Controls and Software

What did we learn with beam in 2008

Chamonix 2009 - Eugenia Hatziangeli

Acknowledgements

- J. Lewis, I. Kozsar, JC. Bau, C. Roderick, R. Billen, G. Kruk, W. Sliwinski, J. Wozniak, M. Gourber-Pace, N. Hoibian, M. Misiowiec, P. Charrue, J. Lauener, N. Stapley, K. Sigerud, V. Tsaplin, M. Buttner, M. Sobczak, M. Arruat, F. Locci, A. Radeva, S. Deghaye, C.H. Sicard, V. Baggiolini, R. Gorbonosov, D. Khasbulatov, M. Zerlauth, N. Trofimov, …
- Collaborators (FNAL, GSI, Dubna): almost all projects – excellent work

LHC Controls Infrastructure

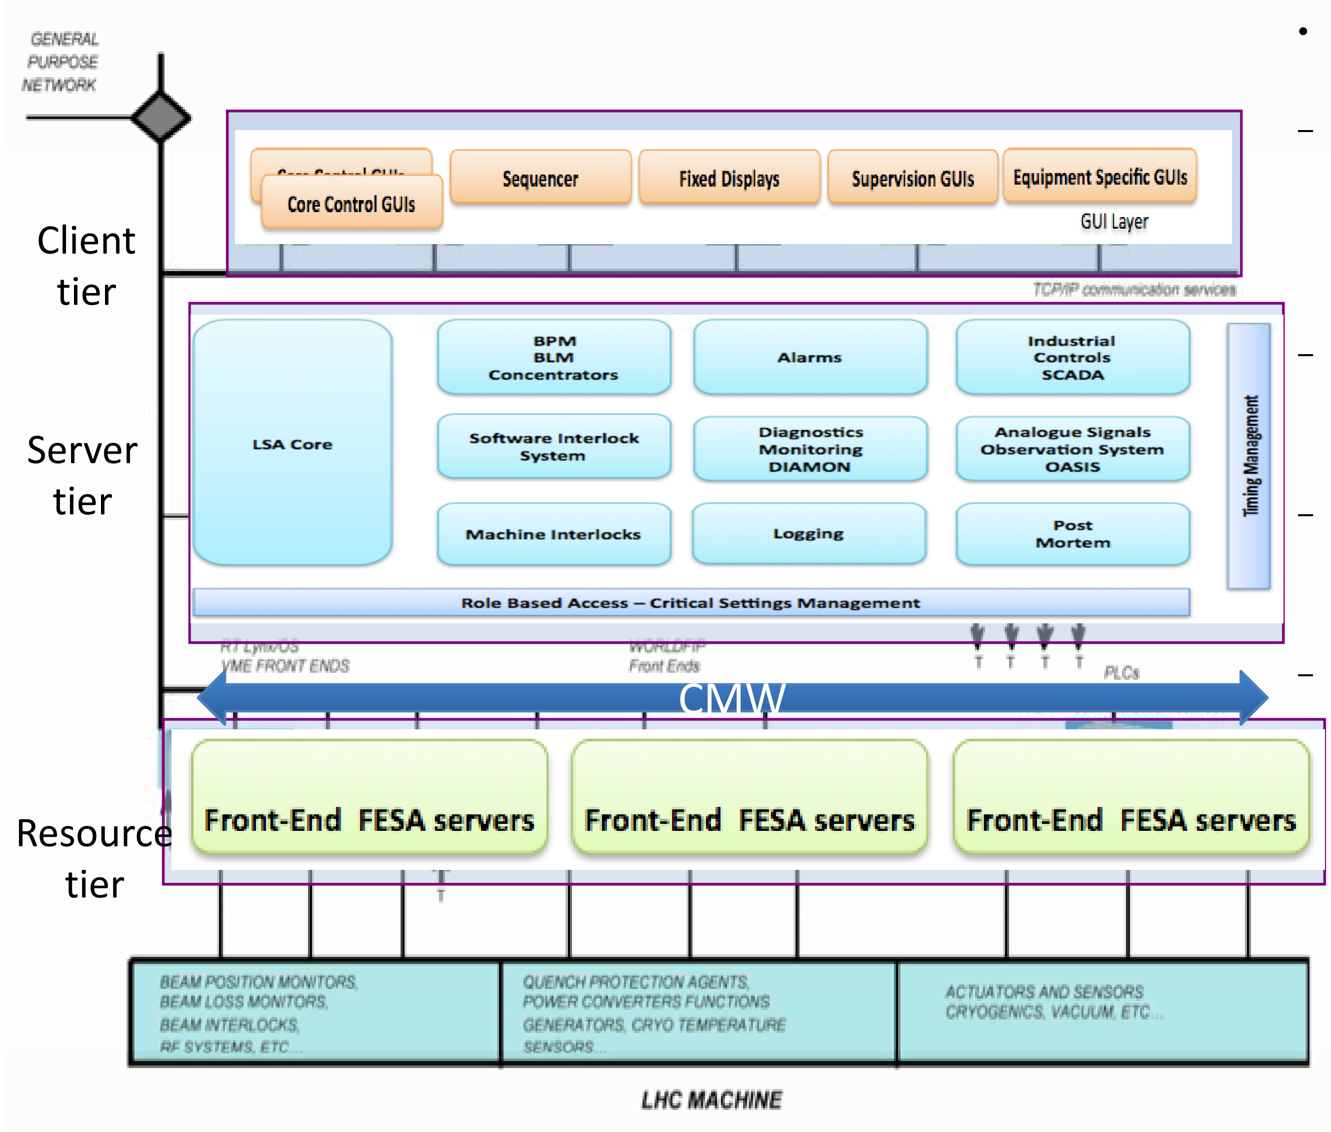

- The 3-tier architecture
  - Hardware Infrastructure
  - Software layers
- Resource Tier
  - VME crates, PC GW & PLC dealing with high performance acquisitions and real-time processing
  - Database where all the setting and configuration of all LHC device exist
- Server Tier
  - Application servers
  - Data Servers
  - File Servers
  - Central Timing
- Client Tier
  - Interactive Consoles
  - Fixed Displays
  - GUI applications
- Communication to the equipment goes through Controls Middleware CMW

