



DOCA Pipeline Language Developer Tool

Table of contents

| | |
|-----------------------|----|
| DPL Nspect | 4 |
| DPL Debugger | 29 |
| DPL Admin | 41 |
| P4 Runtime Controller | 51 |

These pages describe the NVIDIA DOCA Pipeline Language (DPL) Developer Tools, which are used to inspect and debug DPL programs loaded onto the NVIDIA® BlueField® DPU.

For details on creating, compiling, and running DPL programs, refer to the [DOCA Pipeline Language Services Guide](#) page.

Introduction

Developers working with DPL programs require tools for validation and debugging. NVIDIA provides the following DPL developer tools:

- DPL Compiler – converts DPL source code into a binary blob that can be loaded by the DPL Runtime daemon on the DPU
- P4 Runtime Shell – control plane manager that implements the P4Runtime API
- DPL Nspect – command-line interface (CLI) tool for inspecting the current state of tables, keys, counters, and more
- DPL Debugger – user interface tool for debugging packet flow within a DPL program
- DPL Admin – CLI tool for dynamically modifying configuration items in the DPL Runtime daemon

These tools interact with the DPL Runtime daemon (`dpl_rtd`) running on BlueField.

The following pages outline the usage, commands, and options available for each tool.

Prerequisites

1. DPL Runtime Service

1. The DPL Runtime Service container must be set up and running. This service manages the DPL Runtime daemon.
2. `dpl_rtd` must be active on a BlueField configured to operate in DPU mode on the host machine.

2. DPL program compiled for debugging:

1. Compile the DPL program using the DPL Compiler with the debug option enabled (e.g., `dplp4c.sh -g`).

2. Load the compiled binary into `dpl_rtd` using the P4Runtime controller.

Info

For more, refer to section [Compiling DPL Applications](#).

Setup

The developer tools are included in the [DPL Development Container](#), which is publicly available on NGC.

For installation instructions, refer to [DPL Installation Guide](#).

DPL Nspect

This page describes the DPL Nspect tool, which is used to inspect DPL programs loaded onto the NVIDIA® BlueField® DPU.

Introduction

DPL Nspect is a CLI tool with several key use cases:

- Viewing the dynamic state of DPL program components, including tables, entries, keys, actions, and counters
- Recording a debug session and generating an archive that can be opened in the DPL Debugger
- Dumping the state of a DPL program into an archive for offline analysis

DPL Nspect Execution

Running `dpl_nspect.sh` without arguments displays the basic usage information. For a more detailed usage guide, use the `-h` or `--help` argument. This argument also provides command-specific usage details (e.g., `dpl_nspect.sh [command] -h`).

Connecting to a Host

DPL Nspect requires a connection to a local or remote host where the `dpl_nspect_server` is running. The host address and port can be configured using the environment variable. For example: `DPL_NSPECT_SERVER_ADDRESS=127.0.0.1:9560`.

Command Output

The output format of executed commands varies depending on the command. In some cases, the output can be modified using additional arguments to enable integration with automated scripts.

DPL Nspect Commands and Flags

The `dpl_nspect` offers various commands for interacting with the DPL program and BlueField:

- `system-info` – Displays BlueField system information, including hardware model, OS name and version, and OFED version
- `devices` – Lists all configured devices on BlueField, including their ID, name, virtual devices, and the loaded DPL program name

Note

Virtual devices are software abstractions representing real or emulated hardware components. They provide a standardized interface for applications to interact with BlueField resources, regardless of the underlying physical hardware. For example:

- Virtual network interfaces (VNIs) – Logical representations of physical NICs, allowing multiple virtual networks to share a single physical interface
- Virtual processing units (VPUs) – Software-defined processing units within BlueField that can be allocated to applications for tasks like packet processing or security offloading
- Virtual storage devices – Software-defined representations of physical storage, offering flexibility in managing and allocating storage resources

- `counters` – Displays details and values of DPL counters, including the associated table, offset, and value
- `tables` – Lists the P4 tables of DPL programs loaded on BlueField, detailing their keys, actions, counters, and references to the program's source location

- `query` – Lists P4 table entries, showing keys, actions, and the priority of each key set
- `graph` – Generates a DPL pipeline graph using DOT language, which can be rendered with DOT-compatible visualization tools
- `hw-steering` – Displays a hardware steering rules dump, providing a low-level view of hardware entries
- `debug` – starts a debug session

Help Flags

To display details of `dpl_nspect` options usage use any of the following flags:

- `help`, `-h`, `--help` – Displays the help message and exits

Output Control Flags

The following optional flags can be used independently to modify command output format and behavior:

- `-ll`, `--low-level` (default: false) – Displays low-level internal information (DOCA HWS implementation)
- `-ni`, `--non-interactive` (default: false) – Disables interactive output mode. Interactive mode outputs data in chunks similar to the Linux `less` command, making it easier to analyze large outputs. Disabling this mode may be necessary when running `dpl_nspect` commands in automated scripts.
- `-nh`, `--no-hints` – Disables hints for additional information
- `--csv` (default: false) – Outputs command results in CSV format
- `--json` (default: false) – Outputs command results in JSON format
- Flags for the `dpl_nspect counters` command:

- `--top-packets TOP_PACKETS` (default: false) – Displays the top entries with the highest packet counter values
- `--top-bytes TOP_BYTES` (default: false) – Displays the top entries with the highest byte counter values
- Flags for the `dp1_nspect_graph` counters command:
 - `--url` (default: false) – Encodes the output as a URL that can be directly pasted into a web browser for quick rendering using online visualization tools
 - `--type` – Specifies the type of graph to generate

Filter Control Flags

The following flags allow filtering of data based on specific criteria:

- `-d DEVICE_IDS, --device_ids DEVICE_IDS` (default: None) – Specifies a comma-separated list of device IDs to filter, e.g., `1000,2000`. If not specified, data is retrieved for all devices.
- `--table_ids TABLE_IDS` (default: `[]`) – Filters results by specific table IDs
- `--table_names TABLE_NAMES` (default: `[]`) – Filters results by specific table names
- `--counter-names COUNTER_NAMES` (default: `[]`) – Filters results by specific counter names
- `--indices INDICES` (default: `[]`) – Filters results by specific entry indices

Control Flags for Connecting to BlueField

- `-a ADDRESS, --address ADDRESS` – Specifies the `dp1_nspect_server` socket address in the format `[IPv4/IPv6][:port]`.

i Info

For IPv6 addresses, if a port is specified, the address must be enclosed in square brackets. Examples:

- `ipv6:[2607:f8b0:400e:c00::ef]:443`
- `ipv6:[::]:1234`

- This parameter is mandatory, unless the environment variable `DPL_NSPECT_SERVER_ADDRESS` is set with the address and port. For example:

```
DPL_NSPECT_SERVER_ADDRESS=127.0.0.1:9560
```

- `--timeout TIMEOUT` (default: `0`) – Specifies the connectivity timeout (optional)

Misc Flags

- `--version`, `-V` – Displays the program's version number and exits

DPL Nspect Commands

System-info

```
dpl_nspect system-info
```

| DESCRIPTION | VALUE |
|--------------|--------|
| DOCA Version | 2.10.0 |

| DESCRIPTION | VALUE |
|-----------------------------|--|
| Debug Schema Version | Device ID: 1000, doctype: /dplp4c/devtools/base/v6, target_doctype: /dplp4c/devtools/doca/v1 |
| Linux OS version | Ubuntu 22.04 LTS |
| OFED Version | MLNX_OFED_LINUX-24.01-0.1.7.0 |
| DPL Runtime Service version | Not Found |
| RDMA Core Version | 99999P4.2410m1nx54-1.P4.109.gec00b3c78ef6 |

Devices

```
dpl_nspect devices
```

| ID | NAME | PCI ADDRESS | PORT ID | NAME | NUMBER | DPL PROGRAM | FW VERSION | HW MODEL |
|------|------|--------------|---------|---------------|--------|-------------|------------|--|
| 1000 | br9 | 0000:65:00.0 | 0 | enp101s0f0np0 | 50 | alpha2_d2 | 32.38.0588 | BlueField-3 integrated ConnectX-7 network controller |
| | | | 4 | eth3 | 5 | | | |
| | | | 1 | eth0 | 2 | | | |
| | | | 2 | eth1 | 3 | | | |
| | | | 3 | eth2 | 4 | | | |

Counters

dpl_nspect counters

DEVICE: 1000, NV_COUNTER: alpha2_drop2.decap.decap_counter, SCOPE: alpha2_drop2, TYPE: Pkts/Bts, SIZE: 4, IN USE: 3

| IDX | P4 | ENTRY OFFSET | VALUE |
|-----|-------------|--------------|----------|
| 0 | decap_table | 1 | (2, 140) |
| 1 | decap_table | 2 | (0, 0) |
| 2 | decap_table | 0 | (0, 0) |

