CERN IT-DB Services: Deployment, Status and Outlook

Luca Canali, CERN

Gaia DB Workshop, Versoix, March 15th, 2011

Database SERVICES

Outline

- Review of DB Services for Physics at CERN in 2010
  - Availability
  - Incidents
  - Notable activities
- Infrastructure activities, projects, planned changes
  - Outlook and service evolution in 2011

CERN and LHC

CERN – European Organization for Nuclear Research – located at Swiss/French border

LHC – Large Hadron Collider – The most powerfull particle accelerater in the world – launched in 2008

LHC data correspond to about

20 million CDs each year!

RDBMS play a key role for the analysis of LHC data

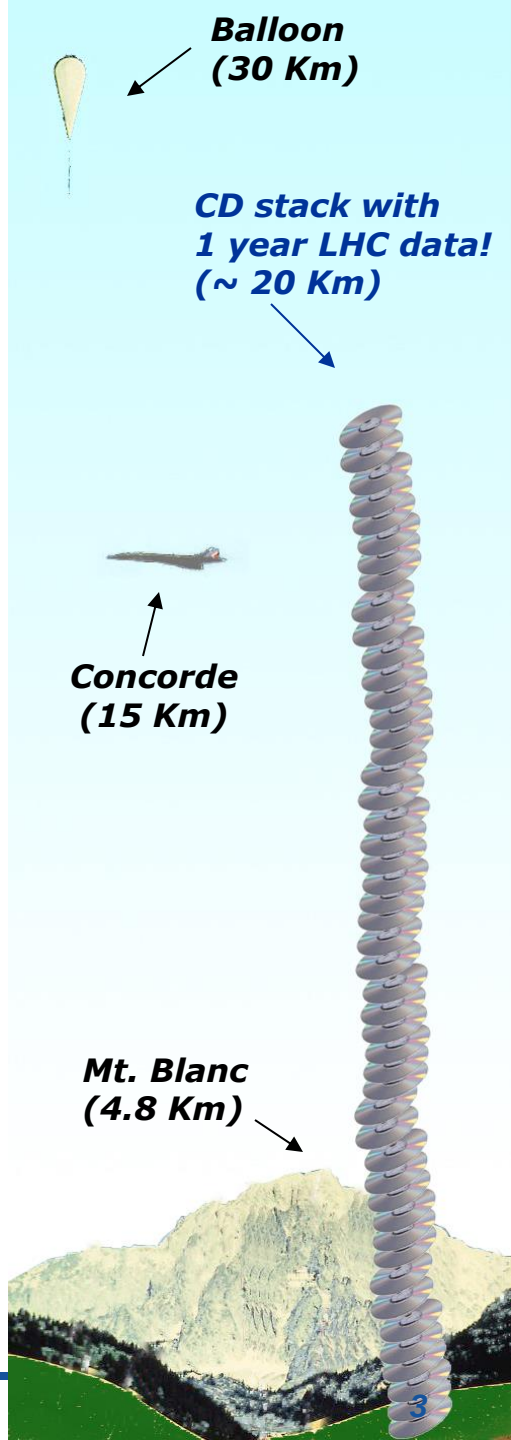

Physics and Databases

- Relational DBs play today a key role for LHC Physics data processing
  - online acquisition, offline production, data (re)processing, data distribution, analysis
    - SCADA, conditions, geometry, alignment, calibration, file bookkeeping, file transfers, etc..
  - Grid Infrastructure and Operation services
    - Monitoring, Dashboards, User-role management, ..
  - Data Management Services
    - File catalogues, file transfers and storage management, …
  - Metadata and transaction processing for custom tape-based storage system of physics data
  - Accelerator logging and monitoring systems

