEStateTotalJobs: 0
GlueCEStateWaitingJobs: 0
GlueCEStateWorstResponseTime: 0
GlueCEStateFreeJobSlots: 1
GlueCEPolicyMaxCPUTime: 0
GlueCEPolicyMaxRunningJobs: 0
GlueCEPolicyMaxTotalJobs: 0
GlueCEPolicyMaxWallClockTime: 0
GlueCEPolicyPriority: 0
GlueCEPolicyAssignedJobSlots: 1
GlueCEAccessControlBaseRule: VO:fermilab
GlueForeignKey: GlueClusterUniqueID=rsgri3.its.uiowa.edu
GlueInformationServiceURL: ldap://rsgri3.its.uiowa.edu:2135/mds-vo-name=local,o=grid
GlueSchemaVersionMajor: 1
GlueSchemaVersionMinor: 2

```

Figure 2: An LDIF representation of a Computing Element entity of the GLUE schema v1.2

```

# Cluster Entity
dn: GlueClusterUniqueID=rsgrid3.its.uiowa.edu,mds-vo-name=local,o=grid
objectClass: GlueClusterTop
objectClass: GlueCluster
objectClass: GlueSchemaVersion
objectClass: GlueInformationService
objectClass: GlueKey
GlueClusterName: rsgrid3.its.uiowa.edu
GlueClusterService: rsgrid3.its.uiowa.edu
GlueClusterUniqueID: rsgrid3.its.uiowa.edu
GlueForeignKey: GlueCEUniqueID=rsgrid3.its.uiowa.edu:2119/jobmanager-condor-cdf
GlueForeignKey: GlueCEUniqueID=rsgrid3.its.uiowa.edu:2119/jobmanager-condor-fermilab
GlueForeignKey: GlueSiteUniqueID=rsgrid3.its.uiowa.edu
GlueInformationServiceURL: ldap://rsgrid3.its.uiowa.edu:2135/mds-vo-name=local,o=grid
GlueSchemaVersionMajor: 1
GlueSchemaVersionMinor: 2

# SubCluster Entity
dn: GlueSubClusterUniqueID=rsgrid3.its.uiowa.edu, GlueClusterUniqueID=rsgrid3.its.uiowa.edu,
   mds-vo-name=local,o=grid
objectClass: GlueClusterTop
objectClass: GlueSubCluster
objectClass: GlueSchemaVersion
objectClass: GlueHostApplicationSoftware
objectClass: GlueHostArchitecture
objectClass: GlueHostBenchmark
objectClass: GlueHostMainMemory
objectClass: GlueHostNetworkAdapter
objectClass: GlueHostOperatingSystem
objectClass: GlueHostProcessor
objectClass: GlueInformationService
objectClass: GlueKey
GlueChunkKey: GlueClusterUniqueID=rsgrid3.its.uiowa.edu
GlueHostApplicationSoftwareRunTimeEnvironment: OSG-0.3.6
GlueHostApplicationSoftwareRunTimeEnvironment: ATLAS_L0C_903
GlueHostApplicationSoftwareRunTimeEnvironment: LCG-2_6_0
GlueHostArchitectureSMPSize: 2
GlueHostBenchmarkSF00: 380
GlueHostBenchmarkSIO0: 400
GlueHostMainMemoryRAMSize: 512
GlueHostMainMemoryVirtualSize: 1024
GlueHostNetworkAdapterInboundIP: FALSE
GlueHostNetworkAdapterOutboundIP: TRUE
GlueHostOperatingSystemName: linux-rocks-3.3
GlueHostOperatingSystemRelease: Rocks Linux
GlueHostOperatingSystemVersion: 3.3
GlueHostProcessorClockSpeed: 1000
GlueHostProcessorModel: Pentium III (Coppermine)
GlueHostProcessorVendor: GenuineIntel
GlueSubClusterName: rsgrid3.its.uiowa.edu
GlueSubClusterUniqueID: rsgrid3.its.uiowa.edu
GlueSubClusterPhysicalCPUs: 1
GlueSubClusterLogicalCPUs: 2
GlueSubClusterTmpDir: /export/grid/data
GlueSubClusterWNTmpDir: /tmp
GlueInformationServiceURL: ldap://rsgrid3.its.uiowa.edu:2135/mds-vo-name=local,o=grid
GlueSchemaVersionMajor: 1
GlueSchemaVersionMinor: 2

```

Figure 3: An LDIF representation of a Cluster and SubCluster entities of the GLUE schema v1.2

