MANAGING INFORMATION FLOW IN ALICE

O. Pinazza, INFN Sezione di Bologna, Bologna, Italy and CERN
A. Augustinus, P. Ch. Chochula, L. S. Jirdén, M. Lechman, P. Rosinský,
CERN, Geneva, Switzerland
G. De Cataldo, INFN Sezione di Bari, Bari, Italy and CERN
A. N. Kurepin, INR RAS – Institute for Nuclear Research of the Russian Academy of Sciences,
Moscow, Russia,
A. Moreno, Universidad Politécnica de Madrid, ETSI Industriales, Madrid, Spain.

Abstract

ALICE is one of the experiments at the Large Hadron Collider (LHC) at CERN in Geneva, Switzerland. The ALICE detector control system is an integrated system collecting 18 different detectors' controls and general services. Is implemented using the commercial SCADA package PVSS. Information of general interest, such as beam and condition data, and data related to shared plants or systems, are made available to all the subsystems via the distribution capabilities of PVSS. Great care has been taken to build a modular and hierarchical system, limiting the interdependencies of the various subsystems. Accessing remote resources in a PVSS distributed environment is very simple and can be initiated unilaterally. In order to improve the reliability of distributed data and to avoid unforeseen and unwished dependencies, the ALICE DCS group has enforced the centralization of global data required by the subsystems. A tool has been developed to monitor the level of interdependency and to understand the optimal layout of the distributed connections, allowing for an interactive visualization of the distribution topology.

INTRODUCTION

ALICE is one of the four general purpose experiments installed around the Large Hadron Collider, at CERN in Geneva. It is composed of 18 detectors, each with its own specific technology choice.

The ALICE collaboration includes more than 1000 physicists and engineers from 105 Institutes in 30 countries. Being developed by several groups in parallel, and being based on different devices and protocols, the controls of the different detectors are heterogeneous too.

The ALICE Detector Control System (DCS), integrates the controls of all detectors in a common environment, exploiting the flexibility of the SCADA PVSS software.

INFORMATION FLOW IN THE ALICE CONTROL SYSTEM

The DCS continuously collects large amounts of data from the various detector devices. Thanks to a strong standardization effort at the beginning of the DCS project it has been possible to limit the variety of devices used and this has been very beneficial to the data collection process. For these devices the hardware diversity is managed through the use of the commercial OPC server protocol.

However, for the very important Front-End Electronics part, which has been custom designed for each detector, there is a large diversity of readout interfaces. Many different field busses and technologies are used such as CANbus, JTAG, Profibus, RS232, Ethernet and custom links. This has called for a more elaborate readout architecture based on the definition of a common high level Front-End Device (FED) which overcomes the hardware diversity.

The use of the TPC/IP based CERN DIM protocol for communication provides furthermore the means for abstracting the data flow from the hardware level. Data is mainly transported using the Ethernet but other protocols are also used.

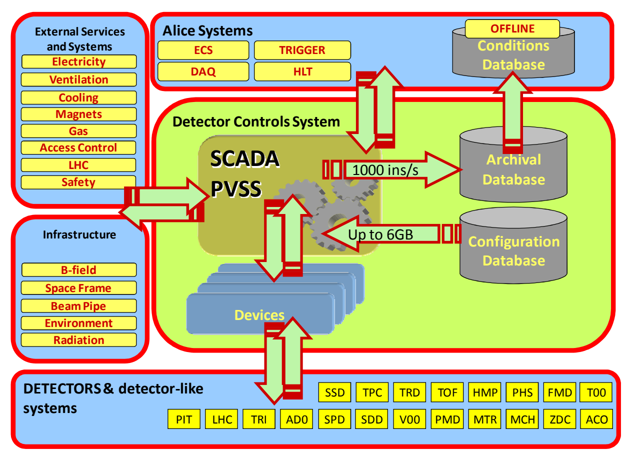

Figure 1: The main data streams managed by the central ALICE DCS.

Amongst the sources of data we count 1200 networked Front-End processors, 300 VME crates and power supplies and 4000 high voltage channels. From the detectors, the ALICE DCS receives data from 180,000 OPC items and 100,000 Front-End devices; eventually, a total of 1,000,000 parameters are supervised by the DCS at a typical rate of 1 Hz (see Fig. 1).

In order to be able to handle such a large amount of data, the ALICE DCS must process and filter the data before they are sent to the archives or transferred to the final users such as the offline analysts and the online displays. More than 100 detector computers represent the detectors side, while a group of 60 backend servers and an Oracle Database service, which is able to process up to 150,000 inserts/s, compose the central side of the control system (see Fig. 2).

