OVS L4 Firewall

Reference Guide

MLNX-15-060471 _v1.2 | January 2022

Table of Contents

Chapter 1. Introduction

L4 Open vSwitch (OVS) firewall is used to perform basic Access Deny List (ACLs) operations. It allows to identify different flows based on L3/L4 headers and execute different actions.

One of the ways to implement OVS L4 firewall is to use the connection tracking module as part of OVS. Connection tracking refers to keeping a record of all currently open connections (stateful inspection).

OVS is an open-source implementation of a virtual switch software layer which resides in a server and supports different switching capabilities. For more information on OVS, please refer to the official documentation for Open vSwitch.

Chapter 2. System Design

The following diagram illustrates packet flow on an NVIDIA® BlueField® DPU connect to the host.

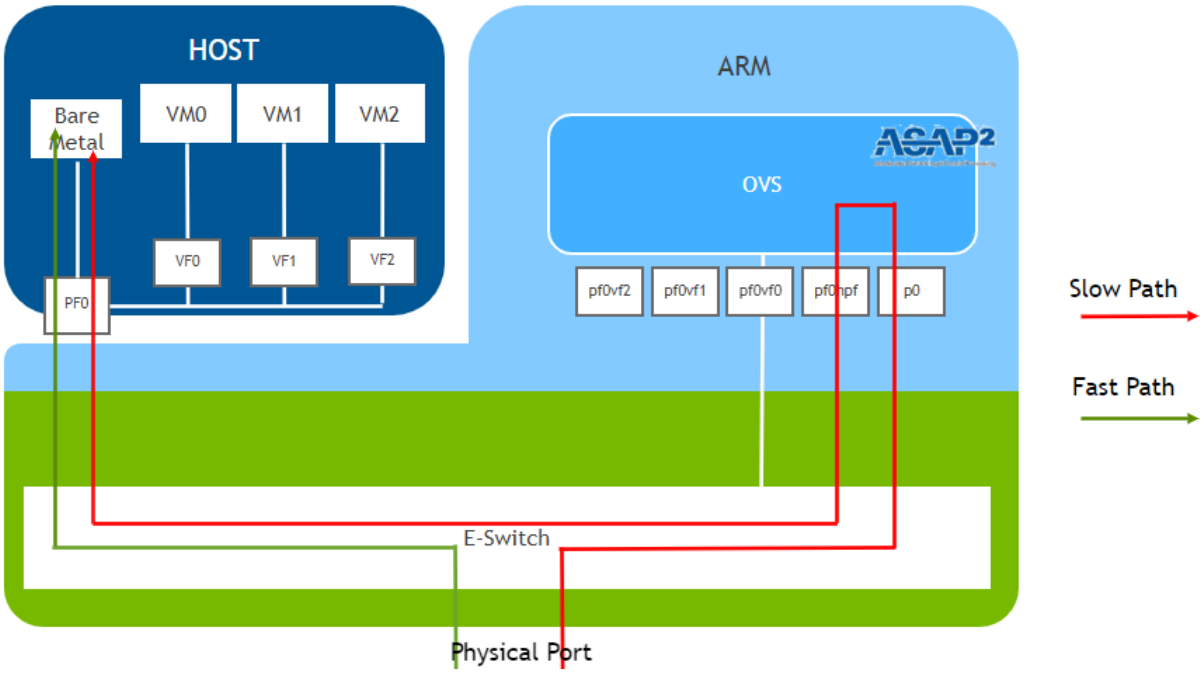

Packet flow steps:

1. Packet is received from physical port on the DPU. If the packet has no match in the eswitch flow table, it is sent to the Arm's `p0`.
2. According to OVS, the packet is directed to its destination. If OVS has no rule for where to send the packet, it is sent to pf0hpf (host representor).
3. `pf0hpf` is associated with `pf0` on the host. The packet is sent from `pf0hpf` on the Arm to `pf0` on the host.
4. The host processes the packet and responds with a packet to the Arm. OVS learns the packet and adds a rule into OVS table. Now ASAP2 adds the same rule to the e-switch.
5. When the next packet from the same flow is sent to the DPU through the physical port, it hits the e-switch flow table and is then passed to its destination.

