Scaling Network Processors and Virtualized Functions to 100 Gbps and Higher

FPGA programmable solutions offer many benefits when dealing with the demanding environment of 100 Gbps computer networks. Netcope Technologies developed Netcope Session Filter (NSF), a solution making traffic processing easier by offloading heavy flows (also known as mega flows in Open vSwitch) directly in hardware. It allows a way to scale the throughput of virtualized functions and solutions based on network processors by providing unprecedented performance.

Network processors are equipped with an ample number of functions, and their aggregate throughput makes them an ideal choice for deployment on high-speed links. However, complex tasks like pattern matching, encryption, or longest prefix match implemented within networks processors do not reflect the requirements of processing the traffic at wire speed. Performing such tasks depends on many factors including traffic composition, packet lengths, bursts, etc. In other words, their performance is not deterministic, putting a company's bottom line at risk when entering a service level agreement (SLA). In such situations, a reliable solution cannot miss any packet. These factors are decisive when comparing the quality of network appliances.

Although FPGA-based solutions provide a substantial advantage in the form of a completely programmable data plane with wire-speed throughput and deterministic processing, advanced development skills are required to build an effective application. Simple tasks like traffic filtering are implemented within an FPGA chip using techniques that guarantee the throughput and latency of data processing. Combining an FPGA-based solution and NPUs can leverage the advantages of both - performance and flexibility required to guarantee SLA.

How Netcope Session Filter Works

Offloading traffic helps network processors to deal with the deluge of data. Suspicious traffic can be processed in much more detail because less traffic burdening network processors means more computational power available. FPGA-based solutions of Netcope Technologies allow a way to dynamically control the level of traffic "zoom-in" based on the current load on network links, and NPU or CPU cores. In other words, they can be used to effectively implement a system which prevents NPUs/CPUs from overloading and produces more predictable results. In addition, the distribution over CPU cores is performed on the level of network flows, provides the possibility of different processing on dedicated NPU/CPU cores and utilizes cache memories effectively. As a result, deeper traffic analysis can be performed and the throughput of a solution can be scaled to hundreds of gigabits.

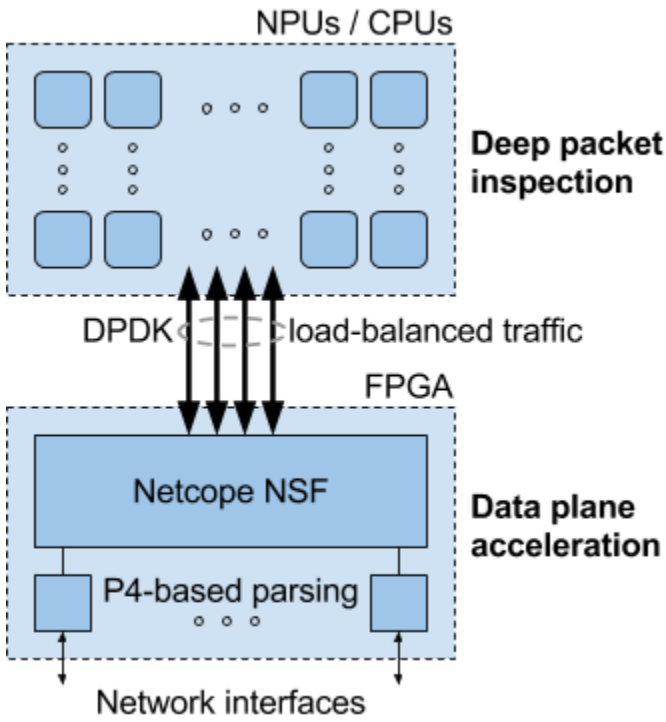

Figure 1: Data plane acceleration with Netcope Session Filter

Netcope Session Filter (NSF) is an FPGA-based solution with ready-to-use firmware and software libraries. Keeping the advantage of wire-speed traffic processing on all packet lengths and full 100 Gbps transfers to software running on NPU/CPU, it processes the traffic at the level of network flows, allowing a way to define a specific action for each network flow (filtering, cropping packets in length, distribution over DMA channels, transmission to an output network interface, gathering flow statistics directly in FPGA, etc.) Network flow definition is programmable, and by default is perceived as a five-tuple of source and destination IPv4/6 addresses, L4 protocol, and source and destination TCP/UDP ports, possibly accompanied with the index of the input network interface. The processing is driven by instructions delivered by software applications or network processors analyzing the traffic. Instructions are issued while the traffic is being analyzed and processing processed in the FPGA chip changes accordingly on the fly. With this approach, no a priori knowledge of the network traffic is necessary. E.g. a flow is filtered out independently of its IP addresses and TCP/UDP ports. This enables the easy detection of attacks coming from unknown IP addresses and non-standard TCP/UDP ports.

