ACCELERATOR DATA FOUNDATION: HOW IT ALL FITS TOGETHER

R. Billen, P. Le Roux, M. Peryt, C. Roderick, Z. Zaharieva, CERN, Geneva, Switzerland

Abstract

Since 2003, a coherent data management approach was envisaged for the needs of installing, commissioning, operating and maintaining the LHC. Data repositories in the distinct domains of physical equipment, installed components, controls configuration and operational data have been established to cater for these different aspects. The interdependencies between the domains have been implemented as a distributed database. This approach, based on a very wide data foundation, has been used for the LHC and is being extended to the CERN accelerator complex.

INTRODUCTION

The use of commercial relational database management systems (RDBMS) at CERN started in the early 1980's. The complex technical aspects that were addressed concerned the construction of the Large Electron-Positron Collider (LEP) at that time, such as project planning, cabling, documents, magnet data, etc. [1]. In the following two decades, data management solutions were implemented for many specific needs in accelerator technology. These efforts were mostly conducted by savvy physicists and engineers, interested in innovative software techniques. Successful database driven systems such as the Proton Synchrotron (PS) Controls and the LEP Alarms were developed by niche teams for specific accelerator domains. This situation remained until the organizational reunification of all accelerator and particle beams activities at the end of 2002.

THE VISION

Although it was not initially foreseen by the management, the idea of a single team with a mandate to cover all accelerator data management was eventually accepted and endorsed. For the first time, a dedicated team of eight to ten database engineers started to work collaboratively with a much wider scope for their design and development work.

In order to tackle the ambitious technical and human objectives and with the LHC in sight, a clear vision had to be set out initially, communicated widely and followed by all. With the accumulated experience of the team members and considering the vastness of the information domain, the eagle's eye view was represented graphically as a puzzle.

The Puzzle

Braking up the complex accelerator domain in a few coherent areas as illustrated in Fig. 1, makes sense from a logical and an organizational point of view. Consequently, the work in the individual sub-domains can be handled by one or two people.

The price to pay is that the interrelationships between the sub-domains are not considered initially, at least not in sufficient depth. Knowing that the development efforts would take several years in any case, this consensus was accepted. However, good communication and common understanding was primed to enable the postponed federation work.

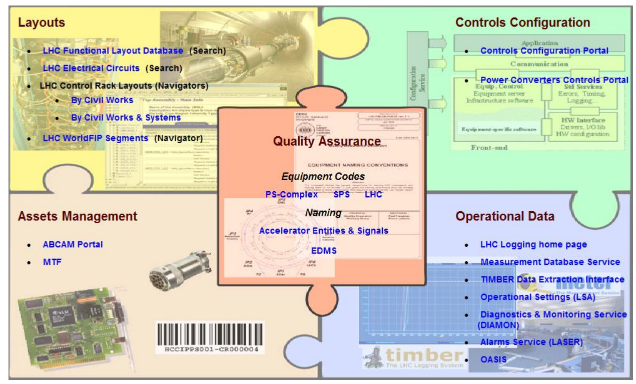

Figure 1: Accelerator data management domains.

The paper will explain first the individual domains, and then demonstrate how the distributed database gets established.

THE STRATEGY

Development Work

Major efforts on requirements gathering, analysis and development work were launched in each of the individual pieces of the puzzle. The respective responsibilities were originally in the hands of experienced senior engineers. Progressively, these responsibilities were taken over by younger software engineers, acquiring the necessary accelerator domain knowledge in the process.

The analysis showed that the legacy databases suffered from a lack of intrinsic quality assurance (QA). Enforcing QA in the database (integrity constraints) and in the data (naming conventions) would also address the anticipated federation between the databases.

Aiming for gradual functionality increase and scope expansion, the schema evolved with carefully planned "big-bang" revolutions and adiabatic modifications.

Technology

The use of Oracle® for the RDBMS is strongly established at CERN, resulting in a huge expertise and knowledge base. The outlined strategy in the ever evolving technology is to stick to the stack of Oracle products for best integration.

Following the trend of control room application software, being developed in Java since the year 2000, the Java Platform (Enterprise Edition) standard was also adopted, namely for database connectivity. Within this context, application's deployments business logic was deployed on the middle-tier Oracle Application Server in the three-tier architecture.

