



Contribution ID: 198

Type: Oral

Design of the data distribution network for the ALICE Online-Offline (O2) facility

Tuesday 5 November 2019 12:00 (15 minutes)

ALICE (A Large Ion Collider Experiment), one of the large LHC experiments, is currently undergoing a significant upgrade. Increase in data rates planned for LHC Run3, together with triggerless continuous readout operation, requires a new type of networking and data processing infrastructure.

The new ALICE O2 (online-offline) computing facility consists of two types of nodes: First Level Processors (FLP): containing a custom PCIe cards to receive data from detectors, and Event Processing Nodes (EPN): compute dense nodes equipped with GPGPUs for fast online data compression. FLPs first buffer the detector data for a time interval into SubTimeFrame (STF) objects. A TimeFrame then aggregates all corresponding STFs from each FLP into the TimeFrame (TF) object, located on a designated EPN node where it can be processed. The data distribution network connects FLP and EPN nodes, enabling efficient TimeFrame aggregation and providing a high quality of service.

We present design details of the data distribution network tailored to the requirements of the ALICE O2 facility based on the InfiniBand HDR technology. Further, we will show a scheduling algorithm for TimeFrame distribution from FLP to EPN nodes, which evenly utilizes all available processing capacity and avoids creating long-term network congestion.

Consider for promotion

No

Author: NESKOVIC, Gvozden (Johann-Wolfgang-Goethe Univ. (DE))

Presenter: NESKOVIC, Gvozden (Johann-Wolfgang-Goethe Univ. (DE))

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

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