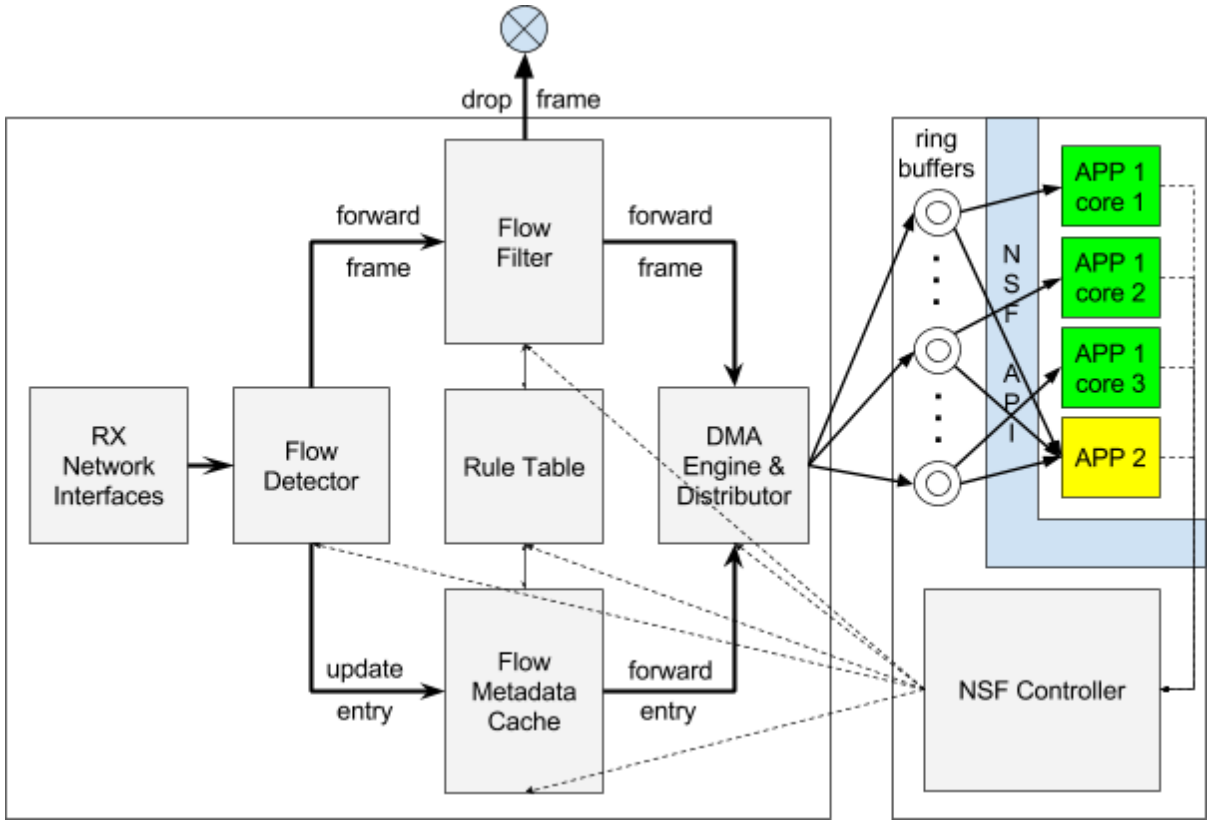

Figure 2: Schematics of NSF firmware and software

NSF in the Field: Practical Examples of DPI and NFV Acceleration

Let us consider Deep Packet Inspection (DPI), a task widely used in IDS and IPS systems like Snort, Suricata and Bro. Real-time entertainment (mostly video streaming services like YouTube and Netflix) is a type of traffic that is little to no threat, but can swallow up to 65% of the bandwidth. Processing it wastes the majority of the application's performance. Netcope Session Filter addresses this issue by filtering unnecessary flows based on the analysis of the first few packets of the flows. When the analysis reveals a flow is DPI-irrelevant, the flow is filtered out.

Another use case is tracking flow information. It is a straightforward task but on high-speed links it can turn out to be challenging. This is another situation where NSF helps. Its firmware features a flow cache able to track hundreds of thousands of flows and keep the number of transferred packets and bytes, initial and last nanosecond timestamp, etc. This information can be used in software to generate NetFlow or IPFIX data for export.

The capabilities of NSF are demonstrated on extracting information of interest related to HTTP traffic like URL domain (.com, .net, etc.), HTTP method (GET, POST, etc.), and HTTP status code. The figure below shows how NSF hardware offloading works. Even though NSF processes all traffic with no packet loss, the drop rate depends on the performance of the

software application. In this particular case, enabling hardware offload reduces the drop rate from 60% to 10%..

