ASGC T1 report:

Computing
 * SLC4 migration for WN still in progress

Castor2
 * Atlas T0 export testing
  * Atlas uses castor1 end point instead of the new castor2 endpoint, we have inform ddm op and received replies from Alexei and Miguel on July 12.
  * testing for replicating castor1 to castor2 atlas data is successful. Data migration away from castor1 will begin after above problem is resolved. Castor1 will be phased out afterwards.
 * CMS
  * a few zero size files found for CMS. Currently investigating possible causes for this problem.
 * Current disk and tape pools status
  * Disk Pool :
   {{{
POOL atlasPrdD0T1 CAPACITY 30.01T FREE 30.00T(99%)
POOL atlasPrdD1T0 CAPACITY 30.01T FREE 30.01T(99%)
POOL cmsCSAD0T1 CAPACITY 42.97T FREE 42.61T(99%)
POOL cmsLTD0T1 CAPACITY 28.64T FREE 28.64T(99%)
POOL cmsPrdD0T1 CAPACITY 44.33T FREE 38.19T(86%)
POOL cmsPrdD1T0 CAPACITY 42.97T FREE 41.03T(95%)
   }}}
  * Tape Pool:
   {{{
atlasPrdtp - - CAPACITY 97.66TB FREE 97.65TB (100.0%)
cmsCSAtp - - CAPACITY 68.36TB FREE 68.36TB (100.0%)
cmsLTtp - - CAPACITY 9.77TB FREE 9.77TB (100.0%)
cmsPrdtp - - CAPACITY 97.66TB FREE 84.84TB (86.9%)
   }}}

Networking
 * 10G deployment from TPE to AMS is complete and router is also upgraded to M320.
 * Discussion is underway with Surfnet to establish 10G link between AMS and CERN to complete the 10G circuit

BNL

None
Problem: BNL dCache/SRM time out errors continued.
Problem: BNL dCache/SRM time out errors continued.

Problem: HPSS experienced an ongoing connectivity problem since Saturday July/07/2007.

Cause: A stuck tape drive caused HPSS performance to deteriorate.

Severity: HPSS performance was affected.

Solution: Restart HPSS on 11:30AM, and resume HPSS service around 12:00noon.

Problem:
dCache time out errors continued. SAMS continues report SRM/dCache unavailable errors.

Problem: SAMS reports that BNL dCache/SRM/SE have time-out errors after 600 second even BNL dCache load is moderate.

Cause: Still under investigation.

Severity: Users and SAM tests experienced time our errors when they copied files into/out of BNL dCache.

Solution: No solution was found yet. The OSG dCache/SRM supporting team was notified about the problem.

Maintenance:
dCache SRM server was restarted around 11:00AM after several SRM parameters were changed.
Maintenance:
30% computing nodes were powered down at 9:00AM because of air conditioning maintenance.

Impact: 40% of computing capacity is not available to users.

Recovery time: All servers were back on-line at 5:00PM.

Improvements:
Some of AOD data files belonging to BNLDISK area were not in disk storage. It caused analysis jobs to fail because of long staging time.

dCache read pools belonging to BNLDISK area were reconfigured to be read-only, which does not accept pool-to-pool transfer and stage-in. It prevents AOD files from being swapping out to make disk space for the newly written or staged-in files.

Maintenance:
A dedicated disk pools with 20 TB were allocated for the USATLAS produced AOD/NTUPLE data files.

Problem: LCG VO box, MyProxy server, and LCG BDII server were off-line since 9:00AM.

Cause: we accidentally powered off these three servers when we shut down farm nodes for maintenance.

Severity: we lost the critical services for one hour.

Solution: the three servers were powered back around 10:00AM.

Problem:
dCache time out errors continued.

Cause: Under investigation.

Solution: The problem was reported to the OSG Storage supporting coordinator who confirmed that this is a general problem to many dCache sites, and forwarded the problem to SRM developers.

Problem: BNL dCache/SRM time out errors continued.

