The ALICE DAQ infoLogger

Sylvain Chapeland, on behalf of the ALICE Collaboration

ALICE detector

ALICE is the heavy-ion detector designed to study physics of strongly interacting matter and the quark-gluon plasma at the at the CERN Large Hadron Collider.

The apparatus includes high resolution tracking (silicon detectors, a large time-projection chamber), particle identification, and triggering elements. It features two large magnets, a main solenoid and a dipole on the Muon arm.

It primarily targets heavy-ion lead-lead collisions, but it also has a substantial physics program with proton-proton and proton-ion collisions.

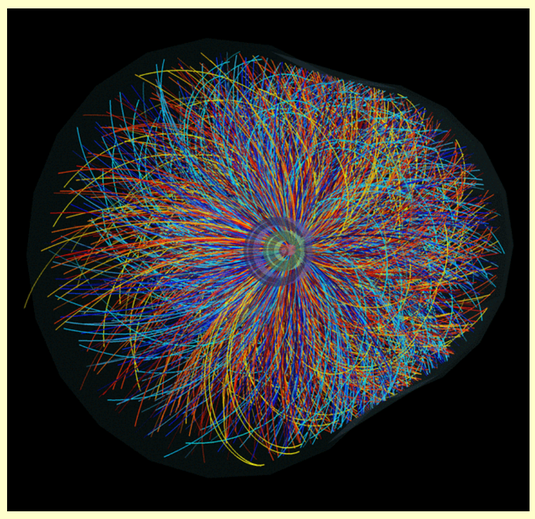

An example of lead ion collision as seen by ALICE

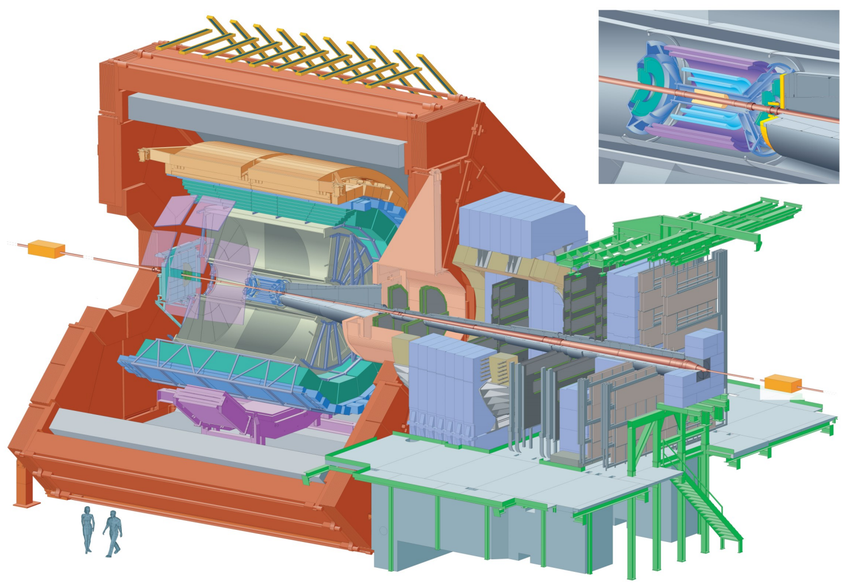

The ALICE detector

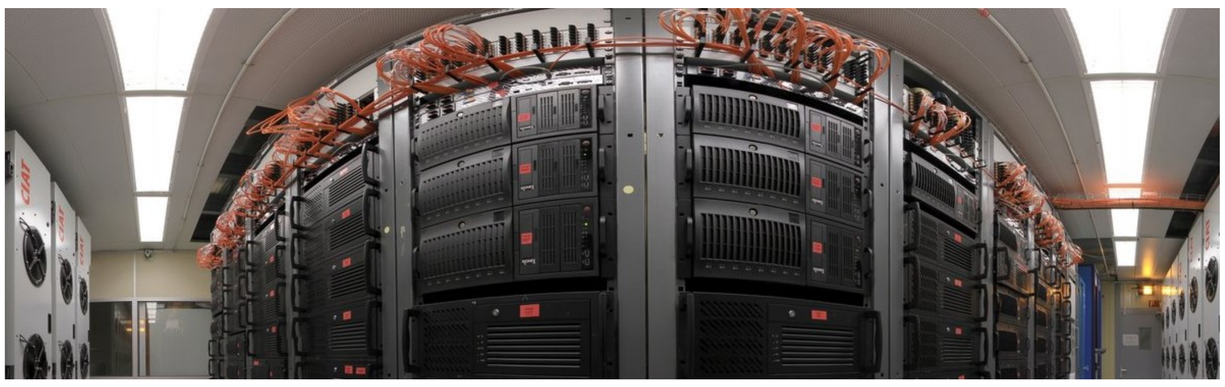

The ALICE DAQ counting room

ALICE data-acquisition system (DAQ)

The DAQ handles the data flow from the detector electronics to the archiving on tape. The DAQ reads out the data through 500 dedicated optical links at an aggregated and sustained rate of up to 10 Gigabytes per second, and stores at up to 2.5 Gigabytes per second. Since 2010, the ALICE DAQ recorded 6.5PB of data to permanent storage. At run time, over 6000 distributed processes are active to handle the data flow on the 320 PCs of the DAQ facilities.

The infoLogger: a package to generate, transport, collect, store, and consult log messages.

