# Overcoming the Bandwidth Limitations of PCI Express

## Introduction

It is estimated that the global Internet traffic was 96 ExaBytes per month in 2016, and the expectation for 2021 is no less than 278 ExaBytes per month. These hard to imagine numbers correlate with the increasing speeds of individual connections. Current 100 Gbps Ethernet will soon get a four times faster brother. As one would expect, the demand for bandwidth drives the advances of the whole industry. CPUs, the brains of our connected world, are getting more processing cores and more RAM channels running at higher frequencies. And as usual, there are some bottlenecks in the ecosystem.

## Details

Current most powerful Xeon Platinum 8180 CPU is a 28-core beast running at 2.5 GHz base frequency. In theory it can execute 70 billion instructions per seconds, if one instruction is issued at every core in each clock cycle. 100 Gbps Ethernet transfers up to 149 million packets per second, so there should be at least 470 clock cycles to process a packet. That sounds reasonable. That CPU also has six DDR4-2666 channels, which results in 1Tbps of throughput, at least in theory. Again, that is well enough for processing several 100 Gbps lines.

The problem: PCI Express. Current CPUs still use third generation of this interconnection system. After very detailed manual fine tuning of PCIe transfers, Netcope Technologies have shown that it is possible to transfer slightly over 110 Gbps of network packet data. But that's about it. Theoretical bandwidth of 16-line PCIe gen3 interface is 128 Gbps. Many applications (such as full packet capture of both directions of a full duplex Ethernet link) are simply not possible.

While using multiple Network Interface Cards (NICs) to deliver packets to the CPU is possible (the CPUs have up to 48 PCIe lines), such setup has several disadvantages: server footprint size, power consumption, maximal achievable port density, management complexity, reliability, cost are obvious. Other issues will arise when more functionality is expected from the (Smart)NIC.

Resources needed for 200G solution - RoI calculation:

| **Card** | **Installation size** | **CAPEX** | **OPEX** |
|---|---|---|---|
| Full height 100G card | 2U | 2 * 10k for 100G card + 2* 5k for server<br>= 30k USD | 2 * 35 W for two 100G cards<br>2 * 150 W for two servers<br>= 370 W |
| Low-profile 200G card | 1U | 15k for card, 7,5k for server<br>= 22,5k USD | 50 W for one 200G card, 250 W for one high-end low profile server<br>= 300 W |
| Savings | 1U of rack space, 100 USD per month<br>= 1200 USD/yr | 7,5k USD | = 70 * 24 * 365<br>= 613,2 kWh * 10c per kWh = 61 USD/yr (not including cooling) |

* Lower power not only means lower electricity bill but also lower cooling requirements and further savings on electricity.

## Meet NFB-200G2QL

Netcope Technologies has recently introduced new member of the Netcope FPGA Boards family - the NFB-200G2QL. This low profile, passively cooled, fully programmable SmartNIC has two QSFP28 interfaces for two 100 Gbps Ethernet full duplex links. At first glance, it has single PCIe gen3 x16 slot, so its bandwidth to the host system is still limited. However, it also has a secondary expansion slot specifically designed to overcome the limitations of PCIe. NFB-200G2QL can be connected to two PCIe x16 slots at the same time.

The card is still young, but it already begins to flex muscles.

## Demonstration of 200 Gbps Traffic Capture in a Single Server

In our new video we show that the card is able to receive full 200 Gbps of Ethernet traffic. That is true even for the shortest allowed, 64 B Ethernet frames.

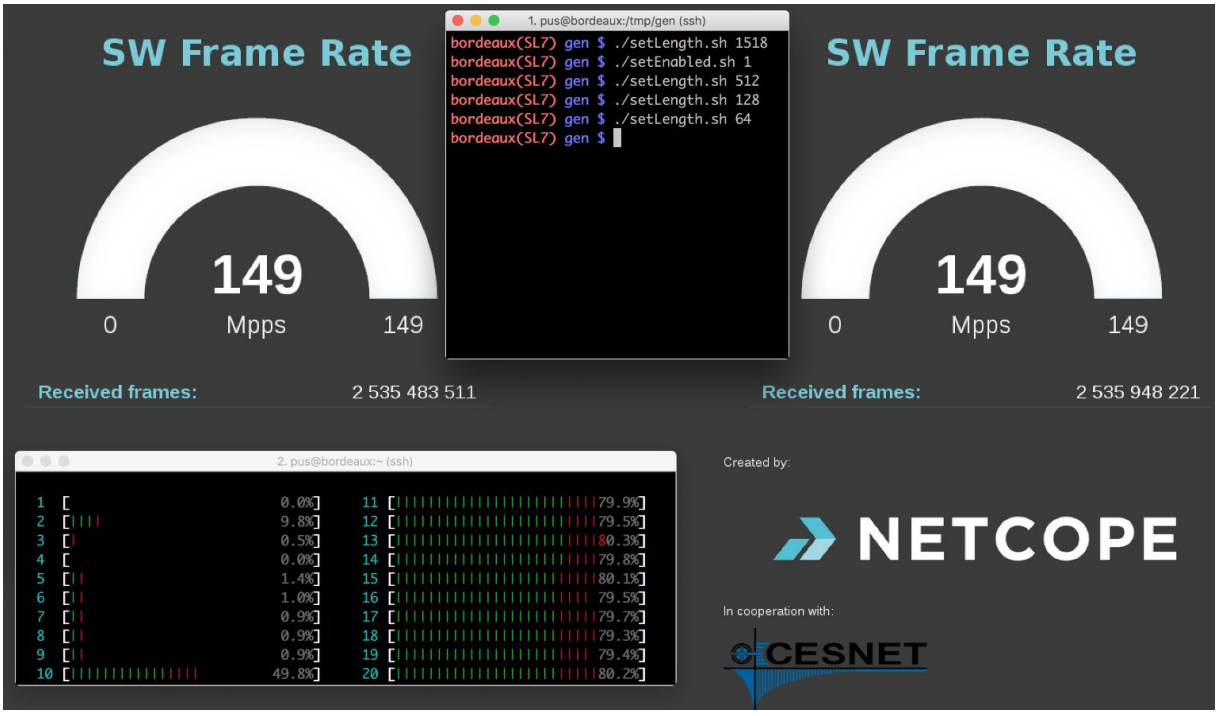


This result was achieved using CPU far slower (and cheaper) than the 28-core giant mentioned earlier. Two 10-core Xeon E5-2660 v3 CPUs were installed, but to make things more interesting, only one of them was actually used to receive data. Yes, that's ***200 Gbps packet capture using one low-profile card and one CPU***. Intel's Recommended Customer Price of this CPU is only $1445. Compare to $10009 for Xeon Platinum 8180. As a bonus, the whole setup fits into 1U rack server.

## Conclusion

Bandwidth requirements of modern high-speed networking systems are not easily met. NFB-200G2QL is the first to the market to offer lossless 200 Gbps packet capture in surprisingly small (low profile, half length) footprint. It is passively cooled and doesn't require external power supply, which makes it perfectly suitable even for large scale

deployment in data centers. In addition to that, the card will be soon supported by the Netcope's unique NP4 Firmware as a Service web portal for adding custom functions. That will add ultimate programmability to the ultimate bandwidth.