

WHAT DID WE LEARN WITH BEAM IN 2008? CONTROLS & SOFTWARE

E.Hatziangeli, CERN, Geneva, Switzerland

Abstract

A large number of components of the LHC controls infrastructure have been intensively exercised during several dry runs and the injection synchronization tests in 2008, in preparation for the LHC start up with beam. All the prior tests were extremely useful to validate the basic functionality, to identify areas of weakness and to prepare the controls to a stable and well functioning infrastructure, which supported the successful LHC start up.

This presentation will summarize the outcome of the controls commissioning with emphasis on major components like LSA, Timing synchronization, middleware, logging, and Role Based Access, the issues observed and the improvements to be implemented before the next LHC run in 2009.

LHC CONTROLS INFRASTRUCTURE

The LHC hardware controls infrastructure is based on the classic 3-tier architecture (Fig.1).

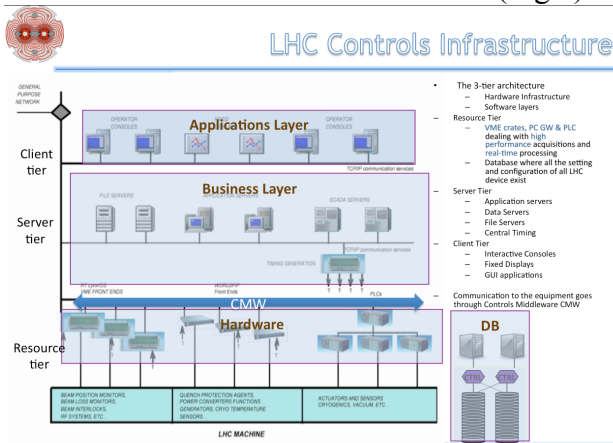


Fig1. The LHC hardware architecture

The LHC software infrastructure, which is also based on a 3-tier architecture, maps exactly on the above hardware controls infrastructure (Fig.2).

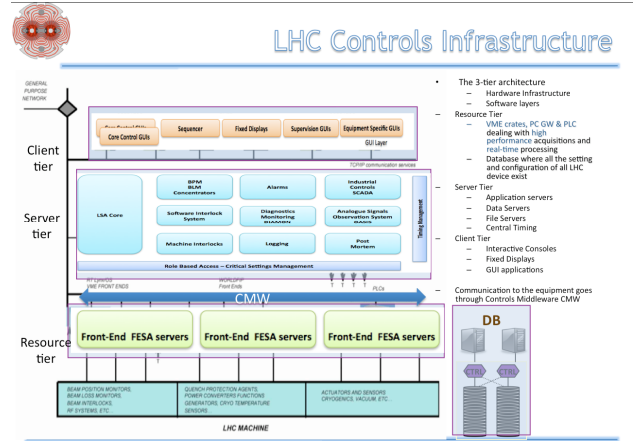


Fig.2 The LHC Software infrastructure

The resource tier contains VME crates, PC gateway & PLC dealing with high performance acquisitions and real-time processing, plus the databases where all the setting and configuration of all LHC devices exist. The Server tier, known as Business layer, is made of the Application, Data and File servers and the Central Timing system. Lastly, the Client or Application tier, supports all the operational GUIs used by operators to control the machine running on the interactive consoles and a series of Fixed Displays which constantly monitor and display important parameters of the accelerators.

PREPARATION FOR BEAM

The objective of controls was to have, if possible, all the infrastructure deployed before the start-up and test repeatedly all the required functionality. Clearly the reason was that when we finally get beam we don't want to waste time solving controls problems that could have been resolved before. This was achieved with high level controls tests done by operators and supported by controls experts, in the CCC using the whole hardware and software infrastructure, checking the functionality of the individual systems, i.e.:

- Several individual high level control tests which are known as dry runs
- Full integration system tests performed with all systems driven through operational cycles

During the dry runs and full systems tests the following tests were executed to commission all controls mission critical systems for the LHC start-up:

- 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
 - Successful first validation of the LHC Timing system
- Full Controls scalability tests May-June 2008
 - Aim to identify areas of concern under the full LHC load
 - Systems tested were 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

START UP 2008 CHALLENGES

Clearly putting all the controls components together to form a coherent operational suite, which will work correctly together, was a challenge, keeping in mind that changes due to fine-tuning of requirements or due to issues meant individual/system and integration tests need to be repeated again.