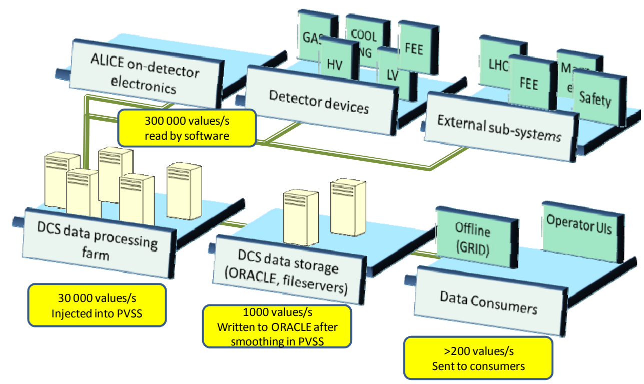

Figure 2. Data flow reduction from the detectors, to the archives and to the final customers.

The data flow is not unidirectional. Up to 6GB of configuration data is uploaded to detectors from the central ORACLE database before a physics run can be started.

Furthermore, the DCS is also the central collector of information of general interest, like the LHC beam data, data from various technical services, environmental data, radiation and magnetic field data, etc.

For example, global environment data is collected and archived and at the same time redistributed to the detectors, to be used for their calibrations or to allow for correlations with their internal variables. Much data is received from the LHC accelerator and managed centrally and distributed to several subsystems in order to define operational states and guarantee safe procedures.

In order to provide a controlled data exchange with the detectors, the ALICE DCS has exploited the distribution properties of PVSS.

The SCADA PVSS is a complex modular system where the different processes (aka managers) can run on a single machine or in a scattered way, i.e. on different machines. In a scattered system, managers run on dedicated machines, in order to distribute the load; in a distributed system, remote, single or scattered, the systems are interconnected through dedicated daemons, so called distribution managers (see Fig. 3); via this connection there is only information sharing and .no sharing of computing resources.

THE ALICE DCS TOPOLOGY

The SCADA PVSS system not only allows to collect, publish and archive the various data. Thanks to its distribution properties and its modular architecture it has been possible to design the ALICE DCS as a distributed system where each detector is a complete control system itself. The detectors have implemented their DCS on more than one machine and are therefore distributed systems themselves.

The smaller detectors have implemented their controls on a few machines which are interconnected in a hierarchical structure. The bigger detectors have implemented their controls over several machines which are interconnected in a mesh, in order to easily share the data inside the detector community. Interactive operation takes place in both cases via a remote user interface on one of the nodes.

A particular group of machines is the LHC cluster, which is responsible for collecting the data from the LHC accelerator and for the publication of ALICE information back to LHC. Some of the these parameters, such as parameters from luminometers and beam monitors, are of direct interest also to the many detectors; they are therefore collected and sent both to LHC and the ALICE detectors. Since its installation, the LHC cluster has represented an interesting testbed for the ALICE distributed topology. A big effort has been devoted to the

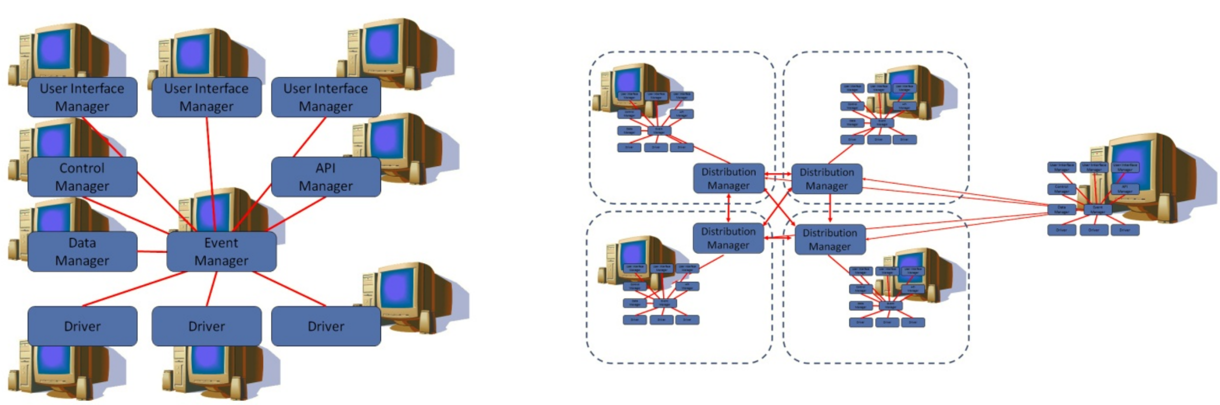

Figure 3: Schema of a scattered system (left) where the PVSS managers of a single project are running on different machines and a distributed system (right) in which the distribution managers handle the connections between remote projects

optimization of the data distribution, trying to limit interconnections, monitoring dependencies, and ensuring the availability and security of the data.

MONITORING THE DISTRIBUTED ENVIRONMENT

Within this distributed scheme, the ALICE DCS has been implemented as a hierarchical system, with the aim of limiting the interdependencies of the various subsystems. All the detectors nodes connect to the central servers (Fig. 4); in this way, ALICE DCS can provide global data, keep the systems synchronized and have a clear monitor of possible interdependencies.

