



Training on **GLUE 2** **information validation**

EGI Technical Forum September


Maria Alandes ²⁰¹³ Pradillo, CERN



Overview



- Part one
 - Introduction
 - glue-validator libraries
 - Command line options
 - Output formats
 - Error messages
- Part two
 - Current status of GLUE 2 validation
 - Future validation process
 - Long term goals
- Part three
 - Most common GLUE 2 errors and how to fix them

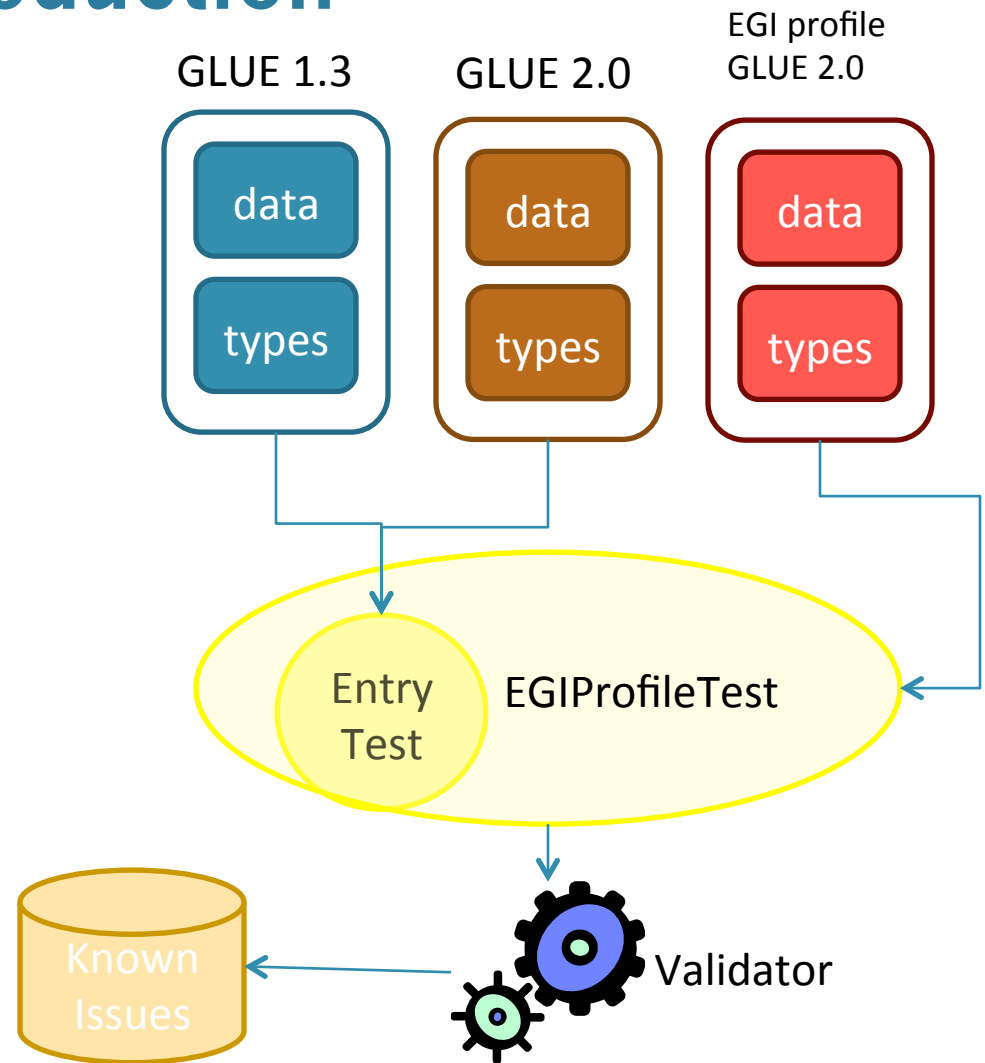


Part one – What is and how to use glue-validator




Introduction

- Glue-validator is a command line tool written in python
- It is able to validate against GLUE 1.3, GLUE 2.0 and EGI profile for GLUE 2.0



