



Compiling DPL Applications

Table of contents

Introduction

Installation

Dependencies

Compiling Applications

Compiling Applications in Debug Mode

This section describes how to use the NVIDIA DPL compiler to compile DPL applications.

Introduction

DOCA Programming Language (DPL) applications are an educational resource provided as a guide on how to program on the NVIDIA® BlueField® networking platform (DPU or SuperNIC) using DPL.

The NVIDIA DPL compiler (dplp4c.sh) is provided as an executable in a self contained docker image. The docker image is available to partners by logging into the [NVIDIA NGC](#) system and downloading the most recent [DPL docker containers](#).

Installation

Please refer to the [DPL Runtime Service](#) for the system requirements and [DPL Installation Guide](#) on how to install the compiler. For information on the development environment, refer to the [DOCA Pipeline Language Model](#) and the [DPL System Overview](#).

Dependencies

- BlueField-3 is required
- Ubuntu 22.04 hosts (x86) or greater
- A P4Runtime controller based on API version 1.3.0 or greater

Compiling Applications

The DPL compiler is provided with all its dependencies in the dpl-dev container. For details on how to obtain and install the docker images see [DPL Installation Guide](#) .

A shell script, dplp4c.sh, is provided for a convenient way to execute the DPL compiler from the dpl-dev container:

```
./dplp4c.sh [--mount dir]* dplp4c_args*
```

- `--mount` – (optional) directory to be mounted into the container

- `dplp4c_args` – arguments to be passed to the compiler (e.g., `sample.p4`)

For example, compiling the "hello_packet" DPL example is as simple as:

```
local-user@vm-1:~/p4-samples/dpu/hello_packet$ dplp4c.sh hello_packet.p4
Generating compiler output in "_out"
updating: MANIFEST.json (deflated 40%)
updating: hello_packet.program.json (deflated 87%)
updating: hello_packet.debug.json (deflated 96%)
```

The following files in this example are produced in the `_out` directory:

<code>compiler.log</code>	Log of any compiler warnings or errors to the specified program
<code>hello_packet.dpl config</code>	Binary blob, containing all the data needed for the DPL Runtime daemon to load the program
<code>hello_packet.p4i nfo.txt</code>	P4Runtime protobuf file, in text format

Additional NVIDIA® BlueField®-specific arguments include:

```
--help                Print this help message
--version             Print compiler version
-I path              Specify include path
(passed to preprocessor)
-D arg=value         Define macro (passed to
preprocessor)
-U arg              Undefine macro (passed to
preprocessor)
-E                  Preprocess only, do not
compile (prints program on stdout)
-M                  Output `make` dependency
rule only (passed to preprocessor)
-MD                 Output `make` dependency
rule to file as side effect (passed to preprocessor)
-MF file            With -M, specify output
file for dependencies (passed to preprocessor)
```

-MG	with -M, suppress errors for missing headers (passed to preprocessor)
-MP	with -M, add phony target for each dependency (passed to preprocessor)
-MT target	With -M, override target of the output rule (passed to preprocessor)
-MQ target	Like -Mt, override target but quote special characters (passed to preprocessor)
-g Nspect	Enable debugging via DPL
--nocpp	Skip preprocess, assume input file is already preprocessed.
--Wdisable[=diagnostic]	Disable a compiler diagnostic, or disable all warnings if no diagnostic is specified.
--Winfo[=diagnostic]	Report an info message for a compiler diagnostic.
--Wwarn[=diagnostic]	Report a warning for a compiler diagnostic, or treat all info messages as warnings if no diagnostic is specified.
--Werror[=diagnostic]	Report an error for a compiler diagnostic, or treat all warnings as errors if no diagnostic is specified.
--maxErrorCount errorCount	Set the maximum number of errors to display before failing.
--target target	Compile for the specified target device.
--odir out_directory	Write output to out directory
--enable feature[,feature]* --help	Enable a feature, or comma-separated list of features
--disable feature[,feature]* --help	Disable a feature, or comma-separated list of features

See section [Compiling DPL Applications](#) for examples of how to use the `dplp4c.sh` to compile.

i Info

The binary `<sample_name>.` is created under `_out/<sample_name>.dplconfig.`

i Info

The P4Runtime file is created under `_out/<sample_name>.p4info.txt.`

Compiling Applications in Debug Mode

To enable packet debugging, you must ensure the following:

- Compile your DPL program with the `-g` flag. This enables parsing for a debug flow.
- Ensure you have at least one `nv_send_debug_pkt()` call in your DPL program. A debug packet will be emitted where `nv_send_debug_pkt()` is placed, and represent the program state at that point.

Please note packet debugging is only available on the RX direction. Debug packets will not be produced on TX. You may optionally include a "cookie", within your debug extern call, i.e.: `nv_send_debug_pkt(8w0x42)`. This cookie may help differentiate between different debug extern calls, and will appear in your debugging output.

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