

ALICE DCS PREPARATION FOR RUN 3

Alexander Kurepin^{1,2, a}, **André Augustinus**^{1, b}, **Peter Matthew Bond**^{1, c},
Peter Chochula^{1, d}, **John Larry Lang**^{3, e}, **Mateusz Lechman**^{1,4, f},
Ombretta Pinazza^{1,5, g}, **Kevin Cifuentes Salas**^{6, h}

¹ CERN, Geneva, Switzerland

² Institute for Nuclear Research, Russian Academy of Sciences, Moscow, Russia

³ University of Helsinki, Finland

⁴ Institute of Physics, Slovak Academy of Sciences, Bratislava, Slovakia

⁵ INFN Sezione di Bologna, Bologna, Italy

⁶ Universidad de Deusto, Bilbo, Spain

E-mail: ^a alexander.kurepin@cern.ch, ^b andre.augustinus@cern.ch, ^c p.bond@cern.ch,
^d peter.chochula@cern.ch, ^e john.larry.lang@cern.ch, ^f mateusz.lechman@cern.ch,
^g ombretta.pinazza@cern.ch, ^h kevin.cifuentes.salas@cern.ch

The ALICE experiment is heavy ion collision detector at the CERN LHC. Its goal to study extreme phase of matter – called quark-gluon plasma. It is collaboration of 41 countries and more than 1800 scientists. A large number of complex subsystems requires supervision and control system. ALICE Control Coordination (ACC) is the functional unit mandated to coordinate the execution of the Detector control system (DCS).

In 2020, the ALICE experiment at CERN will start collecting data with upgraded detector. The ALICE upgrade addresses the challenge of reading out and inspecting the Pb-Pb collisions at rates of 50 kHz, sampling the pp and p-Pb at up to 200 kHz. ALICE O2 project merges online and offline into one large system with ~8400 optical links, data rate 1.1 TB/s, data storage ~60PB/year. From DCS O2 requires continuous data flow with ~100 000 conditions parameters for event reconstruction. Data has to be injected into each 50ms data frame.

DCS-O2 interface consists of electronics and software modules for configuring CRU controllers and provide continuous dataflow to O2 system.

We describe the architecture and functionality of the ADAPOS mechanism. We will discuss the requirements and results obtained during the test campaign. We will also provide a description of a new front-end access mechanism allowing for detector control in parallel to the data acquisition.

Keywords: slow control, DCS, LHC, ALICE

© 2018 André Augustinus, Peter Matthew Bond, Peter Chochula, Alexander Kurepin, John Larry Lang, Mateusz Lechman, Ombretta Pinazza, Kevin Cifuentes Salas

1. The ALICE experiment

The ALICE experiment is one of four big experiments on LHC ring at CERN. Its goal studying one of the states of matter called quark-gluon plasma.

Each subdetector should stay autonomous from the other subsystems and to guaranty this the central distributed system is segmented into subdetector systems. The ALICE detector consists of 19 subdetectors, built with different detection technologies, with different operational requirements and a number of infrastructure services, all supervised by a single operator. The architecture of ALICE Detector Control System (DCS) is based on standards adopted by the LHC experiments at CERN.

During twelve years of ALICE operation, the initial architecture and system implementation proved their reliability and robustness. The DCS provided stable 24/7 services with small interruptions required mainly for the infrastructure modifications (cooling and ventilation upgrades, reorganization of the computer racks etc.) Even during the service breaks, the core systems related to safety remain operational.

1.1 The ALICE Detector Control System

The ALICE experiment is using a commercial SCADA system WINCC OA, extended by CERN JCOP and ALICE software frameworks. It is configured as a large distributed system running on about 100 servers [1] The main advantage of use WINCC OA is that CERN is widely using this system and we take a lot of advantage from the huge knowledge base. To minimize the impact of the human factor, many procedures have been optimized and automatized in order to allow the operator to supervise about 1 000 000 parameters from a single console.

Where ever possible ALICE DCS uses commercial hardware and software standards [2] while for nonstandard devices, like detector custom frontend electronic modules, a client-server mechanism is used based on the DIM communication protocol.