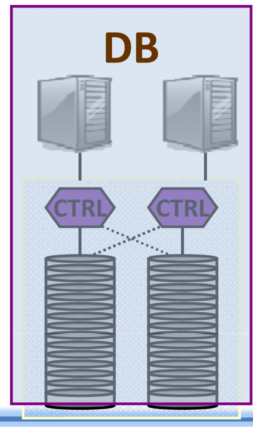

Outline

- **Preparation for beam**
- Start up 2008 challenges and outlook
- What we learnt
- Summary

Preparation for Beam

- All Controls infrastructure deployed and required functionality tested
- Individual high level control tests => dry runs
  - From CCC using controls infrastructure and software
  - Check the functionality of the individual systems
- Full integration tests performed with all systems driven through operational cycles => Full system tests

- Dry runs and TI8/TI2 beam tests –2003 onwards
  - Initial versions LHC Core controls (LSA) tested repeatedly
  - Several CO services put in operations (Fixed Displays, Logging, Software Interlock system)
- LEIR & SPS operational software 2005-06
  - LSA operational deployment
- Injection Dry run Dec 2007
  - Validation of the LHC Timing system – OK
- Full CO scalability tests May-June 2008
  - Aim to identify areas of concern
  - LSA, Logging, Data Concentrators, Middleware, Timing, SIS, LASER, DIAMON, RBAC
- Full dry runs TI8, TI2 31st July, 1st Aug
- Final check out and tests 7-8 Aug
- Five preparatory Injection test Aug – Sep 2008

Outline

- Preparation for beam
- **Start up 2008 challenges and outlook**
- What we learnt
- Summary

LHC Software Architectre - LSA

- LSA is the core control system for LHC, used to
  - Manage equipment configuration & states
  - Manage machine settings
  - Drive the accelerators through their cycle
- **Overall LSA worked very well and provided all required functionality for LHC**
  - It was well established (first used in 2003)
  - Core functionality was tested and improved before LHC first beam
    - LEIR (partial controls, since 2005)
    - SPS and all its transfer lines (full controls, since 2006)
- Major reworking (refactoring) – April '08
  - Introduced new LHC-related concepts **(Beam Process)**
  - Extensions for FGC & new implementations for RF & Collimators
  - **Transactional behaviour** (commit/rollback) for FGC & Collimators
  - **Bottom-up propagation of Trims** to keep the parameter space consistent
  - Mechanism for **trim critical setting** (BLM thresholds)
  - **Data Concentrators** (BLM, BPM)

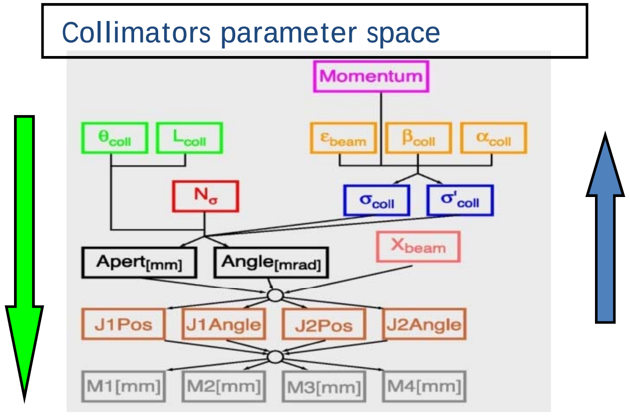

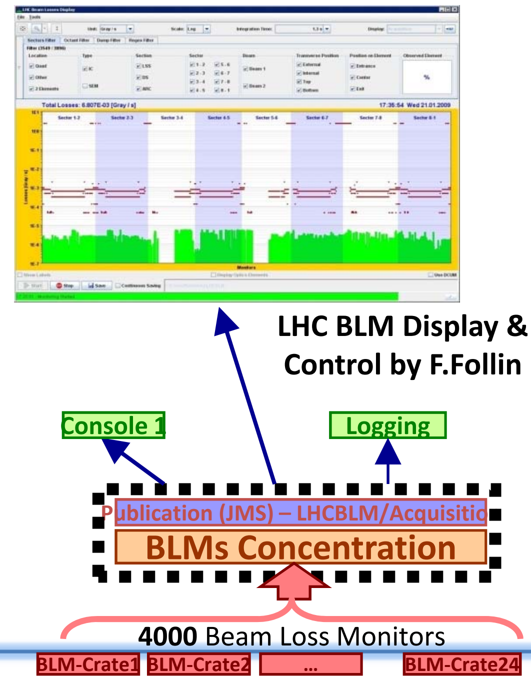

LSA Outlook for 2009

- **Complete online check of any hardware settings**
  - MCS devices
  - Integration with SIS and Sequencer