Where to get glue-validator



- EMI/UMD repositories
- EPEL repositories
 - Obsolete version right now!
 - To be updated in the upcoming weeks
- Midmon server
 - Limited to site validation
 - <https://midmon.egi.eu/nagios>
- Access to CERN AFS?
 - Latest version installed in malandes public area



glue-validator libraries

- The data library contains a description of the GLUE schema:
 - Object Classes
 - Attributes (type, single/multi valued, mandatory or

```
schema = {
  'organization' : {
    'o' : ('String', True, 'Mandatory'),
  },
  'GLUE2Entity' : {
    'dn' : ('DN_t', True, 'Mandatory'),
    'objectClass' : ('ObjectClass', False, 'Mandatory'),
    'GLUE2EntityCreationTime' : ('DateTime_t', True, 'Recommended'),
    'GLUE2EntityValidity' : ('UInt64', True, 'Desirable'),
    'GLUE2EntityOtherInfo' : ('EntityOtherInfo_t', False, 'Optional'),
  },
  'GLUE2Group' : {
    'GLUE2GroupID' : ('URI', True, 'Mandatory'),
    'GLUE2GroupName' : ('URI', True, 'Optional'),
  },
  'GLUE2Extension' : {
    'GLUE2ExtensionLocalID' : ('LocalID_t', True, 'Mandatory'),
    'GLUE2ExtensionKey' : ('String', True, 'Mandatory'),
    'GLUE2ExtensionValue' : ('String', True, 'Mandatory'),
  },
  'GLUE2Location' : {
    'GLUE2LocationID' : ('URI', True, 'Mandatory'),
    'GLUE2LocationAddress' : ('String', True, 'Optional'),
    'GLUE2LocationPlace' : ('String', True, 'Desirable'),
    'GLUE2LocationCountry' : ('String', True, 'Recommended'),
    'GLUE2LocationPostCode' : ('String', True, 'Optional'),
    'GLUE2LocationLatitude' : ('Real32', True, 'Recommended'),
    'GLUE2LocationLongitude' : ('Real32', True, 'Recommended'),
    'GLUE2LocationDomainForeignKey' : ('URI', False, 'Optional'),
  },
}
```



glue-validator libraries

- The type library contains a description of the types as defined by the GLUE schema
 - Enumerations are defined within the OGF GLUE working group
 - New values will be added as needed

```
        return True
    else:
        return False

def is_AppEnvState_t(value):
    states = [
        'installable',
        'installationfailed',
        'installedbroken',
        'installednotverified',
        'installedverified',
        'installingautomatically',
        'installingmanually',
        'notinstallable',
        'pendingremoval',
        'removing'
    ]
    if value in states:
        return True
    else:
        return False

def is_SchedulingPolicy_t(value):
    policies = [
        'fairshare',
        'fifo',
        'random'
    ]
    if value in policies:
        return True
    else:
```



glue-validator libraries

- There are different libraries containing the actual tests:
 - EntryTest: general tests for all attributes

```
def test_single_valued(self):  
    """Verifying single-valued attributes only have one value."""  
    status = True  
    message = ""  
    for obj in self.objects:  
        if obj in self.schema:  
            for attribute in self.schema[obj]:  
                if self.schema[obj][attribute][1] and attribute in self.entry and len(self.entry[attribute]) > 1:  
                    message = message + \  
                        validator.utils.message_generator("WARNING", "W036", self.dn, attribute, self.entry[attribute])  
                    status = False  
    self.assertTrue(status, message)
```