Some of these components are singled out due to their successful deployment or because they posed certain challenges worth mentioning during the commissioning period

LHC Software Architecture – LSA

LSA is the core control system for LHC, used mainly to manage machine settings, equipment configuration and states, and to drive the accelerators through their cycle. Overall the LSA software suite worked very well and it provided all required functionality for LHC. A major factor to this success was that LSA was already well established (first used in 2003), its core functionality was tested and improved before LHC first beam. In particular, LSA was used to provide partial controls for LEIR since 2005 and full controls for SPS and all its transfer lines since 2006.

Nevertheless LSA kept evolving to ensure LHC specific requirements were implemented. As it is usual to consolidate a system with the experience gained from the previous years of usage, in 2008 LSA went through a major refactoring to prepare for LHC with the following outcome:

- Introduced the concept of Beam Process for LHC in addition to existing Super Cycle concept
- Extensions for drive settings for FGC and new implementations for RF to take care of mapping between RF channels and collimators parameter space
- Transactional behaviour (commit/rollback) for FGC and collimators (fig.3)

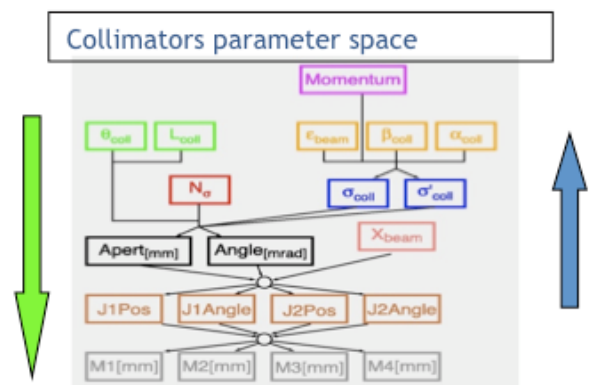


Fig 3. Collimators parameter space

- Bottom-up propagation of Trims to keep the parameter space consistent – a feature particularly interesting for collimators
- Mechanism to trim critical setting (e.g. BLM thresholds)

- Introduction of components known as data concentrators which acquire data from BLM and BPM systems and distribute the data to all clients (fig.4)

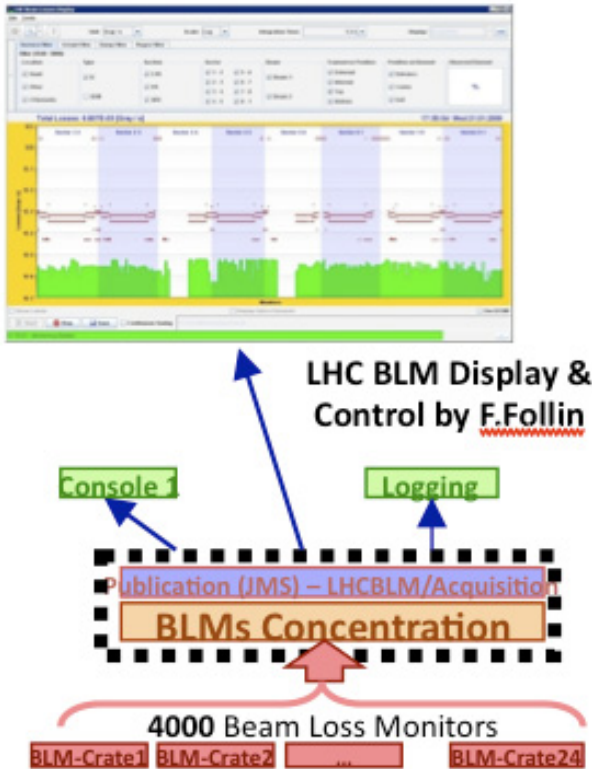


Fig 4. Data Concentrators

LSA Outlook for 2009

In 2009, the main work in the LSA project will be:

