
# Criteria for Deploying gLite WMS and CE

Ian Bird
CERN IT

LCG MB
6th March 2007

# Introduction

- gLite WMS:
  - Effort to push into production readiness started July 2006
  - Performance-wise was ~OK for CSA06
  - But many ongoing issues of reliability and manageability – prevented from making this the production version and replacing the LCG-RB
  - Now, also work on gLite porting, and simplifying dependencies mean that CERN team cannot take the responsibility for driving the WMS improvements.
  - INFN have agreed that this should be their responsibility and that we will agree criteria for taking the WMS back into certification again.
- gLite CE:
  - Assume similar process and define CE criteria

# WMS performance

From discussions with CMS and ATLAS:

| | CMS | ATLAS |
|---|---|---|
| **Performance** | | |
| 2007 Dress rehearsals | Not specified but was 50K jobs/day in CSA06 | 20K successful jobs/day + analysis load |
| 2008 | 200K jobs/day through WMS<br><10 WMS | 100K jobs/day100K<br><10 WMS |
| **Stability** | | |
| | Not specified | <1 restart of WMS or LB every month (== LCG RB) |

The numbers for ALICE and LHCb are understood to be within these requirements

# LCG requirement

- Based on these numbers we propose the following as the **LCG requirements on the WMS:**

- Performance:
  - 2007 dress rehearsals: 50K successful jobs/day
  - 2008: 200K successful jobs/day using <10 WMS entry points

- Stability:
  - <1 restart of WMS or LB every month under this load

# gLite WMS criteria

- A single WMS machine should demonstrate submission rates of at least 10K jobs/day sustained over 5 days, during which time the WMS services including the L&B should not need to be restarted. This performance level should be reachable with both bulk and single job submission.
  - During this 5 day test the performance must not degrade significantly due to filling of internal queues, memory consumption, etc. i.e. the submission rate on day 5 should be the same as that on day 1.
- Proxy renewal must work at the 98% level: i.e. <2% of jobs should fail due to proxy renewal problems (the real failure rate should be less because jobs may be retried).
- The number of stale jobs after 5 days must be <1%.
- The L&B data and job states must be verified:
  - After a reasonable time after submission has ended, there should be no jobs in "transient" or "cancelled" states
  - If jobs are very short no jobs should stay in "running" state for more than a few hours
  - After proxy expires all jobs must be in a final state (Done-Success or Aborted)
- For verifying these criteria the test suite written by Andrea and currently used by Simone and Andrea will be taken as the baseline.

# gLite CE criteria

- Performance:
  - 2007 dress rehearsals:
    - 5000 simultaneous jobs per CE node.
    - 50 user/role/submission node combinations (Condor_C instances) per CE node
  - End 2007:
    - 5000 simultaneous jobs per CE node (assuming same machine as 2007, but expect this to improve)
    - 1 CE node should support an unlimited number of user/role/submission node combinations, from at least 10 VOs, up to the limit on the number of jobs. (might be achieved with 1 Condor_C per VO with user switching done by glexec in blah)
- Reliability:
  - Job failure rates due to CE in normal operation: < 0.5%; Job failures due to restart of CE services or CE reboot <0.5%.
  - 2007 dress rehearsals:
    - 5 days unattended running with performance on day 5 equivalent to that on day 1
  - End 2007:
    - 1 month unattended running without performance degradation

# Summary

- WMS:
  - Propose as LCG requirements – clear statement from CMS, but not from ATLAS (yet …)
  - Discussed with certification team, deployment testers, EIS testers, developers
- CE:
  - Propose these requirements as LCG requirements – based on LCG-CE and deployment experience
  - Discussed with certification team, deployment testers, and developers

- Expect to write similar document for LFC to clarify performance goals