Maintenance:
On Wednesday, July 11 at 8:30 am, the cooling system at the RACF will be partially shutdown for maintenance work. About 30% of the cooling capacity will be unavailable for an estimated 3-4 hours. To cope with the temporary loss of cooling, the RACF will close about 30% of its nodes to Condor jobs at 4 pm today (Tuesday, July 10) to prepare for shutdown of these nodes.

Impact: 40% of computing capacity is not available to users.

Recovery time: July/11/2007.

CERN-PROD T0

Networking services

The AUG test in LHC point 8 were executed last Monday morning between 9:00 and 12:00. Some issues related to the powering of the network and GSM equipment in the sector need to be reviewed in collaboration with TS/EL. An LHC wide upgrade of the UPS powering of the network infrastructure is under investigation by TS/EL.

Due to misconfiguration (use of the wrong Ethernet protocol) of PLC"""""s (SCHNEIDER) from TS/CV several interventions were required by the First Line support and network experts to get the problem solved. Users are

reminded that only Ethernet 802.3 is supported on CERN networks.

The starpoint upgrade project has reached an other milestone last week. Building 40 has been completely upgraded to the new infrastructure (switch and backbone). Regretfully hardware failures with the new equipment required the replacement of several switches. The supplier is investigating the issue.

The "First line support" workload statistics for JUNE 2007 show a significant peak in activity. These can be explained after an extremely quit month of May, increased LHC activities and the arrival of new people like summer students.

CASTOR databases

* On Friday the CASTOR CMS stager suffered a high load, this was traced with the CASTOR team to a statement having a new execution plan. Castor developers have added a hint (comment to the optimiser) to indicate to the Oracle optimiser that another execution plan should be used.

* On Sunday, two issues were worked on with the CASTOR team for CASTOR public stager: an execution plan change (which same solution as for the CASTOR CMS issue mentioned above) and an investigation to identify the reason for a statement becoming less efficient (data distribution change).

* Following a security incident in the Castor Public stager database, we are limiting the number of the machines that can access to the Oracle Listener port. We are also investigating how to use SSL certificates to replace the current username/password authentication methods.

* All the old (before the upgrade to 2.1.3-x) Castor DLF (atlas,cms...) data has been moved to a new server and is ready for use if required.

* Problems in the Castor Name Server due to uncommitted transactions that block other sessions. This problem affects the Castor Stagers.

* All Castor databases have been moved to TSM21 (double copy TSM system)

Fabric Services

- Our gLite production CEs, ce109 and ce112 have been retired and reinstalled as LCG CEs, submitting to our production nodes in SLC4/64. With the release of these CEs, our SLC4 resources are now also available for Grid jobs, which corresponds roughly to a doublication of the available resources.

- on request of ATLAS a 32bit python version has been made available

on the SLC4/64 bit worker nodes. It is available in the standard search path as "python32" on all 64bit worker nodes.

- after the upgrade of the worker node software, a few remaining sick machines have been identified and repaired.

- the upload of the APEL accounting data for June went wrong due to the large amount of data, which caused our rgma node to run out of memory. Attempts to upload the missing data by hand in time failed due to problem with aggregation of the data on the receiving site, followed by problems with the registration server. Due to these problems we were not able to deliver all required data in time, and had to correct the accounting information for June.

- the CA rpms have been updated to the latest version

- Several machines in the Computer Center have been identified, which were missing the latest updates. Users of critical machines have been informed and warned, and a notification for synchronisation was sent on Thursday to clean up this situation

- Prod-bdii upgraded to the new schema

- more LXBATCH SLC3 capacity being migrated to SLC4:

Current status is:

- SLC3: 2650 Ksi2K on public resources from which ~900 can be upgraded to SLC4.

The upgrade will be done little by little in function of the load.
See [8]http://lsfmon.cern.ch/lrf-lsf/queues.php for the current usage of SLC3, SLC4

- SLC4: 4880 Ksi2K on public resources
~1000 Ksi2K that will be dedicated to the CMS pre-CSA07 next week

- Castor operations

-- a problem with Castor"""""s "cleaning daemon" caused accumulations in the stager database SUBREQUEST table. On two instances (Castorcms on Friday, and Castorpublic on Saturday) this led to a change of Oracle execution plan, which unfortunately slowed down the queries, and led to time-outs in the user access to the services.

