

Table of Contents

LCGSCM Database Services Status.....	1
November 4, 2009.....	1
October 21, 2009.....	1
October 7, 2009.....	1
May 27, 2009.....	2
April 28, 2009.....	2
January 21, 2009.....	2
January 14, 2009.....	3
December 10, 2008.....	3
November 26, 2008.....	3
November 05, 2008.....	4
October 15, 2008.....	4
October 1, 2008.....	4
September 17, 2008.....	5
September 3, 2008.....	5
August 6, 2008.....	6
July 23, 2008.....	6
July 09, 2008.....	7
June 25, 2008.....	7
June 4, 2008.....	7
May 21, 2008.....	7
March 19, 2008.....	8
February 27, 2008.....	8
February 13, 2008.....	8
January 23, 2008.....	8
Old Reports.....	9

LCGSCM Database Services Status

November 4, 2009

- LCGR database successfully migrated to new hw (RAC 7) on Wednesday 21.10
- Oracle has released the Critical Patch Upgrade for October 09. Patch deployment on some integration databases already done.
- Dashboards migration to LCGR cluster:
 - ◆ ATLAS and CMS preparation still ongoing.
- Streams Replication:
 - ◆ ASGC database resynchronization using transportable tablespaces with help of Tier1 BNL successfully performed on Wednesday 28.10

October 21, 2009

- Production databases migration to the new hw (RAC 7):
 - ◆ LHCB, CMSR and downstream databases (ATLDSC and LHCBDS) already done
 - ◆ LCGR scheduled for today
- Dashboards migration to LCGR cluster:
 - ◆ Alice and LHCB already done using Data Pump to transfer the data.
 - ◆ ATLAS and CMS being prepared. Migration will be done using Oracle Streams to transfer the data from one cluster to another as the required downtime for Data Pump cannot be allowed.
- Streams replication:
 - ◆ RAL had a severe problem with the database storage which caused some data loss on the databases. Streams were reconfigured in order to recover and resynchronize the databases.
 - ◆ SARA has migrated the ATLAS/LHCB/LFC database to new hw using Data Guard. The operation was not as smooth as with other Tier1 sites due to a problem with the connection strings and the aliases to the database nodes.
 - ◆ ASGC database inaccessible during 3 weeks. Resynchronization will be performed end of this week or beginning next week.

October 7, 2009

- All integration databases were migrated to RHEL5 before the summer. Databases hosted on RAC7 are being installed with RHEL4, despite of the initial plan was to migrate all the DB clusters to RHEL5 by end of summer. The migration had to be rolled back to RHEL4 due to bugs with version 5.3 affecting our hardware infrastructure (only RAC7) and now lack of time to roll-out RHEL5.4 before the LHC start-up.
- Several production databases have been upgraded during last months with the latest security and recommended Oracle patches. The list of patched databases includes: ATLAS online (ATONR), ATLAS offline (ATLR), CMS online (CMSONR), CMS offline (CMSONR), LHCB online (LHCBONR), LHCB offline (LHCBR), Grid central DB (LCGR) and downstream databases (ATLDSC and LHCBDS).
- New applications migrated to production: LCG_SITE_MONITORING (dashboard), LCG_SAM_ATP (grid operations).
- LCG_SAME: bad queries have been re-written in order to improve their performance.
- Streams:
 - ◆ Replication for AMI ("Atlas Metadata Interface") application data between IN2P3 (Lyon) and CERN (ATLAS production database) was moved to production (09.06.09). This is the first setup in production where the master is a Tier1 site and CERN acts as a replica.

- ◆ New Streams test environments: PVSS data replication for ALICE and CMS experiments (now all 4 LHC experiments use streams from online to offline).
- ◆ Tier1s:
 - ◇ RAL has migrated the ATLAS and LHCb databases to new hw from 32bit to 64bit using Data Guard.
 - ◇ BNL and NDGF are preparing similar intervention in the coming weeks.
 - ◇ SARA is also preparing a database migration but using a different approach.
 - ◇ ASGC (Taiwan) is out of the ATLAS replication setup due to the database corruption after a planned intervention. A complete re-instantiation will be needed.
- To do in the next weeks:
 - ◆ Migration of Dashboards to the LCGR cluster. Preparation ongoing.
 - ◆ Migration of several production databases to the new hw RAC7: CMS offline (CMSONR), LHCb offline (LHCBR), Grid central DB (LCGR) and downstream databases (ATLDSC and LHCBDSC). At the moment, they are being prepared as Oracle standby databases. This activity is in preparation to moving those databases following a Oracle Dataguard switchover intervention.
 - ◆ Oracle will release the October Critical Patch Update this month.