1.2 New computing system O2

The interaction rates at LHC in ALICE during the RUN3 period, planned to start in 2021, will increase by a factor of 100. A new Combined Online Offline computing facility (O2) [3] has been developed to cope with the new data processing requirements.

The detector readout will be upgraded and it will provide 3.4 TByte/s of data, carried by 9 000 optical links to a First Level Processing (FLP) farm consisting of 270 servers.

The data taken during the same period of time by individual detectors is merged and sent 1600 Event Processing Nodes (EPN), deploying ~100 000 CPU cores. After initial processing, a compressed volume of 100 GByte/s will be transferred to the computing GRID facilities.

Due to high interaction rates, it is not possible to store the detector data on disk and perform the processing once the run ends. Even the local disk storage of 50PB, which will be installed in ALICE, is not sufficient for the whole data processing. The detector data will therefore be processed immediately on the EPN farm. The event reconstruction will allow for significant reduction of the data volumes before it is passed to the GRID for further analysis. To allow for such analysis, the DCS conditions data needs to be provided together with the detector data.

The new readout electronics will be accessed through the CERN GBT link [4]. The traffic on this link will be shared between the DCS and the data acquisition system, which requires a new concept for the front-end control to be implemented.

The conditions data handling and front-end electronics access represent the two main fields of new DCS developments for the LHC RUN3.

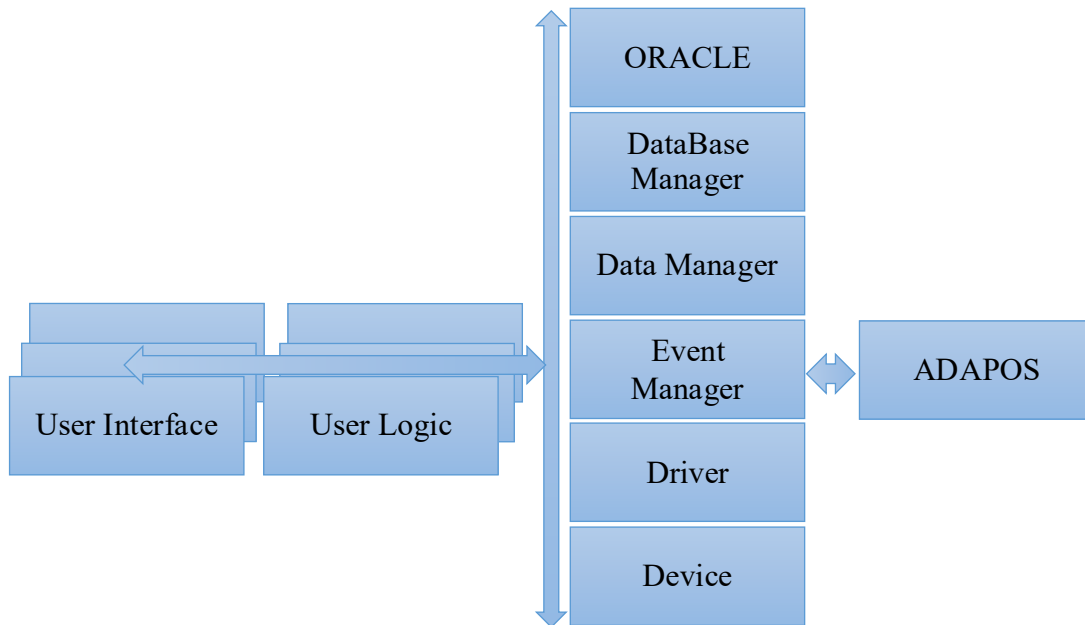
2. Conditions data

The conditions data, currently collected after each run, will be provided continuously to a farm containing 100 000 CPU cores and tens of PB of storage. The typical reading frequency for the DCS devices is ~ 1Hz but of course, it largely depends on device. Usually the DCS data readout is not triggered externally and it is driven by the conditions change. As a result, the data arrival time to the DCS is not regular and cannot be predicted. The DCS devices equipped with commercial OPC servers

usually do not allow for a parameter readout on demand. The firmware of most of the devices performs internal poll of all devices, before it makes the data available to the OPC server.

