



# WLCG Operations Coordination

Andrea Sciabà  
IT/SDC

On behalf of the WLCG Operations Coordination team

GDB  
11<sup>th</sup> December 2013



# Outline

- Previous report on October 9<sup>th</sup>
- News
- Summary of the planning meeting
- Status of task forces
- Experiment plans during the end-of-year break
- Conclusions

# News

- **New task force on multicore deployment approved**
- **Mandate**
  - Find out how experiments plan to **request and use** multicore resources
  - **Agree** on practical multicore resource provisioning **models** for WLCG and **evaluate** them at selected sites with real workflows
  - Define a **deployment strategy**
  - Drive the **deployment process**
- The scope is **Grid resources**, not cloud
  - But interaction with the Cloud WG is highly desirable
- Coordinate also with the Machine/Job features TF
- To be led by **Alessandra Forti** and **Antonio Perez-Calero**

# Recent and future WLCG Operations Coordination meetings

- Planning meeting on November 7<sup>th</sup>
- Until March:
  - December 19
  - January 16, 30
  - **February 11: WLCG operations coordination F2F at Pre-GDB**
  - February 20
  - March 6

# News from the planning meeting

- Cancelled two TFs previously proposed
  - Data access (e.g. Tier-1 WNs in the OPN): limited interest, significant changes to the site network infrastructure
    - Replaced by a [survey](#) among T1s
  - Dynamic Data Placement: postponed until the development phase is concluded for CMS and LHCb

# Experiment plans for the next months (1/2)

- ALICE
  - Finish conversion to CVMFS
  - Rationalisation of SAM tests
    - MonaLISA will inject xrootd and VOBOX test results into message bus
- ATLAS
  - Full reprocessing campaign from April
  - Progress with FAX integration and WebDAV deployment
  - Commission Rucio with sites
  - Work on cloud benchmarking (with HammerCloud?) and multicore resources

# Experiment plans (2/2)

- CMS
  - Extend multicore testing at sites
  - Complete disk/tape separation (3 T1s done)
  - Continue scale testing of xrootd federation and deployment of fallback and remote access
  - Commission Russian T1
- LHCb
  - Incremental stripping campaign in spring (6-8 weeks)
  - Switch to SL6 as default for user analysis from January
  - Consume perfSONAR and WLCG monitoring information in DIRAC

# SL6 task force final report

- Positive outcome of the task force
  - **92.6%** of the resources moved to SL6
  - Creation of WLCG repository
  - Cleaner HEP\_OSlibs RPM
  - EMI-3 WNs tested and usable
  - **15/16** T0/1's done
  - **124/131** T2's done
- **Thanks to Alessandra and all the TF members!**



# CVMFS

- ALICE
  - New CVMFS target deadline: end of 2013
- CMS
  - Only very few sites left!
- Some operational issues to be addressed
  - WN caches becoming stale from time to time, need for better fabric monitoring to spot it
  - Use a SAM probe (developed by Stefan) to check several aspects of the CVMFS deployment at the site
- CVMFS 2.1.15 is the baseline version

# gLExec

- Still about 30 sites to go
  - Some sites still needing to migrate to SL6
  - And some special cases: USATLAS, ARC CE sites, ALICE sites without CREAM
- ALICE and ATLAS still need development
- CMS will make gLExec SAM test critical in January
- LHCb is testing DIRAC pilot with gLExec

# SHA-2

- Since December 1<sup>st</sup> CAs may decide to issue SHA-2 certificates by default
  - CERN CA will do it “when WLCG is ready”, some time early next year
  - OSG will start mid January
- Only very few StoRM and dCache instances still need upgrading
  - Notably FNAL and BNL, to be done by end December
  - EGI sites tracked by EGI operations
- Experiments look ready

# perfSONAR

- All sites must deploy pS 3.3.1 (or later) by April 1<sup>st</sup>
  - The TF provides all the needed support
  - Sites not yet ready are now tracked by tickets
    - pS down, not registered, not updated, no mesh, etc.
- If a pS instance is behind a firewall, it should open access to IPs from CERN+Wigner to allow central operations
- Monitoring will be run by OSG Operations
  - **Need to find someone to take responsibility for the code**
- More information in TF [twiki](#)

# FTS-3

- Service stable in the last 2 months after several bug fixes
- Experiment usage
  - ATLAS: 30% of all production transfers + functional tests at all sites
  - CMS: 30% of debug transfers
  - LHCb: all WAN transfers
- Investigating deployment scenarios
  - Single instance clearly the preferred option
  - Currently running performance tests comparing autoconfiguration vs. manual configuration
  - Need more testing to measure scalability and behaviour when multiple instances run in parallel

# Tracking tools evolution

- Identified which Grid+IT Savannah projects need to be migrated to JIRA
- Message from GGUS team about support over Christmas break:
  - *GGUS is monitored by a system connected to the on-call service. In case of total GGUS unavailability the on-call engineer (OCE) at KIT will be informed and will take appropriate action. If GGUS is available but there is a problem with the workflow, e.g. ALARM to CERN doesn't generate email notification to the operators, then WLCG should submit an ALARM ticket, notifying Site DE-KIT, which triggers a phone call to the OCE*
  - if the web service is unavailable, the GOCDDB has KIT contact details

# Xrootd

- Monitoring status and plans
  - Detailed monitoring for dCache can coexist with SHA-2 compliance only for 2.6 and later
    - Priority should be given to SHA-2 (i.e. 2.2 is still OK)
  - Consolidate GLED collector system and infrastructure
    - FAX: 1 collector in US, 1 at CERN, 1 for EOS
    - AAA: 1 collector in US, 1 for EOS
  - Finalise merging of Dashboard federation monitoring and data popularity monitoring

# Machine/job features

- Recently started a collaboration with Igor Sfiligoi to try to converge on a mechanism to minimise draining waste of CPU time for multicore pilots
  - Based on bi-directional communication between pilot and resource



# IPv6

- Tested a simple recipe to enable IPv6 on SLC6 nodes
  - Will be extended to SLC5 as many experiment services still run on it
- CMS started testing CMSWEB on IPv6 with promising results and is running data transfer tests between IPv6 SEs since months
- ATLAS plans to start DDM transfer tests soon

# WMS decommissioning

- Experiment usage
  - Still used for a small fraction of CMS analysis
    - CMS users strongly encouraged to move to glidein submission
    - Usage is indeed decreasing
  - About 20 LHCb sites need to be moved to direct submission

# Other task force news

- Middleware readiness
  - The kick-off meeting will be tomorrow ([agenda](#))

# Experiment plans for the Christmas break

- ALICE
  - MC production at all sites
  - Investigate SLC6 vs. SLC5 job failure rate and CPU/wallclock efficiencies
- ATLAS
  - MC production: 130M events, ~10 days
  - Reprocessing: 2.2 PB of input, ~30 days from next week for 20% of the T1s
  - Group production: ~5 weeks on 35% of resources
- CMS
  - Production and digitisation-reconstruction of Run2 MC samples
  - Digitisation-reconstruction of 7 TeV MC for 2011 data
- LHCb
  - Mainly MC production
- All experiments are happy with the “best effort” level of support that sites usually provide in this period

# Conclusions

- New task force on multicore deployment approved
- Middleware readiness task force is starting its activities
- Sites are strongly encouraged to participate to both TFs
- The SL6 task force has completed achieving its goal
- Experiments will run during the end-of-year break with best effort support as usual