

Segment Routing at 100G Using FPGA Smart NIC and P4

Viktor Puš, Petr Kaštovský¹, Michal Kekely
{pus, kastovsky, kekely}@netcope.com
Netcope Technologies, a.s.
Sochorova 3232
616 00 Brno
Czech Republic

Pavel Benáček
benacek@cesnet.cz
CESNET, a.l.e.
Žikova 4
160 00 Prague 6
Czech Republic

Abstract:

Service Function Chaining (SFC) is a process of passing network traffic among individual typically virtualized network functions in NFV and SDN infrastructures. The typical SFC can be a series of IDS/IPS, firewalls, WAN optimizers and load balancers that the traffic needs to go through on its way from client to server and vice versa.

One approach to support SFC in network function virtualization infrastructure (NFVI) is to use Segment routing, in particular, the IPv6-based Segment routing (SRv6). This technology is now becoming very attractive and deployed in large networks of the future as demonstrated by SoftBank's recent joint announcement with Cisco as well as other deployments. This raised our attention first during P4.org workshop in May 2017 where teams from Bell Canada, Cisco Systems and Barefoot Networks presented the concept of The Extensible Network - Evolution in Protocol and Data Plane Agility and explained the benefits of SRv6 for SFC.

In our demo we focus on demonstrating the ability to quickly develop SRv6 acceleration using an FPGA-based hardware accelerator and P4 programming language. Similarly to accelerating SRv6, other applications can be accelerated using this approach. Good candidates being processing nodes of Vector Packet Processing (VPP).

In order to perform segment routing, SRv6 router goes through the Segment List in Segment Routing Header (SRH) and uses Segments Left field as an index of the active segment that is copied over to Destination Address of IPv6 header. We accelerate this data plane operations to demonstrate the productivity and flexibility of P4 language combined with FPGAs. We use Netcope P4 Cloud compiler service to generate FPGA firmware bitstream from the P4 description. The compiler was initially built for Xilinx Virtex-7 based Netcope board, but now the cloud service also supports (or will support in near future) newer UltraScale+ and Arria 10 based boards from Netcope and other hardware vendors. In case of this demo, we decided to use the NFB-2002QL card which has two 100 Gbps Ethernet ports and is equipped with a powerful Xilinx UltraScale+ VU7P FPGA. After compiling the firmware, it is tested for both correctness and performance. Wireshark is used to compare packets before and after (see Fig. 1). As for performance, simple command-line script shows throughput of almost 100 Gbps (Fig. 2).

Technical requirements:

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

¹ Presenting author

```

> Frame 3: 357 bytes on wire (2856 bits), 357 bytes captured (2856 bits) on interface 0
> Ethernet II, Src: 00:00:00_00:00:00 (00:00:00:00:00:00), Dst: Broadcast
Internet Protocol Version 6, Src: ::1, Dst: fe80::9618:82ff:fe6f:38eb
  0110 .... = Version: 6
  .... 0000 0000 .... = Traffic Class: 0x00
  .... 0000 0000 0000 0000 0000 0000 = Flow Label: 0x00000000
  Payload Length: 303
  Next Header: Routing Header for IPv6 (43)
  Hop Limit: 64
  Source: ::1
  Destination: fe80::9618:82ff:fe6f:38eb
  [Destination SA MAC: HewlettP_6f:38:eb (94:18:82:6f:38:eb)]
  [Source GeoIP: Unknown]
  [Destination GeoIP: Unknown]
  Routing Header for IPv6 (Source Route)
    Next Header: TCP (6)
    Length: 32
    [Length: 264 bytes]
    Type: Source Route (0)
    Segments Left: 8
    Reserved: 00000000
    Address[1]: 2001:db8:dead::1
    Address[2]: 2001:db8:dead::2
    Address[3]: 2001:db8:dead::3
    Address[4]: 2001:db8:dead::4
    Address[5]: 2001:db8:dead::5
    Address[6]: 2001:db8:dead::6
    Address[7]: 2001:db8:dead::7
    Address[8]: 2001:db8:dead::8
    Address[9]: 2001:db8:dead::9
    Address[10]: 2001:db8:dead::10
    Address[11]: 2001:db8:dead::11
    Address[12]: 2001:db8:dead::12
    Address[13]: 2001:db8:dead::13
    Address[14]: 2001:db8:dead::14
    Address[15]: 2001:db8:dead::15
    Address[16]: 2001:db8:dead::16
  Transmission Control Protocol, Src Port: 20, Dst Port: 80, Seq: 36, Ack: 36

> Frame 3: 357 bytes on wire (2856 bits), 357 bytes captured (2856 bits) on interface 0
> Ethernet II, Src: 00:00:00_00:00:00 (00:00:00:00:00:00), Dst: Broadcast
Internet Protocol Version 6, Src: ::1, Dst: 2001:db8:dead::8
  0110 .... = Version: 6
  .... 0000 0000 .... = Traffic Class: 0x00 (DSCP: CS0)
  .... 0000 0000 0000 0000 0000 0000 = Flow Label: 0x00000000
  Payload Length: 303
  Next Header: Routing Header for IPv6 (43)
  Hop Limit: 64
  Source: ::1
  Destination: 2001:db8:dead::8
  [Source GeoIP: Unknown]
  [Destination GeoIP: Unknown]
  Routing Header for IPv6 (Source Route)
    Next Header: TCP (6)
    Length: 32
    [Length: 264 bytes]
    Type: Source Route (0)
    Segments Left: 7
    Reserved: 00000000
    Address[1]: 2001:db8:dead::1
    Address[2]: 2001:db8:dead::2
    Address[3]: 2001:db8:dead::3
    Address[4]: 2001:db8:dead::4
    Address[5]: 2001:db8:dead::5
    Address[6]: 2001:db8:dead::6
    Address[7]: 2001:db8:dead::7
    Address[8]: 2001:db8:dead::8
    Address[9]: 2001:db8:dead::9
    Address[10]: 2001:db8:dead::10
    Address[11]: 2001:db8:dead::11
    Address[12]: 2001:db8:dead::12
    Address[13]: 2001:db8:dead::13
    Address[14]: 2001:db8:dead::14
    Address[15]: 2001:db8:dead::15
    Address[16]: 2001:db8:dead::16
  Transmission Control Protocol, Src Port: 20, Dst Port: 80, Seq: 36, Ack: 36

```

Figure 1: Comparing packets before and after SRv6 node processing.

```

Every 1.0s: cat /bin/stats.txt

** P4 performance ****
RX INTERFACE
    Frames           : 881181391
    Bytes            : 312819393805
    Data rate        : 94.712 Gbps

TX INTERFACE
    Frames           : 881182465
    Bytes            : 312819775075
    Data rate        : 94.713 Gbps

```

Figure 2: SRv6 processing performance measurement.