

ASGC T1 report

Computing

* SLC4 migration for WN in progress. 200 Nodes have now been upgraded and another 300 nodes are still remaining.

Castor2

* Currently srm 2.2 from castor-srm-1.0-13 is deployed but it depends on old castor2 schema. CERN will provide us with a newer version for installation that will work with the current castor2 DB schema.

* ASGC Castor2 provides different service classes depending on different DN used. We need to confirm with Atlas and CMS which DNs they will use for each type of activity so service class access can be configured properly.

FTS

* Testbed FTS 2.0 server deployed with channels configured

Events:

{{{

Title: sam job submission testing failure

Date: between 22-Jul-2007 21:16:26 and 23-Jul-2007 06:18:36.

Reason: root cause of sam job submission error. due to the confusion of new slc4 wn, that batch server fail to reset the pbs mom to be offline instead of free state, but most of the env need to be carried out from pbs head as well

as lack of know host files updated in ce/batch server. Due to this prob, SAM testing fail in 7 times since "22-Jul-2007 21:16:26", around the same time we restart pbs server.

Severity: totally 8 sam testing failure have been observed in sam monitoring, and all due to the ssh authentication failure, since new wn have new ssh pub key.

Solution: Removed SL4 WN from the batch system head node configuration to prevent further jobs to these nodes.

Details:

simple job submission validation fail due to replication of cache_export_dir, from wn back to batch head,

}}}

{{{

Title: ca pkg expired

Date: 17-Jul-2007 23:11:06 & 18-Jul-2007 00:06:46

Reason: this is due to the expired ca pkg release on wn w-wn0040 and have been fixed at Wed Jul 18 02:00 UTC 2007

Severity: the impact should be limited to this particular wn itself, only sam testing fail with critical testing error, and affect also site availability. the data transfer shouldnt have impact from current critical testing error, since we have sym links to /etc/grid-security/certificates

Solution: force upgrading the ca pkg with --justdb arg

}}}

{{{

Title: job submission failure - sam

Date: 17-Jul-2007 16:21:31, 17-Jul-2007 17:16:35 and 17-Jul-2007 18:21:19

Reason: this is due to the insufficient disk space from free wn, dedicated for monitoring jobs. after belle jobs crash on target wn, say w-wn0219, the temporary io in globus cache export have generated more than 6gb files. that all root have been fillup and not able to pass stagein files on it.

Severity: three sam testing failure. after cleanup the globus cache, we're able to pass generic validation via batch farm as well as from globus job submission with lcgpbs job manager.

Solution: remove the globus cache export dir in belle001 help freeing ~6gb disk space from root.

}}}

CERN-PROD

- Migration of SLC4/64 bit worker nodes to gLite 3.1 WN native build software completed. Several thousands of Grid jobs are seen now on SLC4. The majority of grid jobs still runs on SLC3.
- two LCG-style production CEs have been put into draining mode. They will be converted to submit to the SLC4 worker nodes as soon as all jobs which are served by them have finished. We expect that the migration of these nodes will be finished by the end of next week. More will follow if needed.
- We hope to deploy the latest Castor-2 software release 2.1.3-21 on the production instances in the coming weeks. This bugfix release addresses several important operational issues, in particular the problem of the stager cleaning daemon. On Wednesday we plan to upgrade CASTORALICE.
- SRMv2 endpoint at CNAF has been setup for the preproduction (PPS) service for ATLAS and LHCb

TRIUMF-LCG2

None

USCMS-FNAL-WC1

No problems at USCMS Tier-1 Site

Tier1 GridKA (FZK)

[author : Jos van Wezel]

time: 16/7 22:00 - 17/7 07:00

problem: SRM errors on transfers to FZK/GridKa

cause: known bug in dCache which causes storage paths to attain root permissions under certain conditions

severity: high usual tests do not necessarily pickup the error. this time the problem occurred on an ops vo path.

solution: chown

time:13/7 00:00 - 17/7 00:00

problem: tape access problematic

cause: SAN configuration changes and broken tape drives

severity: CMS had to temporarily halt their T0-T1 transfers to FZK/GridKa

solution: fixed drives and configuration

INFN-T1

Nothing significant to report during this week.

NIKHEF

No major operational issues this week.

SARA

Problem: No job submission due to high load on CE

Solution: restarted torwue on July 16th

Problem: The SAM host-cert-valid test failed on our srm endpoints. The certificates of both srm endpoints ant1 and srm are valid.

Solution: Debug the SAM host-cert-valid test.

See above.

See above.

Problem: autovacuum on postgresdb failed which underlies dcaches pnfs. This caused a shutdown of the postgres db resulting in failures of the lcg-rm tests.

Solution: Now we do a vacuum in a cron job.

NDGF-T1

Scheduled outage for hardware and software upgrade on srm.ndgf.org.

PIC

The most important issue concerning pic this week is the migration to SLC4. We have installed a new CE from scratch which points to a pbs queue with WN's running SLC4. This CE at the moment is configured to support ops and dteam. We have opened as well the access to the sgm users from the vo's cms, atlas and lhcb so that they can test the installation software. It has been really a tedious decision of what to include in the software directory for every VO. Different VO's have different needs. Finally we have decided to define a "brand new" software directory and not to publish any old tag from the other production CE's, nor copy old software. From the atlas point of view this approach is acceptable. They have already installed their latest version of the software in the ce-test and it seems it works fine. Instead from the cms point of view this is not efficient at all. We are still trying to reach a solution which will be ok for all the vo's, or at least for lhcb, cms and atlas.

RAL

Problem: LHCb SL4 jobs are not running correctly
Diagnosis: LHCb SL4 jobs which get submitted to our CE for testing SL4 WNs are not using much cpu time. Initial investigation indicates that problem lies on the LHCb side.
Resolution: None as yet, awaiting feedback from LHCb contact

Problem: An Atlas disk server in our dCache SE crashed with a Machine check exception
Diagnosis: Analysis of the report from the MCE suggested the memory was at fault
Resolution: New memory was received from the supplier and installed and the machine was returned to service

Problem: Alice unable to query GRIS of test CE for SL4 WNs from RAL-LCG2 VO BOX
Diagnosis: There is an iptables rule blocking access to the GRIS port from the Alice VO Box, however an identical rule is in place on our production CE and has never to our knowledge caused problems.
Resolution: The firewall rule denying access to the GRIS port was removed from the test CE.

Castor deployment status: Seperate instances are available for testing, The Atlas instance is being reorganised to match their requested service classes.
Oracle updates plus two hotfixes have been applied to the Castor databases.

IN2P3

- 2007/07/19

Around 17:00, All our SEs failed the SE-lcg-cp test with the strange error below:
>Host too long

```
>/opt/lcg/same/client/sensors/SE/tests/SE-lcg-cp: line 30: 15722 Segmentation fault
>lcg-cp -v --vo $SAME_VO $fn file:`pwd`/testFile.txt
No action has been taken because it seems that it came from SAM tests (most of the T1
were failing the test at same time)
```