

January 16, 2007

- Oracle January security patch was released yesterday. It does not seem critical but we are still analyzing it. More news in the next days.
- In order to free up space for new production servers, the pilot/integration INT3R database was moved to a 64bit installation, same as new production servers. Thanks to the developers for their help. The replacement of validation INT6R database (also a 64bit installation) is scheduled to end of next week.

December 12, 2007

- patch for fragmentation problem of FTS is out. We would like to install on integration, it is a rolling operation
- meeting on December 11 with SAM and Gridview responsables. Discussed better ways to handle duplicate sources of data; encryption possibilities in Oracle and queries that can be improved.
- due to the arrival of new HW, we need to migrate WLCG integration RACs to different machines. This will include move to 64bit platform, which will be the future production HW platform.
 - ◆ If possible we would prefer to start from scratch, only copying data when necessary. I volunteer to contact application owners to see if this is possible.

November 14, 2007

- we have applied the October CPU in a rolling way to LCG RAC and all other production RACs.
- start work with Gridview in order to improve the application by decreasing necessary reads/writes from the database and amount of raw data necessary to keep.

October 17, 2007

- **Oracle Critical Patch Upgrade for October 07**

A new patch was released by Oracle yesterday and it needs to be applied on all our databases as there are some critical security holes patched. The patch seems rolling and first test deployments will be done today. In case everything goes ok on our test RACs, we would like to proceed with deployment on the integration databases early next week.

October 10, 2007

- **Bad plan used by Oracle optimizer affected SAM on 20 September**

For unknown reason the oracle optimizer start using a worst path when getting results for one main query of SAM application. This lead to very slow access to the data and obliged to stop operation of whole SAM application. As the application have not been change recently it was not straight forward to understand the culprit of the problem.

Together with the SAM developers we quickly isolate the problem to a single query. The solution was to add a hint, specifying the optimized to use a specific index. The operation resumed to normal after this.

A service request was open (SR 6498675.994) but meanwhile the problem has disappeared and Oracle analyst could not help further. We suppose the problem is related to 'bind variable peeking', on which the optimized choses the plan based on the first set of values passed to the query. However this plan can be worst in most of the situations.

The close collaboration with SAM team proved very helpful to quickly solve this problem.

- **Drop tables on production schema by SAM developer (on FCR Portal schema) on 27 September**

Judit developed a cleanup script for her application (part of SAM suite) and by mistake ran it against the production database, dropping all tables of the application. Immediately (at 16:30) she contacted Miguel to recover the tables.

This was done in from 17:00 to 19:30 (Thursday 27 Nov) using the on-disk backup (1 day old), which was open as a new database copy in one node and rolled forward with current redo logs. It was a very fortunate case because without on-disk backup the recovery from tapes would have taken at least the full night (current size of DB is >1TB) and much more preparation work would have been required in order to find and configure spare HW to hold the recovery copy of the database.

In order to avoid similar errors in the future we propose to block the 'owner' schemas in the production RAC where these operations are possible.

- **Unavailability of SAM on Tuesday (2 October) from 15:30 to 19:00**

During a regular clean up of a big table of LCG_SAME, in order to set some partitions as READ-ONLY and without need of regular backups, an incomplete command lead to have part of indexes set as 'UNUSABLE' by an Oracle command. The table in question is formed by monthly partitions, each of the partitions have local indexes. The move of partitions from 1H2007 made the local indexes of the same 1H2007 as 'UNUSABLE'.

Even if those partitions are never used by the application, Oracle optimizer thinks that it cannot use any other local index for the queries. This is unexpected behavior and a 'Enhancement Request' was filed for this (SR 6516713.993).

The SAME service was brought down in order to rebuild the most critical local indexes set as 'UNUSABLE', so that the monitoring collection from SAM could be up as quick as possible. This was finished at 19:00 and David Collados from SAM team open the web service and tomcat, activate monitoring collection (critical SAME part). Gridview and FCR components were also activated at same time.

Other indexes were needed to be possible to use 'SAM PORTAL'. In agreement with the SAM team this rebuild started just after (from 19:00) and lasted for the whole night. At 7:30am (Wednesday) the operation had finished and 'SAM Portal' service was open and first connections arrived.

