AUTOMATIC INVENTORY AND CONFIGURATION MANAGEMENT TOOLS FOR THE LHC POWER CONVERTER CONTROLS

S. Page, Q. King, Z. Zaharieva, CERN, Geneva, Switzerland.

Abstract

The LHC has more than 1700 power converters of a variety of different designs, each of which is itself composed of many hardware components. These components must be individually tracked throughout their lifetimes for inventory purposes. Additionally, each component may have associated configuration and calibration values that must follow the component wherever it is installed within the accelerator. The complete history of the inventory and configurations must be retained as well as the calibration values in order to track their evolution. With such a large machine the probability of human error is high, therefore an automated solution is desirable. A system was put in place for the LHC power converter controls that automatically detects the connection of a new device, updates its location and loads its configuration and calibration values from a database within a few seconds. The identification of all key components is remotely readable via the control system. This paper describes the system, detailing its architecture, choices made and results achieved as well as challenges overcome during its implementation.

INVENTORY

Each piece of equipment installed in the LHC – whether a complete power converter or a sub-component – is tagged with a standard barcode containing its LHC equipment code (defining the equipment type and manufacturer) and serial number (defining the individual instance of the type). These give the equipment a unique identifier (LHC ID) that can be used to track its manufacture, installation, repair and modification.

Several challenges remain however. The number of components in the LHC combined with the diversity of their types and sizes result in a system that would be difficult to accurately maintain manually. Scanning barcodes by hand on each intervention would risk to introduce inaccuracies into the system due to human error.

For the LHC power converters, the basic inventory requirement is to link all pieces of equipment that make up a power converter to a logical slot for a converter within the LHC which has a corresponding name. These logical slots are known as 'systems' and they also have a type defining what should be installed within them.

CONFIGURATION

Each LHC power converter has a unique configuration composed of property values and derived from the inventory as illustrated in Fig. 1. These values may be determined by the specific equipment installed in the system, the type of that equipment, the system itself (regardless of installed equipment) or the type of that system. It is critical that the configuration of all of the power converters is correct so that the required level of machine performance can be achieved.

Equipment Configuration

The equipment configuration consists of values that are specific to an individual piece of equipment. These values will typically relate to calibrations that have been performed upon the equipment. It is particularly important that the history of calibration values is retained over time so that trends can be observed and calibration campaigns may be planned to maintain machine performance. The equipment configuration must follow the equipment wherever it is installed within the machine.

System Configuration

The system configuration consists of values that are specific to a slot within the machine. In the case of a power converter, these values are typically those which

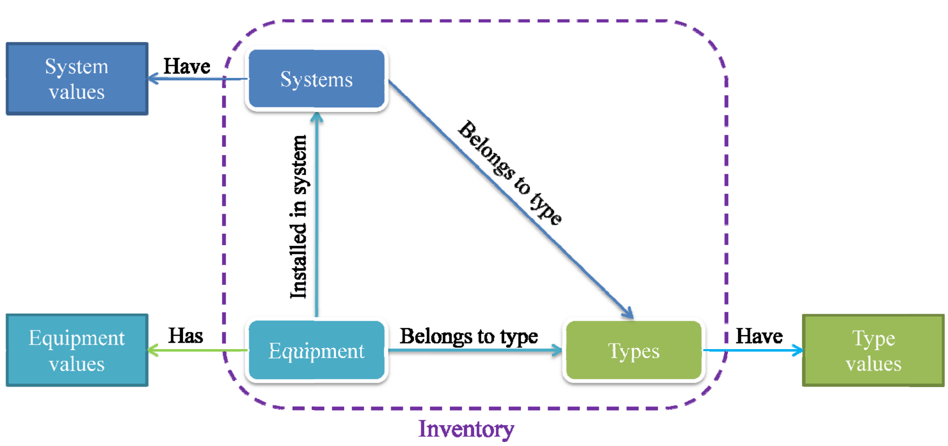

Figure 1: Inventory and configuration relationships.

relate to the circuit and load to which it is connected. The values of these properties do not change when the power converter or any of the equipment of which it is composed is replaced.