- Complete the online check of any hardware settings such as MCS devices and the integration with SIS and Sequencer
- Continue the preparation of the Test environment in order to be able to test devices and configuration for driving settings for key equipment (FGC, Collimators, RF, BLM) with an integration with the CO test bed (FESA, CMW, Timing)
- Improvements in the core logic and generic applications in the fields of performance tuning, better error reporting or diagnostic information to DIAMON
- Apply the Finite State Machine concept to operational processes with concrete

implementations for accelerator mode and beam mode.

The Sequencer

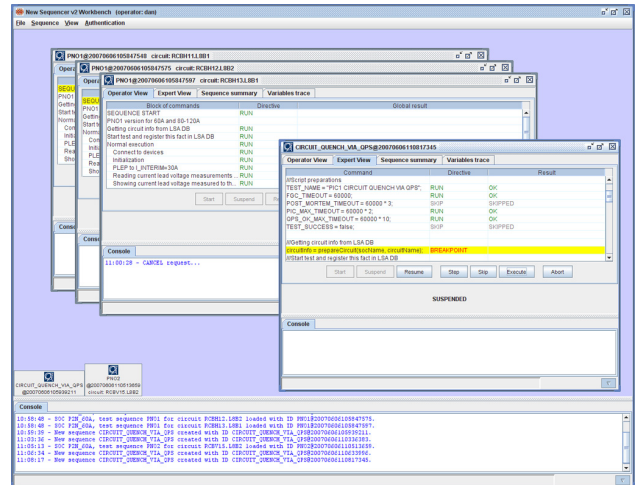


Fig 4. The sequencer

The Sequencer is an important tool, which automates the execution of complex task, known as sequences (injection, ramp...) (Fig. 4)

It helps the operators not to forget any essential steps, as sequences can be complicated and with several layers. The LHC Hardware Commissioning (HWC) operation already relies heavily on the Sequencer with the support for “Powering Groups of Circuits” and “Ramp & squeeze”.

The Sequencer for beam operations was a development in collaboration with Controls and OP and it was first used operationally in SPS in the end of the 2007 run. For the LHC start-up and beam operations all necessary sequences were ready and the Sequencer was used for all controls tests, such as Inject & Dump, Inject and circulate, Commissioning of LBDS.

The Sequencer was proven to be an excellent and vital LHC commissioning tool for operations.

Outlook for 2009

In preparation for 2009 and based on the experience of the commissioning we have seen the most important issue is to conjugate flexibility which is important for commissioning with safety which paramounts with operations and to ensure that specific task cannot be omitted by mistake. Therefore the main development activities for the sequencer are :

- Combine 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

Clearly LSA and Sequencer were one of the many reusable components of the controls infrastructure to be used in LHC. We have had excellent experience in LHC with Controls components deployed previously in other machines.

- Fixed Displays (**LEIR, SPS, HWC, CPS**)
- Software Interlocks System (developed for the SPS and 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 and this coupling takes place only during injection (Fig 5)

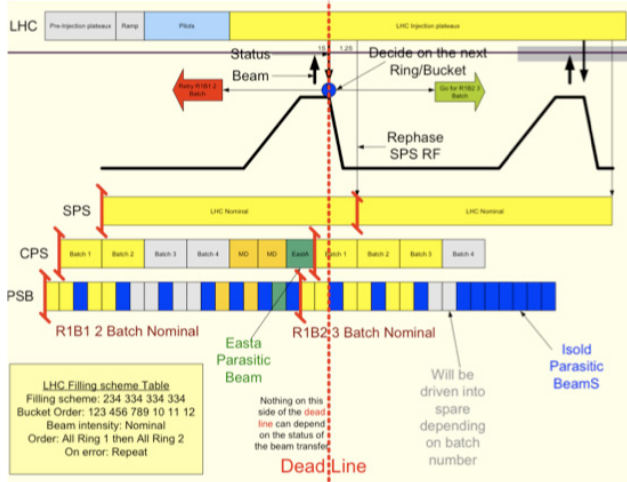


Fig 5 : The complex LHC timing

The data exchange with the injectors during injection is very complex and testing the timing with beam was essential. In December 2007, a timing dry run was prepared and it was successful, while in August 2008, two controls