Chapter 3. Application Architecture

The following diagram illustrates a packet's flow in different OVS tables.

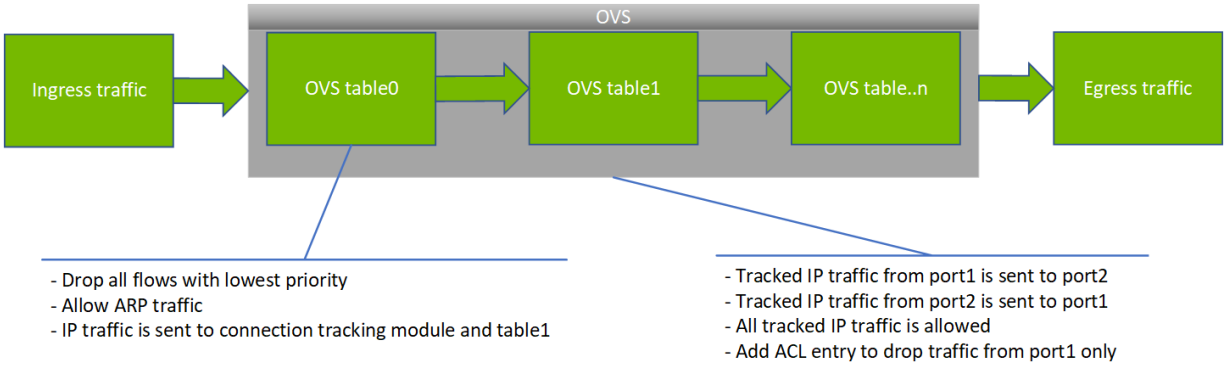

Chapter 4. Configuration Flow

1. Add `table0` entry with priority 1 and action drop.

   The lowest priority entry to drop all flows if no other match
2. Add `table0` entry with priority 10 for ARP traffic with action `normal(=forward)`.

   All ARP traffic will be forwarded.
3. Add `table0` entry with priority 100 for IP traffic with action `ct(=connection tracking)` and forward to OVS `table1`.

   All IP traffic will be set for connection tracking and sent to OVS `table1`.
4. Add `table1` entry for IP and tracked traffic from `port1` with action `port2`.

   All tracked IP traffic from `port1` will be sent to `port2`.
5. Add `table1` entry for IP and tracked traffic from `port2` with action `port1`.

   All tracked IP traffic from `port2` will be sent to `table1`.

   At this point, all IP traffic is tracked by OVS connection tracking module.
6. Add access deny list (ACL) entry to `table1` to drop all tracked IP traffic from `port1`.

   Traffic from `port1` should be blocked, while traffic from `port2` remains allowed.

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 run the application:

   a). Configure the OVS switch, ports, and enable OVS.

   ```
   ovs-vsctl del-br ovsbr1
   ovs-vsctl del-br ovsbr2
   ovs-vsctl add-br ovsbr1
   ovs-vsctl add-port ovsbr1 p0;
   ovs-vsctl add-port ovsbr1 pf0hpf;
   ovs-vsctl set Open_vSwitch . other_config:hw-offload=true;
   systemctl restart openvswitch
   systemctl enable openvswitch
   ```

   b). Configure OVS L4 Firewall rules as detailed in section Configuration Flow.

   ```
   ovs-ofctl add-flow ovsbr1 table=0,priority=1,action=drop
   ovs-ofctl add-flow ovsbr1 table=0,priority=10,arp,action=normal
   ovs-ofctl add-flow ovsbr1 "table=0,priority=100,ip,ct_state=-
   trk,actions=ct(table=1)"
   ovs-ofctl add-flow ovsbr1 "table=1,in_port=1,ip,ct_state=
   +trk,action=ct(commit),2"
   ovs-ofctl add-flow ovsbr1 "table=1,in_port=2,ip,ct_state=
   +trk,action=ct(commit),1"
   ovs-ofctl add-flow ovsbr1 "table=1,in_port=1,ip,ct_state=+trk,action=drop"
   ```

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 NONINFRINGEMENT, 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. The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a world¬wide basis. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

NVIDIA Corporation | 2788 San Tomas Expressway, Santa Clara, CA 95051
http://www.nvidia.com