

Using ALFA for high throughput, distributed data transmission in ALICE O2 system

Thursday 13 October 2016 16:30 (15 minutes)

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 of the computing system, 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 Nodes), 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.

The ALFA framework will standardize and handle software related tasks such as readout, data transport, frame building, calibration, online reconstruction and more in the upgraded computing system.

ALFA supports two data transport libraries: ZeroMQ and nanomsg. This paper discusses the efficiency of ALFA in terms of high throughput data transport. The tests were performed using multiple FLPs, each of them pushing data to multiple EPNs. The transfer was done using push-pull communication pattern with multipart message support enable or disabled. The test setup was optimized for the benchmarks to get the most performant results for each hardware configuration. The paper presents the measurement process and final results –data throughput combined with computing resources usage as a function of block size, and in some cases as a function of time.

The high number of nodes and connections in the final set up may cause race conditions that can lead to uneven load balancing and poor scalability. The performed tests allow to validate whether the traffic is distributed evenly over all receivers. It also measures the behavior of the network in saturation and evaluates scalability from a 1-to-1 an N-to-N solution.

Tertiary Keyword (Optional)

Distributed data handling

Secondary Keyword (Optional)

Network systems and solutions

Primary Keyword (Mandatory)

DAQ

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

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

Session Classification: Posters B / Break

Track Classification: Track 1: Online Computing