

# Discrepancies in GLUE usage seen on LCG systems

Jeff Templon

February 3, 2006

## 1 Introduction

## 2 Discrepancies

### 2.1 GlueSARoot

Recent messages on LCG-ROLLOUT indicate that the lcg-utils package is decoding the SARoot information in order to deduce things about the storage element properties.

First this should go away ASAP since the SARoot attribute is deprecated. Secondly, the recommended syntax for these values (needed in order for lcg-utils to work right) is:

**classic SE** voname:voname

**DPM SE** voname:path

Presumably the syntax for the DPM SE is the same for all SRM SEs. In any case the “voname” should be extracted from the GlueSAAccessControlBaseRule and the “path” should be found in the the GlueSAPath.

It also is not clear why the path is needed for one SE type but not for the other.

### 2.2 GlueSAPath

There are two types of incorrect values observed.

**non-path format** *e.g.*, one of the SEs publishes a value of CMS.

**emulate old, incorrect SARoot** *e.g.*, the following is observed:

biomed:/pnfs/tier2.hep.manchester.ac.uk/data/biomed.

Neither is correct.

### 2.3 GlueHostMainMemoryRAMSize

The schema gives this parameter units of megabytes. I find it unlikely (but not impossible) that all of the following observed values are correct:

```
1 GlueHostMainMemoryRAMSize: 1036836
1 GlueHostMainMemoryRAMSize: 2027952
1 GlueHostMainMemoryRAMSize: 2048128
1 GlueHostMainMemoryRAMSize: 4101216
1 GlueHostMainMemoryRAMSize: 4122160
```

There are similar unlikely values for the `VirtualSize` attribute.

## 2.4 GlueSiteWeb

This value is likely wrong:

```
GlueSiteWeb: h
```

## 2.5 GlueHostOperatingSystemRelease

Obviously the standardizations that were proposed on ROLLOUT have been ignored:

```
1 GlueHostOperatingSystemRelease: 3.0
1 GlueHostOperatingSystemRelease: 3.02
1 GlueHostOperatingSystemRelease: 3.03
1 GlueHostOperatingSystemRelease: 3.2
1 GlueHostOperatingSystemRelease: 3.5
1 GlueHostOperatingSystemRelease: 303
1 GlueHostOperatingSystemRelease: 305
1 GlueHostOperatingSystemRelease: 3_0_4
1 GlueHostOperatingSystemRelease: 5.10
1 GlueHostOperatingSystemRelease: EL3
1 GlueHostOperatingSystemRelease: Sarge
1 GlueHostOperatingSystemRelease: sl3
2 GlueHostOperatingSystemRelease: 3.05
2 GlueHostOperatingSystemRelease: 304
2 GlueHostOperatingSystemRelease: SL
4 GlueHostOperatingSystemRelease: SLC3
5 GlueHostOperatingSystemRelease: 3.04
5 GlueHostOperatingSystemRelease: SL3
7 GlueHostOperatingSystemRelease: 3.0.6
9 GlueHostOperatingSystemRelease: 7.3
16 GlueHostOperatingSystemRelease: 3.0.3
35 GlueHostOperatingSystemRelease: 3.0.5
39 GlueHostOperatingSystemRelease: 3
43 GlueHostOperatingSystemRelease: 3.0.4
```

## 2.6 GlueServiceInformationServiceURL

There are a number of different formats observed in the wild for this attribute. At least some of them are not real URLs. Below is a selection of various types seen.

```
MDS2GRIS:ldap://dcache.zib.de:2135/mds-vo-name=local,o=grid
http://lfc-lhcb-ro.cern.ch:8085/
ldap://epgmo1.ph.bham.ac.uk:2136/mds-vo-name=local,o=grid
ldap://pxy.prd.hp.com:2135/mds-vo-name=local,o=grid
ldap://prod-ce-01.pd.infn.it:2170/mds-vo-name=local,o=grid
MDS2GRIS:ldap://cclcgip01.in2p3.fr:2170/mds-vo-name=local,mds-vo-name=IN2P3-CC,mds-vo
MDS2GRIS:ldap://site-bdii.cern.ch:2170/mds-vo-name=cernlcg2,o=grid
```

The ones prefixed via MDS2GRIS are not URLs.