Type Configuration

The type configuration consists of values that are common to either all pieces of equipment or all systems of a particular type (as defined by its LHC equipment code). These property values can be updated centrally by the people responsible for the equipment. They can then be applied to every instance of the type. Typical examples of properties within the type configuration include gains common to all DCCT heads of a certain type.

IMPLEMENTATION

Identification of Equipment

In order to automatically update the inventory and configuration of each power converter a means of programmatically identifying the installed equipment was required. Dallas Semiconductor Silicon Serial Number chips were integrated into almost all equipment that needed to be identified. DS18B20 thermometer chips were used for analogue measurement equipment and DS2401 ID chips were used for all other equipment. Both types of chip provide a unique 64 bit number that can be used to individually identify equipment. These chips can be read remotely via the Function Generator/Controller (FGC) installed in each power converter. Each FGC has the mappings between all Dallas IDs and the LHC ID of their equipment stored in its non-volatile memory and can therefore return programmatically the barcodes of all equipment installed within the power converter.

Control System Architecture

Each FGC installed in a power converter is connected via a WorldFIP fieldbus to a Linux PC front-end system (with up to 30 FGCs per bus), see Fig. 2. In the LHC, there are 73 front-ends (known as gateways) covering the ~1700 power converters. The front-end systems provide command access to the power converters as well as publishing status data [1]. In order to decouple non-critical background operations such as the inventory and configuration system from the online accelerator controls, a status server is used which receives the status of every FGC at a rate of 0.2Hz and makes it available on request.

Configuration and inventory management is performed by a dedicated process which obtains the state of every FGC from the status server and responds to any that raise a flag to request to be synchronised.

Updating the Inventory

When a piece of equipment is replaced, the technician performing the modification resets the FGC in the power converter as part of the procedure. This causes the FGC to scan the IDs of all installed equipment and then to request to be synchronised with the database. The inventory will be updated and the configuration for the newly installed equipment will be loaded into the FGC within a few seconds.

The procedure is fully automated and avoids the technician needing to have any interaction with the database or to load configuration values by hand.

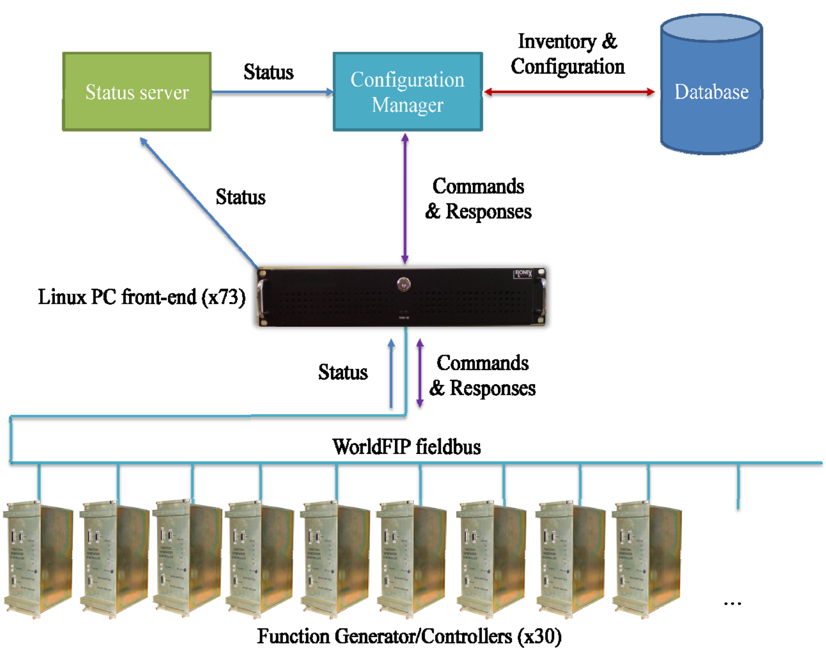

Figure 2: System architecture.

