



# WLCG Worldwide LHC Computing Grid

Middleware Review

Problems

Roadmap

CCRC 08

Markus Schulz  
CERN IT-GD-ITR & EGEE SA3  
Markus.schulz@cern.ch



EGEE-II INFISO-RI-031688



Information Society  
and Media





# Middleware Problems

## Lcg-utils < ----- > classic SE

- Lcg-utils Segfaults when used with a Classic SE
  - Bug #32016 filed 07/12/07 (lcg-cr )
  - Fixed in Patch [#1587](#)
- This introduced lcg-cp problem ( Bug [#32808](#) )
  - Fixed in Patch [#1641](#) ( 2/01/08 )
- A bit later: lcg-rep and Classic SE problem re-appeared
  - Bug [#32999](#)
  - Fixed in Patch [#1658](#) ( 29/01/07)
- **Who was affected?**
  - Some larger sites use dedicated Classic-SE for ops VO to get better SAM results → wrong status for production SE
  - Sites < T2 ( most Vos stage data from WN scratch to T0/1/2 )
- **How could we miss this?**
- During the transition to SL4 we had for a while no Classic SE on the testbed



# Middleware Problems

## Long FQNS ( VO name bug)

- Patch [#1641](#) (GFAL lcg-util) released, fixing 11 bugs (segfaults, SE..)
  - This introduced the VO name bug Bug [#33327](#) (07/02/08)
  - Testbed VOMS server hosts various Vos (including long VO names)
    - This problem had been seen before.....
- Patch was not removed as HEP sites benefited from the fixes
- Patch [#1679](#) fixed the problem
- **How could we miss this?**
- Due to the configuration of the testbed only a subset of Vos are tested
- Clearly a hole in the process
- **Reaction**
- Process has been updated
- Testbed is being reconfigured
- **Who was affected?**
- Sites hosting Vos with DNS style names ( and their Vos)



# Middleware Problems

## FTS Proxy Corruption

- Introduced with the introduction of the proxy delegation code
  - Sharing of credential between several transfers
  - 3 months pre-release
  - 3 months pilot service
  - 6 months production
  - **Hiding for 1 year**
- Reason: Race condition when 2 clients delegate at the same time
  - Key and cert information get mixed up → useless
- **Who was affected?**
  - Everyone using FTS
- **Fix:**
  - A client based workaround ( code change needed) :
  - Hack on server: delete malfunctioning proxy
    - service interruptions (until next transfer request arrives (minutes))
  - Proper fix: Not before end of July ( massive tests needed)



# General Comments

## Release Process

All components and Patches for the CCRC have been released following the documented release process.

This includes FTS and the fast tracked DM components

We have handled CCRC like normal production



# Medium Term Changes

- Since 6 months work has been directed towards an improved certification model
  - Isolate test activities from another
  - Stable physical testbed as a background for testing
  - Testing of new Patches with virtual test beds
- Better documentation of test results
  - Public access
- DM tests have been updated and will be further improved
  - Some problems have occurred with a classic SE as a source
  - Others when used as destination
- Problem: Combinations
  - Castor , DPM, Classic SE , dCache
  - Source/Destination, SRM versions, .....



# Other Problems

- Information system related
- FTS
- Suffers from static information being volatile
  - Many potential reasons
    - Network, site BDII server node overload,
    - Short dropout of the services.....
- Top level BDIIs have short caching times ( several minutes)
  - Behavior has been optimized for WMS
    - Don't send jobs to oversubscribed services
  - Caches ALL data
- **Plan:**
  - 4 hour deep cache for static data
  - August 2008
  - **Now:** Monitoring of ALL BDIIs <http://lxn1191/nagios/>



# Other Problems

- Information system related???
- Several problems with LHCb production at NIKHEF
- Combination of several problems
  - Not a general scalability problem of the information system
- Detailed postmortem has been carried out
  - Results have been widely circulated





# Recommendations

- 1: Update the deployment and trouble shooting documentation explaining clearly the dangers of co-hosting the BDII with other services which could generate a high load
- 2: Produce a new release of the BDII in which the new index is set.
- 3: Improve the efficiency of the lcg-utils commands, in particular as some kind of cache is need to avoid repeated queries to to BDII
- 4: Use the logs gathered from the incident to test the performance of the BDII in such situation and address and performance bottlenecks found.
- 5: Monitor the load on the BDII at NIKHEF. If the load is consistently high, consider adding and additional machine for load balancing.
- 6: Advise the VOs the dangers of fail over methods that can do DOS loops. All fail over should contain some exponential backup.
- 7: Implement, as planned service discovery APIs. These interfaces to the information systems should contain limiter that prevent single threads (better processes) from issuing rapid fire queries. This limiter is needed to prevent accidental DOS attacks that make the whole resource unusable



# Upcoming Releases/RoadMap

- CREAM CE
  - Stress tests (acceptance tests) started at INFN
  - As soon as the internal candidate passes a Patch will be created
    - Optimistic: 2 weeks
  - Certification will follow
    - 2-4 weeks depending on the quality of the YAIM module
  - Needs to pass through PPS
    - 2 weeks
- WMS on SL4
  - Many improvements (fixes) over current prod version
  - 5 day stability and scalability tests have started at CERN
    - Intensive testing at Imperial College + INFN + CERN
  - Assuming success, release to PPS within 2-3 weeks
    - Update of documentation ( description of open issues)
  - Several improvements are already lined up



# Upcoming Releases/RoadMap

- SCAS and glexec on the WN
  - Server version of LCAS has been coded
  - NIKHEF waits for a new version of an interoperability lib.
    - Provided by the globus team
  - Internal tests will start in March
  - Glexec on WN
    - Tests on different batch systems in preparation
  - SCAS will require extensive scalability and stability tests
    - SCAS is a potential SPOF on a site
- Estimate: 3-4 months away from production



# Upcoming Releases/RoadMap

- Implementation of the recommendation of the JPWG
  - This is the job priority system using VO-views
- Yaim changes and info providers are being prepared now
  - Affects YAIM core
  - Will be released with next core release of YAIM
- Functional tests carried out with prototype versions
- Template scheduler configuration only for Maui/PBS available
  - Need this for LSF and Condor
  - No real showstopper since larger sites can handle this
- Estimate 2 – 3 months to production



# Upcoming Releases/RoadMap

- AMGA-Oracle
  - In Certification
  - Can reach production within 4 weeks
- FTS on SL4 (last core component to move)
  - A Patch has been created (early stage)
  - Certification has been difficult in the past (2 months)
  
- SL5, gcc-4.2/4.3 ???
  - SL5 builds with default compiler started
  - We need clarification which combinations are required when?
  - Tentative schedule for SL5 WN September 08
    - Depends highly on TCG/TMB prioritization