- EGIProfileTest: specific tests per attribute

```
def test_GLUE2LocationLongitude_OK (self):  
    message = validator.utils.message_generator("ERROR", "E008", self.dn, "GLUE2LocationLongitude", self.value[0])  
    self.assertTrue( float(self.value[0]) > -180 and float(self.value[0]) < 180, message)
```



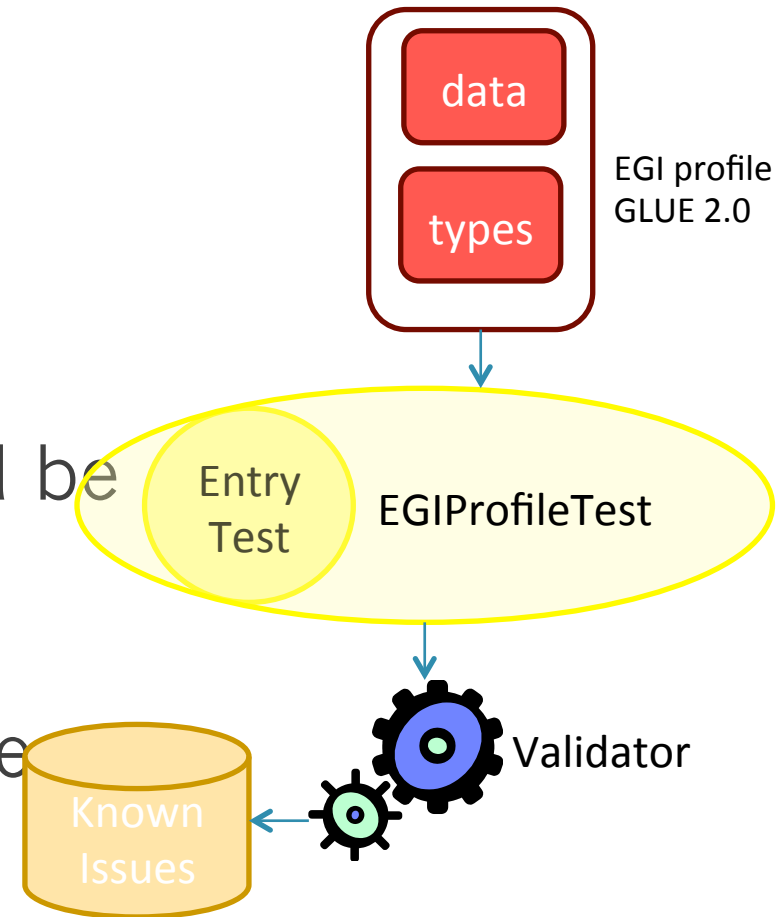

glue-validator libraries

- The KnownIssues library contains a list of tests for GLUE attributes that are wrongly published due to known issues in the middleware

```
test_list = (  
    'test_GLUE2EndpointImplementationVersion_OK',  
    'test_GLUE2StorageServiceCapacityTotalSize_OK',  
    'test_GLUE2StorageServiceCapacityFreeSize_OK',  
    'test_GLUE2StorageServiceCapacityUsedSize_OK',  
    'test_GLUE2StorageServiceCapacityReservedSize_OK',  
    'test_GLUE2StorageShareCapacityTotalSize_OK',  
    'test_GLUE2StorageShareCapacityFreeSize_OK',  
    'test_GLUE2StorageShareCapacityUsedSize_OK',  
    'test_GLUE2EntityValidity_OK',  
    'test_mandatory_attributes',  
    'test_data_types'  
)
```

Scope of this training

- Validation against the EGI profile for GLUE 2
 - It specifies how the information schema should be used in EGI
 - How information should be interpreted
 - What uses are likely
 - How information may be validated