Database SERVICES

At the heart of CERN, LHC and Experiment Operations

CERN IT Department

http://cern.ch/it-dep/db/

IT/DB GROUP

Streams

Experiment Offline Databases

9 Production DBs, 7 Integration DBs
5 Tests DBs, 7 Standby DBs

Streams

Tier-1 Centres

dedicated 10 Gbit/s links

NDGF, GridKa, BNL, RAL, ASGC, SARA-NIKHEF, PIC, FNAL, CCIN2P3, INFN - CNAF, TRIUMF

Data

Experiment Online Databases

4 Production DBs

RAW Data

LHC Experiments

LHCb, CMS, ATLAS, ALICE

CASTOR

CERN Advanced STORage manager

21 Production DBs
4 Development DBs

Middleware

73 Application Servers

Administrative/IT/ Engineering Databases

17 Production DBs
11 Development DBs
4 Ref/Test DBs

Accelerators ACC

12 Production DBs

EDMS

AIS

LHC Operations

PARTNERS:

CERN openlab

ORACLE

CERN

Current status on February 2010

CERN Databases in Numbers

- CERN databases services
  - Global users community of several thousand users
  - ~100 Oracle RAC database clusters (2 – 6 nodes)
  - Currently over 3000 disk spindles providing more than ~3PB raw disk space (NAS and SAN)
- Some notable DBs at CERN
  - Experiments' databases – 14 production databases
    - Currently between 1 and 12 TB in size
    - Expected growth between 1 and 10 TB / year
  - LHC accelerator logging database (ACCLOG) – ~50 TB
    - Expected growth up to 30 TB / year
  - ... Several more DBs on the range 1-2 TB

Updates on LHC

- Successful re-start of LHC operation in 2010
  - 2011 run started mid Feb., beam energy of 3.5 TeV
  - Work going on with the acceleration to increase luminosity (and rate of data collection)