DEVICE: 1000, NV_COUNTER: alpha2_drop2.encap.encap_counter, SCOPE: alpha2_drop2, TYPE: Pkts/Bts, SIZE: 4, IN USE: 3

| IDX | P4 | ENTRY OFFSET | VALUE |
|-----|-------------|--------------|--------|
| 0 | encap_table | 0 | (0, 0) |
| 1 | encap_table | 2 | (0, 0) |
| 2 | encap_table | 1 | (0, 0) |

DEVICE: 1000, NV_DIRECT_COUNTER: alpha2_drop2.gtp.gtp_counter, SCOPE: alpha2_drop2, TYPE: Pkts/Bts, SIZE: 128, IN USE: 3

| IDX | P4 | ENTRY OFFSET | VALUE |
|-----|-----------|--------------|--------|
| 0 | gtp_table | 2 | (0, 0) |
| 1 | gtp_table | 1 | (0, 0) |
| 2 | gtp_table | 0 | (0, 0) |

The value is in the following format: (number of packets, number of bytes)

Tables

dp1_nspect tables

| DEVICE: 1000 | | | | | | |
|------------------|--------------------------------|---|---|------------------------------|---------|------------------|
| ID | TABLE | KEYS | ACTIONS | COUNTER | ENTRIES | SOURCES |
| 449 637 66 | alpha2_drop2.decap.decap_table | std_meta.ingress_port headers.gre.\$valid\$ headers.vxlan.\$valid\$ | alpha2_drop2.decap.send_to_port alpha2_drop2.decap.decap_gre alpha2_drop2.decap.decap_vxlan NoAction | | 4 | alpha2_d2.p4:235 |
| 415 214 46 | alpha2_drop2.encap.encap_table | std_meta.ingress_port headers.ipv4.src_addr | alpha2_drop2.encap.send_to_port alpha2_drop2.encap.encap_gre alpha2_drop2.encap.encap_vxlan NoAction | | 4 | alpha2_d2.p4:173 |
| 343 181 61 | alpha2_drop2.gtp.gtp_table | std_meta.ingress_port headers.gtpv1.teid | alpha2_drop2.gtp.send_to_port alpha2_drop2.gtp.drop | alpha2_drop2.gtp.gtp_counter | 3 | alpha2_d2.p4:102 |

Query

dp1_nspect query

DEVICE: 1000, TABLE: alpha2_drop2.gtp.gtp_table, ID: 34318161, SOURCE: alpha2_d2.p4:102

| PRI O | P4 KEY | VALU E | MAS K | TYP E | P4 ACTION |
|----------|---|------------|----------|------------------------|---|
| 0 | std_meta.ingress_port headers.gtpv1.teid | 0x0 0x1 | | exac t exac t | alpha2_drop2.gtp.send_to_port(port=0x1) |
| 0 | std_meta.ingress_port headers.gtpv1.teid | 0x1 0x1 | | exac t exac t | alpha2_drop2.gtp.send_to_port(port=0x0) |
| 0 | ALWAYS_HIT | | | | alpha2_drop2.gtp.drop() |

DEVICE: 1000, TABLE: alpha2_drop2.encap.encap_table, ID: 41521446, SOURCE: alpha2_d2.p4:173

| PRI O | P4 KEY | VALUE | MASK | TYPE | P4 ACTION |
|----------|--|-----------------------|----------------|----------------------|--|
| 3 | std_meta.ingress_port headers.ipv4.src_addr | 0x0 0x10000 000 | 0xff000 000 | exact terna ry | alpha2_drop2.encap.send_to_port(port=0x4) |
| 2 | std_meta.ingress_port headers.ipv4.src_addr | 0x0 0x20000 000 | 0xff000 000 | exact terna ry | alpha2_drop2.encap.encap_gre(port=0x2, src_mac=0x112233445566, dst_mac=0x10203040506, sip=0x1010101, dip=0x2020202, key=0xabcd) |
| 1 | std_meta.ingress_port | 0x0 | 0xff000 000 | exact | alpha2_drop2.encap.encap_vxlan(port=0x1) |

DEVICE: 1000, TABLE: alpha2_drop2.encap.encap_table, ID: 41521446, SOURCE: alpha2_d2.p4:173

| | | | | | |
|---|-----------------------|------------|--|---------|---|
| | headers.ipv4.src_addr | 0x20000000 | | ternary | port=0x3, src_mac=0x112233445566, dst_mac=0x10203040506, sip=0x1010101, dip=0x2020202, vni=0xabcd) |
| 0 | ALWAYS_HIT | | | | NoAction() |

DEVICE: 1000, TABLE: alpha2_drop2.decap.decap_table, ID: 44963766, SOURCE: alpha2_d2.p4:235

| PRI O | P4 KEY | VALU E | MAS K | TYP E | P4 ACTION |
|-------|---|-------------------|-------|-------------------------|--|
| 0 | std_meta.ingress_port headers.gre.\$valid\$ headers.vxlan.\$valid\$ | 0x3 0x0 0x1 | | exact exact exact | alpha2_drop2.decap.decap_vxlan(port=0x0, src_mac=0xaabbccddeeff, dst_mac=0x112233445566) |
| 0 | std_meta.ingress_port headers.gre.\$valid\$ headers.vxlan.\$valid\$ | 0x4 0x0 0x0 | | exact exact exact | alpha2_drop2.decap.send_to_port(port=0x0) |
| 0 | std_meta.ingress_port headers.gre.\$valid\$ headers.vxlan.\$valid\$ | 0x2 0x1 0x0 | | exact exact exact | alpha2_drop2.decap.decap_gre(port=0x0) |
| 0 | ALWAYS_HIT | | | | NoAction() |

Different Output Option Flags

- `--csv`

```
dpl_nspect --csv tables
```

```
"DEVICE: 1000"  
"ID", "TABLE", "KEYS", "ACTIONS", "COUNTER", "ENTRIES", "SOURCE"  
"44963766", "alpha2_drop2.decap.decap_table", "std_meta.ingress_port  
headers.gre.$valid$  
headers.vxlan.$valid$","alpha2_drop2.decap.send_to_port  
alpha2_drop2.decap.decap_gre  
alpha2_drop2.decap.decap_vxlan  
NoAction","","4","alpha2_d2.p4:235"  
"41521446", "alpha2_drop2.encap.encap_table", "std_meta.ingress_port  
headers.ipv4.src_addr","alpha2_drop2.encap.send_to_port  
alpha2_drop2.encap.encap_gre  
alpha2_drop2.encap.encap_vxlan  
NoAction","","4","alpha2_d2.p4:173"  
"34318161", "alpha2_drop2.gtp.gtp_table", "std_meta.ingress_port  
headers.gtpv1.teid","alpha2_drop2.gtp.send_to_port  
alpha2_drop2.gtp.drop","alpha2_drop2.gtp.gtp_counter","3","alpha2_d2
```

- `--json`

```
dpl_nspect --json tables
```

```
[  
  [  
    {  
      "DEVICE": 1000  
    },  
    {
```

```

    "ID" : 44963766 ,
    "TABLE" : "alpha2_drop2.decap.decap_table",
    "KEYS" : [
        "std_meta.ingress_port",
        "headers.gre.$valid$",
        "headers.vxlan.$valid$"
    ],
    "ACTIONS" : [
        "alpha2_drop2.decap.send_to_port",
        "alpha2_drop2.decap.decap_gre",
        "alpha2_drop2.decap.decap_vxlan",
        "NoAction"
    ],
    "COUNTER" : "",
    "ENTRIES" : 4,
    "SOURCE" : {
        "filename" : "/tmp/p4src/alpha2_d2.p4",
        "line" : 235,
        "column" : 11
    }
},
{
    "ID" : 41521446,
    "TABLE" : "alpha2_drop2.encap.encap_table",
    "KEYS" : [
        "std_meta.ingress_port",
        "headers.ipv4.src_addr"
    ],
    "ACTIONS" : [
        "alpha2_drop2.encap.send_to_port",
        "alpha2_drop2.encap.encap_gre",
        "alpha2_drop2.encap.encap_vxlan",
        "NoAction"
    ],
    "COUNTER" : "",
    "ENTRIES" : 4,

```



```

    "SOURCE": {
      "filename": "/tmp/p4src/alpha2_d2.p4",
      "line": 173,
      "column": 11
    }
  },
  {
    "ID": 34318161,
    "TABLE": "alpha2_drop2.gtp.gtp_table",
    "KEYS": [
      "std_meta.ingress_port",
      "headers.gtpv1.teid"
    ],
    "ACTIONS": [
      "alpha2_drop2.gtp.send_to_port",
      "alpha2_drop2.gtp.drop"
    ],
    "COUNTER": "alpha2_drop2.gtp.gtp_counter",
    "ENTRIES": 3,
    "SOURCE": {
      "filename": "/tmp/p4src/alpha2_d2.p4",
      "line": 102,
      "column": 11
    }
  }
]
]