The infoLogger provides an interface to inject logs, a central repository to store the messages, and user interfaces to display and query them. It is used in the ALICE DAQ and 3 other online systems (Experiment Control System, Trigger, and High-Level Trigger). Any distributed process can create a log message by calling one of the infoLogger library functions or the command-line executable. The message is then sent to the local infoLoggerReader daemon, which collects all the logs of the node where it runs, and sends them to a central infoLoggerServer daemon, where the received messages are stored in a database. The infoBrowser user interface allows reading messages, either queried from the database or received on-line by the server.

infoLogger architecture

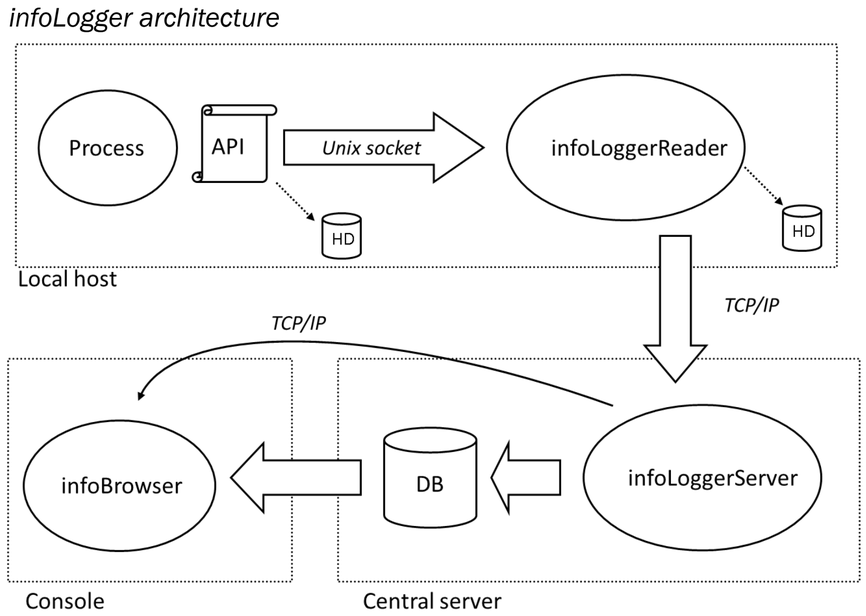

- Library overhead: 4 µs to inject a log message from client to reader.
- Transient disk buffers are used when the log pipeline is unavailable. Flood prevention possible by setting a limit on maximum number of messages in a burst.
- The server is multi-task: reception, message processing, db insert and GUI dispatch run in separate threads.
- MySQL database stores logs. For 2012, 160 million rows, 83GB of MySQL data files, including indexes (53% of db size). Table is partitioned by log time.
- In average, about 500000 messages are logged and archived every day for the DAQ system. The infoLogger has been in production since 2004.