Command Line Options

```
Usage: ./glue-validator [LDIF OPTIONS] [-g] [-s] [Other Options]

Mandatory Arguments:

Server Mode: Obtains LDIF from an OpenLDAP server
-H --hostname      Hostname of the LDAP server
-p --port          Port for the LDAP server
-b --bind          The bind point for the LDAP server

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

Optional Arguments:

GLUE version: Selects the GLUE schema version to be tested
-g --glue-version  The glue schema version to be tested [glue1|glue2|egi-glue2 (default)]

Testsuite type: Selects the set of tests to be executed against the LDIF
-s --testsuite     The testsuite [general|egi-profile (default)]

Other Options:
-k --exclude-known-issues Do not run tests for wrongly published attributes due to known bugs
-t --timeout              glue-validator runtime timeout, default 10s
-v --verbose              Verbosity level 0-3, default 0
-r --separator            Defines the separator for the output messages, default \n
                          This is only available for the verbosity level 3.
-V --version              Prints glue-validator version
-h --help                 Prints glue-validator usage

Examples:

GLUE 1.3 validation: glue-validator -H localhost -p 2170 -b o=grid -g glue1
GLUE 2.0 validation: glue-validator -H localhost -p 2170 -b o=glue -g glue2
EGI profile against GLUE 2.0 validation: glue-validator -H localhost -p 2170 -b o=glue
```



Command Line Options

- Very similar to `ldapsearch`
`glue-validator -H hostname -p port -b binding`
`ldapsearch -x -LLL -h hostname -p port -b binding`
- By default, validation is against the EGI profile for GLUE 2.0
- Some interesting options
 - Verbosity (default is 1)
 - Actually level 0 and 1 are the same → to be fixed



Command Line Options

- Some more interesting options
 - Exclude known issues
 - This is a very useful option for sites
 - Avoids running tests that are known to fail due to bugs in the info providers
 - This option will be always used in production
 - Timeout
 - Useful when validating top BDIIs
 - Separator
 - Useful to manipulate detailed output



Output format

- Nagios output with different verbose options
- Level 0 and 1
 - Number of errors, warnings and info messages

```
CRITICAL - errors 9, warnings 483, info 1825 | errors=9;warnings=483;info=1825
```


- Level 2
 - Details per message type

```
CRITICAL - errors 9, warnings 480, info 1825 | errors=9;warnings=480;info=1825
```

```
Summary per type of error, warning and info message:
```

```
E002 - Obsolete entry (GLUE2EntityValidity): 9
```

```
I012 - Unknown VO name in share (GLUE2EntityOtherInfo): 21
```



Output format

- Level 3
 - Affected DN, attribute and published value
 - I012 Description: Unknown V0 name in share
 - I012 Affected DN:
 - GLUE2ManagerID=ce207.cern.ch_ComputingElement_Manager
 - GLUE2ServiceID=ce207.cern.ch_ComputingElement
 - GLUE2GroupID=resource
 - GLUE2DomainID=CERN-PROD
 - o=glue
 - I012 Affected attribute: GLUE2EntityOtherInfo: Share
 - I012 Published value: na48



Error messages

- Three types of messages
 - ERROR:
 - Values that are definitely invalid
 - WARNING
 - Values that are likely, but not certain, to be wrong
 - INFO
 - Values that may be valid but that are unknown or seem wrong to glue-validator
- Only ERROR messages will raise a CRITICAL error in Nagios
- Twiki giving more details on each error
 - Tips on how to fix the error
 - Bug in the Information provider
 - Misconfiguration of the site
 - Whether there are any known issues



Error messages

Easily identify error number

Error messages

The follow error messages, raise a CRITICAL error in Nagios:

Error message number	Description	Associated test	Tips	Known Issues
001	Creation time in the future	test_GLUE2EntityCreationTime_OK	Site misconfiguration: The test fails when the published time is two minutes more than the time of the machine where the test is run. Check the times in your machines to make sure they are correct	NO
002	Obsolete entry	test_GLUE2EntityValidity_OK	Site misconfiguration: For non static entries (static entries come from static LDIF files and have no GLUE2EntityValidity published) this could be related to BUG:101237 . Update the bdii to version $\geq 5.2.20$ to get rid of ...	NO