```

Graph (requires the --low-level flag)

- No extra flags

```
dpl_nspect -ll graph --type pipeline_low_level
```

```
Device 1000:
digraph hybrid_pipeline {
    graph [compound=true]
        // Legend
        subgraph cluster_legend {
            hal_table [label=<<B>HAL Table ID</B><BR/>
[Implements P4 Objects]> shape=diamond]
            legend_invis [label=legend_invis shape=point
style=invis width=0]
            p4_control [label="P4 Control" color=blue
shape=rectangle]
        }
        alpha2_drop2_start [label=start color=green
shape=none]
        legend_invis -> alpha2_drop2_start [style=invis]
        // Pipeline Stage: main
        subgraph cluster_alpha2_drop2 {
            graph [color=blue label=alpha2_drop2
shape=rectangle]
                24 [label=<<B>24</B><BR/>[alpha2_drop2]>
shape=diamond]
                alpha2_drop2_start -> 24
                48 [label=<<B>48</B><BR/>
[alpha2_drop2.gtp.gtp_table]> shape=diamond]
                DROP [label=DROP fontcolor=red shape=box]
                48 -> DROP
                alpha2_drop2_end [label=end color=red
shape=none]
                48 -> alpha2_drop2_end
                36 [label=<<B>36</B><BR/>[alpha2_drop2]>
shape=diamond]
```

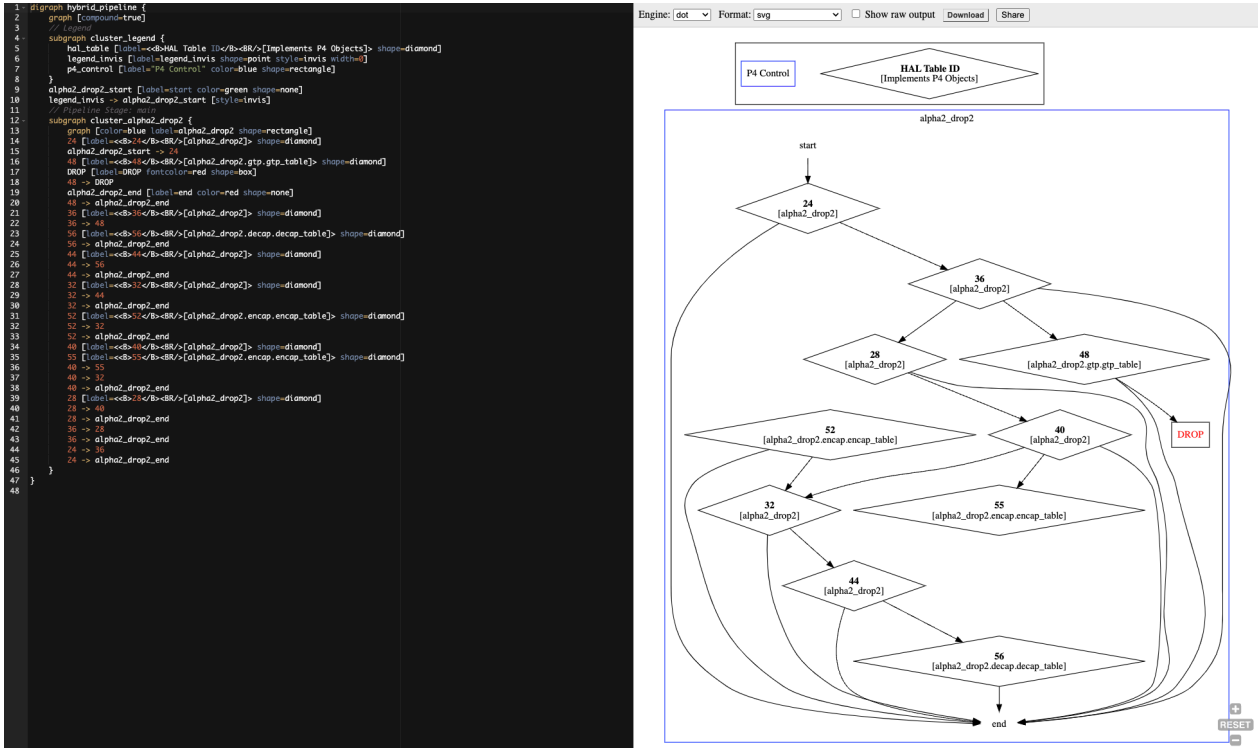
```

36 -> 48
56 [label=<<B>56</B><BR/>
[alpha2_drop2.decap.decap_table]> shape=diamond]
56 -> alpha2_drop2_end
44 [label=<<B>44</B><BR/>[alpha2_drop2]>
shape=diamond]
44 -> 56
44 -> alpha2_drop2_end
32 [label=<<B>32</B><BR/>[alpha2_drop2]>
shape=diamond]
32 -> 44
32 -> alpha2_drop2_end
52 [label=<<B>52</B><BR/>
[alpha2_drop2.encap.encap_table]> shape=diamond]
52 -> 32
52 -> alpha2_drop2_end
40 [label=<<B>40</B><BR/>[alpha2_drop2]>
shape=diamond]
55 [label=<<B>55</B><BR/>
[alpha2_drop2.encap.encap_table]> shape=diamond]
40 -> 55
40 -> 32
40 -> alpha2_drop2_end
28 [label=<<B>28</B><BR/>[alpha2_drop2]>
shape=diamond]
28 -> 40
28 -> alpha2_drop2_end
36 -> 28
36 -> alpha2_drop2_end
24 -> 36
24 -> alpha2_drop2_end
}
}

```

- `--url` – generates an HTTP URL link to the program's control flow graph, in dot format, rendered using <https://dreampuf.github.io/GraphvizOnline/>.

```
dpl_nspect -ll graph --url --type pipeline_low_level
```



- Generate the program's parser graph in dot format:

```
dpl_nspect -ll graph --type parser
```

```
Device 1000:
// Parser packet_parser
digraph packet_parser {
graph [concentrate=true splines=compound]
// Legend
subgraph cluster_legend {
fixed_state [label="Fixed State" style=filled]
```

```

legend_invis [label=legend_invis shape=point style=invis
width=0]
flex_state [label="Flex State"]
default_transition_note [label="All states have an hidden transition to
accept\nEdges are annotated with their match value." shape=none]
legend_invis -> default_transition_note [style=invis]
}
start [label=start color=green shape=point]
default_transition_note -> start [style=invis]
nv_parse_ethernet [label=ethernet id=nv_parse_ethernet
style=filled]
start -> nv_parse_ethernet
nv_parse_vlan [label=vlan id=nv_parse_vlan style=filled]
nv_parse_ethernet -> nv_parse_vlan [label="0x8100" color=gray
id="nv_parse_ethernet-nv_parse_vlan"]
nv_parse_ipv4 [label=ipv4 id=nv_parse_ipv4 style=filled]
nv_parse_ethernet -> nv_parse_ipv4 [label="0x800" color=gray
id="nv_parse_ethernet-nv_parse_ipv4"]
nv_parse_ipv6 [label=ipv6 id=nv_parse_ipv6 style=filled]
nv_parse_ethernet -> nv_parse_ipv6 [label="0x86dd" color=gray
id="nv_parse_ethernet-nv_parse_ipv6"]
nv_parse_mpls [label=mpls id=nv_parse_mpls style=filled]
nv_parse_ethernet -> nv_parse_mpls [label="0x8847" color=gray
id="nv_parse_ethernet-nv_parse_mpls"]
nv_parse_geneve [label=geneve id=nv_parse_geneve
style=filled]
nv_parse_inner_ipv4 [label=inner_ipv4 id=nv_parse_inner_ipv4
style=filled]
nv_parse_geneve -> nv_parse_inner_ipv4 [label="0x800"
color=gray id="nv_parse_geneve-nv_parse_inner_ipv4"]
nv_parse_inner_ipv6 [label=inner_ipv6 id=nv_parse_inner_ipv6
style=filled]
nv_parse_geneve -> nv_parse_inner_ipv6 [label="0x86dd"
color=gray id="nv_parse_geneve-nv_parse_inner_ipv6"]
nv_parse_inner_mpls [label=inner_mpls id=nv_parse_inner_mpls
style=filled]

