

## KIT Service Incident Report

### Description

TSM is the central component of KIT tape storage system which keeps the whole tape metadata information in its database. TSM at KIT uses mirror storage which is managed by TSM itself and not by OS or storage hardware. Database and undo/redo logs are on local storage and on TSM managed mirror copies on SAN storage. The SAN storage reported controller errors (20.04.2016 18:00) after having been repaired the day before. (The repair was done online.) Also TSM reported storage errors. The SAN disk storage was set read/only. In order to fix the problem, TSM was shut down, the SAN mirror remounted RW and the TSM server restarted. TSM did not start properly, probably because database problems. We did not succeed to get TSM running and eventually decided to restore the database to most recent time. This procedure uses the redo files on disk. The restore took extraordinarily long and was therefore stopped because no activity could be observed. We suspected that the redo log files were partially corrupted. The next step was to restore to a certain point in time before the SAN storage problem of 20.04.16 18:00. We supposed that the storage errors corrupted the database and restored to 20.04.16 15:00. The restore succeeded and operations could continue.

Recently it was discovered that some files were no longer on tape. Further investigation showed that the missing files were written directly before the time of restore (20.04.16 15:00). It turns out that the TSM software used the latest full backup of 20.04.16 9:00, although the command suggested it used 20.04.16 15:00. This resulted in the situation that we unknowingly started operations with a 6 hours gap in the database during which files had been written. The files are actually on tape but are lost because they are no longer registered in the TSM database.

### Impact

7185 files are lost for ATLAS, 75 files for LHCb and 2 for CMS, none for other VOs

### Time line of the incident

Date	Task
20.04.2016 09:12	Regular backup of the TSM database completed.
20.04.2016 18:00	Disk storage system problem detected.
21.04.2016 10:00	In order to remount the file systems R/W, TSM was shut down.
21.04.2106 10:00-13:00	Trying to start TSM fails and is crashing with segmentation faults
21.04.2016 14:00	Restore of the TSM Database to the most recent backup before the storage went down.
21.04.2016 21:00	The restore process shows no activity and was stopped after 5 hours
22.04.2016 08:00	Attempt to restore the TSM database to 20.04.2016 15:00. TSM autonomously selected the last full database backup from 20.04.2016 09:00.
22.04.2016	After the successful database restore, dCache administrators rescheduled a flush for all files written to tape since 20.04.2106 15:00:00.
12.01.2017	dCache administrators find out that some files are lost on tape.

## **Analysis**

From the events we concluded that the TSM mirroring did not function as expected. A restore of the DB based on files stored on the local disk did not succeed which led us to believe the mirror function actually left files on the local disk corrupted. The subsequent restore procedure did not prevent the use of a less recent backup image. Obviously the use of the disk based undo/redo log fails without warning and the restore continues with the (older) tape based copy.

## **Follow up actions**

We switched off TSM managed mirroring and research different backup strategies.

## **Summary**

Disk repair resulted in disk storage access problem which resulted in corrupted disk contents. The tape management system (TSM) uses the disk for its database files and therefore the disk contents had to be restored from tape. Requesting restore of a backup image made after the full backup, silently results in a data base restore of the full backup resulting in loss of database entries and therefore loss of files.