May 27, 2009

- April security patch update and Oracle recommended bundle patches have been successfully installed on all production databases. Some issues have been observed during the patching:
 - ◆ ATLAS offline DB (ATLR) has suffered from severe performance degradation between 16:40 and 17:20 (on Tuesday 12th May), due to high load generated by ATLAS DQ2
 - ◆ CMS online DB (CMSONR) was temporarily (around 11:00, 14th May) not accessible to some applications due to wrong connection string.
- Oracle Enterprise Grid Control agents have been upgraded to version 10.2.0.5.
- Streams:
 - ◆ Replication from CERN to IN2P3 (Lyon) was stopped on Sunday 03.05 due to a cooling problem at the computing center. ATLAS, LHCb and LFC replication was resumed on Monday 04.05 around 14:00.
 - ◆ A network problem at RAL affected ATLAS replication from Sunday 03.05 until Tuesday 05.05.
 - ◆ ATLAS requested to stop the apply process at IN2P3 (Lyon) during part of the ATLAS Stress tests. Apply was disabled from Tuesday 12.05 to Wednesday 13.05.

April 28, 2009

- LCG_SAME: performance problems with some of the queries. Working to improve them.
- VOMS: the problem affecting VOMRS (probably related to possible Oracle bug mixing bind-variables) will be investigated with the help of Steve Traylen.
- Security patch: April critical security patch was released and we are deploying it on development and validation databases this week. Together with the CPU APR09 patch we have also decided to apply a series of 'bundle patches' recommended by Oracle. Application managers have been notified by email. The intervention is rolling and no service interruption was/is expected.

January 21, 2009

- Streams: Problem with Taiwan continue (DB corrupted). We will try to work with Jason this week, profiting of his presence at CERN.

- VOMS: one of the two problems seen by Castor team in RAL related to possible Oracle bug mixing bind-variables seems also be present in VOMS, but too infrequent to manage to give enough data to Oracle. It is not too serious to VOMS.
- Security patch: January critical security patch was released and we plan to deploy on development and validation databases next week. An email to the application managers will be sent with the date. This is rolling and no service interruption is expected.

January 14, 2009

- Streams: Atlas streams configuration to Tier1 is being re-created today. The streaming to Tier1 was stopped since last year due a problem in a procedure by Atlas DBAs.
 - ◆ Problem with Taiwan, which database is corrupted. This affects Atlas configuration.
- VOMS: all DB problems are currently solved - both of bind variables and too many short connections. Thanks to Steve Traylen for the help.

December 10, 2008

- Streams: problem created by an Atlas DBA which stopped all the streams configuration. It was fixed immediatly but also reminded that they should not to these operations.
- Meeting with VOMS to discuss last problems with database. The problem of missing bind variables is solved and still to solve problem of many short connections. Being investigated together with Steve Traylen, it is likely to be related to connection pool configuration.
- Xmas shutdown: validation DB as well as experiment online DBs will be left running unattended.

November 26, 2008

- Patches: All production databases at CERN were patched with October Critical Patch, kernel updates and a clusterware patch bundle. This last one included most likely a change which makes nodes to self-reboot quite easily when not responding. During next intervention we plan to increase the threshold to avoid false reboots.
- Meeting with security people to discuss DB security. After a look at the logs we saw no connection to LCG production DB from outside CERN, so we plan to close it in the firewall. This will be discussed further. The password security, account bookkeeping and logging of failed connections already in place were well appreciated.
- Streams:
 - ◆ Another case of missing primary key aborted streams, this time between CMS on-line and off-line DB. The application was removed from the configuration and responsible reminded.
 - ◆ CNAF DB for Atlas streams not working fine today. We have contacted CNAF support.
- A workshop on Distributed Database Operations has been held on November 11th and 12th 2008(<http://indico.cern.ch/event/43856>). It was well attended, with 40 people registered. The workshop has focused on discussing:
 - ◆ DB Deployment and Operation for CASTOR Services at Tier1s
 - ◆ Hardware Requirements and Performance Review
 - ◆ Database Administration, Configuration and Monitoring at Tier0 and Tier1

- Xmas shutdown: we would like to know if the LCG validation services need to running during the shutdown.

