



NVIDIA BlueField DPU Modes of Operation

User Guide

Table of Contents

Chapter 1. Introduction.....	1
Chapter 2. Embedded CPU Function Ownership Mode.....	2
2.1. Configuring ECPF Mode from Separated Host Mode.....	3
Chapter 3. Restricted DPU Host Mode.....	4
3.1. Enabling Host Restriction.....	4
3.2. Disabling Host Restriction.....	4
3.3. Hiding Network Ports from Host.....	5
Chapter 4. Separated Host Mode.....	6
4.1. Configuring Separated Host Mode from ECPF Mode.....	7

Chapter 1. Introduction

The NVIDIA® BlueField® DPU has several modes of operation:

- ▶ [Embedded function \(ECPF\)](#) (ECPF) ownership where the embedded Arm system controls the NIC resources and data path (**default**)
- ▶ [Restricted mode](#) which is an extension of the ECPF ownership with additional restrictions on the host side
- ▶ [Separated host mode](#) (symmetric model)

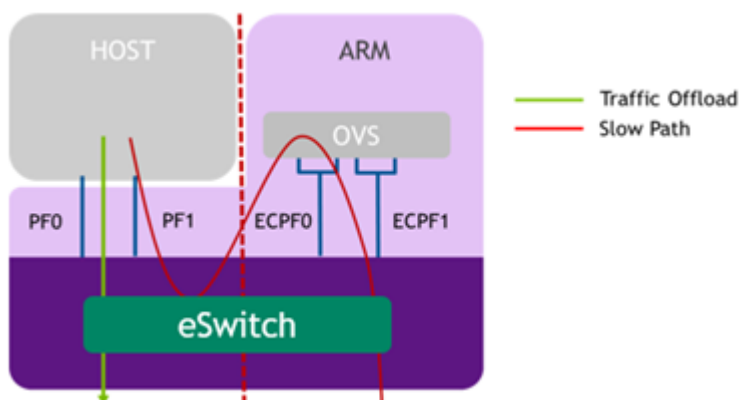
Chapter 2. Embedded CPU Function Ownership Mode

This mode, also known as ECPF or DPU mode, is the default mode for BlueField DPU.

In ECPF mode, the NIC resources and functionality are owned and controlled by the embedded Arm subsystem. All network communication to the host flows through a virtual switch control plane hosted on the Arm cores, and only then proceeds to the x86 host. While working in this mode, the DPU is the trusted function managed by the data center and host administrator—to load network drivers, reset an interface, bring an interface up and down, update the firmware, and change the mode of operation on the DPU device.

A network function is still exposed to the host, but it has limited privileges. In particular:

1. The driver on the host side can only be loaded after the driver on the embedded side has loaded and completed NIC configuration.
2. All ICM (Interface Configuration Memory) is allocated by the ECPF and resides in the embedded host memory.
3. The ECPF controls and configures the NIC embedded switch which means that traffic to and from the host interface always lands on the Arm side.



When the server and DPU are initiated, the networking to the host is blocked until the virtual switch on the DPU is loaded. Once it is loaded, traffic to the host is allowed by default.

There are two ways to pass traffic to the host interface: Either using representors to forward traffic to the host (every packet to/from the host would be handled also by the network

interface on the embedded Arm side), or push rules to the embedded switch which allows and offloads this traffic.

2.1. Configuring ECPF Mode from Separated Host Mode

To enable this mode:

1. Start MST driver set service:

```
$ mst start
```

2. Identify the MST device:

```
$ mst status -v
```

Output example:

```
...
DEVICE_TYPE          MST          PCI          RDMA
NET                  NUMA
BlueField(rev:0)     /dev/mst/<device>.1  37:00.1      mlx5_1
net-ens1f1           0
BlueField(rev:0)     /dev/mst/<device>    37:00.0      mlx5_0
net-ens1f0           0
```

3. Run the following commands on the Arm:

```
$ mlxconfig -d /dev/mst/<device> s INTERNAL_CPU_MODEL=1
```

4. Power cycle the server.



Note: If OVS bridges `ovsbr1` and `ovsbr2` are not created (you may verify that using command `ovs-vsctl show`), make sure field `CREATE_OVS_BRIDGES` is set to "yes" in `/etc/mellanox/mlnx-ovs.conf`.

