

ASGC T1

SAM events:

- * June 11
- * RM failure due to Information system problem on DPM
- * quanta CE failed, however dteam test and samap test passed.

CMS

Squid upgrade to 3.0rc2, including also the migration of squid server. (shuttle to dual core 1U server)

TRIUMF

FTS load high. Known oracle problem requires upgrade. Restart oracle fixes it.
Test 1 node with 5 slots on 4 cores.
LFC swap space filling for 1 week. Restart fixes it but needs watching.
1 node offline for benchmarking
vobox manual patch of ATLAS DDM to enable priorities

BNL T1 report

Monday: June/11/2007

Problem:

One of GridFtp servers node (dcdoor02) crashed on 6:00AM Monday morning.

Cause: an IRQ kernel error.

Severity: the GridFtp server was off-line for four hours and 14% connectivity was lost due to this problem.

Solution:

Our system administrator updated the server since it hadn't been updated in many months. Then, after the update, we rebooted it a second time so it would be running the updated kernel.

Problem:

Two OSG gatekeepers were reported critical at 04:54:56 EDT 2007

Cause:

GUMS server went off-line (but log output had stopped. we restarted Tomcat and it appears to be functioning). The new host certificate generated for the GUMS server was missing an attribute needed for a server. The old one had

X509v3 extensions:
Netscape Cert Type:
SSL Client, SSL Server

while the new one had:

X509v3 extensions:
Netscape Cert Type:
SSL Client, S/MIME

Severity:

Two OSG gatekeepers were impacted for three hours before our administrator intervention.

Solution:

We disabled the GUMS on two OSG gatekeepers and used the static Grid map files to allow BNL CE to be accessible right away. In the mean time, we obtained a new host certificate with a proper server attribute for our GUMS server.

Tuesday: June/12/2007

Maintenance:

Kernel update was performed on all seven GridFtp server nodes one by one. Each GridFtp server has less than an hour downtime. The update is transparent to users. No need to make announcement.

Wednesday: June/13/2007

Problem:

A fraction of data transfers from BNL to other ATLAS Tier 1 sites failed with certificate mismatch errors.

Cause:

A fraction of our dCache read pool nodes have bad certificates that their DNs do not match with their hostnames. The error message is shown as follows:

```
-----  
06/07 12:01:04 Cell(SRM-dcsrcm@srm-dcsrcmDomain) : Authentication failed.  
Caused by GSSException: Operation unauthorized (Mechanism level:  
[JGLOBUS-56] Authorization failed. Expected  
\"/CN=host/acas0203.usatlas.bnl.gov\" target but received  
\"/DC=org/DC=doegrids/OU=Services/CN=acas0399.usatlas.bnl.gov\")
```

Severity: This problem affects the data transfer directly between the remote party and these affected dCache read pool nodes while the data transfer via (GridFtp servers) was not affected. BNL to other Tier 1 data transfer does not use GridFtp server nodes; all data transfer from BNL affected nodes to Tier 1 sites experienced transfer failures. USATLAS production is NOT affected by this problem.

Solution:

We replaced these bad certificates on Wednesday. We notified the ATLAS data operation team to confirm whether the lower performance problem with the data transfers from other Tier 1 sites to BNL. We will add scripts to validate the host certificates.

Thursday: June/14/2007

None.

IN2P3

The top-bdii we use (topbdii.grif.fr) had 2 small interruptions during the week, provoking the failure of the CE-sft-lcg-rm tests.

Report for Tier1 GridKA (FZK):

[author : Jos van Wezel]

A week filled with different problems. Possibly because the number of jobs issued against FZK GridKa raised sharply last weeks, both CE hosts reached their operating limits and many tests failed. We are working on a solution but it will take at least another week before stability is expected to return. The SRM instabilities continued. Restarting is the applied solution. Additionally a corruption with unknown cause in the LFC MySQL DB forced us to take the LFC unscheduled offline for several hours, impacting all users but presumably mostly Atlas. Lastly instabilities in the info system results in timeouts and erroneous returns. It is being investigated. Suggestions how to tackle this problem are welcome.

Announcement:

FZK/GridKa is unreachable on 27/6 from 05:00 UTC to 18:00 UTC. Network connections are restored after 18:00 UTC but maintenance will continue till 28/6 18:00 UTC. Services are impacted during the whole period of 27/6 05:00 UTC till 28/6 18:00 UTC.

NIKEF

We reduced the maximum CPU and wall time for all queues to 24h and 36h, respectively. Recently we observed more and more \"pilot\" jobs by HEP VOs that apparently claim the CPU until they run out of time. Since each of these VOs may claim ~80% of the available CPUs, the estimated response time for all jobs has increased a lot. This reduction was performed after consulting various people in the HEP community. We will observe the system for problems.

SARA

Problem: Problems with site bdii due to a misconfiguration.

Solution: misconfiguration has been corrected

A lot of down time this week because of the sgm poolaccount issue and problems with the node running the SAM tests. This downtime has only affected the ops VO but not the other VOs who where able to do their work as usual.

We have had problems on wednesday with the oracle database underlying the LFC for ATLAS and the FTS. This database was in a inconsistent state. This was solved by dropping the constraints and indices and rebuilding them again and dropping and creating the Cns_symlinks table. We have also hit a bug in the LFC which caused the fact that symlinks could not be deleted. This will be fixed in a next release.

LHCB had write permission problems on the SRM. This has been fixed by changing the ownership of all files and directories from user lhcb to user lhcbprod.

One dCache pool node had a broken disk which caused a file system problem which should not happen on a RAID6 system. We are still investigating this and looking if we can rescue the data.

Problem: Srm problems.

Solution: Service which was very slow has been restarted.

Problem: Problems with the oracle LFC which was very slow.

Solution: Restarted the service with an increased number of threads which fixed the problem.