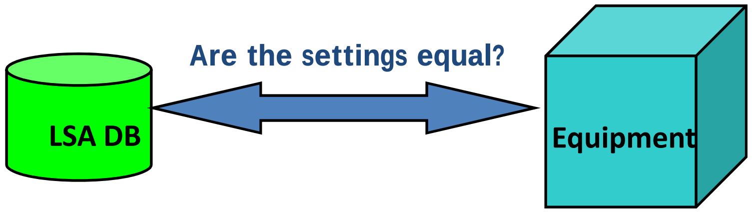

- **Continue preparation of the Test environment**
  - **Test devices and configuration** for driving settings for key equipment (FGC, Collimators, RF, BLM)
  - Integration with the **CO test bed** (FESA, CMW, Timing)

- **Improvements in the core logic and generic applications**
  - **Performance** tuning
  - Better **error reporting**
  - **Diagnostic information to DIAMON**

- **Apply the Finite State Machine concept to operational processes**
  - Implementations for **accelerator mode** & **beam mode**

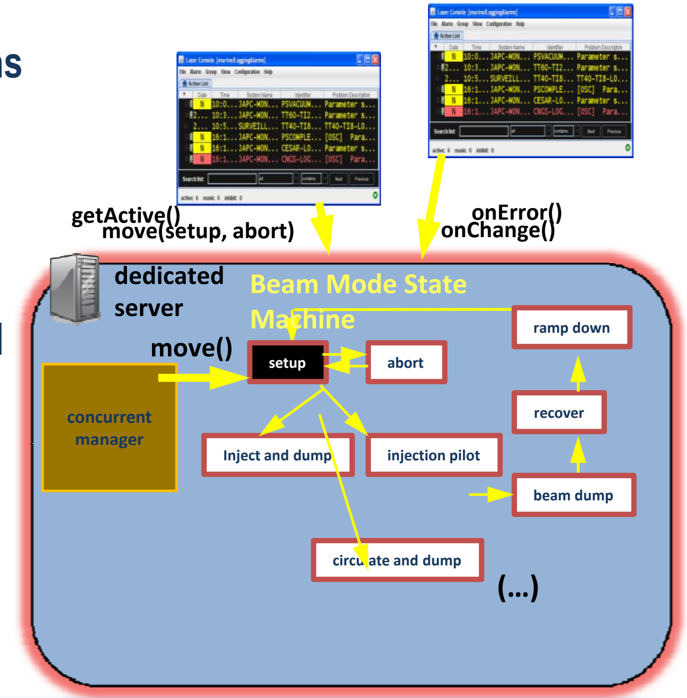

Sequencer

- The Sequencer automates the execution of task sequences (injection, ramp…)
- The Sequencer helps the operators not to forget anything – sequences can be complication and with several steps
  - HWC already relies heavily on the Sequencer
    - Powering Groups of Circuits
    - Ramp & squeeze
- The Sequencer for beam operations was development by a Controls and OP collaboration
  - First used operationally in SPS - end 2007
- The Sequencer for beam operations was used for all controls tests
  - Inject & Dump, Inject and circulate
  - Commissioning of LBDS
  - All necessary sequences for beam were ready
- The Sequencer was proven to be an excellent and vital LHC commissioning tool

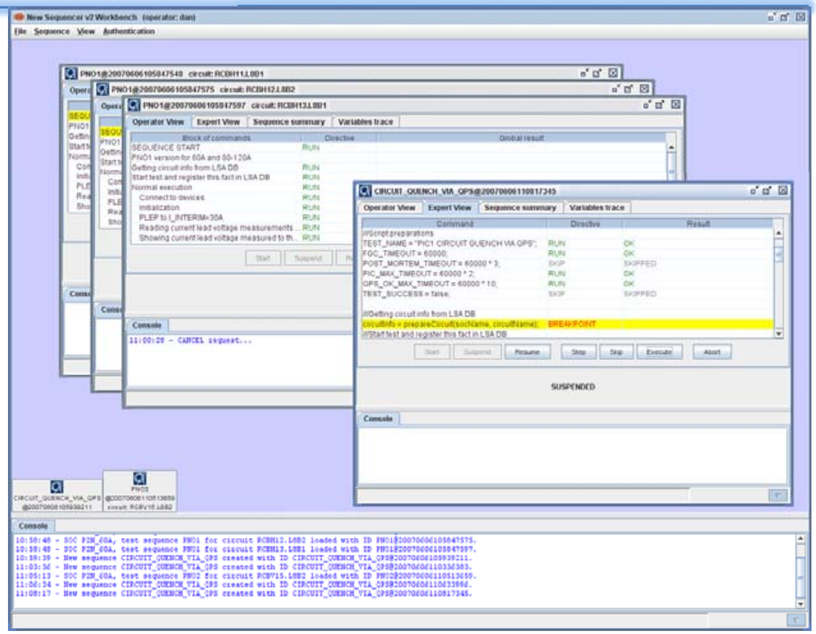

- **Outlook for 2009**
  - Conjugate flexibility with safety
    - Provide different operating modes: **MD (flexible), OP (rigorous)**
    - Ensure some tasks are carried out
  - Improve the User Interface
  - Simplify the sequences and tasks
  - Manage sequences better
  - Integrate sequencer with other applications

Several Reusable Components in LHC

- Excellent experience in LHC with Controls components deployed **previously in other machines**
  - Sequencer (**SPS, HWC**)
  - Fixed Displays (**LEIR, SPS, HWC, CPS**)
  - Software Interlocks System (developed for the **SPS** used operationally since 2007)
  - Alarm system – LASER (**LEIR, SPS, HWC, CPS**)
  - Analogue Observation – OASIS (**LEIR, SPS, CPS**)
  - ….

