



Contribution ID: 161

Type: Oral presentation

A Lossless Network for Data Acquisition

Tuesday, June 7, 2016 4:50 PM (20 minutes)

The planned upgrades of the experiments at the Large Hadron Collider at CERN will require higher bandwidth networks for their data acquisition (DAQ) systems. The network congestion problem arising from the bursty many-to-one communication pattern, typical for these systems, will become more demanding. It is questionable whether commodity TCP/IP and Ethernet technologies in their current form will be still able to effectively adapt to the bursty traffic without losing packets due to the scarcity of buffers in the networking hardware. We continue our study of the idea of lossless switching in software running on commercial-off-the-shelf servers for data acquisition systems, using the ATLAS experiment as a case study. The flexibility of design in software, performance of modern computer platforms, and buffering capabilities constrained solely by the amount of DRAM memory are a strong basis for building a network dedicated to data acquisition with commodity hardware, which can provide reliable transport in congested conditions.

In this presentation we extend the popular software switch, Open vSwitch (OVS), with a dedicated, throughput-oriented buffering mechanism for data acquisition. We compare the performance under heavy congestion of typical Ethernet switches to a commodity server acting as a switch, equipped with twelve 10 Gbps Ethernet interfaces providing a total bandwidth of 120 Gbps. Preliminary results indicate that software switches with large packet buffers perform significantly better, reaching maximum bandwidth, and completely avoiding throughput degradation typical for hardware switches that suffer from high packet drop counts. Furthermore, we evaluate the scalability of the system when building a larger topology of interconnected software switches, highlighting aspects such as management, port density, load balancing, and failover. In this context, we discuss the usability of software-defined networking (SDN) technologies, Open vSwitch Database (OVSDB) and OpenFlow, to centrally manage and optimize a data acquisition network. We build an IP-only leaf-spine network consisting of eight software switches running on separate physical servers as a demonstrator.

We intend to show in this presentation that building a high bandwidth lossless network based on software switches dedicated for data acquisition is feasible and can be considered as a viable solution for future small and large-scale systems based on commodity TCP/IP and Ethernet.

Primary author: JERECZEK, Grzegorz (CERN)

Co-author: PANDURO VAZQUEZ, Jose Guillermo (Royal Holloway, University of London)

Presenter: JERECZEK, Grzegorz (CERN)

Session Classification: Emerging Technologies / Feedback

Track Classification: Emerging Technologies / Feedback on Experience