

NetGPU

Tuesday, August 25, 2020 7:45 AM (45 minutes)

This introduces a working proof-of-concept alternative to RDMA, implementing a zero-copy DMA transfer between the NIC and GPU, while still performing the protocol processing on the host CPU. A normal NIC/host memory implementation is also presented.

By offloading most of the data transfer from the CPU, while not needing to reimplement the protocol stack, this should provide a balance between high performance and feature flexibility.

This presentation would cover the changes needed across the kernel; mm support, networking queues, skb handling, protocol delivery, and a proposed interface for zero-copy RX of data which is not directly accessible by the host CPU. It would also solicit input for further API design ideas in this area.

A paper is planned. This proposal was originally submitted for the main track and was recommended for the networking track instead.

I agree to abide by the anti-harassment policy

I agree

Primary author: LEMON, Jonathan (Facebook)

Presenter: LEMON, Jonathan (Facebook)

Session Classification: Networking and BPF Summit

Track Classification: Networking & BPF Summit