

Service Incident Report for PIC ATLAS Storage Service failure July 21-22

Incident Start: 21st July 2010 22:00 UTC

Incident End: 22nd July 2010 around 08:30 UTC

Description

All active ATLAS dCache pools were full of unremovable data (disk only) so new ATLAS write operations failed with *no space available* error.

Time line of the incident

Incident started to happen on 21st July 2010:

- 21/07/10 22:00 - all ATLAS write attempts start failing with *no space available* or *pool cost too high* like errors.
- 21/07/10 22:05 - ATLAS specific SAM test SRMv2-ATLAS-lcg-cr starts failing. This VO-specific SAM alarm does not trigger a local alarm on the Nagios system at PIC.
- 21/07/10 23:19 - ATLAS shifter opens GGUS Team ticket on PIC reporting transfer problems (GGUS [60349](#))
- 22/07/10 08:00 - PIC's storage expert is notified about the issue
- 22/07/10 08:12 - ATLAS escalates ticket to ALARM ticket (GGUS [60359](#))
- 22/07/10 08:15 - the root cause of the error is that all active ATLAS dCache pools got full of unremovable data (disk only data). Manual triggered data migration process started to ATLAS dCache pools with removable space (disk/tape cached files).
- 22/07/10 08:20 - Write transfers start working but there's a big backlog of data to be written. All ATLAS dCache pools are setup as active
- 22/07/10 08:30 - All ATLAS write transfers are working again.

Details

- ATLAS dCache pool setup consists in two pool groups: new DDN pool servers and old Thumper pool servers. In order to focus activity in the new hardware, dCache links' priority were set that DDN servers had higher priority.
- Having 2 links with different priorities was not working as expected: a pool in the lower priority link would had to handover when no high priority pool was available. This didn't happen and the transfer failed with errors like *no space available*, *pool cost too high* or *Best pool <dc008_1> too high : Infinity*.

Follow-up

- Both ATLAS links have equal priority from now on.
- A ticket has been opened to dCache.org in order to clarify link priority behaviour **#5753: dCache 1.9.5-21@PIC: PoolManager cost issue (ATLAS impacted)**
- Ensure that the relevant VO-specific SAM tests are interfaced to the local Nagios system at PIC (watched by the 24x7 operator).
 - Part of this was already done on Jul 23rd: the result of SAM-ATLAS FCR

tests has been implemented as a sensor at nagios.pic.es
(SAMtest_SRMDISKv2_Atlas @ coesrv02)

- Another part still missing: clarify the situation of visualizing in the local Nagios the SAM-ATLAS tests which are critical for the "dashboard" and not FCR.