November 05, 2008

- Patches: Validation/development system currently with three patches deployed - erroneous ORA-942 error bug; October Critical Patch; Clusterware patch bundle (19 bugs fixed). The deployment in production of the whole set will be done next Thursday 13th November. It is a rolling intervention, service will be always available.
- Streams:
 - ◆ Post-mortem: Following some ATLAS stress tests launched on Friday evening (17 October), on Saturday morning a high load at NDGF caused the database to be totally unresponsive. This caused, on Sunday morning, a "pause" of the capture process at CERN as the queue was full. On Monday morning NDGF had to reboot the host in order to fix the problem with the database. Capture was resumed and apply restart working everywhere by then. Extra monitoring was added to spot this situations and study to increase memory dedicated to streams to avoid capture queue to get full over long weekends. (full details: https://twiki.cern.ch/twiki/bin/view/PSSGroup/StreamsPostMortem#Problem_with_ATLAS_replication)
 - ◆ Missing primary key on a table of Atlas Online database created a problem on the apply on the Offline DB, making necessary to create an ad-hoc index. Users were reminded that all replicated tables must have primary keys defined.
- Workshop: Distributed Database Operations Workshop on 11/12 November mostly dedicated to operations at Tier0 and Tier1 sites (<http://indico.cern.ch/event/43856>)

October 15, 2008

- Released this night the October Critical Patch Update from Oracle. We are currently deploying on our tests systems and we will deploy during the following weeks in the validation databases. The patch for ORA-942 bug currently in validation will be deployed in production together with the October CPU.
- Intervention on LFC for LHCb done last week on each of Tier1. It was found database misconfiguration at RAL, which slowed down the intervention. This was corrected and recommendations were sent to all Tier1 DBAs.
- Task-force for solving high-load problem with Atlas reconstruction jobs at Tier1 progressed with some improvements in the software and tests are ongoing.
- Intervention on main table of LCG SAME yesterday required 5 hours of writing downtime. Archiving and partitioning policy is being discussed to reduce those necessary downtime in interventions like this.

October 1, 2008

- Atlas Tier1 conditions DB getting overloaded issue: 'task-force' created including Atlas developers and IT-DM DBAs to try to find a solution.
- LFC replication to SARA aborted due 'user error', which means that data was not the same as destination as in origin. Root cause not found - usually is write at read-only destination. Problem corrected next morning.

- CNAF is the only Tier1 still running Oracle version 10.2.0.3 for LFC and Conditions database. This version has known streams bugs causing the apply process to abort time to time.
- CMS Online database was not reachable during whole Friday 26th September. During the weekend also streams from online database to offline were not working. This was caused by hardware problems and DNS misconfiguration at the pit. Still necessary to improve the workflow between IT-DM and CMS SysAdmins for systems located at the pit.
- Currently on the validation databases at CERN it is applied a patch against a bug for error ORA-942 and possible SQL execution in wrong schema. Users were requested to see if no misbehavior is seen on validation systems, so that in two weeks we deploy the patch in production.
- Need to define deletion policy for SAM and Gridview databases. Currently all row data is kept, bringing possible performance problems.

September 17, 2008

- LFC replication from LHCb to Gridka is aborted due to error 'connection lost contact'. It was third time it happen in the same setup, Service Request open with Oracle since April (when first occurred) but no solution yet. Workaround exists but it is very painful as it requires to recreate all the stream processes involved in the replication to Gridka.
- Problem on Atlas Tier1 conditions DB being overloaded time to time. Maria Girone and Sasha Vanyashin have worked on this. When ATLAS start a reconstruction set of jobs together and each job opens 20 DB connections the combination of 20 to 30 such jobs together with Oracle streams receiving data into the Tier 1 DB from CERN is sufficient to overload the Oracle backend. ATLAS tested this but without the streaming component and this is now a serious problem for them and other experiments in the case of shared servers.
- The scheduled OS upgrade on the production RACs was done in a rolling fashion during last week. This addressed a problem where Oracle clusterware was crashing.
- For further protection against human errors, disaster recoveries and security attacks we are setting up a standby copy of the production experiment online and offline databases and LCG database.