The most important characteristic of the database infrastructure is its *on-line* usage, even in the critical domain of accelerator controls and operations. Stable, high-available database services and robust applications are prerequisites for success.

Responsibilities

A competence shift has to be recognized and accepted, with clear limits of responsibilities, by all parties involved. Where previously the equipment specialists or application software experts "managed" the data, now database engineers need to understand, model and structure this information. The responsibilities between the stakeholders are distributed as follows:

- *Data Management (DM) development team*: Gather and analyse user requirements, provide database objects, data access APIs, data browsing and data maintenance interfaces, initial data capture.
- *Application development team*: Provide applications using DM APIs for data access.
- *Data owners*: Equipment owners are responsible for their data; maintain it with DM provided interfaces.
- *DB infrastructure service*: Database and application server hosting and administration: installation, monitoring, patch and upgrade maintenance, backup.

SOME TACTICS

No strategy is successful without targeted tactical moves. Those that made a major impact were:

- Involve end-users right from the start, throughout the design and development process
- Communicate constantly on scheduled interventions and their anticipated impact
- Iterate rapidly based on end-user feed-back
- Provide adequate environments for development, unit testing, system testing and production
- Push data ownership to the experts, assist and guide the usage of the data maintenance interfaces.

LAYOUT DATA

Traditionally, layout data concerns the *functional positions* of accelerator components, i.e. machine layout, mainly serving machine physicists in the establishment of a machine design. Due to the complexity of the LHC in terms of installation, controls infrastructure and electrical powering, the scope of the layout database was widely opened up [2].

Machine Layout

A machine layout repository aims initially to simulate particle beam behaviour by simulating the magnetic model and validating the beam optics. At a later stage, the mechanical aspects are used for integration and installation work, taking into account the underground tunnel dimensions. 3-D visualizations were generated from this data for the transport and fitting of the large components in the tunnel.

Adding the information of measured positioning discrepancies and other mechanical non-conformities, allows the establishment of the *as-built model* of the LHC machine.

Controls System Layout

More than 9,000 racks of electronics for the LHC controls that had to be installed in underground areas and surface buildings. The layout description relating racks, crates, modules, fieldbus connections was primordial for the preparation, execution and verification of the controls electronics. This information is also the starting point for the automatic configuration of the front-end computers.

Electrical Circuits Layout

The layout database holds the description of the electrical objects that are part of the powering circuits [3]. The relationship between power converters, superconducting current leads, bus bars, warm cables and magnets can be quite complicated. With this data in the layout database, the circuit description can be generated in XML format, allowing the verification of the circuits, of which a major part is located in the cryostats.

CONTROLS CONFIGURATION DATA

This domain concerns the control system topology from front-end computers (FECs) to control room consoles with all its configuration details. A 20-year old legacy system for the PS-complex is at the basis of the current system [4].

Database

The configuration data for the ~65,000 controls devices for all accelerators at CERN is a major part of the Controls Configuration Database (CCDB). There are several device-property models currently supported in the database (GM, FESA, Hardware, Virtual, SL) due to historic changes in the Controls systems and the diversity of the controlled devices. The database was fully redesigned and data integrity constraints were implemented inside the database thus catering for old, new and evolving controls architectures such as the Role-Based Access to the Controls devices or the Configuration of the Timing System. Today, the Controls Configuration Database (CCDB) covers a wide scope, as depicted in Fig. 2.

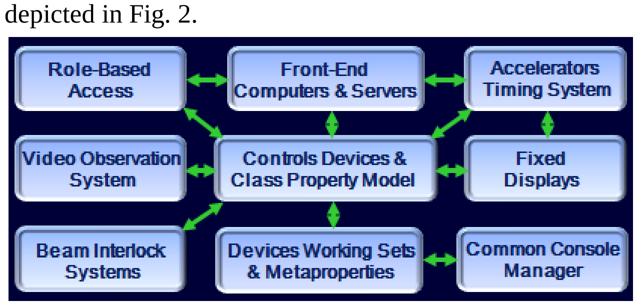

Figure 2: Controls configuration database.

The Power Converters Controls Configuration is another important area, which has been developed into a separate schema. Nevertheless, it is well integrated with the rest of the Controls Configuration data [5].

Interactive User Interfaces and APIs

More than 200 legacy forms-type interfaces were fully redeveloped, using the Java technology, urging equipment owners to take up data responsibility. The new data editing applications greatly expanded on the previously existing functionality and enforced stringent access rights for data modification.