```

```

nv_parse_geneve -> nv_parse_inner_mpls [label="0x8847"
color=gray id="nv_parse_geneve-nv_parse_inner_mpls"]
nv_parse_inner_ethernet [label=inner_ethernet
id=nv_parse_inner_ethernet style=filled]
nv_parse_geneve -> nv_parse_inner_ethernet [label="0x6558"
color=gray id="nv_parse_geneve-nv_parse_inner_ethernet"]
nv_parse_gre [label=gre id=nv_parse_gre style=filled]
nv_parse_gre -> nv_parse_inner_ipv4 [label="0x800" color=gray
id="nv_parse_gre-nv_parse_inner_ipv4"]
nv_parse_gre -> nv_parse_inner_ipv6 [label="0x86dd" color=gray
id="nv_parse_gre-nv_parse_inner_ipv6"]
nv_parse_gre -> nv_parse_inner_mpls [label="0x8847" color=gray
id="nv_parse_gre-nv_parse_inner_mpls"]
nv_parse_nvgre_vsid [label=nvgre_vsid id=nv_parse_nvgre_vsid
style=filled]
nv_parse_gre -> nv_parse_nvgre_vsid [label="0x6558" color=gray
id="nv_parse_gre-nv_parse_nvgre_vsid"]
nv_parse_icmp [label=icmp id=nv_parse_icmp style=filled]
nv_parse_icmpv6 [label=icmpv6 id=nv_parse_icmpv6
style=filled]
nv_parse_inner_vlan [label=inner_vlan id=nv_parse_inner_vlan
style=filled]
nv_parse_inner_ethernet -> nv_parse_inner_vlan [label="0x8100"
color=gray id="nv_parse_inner_ethernet-nv_parse_inner_vlan"]
nv_parse_inner_ethernet -> nv_parse_inner_ipv4 [label="0x800"
color=gray id="nv_parse_inner_ethernet-nv_parse_inner_ipv4"]
nv_parse_inner_ethernet -> nv_parse_inner_ipv6 [label="0x86dd"
color=gray id="nv_parse_inner_ethernet-nv_parse_inner_ipv6"]
nv_parse_inner_ethernet -> nv_parse_inner_mpls [label="0x8847"
color=gray id="nv_parse_inner_ethernet-nv_parse_inner_mpls"]
nv_parse_inner_icmp [label=inner_icmp id=nv_parse_inner_icmp
style=filled]
nv_parse_inner_icmpv6 [label=inner_icmpv6
id=nv_parse_inner_icmpv6 style=filled]
nv_parse_inner_ipv4 -> nv_parse_inner_icmp [label="0x1"
color=gray id="nv_parse_inner_ipv4-nv_parse_inner_icmp"]

```

```

nv_parse_inner_tcp [label=inner_tcp id=nv_parse_inner_tcp
style=filled]
nv_parse_inner_ipv4 -> nv_parse_inner_tcp [label="0x6"
color=gray id="nv_parse_inner_ipv4-nv_parse_inner_tcp"]
nv_parse_inner_udp [label=inner_udp id=nv_parse_inner_udp
style=filled]
nv_parse_inner_ipv4 -> nv_parse_inner_udp [label="0x11"
color=gray id="nv_parse_inner_ipv4-nv_parse_inner_udp"]
nv_parse_inner_ipv6 -> nv_parse_inner_icmpv6 [label="0x3a"
color=gray id="nv_parse_inner_ipv6-nv_parse_inner_icmpv6"]
nv_parse_inner_ipv6 -> nv_parse_inner_tcp [label="0x6"
color=gray id="nv_parse_inner_ipv6-nv_parse_inner_tcp"]
nv_parse_inner_ipv6 -> nv_parse_inner_udp [label="0x11"
color=gray id="nv_parse_inner_ipv6-nv_parse_inner_udp"]
nv_parse_inner_mpls1 [label=inner_mpls1
id=nv_parse_inner_mpls1 style=filled]
nv_parse_inner_mpls -> nv_parse_inner_mpls1 [label="0x0"
color=gray id="nv_parse_inner_mpls-nv_parse_inner_mpls1"]
nv_parse_inner_mpls_next_header
[label=nv_parse_inner_mpls_next_header
id=nv_parse_inner_mpls_next_header style=filled]
nv_parse_inner_mpls -> nv_parse_inner_mpls_next_header
[label="0x1" color=gray id="nv_parse_inner_mpls-nv_parse_inner_mpls_next_header"]
nv_parse_inner_mpls2 [label=inner_mpls2
id=nv_parse_inner_mpls2 style=filled]
nv_parse_inner_mpls1 -> nv_parse_inner_mpls2 [label="0x0"
color=gray id="nv_parse_inner_mpls1-nv_parse_inner_mpls2"]
nv_parse_inner_mpls1 -> nv_parse_inner_mpls_next_header
[label="0x1" color=gray id="nv_parse_inner_mpls1-
nv_parse_inner_mpls_next_header"]
nv_parse_inner_mpls3 [label=inner_mpls3
id=nv_parse_inner_mpls3 style=filled]
nv_parse_inner_mpls2 -> nv_parse_inner_mpls3 [label="0x0"
color=gray id="nv_parse_inner_mpls2-nv_parse_inner_mpls3"]
nv_parse_inner_mpls2 -> nv_parse_inner_mpls_next_header
[label="0x1" color=gray id="nv_parse_inner_mpls2-

```

```

nv_parse_inner_mpls_next_header" ]
nv_parse_inner_mpls4 [label=inner_mpls4
id=nv_parse_inner_mpls4 style=filled]
nv_parse_inner_mpls3 -> nv_parse_inner_mpls4 [label="0x0"
color=gray id="nv_parse_inner_mpls3-nv_parse_inner_mpls4"]
nv_parse_inner_mpls3 -> nv_parse_inner_mpls_next_header
[label="0x1" color=gray id="nv_parse_inner_mpls3-
nv_parse_inner_mpls_next_header" ]
nv_parse_inner_mpls4 -> nv_parse_inner_mpls_next_header
[label="0x1" color=gray id="nv_parse_inner_mpls4-
nv_parse_inner_mpls_next_header" ]
nv_parse_inner_mpls_next_header -> nv_parse_inner_ipv4
[label="0x4" color=gray id="nv_parse_inner_mpls_next_header-nv_parse_inner_ipv4"]
nv_parse_inner_mpls_next_header -> nv_parse_inner_ipv6
[label="0x6" color=gray id="nv_parse_inner_mpls_next_header-nv_parse_inner_ipv6"]
nv_parse_inner_vlan -> nv_parse_inner_ipv4 [label="0x800"
color=gray id="nv_parse_inner_vlan-nv_parse_inner_ipv4"]
nv_parse_inner_vlan -> nv_parse_inner_ipv6 [label="0x86dd"
color=gray id="nv_parse_inner_vlan-nv_parse_inner_ipv6"]
nv_parse_inner_vlan -> nv_parse_inner_mpls [label="0x8847"
color=gray id="nv_parse_inner_vlan-nv_parse_inner_mpls"]
nv_parse_ipsec_esp [label=esp id=nv_parse_ipsec_esp
style=filled]
nv_parse_ipv4 -> nv_parse_inner_ipv4 [label="0x4" color=gray
id="nv_parse_ipv4-nv_parse_inner_ipv4"]
nv_parse_ipv4 -> nv_parse_inner_ipv6 [label="0x29" color=gray
id="nv_parse_ipv4-nv_parse_inner_ipv6"]
nv_parse_ipv4 -> nv_parse_icmp [label="0x1" color=gray
id="nv_parse_ipv4-nv_parse_icmp"]
nv_parse_tcp [label=tcp id=nv_parse_tcp style=filled]
nv_parse_ipv4 -> nv_parse_tcp [label="0x6" color=gray
id="nv_parse_ipv4-nv_parse_tcp"]
nv_parse_udp [label=udp id=nv_parse_udp style=filled]
nv_parse_ipv4 -> nv_parse_udp [label="0x11" color=gray
id="nv_parse_ipv4-nv_parse_udp"]