```

# VOView Entity: Ignored in the mapping with the basic rules
dn: GlueVOViewLocalID=cdf,GlueCEUniqueID=rsgrid3.its.uiowa.edu:2119/
  jobmanager-condor-cdf,mds-vo-name=local,o=grid
objectClass: GlueCETop
objectClass: GlueVOView
objectClass: GlueCEInfo
objectClass: GlueCEState
objectClass: GlueCEAccessControlBase
objectClass: GlueCEPolicy
objectClass: GlueKey
objectClass: GlueSchemaVersion
GlueVOViewLocalID: cdf
GlueCEAccessControlBaseRule: VO:cdf
GlueCEStateRunningJobs: 0
GlueCEStateWaitingJobs: 0
GlueCEStateTotalJobs: 0
GlueCEStateFreeJobSlots: 1
GlueCEStateEstimatedResponseTime: 0
GlueCEStateWorstResponseTime: 0
GlueCEInfoDefaultSE: rsgrid3.its.uiowa.edu
GlueCEInfoApplicationDir: /export/grid/app
GlueCEInfoDataDir: /export/grid/data
GlueChunkKey: GlueCEUniqueID=rsgrid3.its.uiowa.edu:2119/jobmanager-condor-cdf
GlueSchemaVersionMajor: 1
GlueSchemaVersionMinor: 2

# Site Entity: Ignored in the mapping
dn: GlueSiteUniqueID=rsgrid3.its.uiowa.edu,mds-vo-name=local,o=grid
objectClass: GlueTop
objectClass: GlueSite
objectClass: GlueKey
objectClass: GlueSchemaVersion
GlueSiteUniqueID: rsgrid3.its.uiowa.edu
GlueSiteName: UIOWA-ITB
GlueSiteDescription: OSG Site
GlueSiteUserSupportContact: mailto: shaowen-wang@uiowa.edu
GlueSiteSysAdminContact: mailto: shaowen-wang@uiowa.edu
GlueSiteSecurityContact: mailto: shaowen-wang@uiowa.edu
GlueSiteLocation: Iowa City , US
GlueSiteLatitude: 41.67
GlueSiteLongitude: -91.55
GlueSiteWeb: http://rsgrid3.its.uiowa.edu
GlueForeignKey: GlueSiteUniqueID=rsgrid3.its.uiowa.edu
GlueSchemaVersionMajor: 1
GlueSchemaVersionMinor: 2

```

Figure 4: An LDIF representation of a VOView and Site entities of the GLUE schema v1.2. A mapping for VOView and Site is not provided by the basic rules but is provided by the extended rules.

```

# Location Entity: Ignored in the mapping with the basic rules
# OSG_SITE_READ, rsgrid3.its.uiowa.edu, rsgrid3.its.uiowa.edu, local, grid
dn: GlueLocationLocalID=OSG_SITE_READ,
    GlueSubClusterUniqueID=rsgrid3.its.uiowa.edu,
    GlueClusterUniqueID=rsgrid3.its.uiowa.edu,mds-vo-name=local,o=grid
objectClass: GlueClusterTop
objectClass: GlueLocation
objectClass: GlueKey
objectClass: GlueSchemaVersion
GlueLocationLocalID: OSG_SITE_READ
GlueLocationName: OSG_SITE_READ
GlueLocationPath: /export/grid/data
GlueChunkKey: GlueClusterUniqueID=rsgrid3.its.uiowa.edu
GlueSchemaVersionMajor: 1
GlueSchemaVersionMinor: 2

# OSG_SITE_WRITE, rsgrid3.its.uiowa.edu, rsgrid3.its.uiowa.edu, local, grid
dn: GlueLocationLocalID=OSG_SITE_WRITE,
    GlueSubClusterUniqueID=rsgrid3.its.uiowa.edu,
    GlueClusterUniqueID=rsgrid3.its.uiowa.edu,mds-vo-name=local,o=grid
objectClass: GlueClusterTop
objectClass: GlueLocation
objectClass: GlueKey
objectClass: GlueSchemaVersion
GlueLocationLocalID: OSG_SITE_WRITE
GlueLocationName: OSG_SITE_WRITE
GlueLocationPath: /export/grid/data
GlueChunkKey: GlueClusterUniqueID=rsgrid3.its.uiowa.edu
GlueSchemaVersionMajor: 1
GlueSchemaVersionMinor: 2

```

Figure 5: An LDIF representation of Location entities of the GLUE schema v1.2. The basic rules do not provide a mapping to old ClassAd for these entities, but the extended rules do (see Fig.9).