dedicated tests were organized with limited scope and impact as they were done in the shadow of HWC and the Injectors operations. Despite these test, a few serious problems were encountered during the LHC injection tests related with the following domains:

- RF synchronization
- Extraction timing LTIMS
- Timing tables update

RF synchronization issue

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 (Fig 6)

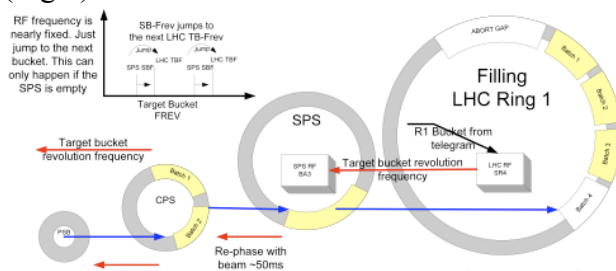


Fig 6. RF Timing synchronisation

Despite expectations, during the 1st injection test the actual implementation sent the RF parameters far too late (by 25ms) so the beam was not correctly transferred from the SPS to the correct Ring and Bucket.

The cause was quickly identified and resolved but it had proven that the tests of the LHC filling use case was not sufficient to identify the deficiencies of the timing systems, as the lack of RF synchronisation went unnoticed without the presence of beam.

Extraction timing LTIMS issue

During the 2nd injection tests, the local timing LTIM in the front-end server became corrupted and disabled resulting therefore that the SPS extraction forewarning (Fig. 7) timing was not sent out.




Fig 7. SPS extraction forewarning event

Timing tables update issue

Errors in the table update protocol between the LHC gateway and the master and slave LHC timing crates (Fig. 8) caused the loading of the Injection timing tables to become erratic.

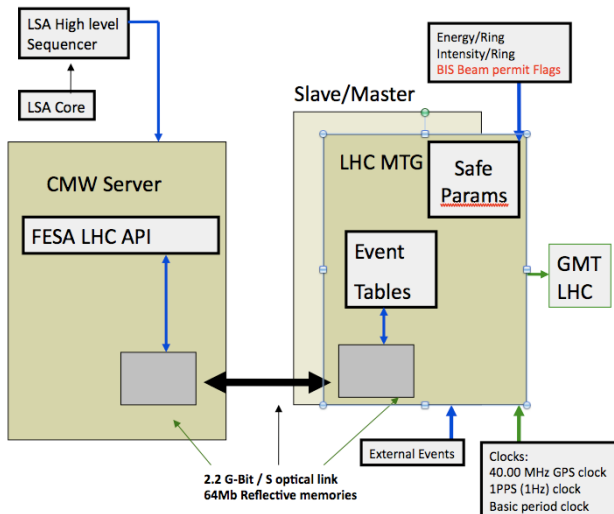


Fig 8. The LHC timing gateway & the master and slave timing crates

Both issues were quickly identified and solved immediately to facilitate the 2nd injection tests.

Outlook for 2009 for the LHC timing

No fundamental changes in the LHC timing are foreseen, as all the important issues have been resolved.

With an aim for reliability, a consolidation of all LHC timing components is being done, together with the deployment of an LHC timing test bed which will be available to allow for intensive testing of the complete system.

Databases and Logging – Hitting the Limits

The control system is highly data centric and the databases are extremely critical components of the control system of all CERN accelerators. They are essential for settings and configuration, short-term measurement data service for one week and data logging service for long-term archive.