I think that the database team (in particular Miguel) did an excellent job in overcoming the problems and restoring normal operations. We will continue to collaborate studying ways preventing such incidents in future.

- **Oracle Instant Client 10.2.0.1 bug when uptime >248 days**

Last Friday 5th October, LFC worker nodes hit a bug from Oracle Instant client (Bug 4612267 OCI CLIENT IS IN AN INFINITE LOOP WHEN MACHINE UPTIME HITS 248 DAYS) when uptime of the machine is bigger than 248 days . The symptoms is a hang on connecting to the database on a Oracle library. This is a known bug from 10.2.0.1 version. Workaround is to reboot the node. We recommend service managers to check the uptime of the machines running with Oracle Instant Client and reboot if found necessary. Solution to this bug is to upgrade Instant Client to version 10.2.0.3, which is available on SLC4 repositories. For SLC3, in case the application *does not* use OCCI driver (ie, uses OCI, JDBC, thin, ...), there should not be a problem and a upgrade is recommended.

August 22, 2007

- There was a error affecting whole WLCG cluster on Friday (August 17th) afternoon for about 1h20. It is not yet understood the cause of the crash. However the cluster was instable since the deployment of

CPU patch on August 15th due an error on deployment in one of the instances (command run as root changed some file permissions). This error was thought to be fixed next day and database access was not showing any problem.

- Given read only access to LFC for LHCb to Barbara (CNAF)
- no updates on FTS fragmentation
- everything ready for tomorrow upgrade of FTS, backup of affected tables was tested and will be done before the intervention.

August 15, 2007

- Critical Patch Update for July is being applied at the moment after testing in the integration systems.
- Fragmentation of FTS table seems to be related to a bug in Oracle Automatic Segment Space Management. A Service Request is going on to try to solve the problem.
- New Request of read only account for LFC for LHCb from CNAF.

July 11, 2007

- Added LHCb LFC moved to downstream capture and replicated to 4 Tier1 sites (RAL, PIC, IN2P3, GridKa)
- Moved Atlas Conditions production to 3D setup: Online->Offline->Tier1
- Studying the reasons of fragmentation of FTS' T_FILE table

June 20, 2007

- FTS defragmentation going on. One table had 130GB of reclaimable space after clean-up of data. Need to see the need also for Tier1-sites.
- SAM problem solved two weeks ago. The cause was connections waiting for more data after the firewall configuration downloaded every hour.
- Database workshop included a recovery exercise. Included recovery from disk/tape and then synchronizing with T0. Four sites successful did it. Two others did not have the backup infrastructure correctly in place.
- LHCb LFC replicas will be set up soon with same configuration of conditions data.

May 30, 2007

- problem on SAM (affected also VOMS) last Thursday due a move of partition which make an index partition unusable (even if it was not used). Due other high loads on the application it was not so easy to detect the root cause and took a bit longer to fix by rebuilding the affected index partitions.

May 9, 2007

- Removed some vulnerable features from Oracle databases. Those features were not used and this way we avoid applying a non-rolling security patch.
- Discussed with Gridview to remove old raw data that was not used anymore. There will be a job to clean-up this data time to time.
- 3D - interventions on other T1 sites are reported
<https://twiki.cern.ch/twiki/bin/view/PSSGroup/LCG3DWiki>
- Some applications requesting possible replication using 3D. Requirements will be collected in 3D meetings and presented at Physics Group Leaders meeting as it requires some resources.

April 25, 2007

- April security patch do not need to be apply. Only one vulnerability could affect our DBs but can be removed by revoking access to some non used features.
- Partitioning of LCG_SAME was done successfully. It took much longer than expected (several days against 10h in tests) due the high activity on this table.
- Need to discuss now with Gridview developers about partitioning.
- No problems found on FTS2 run on pilot.

April 18, 2007

- April security patched released this night. Only one issue might affect our clusters, need to study if it is necessary to apply. The installation implies complete shutdown of the RAC.
- Partitioning tested for LCG_SAME, it took 6h but data always available. Proposed to do on production tomorrow. The idea behind is to move data to read-only tablespaces not needing regular backups.
- Partitioning discussed with FTS2 to avoid creation of History tables. Seems a possible solution
- FTS2 pilot running without big problems for the moment. There is a close look to this.
- Waiting for Gridview developers to be available in order to discuss partitioning and query changes.
- Database for Atlas available at PIC, still need to pass CERN configuration checklist.
- Oracle Client 10.2.0.3: download and installed on the external area of LCG. Already being pick up for testing and validation.