```

### Computing Element Entity (1)
GlueCEHostingCluster = "rsgrid3.its.uiowa.edu"
GlueCEName = "cdf"
GlueCEUniqueID = "rsgrid3.its.uiowa.edu:2119/jobmanager-condor-cdf"
GlueCEInfoGatekeeperPort = 2119
GlueCEInfoHostName = "rsgrid3.its.uiowa.edu"
GlueCEInfoLRMSType = "condor"
GlueCEInfoLRMSVersion = "6.7.13"
GlueCEInfoTotalCPUs = 1
GlueCEInfoJobManager = "condor"
GlueCEInfoContactString = "rsgrid3.its.uiowa.edu:2119/jobmanager-condor"
GlueCEInfoApplicationDir = "/export/grid/app"
GlueCEInfoDataDir = "/export/grid/data"
GlueCEInfoDefaultSE = "rsgrid3.its.uiowa.edu"
GlueCEInfoGRAMVersion = "1.7"
GlueCEStateEstimatedResponseTime = 0
GlueCEStateFreeCPUs = 0
GlueCEStateRunningJobs = 0
GlueCEStateStatus = "Production"
GlueCEStateTotalJobs = 0
GlueCEStateWaitingJobs = 0
GlueCEStateWorstResponseTime = 0
GlueCEStateFreeJobSlots = 1
GlueCEPolicyMaxCPUTime = 0
GlueCEPolicyMaxRunningJobs = 0
GlueCEPolicyMaxTotalJobs = 0
GlueCEPolicyMaxWallClockTime = 0
GlueCEPolicyPriority = 0
GlueCEPolicyAssignedJobSlots = 1
GlueCEAccessControlBaseRule = "VO:cdf"
GlueInformationServiceURL = "ldap://rsgrid3.its.uiowa.edu:2135/mds-vo-name=local,o=grid"
GlueSchemaVersionMajor: 1
GlueSchemaVersionMinor: 2
### Cluster Entity
GlueClusterName = "rsgrid3.its.uiowa.edu"
GlueClusterService = "rsgrid3.its.uiowa.edu"
GlueClusterUniqueID = "rsgrid3.its.uiowa.edu"
GlueClusterTmpDir = "/tmp"
GlueClusterWNTmpDir = "/scratch"
### SubCluster Entity
GlueHostApplicationSoftwareRunTimeEnvironment = "OSG-0.3.6,ATLAS_LOC_903,LCG-2_6_0"
GlueHostArchitectureSMPSize = 2
GlueHostArchitecturePlatformType = undefined
GlueHostBenchmarkSF00 = 380
GlueHostBenchmarkSI00 = 400
GlueHostMainMemoryRAMSize = 512
GlueHostMainMemoryVirtualSize = 1024
GlueHostNetworkAdapterInboundIP = FALSE
GlueHostNetworkAdapterOutboundIP = TRUE
GlueHostOperatingSystemName = "linux-rocks-3.3"
GlueHostOperatingSystemRelease = "Rocks Linux"
GlueHostOperatingSystemVersion = 3.3
GlueHostProcessorClockSpeed = 1000
GlueHostProcessorModel = "Pentium III (Coppermine)"
GlueHostProcessorVendor = "GenuineIntel"
GlueHostProcessorVersion = undefined
GlueHostProcessorInstructionSet = "i686"
GlueHostProcessorOtherDescription = "i686"
GlueSubClusterName = "rsgrid3.its.uiowa.edu"
GlueSubClusterUniqueID = "rsgrid3.its.uiowa.edu"
GlueSubClusterPhysicalCPUs = 1
GlueSubClusterLogicalCPUs = 2
GlueSubClusterTmpDir = "/export/grid/data"
GlueSubClusterWNTmpDir = "/tmp"

```

Figure 6: The first of two old ClassAds representing the GLUE Schema presented before in LDIF format. This ClassAd was generated using the “basic rules” for the mapping. The 3 pound signs (###) are used as comments only and are not part of the ClassAd.

