

Summary of CMS Data Loss Incident at FR-CCIN2P3 tier-1

Date: 26-11-2009

This fall in preparation for data collection CMS was in the process of identifying files that were resident on Tier-1 mass storage, but not tracked in the experiment's data management system. This build-up of untracked and unwanted files is usually caused by quiet failures of the deletion system. Consistency checks are run periodically to determine files that are listed as resident but in fact missing, indicating data loss or a missed failed transfer, and files that are listed as deleted but in fact still resident, indicating untracked unneeded storage use.

CC-IN2P3 was in the consistency process with the help of the Central CMS Data Operations trying to attain the list of files to clean out. On November 11th CMS Data Operations sent confirmation of a list of directories that contained files that were not tracked in the data management system. The CMS instructions were assessed after the fact to be potentially confusing. The instructions were misinterpreted by the local data manager and the entire contents of the listed directories were deleted and not just the files untracked in the data management system. After some exchanges internal to the site for confirming the request, the files in the specified directories were definitely removed from the storage systems (disk and tape) in several stages performed on November 12th and on November 19th. On November 20th CMS began to see the CC-IN2P3 site fail the Tier-1 availability tests, because the sample needed for the SAM tests and Job Robot no longer existed, and the problem was identified. On Saturday November 21st it was confirmed that it was not possible to recover the deleted entries at the level of the MSS namespace.

In the end approximately 600TB of files were deleted, of that 480TB were classified as custodial data on tape. Of the samples approximately 100TB cannot be retransferred from CERN or another Tier-1. The unrecoverable sample is approximately two-thirds simulated events produced before summer 09 and one-third derived data from CMS Cosmic running. All the samples lost could be re-derived at a later date if there was a need from CMS.

CMS completed the process of synchronizing the data management system with the new file catalog at CC-IN2P3 on November 24th the transfer of data from the tier-0 was suspended until this date. The retransfer of lost samples to IN2P3 will proceed slowly as a massive retransfer would place a high load on CERN and the other Tier-1s. The holiday shutdown will be employed for some transfers of old samples. The SAM test samples have been re-replicated and IN2P3 is passing the site readiness tests.

Moving forward: effective immediately CMS will provide a list of every file intended for deletion when a clean-up is requested. On the CC-IN2P3 side, a more stringent multi-step procedure will be now on implemented for any destructive operation requested by and involving data of any LHC experiment:

- An explicit list of files will be requested to the experiment. This list should be attached to the ticket open against the site for requesting the operation.
- The consequences of the request will be assessed by two different individuals within the site. As a result of this step, an exhaustive list containing the full file names, associated file sizes, aggregated amount of data involved in the operation and some relevant characteristic of each

file (i.e. custodial type) will be established and provided to the experiment for final confirmation. The explicit final confirmation will be required to be issued by an authoritative person within the experiment

- If confirmed, the destructive operation will be started by the site not before the next business day

In addition to these already decided measures, there will be a review of the incident, CMS and Tier-1 procedures, and future recovery options at a meeting during CMS week (December 7-11).