Guidelines on what to do to get rid of the error

Whether there are any known bugs affecting the attribute publication



Some examples

- Remember to export the validator libraries in PYTHONPATH

```
export PYTHONPATH=$PYTHONPATH:/afs/cern.ch/user/m/malandes/  
public/glue-validator/usr/lib/python2.4/site-packages/
```

- Site validation

```
glue-validator -H prod-bdii -p 2170 -b o=glue  
glue-validator -H prod-bdii -p 2170 -b o=glue -v 2  
glue-validator -H prod-bdii -p 2170 -b o=glue -v 3  
glue-validator -H prod-bdii -p 2170 -b o=glue -v 3 -r " "  
glue-validator -H lcg-bdii -p 2170 -b GLUE2DomainID=CERN-  
PROD, GLUE2GroupID=grid, o=glue
```



Some examples

- Resource validation

```
glue-validator -H prod-bdii -p 2170 -b  
GLUE2GroupID=resource,o=glue
```

```
glue-validator -H prod-bdii -p 2170 -b "o=glue  
'(objectClass=GLUE2ComputingService)'"
```

```
glue-validator -H prod-bdii -p 2170 -b  
GLUE2ServiceID=ce206.cern.ch_ComputingElement,GLUE2GroupID=res  
ource,GLUE2DomainID=CERN-PROD,o=glue
```

- Top BDII validation

- Do you really want to do this? 😊

```
glue-validator -H lcg-bdii -p 2170 -b o=glue
```

Example with verbosity level 2

```
[lxadm12] /afs/cern.ch/user/m/malandes/public/glue-validator/usr/bin > ./glue-validator -H prod-bdii -p 2170 -b o=
glue -v 2
WARNING - errors 0, warnings 480, info 1811 | errors=0;warnings=480;info=1811
Summary per type of error, warning and info message:
I001 - Unknown Grid Infrastructure name (GLUE2EntityOtherInfo): 1
I010 - Incorrect middleware version (GLUE2EntityOtherInfo): 9
I012 - Unknown VO name in share (GLUE2EntityOtherInfo): 17
I046 - Number of seconds higher than 1 million (GLUE2ComputingShareEstimatedAverageWaitingTime): 184
I047 - Number of seconds higher than 1 million (GLUE2ComputingShareEstimatedWorstWaitingTime): 190
I067 - Number of physical CPUs greater than 10 (GLUE2ExecutionEnvironmentPhysicalCPUs): 2
I095 - Recommended attribute not present (Several attributes): 1408
W025 - Incoherent number of total jobs (GLUE2ComputingShareTotalJobs): 184
W034 - Mandatory attribute not present (Several attributes): 176
W037 - Wrong type (Several attributes): 95
W041 - Rpm used to publish implementation version does not exist (GLUE2EndpointImplementationVersion): 25
```

```
[lxadm12] /afs/cern.ch/user/m/malandes/public/glue-validator/usr/bin > ./glue-validator -H prod-bdii -p 2170 -b o=
glue -v 2 --exclude-known-issues
WARNING - errors 0, warnings 184, info 403 | errors=0;warnings=184;info=403
Summary per type of error, warning and info message:
I001 - Unknown Grid Infrastructure name (GLUE2EntityOtherInfo): 1
I010 - Incorrect middleware version (GLUE2EntityOtherInfo): 9
I012 - Unknown VO name in share (GLUE2EntityOtherInfo): 17
I046 - Number of seconds higher than 1 million (GLUE2ComputingShareEstimatedAverageWaitingTime): 184
I047 - Number of seconds higher than 1 million (GLUE2ComputingShareEstimatedWorstWaitingTime): 190
I067 - Number of physical CPUs greater than 10 (GLUE2ExecutionEnvironmentPhysicalCPUs): 2
W025 - Incoherent number of total jobs (GLUE2ComputingShareTotalJobs): 184
```