```

### Computing Element Entity (2)
GlueCEHostingCluster = "rsgrid3.its.uiowa.edu"
GlueCEName = "fermilab"
GlueCEUniqueID = "rsgrid3.its.uiowa.edu:2119/jobmanager-condor-fermilab"
GlueCEInfoGatekeeperPort = 2119
GlueCEInfoHostName = "rsgrid3.its.uiowa.edu"
GlueCEInfoLRMSType = "condor"
GlueCEInfoLRMSVersion = "6.7.13"
GlueCEInfoTotalCPUs = 1
GlueCEInfoJobManager = "condor"
GlueCEInfoContactString = "rsgrid3.its.uiowa.edu:2119/jobmanager-condor"
GlueCEInfoApplicationDir = "/export/grid/app"
GlueCEInfoDataDir = "/export/grid/data"
GlueCEInfoDefaultSE = "rsgrid3.its.uiowa.edu"
GlueCEInfoGRAMVersion = "1.7"
GlueCEStateEstimatedResponseTime = 0
GlueCEStateFreeCPUs = 0
GlueCEStateRunningJobs = 0
GlueCEStateStatus = "Production"
GlueCEStateTotalJobs = 0
GlueCEStateWaitingJobs = 0
GlueCEStateWorstResponseTime = 0
GlueCEStateFreeJobSlots = 1
GlueCEPolicyMaxCPUTime = 0
GlueCEPolicyMaxRunningJobs = 0
GlueCEPolicyMaxTotalJobs = 0
GlueCEPolicyMaxWallClockTime = 0
GlueCEPolicyPriority = 0
GlueCEPolicyAssignedJobSlots = 1
GlueCEAccessControlBaseRule = "VO:fermilab"
GlueInformationServiceURL = "ldap://rsgrid3.its.uiowa.edu:2135/mds-vo-name=local,o=grid"
GlueSchemaVersionMajor = 1
GlueSchemaVersionMinor = 2
### Cluster Entity
GlueClusterName = "rsgrid3.its.uiowa.edu"
GlueClusterService = "rsgrid3.its.uiowa.edu"
GlueClusterUniqueID = "rsgrid3.its.uiowa.edu"
GlueClusterTmpDir = "/tmp"
GlueClusterWNTmpDir = "/scratch"
### SubCluster Entity
GlueHostApplicationSoftwareRunTimeEnvironment = "OSG-0.3.6,ATLAS_LOC_903,LCG-2_6_0"
GlueHostArchitectureSMPSize = 2
GlueHostArchitecturePlatformType = undefined
GlueHostBenchmarkSF00 = 380
GlueHostBenchmarkSI00 = 400
GlueHostMainMemoryRAMSize = 512
GlueHostMainMemoryVirtualSize = 1024
GlueHostNetworkAdapterInboundIP = FALSE
GlueHostNetworkAdapterOutboundIP = TRUE
GlueHostOperatingSystemName = "linux-rocks-3.3"
GlueHostOperatingSystemRelease = "Rocks Linux"
GlueHostOperatingSystemVersion = 3.3
GlueHostProcessorClockSpeed = 1000
GlueHostProcessorModel = "Pentium III (Coppermine)"
GlueHostProcessorVendor = "GenuineIntel"
GlueHostProcessorVersion = undefined
GlueHostProcessorInstructionSet = "i686"
GlueHostProcessorOtherDescription = "i686"
GlueSubClusterName = "rsgrid3.its.uiowa.edu"
GlueSubClusterUniqueID = "rsgrid3.its.uiowa.edu"
GlueSubClusterPhysicalCPUs = 1
GlueSubClusterLogicalCPUs = 2
GlueSubClusterTmpDir = "/export/grid/data"
GlueSubClusterWNTmpDir = "/tmp"

```

Figure 7: The second of two old ClassAds representing the GLUE Schema presented before in LDIF format. This ClassAd was generated using the “basic rules” for the mapping. The 3 pound signs (###) are used as comments only and are not part of the ClassAd.

