



Service Configuration

Table of contents

Configuration of the general behavior - `dpl_rt.conf`

P4RT_RPC_SERVER

DPL_ADMIN_RPC_SERVER

DPL_NSPECT_RPC_SERVER

Performance fine tuning - `system.conf`

Device level configuration - `devices.d/<device-id>.conf`

DPL Port ID Assignment

DPL Port ID restrictions

Example DPL Device Configuration File

In the current release of the DPL Runtime Service, there are three types of configuration files.

Each format is similar in nature to INI files but they allow repeated sections (e.g. `[section]`) and the comments are marked by `#` rather than `;`

Configuration of the general behavior - `dpl_rt.conf`

The DPL Runtime Service searches for this path:

```
/etc/dpl_rt_service/dpl_rt.conf
```

The contents typically look like this:

```
# Example of a possible DPL RT Service GENERAL configuration file

[LOGGING]
log_file_path=/var/log/doca/dpl_rt_service/dpl_rtd.log
log_level=INFO
# Possible log_level values (case insensitive):
# DISABLE
# CRITICAL
# ERROR
# WARNING
# INFO
# DEBUG
# TRACE

[P4RT_RPC_SERVER]
server_address=[ : : ]      # IPv6 "ANY" allows IPv4 connections
server_tcp_port=9559

[DPL_ADMIN_RPC_SERVER]
server_address=[ : : ]      # IPv6 "ANY" allows IPv4 connections
server_tcp_port=9600

[DPL_NSPECT_RPC_SERVER]
server_address=[ : : ]      # IPv6 "ANY" allows IPv4 connections
```

```
server_tcp_port=9560
```

The default logging verbosity that is configured here will be effective when the DPL Runtime Service is started. When the DPL Runtime Service is running, the logging level can be modified with the [DPL Admin client](#), provided in the DPL Dev container. Modifying the logging level from the DPL Admin tool does NOT modify the configuration file, so keep in mind that when the DPL Runtime Service is restarted, the log verbosity will be as specified in the configuration file.

This file also controls the TCP binding of three gRPC servers. It allows you to specify any address (allowing for remote connections from any accessible network interface of the system) or to limit access to a specific IP address that is dedicated for management. Choosing a non-default TCP port for any of the gRPC servers is also possible.

P4RT_RPC_SERVER

This is the server that listens for clients implementing the [P4Runtime protocol](#).

An [open-source client](#) is provided in the DPL Dev container.

DPL_ADMIN_RPC_SERVER

This is the server that listens for the p4admin client that is provided in the DPL Dev container.

DPL_NSPECT_RPC_SERVER

This is the server that listens to the DPL Nspect client/debugger.

Performance fine tuning - system.conf

The DPL Runtime Service searches for this path:

```
/etc/dpl_rt_service/system.conf
```

```
# Example of a possible DPL RT Service system configuration file
```

```
[HAL]  
queue_size=1024  
queues_num=1  
burst_size=32
```

These parameters control the internal behavior of how the DPL Runtime Service interfaces with the underlying hardware. Manipulating these values may affect the maximum rate of rule insertion/deletion and its latency. Tuning these parameters for optimal values can be a complex process and is dependent on the use case and configuration of the DPU. For the current release, this file is for internal use and it is not advised for end users to change the default values, unless under the direct guidance of NVIDIA technical support.

Device level configuration - devices.d/ .conf

The DPL Runtime Service searches for this path:

```
/etc/dpl_rt_service/devices.d/<device-id>.conf
```

e.g. `/etc/dpl_rt_service/devices.d/1000.conf`

A template is available here:

```
/etc/dpl_rt_service/devices.d/NAME.conf.template
```

In the template, you will also find internal documentation about the meaning of each field.

Info

If you use Scalable Functions (SFs) or SR-IOV Virtual Functions (VFs), be sure to refer to their representors in the configuration file.

The `mac` and `mtu` settings are for future implementation and currently have no effect.

The values are returned back as they appear in the file when queried with the `dpl_admin` client but they should not be relied upon.

DPL Port ID Assignment

The DPL Port IDs are assigned by the user. The user decides which DPL Port ID is assigned to which DPU interface. This mapping is critical for achieving the desired results when adding P4 table entries. A DPL Port can be:

- Uplink net device interface
- Host PF representor net device interface (`pf<X>hpf`)
- VF representor net device interface
- SF representor net device interface

Note

Make sure all representor ports added to the configuration file belong to the same uplink port used.

DPL Port ID restrictions

The following restrictions must be considered when assigning DPL Port IDs:

- ID of value `UINT32_MAX` is reserved
- For P4 device, the ID must be an integer number greater than zero
- For interfaces, the DPL ID must be an integer number between zero and `UINT32_MAX`
- Currently, only one Uplink port can be added to the configuration file

Example DPL Device Configuration File

```
# Example of a possible DPL RT Service Device configuration file:
#
# This configuration file specifies the DPL device and its interfaces
# and their DPL Port IDs that will be used by a DPL program.
#
# The DPL Port IDs are assigned by the user. The user decides which
# DPL Port ID is assigned to which ConnectX/DPU interface. This mapping
# is critical for achieving the desired results when adding table entries.
# For DPL device, the ID must be an integer number greater than zero.
#
# The configuration file consists of following sections:
# - [DEVICE] section: Must appear only once.
# - [P4_RT_CONTROLLER] section: Must appear only once.
# - [INTERFACE] section: Must be repeated for each DPL Port (network interface).

[ DEVICE ]
# The DPL Device ID, used for connecting a controller to manage this device's tables.
dpl_device_id=1000
# Cache counter - decrease HW accesses - when expired an HW access will occur upon request.
dpl_counter_cache_timeout=0

[ P4_RT_CONTROLLER ]
# Packets delivered to the DPL RT Service from a controller will have this source DPL Port ID.
# So, this ID can be used for matching traffic originated from the controller.
p4_controller_port_id=9876

[ INTERFACE ]
# Interface name on the system to attach to this DPL device.
interface=p0
# DPL Port ID, used to reference this port by the DPL program and/or when updating table entries.
dpl_logical_port_id=0
# Ethernet frame size.
mtu=1514
# Only uncomment and provide this if you wish to override the interface's MAC address.
# mac=00:00:00:00:00:00
```

```
[INTERFACE]
interface=pf0hpf
dpl_logical_port_id=65535
mtu=1514
# mac=00:00:00:00:00:00
```

```
[INTERFACE]
interface=pf0vf0
dpl_logical_port_id=1
mtu=1514
# mac=00:00:00:00:00:00
```

```
[INTERFACE]
interface=pf0vf1
dpl_logical_port_id=2
mtu=1514
# mac=00:00:00:00:00:00
```

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 (“NVIDIA”) makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes 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 and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

© Copyright 2025, NVIDIA. PDF Generated on 04/24/2025