



NVIDIA DOCA Simple Forward VNF

Reference Application

Table of Contents

Chapter 1. Introduction.....	1
Chapter 2. System Design.....	2
Chapter 3. Application Architecture.....	5
Chapter 4. Configuration Flow.....	6
Chapter 5. Running Application on BlueField.....	8
Chapter 6. Running Application on Host.....	10
Chapter 7. References.....	11

Chapter 1. Introduction

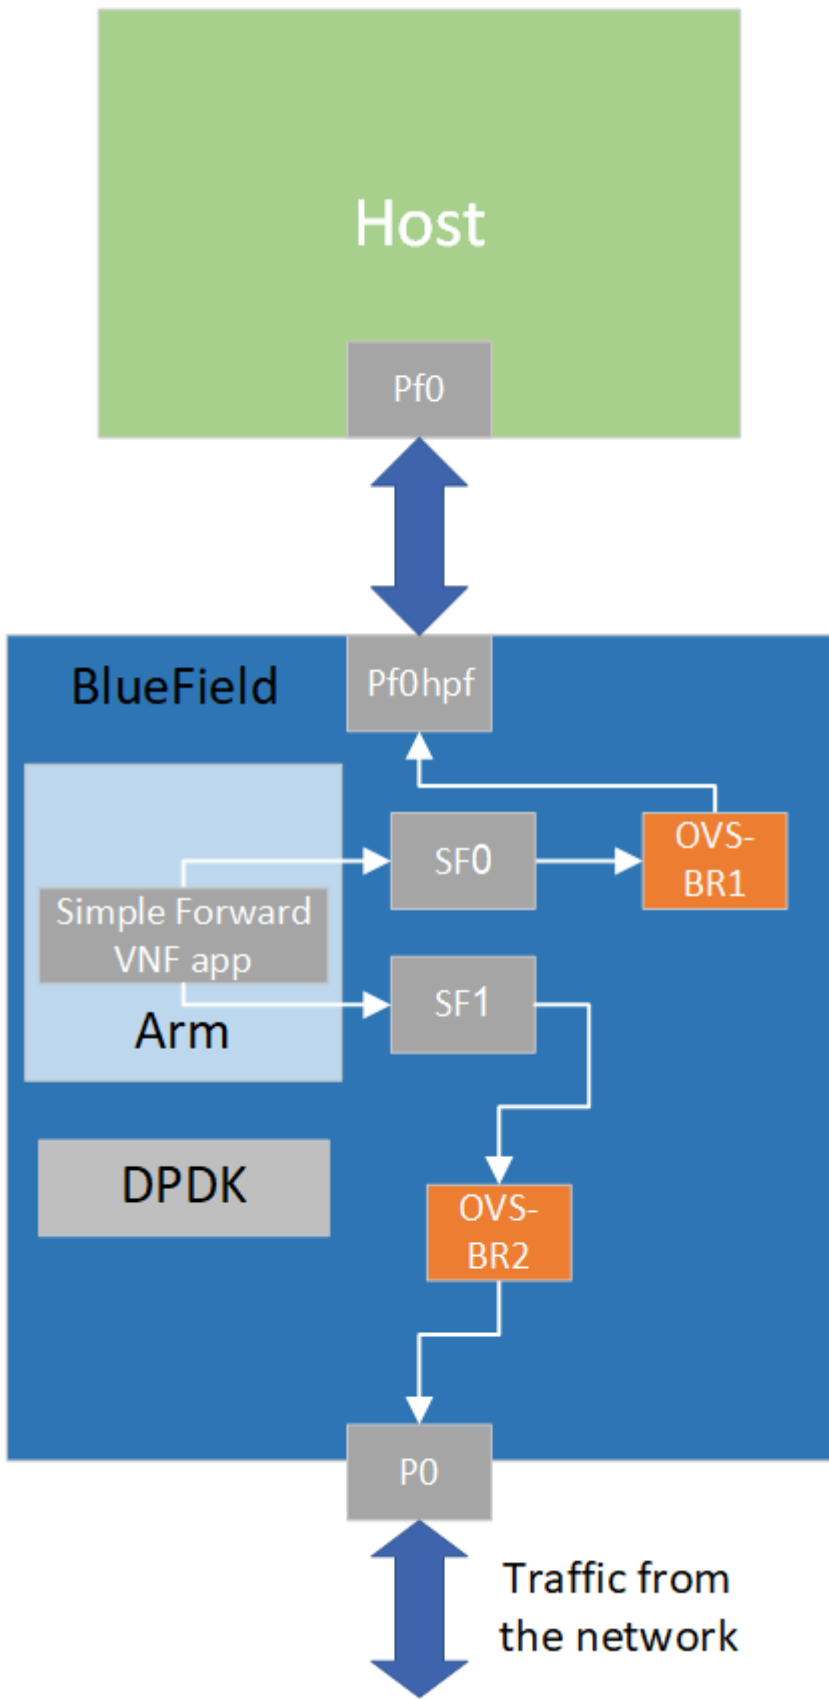
Simple forward is a forwarding application which takes VXLAN traffic from a single RX port and transmits it on a single TX port.