Example with verbosity level 3

```
[lxadm12] /afs/cern.ch/user/m/malandes/public/glue-validator/usr/bin > ./glue-validator -H prod-bdii -p 2170 -b o=
glue -v 3 --exclude-known-issues
WARNING - errors 0, warnings 184, info 401 | errors=0;warnings=184;info=401

W025 Description: Incoherent number of total jobs
W025 Affected DN: GLUE2ShareID=grid_sixt_vo.sixt.cern.ch_ce203.cern.ch_ComputingElement
                  GLUE2ServiceID=ce203.cern.ch_ComputingElement
                  GLUE2GroupID=resource
                  GLUE2DomainID=CERN-PROD
                  o=glue
W025 Affected attribute: GLUE2ComputingShareTotalJobs
W025 Published value: 0
W025 Additional information: GLUE2ComputingShareRunningJobs=0 GLUE2ComputingShareWaitingJobs=215202 Difference is
                             215202

W025 Description: Incoherent number of total jobs
W025 Affected DN: GLUE2ShareID=grid_sixt_vo.sixt.cern.ch_ce205.cern.ch_ComputingElement
                  GLUE2ServiceID=ce205.cern.ch_ComputingElement
                  GLUE2GroupID=resource
                  GLUE2DomainID=CERN-PROD
                  o=glue
W025 Affected attribute: GLUE2ComputingShareTotalJobs
W025 Published value: 0
W025 Additional information: GLUE2ComputingShareRunningJobs=0 GLUE2ComputingShareWaitingJobs=215202 Difference is
```

Example with separator

```
[lxadm12] /afs/cern.ch/user/m/malandes/public/glue-validator/usr/bin > ./glue-validator -H prod-bdii -p 2170 -b o=glue -v 3 --exclude-known-issues -r " "
WARNING - errors 0, warnings 184, info 401 | errors=0;warnings=184;info=401
W025 Description: Incoherent number of total jobs W025 Affected DN: GLUE2ShareID=grid_sixt_vo.sixt.cern.ch_ce203.cern.ch_ComputingElement,GLUE2ServiceID=ce203.cern.ch_ComputingElement,GLUE2GroupID=resource,GLUE2DomainID=CERN-PROD,o=glue W025 Affected attribute: GLUE2ComputingShareTotalJobs W025 Published value: 0 W025 Additional information: GLUE2ComputingShareRunningJobs=0 GLUE2ComputingShareWaitingJobs=215202 Difference is 215202

W025 Description: Incoherent number of total jobs W025 Affected DN: GLUE2ShareID=grid_eela_prod.vo.eu-eela.eu_ce205.cern.ch_ComputingElement,GLUE2ServiceID=ce205.cern.ch_ComputingElement,GLUE2GroupID=resource,GLUE2DomainID=CERN-PROD,o=glue W025 Affected attribute: GLUE2ComputingShareTotalJobs W025 Published value: 0 W025 Additional information: GLUE2ComputingShareRunningJobs=0 GLUE2ComputingShareWaitingJobs=215202 Difference is 215202

W025 Description: Incoherent number of total jobs W025 Affected DN: GLUE2ShareID=grid_alice_alice_ce208.cern.ch_ComputingElement,GLUE2ServiceID=ce208.cern.ch_ComputingElement,GLUE2GroupID=resource,GLUE2DomainID=CERN-PROD,o=glue W025 Affected attribute: GLUE2ComputingShareTotalJobs W025 Published value: 0 W025 Additional information: GLUE2ComputingShareRunningJobs=0 GLUE2ComputingShareWaitingJobs=215243 Difference is 215243

W025 Description: Incoherent number of total jobs W025 Affected DN: GLUE2ShareID=grid_eela_prod.vo.eu-eela.eu_ce204.cern.ch_ComputingElement,GLUE2ServiceID=ce204.cern.ch_ComputingElement,GLUE2GroupID=resource,GLUE2DomainID=CERN-PROD,o=glue W025 Affected attribute: GLUE2ComputingShareTotalJobs W025 Published value: 0 W025 Additional information: GLUE2ComputingShareRunningJobs=0 GLUE2ComputingShareWaitingJobs=215202 Difference is 215202
```