2.1 New Conditions data flow

The current DCS assembles the conditions data block using the ORACLE archive after each run, new requirement results in the increase of the DCS data-publishing rate by a factor of 5000.



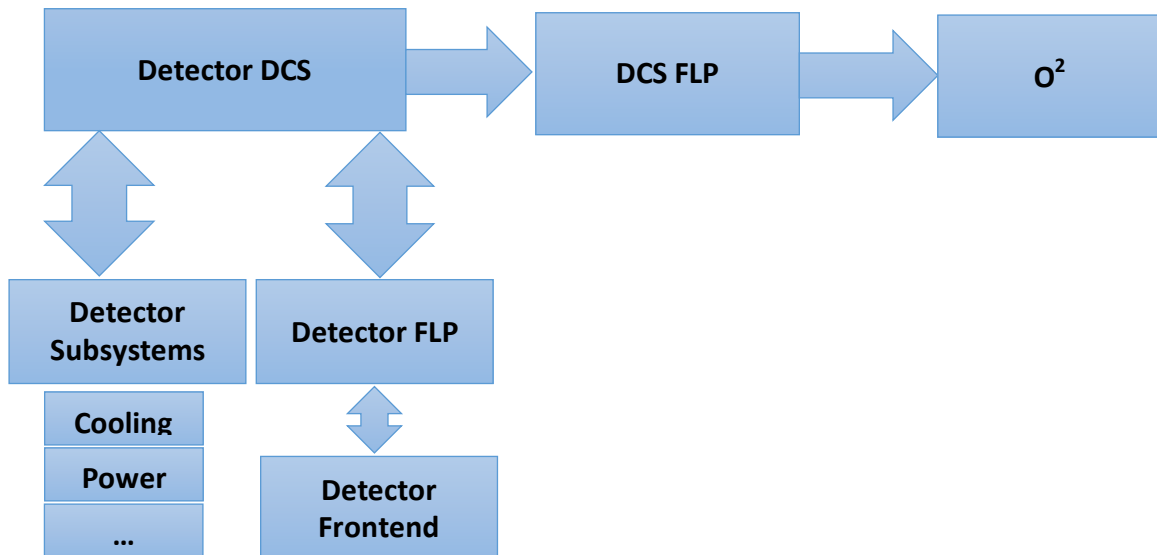
The newly developed approach is based on a process image stored in the computer memory. The upgraded system must ensure a steady streaming of all conditions, each 50 ms data frame of data taking must be complemented with a block of about 100 000 DCS parameters for reconstruction in the O2 facility. Each of the 1500 EPN nodes will receive the same conditions data for each 50ms data frame. This large formatted data block contains all information about the monitored channels required for the O2 processing. This information consists of the channel name, timestamped values and various flags describing the data type and quality information. The DCS creates this block at the startup time and populates it with the latest known values. As the new data arrives to the DCS, the process image is updated. At each period of time the process image contains the actual status of all condition parameters.

The conditions data flow from the DCS to O2 is managed by the Alice DATAPOINT Service (ADAPOS). Its task is to collect the DCS data and assemble a memory block with conditions data required for the further stages of reconstruction.

Data to be published to the O2 is scattered across ~120 WINCC OA systems. At the startup time, the ADAPOS server first needs to locate the data and subscribe to each value. The present implementation is based on DIM protocol [5], supported on the WINCC OA platform at CERN. For each value a corresponding DIM service is created and the published value is updated on each change. Using the DIM DNS service ADAPOS finds the values and establishes a peer-to-peer connection to each publishing server.

3. New frontend access

Part of the DCS information is produced by the detector frontend modules. The firmware of the receiver cards extract the DCS information off the data stream and publishes it to the DCS clients implemented in WINCC OA. Each client subscribes to a required subset of published values without the need to know the details on physical configuration of the frontends and FLPs. A common name service will handle the redirections of subscription requests. One of the already existing technologies supporting this mode of operation is DIM.



