

BNL-LCG2

Friday:

We continue to observe sluggish Tier 0 data export from CERN to several Tier 1 sites. Our PNFS should high load generated by USATLAS production and Tier 0 data exports. We saw many time-out in data transfer.

Saturday, June/02/2007

Problem: SRM server in BNL was off-line between 4:00PM and 7:00PM.

Cause: The SRM server, GridFtp door server, and write pool server certificates expired around 4:00PM.

Severity: the data transfer completely stopped during 4:00PM and 7:00PM. Tier 0 data export to BNL, USATLAS Production and AOD/NTUPLE data replication were impacted.

Solution: We renewed the certificates around 6:00PM.

Problem: USATLAS production has problems to write data into BNL dCache system at 8:48PM. No data can be written into the subdirectories of the dCache root directory: /pnfs/usatlas.bnl.gov/AOD01, RDO01, etc.

Cause:

We changed the ownership of directories to the production account ???usatlas1???. But did not attach storage tag, which means that there were no write pool resources assigned to these sub-directories. The production could not create subdirectories and write files.

Severity: The USATLAS production was affected for two hours.

Solution: attach the storage tag into subdirectories, as we agreed on the morning meetings.

Problem: The problem happened on Wednesday continued. USATLAS production manager reported the performance of the data export from BNL to some Tier 2 sites was sluggish.

Cause: Production requests data files that are only in HPSS, but not in disk areas. It takes long time to stage-in files into disks and transfer them to the remote sites.

Severity: many Tier 2 sites are running out of data files and waiting for data transfer since Tuesday morning. The resource utilization for USATLAS decreased.

Solution:

We reallocated 10 HPSS drives to speed up the data stage-in. The USATLAS production manager gave us a list of files. Our data management expert created a script to stage-in the input files from HPSS to disks.

Sunday: June/03/2007.

Problem:

The Panda monitor (USATLAS production monitoring server) crashed. The kernel crashed from a stack overflow (do_IRQ: stack overflow: 488), probably because it was not able to keep up with the increasing memory pressure.

Cause: from the message that is still visible on the console and the system logs, it looks like that apache was using too much memory, triggering the kernel's out-of-memory killer several times in the hour prior to its crashing:

```
Jun  3 00:12:33 gridui02 kernel: Out of Memory: Killed process 27405 (httpd).
```

```
Jun  3 00:29:04 gridui02 kernel: Out of Memory: Killed process 28348
```

there could be a memory leak in the panda monitor or the version of apache it is using.

Severity:

USATLAS production dashboard is off air for twelve hours. Production runs blindfolded.

Solution: Reboot the server, and add a memory warning watermark in BNL Nagios monitoring page.

Problem: USATLAS production manager reported the performance of the data export from BNL to some Tier 2 sites was sluggish.

Cause: Production requests data files that are only in HPSS, but not in disk areas. It takes long time to stage-in files into disks and transfer them to the remote sites.

Severity: many Tier 2 sites are running out of data files and waiting for data transfer since Tuesday morning. The resource utilization for USATLAS decreased.

Solution:No Solution was given yet.

problem: data export from Tier 0 and USATLAS production data transfer suddenly failed around 1:00PM.

Cause:

The dCache map file is updated around 1:00PM. Both update scripts and SRM read/write from/to the same dCache grid user map file. SRM reads only a fraction of the

file around 1:00PM, and two accounts are missing from the file: USATLAS1 and USATLAS4

Severity: Both USATLAS production and Tier 0 data export have been off-line for twenty minutes.

Solution:

Our dCache administrators redirected the output of the dCache grid map file updating script to a temporary file. After the update is finished and validated, then the script swaps the new grid map file with the existing dCache grid map file.

FNAL

No report

CERN

No report

INFN

No report

Triumf-lcg2

New CA rpm installed.

IN2P3

No report

GridKa

No report

NIKHEF

No major operational issues this week.

NDGF-T1

Our FTS server ran out of tablespace, it took a while to get this extended. T1-T1 transfers were halted during this time.

PIC

No report

RAL-LCG2

A new subcluster and queue has been added to our CE for SL4, currently access is limited to dteam and there is only 1 worker node installed with the SL3 on SL4 WN install, we will be installing a native SL4 WN next week

Maintenance of RB completed successfully.

Testing of FTS2 instance continues. No problems observed yet.