In 2008 we have seen both hardware and software limits being reached for both the measurement and logging service.

Hardware limit

End of 2007, the measurement and logging service hit the I/O limit due to being run on the same ageing hardware. New hardware has been purchased in March 2008 and the opportunity was used to separate the existing service into 3 high-availability database services. In addition, having seen the speed by which the disk space was consumed, additional logging storage was installed at that time with a second disk installation towards the end of last year (Fig 9).

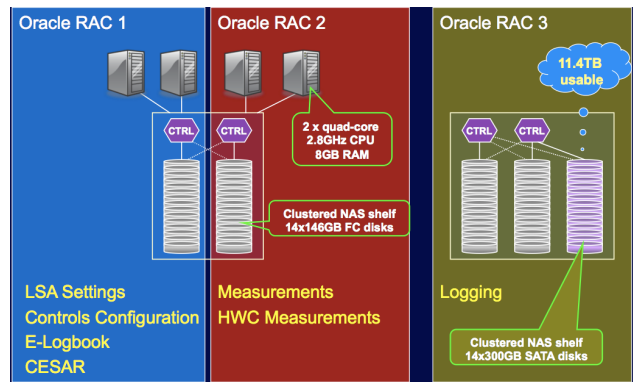


Fig 9. The Controls Configuration, Measurement and Logging database services

During the LHC commissioning, the CPU limit was reached as simultaneous data writing and extraction of the same (most recent) data from the measurement database to the logging database. This resulted into occasional delayed data transfer. An improved filtering code and a rationalized MDB backup time resolved this issue.

Software limit

From the software side (Fig 10), as soon as the Beam Loss Monitors (BLM) were deployed in mass, the logging clients could not cater efficiently the increasing client requirements both in data volume and data rate. Namely for the LHC BLM this amounted to ~4000 monitors and 100 000 variables/second. A reduction of the data volume by 30% (6/9 running sums were logged) and a reduction of the data rate by a factor 2 (every 2sec) solved temporarily this issue.

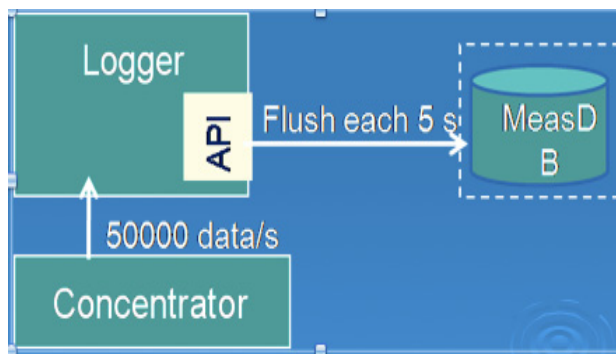


Fig 10. The BLM data concentration and logging

The logging client API was showing a bottleneck at the serialization of the data input. To overcome this a massive parallelisation was introduced in the data entry API, with the possibility to log on fixed frequency and on change. In addition to safeguard against temporary database unavailability a new functionality became available for the LHC start-up, which allowed for measurement data to be stored on local disk (max. 1 day) if the central database is not available.

The estimation of the volume of the data being logged up to the incident in September was 10TB/year during LHC exploitation (Fig 11).

With the new requests for BLM, BPM, QPS without filtering at 5Hz, QPS++, the data volume will increase to 25TB/year. At present we have arrived to 11TB of data stored in the logging database.