A set of CCDB reporting tools, comprised of some 160 different reports, were redesigned and vastly improved during the last 3 years [6].

A number of Java-based and Pro*C-based APIs provide access to the configuration data to the various areas of the Controls Systems. The Pro*C APIs, primarily the ones for the drivers generation for the FECs, are being renovated, using modern Oracle-based technology to produce standardized XML file output.

OPERATIONAL DATA

The remote manipulation of particle beams in the accelerator chain needs a heuristic approach to ensure performance and reproducibility of live operations.

Settings

The accelerator settings parameter space has grown to a complex but complete database model [7] as basis for the LHC-era Software Architecture (LSA). For each accelerator zone, beam processes are to take place, corresponding to a sequence of (magnetic) cycles. The evolution of parameter settings at high-level physics and hardware level is kept to eventually drive the low-level equipment devices. The theoretical optics model and the measured model of the individual magnets are taken into account. Specific (critical) settings are dealt with explicitly such as Beam Loss Monitor (BLM) thresholds, collimator positions, RF-cavity calibrations, power converter limits.

Measurements and Logging

Mainly for accelerator performance monitoring, time-series data on beam and equipment measurements is captured and stored. Already for LHC hardware commissioning, the ability to retrieve, visualize and analyze this information proved to be extremely valuable. A lot of this data will be kept available on-line for the lifetime of the LHC. Currently estimated at 15TB/year, the required infrastructure for this service is getting quite demanding and huge [8].

Alarms

The Alarms System (LASER – LHC Alarms Service) [9] is a centralized service ensuring the capturing, storing and notification of anomalies for the whole accelerator chain as well as for the technical infrastructure at CERN. The underlying database holds the pre-defined alarm definitions and the time-stamped run-time alarm events. The alarm archives are kept since 2005.

ASSETS DATA

For the physical components, a CERN-centralized commercial asset management system is in place, completed with home-made web-deployed user-friendly interfaces [10]. Each asset is uniquely identified by a CERN-wide serial number called equipment *part identifier.* The complete lifecycle of the asset can be recorded, from its construction, testing, installation, movements, repairs to its final destruction. For historical reasons, this system is called MTF (Manufacturing and Test Folder). Despite the fact that this *electronic traveller* also covers the subsequent phases of the equipment lifecycle, the acronym MTF continues to be used.

QUALITY ASSURANCE

The LHC started with a detailed plan on Quality Assurance (QA) covering all activities from design, development, production, installation, servicing and documentation. The supporting QA tools for document and asset management have reached a mature state and are being used for other CERN accelerators and projects as well. Identification and naming is always a starting point for QA on the concerned equipment. Equipment code catalogues, the official references, are maintained and publicly available as well as a ‘dictionary’ list of the signals and parameters used in the various systems of the accelerators. This functionality is implemented in the so called Accelerators Entities and Signals Naming database.

The imposed naming conventions, clear responsibilities and procedures to be followed by all actors is the basis of the data federation, making intercommunication possible between the different data domains of the puzzle.

DATA FEDERATION

The ideal situation for a perfect integration is to have a unique identifier (primary key) for each and every object throughout all data domains. This is not the approach that has been taken, so how do the domains communicate, share and propagate data? In fact, several solutions have been put in place that fit each individual data sharing requirement. This is best illustrated by the following concrete examples.

Assets Installed in Layout Slots

Functional positions (also known as *slots*) in the layout database can be imported in MTF if required. The slots that are eligible for import are flagged with a dedicated attribute in the table that holds the slots. A materialized view (MV) processes and combines all the necessary data for MTF for the flagged slots. The MTF administrators import the new slots information at their earliest convenience, by reading out the MV and propagating the data in their local database objects.

Via an MTF web interface, equipment owners can assign an asset – by means of its part identifier – to the slot in which the physical equipment is installed.

The link between layout and assets is bidirectional. MTF provides an MV representing the slot-to-asset mapping. The data propagation towards the layout database is fully automated by a database process scheduled to run regularly in the layout database. Every three hours, the MTF views are read out in order to synchronise the mapping in the layout database.