Chapter 3. Restricted DPU Host Mode

Restricted mode is a specialization of Embedded mode and implements an additional layer of security where the host system administrator is prevented from accessing the DPU from the host. Once Restricted mode is enabled, the data center administrator should control the DPU entirely through the Arm cores and/or BMC connection instead of through the host.

For security and isolation purposes, it is possible to restrict the host from performing operations that can compromise the DPU. The following operations can be restricted individually when changing the DPU host to Restricted mode:

- ▶ Port ownership – the host cannot assign itself as port owner
- ▶ Hardware counters – the host does not have access to hardware counters
- ▶ Tracer functionality is blocked
- ▶ RShim interface is blocked
- ▶ FW flash is restricted

3.1. Enabling Host Restriction

To enable host restriction:

1. Start the MST service.

```
$ mst start
```

2. Set restricted mode. From the Arm side, run:

```
$ mlxprivhost -d /dev/mst/<device> r --disable_rshim --disable_tracer --  
disable_counter_rd --disable_port_owner
```



Note: If RShim is disabled, power cycle is required.

3.2. Disabling Host Restriction

To disable host restriction set the mode to privileged mode:

```
$ mlxprivhost -d /dev/mst/<device> p
```

The configuration takes effect immediately.



Note: If reverting from "rshim-disabled" mode, system power cycle is required.

3.3. Hiding Network Ports from Host

1. Make sure the DPU is operating in embedded mode. Run:

```
# mlxconfig -d /dev/mst/<device> q | grep -i cpu
INTERNAL_CPU_MODEL EMBEDDED_CPU(1)
```

If the output is `INTERNAL_CPU_MODEL SEPERATED_HOST(0)`, then it means your device is operating in separated mode. See section [Configuring ECPF Mode from Separated Host Mode](#) to change this configuration.

2. To hide network ports from host, run:

```
mlxconfig -d /dev/mst/<device> s NUM_OF_PF=0
```



Note: Power cycle is required for this configuration to take effect.

3. To re-expose network ports to host, run:

```
mlxconfig -d /dev/mst/<device> s NUM_OF_PF=2
```



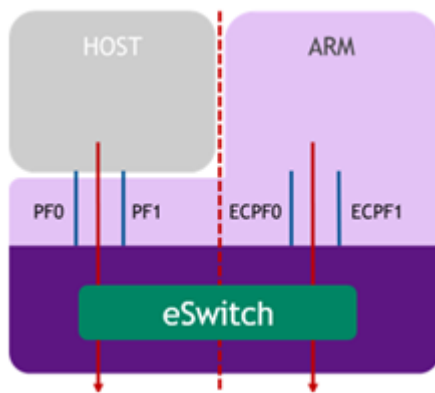
Note: Power cycle is required for this configuration to take effect.

Chapter 4. Separated Host Mode

In Separated mode, a network function is assigned to both the Arm cores and the x86 host cores. The ports/functions are symmetric in the sense that traffic is sent to both physical functions simultaneously. Each one of those functions has its own MAC address, which allows one to communicate with the other, and can send and receive Ethernet and RDMA over Converged Ethernet (RoCE) traffic. There is an equal bandwidth share between the two functions.

There is no dependency between the two functions. They can operate simultaneously or separately. The host can communicate with the embedded function as two separate hosts, each with its own MAC and IP addresses (configured as a standard interface).

In Separated mode, the host administrator is a trusted actor who can perform all configuration and management actions related to either network function.



This mode enables the same operational model of a SmartNIC (that does not have a separated control plane). In this case, the Arm control plane can be used for different functions but does not have any control on the host steering functions.

The limitations of this mode are as follows:

- ▶ Switchdev (virtual switch offload) mode is not supported on either of the functions
- ▶ SR-IOV is only supported on the host side

4.1. Configuring Separated Host Mode from ECPF Mode

On the server host, follow these steps:

1. Enable separated host mode. Run:

```
$ mst start  
$ mlxconfig -d /dev/mst/<device> s INTERNAL_CPU_MODEL=0
```

2. Power cycle.

3. Verify configuration. Run:

```
$ mst start  
$ mlxconfig -d /dev/mst/<device> q | grep -i model
```

4. Remove OVS bridges configuration from the Arm-side. Run:

```
$ ovs-vsctl list-br | xargs -r -l
```

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. 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

© 2022 NVIDIA Corporation & affiliates. All rights reserved.