April 4, 2007

- Lowered the FTS load on production RAC by adding two indexes and strict agents to run on one node. This was done after high load been reported on the machines. Being checked with fts developers if this applies also to new FTS version.
- FTS2 pilot running on INT6R. Load is too low to see potential problems.
- Proposed changes to Gridview queries. Still working with developers.
- Partitioning of big tables of Gridview and SAM starts to be necessary (already a 58GB table!). Suggestions will be sent to developers.

March 21, 2007

- DBAs meeting with T1 sites in Amsterdam going on (21/21 March)

March 14, 2007

- Following some hangs of LCGR instance we have applied the workaround suggested by Oracle. This solved the situation so far as no more hangs were seen.
- New pilot RAC prepared for WLCG service. There are presently two 2-node RAC for pilot and tests of WLCG applications. It is now necessary to decide which accounts are kept on each of the RACs.

February 28, 2007

- LCG production RAC running on 8 nodes with RHES4 and ORacle 10.2.0.3. Intervention successful and smooth.
- LHCB RAC also expanded and upgraded.
- ITRAC hangs affected LCG two times this week. Oracle recommends to apply a workaround. Plan to do it soon on all production RACs. Transparent and rolling operation.

- LCG_GRIDVIEW service suffered some instabilities during 27th Feb, after first hang of one LCGR instance. Understood as bug of Oracle tool to modify services ('srvctl modify' command). Service needed to be recreated. Thanks for the cooperation of lcg_gridview team. Service request opened with Oracle support. Internal procedures changed in order not to use the buggy command.

February 21, 2007

- expansion of WLCG RAC scheduled for Thursday February 22 at 8am. Intervention plan done, coordination of Giacomo Govi on PSS side.
- No news on the service request related to last week' problem.
- LHCB RAC intervention will be done Monday (Feb 26th) from 8am. LFC for LHCB will be affected.
- Workaround for the ITRAC hangs seems to solve the problem with some tradeoff on performance. It will not be added to new LCG RAC nodes.

February 14, 2007

- expansion of WLCG RAC was not done due a (what seems) Oracle bug on the recreation of redo log files. This step had as aim to keep a convention across all RACs. For future migrations we will skip the step until we understand from Oracle exactly what has failed and how to avoid.
- possible to reschedule to anytime from tomorrow.
- LHCB RAC upgrade scheduled to W8 (no specific date set yet).

January 31, 2007

- removed access to vulnerable package on production RAC on Tuesday 30 January. Transparent intervention.
- ATLAS online and offline RACs expanded, moved to new hardware and upgraded to 10.2.0.3 this week.
- Proposed date for expansion of WLCG RAC: 13 February (Tuesday) from 9 to 12 CET (downtime).

January 17 2007

- 10.2.0.3 deployed on integration (int3r) services since beginning of December 06.
 - ◆ new patches need to be applied over 10.2.0.3. It will be announced in Thursday morning meeting for next day intervention. Expected downtime of 30 minutes.
 - ◆ no feedback from applications on 10.2.0.3 testing (apart from VOMS)
- Expansion of LCGR on W7 (to 8 node RAC and new HW)
 - ◆ deployment of 10.2.0.3 in production together with the expansion
 - ◆ expected downtime < 3 hours
- Other LHC experiment expansions affecting WLCG services
 - ◆ ATLAS and PDBR RAC (ALICE) on W5
 - ◆ CMS RAC on W7
 - ◆ LHCB RAC on W8
- Streams (distributed databases) status
 - ◆ LFC for LHCB running 24/7 to CNAF.
 - ◆ ATLAS streaming to four T1 sites (not production)

This topic: LCG > LcgScmStatusDb2007

Topic revision: r1 - 2008-10-15 - SteveTraylen



Copyright &© 2008-2020 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
Ideas, requests, problems regarding TWiki? Send feedback