

# Chapter 18

## Controls technologies

*J. Serrano*<sup>1\*</sup>

<sup>1</sup>CERN, Accelerator & Technology Sector, Switzerland

\*Corresponding author

### 18 Controls technologies

#### 18.1 Overview

By the time of the commissioning and subsequent operation of the HL-LHC, many of the physical elements of the control system will have been upgraded due to obsolescence. This applies particularly to the front and back end CPUs and storage. It is not, however, foreseen that the overall control system strategy and architecture will change in its conceptual structure during this period, and many parts of the current controls infrastructure will still be sufficient for the HL-LHC needs. Nevertheless, three areas have been identified as having to be addressed so that the control system can respond to the new challenges presented by the HL-LHC upgrade project.

During operation of the HL-LHC there will be an increase of radiation in some areas, which will trigger re-designs and relocation of electronics currently installed close to beam line elements (see Chapter 10 and 19). The HL-LHC project will also require the installation of new, more powerful Nb<sub>3</sub>Sn magnets that will raise the need for more diagnostics data (i.e. higher data rates) in subsystems such as the quench detection system and the cold powering system. Higher data rates will also be needed during the commissioning of the HL-LHC, as equipment groups will need to fine-tune their systems and will therefore require access to the full diagnostics capabilities of their systems. In order to assure correct functionality up to the end of the HL-LHC operational period with ultimate performance, it is important to be conservative regarding the design choices and to share proven solutions as much as possible. This approach assures that proven solutions persist and that all design efforts can be concentrated on making a few designs very robust instead of spreading efforts into a large number of sub-optimal designs.

The increase in data bandwidth needs triggered by the HL-LHC is an overarching theme in this work package. It has an effect on the electronics interfacing to the accelerator components, on the communication technologies used to get the data out of those crates towards higher layers of the control system, and finally on the solutions used to store the data in the logging system and, later, to extract it and analyse it in an efficient way.

#### 18.2 Control technologies

##### 18.2.1 Data logging

Development of the next generation accelerator logging systems (NXCALS) has started, and the aim is to have a new Logging system with equivalent functionality as today but with significantly easier scalability, faster data extraction and analysis performance.

Following past experience, we are expecting a massive increase in the volume of data to be logged as a direct result of installing, commissioning and then operating the HL-LHC as well as of numerous consolidation actions across the various equipment systems presently installed in the LHC. Figure 18-1 shows the growth in

data logged post LS1 as a result of updates to various LHC systems, most notably the LHC quench protection system (QPS), and the need to increase data rates to better understand the operational behaviour of the machine.

Based on user input and given the extensive hardware and operational changes foreseen for the HL-LHC, we must be prepared for similar or even greater increases in the amount of data collected in the logging system in the HL-LHC era. We foresee the need to install additional, properly dimensioned, hardware in 2025, to be ready for use in 2026 to support the HL-LHC hardware commissioning followed by beam commissioning.

In addition to the need for additional storage, it is critical to develop and deploy a new software infrastructure to properly support the users of the foreseen data sets. In order to properly validate the behaviour and performance requirements with realistic data and use cases *prior* to the HL-LHC commissioning period, the deployment of this software should take place before the start of Run 3.

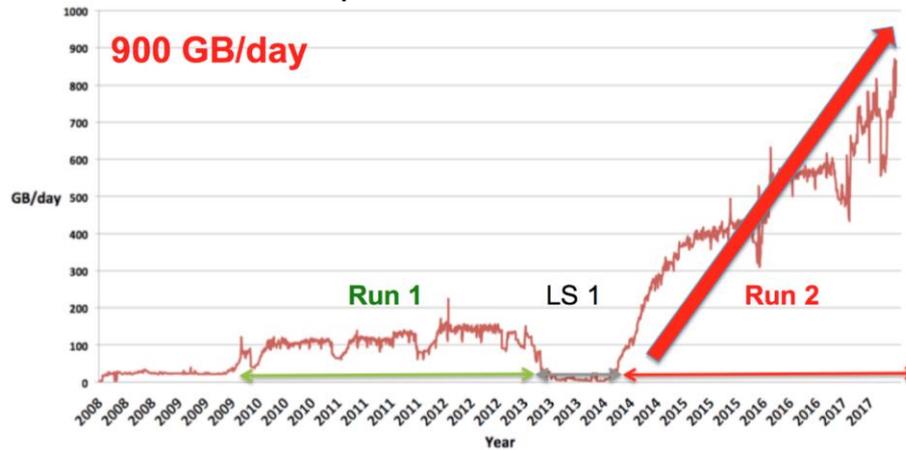


Figure 18-1: Storage evolution.