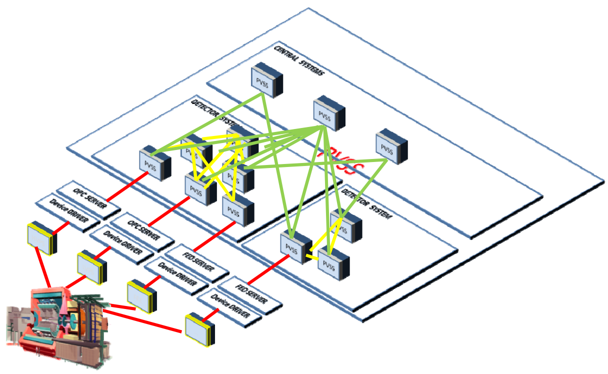

Figure 4. Central servers connect to all detectors nodes. Different detectors are not connected directly.

Connecting remote systems in a distributed environment is very simple. However, this can easily lead to unforeseen and unwished inter dependencies and unnecessary or redundant data flows. Furthermore, a distributed connection can be initiated unilaterally, even without the approval of the target system.

The central ALICE DCS group has therefore initiated a project to map, understand and monitor the data flow between central systems and detectors and between the detectors themselves. The aim is not only to prevent or block unauthorized connections, but mainly to better understand the needs of the different systems and to find the best ways to provide the data they need.

A monitoring tool has been developed which compares a list of authorized, local connections, with the actual connections observed. The authorization list can be maintained and updated from PVSS directly (Fig. 5). The tool has been built using simple PVSS control scripts.

The project provides an interactive graphical representation of the state of connections (Fig. 6). From here, it is easy to assess the degree of interdependency between the systems and to understand the optimal layout of the distributed connections. Used regularly this tool helps to optimize the flow of data between systems and to reduce the risk of faulty interactions between different systems.

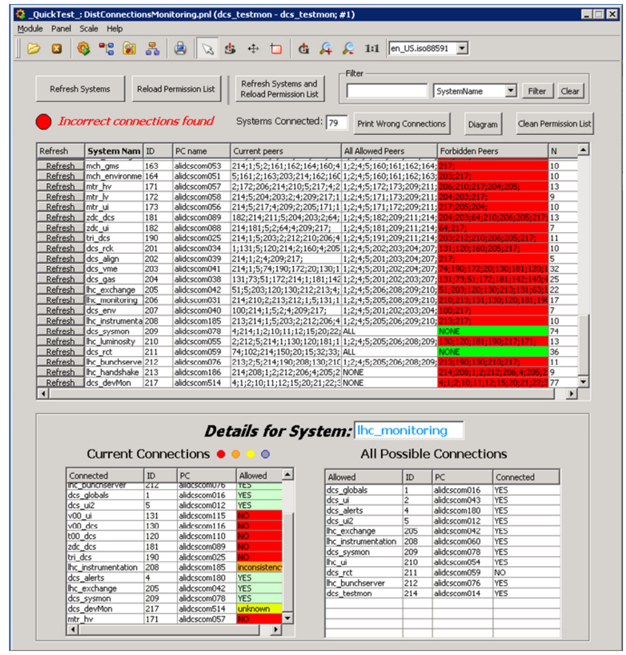

Figure 5. Expert panel to monitor the distributed systems.

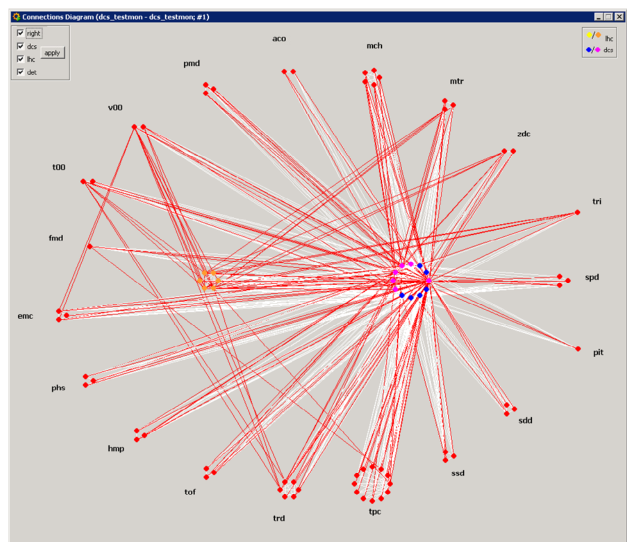

Figure 6. Graphical representation of the ALICE DCS interconnections.

CONCLUSIONS

The ALICE DCS has made a big effort to standardize, simplify and optimize the flow of data between computer systems in the ALICE experiment.

For this purpose a tool has been developed to analyze the network connections between distributed systems. This monitoring tool is raising interest inside the CERN JCOP Framework community, and will hopefully be integrated in a more general monitoring context, to give to the PVSS central developers another tool to optimize and simplify the data flows in the operational networks.