

Table of Contents

Oracle physics databases (T0 and T1) upgrades to Oracle 11gR2 report.....	1
Upgrades at T0.....	1
Databases.....	1
Timeline.....	1
Streams environments.....	1
Problems identified during testing phase.....	1
Timeline.....	2
Problems faced during the upgrades.....	2
Problems faced after upgrade.....	3
Upgrades at T1.....	4
RAL.....	4
TRIUMF.....	4
KIT.....	4
SARA.....	4
CNAF.....	4
PIC.....	5
BNL.....	5

Oracle physics databases (T0 and T1) upgrades to Oracle 11gR2 report

Upgrades at T0

Databases

- ATLDSC, LHCBDSC, ATLARC, ATONR, ATRR, ADCR, CMSARC, CMSR, CMSONR, LHCBR, LHCBONR, LCGR, ALIONR, PDBR, COMPR databases upgraded successfully.
- 4 hours downtime for each database. Also hw migration using Oracle DataGuard. No issues found.
- LCGR upgrade took 15 minutes more than scheduled due to a missconfiguration with the scan listeners.
- After LCGR migration to new hw, a problem with the eth2 interface was found on node number #2. Node was taken down for further investigations and back on 16th February.

Timeline

ATLARC	24-11-2011
ATLDSC and LHCBS	14-12-2011
ATONR (ATLAS online db)	11-01-2012
ATLR (ATLAS offline db)	17-01-2012
ADCR (ATLAS offline db)	17-01-2012
CMSARC	17-01-2012
LHCBR (LHCb offline db)	24-01-2011
COMPR	27-02-2012
LCGR	31-01-2012
CMSR	01-02-2012
CMSONR	02-02-2012
LHCBONR (LHCb online db)	06-02-2012
PDBR	08-02-2012
ALIONR (ALICE online db)	08-02-2012

Streams environments

Problems identified during testing phase

- DDL failing in cascade streams replication from 10g to 11g.
 - ◆ problem reported to Oracle (SR# 3-3378025701) and fixed by patch 12593103 which is included 11.2.0.3
- Redo logs from primary database are not registered on downstream capture database and cannot be removed by RMAN * According to Oracle Support (SR# 3-4893515191) it is a feature and customer should implement its own solution for redo removal * Cron jobs were implemented to keep reasonable redo window on downstream capture databases
- Single capture process (and source queue) cannot serve 10g and 11g targets in parallel
 - ◆ reported to Oracle (#3-4497346601)
 - ◆ WORKAROUND: Separation of 10g and 11g targets by creation of new capture process and source queue dedicated for 11g targets
- 11g propagation processes randomly stop sending messages to 10g targets
 - ◆ reported to Oracle (#3-4497346601)
 - ◆ WORKAROUND: implementation of propagation processes restarting job

Timeline

Date	Action
14-Dec-2011	Upgrade and migration of downstream capture databases (plus implementation of all identified problems workarounds)
11-Jan-2012 - 17-Jan-2012	ATLAS online-offline upgrade
18-Jan-2012 - 26-Mar-2012	ATLAS T0 - T1s upgrade: RAL 18-Jan-2012, TRIUMF 23-Jan-2012, KIT 24-Jan-2012, BNL 22-Feb-2012, IN2P3 26-Mar-2012
24-Jan-2011 - 06-Feb-2012	LHCB online-offline upgrade
30-Jan-2012 - 19-Mar-2012	LHCB T0 - T1s upgrade: PIC 30-Jan-2012, CNAF 31-Jan-2012, RAL 31-Jan-2012, SARA 01-Jan-2012, KIT 31-Jan-2012, IN2P3 19-Mar-2012
01-Feb-2012 - 02-Feb-2012	CMS online - offline upgrade
08-Feb-2012	ALICE online - offline upgrade
08-Feb-2012 - 27-Feb-2012	COMPASS online - offline upgrade

