



Contribution ID: 475

Type: **Poster**

Feasibility tests of RoCE for the cluster-based event building in LHCb

Thursday, 7 November 2019 16:15 (15 minutes)

This paper evaluates the utilization of RDMA over Converged Ethernet (RoCE) for the Run3 LHCb event building at CERN. The acquisition system of the detector will collect partial data from approximately 1000 separate detector streams. Total estimated throughput equals 40 terabits per second. Full events will be assembled for subsequent processing and data selection in the filtering farm of the online trigger. As a result, inter-node large-throughput transmissions with a combination of 100 and 25 Gigabit-per-second will be essential features of the system. Therefore, the data exchange mechanism of the cluster must utilize memory-lightweight data transmission protocols.

In this work, the RoCE high-throughput kernel bypass Ethernet-based protocol is benchmarked as an applicable technology for the event building network. CPU and memory bandwidth utilization for RoCE-based data transmissions is investigated and discussed. A comparison of RoCE with InfiniBand protocol is presented. Preliminary performance results are discussed with the selected network hardware supporting the protocol. Relevant utilization and interoperability issues are detailed along with lessons learned along the road.

Consider for promotion

No

Primary authors: KRAWCZYK, Rafal Dominik (CERN); COLOMBO, Tommaso (CERN); NEUFELD, Niko (CERN); PISANI, Flavio (Universita e INFN, Bologna (IT)); VALAT, Sebastien (CERN)

Presenter: KRAWCZYK, Rafal Dominik (CERN)

Session Classification: Posters

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