



## **Loading DPL Applications**

# Table of contents

Introduction

---

Prerequisites

---

Loading the Application

---

Using p4runtime\_sh.sh Script

---

Example Command

---

P4Runtime Optimizations

---

NVIDIA DPL programs are deployed to the NVIDIA® BlueField® networking platform (DPU or SuperNIC) using the P4Runtime API. This allows platform-independent, standards-based integration with SDN controllers.

## Introduction

The DPL compiler generates a pipeline binary optimized for execution on BlueField. Pipeline loading and control—such as installing the program and populating P4 tables—are handled via the P4Runtime API, an open and well-defined interface.

The P4Runtime server, running on the BlueField device, enables a P4 Controller to:

- Connect over gRPC
- Set the `ForwardingPipelineConfig` (i.e., install the compiled DPL binary and `p4info`)
- Query the device for its current pipeline config and table state
- Maintain runtime P4 tables as defined in the DPL source

This model enables integration with open-source, proprietary, or custom-built controllers in a standardized way.

## Prerequisites

Before loading a DPL application, ensure the following services and components are properly set up:

- [DPL Runtime Service](#) is running and configured on the BlueField (Arm side). See the [Container Deployment](#) page for setup instructions.
- [DPL Development Container](#) and the `p4runtime_sh.sh` launch script are installed on the host. See the [DPL Installation Guide](#) for more details.

## Loading the Application

The a DPL application can be loaded using:

- A custom P4Runtime controller

- The the [NVIDIA-supplied Python controller](#)
- An [open source controller](#)

In the following example we'll be using the P4Runtime controller bundled within the [DPL Development Container](#).

## Using p4runtime\_sh.sh Script

Running the script with no arguments displays the usage information:

```
usage: p4runtime_sh.sh -i <docker image> -p <program_folder> -a
<dpl_rtd_ip:port> [OPTIONS]
```

## Example Command

```
p4runtime_sh.sh -i doca_p4_dev:latest -p /root/hello_packet/_out
-a 192.168.1.100:9559 --device-id 1000
```

Arguments:

- `-i` - The pulled [DPL Development Container](#) image.
- `-p` - Directory that holds the DPL program compilation outputs ([Compiling DPL Applications](#)).
- `-a` - Address of the DPL Runtime daemon and the P4Runtime port as specified in the DPL Runtime [Service Configuration](#).
- `--device-id` - (Optional) ID of the target device

After successful loading, the script launches an interactive Python shell connected to the P4Runtime server. From here, you can inspect and manipulate tables (see [p4runtime\\_sh Usage](#)).

## P4Runtime Optimizations

DPL table entries are added via a P4Runtime controller, which may run remotely or locally on BlueField.

- The standard gRPC-based controller model supports ~50K rule insertions per second
- For use cases requiring high-speed (1M+) rule insertions, NVIDIA is introducing a bulk insertion API extension to the P4Runtime protobuf specification

### Note

This feature is planned for future releases.

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.<br/><br/><br/><br/>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.<br/><br/><br/><br/>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.<br/><br/><br/><br/><b>Trademarks</b><br/><br/><br/>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.<br/>

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