September 3, 2008

- LHCb database streams configuration has been completed including the new LHCb Online database. Their setup includes the online and offline conditions data to be replicated to six Tier-1 sites.
- Few problems still seen in new LHCb Online database setup with the storage and delaying the move of remaining applications to this setup. LHCb SysAdmins are looking at it.
- Also new Atlas schemas are now being streamed to Tier-1 sites. A new bug was found and workaround applied, which includes to set up rules at the destination sites for some operations.
- Corruption was found on the Atlas online server two weeks ago has not affected DB services and data availability. However the corruption has caused a serious risk of further spreading and affecting services in case of subsequent HW failure. Therefore a standby database has been set up with Oracle Dataguard technology to further protect the database in case of data loss. An intervention has been performed on Wednesday 20th August afternoon when the standby database has been activated for production of Atlas online. This has fixed the corruption issue. There are strong indications that the 'silent corruption' was generated by a faulty drive. We have also added additional monitoring at the

DB level to allow early detection of corruption issues.

- All production databases are now patched with the latest Oracle security upgrade. This was done without disturbing the service.
- Following a recent update from Oracle support we have applied an OS (kernel and glibc) upgrade on all the integration databases. This was done without disturbing the service.
 - ◆ Details on the issues addressed by the OS upgrade:
 - ◇ Oracle crashes caused by the CCSD demon (clusterware) have been reported in production, in particular on ATR, since the upgrade to 10.2.0.4. This issues has been diagnosed as a consequence of a bug in mlockall() fixed in RHEL4 Update 7 (in particular in glibc-2.3.4-2.41).
 - ◇ RHEL4 Update 6 kernels (in particular the kernel used in PDB production, 2.6.9-67.0.20.Elsmp) have a bug in sys_times() which may cause system crashes (although not yet seen in PDB production). RHEL4 Update 7 kernel 2.6.9-78.0.1.Elsmp has a fix for this issue.
 - ◆ OS upgrade has been already scheduled for the production databases:
 - ◇ ATR and ATONR: Wednesday 3rd September from 14:00 till 17:00
 - ◇ CMSR and CMSOR: Monday 8th September between 15:00 and 17:00
 - ◇ PDBR: Monday 8th September between 15:00 and 17:00
 - ◇ LCGR: Tuesday 9th September between 10:00 and 12:00
 - ◇ ATLDSC and LHCBDS: Tuesday 9th September between 15:00 and 17:00
 - ◇ LHCBR: Wednesday 10th September (to be confirmed)
- Work going on with SAM and Gridview team in order to prepare the move of common tables to a Topology database. First tests in development database are being carried on.

August 6, 2008

- Oracle Critical Patch Update for July is deployed since 25th July on the development and validation databases of LCG. Following our security and change management policy, we will deploy next Monday morning (11th August) in the production database. All other production databases will have the patch deployed between tomorrow and next week.
- Last week there was a problem with the Stream setup which was several days behind due a missing archive logs. This seems very difficult to monitor. Problem is solved and a post-mortem was created at <https://twiki.cern.ch/twiki/bin/view/PSSGroup/StreamsPostMortem>
- Still a problem with LFC propagation to Gridka, which had some database problems. The site needs to be re-instantiated and should be in synch today.
- Last week there was a problem with FTS caused by an Oracle bug that was triggered when we created a copy of the schema in preparation for the next FTS upgrade. This was solved the day after and a post-mortem was created by the FTS deployment team.

July 23, 2008

- Oracle Critical Patch Update for July is out and we will deploy on the validation RACs of LCG this Friday morning. The installation is rolling so it should be transparent to the users. Following our security and change management, we plan to deploy this patch on the production database during the week of 11th of August.
- On Tuesday 15th July a switch problem affected the LCG validation RACs to which the access was

very slow. This problem was understood and fixed by CS on the same day.

July 09, 2008

- Working with Gridview people to improve performance. Developers looking at new ways to aggregate the data.
- Database developers workshop yesterday, good attendance of about 50 people. Still not easy to get some developers present.
- Meeting with LCG developers on 25th June, minutes at <http://indico.cern.ch/conferenceDisplay.py?confId=36783>
- CNAF computer center a bit shaky, may affect streams and is running in a separate configuration
- RAL LFC stream was aborted, seems they were writes at destination database. Account was found open, after being locked by us.