Timing and Synchronisation

The LHC timing is 'loosely' coupled to the injector chain **only during injection**

- The **data exchange** with the injectors during injection is complex
- **Testing with beam was essential**
  - Timing dry run - Dec 2007
  - Two controls tests – Aug 2008
    - limited impact
    - in the shadow of HWC & Injectors
  - Injection tests Aug – Sep 2008
    - Issues
      - RF synchronization
      - Extraction timing LTIMS
      - Timing tables update

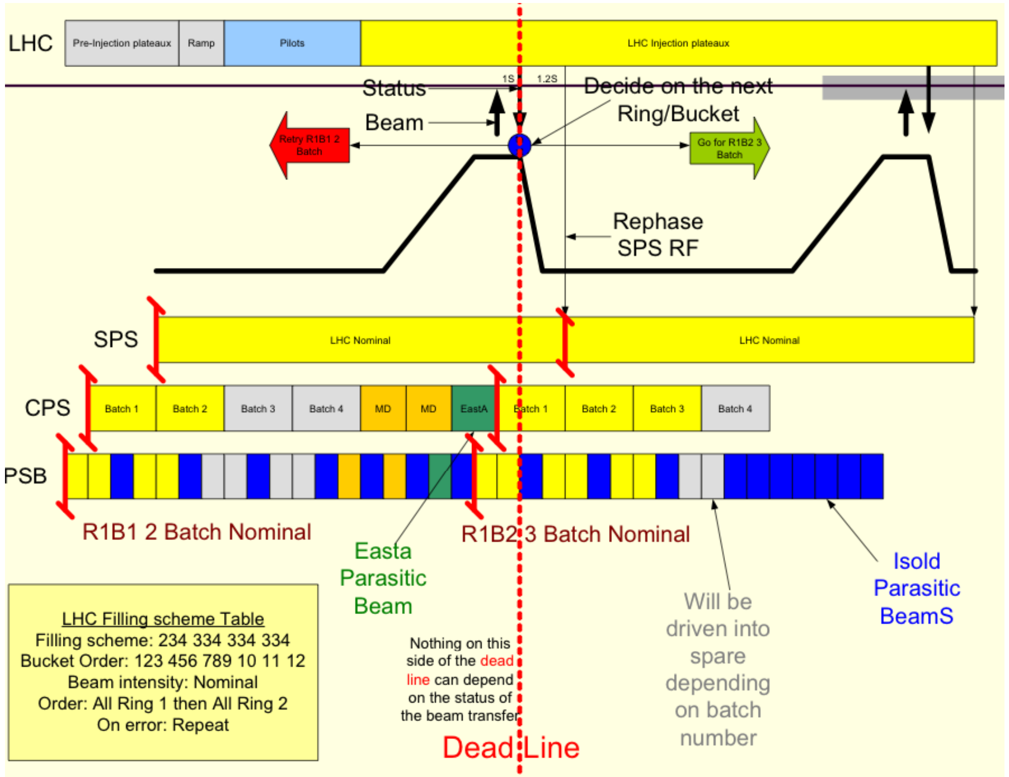

1st Injection Test - RF Synchronisation

- **RF Injection** parameters must be established in the RF system 450ms before the first CPS batch is extracted towards the SPS
- They must also remain stable for at least 1 sec after extraction towards the SPS has taken place

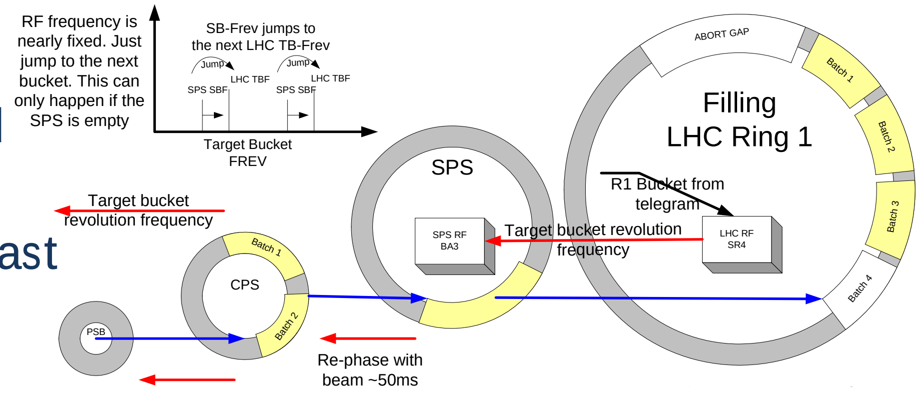

- The actual implementation sent the RF parameters **far too late (25ms)** so the **beam was not correctly transferred** from the SPS to the correct Ring and Bucket
- The cause was **quickly identified** and **resolved on the spot**!!!
  - The test of the LHC filling use case was not sufficient – lack of RF synch went unnoticed as we did not have beam

