



Contribution ID: 214

Type: **Poster presentation**

Benchmarking message queue libraries and network technologies to transport large data volume in the ALICE O2 system

Tuesday, June 7, 2016 3:00 PM (1h 30m)

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

ALICE has been successfully collecting physics data of Run 2 since spring 2015. In parallel, preparations for a major upgrade, called O2 (Online-Offline) and scheduled for the Long Shutdown 2 in 2019-2020, are being made. One of the major requirements is the capacity to transport data between so-called FLPs (First Level Processors), equipped with readout cards, and the EPNs (Event Processing Node), performing data aggregation, frame building and partial reconstruction. It is foreseen to have 268 FLPs dispatching data to 1500 EPNs with an average output of 20 Gb/s each. In overall, the O2 processing system will operate at terabits per second of throughput while handling millions of concurrent connections. To meet these requirements, the software and hardware layers of the new system need to be fully evaluated.

In order to achieve a high performance to cost ratio three networking technologies (Ethernet, InfiniBand and Omni-Path) were benchmarked on Intel and IBM platforms.

The core of the new transport layer will be based on a message queue library that supports push-pull and request-reply communication patterns and multipart messages. ZeroMQ and nanomsg are being evaluated as candidates and were tested in detail over the selected network technologies.

This paper describes the benchmark programs and setups that were used during the tests, the significance of tuned kernel parameters, the configuration of network driver and the tuning of multi-core, multi-CPU, and NUMA (Non-Uniform Memory Access) architecture. It presents, compares and comments the final results. Eventually, it indicates the most efficient network technology and message queue library pair and provides an evaluation of the needed CPU and memory resources to handle foreseen traffic.

Primary author: WEGRZYNEK, Adam Tadeusz (Warsaw University of Technology (PL))

Co-authors: FUCHS, Ulrich (CERN); CHIBANTE BARROSO, Vasco (CERN)

Presenter: WEGRZYNEK, Adam Tadeusz (Warsaw University of Technology (PL))

Session Classification: Poster session 1

Track Classification: Fast Data Transfer Links and Networks