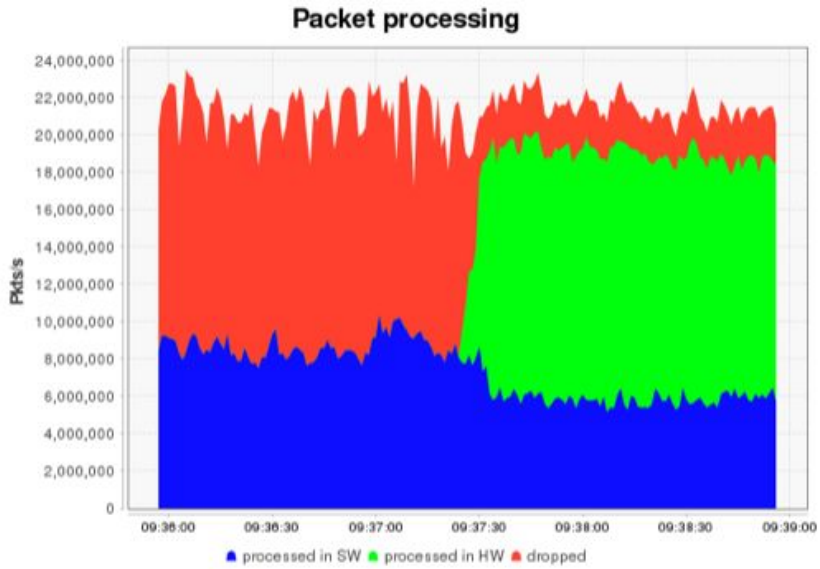

Figure 3: Reduction of the volume of unprocessed traffic (red) after enabling NSF hardware offloading (green) in the middle of the time frame. The rest is traffic processed in software (blue).

A similar situation emerges with Network Functions Virtualization (NFV). Virtualized Network Functions (VNF) are prone to be overloaded as their functionality is implemented purely in software. This can lead to severe SLA violations. It is now clear that hardware acceleration is a must if VNF are to be effectively deployed on the highest-speed links. NSF addresses this issue by offloading traffic of no interest in hardware precisely as described in the previous paragraphs.

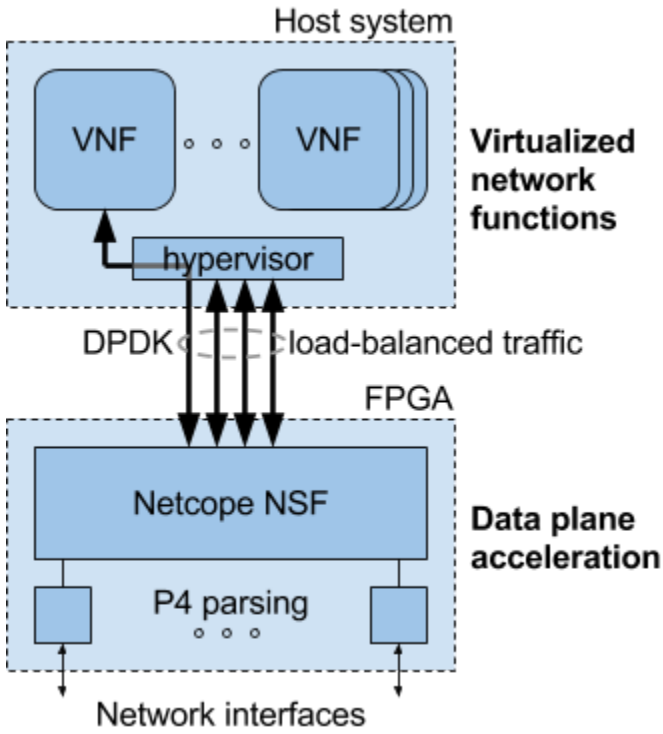

Figure 4: NSF deployment in NFV environment

Conclusion

To sum up, FPGA-based co-processing solves the problem of traffic overload in 100 Gigabit networks. In addition, programmable FPGAs excel in feature flexibility. Unlike traditional purpose-built co-processors, FPGAs enable the implementation of new tasks with changing requirements on traffic processing.

NSF helps network processors and VNFs to keep pace with the growing speed of network links. It reduces the amount of traffic reaching the processing applications and leaves more room for processing the traffic of interest by offloading heavy flows directly in hardware.

"Netcope's 100G-ready NSF allows for the early classification and hardware offload of sessions that are not relevant to Intrusion Protection using all remaining power to analyse the remaining traffic for Intrusion candidates. Netcope's Hardware Session-based Offloading coupled with inter-adapter load-balancing of full-duplex links, means that we can deliver 200 Gbps of DPI analyses using standard server configurations." says Martin Hayes, CTO of Picomass.