



Contribution ID: 151

Type: Oral

Let's get our hands dirty: a comprehensive evaluation of DAQDB, key-value store for petascale hot storage.

Tuesday, 5 November 2019 16:45 (15 minutes)

Data acquisition (DAQ) systems are a key component for successful data taking in any experiment. The DAQ is a complex distributed computing system and coordinates all operations, from the data selection stage of interesting events to storage elements.

For the High Luminosity upgrade of the Large Hadron Collider (HL-LHC), the experiments at CERN need to meet challenging requirements to record data with a much higher occupancy in the detectors. The DAQ system will receive and deliver data with a significantly increased trigger rate, one million events per second, and capacity, terabytes of data per second.

An effective way to meet these requirements is to decouple real-time data acquisition from event selection. Data fragments can be temporarily stored in a large distributed key-value store. Fragments belonging to the same event can be then queried on demand, by the data selection processes.

Implementing such a model relies on a proper combination of emerging technologies, such as persistent memory, NVMe SSDs, scalable networking, and data structures, as well as high performance, scalable software.

In this paper, we present DAQDB, an open source implementation of this design that was presented earlier, with an extensive evaluation of this approach, from the single node to the distributed performance. Furthermore, we complement our study with an in-depth comparison with the state-of-the-art solutions, a description of the challenges faced and the lessons learned.

Consider for promotion

No

Primary authors: CICALESE, Danilo (CERN); LE GOFF, Fabrice (CERN); LEHMANN MIOTTO, Giovanna (CERN); Mr JERECZEK, Grzegorz (Intel Corporation); RADTKE, Jakub (Intel); Mr SCHMIEGEL, Jakub (Intel); LOVE, Jeremy Robert (Argonne National Laboratory (US)); MACIEJEWSKI, Maciej; Ms SZYCHOWSKA, Malgorzata (Intel); MOMMSEN, Remi (Fermi National Accelerator Lab. (US))

Presenter: Mr JERECZEK, Grzegorz (Intel Corporation)

Session Classification: Track X – Crossover sessions

Track Classification: Track X – Crossover sessions from online, offline and exascale