Service Incident Report for the accidental ATLAS file deletion at PIC on October 2012

Description

205.544 files were deleted as a result of a mistake in dark data removal procedure. Affected files were from PIC atlasdatadisk/ following directories: data09_cos , data09_2TeV, data09_900GeV, mc09_10TeV, mc09_14TeV.

Impact

Out of 205.544 lost files: 9 files were not in LFC, 2.709 files were recovered and 202.826 files could not be recovered from elsewhere and hence were removed from dataset definitions. This summary is available from the ATLAS consistency service historic records:

http://bourricot.cern.ch/dq2/consistency/history1/

Time line of the incident

Preliminary check for dark data was done before October 9th 2012. A script to check if the files in the directories named *09* were registered in the LFC was developed. A list of 607.547 potential orphan files was obtained.

On October 9th deletion was started, using lcg-del command for every file in the list. During deletion one of the datasets was crosschecked using "dq2-list-datatset-replicas" tool and was found to be a valid dataset. The developed script to detect orphan files had a problem, and the files detected were actually correctly registered in the LFC. After detecting this, the deletion was stopped and PIC internal RT ticket was filed at 21:05 to follow up the details of the accident.

October 10th at 10:10, DQ2-DDM Operations Savannah Bug report #98070 was started. Next day files were declared as lost to ATLAS consistency service and the accident reported on the daily WLCG Operations meeting.

Analysis

The error was introduced when in the PIC SRM host endpoint was changed by mistake in the script from "srmatlas.pic.es" to network alias "srm.pic.es". This change leaded to the output of "lcg-lr" command to be the same as in case of files not in LFC, since srm was not known to ATLAS LFC. On the other hand deletion of the files was successful since network name of the srm server was correct.

Follow-up actions

Custom procedure of individual directories dark data check up was retired. Proper procedure now includes usage of complete LFC dump (provided by ATLAS DDM Operations) and ATLAS scripts.