In recent years, the so-called “Big Data” technology landscape has evolved significantly to support large-scale data logging and analysis, opening up new possibilities to perform efficient analysis of large data sets. To gain experience with these technologies and help choose a direction for NXCALS, a Proof of Concept (PoC) Logging System was developed in collaboration with IT-DB in early 2016. This PoC was based on the open-source Apache Hadoop technology - as a replacement for the current Oracle-based CALS [1] system. The PoC work clearly demonstrated the potential to successfully replace the current system and improve performance and scalability for an overall lower hardware cost than the current system (not considering the Oracle Licensing costs). Subsequently, the approval by CERN management led to the full-scale development of NXCALS.

The NXCALS system is based on a microservices architecture. The aim of this is to be able to more easily upgrade or replace different aspects of the system in the future as necessary, without being forced to put in place a completely new system. From a technology perspective, NXCALS is based on in-house developments combined with open-source software such as Hadoop (HDFS and HBase), Kafka, Spark, and Jupyter notebooks.

The core technologies used in NXCALS are based on the concept of “horizontal scalability”, which essentially means the ability to increase performance by adding more resources to the underlying infrastructure. From this perspective, the NXCALS system has the potential to adapt to the required performance needs of the future, provided sufficient resources can be financed and that sufficient physical hosting capacity is available.

In terms of potential data analysis performance, a key difference in NXCALS with respect to the CALS system is a change in paradigm. With the CALS system, users first extracted the data and then performed the analysis on their local machines. With the NXCALS system, users seeking high levels of data analysis performance need to submit their analysis algorithms to be executed directly on the NXCALS cluster, using

Spark, and then retrieve only the results. This change of paradigm has already revealed cases where analysis times can be reduced from several days to less than an hour.

### 18.2.2 New distributed I/O tier

The HL-LHC will place challenging demands on data acquisition to/from the accelerator components which need to be controlled and diagnosed, such as the new Nb<sub>3</sub>Sn magnets. The need for larger amounts of diagnostics information will result in a requirement for more throughput in the lower layers of the control system and will therefore affect the electronics in this tier and the communication links used to send the information up the controls stack. The current controls architecture has front-end computer systems (VME or PICMG 1.3) with a large variety of reusable electronic cards to control accelerator components by sending and receiving data and carrying out calculations in real-time. In the LHC, these front-end computers typically drive some kind of fieldbus, which connects to Input/Output (I/O) modules sitting close to the accelerator, as shown in Figure 18-2. Historically, there has been much less sharing and reuse of design effort in this lower Distributed I/O Tier (DI/OT) than in the front-end tier.

For the HL-LHC, the proposal is to extend the sharing model of the front-ends to the DI/OT layer. The electronics in this layer is designed to transmit data as fast as possible to/from actuators and sensors attached to accelerator components. These I/O modules are connected to a smaller number of high-performance front-end computers which process the data and perform the necessary calculations. By collaborating with equipment groups and providing a service in this I/O layer analogous to that of the front-end tier, we will ensure a uniform level of quality and increase overall availability of electronics deployed in this tier, including those subject to radiation [2].

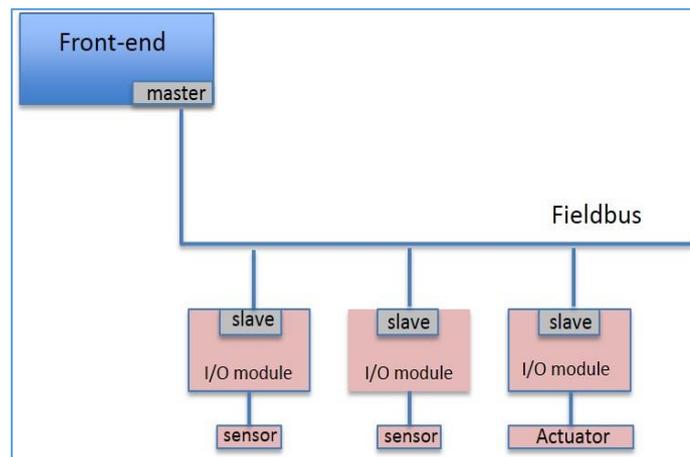


Figure 18-2: Proposed controls architecture.