For a packet received on an RX port, simple forward will create a flow based on the packet's tunnel and 5-tuples. For the following packets with the same key, simple forward checks the packet's keys. If it finds that the packet matches the existing flow, then it does not create a new flow. Otherwise, a new flow is created. And then the packets are forwarded on the other port.

Simple forward should be run with dual ports. By using a traffic generator, the RX port receives the VXLAN packets and forwarding forwards them back to the traffic generator.

Chapter 2. System Design

The following diagram illustrates simple forward's packet flows. It receives traffic coming from the wire and passes it to the other port.



Chapter 3. Application Architecture

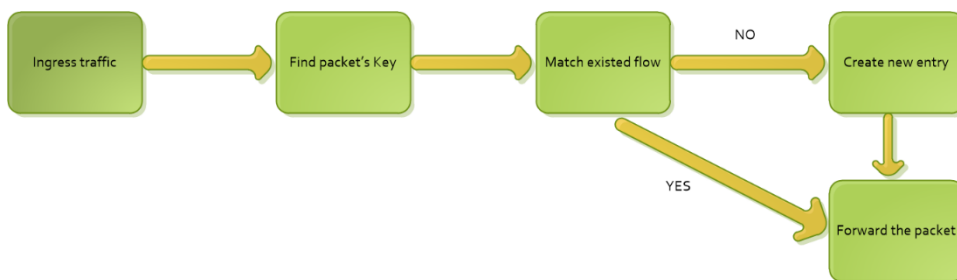
Simple forward initializes DPDK first, during which time simple forward handles the incoming packets.

The following diagram illustrates the initialization process.



1. `DPDK_init` – EAL init, parse argument from command line and register signal.
2. Start port – `mbuf_create`, `dev_configure`, rx/tx/hairpin queue setup and start the port.
3. `simple_fwd INIT` – create flow tables, build default forward pipes.

The following diagram illustrates how to process the packet.



1. Based on the packet's info, find the key values (e.g. src/dst ip, src/dst port, etc).
2. Traverse the inner flow tables, check if the keys exist or not.
 - ▶ If yes, update inner counter
 - ▶ If no, a new flow table is added and new pipes are configured on the DPU
3. Forward the packet to the other port.

Chapter 4. Configuration Flow

1. DPDK initialization.

```
dpdk_init(&argc, &argv, &nb_queues, &nb_ports)
```

2. Port initialization and start.

```
simple_fwd_start_dpdk_port(&port_cfg);
```

- ▶ Mbuf pool created using `rte_pktmbuf_pool_create`
- ▶ Driver initialization – use `rte_eth_dev_configure` to configure the number of queues
- ▶ Rx/Tx queue initialization – use `rte_eth_rx_queue_setup` and `rte_eth_tx_queue_setup` to initialize the queues
- ▶ Rx hairpin queue initialization – use `rte_eth_rx_hairpin_queue_setup` to initialize the queues
- ▶ Start the port using `rte_eth_dev_start`

3. Simple forward initialization.

```
simple_fwd_init((void *)&port_cfg);
```

- ▶ `simple_fwd_create` - create flow tables using `doca_ft_create`
- ▶ `simple_fwd_init_ports_and_pipes` - initialize DOCA port using `simple_fwd_init_doca_port` and build default pipes for ports using `build_fwd_pipe`

4. Main loop.

```
simple_fwd_process_packets(void *p);
```

- ▶ Receive packets using `rte_eth_rx_burst` in a loop
- ▶ Process packets using `simple_fwd_process_offload`
- ▶ Tx_burst the packets on the other port. Or mbuf free if `rx_only` is set to true.