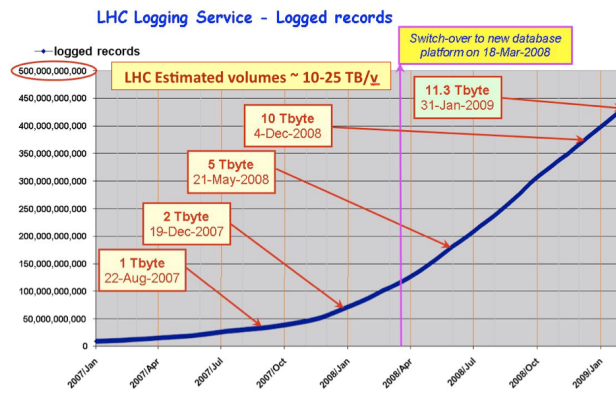


Fig 11. Evolution of the Logging – Data Volume

Outlook for 2009 for the Databases and Logging

Address performance issues to fulfill nominal logging requirements:

The following solutions will address the client requirements in terms of data rates (BLM, BPM, QPS, ...):

- Increased processing power of the measurement database by introducing better and/or more performing nodes in the cluster. This is planned on Feb 2009
- Increase parallelism and spread the concentrators and logging processes on several dedicated servers
- Increased hardware performance by procuring new more powerful database hardware.

Provide maximum availability

The present infrastructure has already high-redundancy for high-availability:

- Most interventions are fully transparent, while a minimum of scheduled security related interventions
- New functionality is implemented to persist data in local store if the measurement database is not available (1 day max) with automatic data transfer to the database once it becomes available again.

Scale in function of data volumes

At present with the existing requests and the foreseen new clients for logging the LHC data volume is estimated between 10 - 25TB/year. To ensure we have the necessary disk space available:

- Disks were installed in time for the startup in 2008
- Extra disks will be required before end of 2009

Organize a discussion with data providers to clarify and decide if all this data volume, which is requested to be logged, is really needed.

The Controls Middleware CMW

Several issues were discovered in the CMW infrastructure during its massive deployment in LHC, such as blocked or loss of subscriptions or blocked get/set operations, inability to handle high load which affected FESA servers (BI, BT, Collimators, Timing) and FGC front end servers.

Several errors were discovered and solved at the CMW RDA library level and few were found in the FESA framework. In total 6 formal releases were made in the space of 5 weeks which fixed more than 10 serious faults. In addition, in order to protect the front-end servers against several clients accessing the same information, CMW gateways were put in place to serve simultaneous accesses to Beam Instrumentation front ends.

Outlook for 2009 for CMW

The FESA/CMW servers are now stabilised and all the critical issues have been resolved with only a few remaining RDA issues which are in progress to be completed.

The Java Messaging Service - JMS

The usage of the JMS service during 2008 was growing beyond expectations (more than 100 connected physical clients, 25 from the CCC, the rest from offices, GPN, ...). The data concentrators for BLM and BPM were the two critical systems in LHC which relied on JMS for data distribution to the logging database and to operational GUI application in the CCC.

During the LHC commissioning and beam injection tests, we have discovered that the present throughput was insufficient with 10-15 MB per second of actual data delivered with around 4 million messages a day.

This resulted in unstable and erratic behaviour from the JMS broker with several client connections frozen and unable to receive data. As the problem was located in the actual configuration of the JMS brokers being insufficient for the final load, the solution taken for the LHC start up was:

- Fine tuning of the configuration,
- Upgrade of the hardware platform
- Upgrade to the latest better performing version of the JMS software
- Install protection of resources for CCC with the deployment of a primary broker for CCC and a secondary broker for GPN

Outlook for 2009 for JMS

In order to be able to provide the appropriate level of JMS service for all the mission critical system that rely on it, the following solutions were put in place:

- The manpower for support has been increased
- Better monitoring facilities of the JMS brokers and the internal system parameters are being implemented
- A consolidated set of test to test future JMS broker updates and configurations
- Investigate and install clustered JMS 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 (by Authenticate and Authorise). It is integrated into the BE controls middleware and is the result of a collaboration project with LAFS (LHC software collaboration with Fermilab).

RBAC was first deployed in 2008 for the LHC dry-runs and start up with the objective of having RBAC in all LHC operational applications. Today, 75 % of all applications were RBACed (Fig. 12)

Most of the LHC equipment were also RBACed (PO, BI, BT) while RF and Collimators were in the stage of defining their own ROLES and RULES. Equipment groups quickly embraced RBAC as there is high

expectation to get protection from unauthorized access to their equipment.