Modification of the Configuration

The property values that make up the configuration of a power converter are directly accessible through the standard control system. The values can be updated by setting them in the same way as would be done for any other property. Initially, the new value will be stored only in the FGC in the power converter and the configuration will be marked as being out of synchronisation. In this intermediary state, the new configuration can be tested before it is made permanent or rolled-back. To permanently store the updated configuration a special property is set on the FGC to trigger a synchronisation of the configuration to the database (known as a 'SYNC_DB'). Alternatively the modifications can be reverted by triggering a synchronisation in the other direction (from the database to the FGC, known as a 'SYNC_FGC').

Synchronisation Process

An FGC which needs to be synchronised sets a flag in the status data that is sent by its gateway to the status server. The FGC Configuration Manager surveys the value of this flag for all FGCs every 5 seconds. When a flag is found to be set, the Configuration Manager sends commands to the FGC to read the type of synchronisation required and follows the appropriate procedure.

Synchronisation from the Database to an FGC – SYNC_FGC

The first step when synchronising from the database to the FGC is to establish the list of equipment installed in the system and to update the inventory in the database. The Configuration Manager reads the list of installed equipment from the FGC. Once this has been done, the configuration relating to the system, the installed equipment and their types are read from the database and the related properties set on the FGC. On completion of the process, a property is set on the FGC to indicate that it is now synchronised with the database.

Synchronisation from an FGC to the Database – SYNC_DB

Synchronising the configuration from the FGC to the database does not update the inventory, since the software in the FGC enforces that this procedure cannot be launched until after a successful SYNC_FGC (which should have taken place at start-up). The Configuration Manager reads the list of properties that make up the configuration from the database then gets the same properties from the FGC. Any modifications are then written back to the database. On completion of the process, a property is set on the FGC to indicate that it is now synchronised with the database.

Configuration and Inventory History

A full history is kept in the database of every configuration property and the systems in which all pieces of equipment have been installed. This history can be browsed through a web interface that also allows the data to be extracted to files for further analysis.

CHALLENGES

Mapping IDs to Barcodes

The Dallas ID chips attached to each piece of equipment provide a unique ID value, however this value itself has no meaning and must be mapped to a standard LHC barcode. For most equipment, this mapping was retrieved automatically as part of the testing process following manufacture [2]. However the diversity of the equipment covered as well as the number of different manufacturers resulted in varying levels of accuracy and coverage for the Dallas ID to barcode mappings. Any ambiguities and inconsistencies in the mappings between Dallas IDs and barcodes were removed from the data set and the devices considered unknown. A campaign was carried out during LHC hardware commissioning to retrieve the unknown IDs.

STATUS

The use of the configuration and inventory management system is well advanced [3]. It has been particularly important to LHC hardware commissioning, during which components are frequently interchanged and configurations are updated and refined. At the time of writing, 1975 power converters are being managed by the system, covering 49251 pieces of equipment of 342 different types. Across that hardware, 99642 property values are being tracked and synchronised.

In addition to the LHC, the system is now also being used for the CERN Proton Synchrotron Main Power Supply following the upgrade of the control electronics to an FGC at the beginning of 2009. The next generation of FGC controls hardware will be used to renovate other areas of the PS Complex and will also be covered by the configuration and inventory management system.

REFERENCES

[1] S. Page, “Integration of the LHC Power Converters within the High-Level LHC Control System”, ICALEPCS, Geneva, October 2005, PO1.043.6.

[2] P. Fraboulet, “Application of Advanced Hardware Testing and Inventory Management Methods for the LHC Power Converter Control System, ICALEPCS, Geneva, October 2005, PO2.038-3.

[3] Q. King, “Status of the LHC Power Converter Controls”, ICALEPCS, Kobe, October 2009, MOB003.