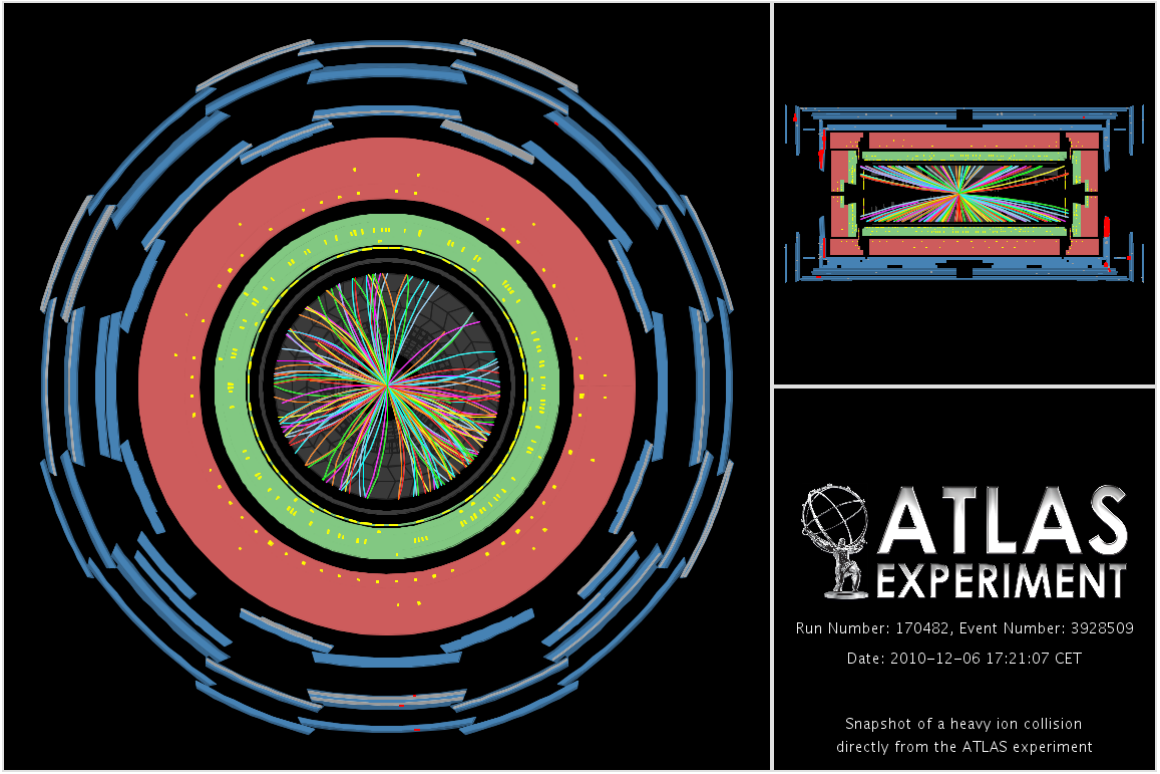

Status of the DB Services for Physics

Service Numbers

- Infrastructure for Physics DB Services
  - ~115 quadcore machines
  - ~2500 disks on FC infrastructure
- 9 major production RAC databases.
- In addition:
  - Standby systems
  - Archive DBs
  - Integration systems and test systems
  - Systems for testing streams and 11.2

Services and Customers

- Offline DB Service of LHC experiments and WLCG
- Online DB Service
- Replication from online to offline
- Replication from offline to Tier1s
- Non-LHC
  - biggest user in this category is COMPASS
  - and other smaller experiments

DBA Support

- 24x7 support for online and offline DBs
  - Formalized with a ‘CERN piquet’
  - 8 DBAs on the piquet
  - Temporary reduced personnel in Q3 and Q4:
  - Note on replication from offline to Tier1s
    - is ‘best effort’, no SMS alert (only email alert)
    - on-call DBA checks email 3 times per day

Service Availability

- Focus on providing stable DB services
  - Minimize changes to services and provide smooth running as much as possible
  - Changes grouped during technical stops
    - 4 days of stop every ~5 weeks
    - Security patches, reorg of tables
    - Major changes pushed to end-of-the-year technical stop (~2 months of stop)
- Service availability:
  - Note these are averages across all production services
  - Offline Service availability: 99.96%
  - Online Service availability: 99.62%

Notable incidents in 2010 1/2

- Non-rollingness of April Patch
  - Security and recommended patch bundle for April 2010 (aka PSU 10.2.0.4.4)
  - Contains patches marked as rolling
  - Passed tests and integration
- Two issues show up when applied in production
  - Non rolling on clusters of 3 or more nodes with load
  - On DBs with cool workload
    - Symptoms: after ora-7445 and spikes of load appear
- Ora-7445
  - Reproduced on test and patch available from Oracle
  - Thanks to persistency team for help
- Non-rollingness
  - Reproduced at CERN
  - Related to ASM

Notable incidents in 2010 2/2

- Two issues of unscheduled power cut at LHCB online pit
  - ~5 hours first occurrence (9/8)
  - ~2 hours for second occurrence (22/8)
- In first incident DB became corrupted
  - Storage corruption
  - Lost write caused by missing BBUs on storage after previous maintenance
  - Restore attempted from compressed backup, too time consuming
  - Finally switchover to standby performed
    - See also further comments on testing standby switchover in this presentation
- Another instance of corrupted DB after power cut
  - 18-12-2010, archive DB for Atlas corrupted
  - Recovery from tape: about 2 days

Notable recurring issues

- Streams
  - Several incidents
  - Different parts of replication affected
  - Often blocks generated by users workload and operations
- High loads and node reboots
  - Sporadic but recurrent issues
  - Instabilities caused by load
  - Run-away queries
  - Large memory consumption makes machine swap and become unresponsive
  - Execution plan instabilities make for sudden spikes of load
  - Overall application-related. Addressed by DBAs together with developers

Activities and Projects in 2010

Service Evolution

- Replaced ~40% of HW
  - New machines are dual quadcores (Nehalem-EP)
    - Old generation was based on single-core Pentiums
  - New storage arrays use 2TB SATA disks
    - Replaced disks of 250GB
- New HW used for standby and integration DBs
  - New HW (RAC8+RAC9): 44 servers and 71 storage arrays (12 bay)
  - Old HW (RAC3+RAC4): 60 servers and 60 storage arrays (8 bay)

Consolidation of Standby DBs

- New HW installed for standby DBs
  - Quadcore servers and high-capacity disks
    - This has increased resources on standby DBs
    - Provided good compromise cost/performance in case of switchover operation (i.e. standby becomes primary)
  - Installed in Safehost (outside CERN campus)
    - Reduce risk in case of disaster recovery
    - Used for stand by DBs when primary in CERN IT

