Flexible OVS acceleration with P4 and low profile FPGA card

Matej Vido, Pavel Benáček[1], Tomáš Martínek
{vido, benacek, martinek}@cesnet.cz
CESNET, a.l.e.
Zikova 4
160 00 Prague 6
Czech Republic

Viktor Puš, Petr Kaštovský, Michal Kekely
{pus, kastovsky,kekely}@netcope.com
Netcope Technologies, a.s.
Sochorova 3232
616 00 Brno
Czech Republic

Technical requirements:

large table; at least two power plugs; LCD screen with VGA connector; possibly a poster stand

Abstract:

Modern data center networking revolves around several widely accepted and supported concepts and standards. Server and network virtualization requires virtual packet switch to be placed in between the server physical network interface and the virtual machines' virtual network interfaces. Open vSwitch (OVS) is de facto standard open-source implementation of a virtual switch. Especially in communication-intensive workloads, the performance of virtual switch is crucial for the performance of the whole compute system. That's why the OVS is often used in conjunction with Data Plane Development Kit (DPDK), which provides fast packet receive/transmit API.
With the ever-increasing CPU speed and core count, as well as broadening adoption of 100 Gbps Ethernet, further improvements in OVS performance are highly desirable. Next step in this direction is FPGA-based acceleration. FPGA-based Network Interface Cards (NICs) certainly achieve high throughput, but their programming is often very low-level and cumbersome.
P4 language provides very good match for implementing hardware offload of the most timing-critical tasks of OVS. If most of the traffic can be decided and forwarded by the FPGA accelerator (based on the P4 description and runtime rules obtained from the OVS), then significant improvement in throughput can be achieved.

To bring this flexibility into FPGAs, we have created a compiler from P4 language to VHDL. The generated processing core in VHDL is then deployed on the FPGA-based card. In our case, we decided to use the novel NFB-2002QL card which is equipped with a powerful Xilinx UltraScale+ FPGA. Our design uses the P4-generated core for the preprocessing of incoming packets (parsing, flow classification). After that, preprocessed flows are sent via the DPDK PMD for further processing in the OVS user space datapath. The OVS is capable to offload entries from internal flow tables to the hardware accelerator using the rte_flow DPDK API. Therefore, the OVS can populate the generated P4 core with rules which allows the packet to skip the classification stage in software datapath. If the packet can't be classified in the firmware (due to the missing rule in the P4 core), it traverses through the whole software datapath without any acceleration. As the next step, the missing flow entry can be offloaded to the generated P4 core and then the next packets of the same flow are preprocessed in hardware.

The demo will show the usage of P4 language for the acceleration of OVS infrastructure. The whole design consists of the standard OVS environment and processing engine generated from the P4 source code. Firstly, we visualize monitored performance metrics (e.g., the number of processed packets) without any hardware acceleration. Secondly, we enable the hardware processing in generated P4 engine which will lead to the change of monitored performance metrics.

[1] Presenting author

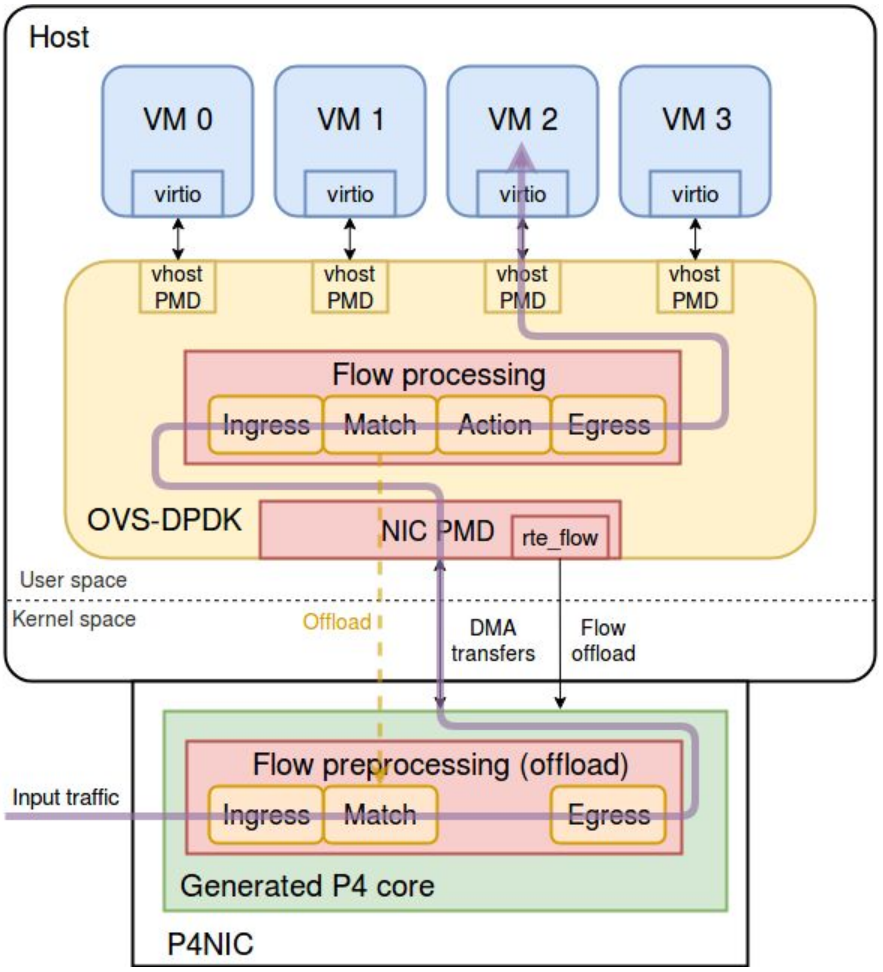

Figure 1: Architecture of the accelerated OVS infrastructure.