Part two – How to improve things with glue-validator



Current status of GLUE validation

- Monthly reports since March 2013
 - Only for WLCG sites for practical reasons
- Manual review of the glue-validator results
 - Ticketing sites
- This approach helped tuning glue-validator
 - And already improved the overall quality
- Some improvements so far but...



Future validation process



- Deploy glue-validator as a Nagios probe
 - Automatic and stable validation process
 - glue-validator already deployed in Midmon
 - As soon as probe is validated it will become a production probe
 - Sites will get tickets from ROD team for critical errors if not fixed after 24h
- glue-validator will be also used in the EGI middleware acceptance tests


Long term goals



- Integration of glue-validator in the resource BDII
 - Enforce early validation in the development stage
 - Requires agreement and coordination with product teams
 - Change of current way of working
 - Is it better to publish nothing than something wrong?



Part three – Most common GLUE 2 errors and how to fix them



Common errors



- Operating System Information
 - Operating system names and versions:
https://wiki.egi.eu/wiki/HOWTO05_How_to_publish_the_OS_name
 - Easy to fix in YAIM:
 - CE_OS →
GLUE2ExecutionEnvironmentOSName
 - CE_OS_RELEASE →
GLUE2ExecutionEnvironmentOSVersion

Common errors



- Batch system attributes
 - In many places default values are published
 - It is OK as far as this is what you want!
 - Configuring the batch system seems to be a complex task
 - Some guidelines here:
 - https://twiki.cern.ch/twiki/bin/view/EGEE/GLUEMonitoring#GluePolicy_GLUE2ComputingShare_a
 - <https://wiki.egi.eu/wiki/Tools/Manuals/TS59>
 - Many GLUE attributes depend on the batch system configuration!



Common errors

- VO, WLCG and Grid Infrastructure names
 - VO names:
 - <http://operations-portal.egi.eu/vo>
 - WLCG names:
 - <http://rebus.cern.ch/apps/topology/>
 - Grid Infrastructure names:
 - https://wiki.egi.eu/wiki/MAN01_How_to_publish_Site_Information
 - What to do if you still want to publish a value that does not exist in any of the above?
 - Please, let us know!



Pending known issues

- Storage related errors are done
- Computing-related errors still to be evaluated
 - Marked as “?” in Error Twiki
 - https://twiki.cern.ch/twiki/bin/view/EGEE/GLUEMonitoring#444444_waiting_jobs
 - List of known issues may be modified
- A final version ready for validation will include these ones as well!

Feedback



- Glue-validator not yet used in production
 - Feedback for the tests
 - Are the tests useful?
 - Are they reporting properly?
 - This all may have an impact on GLUE 2 profile too!
 - Feedback for the error messages
 - Are the tips useful?
 - Feedback for the known issues
 - Are there any more known issues to be added?
 - If the sites have nothing to do for a certain error, it should be a known issue!
 - And feedback for anything else! (bugs, usability, etc)
 - Please use GGUS

Useful links

- glue-validator guide
<http://gridinfo.web.cern.ch/glue/glue-validator-guide>
- glue-validator code
<https://svnweb.cern.ch/trac/gridinfo/browser/glue-validator/>
- EGI profile for GLUE 2.0
<http://go.egi.eu/glue2-profile>
- Error messages
<https://twiki.cern.ch/twiki/bin/view/EGEE/GLUEValidatorErrorCodes>
- GLUE 2 validation monitoring
<https://twiki.cern.ch/twiki/bin/view/EGEE/GLUEMonitoring>