



WMS & CE

Metrics and performance

Ian Bird

GDB

May 2nd 2007

Criteria



- At MB of 6th March agreed criteria for acceptance testing of gLite WMS and gLite CE:

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.

WMS status



- 2 weeks ago achieved:
 - 15,000 jobs per day for 7 days - 320 jobs (0.3%) did not run
 - These 320 were later restarted and succeeded - still should fix the underlying issue
 - The restart can be automated
 - All jobs in final states
 - Stress test is not yet complete - do not yet know the limit of scalability (30k jobs per day?)
 - Proxy renewal is working
 - L&B can scale much higher than WMS - not a limiting factor
 - N.B. to achieve the goals, DAG-based collections were removed and a new collection mechanism introduced
- WMS (3.1) now in standard certification process
 - Meanwhile gLite WMS nodes at CERN are updated to this version for production use
 - Next release with this version will allow full deployment and replacement of old LCG-RB

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

CE Status



- **Current status of gLite CE**
 - Close to 100% success of job submission - after resolving a number of timing issues with Condor
 - Submissions of 6000 jobs to a CE (max ~3000 at any time)
 - Do not know scaling of Condor-C instances yet
 - *Version with Condor-C instance per VO being set up - not tested yet*
 - Several Condor issues were found - not yet clear on a timescale for achieving the criteria
- **Fallback proposal**
 - Keep the LCG-CE "as-is" - there is no effort to port to SL4 (which implies GT4 and potentially many issues)
 - Deploy either on SL3 nodes (or SL4 with Xen/SL3)
- **Should set up a CREAM instance in parallel and subject it to the same testing procedure**