

Report sent on June 15th to: wlcg-scod@cern.ch

Type of Incident: Outage of import / export of LHC Data

Location: IN2P3-CC

Duration: 7 hours

Date: June 10th 2009 03:00 – June 10th 10:30

Description: Local GridFTP servers unavailable

Impact: SRM file transfers failure

Timeline of the Incident

03:30 SAM tests of the SRM service start failing

08:30 Dcache master notified and problem investigation

09:50 LHC-alarm ticket received

https://gus.fzk.de/ws/ticket_info.php?ticket=49392

10:00 Attempt to publish a downtimes announcement.

10:30 Service SRM is restarted and up.

Announcements to end-users

1 news posted on to <http://cc.in2p3.fr> at 11:00 on June 10th.

Broadcast and mail to wlcg-operations@cern.ch at 14:00.

Analysis and Follow-up

During the night, since 3:00 AM, all FTS transfers and file copy commands using the SRM protocol were failing. Many “production” jobs running for ATLAS and CMS have failed.

In the timeline 8:30-10:00 the problem has been understood.

Here is a short report of what happened :

The problem occurred when all GridFTP servers were overloaded. The monitoring script tried to get a connection to each of them unsuccessfully. It turns out a bug in the monitoring script has disabled all GridFTP doors. Consequently, GridFTP servers were unavailable and all requests failed.

In a while, this has been detected by the SAM tests which are only linked to the dCache monitoring dashboard and not to the local monitoring system in which case the oncall engineer should have been informed .

The servers have been re-enabled and up since 10h30.

Right after the incident, dCache masters corrected this monitoring script so that the GridFTP doors cannot be disabled all at the same time, ever again. Given the problem type, at short term it is not envisioned to have any specific modifications to our local monitoring system but at mid term we are going to improve it.

Report sent on June 15th to: wlcg-scod@cern.ch

A downtime announcement has been tentatively done on-to GOCDB portal – in failover mode then; but a bug prevented us to publish the “downtime”.

A ticket has been posted about this :

https://gus.fzk.de/ws/ticket_info.php?ticket=49383

However a broadcast has been sent through the CIC portal to the LHC community, later on 14:00 PM:

[https://cic.gridops.org/index.php?](https://cic.gridops.org/index.php?section=rc&page=broadcastretrieval&step=2&typeb=C&idbroadcast=41103)

[section=rc&page=broadcastretrieval&step=2&typeb=C&idbroadcast=41103](https://cic.gridops.org/index.php?section=rc&page=broadcastretrieval&step=2&typeb=C&idbroadcast=41103)

Next time, we will take utmost attention to communicate at least through broadcasts sent out to the community earlier in the incident treatment process.