



Information System Status and Evolution

Maria Alandes Pradillo, CERN
CERN IT Department, Grid Technology Group
GDB 14th November 2012



Overview

- Current status
 - Development
 - Quality of Service and Data
 - Requirements from experiments
 - Statistics on current use
- Evolution
 - Service Discovery
- Conclusions




Current status of the Information System




Development

BDII development status

Information System GGUS tickets
May to November 2012



Open/Closed tickets since May
2012



- No major incidents
- No releases since August
- EMI 1, EMI 2 and EMI 3 releases are all aligned
 - Note that EMI 1 Standard updates are over
 - Direct upgrade path from EMI 1 to EMI 2 for the BDII
- Details on EMI releases:



http://malandes.web.cern.ch/malandes/infosys/bdii_emi.html

BDII development status

- Next release in December:

- EPEL compliance
- EMIR integration
- ARC integration
- glue-validator improvements
- Service information provider bug fixes



- Continuous improvement of the Documentation


https://tomtools.cern.ch/confluence/download/attachments/983044/EMI_BDII_sysadmin.pdf

- Latest changes:

- glite-info-update-endpoints configuration
- Site BDII configuration aligned with EGI guidelines

- Upcoming changes:

- Nagios probes documentation
- EMIR publisher configuration



Quality of Service

Top level BDIIs and failover



- Initial proposal to configure failover

https://tomtools.cern.ch/confluence/display/IS/WLCG_Support_Proposal

- EGI recommends sites to use the NGI top level BDII:

https://wiki.egi.eu/wiki/Top-BDII_list_for_NGI


- EGI monitors Reliability&Availability of top level BDIIs:

<https://documents.egi.eu/public/ShowDocument?docid=1429>

NGI top level BDII availability

October 2012

■ 99 to 100% ■ NA ■ Less than 99%




Follow up actions by EGI

- Underperforming top level BDIIs in small NGIs are asked to be decommissioned
- Catch all top level BDII provided by EGI or good performing neighbour NGI top level BDII used instead
- NGIs contacted monthly to explain reasons for underperformance.

Top level BDIs and failover

- Configuration of LCG_GFAL_INFOSYS variable hasn't changed since June 2012
 - Most sites still configure one BDII only
 - Good news is that the most configured BDIs have 100% availability




- **Next steps:**

- **Work together with EGI:**

- No clear policy on LCG_GFAL_INFOSYS configuration
 - Good quality of service of NGI top level BDIs
 - To be monitored together with EGI and together with them discuss further actions

Number of BDIs configured
in LCG_GFAL_INFOSYS
Oct-12



(*) The NA in the graph is due to a bug in SAM already fixed.



Data Quality


Glue 2.0

- More on Stephen Burke's presentation
 - EGI Glue 2.0 profile
 - Deployment of Glue 2.0
- From the Information System point of view
 - Improve glue-validator to check compliance
 - Collaborate with EGI to check and monitor Glue 2.0 deployment




To be discussed with next presentation...

GLUE 2.0 Deployment in WLCG



WLCG sites not publishing GLUE 2.0





Requirements from experiments

Multicore support


- Is there anything else needed from IS?
 - Attempt to restart discussion in WM TEG ML in August
 - No more reactions
 - As reported in the last September GDB
 - <https://indico.cern.ch/getFile.py/access?contribId=7&sessionId=1&resId=1&materialId=slides&confId=155072>
 - Information System is ready
 - We can address any requirements




Requirements from experiments

Who?	Requirements	Addressed by
ALICE	<ul style="list-style-type: none">Reliable information published by the resource BDII in the CREAM CE	<ul style="list-style-type: none">Follow up on information provider issues
ATLAS	<ul style="list-style-type: none">Reliable service discovery (GOCDDB, OIM and BDII in sync)Better SE space utilization information and in sync with Storage AccountingReliable client tools	<ul style="list-style-type: none">Cached top BDII
CMS		<ul style="list-style-type: none">Cached top BDIIUnderstand Storage Accounting and what would be needed from the ISginfo
LHCb	<ul style="list-style-type: none">Reliable service discovery of CEs	<ul style="list-style-type: none">Cached top BDII

- Service Discovery is the most important use case for experiments





Statistics on current use

Top BDII query statistics

- Analyzed LDAP logs from production top BDIIs
 - Using the ldap-stats.pl script
 - <http://prefetch.net/code/ldap-stats.pl.html>
- Thanks to:
 - CERN – Ulrich Schwickerath
 - TRIUMF - Di Qing
 - PIC - Arnau Bria, Carles Costa



