



Enabling Grids for E-science

Status of the WMS

Salvatore Monforte (INFN CT)

JRA1 All Hands Meeting 18-20 June 2007, Helsinki

www.eu-egee.org



- **Details about released patches and relevant bugs fixed**
- **Status of testing for certification**
- **Status of bulkmm implementation**

- **gLite WMS patch: #1140 released on 2007-04-24**
 - bug #**25653**: Match-Making should handle 'DENY' prefix in the ACBR of CE/Views
 - bug #**25767**: Purging does not work on collections
 - Added handling of the new EDG_WLL_STAT_COLLECTION tag provided by the LB for supporting the collections
 - bug #**25610**: Error in the template.sh on the JobWrapper
 - results in uberftp failing when the worker-node tries to submit the output back to the WMS when the job has finished.

- **gLite WMS patch: #1167 released on 2007-05-15**
 - *bug #23401: Job failure due to gethostbyname error in condor submission*
 - It was discovered that the job is not resubmitted and goes immediately to aborted state. Problem results from a sporadic network transient that randomly affects 1 condor submission every thousands...
 - *bug #26157: The WM dies while processing a collection with pending nodes*
 - Invalid sequence code used while creating the request inserted in the TQ for further processing results in RH throwing an uncaught LBCTX exception inducing all WM threads to terminate not properly.
 - *bug #26208: LM stopped on bad SizeFile object*
 - This is a side effect of other problems on condor/WM. Sometimes the same jobs is submitted two or more times so the counter inserted in the SizeFile object is not computed correctly.
 - **bug #22795: timer-log file are not removed by LM**

- *bug #26269: JC locks the filelist without giving to WM possibility to submit new requests*

- *bug #26267: The purger does not work: creation of lb context always fails*
 - Problem related to the use of hard-coded invalid-sequence code. removal of any hard-coded occurrence and agreed on the definition of relevant constant supplied with the LB API.

- *bug #26250: WMS client timeout approach does not work properly*

- **Most relevant bug fixed. Included in next tag**
 - *bug #26705: The standalone purger does not work anymore, due to the lack of proxy file.*
 - Due to the removal of NS from gLite 3.1 there was no longer any host proxy the purger refers to while working.
 - *bug #26913: The MM does not use information about previous matches retrying the same CEs*
 - This features erroneously disabled while performing some profiling memory leaks tracking.
 - *bug #26857: The "max-rank" selection algorithm for collection under some particular circumstances does not work properly*
 - If some nodes in the collection is assigned with negative or zero rank values, the standard selection algorithm SIGSEGV killing the WM.
 - *bug #25680: job submission --nodes-resource option does not work for dags/collections*

- *bug #26654: The ISM purchaser for BDII in some circumstances fails parsing VoView information*
 - Using ViewID containing “=” induce the purchaser to fail extracting the value from the dn entry

- *bug #26737: WMPProxy server hangs when managing signals*
 - WMP wrongly handles signals and may hangs while manipulating some SIG.

- *bug #26237: WMPProxy may register wrong DN to LBProxy*
 - when setting the DN of the User within the LBProxy context, the WMP may set it improperly.

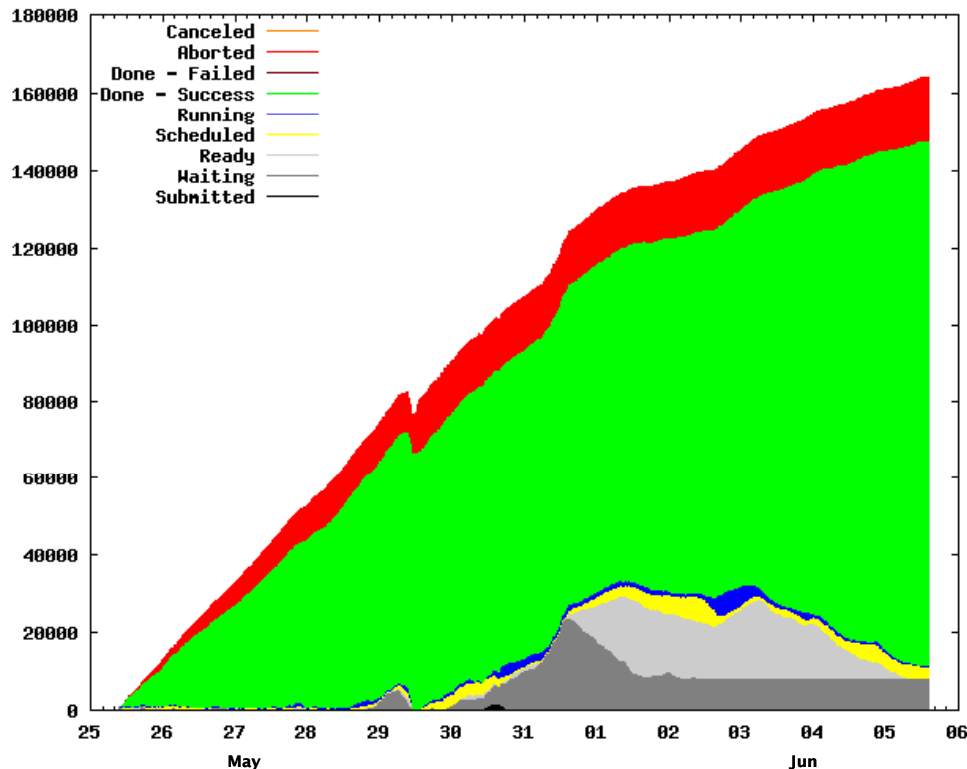
- **Experimental Services were introduced**
 - WMS/LB at CNAF, gLite CE at CERN
 - Instances of the services attached to the production infrastructure JRA1 patches are installed immediately (before the certification)
 - Maintained by SA1 and SA3
 - We planned regular phone conferences every two weeks
 - Testing done by selected application users
 - **ATLAS** (*Simone Campana*), **CMS** (*Vincenzo Miccio and Andrea Sciaba*) and **Alice** (*Patricia Mendez Lorenzo*)
 - Process controlled by the EMT
- This results in a rapid improvement of the **WMS** and **LB**
- Most **gLiteCE** problems ironed out

