STATUS REPORT OF THE MEASUREMENT SERVICE FOR THE CERN ACCELERATOR LOGGING

Marine Gourber-Pace, Grzegorz Kruk, Marek Misiowiec, CERN, Geneva, Switzerland

Abstract

The LHC (Large Hadron Collider) Logging service is aimed to satisfy the requirement of capturing and storing any relevant accelerator data to track its variation over time. This service is presently operational on the whole CERN accelerator complex, from ion and proton sources to LHC, and has become a critical component of the CERN control systems.

The focus is given to the measurement part of this service, which is responsible for the data acquisition and preparation (processing, filtering, concentration) prior to its storage in database and file systems. Incoming data is often processed by a concentration layer, the processes that transform data of multiple devices into single values according to well defined rules and then publish them further on, to the LHC Logging among others.

The paper describes the architecture and presents the solutions to the very challenging requirements imposed by the LHC in terms of overall performance and reliability. The efficiency of the data acquisition and filtering as well as the flexible software design are highlighted.

OVERVIEW

The CERN Accelerator logging service was designed in 2003 in view of the LHC Operation, with the aim of capturing and storing any relevant accelerator data to track its variation over time. The main objectives are twofold: (1) to reconstruct, from the stored data, the exact conditions of a particular event (as a beam loss) which occurred in the machine, (2) to keep time-series data from equipment and beam-related parameters on-line for the lifetime of the LHC and to allow data comparisons from previous years. From the first operational version in 2004 deployed for the LHC beam commissioning tests, the usefulness of such a service became obvious. Since then, its scope has expanded to the whole CERN accelerator complex, thanks to a modular and generic software design and data model.This service is presently fully operational, deployed on all sub-systems commissioned for the LHC and its pre-injectors and transfer lines.

DESIGN AND IMPLEMENTATION

Figure 1 presents a simplified view of the architecture highlighting the most important components of the logging service.

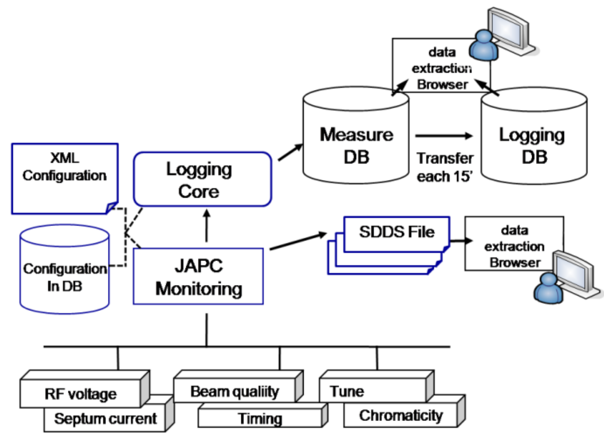

Figure 1: Architecture of logging service.

Main Components

The JAPC (Java Api for Parameter Control) [3] Monitoring is responsible for acquiring data from equipment parameters and distributing the data to the registered clients, the DB (DataBase) logger and the SDDS (Self Describing Data Sets) [4] logger. Therefore numerous data acquisition systems on front-end computers, ranging from accelerator equipment (as the RF cavity voltage) to beam-related parameters (like the machine tune) are monitored and provide the time-stamped data which is processed by one of the registered loggers. The SDDS service is used to handle complex data structures like beam captured image or profile. It comprises a logger which writes each acquisition into a SDDS file, an API (Application Programming Interface) for an efficient retrieval of data, and a file viewer to browse and visualize the content of the files. Complementary to the SDDS service, the DB service captures and stores critical data in a database management system and allows an easy reconstruction of the data variation over time as well as the correlation of data from different systems.

The Measurement DB acts as a high throughput, short term data store (7 days). It is fed by the DB Logger with raw time series data, measured at up to 2 Hz. Every 15 minutes, the Measurement DB transfers filtered data of interest to the Logging DB which serves as a high capacity, long term storage. Data extraction applications are available to provide statistical summaries of time series data, to extract data to a file in various formats, or visualize it on interactive charts [2].

DB Logger

Each DB logger is a process which runs continuously and relies on the logger core providing generic functionalities. It is configured with a custom set of attributes, pre-defined by the user in a standard format. These attributes are typically the list of accelerator and beam parameters to be logged, their monitoring frequency, the conversion rules to transform the raw parameter values received into the final logged values. The logger core implements the following business logic:

- capture from the JAPC Monitoring module the raw values as published by the parameters
- stamp the raw values with a proper time stamp. The correlation of logged data from different systems being a key functional requirement, it is primordial to ensure a coherent data time-stamping. For cycling machines like the LHC injectors, all parameters are stamped with the cycle stamp, representing the cycle's start [UTC]. The usage of this stamp guarantees that all acquisitions from various pieces of equipment which share one cycle stamp value actually correspond to the same physical beam. For the LHC, which is not a cycling machine, the acquisition stamp is used instead.
- perform data processing, filtering or concentration of certain values, according to the configured conversion rules. Transformation rules present a wide range of diversity and complexity from basic converters (like data multiplier) to complicated concentration mechanisms compliant with the very demanding requirements of LHC in terms of data volume per unit of time. The case of the concentrator for LHC BLM (Beam Loss Monitors) and BPM (Beam Position Monitors) is described in detail below.
- buffer and send at a regular frequency the processed values to the Measurement DB.
- report faults or failures which may have occurred in the data transfer to the alarm system, which in turn publishes them in the control room and notifies the system responsibles.