```



```

nv_parse_ipv4 -> nv_parse_gre [label="0x2f" color=gray
id="nv_parse_ipv4-nv_parse_gre"]
nv_parse_ipv4 -> nv_parse_ipsec_esp [label="0x32" color=gray
id="nv_parse_ipv4-nv_parse_ipsec_esp"]
nv_parse_ipv6 -> nv_parse_inner_ipv4 [label="0x4" color=gray
id="nv_parse_ipv6-nv_parse_inner_ipv4"]
nv_parse_ipv6 -> nv_parse_inner_ipv6 [label="0x29" color=gray
id="nv_parse_ipv6-nv_parse_inner_ipv6"]
nv_parse_ipv6 -> nv_parse_icmpv6 [label="0x3a" color=gray
id="nv_parse_ipv6-nv_parse_icmpv6"]
nv_parse_ipv6 -> nv_parse_tcp [label="0x6" color=gray
id="nv_parse_ipv6-nv_parse_tcp"]
nv_parse_ipv6 -> nv_parse_udp [label="0x11" color=gray
id="nv_parse_ipv6-nv_parse_udp"]
nv_parse_ipv6 -> nv_parse_gre [label="0x2f" color=gray
id="nv_parse_ipv6-nv_parse_gre"]
nv_parse_ipv6 -> nv_parse_ipsec_esp [label="0x32" color=gray
id="nv_parse_ipv6-nv_parse_ipsec_esp"]
nv_parse_mpls1 [label=mpls1 id=nv_parse_mpls1 style=filled]
nv_parse_mpls -> nv_parse_mpls1 [label="0x0" color=gray
id="nv_parse_mpls-nv_parse_mpls1"]
nv_parse_mpls_next_header [label=nv_parse_mpls_next_header
id=nv_parse_mpls_next_header style=filled]
nv_parse_mpls -> nv_parse_mpls_next_header [label="0x1"
color=gray id="nv_parse_mpls-nv_parse_mpls_next_header"]
nv_parse_mpls2 [label=mpls2 id=nv_parse_mpls2 style=filled]
nv_parse_mpls1 -> nv_parse_mpls2 [label="0x0" color=gray
id="nv_parse_mpls1-nv_parse_mpls2"]
nv_parse_mpls1 -> nv_parse_mpls_next_header [label="0x1"
color=gray id="nv_parse_mpls1-nv_parse_mpls_next_header"]
nv_parse_mpls3 [label=mpls3 id=nv_parse_mpls3 style=filled]
nv_parse_mpls2 -> nv_parse_mpls3 [label="0x0" color=gray
id="nv_parse_mpls2-nv_parse_mpls3"]
nv_parse_mpls2 -> nv_parse_mpls_next_header [label="0x1"
color=gray id="nv_parse_mpls2-nv_parse_mpls_next_header"]
nv_parse_mpls4 [label=mpls4 id=nv_parse_mpls4 style=filled]

```

```

nv_parse_mpls3 -> nv_parse_mpls4 [label="0x0" color=gray
id="nv_parse_mpls3-nv_parse_mpls4"]
nv_parse_mpls3 -> nv_parse_mpls_next_header [label="0x1"
color=gray id="nv_parse_mpls3-nv_parse_mpls_next_header"]
nv_parse_mpls4 -> nv_parse_mpls_next_header [label="0x1"
color=gray id="nv_parse_mpls4-nv_parse_mpls_next_header"]
nv_parse_mpls_next_header -> nv_parse_ipv4 [label="0x4"
color=gray id="nv_parse_mpls_next_header-nv_parse_ipv4"]
nv_parse_mpls_next_header -> nv_parse_ipv6 [label="0x6"
color=gray id="nv_parse_mpls_next_header-nv_parse_ipv6"]
nv_parse_nvgre_vsid -> nv_parse_inner_ethernet [label="0x0"
color=gray id="nv_parse_nvgre_vsid-nv_parse_inner_ethernet"]
nv_parse_psp [label=psp id=nv_parse_psp style=filled]
parse_gtp [label=gtpv1 id=parse_gtp style=""]
nv_parse_udp -> parse_gtp [label="0x868" color=black
id="nv_parse_udp-parse_gtp"]
nv_parse_vxlan [label=vxlan id=nv_parse_vxlan style=filled]
nv_parse_udp -> nv_parse_vxlan [label="0x12b5" color=gray
id="nv_parse_udp-nv_parse_vxlan"]
nv_parse_vxlan_gpe [label=vxlan_gpe id=nv_parse_vxlan_gpe
style=filled]
nv_parse_udp -> nv_parse_vxlan_gpe [label="0x12b6" color=gray
id="nv_parse_udp-nv_parse_vxlan_gpe"]
nv_parse_udp -> nv_parse_geneve [label="0x17c1" color=gray
id="nv_parse_udp-nv_parse_geneve"]
nv_parse_udp -> nv_parse_inner_mpls [label="0x19eb" color=gray
id="nv_parse_udp-nv_parse_inner_mpls"]
nv_parse_udp -> nv_parse_ipsec_esp [label="0x1194" color=gray
id="nv_parse_udp-nv_parse_ipsec_esp"]
nv_parse_udp -> nv_parse_psp [label="0x3e8" color=gray
id="nv_parse_udp-nv_parse_psp"]
nv_parse_vlan -> nv_parse_ipv4 [label="0x800" color=gray
id="nv_parse_vlan-nv_parse_ipv4"]
nv_parse_vlan -> nv_parse_ipv6 [label="0x86dd" color=gray
id="nv_parse_vlan-nv_parse_ipv6"]

```

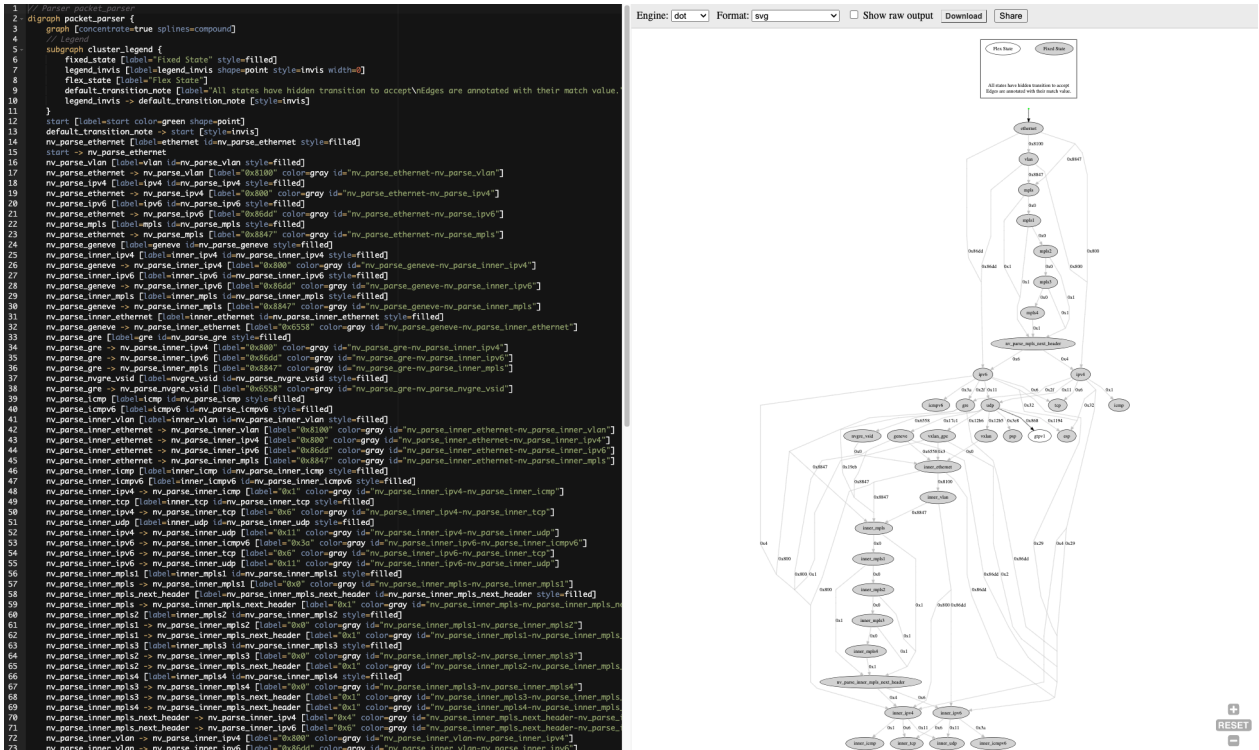
```

nv_parse_vlan -> nv_parse_mpls [label="0x8847" color=gray
id="nv_parse_vlan-nv_parse_mpls"]
nv_parse_vxlan -> nv_parse_inner_ethernet [label="0x0"
color=gray id="nv_parse_vxlan-nv_parse_inner_ethernet"]
nv_parse_vxlan_gpe -> nv_parse_inner_ipv4 [label="0x1"
color=gray id="nv_parse_vxlan_gpe-nv_parse_inner_ipv4"]
nv_parse_vxlan_gpe -> nv_parse_inner_ipv6 [label="0x2"
color=gray id="nv_parse_vxlan_gpe-nv_parse_inner_ipv6"]
nv_parse_vxlan_gpe -> nv_parse_inner_ethernet [label="0x3"
color=gray id="nv_parse_vxlan_gpe-nv_parse_inner_ethernet"]
}

