


Update on the WLCG Information System

Lorenzo Dini – WLCG Information Officer

May 11, 2011



Agenda

- Introduction
- Use cases document
- Quality control
- Cached BDII
- Reliable Top-level BDII deployments
- Conclusions

Introduction

- Last presentation: Dec 2010 at MB
- I replaced Flavia as WLCG Information Officer in April
- Most of the progress on the use cases and quality control was achieved by Flavia
- I focused on the cached BDII testing
- For any information sharing, suggestions or feedback: Lorenzo.Dini@cern.ch

Use Cases Document

- Collection of use cases to detail the way the Information System is used by all dependent WLCG services
 - + Focus on the useful information
 - + SLA with the users
 - It defines what is currently used, not what is planned to be used or planned to be discontinued
- https://twiki.cern.ch/twiki/pub/LCG/WLCGISArea/WLCG_IS_UseCases.pdf

Main Use Cases - 1

ALICE

- Regulate the flow of jobs to a particular site
- Identify CEs in production mode

ATLAS

- Panda bootstrap-time self configuration
- Monitor the resource usage (DDM)
- Via FTS
- Frontier-Squid SysAdmin contacts

CMS

- Via *glite-wms-job-list-match* in the JDL
- CMSSW for automatic software installation at sites
- CMS Site Dashboard
- CMS PhEDEx via FTS

LHCB

- DIRAC Pilot Job delivery
- LHCb CS

Main Use Cases - 2

FTS

- Produces hourly a SE list in XML
- Queries for SEs if not in the list

Dashboard


- Site Latitude and Longitude

Installed capacity

- Cluster and SubCluster hardware specifications
- CE job slots, capabilities and shares
- VoView assigned job slots
- SE and SA sizes and status
- Access and Control Protocols

Quality Control

- Monitoring the most important attributes
- Currently 32 attributes under test
- Heuristics to identify problems
- 63 tickets GGUS currently open



<http://grid-deployment.web.cern.ch/grid-deployment/flavia/ISQuality/ISresults.html>

Lessons learned

- Manual opening and follow-up of tickets is very time consuming
- Sites demand precise information to be provided in the ticket
- While followed attributes rapidly improve, all others slowly degrade
- Need for automation

Quality Control Automation

- BDII development team is preparing an extensible GLUE schema validation tool
- Could be integrated in:
 - New GSTAT view for BDII administrators
 - Local tool to be run during BDII setup
 - Nagios probes to automatically open GGUS tickets
- Rule-set under preparation to be run as the WLCG profile

Why the need for a cached BDII

- Most of the important information queried rarely changes
- Possible causes of instabilities for semi-static information:
 - BDII service problems
 - Publishing service problems
 - Network problems
 - Planned downtimes
- This causes temporary object drops


Cached BDII

- New BDII delayed removal of objects
- Already available in production
- Disabled by default `BDII_DELETE_DELAY=0`
- Marks the objects *toBeDeleted* as:
 - GlueCEStateStatus = “Unknown”
 - GlueSEStatus = “Unknown”
 - GlueServiceStatus = “Unknown”
- Need to carefully validate with users that the behavior does not cause any issue with services

Cached BDII Testing


- 2 VMs deployed from the GT test-bed
 - Same VM specs (4 cores, 4 GB RAM)
 - Same BDII versions
 - *BDII_DELETE_DELAY=345600* (4 days) as cache
 - Waited 4 days for cache to be stable
 - Every 15 minutes run tool performing delta operations of BDII content in both cached and non-cached
- Production-like hardware received only yesterday for testing

Total number of objects




- From 152.000 to 161.000 objects
- 5% size increase
- The cache smoothes down the instabilities

Unstable objects – non-cached




- GlueLocationLocalID highly unstable
- Removed from next plots

Unstable per site – non-cached




- AGLT2 seems to have caused the major instabilities
- This week they have disappeared

Cached BDII improvements



- 1 order of magnitude improvement with cached BDII
- Even not considering GlueLocationLocalID
- The unstable objects are often the same and the data is stabilized by the cache

Unstable per site - cached



- Minor daily inconsistencies
- Possibly due to planned downtimes or incidents

Reliable Top-level BDII

- Top-level BDII deployment specifications:

https://twiki.cern.ch/twiki/pub/LCG/WLCGISArea/BDII_Deployment_Plan.pdf

- Sites volunteering to host service:
 - Taiwan-LCG2, NL-T1, FZK-LCG2, OSG, PIC, RAL, IN2P3-CC, TRIUMF-LCG2
- In contact with OSG. Waiting for the cached BDII to be ready for production
- Site-level BDII: migration to LDAP 2.4

Conclusions

- Use Case document to be completed
 - Defines the SLAs with the users
 - Helps to focus on important data
 - Will review the existing document with users
- Quality Control to be automated
 - 32 attributes, 63 GGUS tickets open, situation improved
 - Automation under preparation
- Cached BDII improves stability
 - 4-day cache radically improves object stability
 - GlueLocationLocalID the most unstable object
 - Need to address some site temporary instabilities
 - Soon available statistics on performance, object types and sites
- Sites are ready for a reliable deployment