```

### Computing Element Entity (1)
GlueCEHostingCluster = "rsgrid3.its.uiowa.edu"
GlueCEName = "cdf"
GlueCEUniqueID = "rsgrid3.its.uiowa.edu:2119/jobmanager-condor-cdf"
GlueCEInfoGatekeeperPort = 2119
GlueCEInfoHostName = "rsgrid3.its.uiowa.edu"
GlueCEInfoLRMSType = "condor"
GlueCEInfoLRMSVersion = "6.7.13"
GlueCEInfoTotalCPUs = 1
GlueCEInfoJobManager = "condor"
GlueCEInfoContactString = "rsgrid3.its.uiowa.edu:2119/jobmanager-condor"
GlueCEInfoGRAMVersion = "1.7"
GlueCEStateFreeCPUs = 0
GlueCEStateStatus = "Production"
GlueCEPolicyMaxCPUTime = 0
GlueCEPolicyMaxRunningJobs = 0
GlueCEPolicyMaxTotalJobs = 0
GlueCEPolicyMaxWallClockTime = 0
GlueCEPolicyPriority = 0
GlueCEPolicyAssignedJobSlots = 1
GlueInformationServiceURL = "ldap://rsgrid3.its.uiowa.edu:2135/mds-vo-name=local,o=grid"
GlueSchemaVersionMajor: 1
GlueSchemaVersionMinor: 2
### Attributes overwritten by VOView Entity
GlueVOViewLocalID: "cdf"
GlueCEAccessControlBaseRule = "VO:cdf"
GlueCEStateRunningJobs = 0
GlueCEStateWaitingJobs = 0
GlueCEStateTotalJobs = 0
GlueCEStateFreeJobSlots = 1
GlueCEStateEstimatedResponseTime = 0
GlueCEStateWorstResponseTime = 0
GlueCEInfoDefaultSE = "rsgrid3.its.uiowa.edu"
GlueCEInfoApplicationDir = "/export/grid/app"
GlueCEInfoDataDir = "/export/grid/data"

```

Figure 8: The first half of the first of two old ClassAds representing the GLUE Schema presented before in LDIF format. This ClassAd was generated using the “extended rules” for the mapping. The 3 pound signs (###) are used as comments only and are not part of the ClassAd.

```

### Cluster Entity
GlueClusterName = "rsgrid3.its.uiowa.edu"
GlueClusterService = "rsgrid3.its.uiowa.edu"
GlueClusterUniqueID = "rsgrid3.its.uiowa.edu"
GlueClusterTmpDir = "/tmp"
GlueClusterWNTmpDir = "/scratch"
### SubCluster Entity
GlueHostApplicationSoftwareRunTimeEnvironment = "OSG-0.3.6,ATLAS_LOC_903,LCG-2_6_0"
GlueHostArchitectureSMPSize = 2
GlueHostArchitecturePlatformType = undefined
GlueHostBenchmarkSF00 = 380
GlueHostBenchmarkSI00 = 400
GlueHostMainMemoryRAMSize = 512
GlueHostMainMemoryVirtualSize = 1024
GlueHostNetworkAdapterInboundIP = FALSE
GlueHostNetworkAdapterOutboundIP = TRUE
GlueHostOperatingSystemName = "linux-rocks-3.3"
GlueHostOperatingSystemRelease = "Rocks Linux"
GlueHostOperatingSystemVersion = 3.3
GlueHostProcessorClockSpeed = 1000
GlueHostProcessorModel = "Pentium III (Coppermine)"
GlueHostProcessorVendor = "GenuineIntel"
GlueHostProcessorVersion = undefined
GlueHostProcessorInstructionSet = "i686"
GlueHostProcessorOtherDescription = "i686"
GlueSubClusterName = "rsgrid3.its.uiowa.edu"
GlueSubClusterUniqueID = "rsgrid3.its.uiowa.edu"
GlueSubClusterPhysicalCPUs = 1
GlueSubClusterLogicalCPUs = 2
GlueSubClusterTmpDir = "/export/grid/data"
GlueSubClusterWNTmpDir = "/tmp"
### Location Entity
GlueLocationLocalID = "OSG_SITE_READ,OSG_SITE_WRITE"
GlueLocationName = "OSG_SITE_READ,OSG_SITE_WRITE"
### Site Entity: at least these two attribute must be present
GlueSiteUniqueID: rsgrid3.its.uiowa.edu
GlueSiteName: UIOWA-ITB

```

Figure 9: The second half of the ClassAd in figure 8