A random visit to one of the http links above results in XML containing

```
<faultcode>SOAP-ENV:Client</faultcode>
<faultstring>HTTP GET method not implemented</faultstring>
```

## 2.7 GlueSubClusterWNTmpDir

Need to think about whether this is allowed:

```
GlueSubClusterWNTmpDir: managed by condor
```

## 2.8 GlueSESizeTotal

This attribute should have units of gigabytes according to the 1.2 specification. I suspect that many of the entries seen below are not correct:

```
1 GlueSESizeTotal: 1580000000
1 GlueSESizeTotal: 15920000000
1 GlueSESizeTotal: 2528520000
1 GlueSESizeTotal: 2750000000
1 GlueSESizeTotal: 3150000000
1 GlueSESizeTotal: 999999999
2 GlueSESizeTotal: 00
2 GlueSESizeTotal: 100000000000
8 GlueSESizeTotal: 1000000000000
174 GlueSESizeTotal: 0
```

The fact that 174 are seen with size 0 probably means that this is what YAIM is setting.

## 2.9 Mail addresses

Mail addresses are seen with two different formats; we should either decide on a single format, or we should mandate that all software needs to be able to handle both.

```
1 GlueServicePrimaryOwnerContact: hengner@mail.sdu.edu.cn
1 GlueServicePrimaryOwnerContact: mailto:hep-service-sc-level2@cern.ch
```

The difference here, in case you didn't spot it, is the `mailto:` prefix.

## 2.10 InformationServiceURL

The SE block has stuff that looks like this:

```
GlueSESizeFree: 2870
GlueSEArchitecture: srm_v1
GlueInformationServiceURL: ldap://se2.itep.ru:2135/mds-vo-name=local,o=grid
```

Note that all the others begin with GlueSE but not the InformationServiceURL. Looking at the schema definition; InformationServiceURL is defined as an attribute in exactly the same way as the others, it is not at all obvious that it should have Glue instead of GlueSE as a prefix. Apparently what is published is correct, so the schema definition needs to be clarified.

## 2.11 GlueSAPolicyMaxNumFiles

138 separate Glue attributes seen

For your viewing pleasure: GlueSAPolicyMaxNumFiles

```
freq      value
   6  GlueSAPolicyMaxNumFiles: 99999999999999
   9  GlueSAPolicyMaxNumFiles: 999999
  63  GlueSAPolicyMaxNumFiles: 0
  98  GlueSAPolicyMaxNumFiles: 00
1499  GlueSAPolicyMaxNumFiles: 10
```

Firstly the definition of this attribute is “the maximum number of files that a single job is allowed to write to this SA”. I don’t think there is any SE that has a limit on this, and I don’t think there is any software that limits this either. I suggest that sites should cease publishing this number unless there really is a limit.

## 2.12 GlueCESEBindAccessPoint

This should be dropped unless the center is providing “file” access to the SE from WNs. I don’t even know if this is supported in the current software stack.

## 2.13 GlueCEInfoLRMS

The schema definition says there are two associated parameters: “type” and “version”. For “type” we see:

```
   6  GlueCEInfoLRMSType: LSF
   8  GlueCEInfoLRMSType: sge
   9  GlueCEInfoLRMSType: bqs
  23  GlueCEInfoLRMSType: maui
  36  GlueCEInfoLRMSType: pbspro
  47  GlueCEInfoLRMSType: lsf
  57  GlueCEInfoLRMSType: lcgpbs
 139  GlueCEInfoLRMSType: condor
 272  GlueCEInfoLRMSType: pbs
 681  GlueCEInfoLRMSType: torque
```

Problems:

- LSF vs. lsf
- Maui is not an LRMS, it's a scheduler
- lcgpbs is not an LRMS, it's a job manager

For version, several different types are seen:

```
1  GlueCEInfoLRMSVersion: 6.6.8
1  GlueCEInfoLRMSVersion: torque_1.1.0p5
5  GlueCEInfoLRMSVersion: 2.0.0p3-snap.1132842037
16 GlueCEInfoLRMSVersion: @GlueCEInfoLRMSVersion@
27 GlueCEInfoLRMSVersion: not defined
36 GlueCEInfoLRMSVersion: PBSPro_5.4.2.43350
```

'not defined' and the macro insertion are clearly wrong. Otherwise I think including the product name (e.g. torque or PBSPro) is also wrong, that should be the "type".

## 2.14 GlueCEPolicyMaxTotalJobs

The most commonly seen value for this attribute is zero (0). This may mean "no limit" but if so this should be documented so software can rely on it.

## 2.15 GlueSEName

Most values look like NIKHEF-ELPROD:disk now. Apparently software is using the part after the colon to decide on what control protocol is being used. This is wrong. Suggest changing this to a free string ASAP.

## 2.16 GlueSEAccessProtocolEndpoint

Of all the SEs on the Grid, only four are publishing values for this attribute.

## 2.17 GlueCEInfoDefaultSE

There are a couple of instances that look like this, indicating a configuration problem.

```
GlueCEInfoDefaultSE: GlueCEInfoApplicationDir: /afs/in2p3.fr/group/cdf/s13
```

## 2.18 GlueSEArchitecture

Currently some SEs are publishing values of `castor`, `srm`, and `srm_v1`. These are not architectures, these are control protocols. Proper values of `Architecture` are `tape`, `disk`, `multidisk`, and `other`.

## 2.19 GlueSEAccessProtocolType

Some sites are publishing `srm` here which is a control protocol type.

## 2.20 GlueHostProcessorModel

There are a wide variety of conventions observed, that is several distinct forms of the same processor type. Is anyone using this information? If so, we should consider a standard set of Models.

freq	value
1	GlueHostProcessorModel: @PIII@
1	GlueHostProcessorModel: ATHLON2000+
1	GlueHostProcessorModel: AthlonMP2000
1	GlueHostProcessorModel: Elonex_2800
1	GlueHostProcessorModel: Itanium 2
1	GlueHostProcessorModel: Pentium III (Coppermine)
1	GlueHostProcessorModel: p3
1	GlueHostProcessorModel: pentiumIII
2	GlueHostProcessorModel: Athlon
2	GlueHostProcessorModel: DUAL-XEON
2	GlueHostProcessorModel: Itanium2
2	GlueHostProcessorModel: opteron
3	GlueHostProcessorModel: IA64
3	GlueHostProcessorModel: xeon
11	GlueHostProcessorModel: XEON
22	GlueHostProcessorModel: P4
25	GlueHostProcessorModel: Opteron
28	GlueHostProcessorModel: Xeon
33	GlueHostProcessorModel: PIV
45	GlueHostProcessorModel: PIII

## 2.21 GlueCEPolicyMaxRunningJobs

The most frequently occurring value is 0. The schema does not give this value special status. I expect the idea is that 0 means “unlimited” which is also rather unrealistic. This value should be checked to contain reasonable information.

## 2.22 GlueSAPolicyMaxFileSize

The values published here are in general not realistic. The unit is bytes, so we see values of 100 Tbyte, 4 Gbyte (this one is probably correct), 100 kbyte, 2 Mbyte, and 10 kbyte, the last one being the most popular!

6	GlueSAPolicyMaxFileSize: 999999999999999
28	GlueSAPolicyMaxFileSize: 4000000000
30	GlueSAPolicyMaxFileSize: 100000
121	GlueSAPolicyMaxFileSize: 2000000
1622	GlueSAPolicyMaxFileSize: 10000