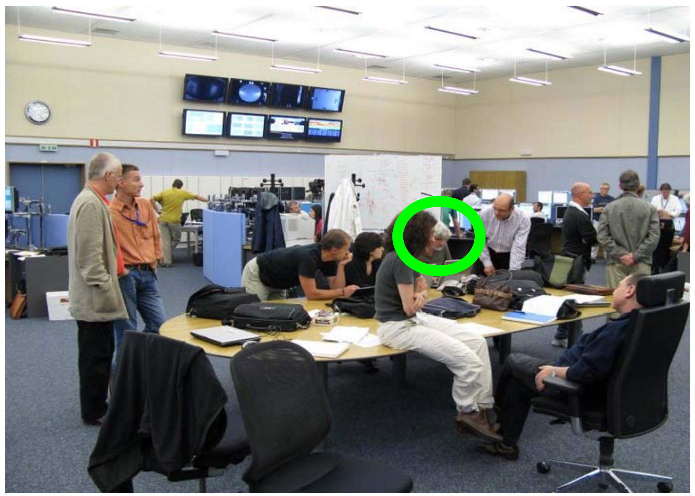

2nd Injection Test - Extraction Timing & Timing Tables

- Local timing **LTIM** in the **FE got corrupted & disabled**
  - The SPS extraction forewarning timing was not sent
- Errors in the table **update protocol** between the **LHC gateway** and the **master and slave LHC timing** crates
  - Injection timing table was not loaded

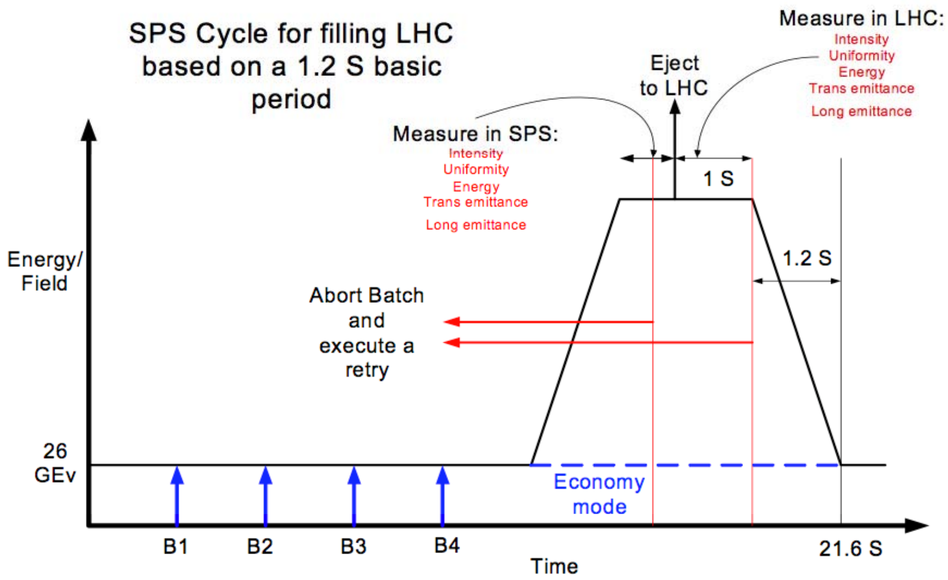

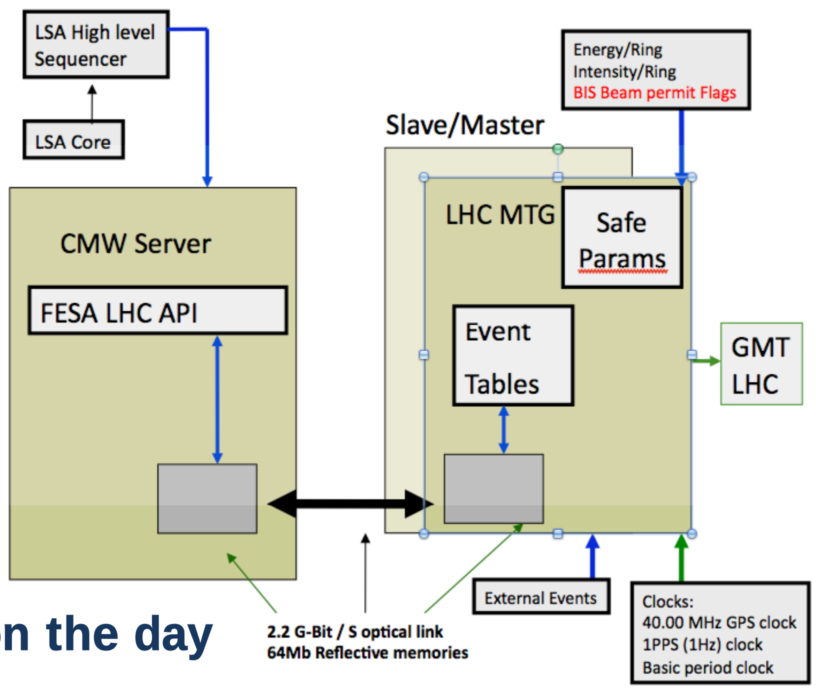

A quick solution was provided for both issues on the day

- The problems correctly fixed the following day

Timing Outlook 2009

- No fundamental changes in the LHC timing
  - All important issues have been resolved
- Consolidation of all LHC timing components – aim for reliability
- LHC timing test bed available to be deployed
  - Intensive testing of the complete system

Databases & Logging - Hitting the Limits

Hardware limit

- **I/O** (end 2007)
  - MeasDB and LogDB on same ageing hardware
  - Purchase **new hardware** & separate into **3 high-availability DB services** – *Mar 2008*
  - Install **additional LDB storage** *Mar-Dec 2008*