As shown in Figure 18-3, the DI/OT kit will be modular, allowing different applications to benefit from common infrastructure at different levels. It will consist of a 3U Europe crate and a passive backplane conforming to the CompactPCI Serial standard. The standard specifies PCIe as the main protocol to be used in communication through the backplane. PCIe is, however, unnecessarily complex for our needs and an attempt to implement it for the DI/OT system could compromise radiation tolerance. We therefore decided to use the basic physical infrastructure of CompactPCI Serial without following further prescriptions on protocols. A simple serial protocol (such as high-speed SPI) will be used instead, with support for automatic discovery of hardware modules. The controller slot in the crate will host the so-called system board, which communicates with other boards through the backplane and with the upper layers of the control system through a fieldbus interface. In order to support different fieldbus technologies, the system board features an FMC (VITA 57) slot, and different communication mezzanines can be plugged in that slot. There will be different boards for radiation and non-radiation areas. Those meant to operate in radiation environments will be optimised for radiation tolerance, so simplicity will be a major design goal, at the expense of performance. The system board variant meant to operate in non-radiation environments will be more complex and capable. The project includes

the development of a radiation-tolerant (rad-tol) switching AC/DC power converter, whose design will be made generic enough so that parts of it can be reused in other projects (e.g. FGCs).

The modularity of this kit caters for different needs in equipment groups. Survey, for example, will use the full kit, including the crate, rad-tol system board and WorldFIP communication mezzanine. They will design their own add-in boards in 3U Europe format to interface with their sensors and actuators. The BLM, BPM and other systems have a need for their own dedicated crate and system board, and they will insert one of the communication mezzanines on it for basic remote diagnostics and slow control. The non-radiation-tolerant variant will be used in the Full Remote Alignment System. A short summary of foreseen uses can be seen in Table 18-1.

Table 18-1: Foreseen uses of DI/OT electronics.

System	Components	Locations
Survey	50 full DI/OT crates and 30 crates with PSU only	UA galleries and RRs
FRAS	18 racks populated with electronics	UR15, US15, UR57, UL557
WIC	50 DI/OT crates	TI2, TI8, TT40, TT41
PIC	36 DI/OT crates	RR13/17, RR53/57, RR73/77, UA23/27, UA43/47, UA63/67, UA83/87
BLM	60 radiation-tolerant WorldFIP mezzanines	SPS
BPM	Potentially 500 radiation-tolerant WorldFIP mezzanines	LHC arcs and Dispersion Suppressors

The DI/OT kit will offer general services like remote monitoring of the platform (temperatures, fan speeds, voltages, and currents...) and remote re-programming of the Field Programmable Gate Arrays (FPGA). It will also benefit from a specific effort on increasing reliability and availability. Another important part of the monitoring infrastructure will be the measurement and reporting of radiation in the location of each crate, through the inclusion in each chassis of a generic radiation monitoring module supported by the R2E working group.

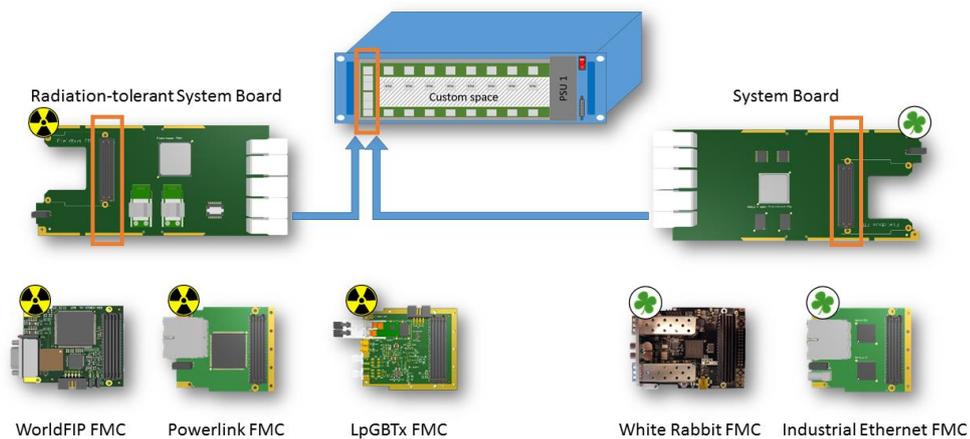


Figure 18-3: Modular DI/OT kit, including radiation and non-radiation-tolerant variants.

### 18.2.3 A new high-speed radiation-tolerant fieldbus

For HL-LHC the QPS system needs to accommodate the new Nb<sub>3</sub>Sn magnets in points 1, 5 and 7. The newly developed Universal Quench Detection System has much increased data acquisition capabilities, rendering the current solution based on WorldFIP sub-optimal, due to the very low available bandwidth (2.5 Mbps maximum) for the transmission of the acquired logging and Post mortem data.

The QPS electronics in points 1 and 5 will be installed in radiation-free zones, however the RR alcoves in point 7 may pose some low ( $< 1\text{--}2$  Gy/year) radiation-tolerance constraints, which should be considered when proposing a new fieldbus for the HL-LHC era.