- **Currently two types of tests are performed:**
 - Submission of collection of 100 jobs
 - Submission of simple jobs with parallel threads

- **Target is to satisfy the acceptance test requirements**
 - 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
 - During this 5 day test the performance must not degrade significantly due to filling of internal queues, memory consumption, etc.
 - *The submission rate on day 5 should be the same as that on day 1*
 - The number of stale jobs (zombie) after 5 days must be <5%.

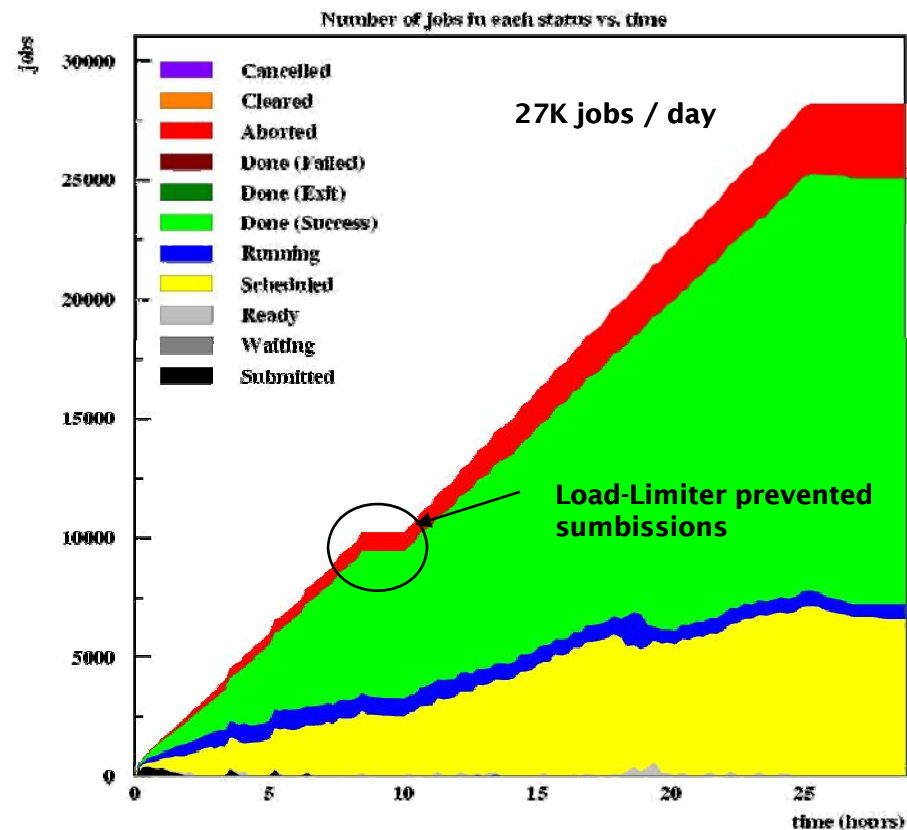
- WMS reaches the required quality of service:
 - **115.000 jobs submitted and processed in 7 days**
 - *a bit more than 16K jobs/day.*
 - **All jobs have been processed with the exception of 320 (0.28%)**

- **Stress test performed on WMS with JobDir enabled**
 - Submission of single jobs with parallel threads sustained for about 2 weeks
 - rates reached peaks of 21kjobs/day
 - WMS limiter disabled



- Probably due to the activity of some cron job started during the test we had a jam around May 31
 - jobs remained in "Waiting" status were processed
 - submitted in a time-range between May 29 and June 1
 - not all jobs submitted in that period stalled
 - The jam was recovered by the end of the first week of June

- **Stress test performed on WMS with JobDir enabled**
 - Submission of collections of 100 Jobs
 - rates reached peaks of 27kjobs/day
 - WMS load-limiter enabled



- **Bulk submission support is in place**
 - possibility to submit a bunch of jobs in one single interaction with the WMS
 - **(possibly) heterogeneous jobs resulting in a collection**
 - **homogeneous jobs modeled by a parametric job**
 - Jobs no longer transformed into a DAG and managed with Condor DAGMan, but handled directly by the WM
 - **Leads to a much smoother behavior**

- Bulk MM: natural completion of bulk submission
 - Match “equivalent” jobs in one shot, i.e. with one single MM operation
 - Jobs are clustered depending on their significant attributes
 - ***Two jobs are equivalent if their significant attributes match lexically***
 - **Significant attributes are specified by the user in the JDL**
 - *Requirements, Rank, FuzzyRank, InputData etc...*

- **Bulk submission limits and future works**
 - Although the submission of collections is in place the WMS does not provide
 - support for the cancellation of the whole collection
 - cancellation of a specific node of the collection
 - *Nevertheless in most cases the single job cancel operation performed on collection nodes works fine.*
 - Still investigating on a possible solution for the recovery of an unexpected bulk submission interruption
 - If the WM dies while processing a collection, the whole collection will be re-processed on service startup
 - *duplicated submission*