June 25, 2008

- Power cut test this morning caused some confusion. Tim Wheibley said our machines would not been affected. Machines not in critical power were restarted and seem some fiber channel switches also failed. All the application services were then running on nodes on critical power and a manual operation needed to relocate them. One node of LHCb got corrupted on OS level.
- Meeting with LCG developers this afternoon.

June 4, 2008

- Found that switches of production LCG, CMS and LHCb RAC were connected to Physics power. This made the database to be unavailable during power cut and also brought some other problems as nodes could not communicate among them. New power bar is now in place but it will be necessary to move the power cord from one to another. This should be a 1/2 minute intervention which might be transparent for most of the applications. We forecast to do it on 25th June, but we will confirm later.
- Intervention for upgrade of Oracle on production LCG database scheduled for 16th June.
- Other upgrades will be: LHCb 5th June (affecting LFC for LHCb), Streams DBs 5th June (affecting LHCb and Atlas propagation to Tier1), Atlas: 18th June, CMS: 25th June.
- Soon will be new meeting with LCG developers for discussing elements to improve on the production and best practices that should be implemented.
- there will be a review of cms_dashboard application soon in view of better resource usage.

May 21, 2008

- cms_dashboard problem yesterday, they were running two instances at same time which overloaded the database. We will propose meeting with developers for a review and also to try to avoid changes on the production instances. Furthermore we will propose to lock 'owner' accounts on the production databases.
- New Oracle version will be deployed tentatively during second week of June. It has been running for 7 weeks in the validation databases without new problems.
- Smooth run of all experiment (and WLCG) databases since the start of CCRC May'08. No higher load seen.

March 19, 2008

- cms_dashboard problem yesterday, when users did wrong command for merging partitions. Downtime of couple hours.
- 10.2.0.4 to be deployed on integrations, proposed next week. Need deep testing before we can deploy in production
- new hardware being installed, need to test apps, as we pass to a different configuration
- Monday meeting about space usage on LCG database
 - ◆ space pressure on the storage arrays
 - ◆ datafiles with 650GB of wasted space
 - ◆ agreed to drop some detailed data
 - ◆ regular meeting with main lcg developers every month

February 27, 2008

- Intervention scheduled for March 4th on production LCG database:
 - ◆ apply of latest Critical Patch for Oracle (30 minutes downtime)
 - ◆ do clean up on SAM and Gridview schemas, partition and move old data (2 hours downtime)
- high load on database due FTS query taking 0.5 seconds instead of 0.01. Application does not seem affected by it. This problem arises after automatic database statistics are gathered for one table. Statistics for this table were locked until end of CCRC. The origin might be another Oracle bug, as this table is getting fragmented very fast. Service request with Oracle will be open.

February 13, 2008

- Oracle January Critical Patch was applied on all experiment production databases. Schedule for W-LCG database after the February CCRC 08.
- Request from Gridview to perform a clean up of a big table. This can be best performed with scheduled downtime of the application. Schedule after February CCRC 08.
- We propose the date 3rd of March 2008 for these interventions.

January 23, 2008

- Oracle January Critical Patch Upgrade was decided to be applied on all databases. Announced to LCG and experiments to be applied this week on the integration services (int1 1r thursday at 10UTC). Following our policy, we would like to apply on the production databases in two weeks time. This is a rolling intervention and no service interruption is expected.
- **Streams:** patch for bug blocking capture process on ATLAS setup took one week to be released, even if it was with the highest priority possible. It took then about 5 hours to recover the backlog to tier1.
- Streams setup for VOMS replication to CNAF is ready.
- Several interventions on Tier1 sites:
 - ◆ TRIUMPH today for storage re-configuration
 - ◆ CNAF wants to migrate streams to different hardware this or next week. LHCb Conditions streams will be done tomorrow (Roberto agreed). LFC needs to be scheduled.
 - ◆ SARA has a service intervention of 7days for moving cluster from room to another.
- These kind of long interventions are bigger than our window to recover backlog. To re-instantiate the data we need then to do a full export/import and that is getting more difficult due the growing amount of data. Would be nice to use transportable tablespaces. However we should not use T0, as we might have more data that is filtered in the streams. We are looking to make the exercise with two Tier1 and then announce in the next 3D workshop.

Old Reports

- LcgScmStatusDb2007
-

This topic: LCG > LcgScmStatusDb

Topic revision: r57 - 2009-11-03 - EvaDafonte



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