A re-usable standards-based solution will serve also other users and will help increase overall quality and therefore availability. An industrial, Ethernet-based solution with 100 Mbps bandwidth,  $\mu\text{s}$  synchronization and supporting 50 slaves/segment is proposed, so it can be considered as a candidate for the QPS electronics and any other subsystem needing faster data transfer rates than what WorldFIP can provide. After a market review including leading Industrial Ethernet technologies such as Profinet and EtherCAT, we decided to design a radiation-tolerant implementation of Ethernet Powerlink. This is the only popular fieldbus technology featuring an open-source implementation of its stack unencumbered by patents and other obstacles. It is important to have full control of these sources because radiation-tolerance is achieved through logic triplication and voting in a flash-based FPGA. This requires a certain degree of introspection in the design so as to be able to test different strategies for triplication and evaluating which ones work best.

The basic technologies used for rad-tol digital design in this work package are mature and well tested for doses of a few hundred Gy. For pure logic, triplication inside a flash-based FPGA, followed by voting, can effectively mitigate the effects of Single Event Upsets (SEUs). Systems involving soft-cores running software are a bit more involved because of the various places at which redundancy can be inserted. Figure 18-4 shows the most likely scenario for the implementation of a rad-tol Powerlink stack inside an FPGA. A RISC-V [3] core is triplicated and runs software stored in Error-Correcting Code (ECC) memory. The data for the program resides in a separate ECC RAM block. This basic building block can then be re-used in other projects needing a small microcontroller running software of moderate complexity in a radiation environment. The main challenge in our context is to implement the open-source Powerlink stack, originally developed to run in desktop systems with no memory limitations, in the amount of memory available in typical flash-based FPGAs.

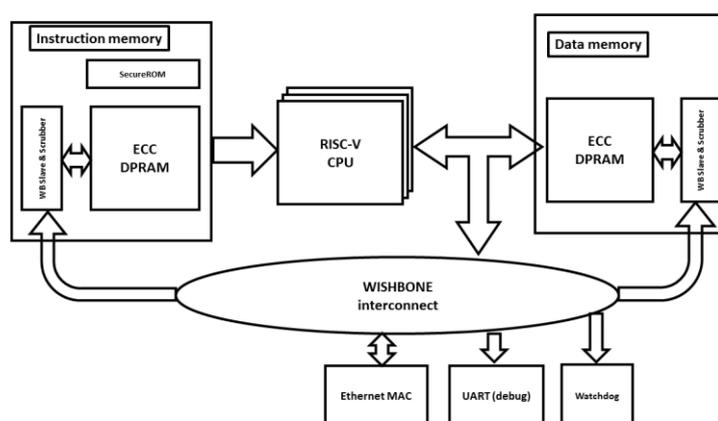


Figure 18-4: Simplified block diagram of the rad-tol Ethernet Powerlink implementation.

The Powerlink mezzanine (Figure 18-5) will bridge the gap between users who are satisfied with the limited bandwidth of WorldFIP and streaming-like multi-Gb/s applications, which will use the LpGBTx chip designed in the EP Department at CERN. Powerlink being a standard, it is relatively easy to find commercial off-the-shelf solutions for the master side of the fieldbus, including PCIe add-in boards hosted in Linux PCs and also bus masters in Programmable Logic Controllers (PLCs). This illustrates a common theme in this work package: using industry standards as far as possible to benefit from a set of verified solutions and customising them only as needed.

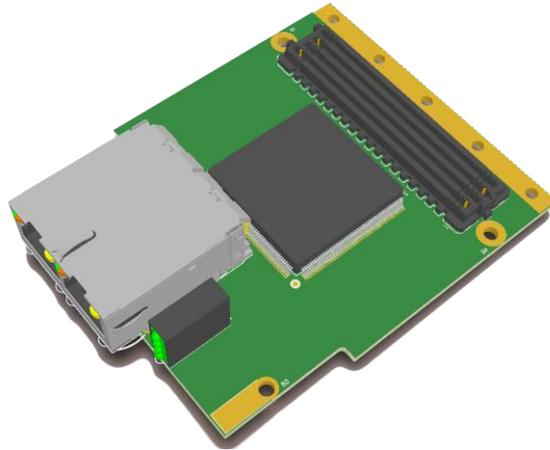


Figure 18-5: 3D model of the Powerlink mezzanine.

### 18.3 References

- [1] The CERN Accelerator Logging Service - 10 Years in Operation: A Look at the Past, Present and Future C. Roderick, *et al.* 14th International Conference on Accelerator & Large Experimental Physics Control Systems, San Fran, USA, 2013, CDS: [tuppc028](#).
- [2] Plans at CERN for electronics and communication in the Distributed I/O Tier\_G. Daniluk *et al.*, ICALEPCS 2017, Barcelona, Spain, 2017, DOI: [10.18429/JACoW-ICALEPCS2017-THPHA071](#).
- [3] Open hardware repository [web page](#).