```

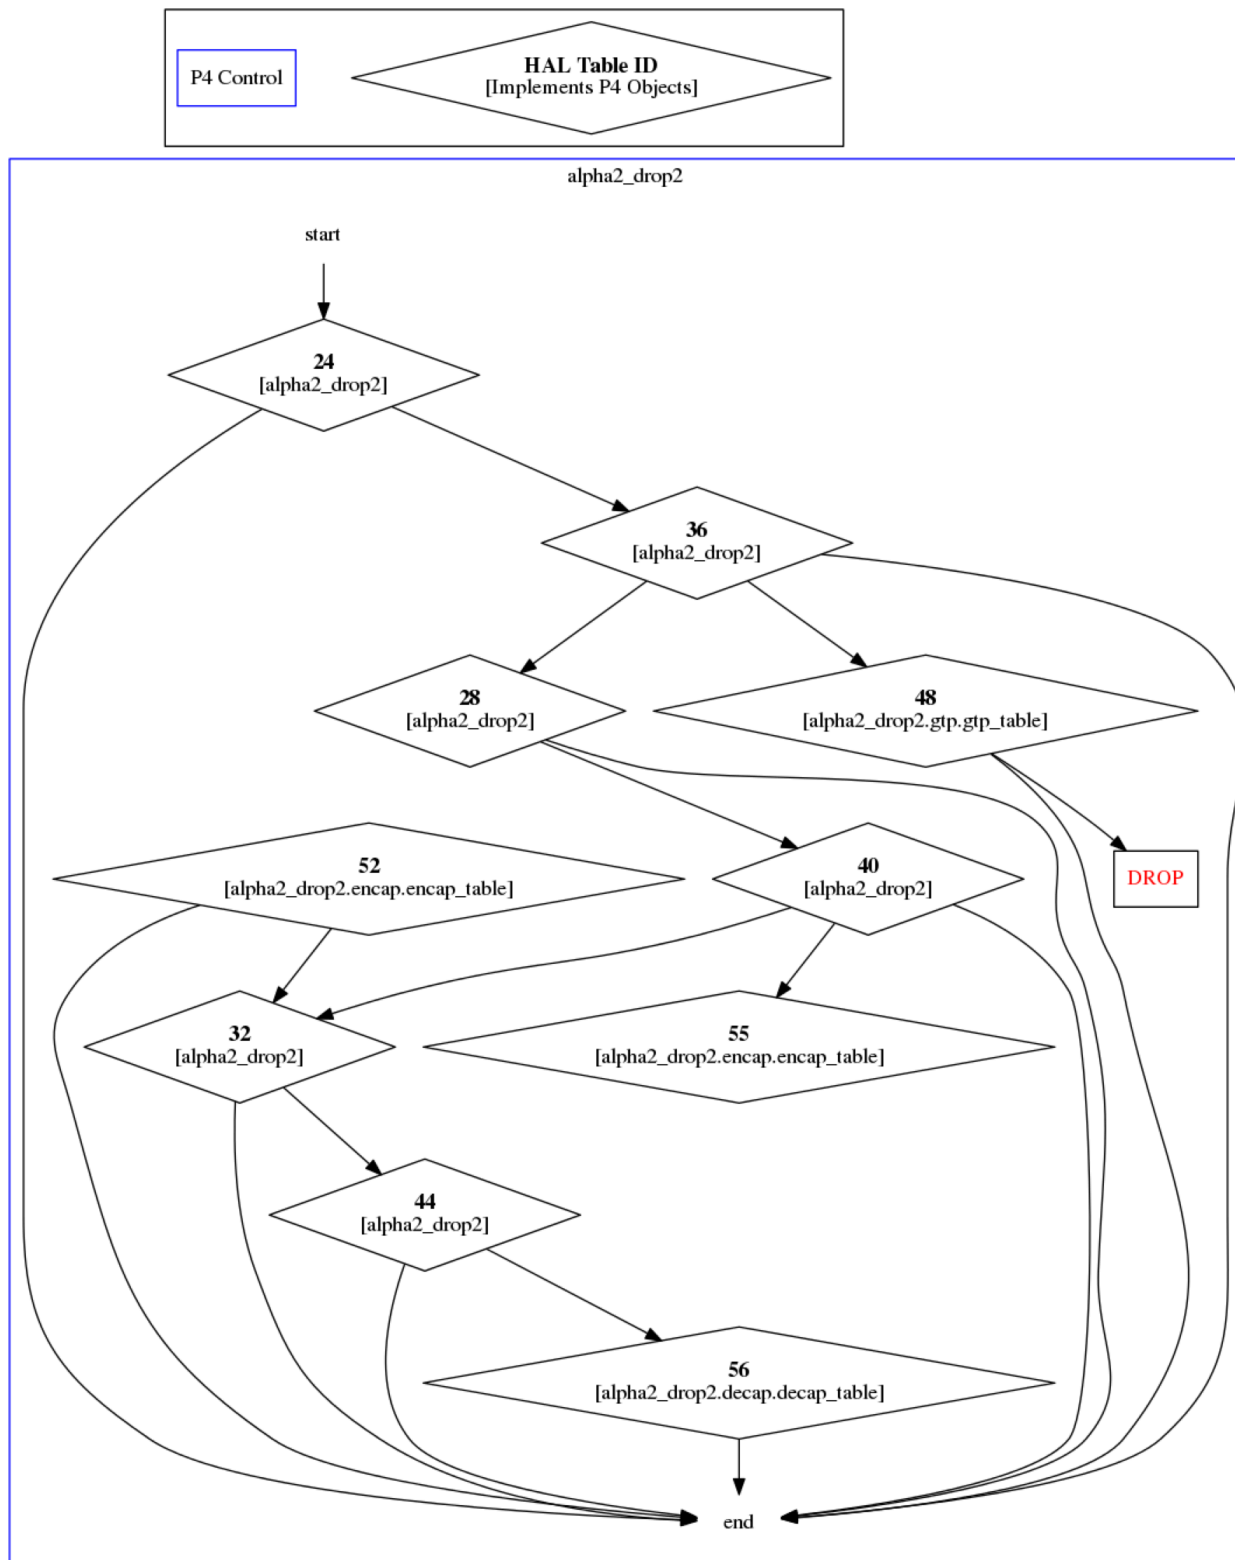
- `--url --type parser`

```
dpl_Inspect -ll graph --url --type parser
```



Illustration

It is possible to render the graph to get the following picture:



Commands and Flags Summary

Some command flags are universally applicable to all commands, but others are applicable to only specific commands.

The following table outlines commands and applicable flags:

| | system-info | devices | counters | tables | query | graph | hw-steering | debug |
|-------------------|-------------|---------|----------|--------|-------|-------|-------------|-------|
| --address | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| --time-out | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| --non-interactive | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| --no-hints | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| --csv | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| --json | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| --device-ids | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| --low-level | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| --table-ids | | | ✓ | ✓ | ✓ | | | |
| --table_names | | | ✓ | ✓ | ✓ | | | |
| --counter-names | | | ✓ | | | | | |
| --indices | | | ✓ | | | | | |
| --top-packets | | | ✓ | | | | | |
| --top-bytes | | | ✓ | | | | | |
| --url | | | | | | ✓ | | |
| --type | | | | | | ✓ | | |
| --output-file | | | | | | | | ✓ |
| --debug-device-id | | | | | | | | ✓ |

DPL Debugger

This page describes the DPL Debugger tools, which are used to debug pipeline packets processed by DPL programs on the NVIDIA® BlueField® DPU.

Introduction

The DPL Debugger is a GUI application that can run a live debug session or open a pre-recorded debug session file captured using the `dpl_nspect debug` command.

A debug session visualizes the flow of network packets through a DPL program's pipeline, allowing developers to trace each packet's path in detail. For more information, refer to the documentation on the [nv_send_debug_pkt](#) extern function.

A debug session visualizes the flow of network packets through a DPL program's pipeline, allowing developers to trace packet details throughout its course. For more information, refer to the documentation on the [nv_send_debug_pkt](#) extern function.

Note

The current version of the debugger only displays debug packets received by the DPU (Rx path) or packets arriving from the primary wire port (PO) in the network-to-host direction.

Executing the DPL Debugger

Running `dpl_debugger . sh` without any arguments launches the GUI application.

To open and debug a pre-recorded session specify the session archive when executing the program:


```
dpl_debugger.sh my-debug-archive.tar.gz
```