- **CPU** (Sep-2008)
  - Simultaneous data writing and extraction MDB→LDB of the same (most recent) data delayed occasionally data transfer
  - Improved filtering code
  - Rationalized MDB backup time

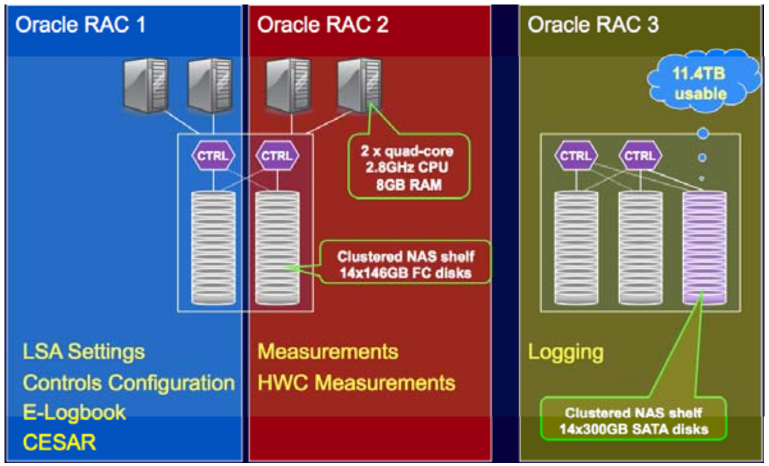

Software limit

- Logging client cannot cater for increasing client requirements, namely LHC BLM
  - ~4000 monitors, **100 000 variables/second**
  - BLM data volume reduced by 30% (*6/9 r.sums*)
  - Data rate reduced by a factor 2 (every 2sec)

- Client **API bottleneck** at serialization of data input
  - Massive // in the data entry API
  - Possibility to log on fixed frequency + on change
  - Dump data on disk if DB down

