

KIT Service Incident Report

Description

Files on 3 tapes were not readable and after further investigations we have identified an error that possibly yields such files. Using that knowledge, we went through occurrences of said error in the past and found 28 affected cartridges written since 2011, 12 of which were already found to be broken earlier with other issues. Unfortunately, we cannot tell whether this error happened for any tapes before 2011.

Impact

396 files are lost for ATLAS, 28 files are lost for CMS.

Time line of the incident

30.09.2014	First broken tape (int. identifier "UR5298") discovered
01.10.2014	Another two broken tapes ("UR5290", "dUR5291") discovered
14.10.2014	Completed checks for all affected tapes, ultimately 424 files were lost

Analysis

We now have found a probable correlation between a TSM error and not readable files [1]. It seems that TSM writes an "EndOfData" marker at a wrong position on tape, which renders all files following it unreadable. With successive attempts to write on an affected tape volume, TSM reports an error indicating a database problem [2]. There are no hardware errors seen on any involved servers or tape drives. However, not all tapes for which this error occurred are actually broken – we found 5 tapes with this error condition, that were still entirely readable.

[1] Further information on ANR8302E:

<http://www-01.ibm.com/support/docview.wss?uid=swg21239786#>

[2] Further information on ANR0102E:

<http://pic.dhe.ibm.com/infocenter/tsminfo/v7r1/index.jsp?topic=%2Fcom.ibm.itsm.msgs.server.doc%2Fmsg-ANR0102E.html>

Follow up actions

Unfortunately, our TSM server version is no longer officially supported and an update is not feasible at the moment. Therefore we have no possibility to prevent loss of data in the future due to this error. We will combat it by monitoring TSM for these errors and immediately react to it, i.e. migrate all data to another scratch tape.

Summary

Upon finding several tape cartridges with partially unrecoverable content, we were able to identify a correlation between these tapes and a specific failure reported for them. A total of 28 cartridges showed this problem since 2011 and we checked all of them for persistent issues. 5 tapes proved to be a wrong positive and all their content was entirely readable. The data on the remaining tapes was no longer retrievable and we had to declare 424 file lost towards ATLAS and CMS.

Because our version of the Tivoli Storage Manager is no longer officially supported and the update is not feasible for the time being, we cannot guarantee to avoid such situations in the future. The only option we have right now is to monitor our systems closely and react immediately to this error by migrating all affected data onto another scratch tape.