5. Process packets.

```
simple_fwd_process_offload(mbufs[j]);
```

- ▶ Parse the packet's `rte_mbuf` using `simple_fwd_parse_packet`.
- ▶ Handle the packet using `simple_fwd_handle_packet`. If the packet's key does not match the existed the flow entry, create a new flow entry and PIPE using `simple_fwd_handle_new_flow`. Otherwise, increase the counter.

6. Simple forward destroy.

```
simple_fwd_destroy();
```

7. Simple forward close port.


```
simple_fwd_close_port(port_id)
```

Chapter 5. Running Application on BlueField

1. Please refer to the [DOCA Installation Guide](#) for details on how to install BlueField related software.
2. To build the application

a). The binary of simple forward example is located under `/opt/mellanox/doca/examples/simple_fwd_vnf/bin/doca_simple_fwd_vnf`. To re-build the simple forward sample, run the following:

```
cd /opt/mellanox/doca/examples/simple_fwd_vnf/src
meson /tmp/build
ninja -C /tmp/build
```

`doca_simple_fwd_vnf` will be created under `tmp/build`.

b). The build process depends on the `PKG_CONFIG_PATH` environment variable to locate the DPDK libraries. If the variable was accidentally corrupted, and the build fails, run the following command:

▶ For Ubuntu:

```
export PKG_CONFIG_PATH=$PKG_CONFIG_PATH:/opt/mellanox/dpdk/lib/aarch64-linux-gnu/pkgconfig
```

▶ For CentOS:

```
export PKG_CONFIG_PATH=$PKG_CONFIG_PATH:/opt/mellanox/dpdk/lib64/pkgconfig
```

c). The simple forward example is based on DPDK libraries. Therefore, the user is required to provide DPDK flags, and allocate huge pages. Run:

```
echo 1024 > /sys/kernel/mm/hugepages/hugepages-2048kB/nr_hugepages
sudo mkdir /mnt/huge
sudo mount -t hugetlbfs nodev /mnt/huge
```

3. To run the application:

```
doca_simple_fwd_vnf [dpdk_flags] -- [additional application flags]
```



Note: SFs must be enabled according to [Scalable Function Setup Guide](#).

For example:

```
/opt/mellanox/doca/examples/simple_fwd_vnf/bin/doca_simple_fwd_vnf -a  
auxiliary:mlx5_core.sf.4 -a auxiliary:mlx5_core.sf.5 -- --nr_queues=4 --  
stats_timer=2
```



Note: The flag `-a auxiliary:mlx5_core.sf.4 -a auxiliary:mlx5_core.sf.5` is mandatory for proper usage of the application. Modifying this flag will result unexpected behavior as only 2 ports are supported. The SF number is arbitrary and configurable. For additional information on available flags for DPDK, use `-h` before the `--` separator.

The command `quit` terminates the application.

Supported EAL options:

- ▶ `--log_level` - set log level
- ▶ `--stats_timer` - set interval to dump stats information
- ▶ `--nr_queues` - set queues number
- ▶ `--rx_only` - set `rx_only` (0 or 1)
- ▶ `--hw_offload` - set HW offload (0 or 1)
- ▶ `--hairpinq` - set forwarding to hairpin queue

Chapter 6. Running Application on Host

Please refer to [Running Reference Applications Over Host Guide](#).

Chapter 7. References

- ▶ `/opt/mellanox/doca/examples/simple_fwd_vnf/src/simple_fwd.c`

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation nor any of its direct or indirect subsidiaries and affiliates (collectively: "NVIDIA") make no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assume no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, and Mellanox are trademarks and/or registered trademarks of Mellanox Technologies Ltd. and/or NVIDIA Corporation in the U.S. and in other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2021 NVIDIA Corporation & affiliates. All rights reserved.