LHC BLM/BPM Concentrators

A LHC concentrator is a process that periodically reads data coming from multiple devices of the same type and transforms it into a single concentrated value. Concentrators combine raw data from numerous devices, apply complex transformation algorithms and finally distribute concentrated values to other systems via JAPC infrastructure. They allow to move common data conversion mechanisms from diverse client applications to a single place and perform this job once for all. Concentrators allow representing many devices as a single entity. They not only provide a way of logical conversion of data, but also protect devices from multiple accesses. Thanks to data concentration, client subscriptions can be moved from front-end machines to more capable server computers those processes operate on. Low level hardware is therefore relieved from providing the same data to multiple users at the same time. LHC BLM/BPM concentrators act as a substantial source of data for the Logging system. For BLM devices whose overall data production rate is 1MB/sec, concentrated values are fed directly to the Measurement DB. In other cases, Logging modules subscribe to concentrated values via JAPC API. LHC concentrators operate, as the Logging system, on a 24/7 basis.

CHALLENGES

This section presents the operational constraints in terms of performance, scalability as well as availability and reliability. These constraints have become extremely challenging with the commissioning of LHC. Moreover, a number of critical applications, exclusively relying on logged data, impose demanding requirements in terms of availability and reliability.

Scalability and Performance

Even if scalability was taken into account within the initial design, it rapidly appeared as an outstanding issue when the LHC requirements revealed to be much more stringent than anticipated in terms of data volume and rate. The initial figures in 2003 for LHC were estimated at 1TB/year during normal operation while presently data rates are already one order of magnitude higher, and LHC is not yet operating with beam. Presently, up to 35000 database signals receive data at frequencies varying from 0.02 to 2 Hz.

Intensive scalability tests performed in 2008 revelead a number of points of failure, which have been addressed since then. Among the implemented solutions, an appropriate *on change* filtering at the Measurement DB entry has been put in place: by eliminating repetitive values, it has allowed to significantly reduce the data volume for heavy clients like LHC power converters which continuously publish 1750 variables at 2 Hz.

Availability and Reliability

High-availability is another essential requirement for the Logging service, making the tolerance to logged data loss close to zero for certain critical applications. This is the case of the application running at the CNGS (CERN Neutrino to Gran Sasso) experiment, which records the rare muon events generated from neutrino beam interactions. A surveillance mechanism of critical data has been deployed: these data are surveyed by a dedicated process, which regularly reads the latest values saved in the Measurement DB and generates an alarm in case no values were written over a period of time, contrary to the expectations.

Monitoring and Instrumentation

To increase the reliability up to the requested level, instrumentation and adequate monitoring tools at all system levels are of the highest importance.

All operational logging processes are surveyed by the CERN controls-standard DIAMON/LASER service [1]. It captures the problems reported by the process in case of errors and detects failures of the process itself. It generates an alarm displayed in the control room and notifies by mail or sms the system's responsibles. In case of a process's crash, it gets immediately restarted in an automated and transparent way. Complementary diagnostic tools provide detailed performance metrics which show the activity of all processes and identify the stressed resources.

As we are still not in a stable operation phase for the LHC, logging on new devices need to be tested without interfering with the operational infrastructure. A replica of the logging operational infrastructure, including a copy of the Measurement DB has been deployed and is extensively used by equipment experts. This guarantees an exhaustive validation of the logging chain in the test environment before switching to the operational one.

Instrumentation on Measurement DB has also heen enabled: all applications accessing the DB need to identify themselves before being granted a connection, application transactions are monitored and registered in terms of number of variables and data volume, etc…
A detailed description is provided in [2]. The deployment of such an instrumentation has proven to be very useful to identify offending processes:

Backup Service

A powerful backup service has been put in place to prevent any loss of data in case of unexpected unavailability of the Measurement Database. The service detects the problem and, in an automated and transparent way, starts saving the data to an alternative independent support (file system). Data are then transferred back to the Measurement Database as soon as it gets operational again.

OUTLOOK

The logging service is anticipated to continue to significantly grow along two dimensions: (1)increasing volume and rate of device data, (2)increasing complexity of data converters. Given the mission-criticality, the system's performance, scalability and availability will have to be continuously re-evaluated.

Moreover, new requirements for extended functionalities have been expressed in 2008, the most important being:

- on line monitoring service: an interactive graphical interface will allow the user to define an arbitrary set of parameters, enable the logger, real time visualize the acquisitions' variation over time. The design of the core has to be reviewed to integrate the requested extensions: dynamic configuration and activation of the logger, immediate availability of the acquisitions This is under analysis now.
- a database driven configuration: the current XML logger configuration will be replaced by a database model, unified with the Measurement and Controls Configuration databases. The preparation of the configuration will significantly gain in automation and efficiency. This development is undergoing and foreseen to be released end 2009.
- an API to convert logged data into a JAPC parameter value [3]: this will allow to retrieve a time-series of values from the Measurement or Logging DB in a JAPC compliant format. The requirement is currently being evaluated with the clients.

CONCLUSION

The CERN accelerator logging is considered as being an essential component of the CERN control system, serving as the main information provider for the post-mortem analysis of any past event. Its powerful and reliable software infrastructure has proven to be able to address the demanding requirements from the LHC. Nevertheless the challenge to keep up with the increasing expectations is still ahead with the LHC entering in operation in the coming weeks.

REFERENCES

[1] P. Charrue, "Diagnostic and Monitoring CERN Accelerator Controls Infrastructure. The DIAMOND Project, "ICALEPCS'09, Kobe, Japan, October 2009.

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

[3] V. Baggiolini et al., "JAPC - the Java API for Parameter Control", ICALEPCS'05, Geneva, Switzerland, October 2005.

[4] M. Borland, "User's Guide for SDDS Toolkit"; http://www.aps.anl.gov/asd/oag/manuals/SDDStoolkit/SDDStoolkit.html.