Oracle Evolution

- Evaluation of 11.2 features. Notably:
  - Evaluation of Oracle replication evolution:
    - Streams 11g, Goldengate, Active Dataguard
  - Evolution of clusterware and RAC
  - Evolution of storage
    - ASM, ACFS, direct NFS
  - SQL plan management
    - for plan stability
  - Advanced compression
- Work in collaboration with Oracle (Openlab)

10.2.0.5 Upgrade - Evaluation

- Evaluation of possible upgrade scenarios
  - 11.2.0.2, vs 10,2.0.5, vs staying 10.2.0.4
  - 11g has several new features
    - Although extensive testing is needed
    - 11.2.0.2 patch set came out in September and with several changes from 11.2.0.1
  - 10.2.0.4 will go out of patch support in April 2011
  - 10.2.0.5 supported till 2013
    - 10.2.0.x requires extended support contract from end July 2011
  - Decision taken in Q3 2010 to upgrade to 10.2.0.5 (following successful validation)

10.2.0.5 Upgrade - Review

- Testing activity
  - Several key applications tested
  - No major issues found
  - Very difficult to organize a 'real world' tests
- Upgrade of production during January 2011
  - Technical stop for the experiments
  - Mostly a smooth change
    - Some minor issues found only when switching to production
    - A few workaround and patches add to be added

Activities on Backup

- Backups to tape using 10gbps
  - have been successfully tested
  - Speed up to 250 MBPS per 'RMAN channel'
- First phase of production implementation
  - Destination TSM at 10gbps
  - Source multiple RAC nodes at 1gbps
    - Typically 3 nodes
  - In progress (~30% of DBs by Q1 2011)
- Other activities
  - Moving backup management to a unified tool inside the group
  - Unified tool for routine test of DB recoveries from tape

Activities on Monitoring

- Improvements to custom streams monitoring
  - Added Tier1 weekly reports
  - Maintenance and improvements to streammon
    - DML activity per schema, PGA memory usage
- OEM 11g
  - Currently deployed at CERN
  - Several issues needed troubleshooting
    - Notably a memory leak triggered by browser
- Internal activities on monitoring
  - We are unifying monitoring infrastructure across DB group

Activities on Data Lifecycle

- Goal: avoid that DB growth impact manageability and performance
  - Activity launched in 2008
  - Partitioning and data movement main tools
    - Compression used too
  - In 2010 more applications modified to allow partitioning
    - Data start to be moved to archive DBs
    - Joint work DB group and experiments/development

Activities on Security

- Internal application developed
  - To track access to DBs
  - Mining audit data
  - Allows to spot unusual access patterns
  - Can be source of info for defining white lists
- Firewall active on DB servers
  - Further discussion on activating white lists going on

Activities for Online Setups

- ALICE, LHCB and CMS online
  - Installations of the DBs at the experiments' pits
  - HW is managed by experiments
- HW warranty expiring
  - Replacement under way
  - IT discussed with experiments on HW replacement
  - Goal of having similar HW at the pit as in IT to reduce maintenance effort and complexity
  - Deployment of new HW expected in Q1 2011

Standby Tests

- Standby DBs and switchover operation
  - Tested and documented
  - Ideally a test switch should be performed on all production DBs
    - Activity needed to validate the disaster recovery infrastructure
  - During technical stop in Q1 2011
    - Scheduled test of Atlas online standby DB
    - Downtime ~.5h to switch to standby and ~.5h to switch back

Outlook and Activities for 2011

Major Scheduled Changes

- Upgrade to 11gR2
  - Upgrade of Oracle to 11gR2 in Q1 2012 (11.2.0.3?)
- Replacement of 2/3 of production HW
  - New servers and storage
- Combined change
  - 'Swing upgrade': upgrade of standby built on new HW
  - Production DBs and constraints from experiments
    - Maintenance window limited to 'extended technical stop' for many DBs, i.e. Q1 2012