Using the DPL Debugger

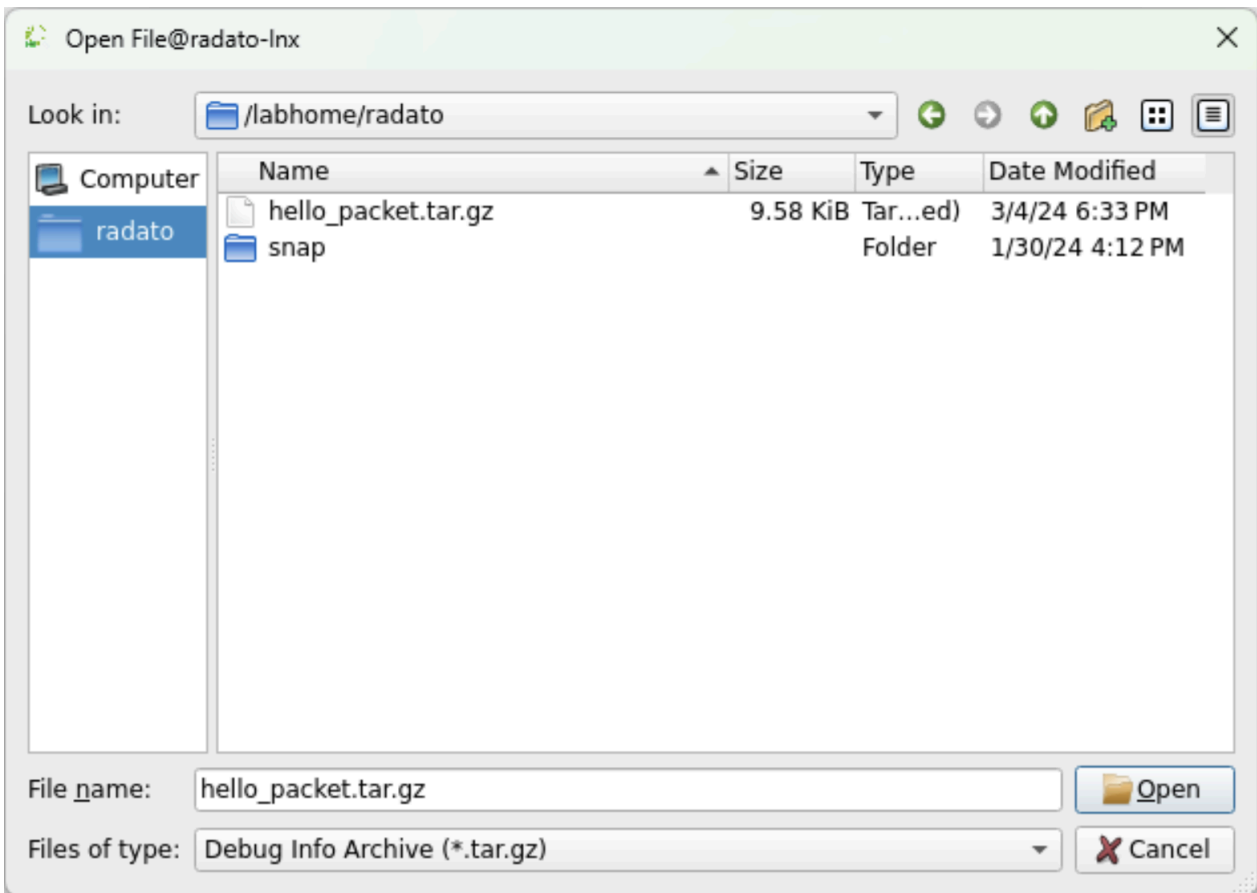
Open Debug Session File

To open a pre-recorded debug session in the DPL Debugger GUI:

1. Select the **File** → **Open Debug Session File**.



2. Choose the appropriate debug session file (`hello_packet.tar.gz` in this example).



Info

The `tar.gz` archive should be the output of a recorded debug session created using the `dpl_nspect debug` command.

Packets List

The left pane of the DPL Debugger displays the **Packets** list, which includes the following fields:

- Date Time – The timestamp when the packet entered the DPL pipeline
- Time Since Start – The elapsed time since the first packet entered the DPL pipeline
- Size – The packet size

- Ingress Port – The port through which the packet entered the pipeline
- Protocol – The network protocol associated with the packet

The screenshot displays the P4Debugger interface with several key components:

- Packets List:** A table showing packet details such as Date Time, Time Since Start (ms), and Size. The selected packet (ID 196) has a size of 109 bytes.
- Debug Points:** A table listing P4 objects and their locations, such as 'Control myprogram.p4:225' and 'Table myprogram.p4:196'.
- P4 Variables:** Sections for 'P4 Variables' and 'Raw Variables' showing values for variables like 'tenant_id' and 'table_d_key'.
- Dissector:** A table showing the protocol stack for the selected packet, including 'ETHERNET', 'IPV4', and 'UDP'.
- Hexdump:** A hex dump of the packet data, showing values like 'a0 88 c2 75 6b 90 00 00 00 00 00 00 00 00 00 00'.
- Pipeline Parser:** A complex flow diagram showing the packet's path through various pipeline stages like 'ip6', 'gre', 'udp', 'vxlan', 'inner_ethernet', and 'inner_vlan'.
- Code Snippets:** A dark-themed window on the right shows C++ code for a 'NvDirectCounter' and 'table' definitions, with a blue highlight on the 'nv_send_debug_pkt()' function call.

Filter

The **Packets** list can be filtered using the fields located at the bottom left of the GUI.



To apply a filter:

1. Enter a condition (e.g., "packet ID is not 0").
2. Click the **+** button to add the filter.
3. Click the **Apply** button to apply the filter.

Packets

| Date Time | Time Since Start (ms) | Size | Ingress Port | Protocol |
|---------------------|-----------------------|------|--------------|----------|
| 2025-01-13 20:55:49 | 0.000000 | 109 | 256 | VXLAN |
| 2025-01-13 20:56:07 | 18408.245433 | 109 | 256 | VXLAN |

Filters

| | | | |
|--------------------------------------|--------|--|---|
| ID | Is Not | 0 |  |
| ID | Is Not | 0 |  |
| <input type="button" value="Apply"/> | | <input type="button" value="Clear All"/> | |

Showing 2 out of 2 packets


To reset the **Packets** list to its original state:

1. Click the **Clear All** button.
2. Click **Apply**.

Parser Graph

The **Parser** graph, located at the bottom of the GUI, outlines the possible routes a packet can take within the DPL pipeline.

- Blue nodes (`inner_ethernet` in the example) indicate the nodes the packet has traversed.
- The graph consists of two types of nodes:
 - Flex nodes – Defined by the DPL programmer
 - Static (fixed) nodes – Built-in components of the hardware



This visualization helps developers understand the packet's path through the parsing stages.

Debug Points

Selecting a packet from the **Packets** List displays its pipeline stages in the Debug Points view at the top of the GUI:

- Each row corresponds to a specific source location in the DPL program code

- The table cells are interactive—clicking a cell highlights the relevant source location, when applicable

| Debug Points | | | |
|--------------|----------|------------------|------------------------------------|
| P4 Object | Location | Table[Entry] | Match Key = Value/Mask (Type) |
| 1 | Control | myprogram.p4:225 | |
| 2 | Table | myprogram.p4:196 | myprogram.tenant[1] |
| | | | headers.vxlan.vni = 0x10 (exact) |
| | | | tenant_id = 0x1 (exact) |
| | | | headers.vlan.vlan_id = 0x0 (exact) |

This allows developers to trace a packet's processing path and correlate it with the corresponding code sections.

Variables

Selecting a pipeline stage (either a single cell or an entire row) displays variable information in two tabs:

- P4 Variables – Shows DPL variables defined in the program and their values at the selected pipeline stage.
- Raw Variables – Provides a low-level view of internal variables, such as registers, samplers, and other hardware-related data.

| P4 Variables | | Raw Variables | |
|--------------|---------------|---|--|
| Name | Value | | |
| 1 | table_d_key 0 | <ul style="list-style-type: none"> ▸ Registers ▸ Samplers Controller Meta 0x00000000 Parser Bitmap 0x00000000 | |
| 2 | session_id 0 | | |
| 3 | tenant_id 0 | | |

This helps developers inspect both high-level program variables and hardware-level details while debugging.

Source Code

Selecting a cell in a pipeline stage opens the DPL program source code in a new tab on the right-hand side of the DPL Debugger GUI. When applicable, the relevant source location is highlighted.

The debugger locates the DPL source file using:

- Search paths defined in the **Settings**
- Manual file selection, if the source file is not found automatically

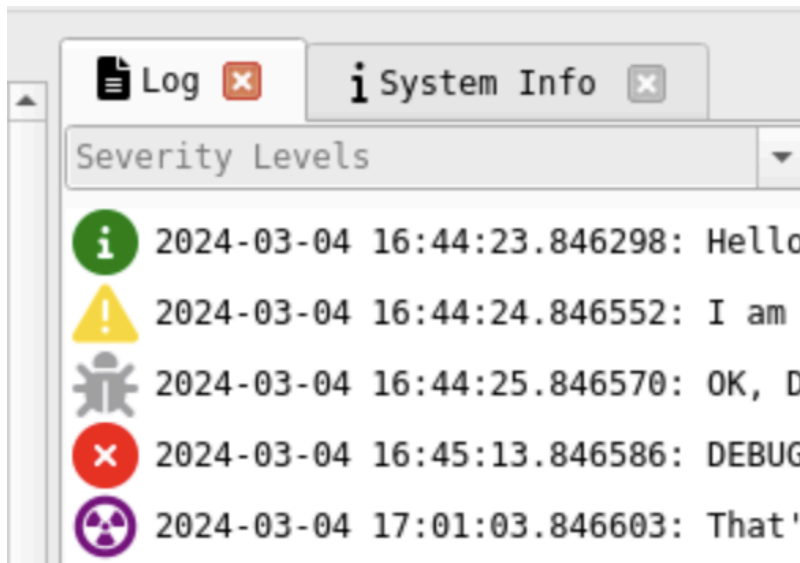
```
Log x i System Info x myprogram.p4 x
214
215     apply {
216         miss_data.session_id = 0;
217         tenant_id = 0;
218         session_id = 0;
219         table_a_key = 0;
220         table_b_key = 0;
221         table_c_key = 0;
222         table_d_key = 0;
223
224         if (headers.vxlan.isValid()) {
225             nv_send_debug_pkt();
226             tenant.apply();
227             if (flow_table.apply().miss) {
228                 route.apply();
229             }
230             session.apply();
231
232             table_a.apply();
233             if (table_b.apply().miss) {
234                 table_a_key = 0xff;
235                 table_a.apply();
236                 if (table_b.apply().miss) {
237                     nv_send_to_controller(miss_d
238                 }
239             }
240             if (table_c.apply().miss) {
241                 table_a_key = 0xffff;
242                 table_a.apply();
243                 if (table_b.apply().miss) {
244                     nv_send_to_controller(miss_d
245                 }
246                 table_c.apply();
247             }
248             table_d.apply();
249         }
250         // On miss the packet goes to the contro
251         nv_send_to_controller(miss_data);
```

Log

The **Log** tab, located on the right side of the GUI, displays internal logs from the DPL Runtime daemon.