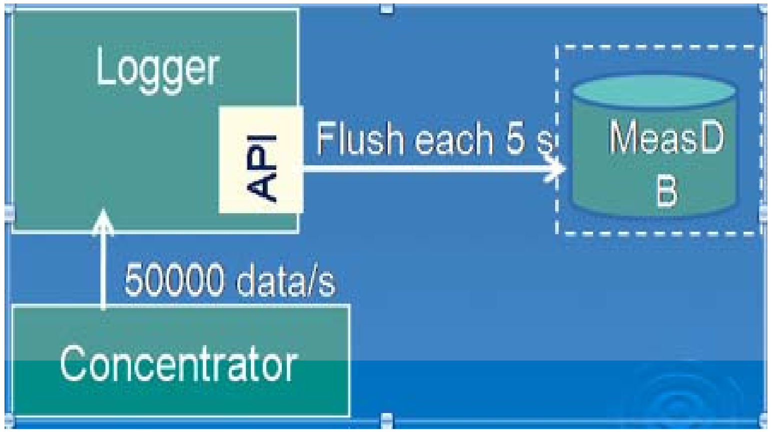

Evolution of the Logging - Data Volume

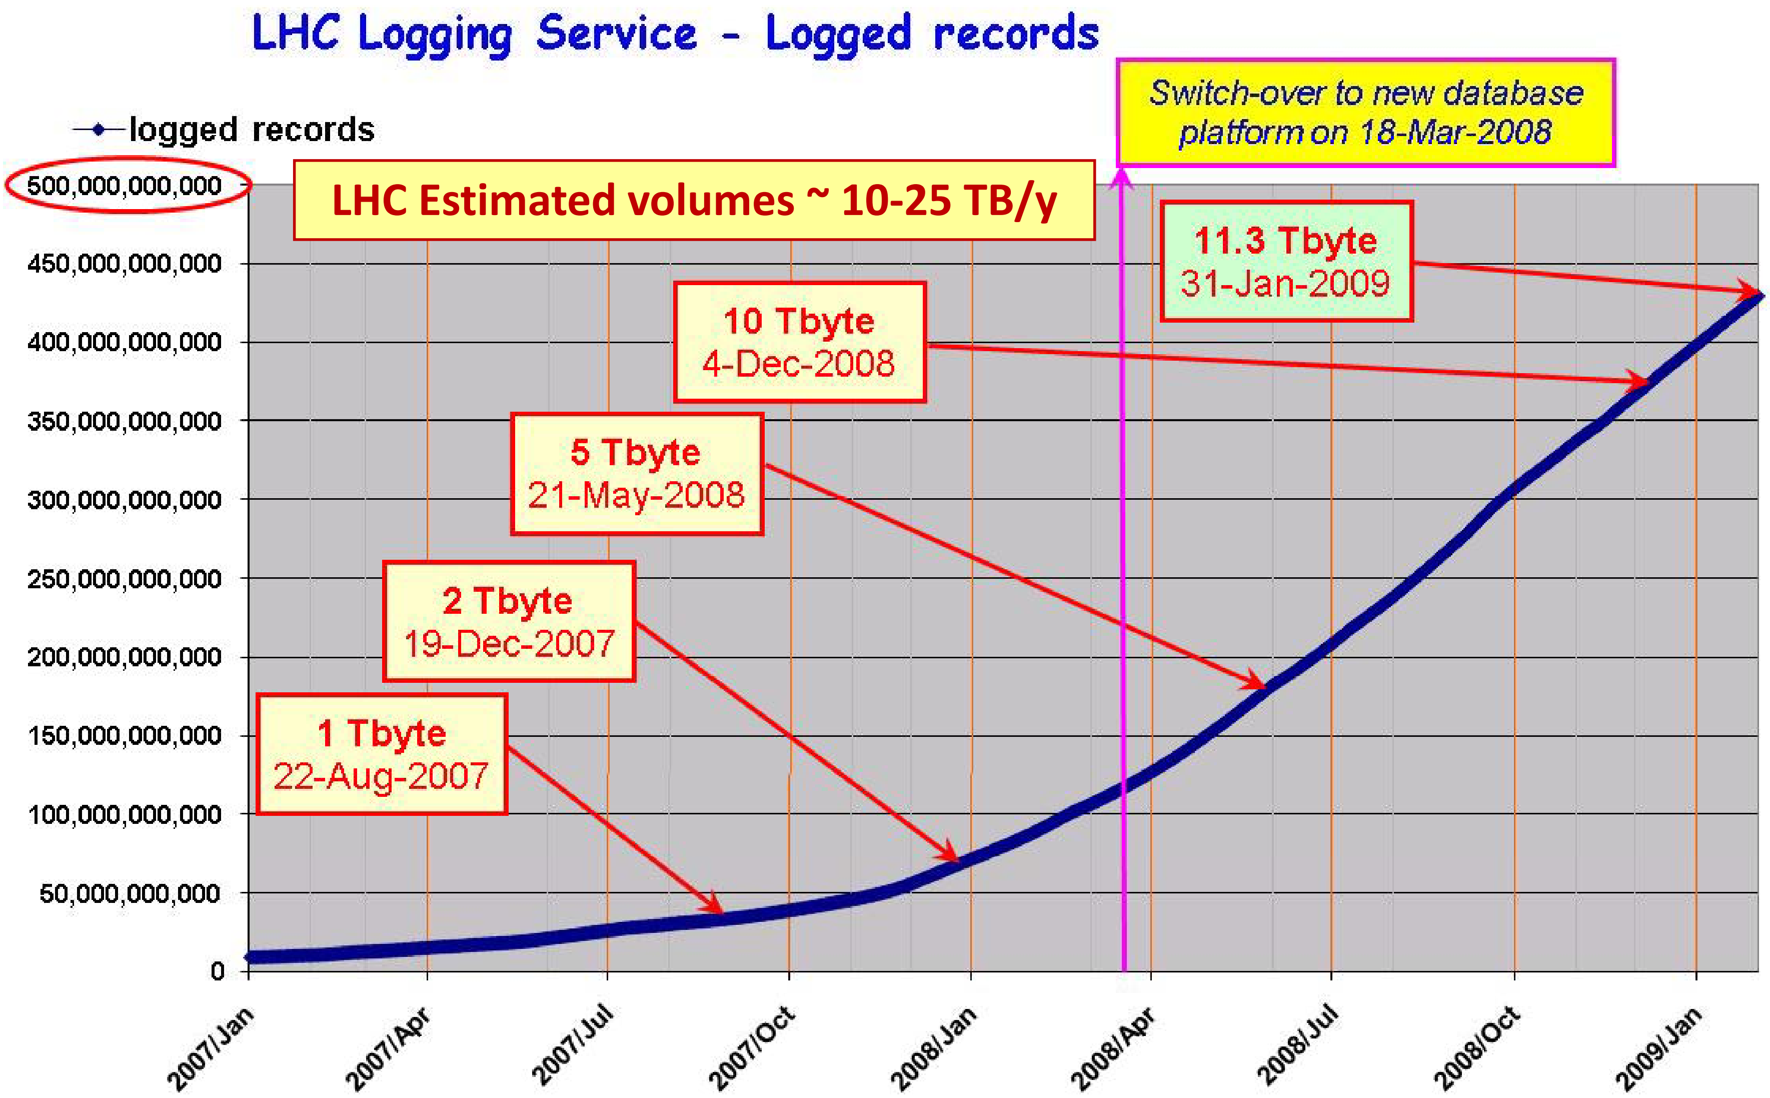

Databases & Logging Outlook 2009

- Address **performance** issues to fulfil nominal logging requirements

Client requirements in terms of data rates (BLM, …)

- **Increased processing power** MDB => better and/or more performing nodes in the cluster – Feb 2009
- **Increase //** => Spread the concentrators & loggers on several dedicated servers
- **Increased hardware performance** => New powerful DB HW

- Provide **maximum availability**

New infrastructure has high-redundancy for high-availability

- **Most interventions are fully transparent**
- **Persist data in local files** if MDB not available (1 day max) with automatic data transfer to MDB

- Scale in function of **data volumes**

LHC Data rates estimated ~10 - 25TB/year

- Disks were installed just-in-time (Mar-Dec 2008)
- Extra disks will be required before end of 2009
- **Do we really need to log all this data?**

Middleware CMW & JMS

CMW

- Blocked/loss of **subscriptions** or blocked **get/set** operations, unable to handle **high load** affected FESA servers (BI, BT, Collimators, Timing,..) and PO
  - Several bugs were discovered and solved at the RDA level and few in the FESA framework
  - In total **6 formal releases** were made in the space of 5 weeks which fixed >10 serious bugs
  - **CMW GW** were put in place to protect simultaneous accesses to (BI) FE

Outlook for 2009

- FESA/CMW servers are in a much better shape now than ever
- Few known RDA issues remain to be solved

JMS - victim of its own success

The **usage of JMS** was growing beyond expectations (>100 connected physical clients, 25/CCC rest from offices, GPN,..)