Software changes preparation

- Testing of existing applications on 11gR2
  - In collaboration with the experiments
  - Our experience is that some issue are only seen under load
  - Load-based testing necessary
  - Investigating Oracle RAT?
- Other software changes
  - Investigating RHEL6
  - rpm-based installation (integrated with CERN OS installation, i.e. quattor)
- Overall goal of unifying procedures across IT-DB group

11gR2 Features

- We will further testing and prepare to deploy new 11gR2 features of interest
- Very interested in
  - Active Data Guard
- Other notable features of interest for us:
  - Improvements to streams
  - Improvements to ASM and NFS
  - Improvements to clusterware
  - SQL plan management
  - Interval partitioning

Activities on Architecture

- New HW acquisition
  - In 2011 a large group of production machines goes out of warranty
  - HW renewal and occasion to profit from more recent HW for performance and capacity
  - Several technologies have been evaluated
    - Use of SSD for caching
    - 10gbps Ethernet for interconnect
    - 8gbps Fiber Channel storage access
    - NAS at 10gbps with flash cache
  - Evaluation of Oracle VM

Current Model

- Dual-socket quad-core DELL blade servers, 24GB memory, Intel Xeon “Nehalem”; 2.27GHz
- Dual power supplies, mirrored local disks, redundant 1GigE, dual HBAs, “RAID 1+0 like” with ASM and JBOD