The Beam Loss Monitoring System

The electronics chain of the LHC Beam Loss Monitoring system (BLM) is an example of a very large and distributed structure. The individual physical parts have been identified as assets in MTF. The installation of these components, their testing and commissioning, have produced and collected a lot of information, uploaded in MTF via MS-Excel files. This information is made available in the Layout database via the MV.

This raw information is now available via the MV, but needs to be transformed by PL/SQL code. Running automatically every night or on demand, the data is parsed, checked and translated into different structures. The resulting format allows modelling and visualization of the complete electronics connectivity chain.

With this structured and consistent BLM data in the layout tables, an MV is populated, accessible from the Settings database. The MV combines information related to the physical assets (serial numbers, firmware versions), required for on-line operation of the LHC.

A dedicated PL/SQL package contains procedures which make use of this view, to allow BLM experts to synchronize, on demand, data changes in the MTF and Layout databases with the Settings database. Moreover, the data is used to cross check the database image with respect to the actual configuration. In case of discrepancy, LHC beam injection will be delayed.

The BLM system experts can complete the on-line configuration of the BLM system with the beam interlock thresholds. Due to their energy and integration time dependency, these result in more than 1.5 million values.

Commissioning Sets of Circuits

The electrical circuits and power converter design information are stored in the Layout database (including operational design limits, and the relationship between circuits, converters, and magnets). This data is imported into the Settings database on demand via a dedicated PL/SQL package. Once in the settings database, the data are used for operation (knowing which converters power which magnets, and what are the operational limits), and for hardware commissioning (HWC). For HWC, testing procedures are defined for Sets of Circuits (SOC), and the tests are conducted by interacting with the related converters.

When a test is started on a SOC via a control room application, a database trigger fires, which calls a PL/SQL procedure. This procedure generates dynamically the metadata in the Measurement database describing a group of variables to be logged (identified by SOC name), for use by the Measurement client. The Measurement variables correspond to pre-defined measurable quantities of the converters (current, voltage) which are related to each of the circuits in the SOC.

Configuring Front-End Computers

In line with our strategy, front-end computers (FEC) are now fully described in Layout prior to their configuration in the CCDB. The breakdown structure is modelled from the hosting rack to the crate with its electronic modules precisely identified, classified and positioned. For the WorldFIP fieldbus gateway, the connection chain is described from the manager board, via optical repeaters, tap box, to the agents connected to sensors and actuators.

The CCDB has a completely different set of capabilities for the ~3,000 FECs and servers used by the controls systems. This data allows for automatic generation of start-up files, hardware configuration including physical addressing and interrupts of the modules in the FECs, as well as software configuration of the start-up sequences and remote reboot of the FECs.

The hardware and/or software configuration of the FECs, is based on the type and position of the crates and modules that built up those FECs. This basic data is made available from the Layout database by means of MVs accessible by the CCDB via a database link. A stored PL/SQL procedure takes care of importing the data.

Some of this data is further used for the configuration of the Diagnostics and Monitoring (DIAMON) tools [11].

Driving Settings of Software Devices

The controls device-property models defined in the CCDB must be represented in the Settings database for on-line operation of the CERN accelerators. The CCDB covers several different implementations of the device-property model, which at the level of on-line Settings management are abstracted into a single device-parameter model. To reflect this, and help facilitate the integration of this data into the Settings database from the CCDB, the descriptions of the various device-property models are unified in database views, which are accessible from the Settings database. In the Settings database, a dedicated PL/SQL package contains a procedure which uses these views, to allow operators to import, on demand, device-property data into the Settings database. Once imported, device-parameter model data can be generated, together with setting values, for on-line operation.

Generating Alarm Definitions

From the information on controls devices in the CCDB, alarm definitions can be configured by the equipment experts (currently some 41,000). These alarm definitions are formatted for the Alarms Service through a chain of 20 views, residing on the CCDB and the alarms database. The PL/SQL import procedure is executed automatically at regular intervals.

Maintaining Data Consistency

The most important goal of data management achieved by the data federation is to have a single source of data! The examples above have shown that in many cases, data propagation is neither immediate nor automatic. Correct execution and verification of the data synchronization procedures is mandatory.

