



NVIDIA DOCA DPU CLI

Guide

Table of Contents

Chapter 1. Introduction.....	1
Chapter 2. General Commands.....	2
Chapter 3. DPU/DOCA Commands.....	4

Chapter 1. Introduction

This guide provides a concise guide on useful commands for DOCA deployment and configuration.


This tables in this guide provide two categories of commands:


- ▶ General commands for Linux/networking environment
- ▶ DOCA/DPU-specific commands



Note: For more information about these commands, such as usage instructions, flag options, arguments and so on, use the `-h` option after the command or use the manual (e.g., `man lspci`).

Chapter 2. General Commands

Command	Description
<code>ifconfig</code>	<p>Used to configure kernel-resident network interfaces. It is used at boot time to set up interfaces as necessary. After that, it is usually only needed when debugging or when system tuning is needed.</p> <p>If no arguments are given, <code>ifconfig</code> displays the status of the currently active interfaces. If a single interface argument is given, it displays the status of the given interface only. If a single <code>-a</code> argument is given, it displays the status of all interfaces, even those that are down. Otherwise, it configures an interface.</p>
<code>ethtool <devname></code>	<p>Used to query and control network device driver and hardware settings, particularly for wired Ethernet devices.</p> <p><devname> is the name of the network device on which <code>ethtool</code> should operate.</p> <div style="background-color: #f0f0f0; padding: 5px;"> Note: This command shows the speed of the network card of the DPU.</div>
<code>lspci</code>	<p>Displays information about PCIe buses in the system and devices connected to them. By default, it shows a brief list of devices.</p>
<code>tcpdump</code>	<p>Dump traffic on a network. Usage: <code>tcpdump -i <interface></code> where <interface> is any port interface (physical/SF rep/VF port rep).</p>
<code>ovs-vsctl</code>	<p>Utility for querying and configuring <code>ovs-vswitchd</code>. The <code>ovs-vsctl</code> program supports the model of a bridge implemented by Open vSwitch in which a single bridge supports ports on multiple VLANs.</p>

Command	Description
<code>mount 10.0.0.10:/vol/myshare/ myshare/</code>	Used for mounting a work directory on the DPU.  Note: Must be used after creating a new directory named <code>myshare</code> under root (i.e., <code>mkdir /myshare</code>)
<code>scp</code>	Secure copy (remote file copy program). Useful for copying files from BlueField to the host and vice versa.
<code>iperf</code>	Used for server-client connection. Useful to check if the network connection achieves the speed of the network card on the DPU (line rate).

Chapter 3. DPU/DOCA Commands

Command	Description
<code>ibdev2netdev</code>	Displays available mlnx interfaces
<code>mst</code>	Used to start MST service, to stop it, and for other operations with NVIDIA devices like reset and enabling remote access
<code>cat /etc/mlnx-release</code>	Displays the full BlueField OS (bfb) version
<code>cat /etc/os-release</code>	Displays the details of the underlying OS installed on BlueField
<code>ibv_devinfo</code>	Displays the current InfiniBand connected devices and relevant information. Useful for checking current firmware version.
<code>ipmitool power cycle</code>	Power cycle
<code>echo 1024 > /sys/kernel/mm/hugepages/hugepages-2048kB/nr_hugepages</code>	DPDK setup. Allocates hugepages for DPDK environment abstraction layer (EAL).
<code>mlxdevm tool</code>	The <code>mlxdevm</code> tool is found under <code>/opt/mellanox/iproute2/sbin/</code> . With this tool it is possible to create an SF and set its state to active, configure a HW address and set it to trusted, deploy the created SF and print info about it.
<code>/opt/mellanox/iproute2/sbin/mlxdevm port add pci/<pci_address> flavour pcisf pfnum <correspondig_physical_function_number> sfnun <unique_sf_number></code>	Creates an SF in the flavor of the given PF with the given unique SF number. Example: <code>/opt/mellanox/iproute2/sbin/mlxdevm port add pci/0000:03:00.0 flavour pcisf pfnum 0 sfnun 4</code>
<code>/opt/mellanox/iproute2/sbin/mlxdevm port show</code>	Displays information about the available SFs
<code>/opt/mellanox/iproute2/sbin/mlxdevm port function set pci/0000:03:00.0/<sf_index> hw_addr <HW_address> trust on state active</code>	Configures SF capabilities such as setting the HW address, making it "trusted", and setting its state to active. <code><sf_index></code> the SF. To obtain this index, you may run <code>mlxdevm port show</code> . Example: <code>/opt/mellanox/iproute2/sbin/mlxdevm port function set pci/0000:03:00.0/229377 hw_addr 02:25:f2:8d:a2:4c trust on state active</code>

Command	Description
<pre>\$ echo mlx5_core.sf.<next_serial> > /sys/bus/auxiliary/drivers/ mlx5_core.sf_cfg/unbind \$ echo mlx5_core.sf. <next_serial> > / sys/bus/auxiliary/drivers/mlx5_core.sf/ bind</pre>	<p>These two commands deploy the created SF. The first command unbinds the SF from the default driver, while the second command binds the SF to the actual driver. The deployment phase should be done after the capabilities of the SF are configured. The SF is identified by <code><next_serial></code> which can be obtained by running the command below.</p>
<pre>ls /sys/bus/auxiliary/devices/ mlx5_core.sf.*</pre>	<p>Displays additional information about the created SFs and their "next serial numbers".</p> <p>For example, if <code>mlx5_core.sf.2</code> exists in the output of the command, then running <code>cat /sys/bus/auxiliary/devices/mlx5_core.sf.2/sfnum</code> would output the sfnum that is related to <code>mlx5_core.sf.2</code>.</p>
<pre>/opt/mellanox/iproute2/sbin/mlxdevm port function set pci/<pci_address>/ <sf_index> state inactive /opt/mellanox/iproute2/sbin/mlxdevm port del pci/<pci_address>/<sf_index></pre>	<p>These two commands must be executed to delete a given SF. First, users must set the state of the SF to inactive, and only then should it be deleted.</p>
<pre>/opt/mellanox/iproute2/sbin/mlxdevm port help</pre>	<p>Displays additional information about operations that can be used on created SF ports</p>
<pre>crictl pods</pre>	<p>Displays currently active K8S pods, and their IDs (it might take up to 20-30 seconds for the pod to start)</p>
<pre>crictl ps</pre>	<p>Displays currently active containers and their IDs</p>
<pre>crictl ps -a</pre>	<p>Displays all containers, including containers that recently finished their execution</p>
<pre>crictl logs <container-id></pre>	<p>Examines the logs of a given container</p>
<pre>crictl exec -it <container-id> /bin/bash</pre>	<p>Attaches a shell to a running container</p>
<pre>journalctl -u kubelet</pre>	<p>Examines the Kubelet logs. Useful when a pod/container fails to spawn.</p>
<pre>crictl stopp <pod-id></pre>	<p>Stops a running K8S pod</p>
<pre>crictl stop <container-id></pre>	<p>Stops a running container</p>
<pre>crictl rmi <image-id></pre>	<p>Removes a container image from the local K8S registry</p>

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.