The DCS data is extracted from the data stream and sent to ALF (Alice Low Level Front-end) interface, which publishes the data to the upper layers of the software. ALF can also receive commands and converts them to data words to be sent to the front-end electronics. To keep the ALF detector neutral, its functionality is restricted to the basic I/O operations. In the current implementation, the ALF can read/write registers implemented on the front-end modules and publish the data using a DIM service. The data published by ALF could be single values, or blocks of data prepared by the electronics modules.

Communication between the WINCC OA systems and the ALFs is managed by the Front-End Device (FRED) module. This layer provides the necessary translation of high level WINCC OA commands to simple ALF transaction and unpacking the ALF data before it is published to the WinCC OA system.

The ALF-FRED architecture decouples the front-end details from the high level SCADA system. Separating this task into 3 layers of software – the drivers, the ALF and the FRED, brings clear advantages – the ALF remains detector independent and can be deployed to all detectors. The FRED layer provides highly customizable detector-specific module which implements all resource intensive calculations, such as data unpacking and first level filtering.

The WINCC OA system implements the standard controls functionality covering the control, monitoring, alert handling and data visualization and archival. From the WINCC OA perspective, the ALF-FRED behaves as any other standard device.

The present ALF-FRED implementation is based on DIM protocol. It allows for easy integration of complex detector granularity into a coherent system. The ALF modules are installed for each FLP belonging to the same detector and are serviced by a dedicated FRED. The SCADA system recognizes the FRED as a device that recognizes high level commands (such as configure, turn on/off) and publishes its data as single services. It is the task of FRED to translate the high-level commands into a sequence of atomic actions to be carried out by ALF.

The separation of the commands and management of their complexity through the different ALF-FRED layer brings an additional advantage. At the lowest layer, the ALF is interfaced to the CRU through a driver developed in ALICE. It takes care of CRU transactions over the GBT link. [6] Replacing this driver allows for use of a different field bus (such as CANbus) without the need of extensive software modifications. The ALF with modified driver will provide the same functionality to FRED and the modification of the front-end access will remain transparent. This approach is used in ALICE for improving the redundancy. Some detectors implement CANbus in their front-end. During the downtime of the GBT link (for example during the FLP maintenance), the CANbus will take over the controls functionality and will eliminate the need for shutting down the operations. The communication speed will be strongly reduced, however it will remain sufficient for assuring the detector safety and will even allow for less demanding applications (such as detector calibration or debugging of the operational procedures).

4. Conclusions

The ALICE upgrade is a challenging task for the whole collaboration. Upgraded detector hardware and change of the data processing strategy required the redesign of the DCS data flow. The data will be provided to the O2 processing facilities in streamed mode for which the DCS data stream has to be diverted from its standard path.

Access to new front-end modules will be shared between the DCS and the Data acquisition. A new device access strategy covers the complexity of the new hardware and provides a flexible mechanism for transferring the information between the DCS and the hardware. Separation of the functionalities using a three layer architecture allows for splitting of the common and detector specific tasks. This approach improves the maintenance and allows for shared developments.

References

- [1] P.Chochula et al., Operational experience with the ALICE Detector Control System // Proceedings ICALEPCS 2013, San Francisco 2013
- [2] P.Chochula et al., Control and monitoring of the front-end electronics in ALICE // 9th Workshop on Electronics for LHC Experiments, Amsterdam 2003.
- [3] P. Moreira, A. Marchioro and K. Kloukinas, The GBT: a proposed architecture for multi-Gb/s data transmission in high energy physics // Topical Workshop on Electronics for Particle Physics, Prague Czech Republic September 3–7 2007, pg. 332 [CERN-2007-007.332].
- [4] M. Richter, A design study for the upgraded ALICE O2 computing facility // J. Phys.: Conf. Ser. 664 082046.
- [5] C. Gaspar, M. Donszelmann, DIM – A distributed information management system for the DELPHI experiment at CERN // Proceedings of the 8th Conference on Real-Time Computer applications in Nuclear, Particle and Plasma Physics, Vancouver, Canada 1993.
- [6] A. Caratelli et al., The GBT-SCA, a radiation tolerant ASIC for detector control and monitoring applications in HEP experiments // Topical Workshop on Electronics for Particle Physics 2014, Aix En Provence, France, 22 - 26 Sep 2014, pp.C03034.