Problems faced during the upgrades

- 10g apply process served by 11g source database stops applying messages after reboot of hosting instance. Messages are removed and lost and resynchronization of data is required.
 - ◆ no workaround available, lost messages have to be resend from T0 (by manual reconfiguration of capture process)
 - ◆ problem appeared for:
 - ◇ KIT (14-Dec-2011 LHCB, 20-Dec-2011 ATLAS cond) , RAL(11-Jan-2012 LHCB LFC), PIC (23-Jan-2012 LHCB LFC)
- Capture process has difficulties (and in most cases crashes) while mining redo logs generated during database upgrade.
 - ◆ observed on all source databases
 - ◆ SOLUTION: recreation of capture process to start just after execution of upgrade scripts - no data are lost as upgrade is done in restricted mode and users cannot access the database to modify the data
- Capture process aborts because unsupported changes were made in one of replicated schemas (ORA-26767) and/or apply process fails with ORA-26714 and ORA-00942
 - ◆ problem was observed after all source database upgrades (starting from ATONR)
 - ◆ new feature of segment advisor in 11g creates temporal compressed tables that are not supported by Streams.
 - ◆ SOLUTION [see metalink note 1082323.1]: additional filter rules were added to capture processes to prevent that changes on those tables (temp tables have fixed names: DBMS_TABCOMP_TEMP_CMP and DBMS_TABCOMP_TEMP_UNCMP) were captured.
- Changes on compressed tables are not replicated when source database is 11.x.x and compatible parameter is 10.x.x (which is the case just after the upgrade)
 - ◆ observed only during ATLAS upgrade, few replicated schema had old partitions compressed on few tables. Changes stopped to be replicated on those tables since ATONR was upgraded (but parameter stayed unchanged).
 - ◆ SOLUTION: old partions from ATONR were dropped (copy of them are on ATLR) and tables have been reinstatiated. Compatible parameter has been increased to 11.2.x few weeks after the database upgrades.
- Logminer/Capture process crashes with ORA-600 error when attempts to read archived logs generated by database with different compatible version

- ◆ problem is triggered by logminer redo threads detection algorithm when at least one of database instances is not up when change of compatible parameter is done. Logminer is trying to read last or next archived log (depends on the point in time when capture process has to start) created by unavailable instance which possibly has different compatible version than database current one. This will lead to fatal error in logminer process.
- ◆ observed two instances of the problem during upgrades of:
 - ◇ ATR on downstream capture (17-Jan)
 - in this case capture process needed to recapture changes before changing the parameter and crashed because of reading redo with newer compatibility value (from instance that was down during intervention)
 - recreation of capture process after the intervention was enough to solve the problem.
 - ◇ CMSONR (02-Feb)
 - capture processes were recreated after modification of parameter value when 1 out of 2 instances were still down. Even though there were no data to be extract from past logminer was making an attempt to open the last redo log written by not running instance (with old compatible parameter value) which was resulting in fatal error in logminer.
 - problem was solved by manual editing of problematic file header and changing the version to match current database compatible parameter value.
- ◆ according to Oracle (SR 3-5268527771) the inability of logminer to read redo files containing different compatible versions in header causing fatal error is not a bug. Changing of compatible parameter should be done on all instances at once.
- ◆ emergency WORKAROUND (not recommended): modify redo log files headers in order to have same compatible version for all logs that logminer wants to access.
- ◆ permanent WORKAROUND: compatible parameter has to be changed on all instances at once (set desired compatible value in spfile and bounce all instances in parallel)
- No redo continuity provided for standby and downstream databases when primary database is being failovered.
 - ◆ Most likely it is a bug as it is fully reproducible. It is being followed up with Oracle Support (SR# 3-5370026131).
 - ◆ WORKAROUND: Recreation of downstream capture process with start SCN after the failover. This can lead to lost of transactions on replicas if old capture process was not up to date with capturing data changes from old primary database.
 - ◆ observed during upgrade of LHCBR (24-Jan). Capture process could not start as one log was missing. After recreation of capture process all LHCBR conditions T1s replicas aborted (18:35) as few transaction were lost. Missing data has been resync manually using source database at 21:15.