COVERING ALL ACCELERATORS

The overall coherent data management architecture, design and implementation were set out with the LHC in mind, having the highest priority as CERN's flagship. An additional advantage was the ability to start from a clean sheet. Nevertheless, a retrofit to existing accelerators is in progress. Existing legacy data is integrated in the existing model. Inevitably, the model is gradually extended to cater for the specifics of the other accelerators. Possibly the most complex area concerns the high level controls and settings management of the PS-complex. A convergence towards the LHC approach and consequent renovation has crystallized in an important project [12] in which the data management is again essential.

EXPORTABILITY

The ideas and developments that have been presented in this paper were materialized at CERN, and continue to be carried out. It does raise the question if this can be used at other laboratories? The answer is "Yes, it can". In the frame of collaboration agreements with GSI, the LSA Settings management and FESA Controls Configuration have been exported in view of integration into the existing software controls of GSI/FAIR [13].

Clearly, the availability of an Oracle RDBMS is a prerequisite to be able to export the deployment, due to our strategic choice to implement data-centric code in the proprietary procedural language Oracle PL/SQL. Without the Oracle database, only the concepts can be exported.

Table 1: Content and Complexity of Databases

	tables	constraints	code	volume
Layout	134	495	55,602	5.3 GB
Configuration	514	1,524	30,326	9.7 GB
Settings	281	1,392	9,026	14.6 GB
Logging	55	103	14,431	+17 TB
Alarms	207	191	24,915	62.8 GB

HUMAN RESOURCES

One should not overlook the most important asset in the process of analysis, design, development and maintenance, namely the people. In addition to their specific technical competence, the core team of database engineers have to:

- Follow and use the technology effectively
- Acquire specific domain knowledge
- Show flexibility in changing user requirements
- Impose data access methods to application developers

Table 1 gives some figures indicating the complexity and size of some of the databases. This also demonstrates the resulting liability of the responsible team members.

REFERENCES

[1] J.P. Koutchouk and J. Poole, "Data Management for LEP Control", PAC 1987, Washington D.C., USA*.

[2] P. Le Roux, R. Billen and J. Mariethoz, "The LHC Functional Layout Database as Foundation of the Controls System", ICALEPCS'07, Knoxville, USA, Oct-2007, RPPA03, p. 526 (2007)*.

[3] M. Zerlauth *et al.*, "The Electrical Circuit Description for the LHC", EPAC 2002, Paris, France, Jun-2002, MOPDO014, p. 2058 (2002)*.

[4] J. Cuperus, R. Billen and M. Lelaizant, "The Configuration Database for the CERN Accelerator Control System", ICALEPCS'03, Gyeongju, Korea, Oct-2003, WE114, p. 309 (2004)*.

[5] Q. King, S.T. Page and Z. Zaharieva, "Automatic Inventory and Configuration Management Tools for the LHC Power Converter Controls", ICALEPCS'09, Kobe, Japan, Oct-2009, TUA004.

[6] Z. Zaharieva, R. Billen, "Rapid Development of Database Interfaces with Oracle APEX, used for the Controls Systems at CERN", ICALEPCS'09, Kobe, Japan, Oct-2009, THP108.

[7] C. Roderick and R. Billen, "The LSA Database to Drive the Accelerator Settings", ICALEPCS'09, Kobe, Japan, Oct-2009, WEP006.

[8] C. Roderick *et al.*, "The LHC Logging Service: Handling Terabytes of On-line Data", ICALEPCS'09, Kobe, Japan, Oct-2009, WEP005.

[9] P. Sollander *et al.*, "Alarms Configuration Management", ICALEPCS'07, Knoxville, USA, Oct-2007, RPPB03, p. 606 (2007)*.

[10] S. Mallon *et al.*, "CERN's Global Equipment Data Repository", ICALEPCS'09, Kobe, Japan, Oct-2009, TUB004.

[11] M. Buttner *et al.*, "Diagnostic and Monitoring CERN Accelerator Controls Infrastructure", ICALEPCS'09, Kobe, Japan, Oct-2009, THA004.

[12] S. Deghaye *et al.*, "CERN Proton Synchrotron Complex High-Level Controls Renovation", ICALEPCS'09, Kobe, Japan, Oct-2009, TUP019.

[13] R. Mueller, J. Fitzek and D. Ondreka, "Evaluation of the LHC Software Architecture for Data Supply and Setting Management within the FAIR Control System", ICALEPCS'09, Kobe, Japan, Oct-2009, THP012.

* Published at the Joint Accelerator Conferences Website (JACoW) http://www.JACoW.org