Contribution ID: **337** Type: **Oral**

Scalable monitoring data processing for the LHCb software trigger

Monday, 4 November 2019 15:00 (15 minutes)

The LHCb high level trigger (HLT) is split in two stages. HLT1 is synchronous with collisions delivered by the LHC and writes its output to a local disk buffer, which is asynchronously processed by HLT2. Efficient monitoring of the data being processed by the application is crucial to promptly diagnose detector or software problems. HLT2 consists of approximately 50000 processes and 4000 histograms are produced by each process. This results in 200 million histograms that need to be aggregated for each of up to a hundred data taking intervals that are being processed simultaneously. This paper presents the multi-level hierarchical architecture of the monitoring infrastructure put in place to achieve this. Network bandwidth is minimised by sending histogram increments and only exchanging metadata when necessary, using a custom lightweight protocol based on boost::serialize. The transport layer is implemented with ZeroMQ, which supports IPC and TCP communication, queue handling, asynchronous request/response and multipart messages. The persistent storage to ROOT is parallelized in order to cope with data arriving from a hundred of data taking intervals being processed simultaneously by HLT2. The performance and the scalability of the current system are presented. We demonstrate the feasibility of such an approach for the HLT1 use case, where real-time feedback and reliability of the infrastructure are crucial. In addition, a prototype of a high-level transport layer based on the stream-processing platform Apache Kafka is shown, which has several advantages over the lower-level ZeroMQ solution.

Consider for promotion

No

Authors: AAIJ, Roel (Nikhef National institute for subatomic physics (NL)); MATEV, Rosen (CERN); PETRUCCI, Stefano (University of Edinburgh, CERN)

Presenter: PETRUCCI, Stefano (University of Edinburgh, CERN)

Session Classification: Track 1 –Online and Real-time Computing

Track Classification: Track 1 –Online and Real-time Computing