



WLCG Service Report

Dirk.Duellmann@cern.ch

~ ~ ~

WLCG Management Board, 28st April 2009


Introduction

- This report covers the service since last week's MB
- Normal week for operations
 - One Service Incident Report from CC-IN2P3


GGUS Summaries

| VO concerned | USER | TEAM | ALARM | TOTAL |
|--------------|------|------|-------|-------|
| ALICE | 2 | 0 | 0 | 2 |
| ATLAS | 13 | 37 | 0 | 50 |
| CMS | 0 | 0 | 0 | 0 |
| LHCb | 17 | 13 | 0 | 30 |
| Totals | 32 | 51 | 0 | 82 |


GGUS tickets per VO




Site Availability using WLCG Availability (FCR critical)




Site Availability using WLCG_SRM2



Site Availability



Site Availability using LHCb Critical Availability



Other site issues/highlights

- ASGC cleanup for ATLAS
 - LFC: purged, isolating ATLAS, adding redundancy
 - Tier 2 cleanup ongoing with help from DPM support
 - Sites should not hesitate to contact service support for help
- CC-IN2P3
 - Service Incident Report received (see next slide)
- CERN
 - Significant performance boost in CASTOR after memory upgrade (8->16GB/node) for repack stager database
 - All sites should regularly review DB clusters for memory utilisation / swapping during increased load periods
 - Common DB monitoring proposed at DB workshop @ PIC

IN2P3 MSS Outage 20th April

Duration: 12 hours

Date: April 20th 2009

Description: Hardware failure of the robotic library inducing a global outage of the MSS.

Impact: MSS

Batch was unavailable for any job depending on MSS.

Local backup service interrupted during the outage.

Estimated 18% shortfall of running (non-grid) jobs during outage (jobs locked in queue).

Timeline of the Incident

Monday April 20th

11:12 incident report opened against our robotic supplier

12:15 robotic supplier on site – hardware and software checks start

18:00 final tests and diagnosis

19:30 hardware change

22:30 robotic library in operational state

23:10 MSS system accessible in degraded mode

Tuesday April 21st

11:30 Full pledge usage of MSS

Full report at <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGServiceIncidents>

Update: Oracle BIGID problem in CASTOR

- Better understanding of the pre-conditions under which the problem occurs
 - Affects Oracle server side data buffer when using “returning into” from OCCI using bulk operations
 - Can’t exclude yet a more general Oracle consistency issue
 - Test case has been developed - reproducible now in minutes!
- More Big ID cases observed at CERN and T1s
 - So far, failed CASTOR requests, but no data loss
- Further investigations with Oracle development
 - Several workaround scenarios being discussed
 - Finding the optimal (risk & effort) needs detailed (Oracle side) knowledge about the corruption process
- Using CERN direct contact into Oracle development to speed up information exchange
 - May need to raise priority for Oracle fix

Other service issues / news

- LHCb : LFC use via CORAL
 - Issues with DB credential look-up from LFC
 - discussion between LHCb, LFC and CORAL team on possible optimisations
 - Is LHCb accessing the correct LFC instance?
 - read-write versus read-only LFC?
 - site local LFC versus Tier 0 instance?
- ALICE : pre-certification WMS 3.2 deployed at two sites (GRIF)
 - issues with job wrapper observed and fixed

DB Workshop @ PIC

- Chaired by Maria Girone
 - <http://indico.cern.ch/conferenceOtherViews.py?view=standard&confId=54037>
- Main outcomes
 - Allocated DB resources for STEP'09 and for the 2009/2010 run
 - Positive test and deployment reports for DB set-ups on RHEL 5 - CERN upgrade by summer 2009
 - Tracking and announcing procedures reviewed and confirmed
 - Good progress with AMI replication preps between IN2P3 and CERN
 - Procedures and responsibilities for re-instantiation of DB between Tier 1 sites confirmed
 - Final NDGF database setup being setup in Oslo - DataGuard based migration from Helsinki
 - Pilot queries being proposed by ATLAS to throttle DB workload
 - Failover LFC database based on DataGuard between CNAF and ROMA1
 - No plans to upgrade to Oracle 11 before the LHC run - but 11.2 beta evaluations ongoing
 - Several castor DB operational issues and procedures discussed
 - Define a standard castor DB config between T1 sites (test setup at T0)

Summary

- Normal operation and number of GGUS tickets
- Main problems (fixed)
 - MSS outage at IN2P3
 - CNAF/CASTOR publishes incorrectly : FTS problem
 - CERN/CASTOR single user unintended "DoS attac"
 - CASTOR user monitoring in 2.1.8 will address this area
- Test case for Oracle "Big-ID" problem found
 - CASTOR risk assessment and search for workaround ongoing with Oracle development