

Description

Discrepancy in an Oracle database table replicated between destination and its source database via oracle streams.

Impact

Transactions (inserts) missing in a destination table between March 14, 2010 and February 17, 2011. So far no user reported problems related to this issue.

Time line

- Tuesday, September 6, 2011. 14:02 UTC. Tier 0 IT-DB administrator reported the inconsistency between the source table and the destination table at BNL, identification of the missing rows was also provided.
- Tuesday, September 6, 2011. 15:22 UTC. Inconsistency was reported within the experiment distributed computing operation channels
- Tuesday, September 6, 2011. 18:19 UTC. Synchronization of the table was done by IT-DB administrator per Tier 1 DBA request. Streams process was disabled during this operation that has taken 5 minutes to complete.
- Tuesday, September 6, 2011. 18:30 UTC. Per table synchronization latest apply transactions were verified.
- Tuesday, September 6, 2011. 18:44 UTC. Full availability of this table and schema was reported within the experiment distributed computing operations channels
- Tuesday, September 6, 2011. 21:00 UTC. Affected table had instantiation problems in 2010 after a hardware migration. Tier 1 DBA checked between source and destination the 9 tables affected in 2010 and found that this problem was isolated to this table and different from the one observed in 2010. This information was also confirmed by Tier0 IT-DB administrators.
- Tuesday, September 6, 2011. 22:00 UTC. No trace files or information associated with this problem was found in the local database.
- Monday-Wednesday, September 12-14, 2011. Posterior verification by row count comparison across both source and destination databases. No discrepancy was found.
- Thursday, September 15, 2011. Preliminary summary of this incident was presented in the T1SCM
- September 19-26, 2011. Final analysis

Analysis

Tier 0 IT-DB administrator suspects that the transactions were skipped during the applying process. The transactions SCN was lower than the table instantiation SCN or the table had been prepared after import was done. In summary 1 table (ATLAS_COOLONL_TRIGGER.COMP200_F0032_IOVS) was affected which corresponds to 0.19% of the entire number of tables and 0.5% of the rows affected in this schema, approximately 0.1% of the total number of rows in the database.

To re-synchronize the table data was copied back by using database links and a temporary table holding the information missed due to the table topology. This procedure took 5 minutes, streams replication had to be

stopped during this procedure. This procedure was done at the Tier 0. Verification of apply transactions was done by the local DBA (at BNL) after the table re-synchronization.

A more extended check was done which included verification of the affected table and the row count comparison between source (ATLR database) and destination (BNL Database). No discrepancy was found with the local schemas being replicated and compared.

Follow up

A preliminary summary of this incident was presented at the T1SC on Sep 15, 2011. Input from Tier 0 experts was used to complement this SIR report. Most relevant items:

1. The creation of a Service Request to Oracle. It was suggested/proposed to let the CERN-IT-DB lead this effort and create this SR due their technical expertise with this technology. Oracle should provide an explanation and recommendation related to this specific issue.
2. The need of a consistency verification mechanism that allows to verify efficiently source and replica (if required), most probably to be included in 11g due to technical restrictions in 10g. Oracle should provide such internal mechanism.

Final analysis (19/09/2011 - 23/09/2011)

Posterior database analysis and by crosschecking the time of the past streams interventions at Source (Tier 0) and Destination (Tier 1) database allowed to understand and complete this SIR as:

1. The SCN for the affected table was the same as the other 8 tables fixed in 2010. The table instantiation SCN was properly set.
2. Due to a streams intervention at Tier 0 involving the creation of new streams environment the entire streams dictionary was re-sent to Tier 1s databases in 08/02/2011.
3. The time period of the gap of the missing rows in the affected table corresponds approximately to the time between the first attempt to fix the table in Feb 2010 and the time after the entire streams dictionary was re-sent to the Tier 1s databases as previously mentioned.

Conclusion:

The 1 table discrepancy was due to a failed attempt to recreate the missed streams data dictionary for the table in Feb 2010 while executing an internal procedure advised by Oracle support from Tier 0. This procedure permitted to select individual parts of the streams dictionary by object id to be sent to the Tier 1 database (BNL). However, it was established that resending the entire data dictionary could have avoided this issue as it minimizes the probability of a human error while applying this procedure to individual objects in the database.

As for the items mentioned in the follow up section of this report, it is to be said that item 1 is no longer needed since the root cause of this issue seemed to be identified. As for item 2 to help detect and or prevent similar issues in the future, CERN IT-DB administrators have written a database procedure to crosscheck the content of the streams dictionary between T0 and T1s. It has been integrated in each LCG replication environment deployed as a database job, which is running every week. Nevertheless, Oracle should provide a detailed crosscheck internal mechanism to be implemented in Oracle 11g.