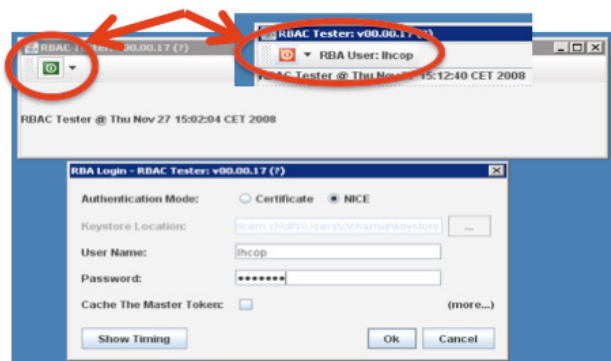


Fig 12: application with the RBAC component

Outlook for 2009 for RBAC

The work in RBAC in 2009 is to ensure all operational GUIs and equipment are now fully RBACed. In addition via the LHC Control Security Panel (LCSP) to make sure the RULES and ROLES are correct and contained; avoiding needless proliferation of ROLES and RULES.

One important change with respect to the LHC start-up in 2008, is that it is planned to start with the default RBAC access mode in STRICT for 2009. This implies that each operational application will need to have a valid RBAC token.

WHAT WE LEARNT

Clearly one important lesson is that testing is essential to success.

The majority of Controls infrastructure was thoroughly tested over many years starting in 2003 onwards. This was done in the form of dry runs, cold check out, and full machine checkouts.

The core controls (LSA, SIS, ...) and services (Sequencer, Logging, ...) were used operationally in other machines than LHC and prior to LHC start up. This allowed for constant exercising of the software resulting in most of the errors or inadequacies to be surfaced. In addition, their functionality was well tuned through their continuous usage. Refactoring made software robust and simpler each time.

Nevertheless we had few issues with systems used for the first time with beam or in this scale. Keeping in mind that Controls is the first infrastructure needed to be available, this makes

it rather difficult to schedule time from the Controls users to test the Controls infrastructure once it is deployed.

Moreover for certain system testing without beam was not sufficient; only the presence of beam revealed the problem (Timing).

In terms of testing, it is clear that better controls integration testing facilities will ensure that future release of critical service libraries like CMW, Timing, FESA will not adversely affect client equipment servers or operational applications.

While testing is of primordial importance, testing without correct initial conditions will not reveal all the problems. Scalability tests gave us some indications of the weak components but when the full LHC load hit the controls system the areas of bottlenecks were clearly revealed.

To avoid the mistakes of the past, an investment in proper test facilities and in a controls-dedicated test bed is underway. Moreover dedicated Controls testing time in the Accelerator schedule should be foreseen (already planned for Injectors) to ensure the controls infrastructure is exercised adequately and under all feasible scenarios before it is handed over to Operations.

SUMMARY

Overall the Controls infrastructure has worked well, thanks to early deployments, software reusability across machines, cold checkouts, dry runs, and full machine checkouts which were extremely useful to identify and correct issues or shortcomings before the LHC start-up with beam.

A lot of work has been invested in providing the requested functionality for the start-up with an emphasis on safety (RBAC, MCS, SIS).

As a results all necessary functionality for beam was there and several core systems have been protected against damage by ignorance.

In 2009 there is still an important amount of work to fulfil the new requirements and to provide the enhancements in existing functionality based on our experience of the start-up. Strong investment will be made in reliability of the controls infrastructure.

A monitoring infrastructure DIAMON (in unison with LASER) is put in place to aid operations and CO experts to diagnose/repair controls problems efficiently. Individual test set ups in labs for Timing and LSA are being prepared in parallel with the Controls (integration) Test Bench.

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): involved in several controls projects

REFERENCES

Detailed and up-to-date information on LHC Controls Infrastructure <http://wikis.cern.ch>