- Insufficient throughput
  - Data delivered: 10-15 MB per second
  - Around 4 million messages a day
- Broker **unstable** or connections **frozen**

Mainly reactive approach in 2008

- Tune configuration, upgrade hardware & JMS software,...)
- Protection of resources for CCC
  - Primary broker for CCC, secondary broker for GPN

Outlook for 2009

- Increased manpower for support
- Provide better **monitoring** of the JMS Brokers & internal system parameters
- Create a consolidated set of test to **test** future JMS broker updates and configurations
- Investigate and install **clustered broker** sets for increased reliability (failover)

Role Based Access Control - RBAC

- RBAC works by giving people **ROLES** and assigning **PERMISSIONS** to ROLES to access device properties (**Authenticate & Authorise**)
- Integrated into the BE controls middleware and a collaboration project with LAFS

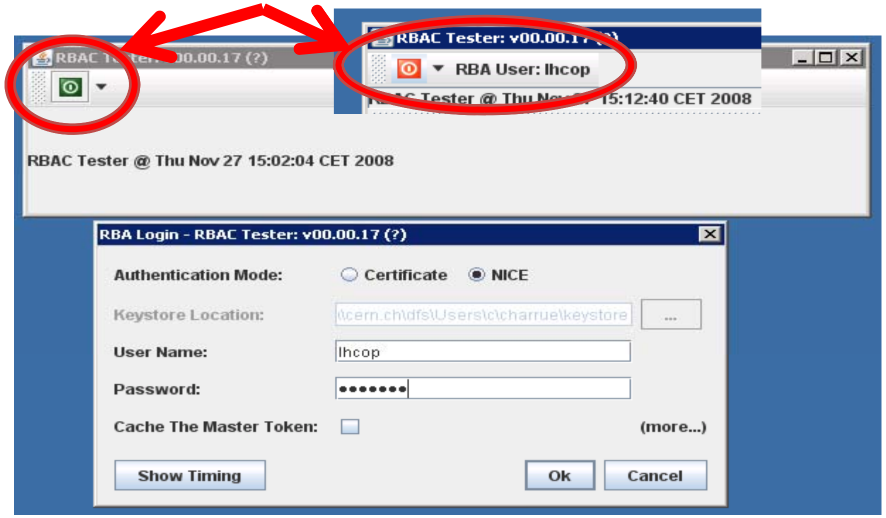

- Deployed in **2008** for the LHC dry-runs and start up
  - Objective: RBAC all LHC Applications
- **75 %** of all applications were RBACed
- **Most of the equipment** were RBACed (PO, BI, BT)
  - RF & Collimators were getting started
- RBAC was quickly embraced by equipment groups
  - High expectation to get protection from unauthorized access to their equipment

Outlook for 2009

- Ensure all operational GUIs and equipment are RBACed
- Make sure the RULES and ROLES are correct and contained (LHC Control Security Panel)
- Default access mode will be **strict** in 2009
  - Each application will need an RBAC token

Outline

- Preparation for beam
- Start up 2008 challenges and outlook
- **What we learnt**
- **Summary**

So … what did we learn?

- **Testing is essential to success**
  - The majority of Controls infrastructure was thoroughly tested over many years 2003 ->
    - Dry runs, cold check out, full machine checkouts
  - The core controls (LSA, SIS, …) and services (Sequencer, Logging, …) were used operationally in other machines
    - Constant exercising of the software – most bugs/inadequacies were surfaced
    - Functionality was well tuned
    - Refactoring made software robust and simpler each time
- Nevertheless we had few issues with systems used **1st time with beam** or **in this scale**
  - Controls is the first infrastructure needed to be available, difficult to *steal* time from user to test our own
  - For certain system only the beam revealed the problem (Timing)
  - Not enough integration testing facilities – room for improvement
    - Service libraries (CMW, Timing, FESA, …) that affect several systems (BI, XTIM, BT,..)
  - Testing without correct initial conditions will not reveal all the problems
    - Scalability tests gave us some indications of the weak components
      - When full LHC load hit the controls system we saw clearly the bottlenecks
- To avoid the same mistakes again
  - **Invest in proper** test facilities – Controls test bed
  - Ask for **dedicated Controls testing time** in the Accelerator schedule
    - already planned for Injectors

Summary

- Overall **the Controls infrastructure has worked well**
  - Thanks to early deployments, reusability across machines, cold checkouts, dry runs, full machine checkouts which were extremely useful to identify and correct bugs/features before startup with beam
- A lot of work has been **invested** in **functionality** and **safety** (RBAC, MCS, SIS)
  - All necessary functionality for beam was there
  - Several core systems have been protected against damage by ignorance
- Still many new requirements/functionality to be implemented in 2009
- **Strong investment in reliability** of the Infrastructure
  - A monitoring infrastructure **DIAMON** (in unison with LASER) is put in place to aid operations and CO experts to diagnose/repair controls problems
  - Individual test set ups in labs (timing, LSA)
  - Controls (integration) **Test Bench** in preparation
- Not to forget the **Injector Renovation**
  - The work to renovate the LHC Injectors with the present Controls Infrastructure is advancing and the first validation of the architecture is positive

Thank You