



Contribution ID: 144

Type: **Oral**

Event-driven RDMA network communication in the ATLAS DAQ system with NetIO

Thursday 7 November 2019 14:00 (15 minutes)

NetIO is a network communication library that enables distributed applications to exchange messages using high-level communication patterns such as publish/subscribe. NetIO is based on libfabric and supports various types of RDMA networks, for example, Infiniband, RoCE, or OmniPath. NetIO is currently being used in the data acquisition chain of the ATLAS experiment.

Major parts of NetIO were recently rewritten using a novel, event-driven approach. All actions are processed asynchronously by a single-threaded central event loop. The event loop is backed by the Linux epoll system. The event-driven design implies that software written with NetIO uses callbacks to react to events.

The motivation for the architectural modifications to NetIO was to improve processing efficiency. Initial benchmarks show that the updated NetIO implementation yields the same or higher throughput, while the CPU resource utilization is reduced by an order of magnitude. The cause for this efficiency gain is largely due to significantly reduced thread synchronization, that became obsolete in the event-driven approach.

The paper will show this architecture is very suitable for IO-heavy workloads that are typically found in DAQ systems of High-Energy Physics experiments. The event-driven architecture will be explained in detail and compared with the original NetIO. The challenges of writing event-driven code are identified. A performance study of the event-driven NetIO in comparison with the original implementation as well as other RDMA networking solutions like MPI will be given.

Consider for promotion

No

Author: SCHUMACHER, Jorn (CERN)

Presenter: SCHUMACHER, Jorn (CERN)

Session Classification: Track 5 –Software Development

Track Classification: Track 5 –Software Development