




# WLCG Operations F2F summary

Simone Campana IT/SDC

on behalf of WLCG Operations Coordination

12 February 2014



# Introduction

- Roughly 30 people in the room
  - +10 remote
- Good representation of experiments and site
  - Using pre-GDB slot helped, will do also in future
- Tried to favor discussion
  - With some presentations to initiate/drive it
  - Mixture of technical aspects and planning
- I will try to give highlights and summarize the main conclusions

# gLexec deployment

- gLexec CMS SAM test not “critical” yet
- Mostly due to scheduling issues of SAM tests at T1s
  - A dedicated discussion right after this presentation
- Still 20 sites did not deploy gLexec

# perfSONAR

- Look at Ian Bird slide 5 at the LHCONe workshop

➤ <http://indico.cern.ch/event/289679/>

- 20 sites w/o perfSONAR installed
- Another 20 with old releases
- Will be brought to the attention of MB and experiments

## perfSONAR deployment

- WLCG agreed on perfSONAR as the core toolkit for network monitoring in the infrastructure
  - Strong push came from experiments
- Deployment of perfSONAR has been (and still is) sometimes problematic
  - Some sites refuse to install it at all
  - Some sites still run very old versions
- perfSONAR needs to be treated as any other service in WLCG
  - Including the level of commitment in installing, configuring, operating it.

# Tracking Tools Evolution

- Some functionalities in JIRA for the Savannah-to-JIRA migration in GGUS still not there
- After many months of waiting, it has been agreed we can leave without such functionalities for GGUS
- Experiments still miss many functionalities in JIRA (not in the scope of our TF)

# SHA-2 migration

- Progress in providing VOMS-admin instance
  - Still, will need experiments adapting to it
  - In the meanwhile, new SHA-2 certificates need to be register in the old system “by hand”
    - Involving experiment VO admins or VOMS service managers
- In the next days, the campaign for distributing the new VOMS server host cert will start
- After the SHA-2 migration is over, we should tackle back the RFC proxy compliance

# Machine/Job Features

- Prototype ready for “bare metal”
  - Experiments testing it
- Different solutions being looked at for clouds
  - E.g. key-value stores
- Good collaboration with OSG for integrating with Igor’s proposal
  - Bidirectional machine/application communication

# Middleware Readiness

- Proposed model will rely on
  - Experiments and their testing framework (e.g. Hammercloud/SAM)
  - Sites, willing to deploy testing instances of some service
  - Monitoring, being able to distinguish testing from production
- A possible full chain will be prototyped starting with ATLAS
  - For some service and for clients
- At the next MB we will discuss possible ways of rewarding sites
  - Relevant also for IPV6 deployment discussion



# Enforcement of Baseline Versions

- Clients and services Baseline Versions in WLCG are defined but not enforced
  - Need first to monitor the release versions at sites
    - Use existing tools/systems when needed
  - Use the information to run ad-hoc campaigns of upgrade
  - Automate
- Leverage experience and operations in EGI/OSG
- Pakiti seems a useful complementary tool for the purpose
  - Need to understand deployment scenario and possible usage by WLCG operations

# WMS decommissioning

- A handful of users left pointing to CERN WMSes
  - Warned several times by experiment operations
- Shared and CMS WMS clusters will be shut down end of April
  - SAM WMS cluster at the end of June
- At this point the WMS will not be a WLCG supported service any longer
  - Including from the perspective of WLCG operations

# Multicore Deployment

- ATLAS and CMS foresee very different usage of MC resources
  - ATLAS: schedule a run a mixture of single/multi core jobs in batch system, 1 payload per job slot
  - CMS: schedule only multicore jobs in batch system , run mixture of single/multi core payloads in the same job slot
- We do not discuss if a model is better than the other
- We instead try various prototypes, optimizing resource usage in the two scenarios
  - Depend on batch system type, which VOs the site supports, etc ..
- Torque/Maui support is great concern
  - Discussed at next pre-GDB

# FTS3 deployment

- Deployment model has been discussed and got consensus
  - Limited number of FTS3 servers needed to handle full scale (potentially even 1)
  - Multiple servers useful for service resilience
- Deploy 3 to 4 servers in production
  - Share configuration, not state
- The 3 LHC VOs distribute their load among the servers according to their policy, *e.g.*
  - ATLAS, CMS round-robin between all servers
  - LHCb uses CERN FTS3 100%, with RAL as backup

# Experiment Computing Commissioning

- Experiments presented current plans for 2014 Computing Commissioning
  - Focused on experiments frameworks rather than WLCG infrastructure
- Conclusion: we do not need a Common Commissioning Exercise (like STEP09)
- WLCG ops will offer the needed support for commissioning activities
  - Several aspects already in hand of task forces
  - Will be followed up as regular WLCG ops items
  - Some concurrent testing for specific aspects would be beneficial
    - We will keep an eye to the schedule and understand if there is the opportunity

# Conclusions

- Many items have been discussed
  - We converged on a conclusion
  - We need formal agreement at the MB
- Some situations need attention of the MB as well
  - Especially concerning deployment