```

### Computing Element Entity (2)
GlueCEHostingCluster = "rsgrid3.its.uiowa.edu"
GlueCEName = "fermilab"
GlueCEUniqueID = "rsgrid3.its.uiowa.edu:2119/jobmanager-condor-fermilab"
GlueCEInfoGatekeeperPort = 2119
GlueCEInfoHostName = "rsgrid3.its.uiowa.edu"
GlueCEInfoLRMSType = "condor"
GlueCEInfoLRMSVersion = "6.7.13"
GlueCEInfoTotalCPUs = 1
GlueCEInfoJobManager = "condor"
GlueCEInfoContactString = "rsgrid3.its.uiowa.edu:2119/jobmanager-condor"
GlueCEInfoApplicationDir = "/export/grid/app"
GlueCEInfoDataDir = "/export/grid/data"
GlueCEInfoDefaultSE = "rsgrid3.its.uiowa.edu"
GlueCEInfoGRAMVersion = "1.7"
GlueCEStateEstimatedResponseTime = 0
GlueCEStateFreeCPUs = 0
GlueCEStateRunningJobs = 0
GlueCEStateStatus = "Production"
GlueCEStateTotalJobs = 0
GlueCEStateWaitingJobs = 0
GlueCEStateWorstResponseTime = 0
GlueCEStateFreeJobSlots = 1
GlueCEPolicyMaxCPUTime = 0
GlueCEPolicyMaxRunningJobs = 0
GlueCEPolicyMaxTotalJobs = 0
GlueCEPolicyMaxWallClockTime = 0
GlueCEPolicyPriority = 0
GlueCEPolicyAssignedJobSlots = 1
GlueCEAccessControlBaseRule = "VO:fermilab"
GlueInformationServiceURL = "ldap://rsgrid3.its.uiowa.edu:2135/mds-vo-name=local,o=grid"
GlueSchemaVersionMajor = 1
GlueSchemaVersionMinor = 2
### VOView Entity for VO:fermilab is not defined: nothing to overwrite.

```

Figure 10: The first half of the second of two old ClassAds representing the GLUE Schema presented before in LDIF format. This ClassAd was generated using the “extended rules” for the mapping. The 3 pound signs (###) are used as comments only and are not part of the ClassAd.

```

### Cluster Entity
GlueClusterName = "rsgrid3.its.uiowa.edu"
GlueClusterService = "rsgrid3.its.uiowa.edu"
GlueClusterUniqueID = "rsgrid3.its.uiowa.edu"
GlueClusterTmpDir = "/tmp"
GlueClusterWNTmpDir = "/scratch"
### SubCluster Entity
GlueHostApplicationSoftwareRunTimeEnvironment = "OSG-0.3.6,ATLAS_LOC_903,LCG-2_6_0"
GlueHostArchitectureSMPSize = 2
GlueHostArchitecturePlatformType = undefined
GlueHostBenchmarkSF00 = 380
GlueHostBenchmarkSI00 = 400
GlueHostMainMemoryRAMSize = 512
GlueHostMainMemoryVirtualSize = 1024
GlueHostNetworkAdapterInboundIP = FALSE
GlueHostNetworkAdapterOutboundIP = TRUE
GlueHostOperatingSystemName = "linux-rocks-3.3"
GlueHostOperatingSystemRelease = "Rocks Linux"
GlueHostOperatingSystemVersion = 3.3
GlueHostProcessorClockSpeed = 1000
GlueHostProcessorModel = "Pentium III (Coppermine)"
GlueHostProcessorVendor = "GenuineIntel"
GlueHostProcessorVersion = undefined
GlueHostProcessorInstructionSet = "i686"
GlueHostProcessorOtherDescription = "i686"
GlueSubClusterName = "rsgrid3.its.uiowa.edu"
GlueSubClusterUniqueID = "rsgrid3.its.uiowa.edu"
GlueSubClusterPhysicalCPUs = 1
GlueSubClusterLogicalCPUs = 2
GlueSubClusterTmpDir = "/export/grid/data"
GlueSubClusterWNTmpDir = "/tmp"
### Location Entity
GlueLocationLocalID = "OSG_SITE_READ,OSG_SITE_WRITE"
GlueLocationName = "OSG_SITE_READ,OSG_SITE_WRITE"
### Site Entity: at least these two attribute must be present
GlueSiteUniqueID: rsgrid3.its.uiowa.edu
GlueSiteName: UIOWA-ITB

```

Figure 11: The second half of the ClassAd in figure 10