

ASGC T1

Castor2:

- * upgrade of Castor services, memory replacement, LSF upgrade are complete
- * currently testing newly created storage classes for Atlas and CMS. Each VO will have D0T1 and D1T0 storage classes available.
- * still need to migrate DB storage to dedicated hardware and migrate CMS data back to Castor2

Network:

6/29 10:09 UTC TPE-AMS 2.5 link is disabled which will be replaced by 10G link in July

- * Current traffic is rerouted via TPE-CHI-AMS 2.5 link in the mean time.

CMS

- * ntpd configuration error on Vobox fixed
- * CMS phedex transfer error on June 26
- * This was fixed by adding CMS as the default acl group to TW-FTT DPM server June 26 21:00 UTC

Events:

```
{
{
{
Title: NFS problem on quanta.grid.sinica.edu.tw CE
Date: June 26
Reason: rpciod failure
Severity: caused soft experiment directory test to fail, but this did
not cause failure of jobs
Solution: fixed by removing unnecessary nfs mount pt and restarting nfs
daemon
Details: none
}}
}}
```

FNAL

Nothing to report

CERN

Fabric Services

=====

-- The new SLC4 native build worker node software is now being tested. The release notes of this version (which is gLite 3.1 based) recommend a reinstallation of the worker node when doing an update of the previous version. At CERN that would imply draining and reinstalling the whole capacity which is available on SLC4/64, which may require several week because we have long running jobs there, and the upgrade would have to be done in a rolling way, working only on a subset of machines at a time.

We have a pressing request from some VOs to release SLC4 for the grid, and therefore we are preparing a work-around, which should allow an update of the worker nodes without having to put them out

of production.

The workaround implies a manual intervention: the old WN software needs to be uninstalled, any remnants of it removed by hand, and then the upgrade can take place. This procedure is now being tested, and we plan to do the update of the SLC4 worker nodes before the weekend, so that the experiments can do tests of the new software via the preproduction CE nodes (ce110 and cd111). If everything works as expected we plan release two production CEs submitting to SLC4/64 by the end of next week. These CEs will be ce109 and ce112 which are currently configured as gLite CEs.

WLCG Transfer Service:

- * Transfer ranging from 20 to 920 MB/s, averaging around 500MB/s per day.
- * Involving all major T1 sites.
- * Mostly traffic from Atlas, less from CMS.
- * 2 tickets have been submitted to sites this week for different problems, 4 tickets have been moved from last week, 1 problem has been solved, and 2 tickets have been closed.
- * Throughput plots:
<http://gridview.cern.ch/GRIDVIEW/>

CERN Grid Operations:

- * One CMS user sent jobs on the 3 CMS LCG RBs, with a bad JDL (2 anyMatch clauses in it), causing the workload manager service to crash again and again on these machines.
- * Creation of two new subclusters gridwms and gridlb (belonging to a new cluster grid) which will contains the gLite WMS 3.1 and gLite LB 3.1 respectively. These new clusters have been created in collaboration with FIO.
- * The Atlas VOBOX still managed by GD has been removed from production.

GridView :

- * The new schema for GOCDB, the GOCDB3 is deployed in production, so we have deployed the new Synchronization module to synchronize the topology data in Gridview from GOCDB3 database.
- * Testing of new continuous scale service availability summarization module.

Physics Database Services

All the Tier1 sites have been connected to the CERN streams set-up and are now in full production phase. Procedures for maintenance and support of the streams set-up are being spread over the DBA team.

A review of the FTS deployment procedures has happened yesterday in the 3D meeting between the Tier0 and Tier1 database administrators and the FTS developers. A document summarizing the discussions will soon be made available on the 3D wiki pages.

IN2P3

No report

Report for Tier1 GridKA (FZK):

instabilities of CE. (infosystem) Under investigation
instabilities of CE. (infosystem) Under investigation
instabilities of CE. (infosystem) Under investigation
Scheduled downtime
Scheduled downtime

NDGF T1

A Finnish colleague suggested the following report: "Dear WLCG, everything is fine. Best regards". That pretty much sums up the NDGF operations week. T0 to T1 tests were running smoothly, too.

NIKHEF

Recently we saw frequent failures of SAM's replica management tests. Most of them failed on a time out. We could not find indications of problems at the site, nor did the sub-tests take excessive amounts of time. Therefore, we assume that the source of the problems is somewhere else.

This kind of error is not conclusive about its cause. We already have a tendency to ignore such failing tests, assuming that the problem lies somewhere else. That raises the question whether this test is useful at all (in particular triggering an error on a time out).

SARA

There were no serious issues this week.

PIC Tier1 availability report:

Date: From 22/06/2007 evening until 27/06/2007 evening.

Problem: The FTS agents failed on 22nd June, for unknown reasons. Trying to solve them, an update of the system was issued, which ended up in installation of the FTS-2.0. We understand this was a mistake of the FTS deployment, since the FTS-2.0 rpms should not have been put in the production repository.

Severity: High. The PIC FTS service was down during all the period. Luckily, the experiments were not transferring files between T1 and T2s during those days. The only activity was CERN-PIC (from CMS, and not very intense) which is driven by CERN FTS.

Solution: The downgrade was extremely painful due to the way the meta-rpms are built. With the help of fts-support@cern, the downgrade could be finally done manually and the service restored with the original version.

--

Date: 24/06/2007

Problem: Two CE-sft-lcg-rm spurious errors at 00h and 08h due to one misconfigured WN pointing to a default_SE recently migrated, with the Information Provider still to be configured.

Severity: Low. The problem appeared in only one of the WNs of the farm.

Solution: The bad WN was detected and correctly configured on Monday 25 June at 11h.

--

Date: From 28/6/2007 at 19:00 until 29/06/2007 at 12:00

Problem: The CE service fails Job Submission due to a human error when changing the batch scheduler configuration for testing one WN with SL4. The exact problem was with a wrong maui reservation that was aimed to assign a test WN to a test queue. Adding a GROUPLIST tag to this maui reservation, resulted in a wrong behaviour.

Severity: Low. This configuration problem was only affecting the OPS and DTEAM vos, since it was a test. Job submission for ATLAS, CMS and LHCb continued with no problems.

Solution: The problem was detected and solved by deleting the reservation, and substituting it by associating nodes to queues using a Torque directive "resources_default.neednodes".

--

Date: From 28/6/2007 at 19:00 until 29/06/2007 at 12:00

Problem: SRM-disk service fails SAM tests because the reconfig of the info provider publishes the hostname (dcsrm02) by default, and not the alias of the service.

Severity: Medium. We see the FTS transfers CMS and ATLAS continued, since these transfers seem not to depend on the info published.

Solution: Solved by hardcoding the alias in the static info published by the dCache Information Provider.

--

Comments:

-Next week we plan to upgrade in all our production nodes the new yaim including the patch for the removal of the default configuration of the VOViews.

RAL

We are deploying a new Resource Broker for Alice to hopefully alleviate the issues reported last week. We are also deploying a new CE to frontend a trial SL4 worker node service. One of our top-level BDII's was installed with the BDII with indexing and we noticed an immediate reduction in system load, the update will be rolled out to the others shortly. CMS are load testing their new Castor instance. An instance for Atlas is being installed.