The problem on Castorpublic could only be solved on Sunday afternoon.
The Compass data acquisition suffered most, as their CDR disk buffer had filled up by 14:00. CDR was stopped for several hours.

The problem of the cleaning daemon is being addressed in the upcoming bug fix release. Until then, we will manually clean the SUBREQUEST tables.

-- On three occasions a blocked thread to the Castor nameserver locked the database. This caused time-outs in calls to the nameserver, degrading services. Investigations into the cause of the problems are ongoing.

We are adding a 4th machine to the loadbalanced DNS alias that serves the nameserver request, although we are not sure that the problem is load-related...

-- We have deployed new passwords for all the Castor stager and dlf databases, using SINDES-based mechanisms.

-- Another 30 diskservers have been reinstalled with SLC4.

LCG deployment:
WLCG Transfer Service:
* Transfer ranging from 20 to 540 MB/s, averaging around 300MB/s per day.
* Involving all major T1 sites.
* Mostly traffic from CMS.
* 1 open ticket in total
* Throughput plots: [13]http://gridview.cern.ch/GRIDVIEW/

Physics Database Services

Backup tape storage for the compass database cluster has been moved from TSM51 to TSM22 - thanks to FIO. The change was necessary as after recent TSM51 instabilities some corrupted and missing backup sets were spotted by our test recoveries.

The setup of LFC back-end database replicas for LHCb has been completed as planned between CERN, RAL, GridKA, IN2P3 and PIC. The intervention has been performed online without stopping the database or LFC services at any of the participating sites. The last LFC db replica at SARA will follow soon, as the recent instabilities of the SARA database cluster have been removed by a recent network and storage reconfiguration.

An urgent data recovery in the ATLAS production database has be performed successfully following an erroneous "drop table" command

issued by an user. The affected application was the ATLAS
distributed data management system.

LCG software:
---
new s/w packages installed in the LCG afs area
(/afs/cern.ch/sw/lcg/external) for all LCG platforms

kcachegrind 0.4.6
frontier_client 2.7.4
sqlite 3.4.0
dcache_client 1.7.0.35
oracle 10.2.0.3-full (includes extra char set library)
mysql_python 1.2.2

Report for Tier1 GridKA (FZK):

[author : Jos van Wezel]

06/07/2007: The tape connection to dcache is problematic this week. A broadcast was send to notify users.

INFN-T1

During this week we upgraded:
* Castor latest version (a downtime was scheduled)
* glite software to update 28
* CAs to version 15

Minor problems were reported and immediately solved. Site is running ok now

NIKHEF

Upgrade of the top level BDIIs with the new bdii release and Glue schema. No problems encountered.
Upgrade of the rest of the services with gLite update 28. The upgrade of the database schema for DPM failed. The upgrade script was aborted (by Yaim?) after 1800 seconds, leaving a partly upgraded database. Consequently, both the Yaim script and the DPM upgrade script could not recover this situation and manual action was required.
The schema upgrade failed because of an inefficient implementation of the upgrade script, which performed one database transaction per row for each of the > 3.7 million rows in the table. Instead, one single update could have been used. We have informed the DPM developers of this problem and its solution and we recommended them to test such scripts with large databases before releasing them.

PIC

No major problems this week. We have finished to define the quattor configuration in order to have WN"s running SLC4 with glite 3.1. Our CE ce-test.pic.es has a queue specially for this purpose. We found a bug in the globus configuration in the quattor profile and now the WN are properly configured. Next week hopefully all the WN will be migrated to SLC4. For the moment this CE is opened just to ops and dteam VO, but we have contacted 2 members of atlas and cms to make them test the queue and the WN configuration.

RAL

SL4 testing continues, we have now reserved 2 job slots per VO (out of 14) after a large influx of Atlas work to the test CE. Hardware is now being obtained for a production CE for SL4.

CMS and Atlas are continuing to test their Castor instances, LHCb"s is now ready for testing.