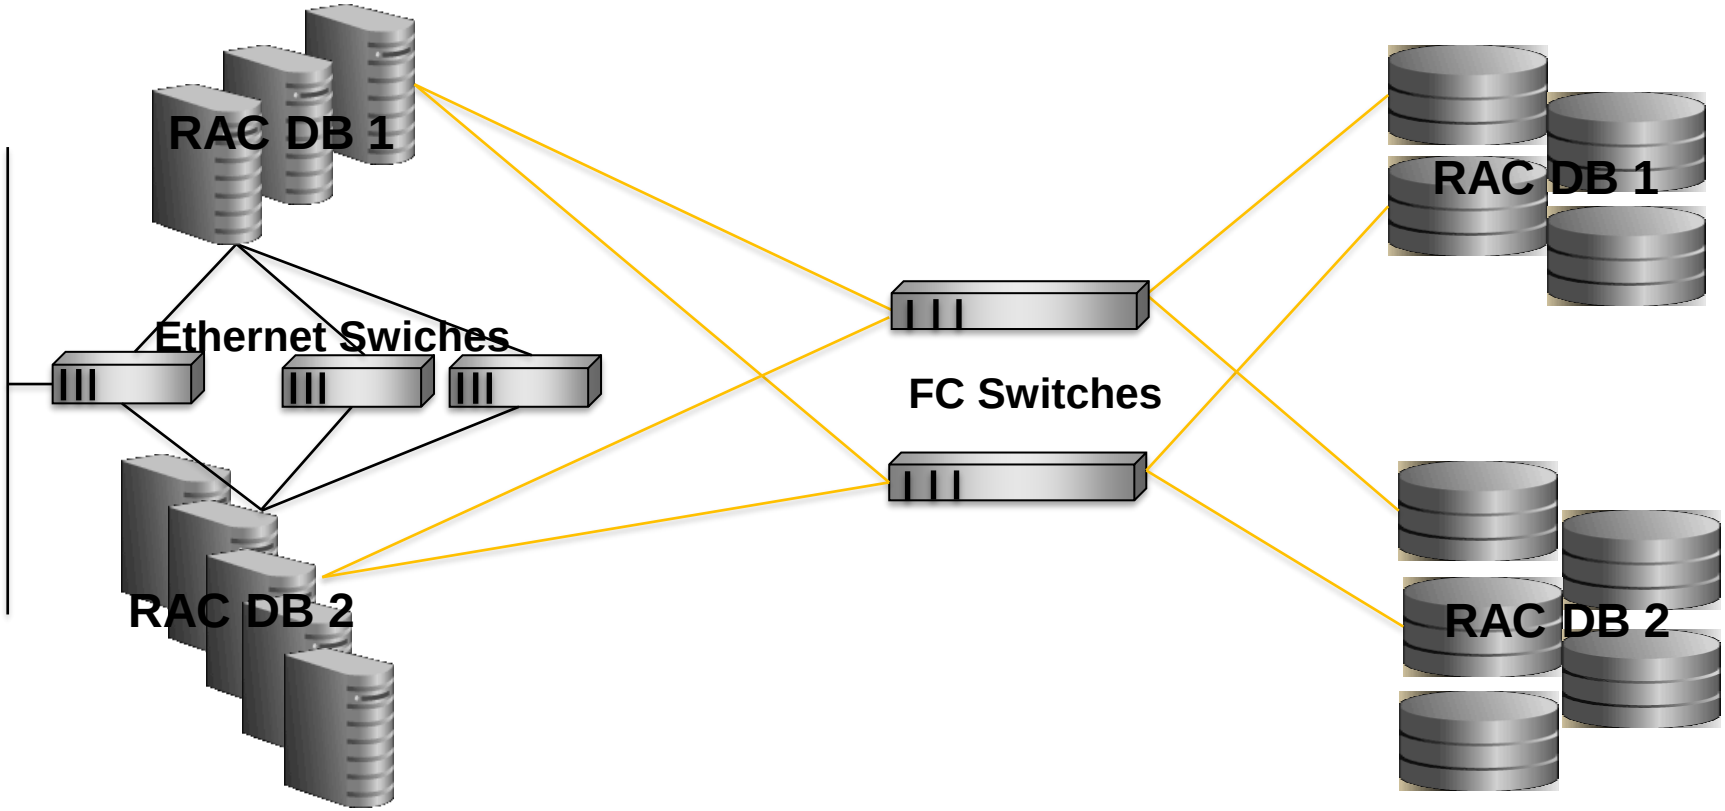

NAS and NetApp at CERN

- Evolution of technology is very interesting
- Performance and capacity
  - 10 GigE connectivity
  - SSD cache to boost IOPS (PAM modules)
  - Allow large DBs with 2TB SATA disks
- Maintenance and reliability
  - 'Filesystem Snapshots' to be used as backup against logical corruption
  - Mature OS and filesystem for stability
  - Redundant controllers for transparent rolling maintenance
  - Support from a major storage vendor
  - Experience at CERN

Solid State Storage

- A revolution in storage
  - Many physics applications spend significant time on random IO read
  - SSD for large increase in IOPS and reduced latency
- Areas of interest
  - Flash-based cache in NAS controllers
  - DB Flash cache feature in 11gR2
    - To be further investigated
  - Entire DB on SSD?
    - For the moment on hold because of cost for multi-TB DBs
- Area in evolution
  - at present a multi-TB DB on SSD is very expensive

Server Specs

- New HW acquisition allows us to profit of technology evolution
- Latest CPUs
  - Although probably 4-cores still best choice for licensing reasons
- More RAM
  - Enlarge Oracle cache to reduce random IO
  - For example servers with 48 GB of RAM
- Faster interconnect:10GigE
  - For storage access
  - Backup to tape

Database technology

- Future needs regarding DB services
  - Review for medium to long term in collaboration with physics experiments and users community at CERN
- NoSQL DBs
  - Preliminary talks and interest from the experiments

Conclusions

- Focus on stability of DB services in 2010
  - Following several years of preparation
  - Infrastructure activities on improving backups, archive, application testing, HW testing
  - Upgrade to 10.2.0.5 performed before 2011 run
- Continuity of DB operations in 2011
  - Priority on running smooth services during data acquisition
  - Preparation for 11gR2 upgrade
  - New HW acquisition, evolution of storage and servers
  - Investigations of new technologies HW and SW

Acknowledgments

- CERN IT-DB group and in particular:
  - Zbigniew Baranowski, Marcin Blaszczyk, Eva Dafonte, Kate Dziedziniewicz, Przemyslaw Radowiecki, Jacek Wojcieszuk, Dawid Wojcik.

- More info:

http://cern.ch/it-dep/db

http://cern.ch/canali