Problems faced after upgrade

- data propagation and capture freezes after network disruption with "ORA-2685: Unable to connect to apply "" because it has connected to another capture".
 - ◆ propagation cannot re-establish broken connection to apply process as it is still holding old connection handler. Restart of apply process cleans old process context and brings replication back on track.
 - ◆ IMPORTANT: If capture processes is restarted before apply, due to problem with communication to replica, an oldest checkpoint SCN is picked up as first one and logminer has to mine all redo logs from retention window.
 - ◆ WORKAROUND: Restart of apply process. Detection of broken connections and automatic restart were implemented in streams monitoring tool.
 - ◆ problem observed during interventions on CERN central firewall (all outgoing connections were being disconnected) on week 16-22nd of April.

Upgrades at T1

RAL

- ATLAS database successfully upgraded on the 18th January.
- Intervention did over-run for few reasons:
 - ◆ the backup of the database did over-run (block change tracking was not enabled and incremental level 1 backup took more than 1 hour of the time allocated for the intervention).
 - ◆ pre-check was run the day of the intervention: found out that some objects were invalid and took ~1 hours to recompile/rebuild/drop them.
 - ◆ clufy (cluster verification tool) run the same day of the intervention: found some problems, for example Voting Disk configuration (one voting disk was on NFS and clufy didn't like it so Carmine had to drop it).
- Unfortunately these problems were not found during the testing because RAL's testing HW configuration is different from the production one (so couldn't spot the VD problem) and the database is much smaller than the OGMA one so the backup was faster and didn't have any invalid object.
- Lesson learned: run the checks well in advance!
- LHCb database upgrade went smoothly.

TRIUMF

- ATLAS database successfully upgraded on the 23rd January.
- Database upgrade had to be re-started as `shared_pool_size` and `java_pool_size` parameter were not set to the minimum value. This delayed the upgrade by 15 minutes.

KIT

- ATLAS database successfully upgraded on the 24th January.
- Intervention took longer than scheduled due to some firewall problems.
- LHCb databases migrated and upgraded on the 1st February. DataPump export and import was used instead of DataGuard which is the procedure recommended by CERN.
- Intervention took 12 hours to be completed:
 - ◆ DataPump requires more downtime than DataGuard.
 - ◆ Configuration issues found with the new hw.

SARA

- LHCb database successfully upgraded on the 1st February.
- The upgrade of the clusterware (from 11.1 to 11.2.0.3) took about 4 hours.
- The database upgrade itself took about 2 hours.
- No problems were found, but the process was slower than expected.
- Used CERN documentation as a guide. References to local CERN scripts generated some confusion.
- Lesson learned: CERN documentation is provided only as a reference. Documentation and stepst must be prepared beforehand.

CNAF

- LHCb and LFC databases successfully upgraded on the 31st January.
- Slow progress due to the lack of preparation (no test system) and a big snowstorm which difficulted Alessandro to reach CNAF offices - 26 hours in total (7 hours scheduled).

PIC

- LHCb database upgraded successfully on the 30th January.
- Intervention took a bit longer due to some problems found during the post-intallation steps while registering the database with the clusterware (mistake with instances configuration).
- Elena has also found some problems while deinstalling the old Oracle software. Fixed using the note: "RAC Instance Fails to Start With Error ORA-27504 [ID 1059831.1]"

BNL

- ATLAS database successfully upgraded on the 22nd February. Also hardware Migration and CPU JAN 2012 applied. Using DataGuard. No issues were found. Took about 2 hours and 10 minutes.
- Upgrades still pending due to the issues found with the Oracle installer and the database storage.
- Bug 13731278 prevents the CRS upgrade from 10.2.0.5 to 11.2.0.3.
- Service Request (3-5225032577) open and escalated. According to the support, there is no way to pass over the bug without a patch, being developed by Oracle (no timeline has been provided)

-- ZbigniewBaranowski - 29-Feb-2012

This topic: DB > PhysicsDatabase11gUpgradeReport

Topic revision: r10 - 2012-04-23 - ZbigniewBaranowski



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use Discourse or Send feedback