Top BDII query statistics

Site	Queries in one week
CERN	3223758
TRIUMF	397701
PIC	36007

Most queried attributes

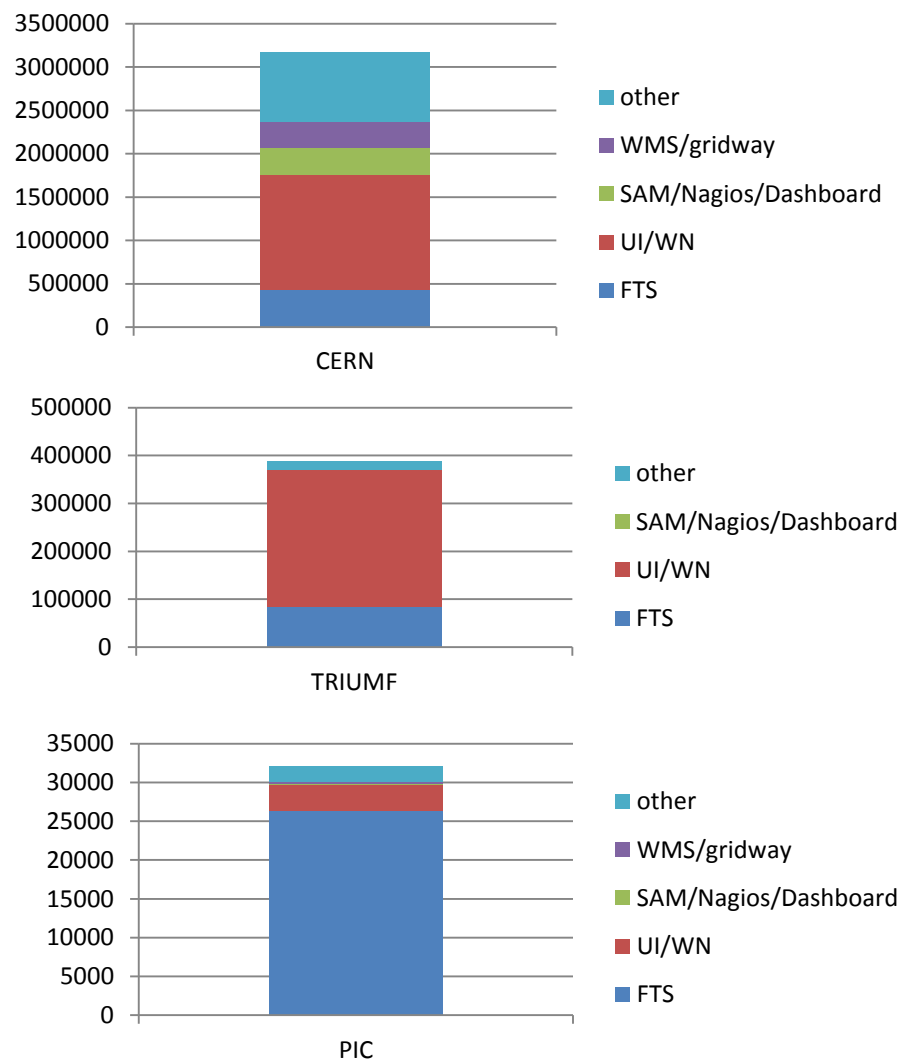
No attribute explicitly requested

GLUEService (version, type, endpoint)

GLUESE (port, status)


GLUESA (root, accesscontrolbaserule)

GLUECE (name, status, uniqueID)



Top BDII query statistics


- Difficult to extract conclusions
 - No common pattern
 - Analysis from only one week!
 - This would require a longer time window -> Log file increases fast (>15 GB in one week)
- It was an interesting exercise though
 - BDII heavily used
 - Stability, performance and robustness very important
 - Which attributes are most queried



Evolution of the Information System

Service Discovery

- Focus on Service Discovery
- Cached top level BDII has improved the perceived stability
- How can we make it even better?
 - Monitor quality of the published information
 - As part of the overall GLUE 2.0 effort
 - Provide a reliable client tool to query the IS
 - For both EGI and OSG resources
 - OSG is not publishing in GLUE 2.0 though!
 - Study other possible implementations: EMIR



ginfo

- Simple command line tool
 - Queries BDII GLUE 2.0 Service Endpoints
- Available in EPEL for SL5 and SL6
 - <http://fedoraproject.org/wiki/EPEL>
 - yum install ginfo
- Ready to be used!
 - Feedback is very much welcome!
 - project-grid-info-support@cern.ch
 - Already some feedback from CMS and ATLAS
 - Missing other GLUE objects currently available in lcg-info
 - Query OSG resources?


GLUE 2 attribute filters

EndpointCapability
ServiceAdminDomainForeignKey
EndpointID
ServiceID
EndpointImplementationName
EndpointImplementationVersion
EndpointInterfaceName
EndpointInterfaceVersion
EndpointQualityLevel
EndpointTechnology
ServiceType
EndpointURL
PolicyRule



EMIR

- EMI Service Registry
- Provides References to Services
 - Endpoints
 - Associated Metadata
 - Service Type, Supported VO, Capability
- Automated Information Management
 - Services are the authoritative source
- Components
 - Service Publisher
 - Domain Service Registry
 - Global Service Registry
- A domain could be a site, a country, a region...
 - It's a collection of services
- EMIR pilot testing is ongoing
 - Some robustness issues detected but overall feedback is positive
 - Not ready for production deployment yet



Conclusions

- Current status
 - BDII performing well in production
 - On going development:
 - EPEL compliance
 - EMIR integration
 - ARC integration
 - glue-validator improvements
 - Service information provider bug fixes
- Future work
 - Focus on Service Discovery
 - Follow up Glue 2.0 deployment and evaluate published information
 - Follow up EMIR pilot and ginfo