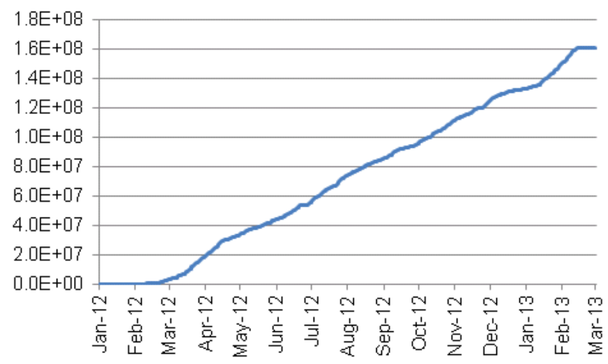

Number of log messages collected over time by infoLogger in 2012-2013

In the control room

At the DAQ control position, the 2 screens on the right are used to display incoming log messages. The top position shows all messages, whereas the bottom one keeps only error messages for better impact on the operator.

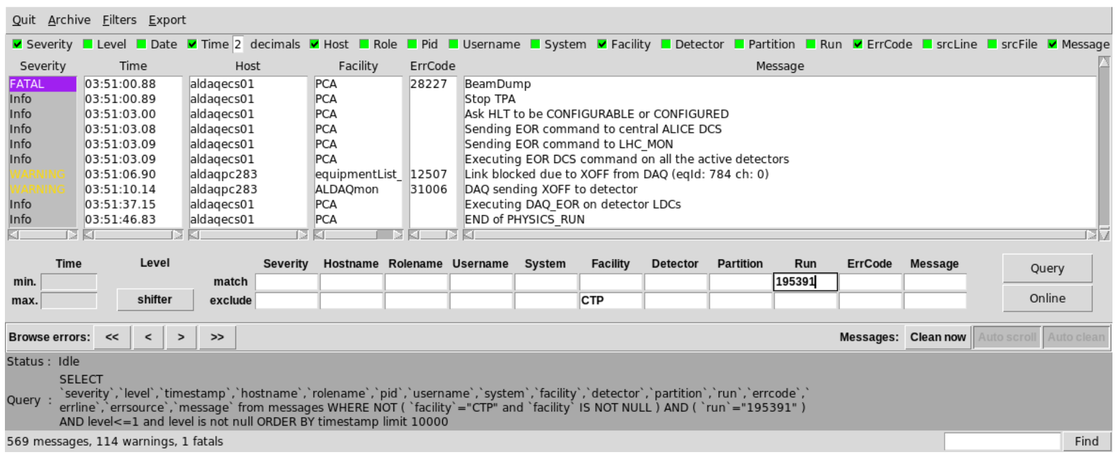

Example view of the infoBrowser

This is the log display and query interface. Here is presented the Tcl/Tk version. A Web query interface is also available (PHP / MySQL / Apache / Shibboleth / CERN Single Sign-ON).

Documentation on errors is directly available from the GUI. A TikiWiki CMS groupware holds the DAQ operation knowledge base, from which is extracted the contextual display: description, troubleshooting procedure, severity, contact, etc.

Log fields

The following tags are used to structure the log information:

Severity, Level, Timestamp, Host name, Role name, Process ID, System, Facility, Detector, Partition, Run number, Error code, Source line, Source file, Message.

This list can easily be extended according to needs. The message structure format is versioned.

Contact: sylvain.chapeland@cern.ch

Conference on Computing in High Energy Physics, October 2013, Amsterdam, The Netherlands