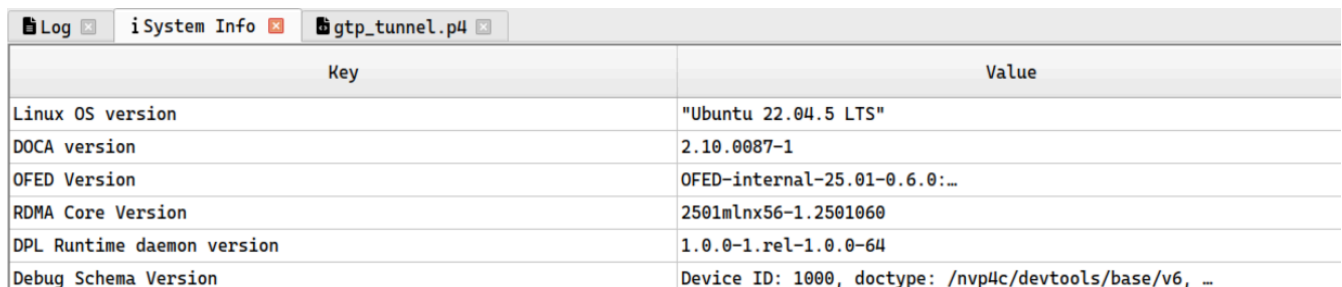
Note

Currently, no log messages are being sent.



System Info

The **System Info** tab, located on the right side of the GUI, displays details about the system where debugging is performed.



The screenshot shows a window titled 'System Info' with a table of system details. The table has two columns: 'Key' and 'Value'. The data is as follows:

| Key | Value |
|----------------------------|--|
| Linux OS version | "Ubuntu 22.04.5 LTS" |
| DOCA version | 2.10.0087-1 |
| OFED Version | OFED-internal-25.01-0.6.0:... |
| RDMA Core Version | 2501mlnx56-1.2501060 |
| DPL Runtime daemon version | 1.0.0-1.rel-1.0.0-64 |
| Debug Schema Version | Device ID: 1000, doctype: /nvp4c/devtools/base/v6, ... |

Packet Dissector

The **Dissector** view, located in the center of the GUI, analyzes a packet's structure when a pipeline stage is selected.

Dissector (displaying 109 bytes)

| Field | Value |
|------------|------------------------------------|
| ▼ ETHERNET | |
| dst_addr | a0:88:c2:75:6b:90 |
| src_addr | 00:00:00:00:00:00 |
| ether_type | Internet Protocol version 4 (IPv4) |
| ▶ IPV4 | |
| UDP | |

Hex Dump

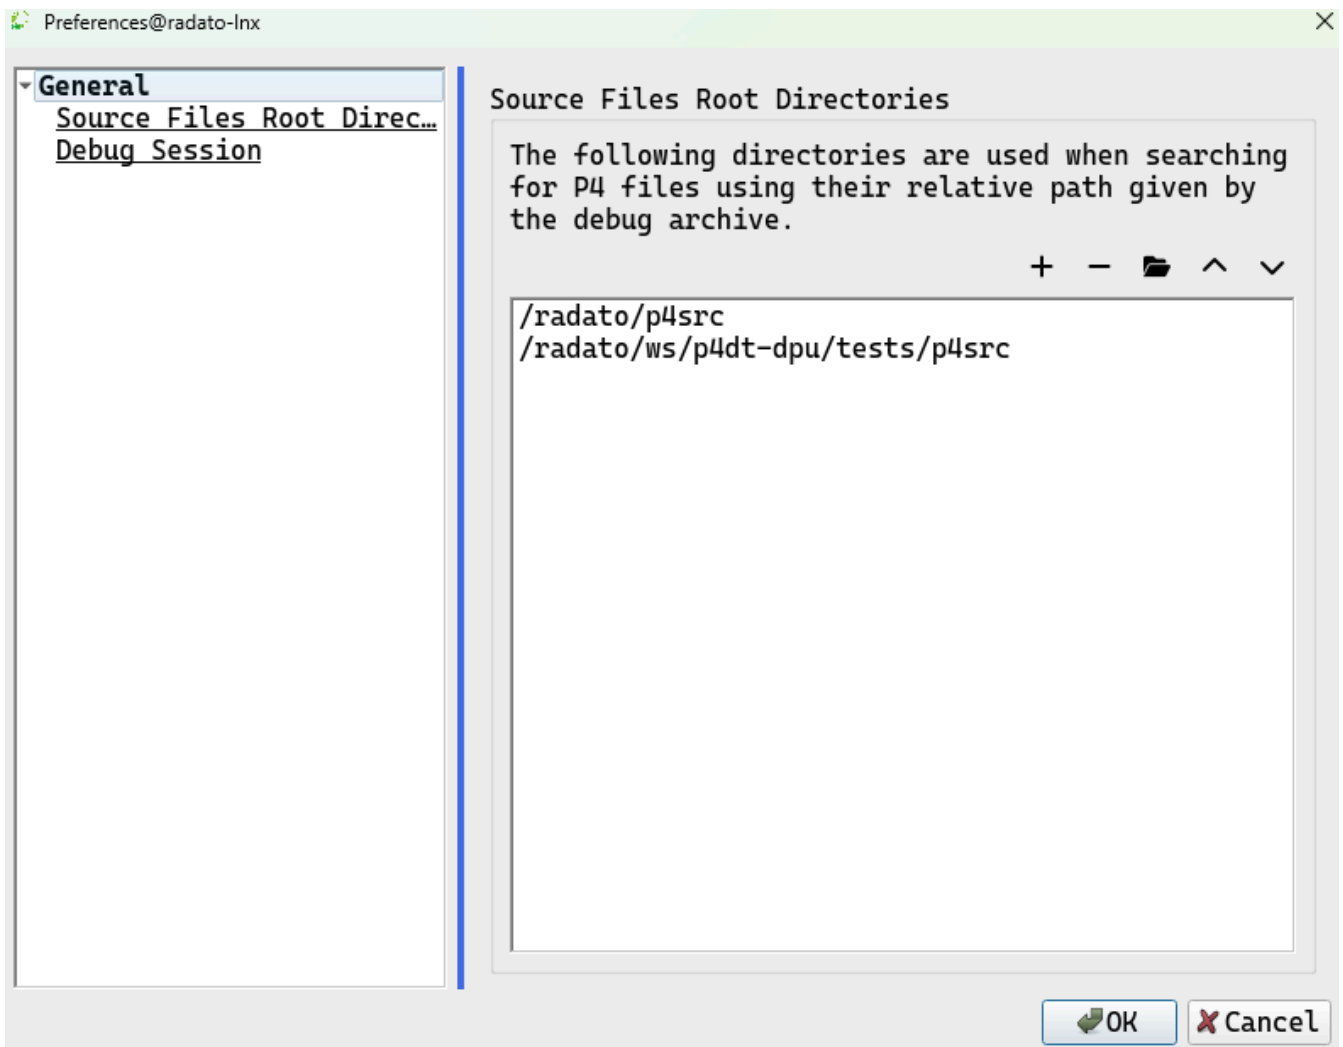
The **Hex Dump** view, located in the center of the GUI, displays a raw hexadecimal representation of the packet when a pipeline stage is selected.

Hexdump (displaying 109 bytes)

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 |
|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|
| a0 | 88 | c2 | 75 | 6b | 90 | 00 | 00 | 00 | 00 | 00 | 00 | 08 | 00 | 45 | 00 |
| 00 | 5f | 00 | 01 | 00 | 00 | 40 | 11 | 7c | 8b | 7f | 00 | 00 | 01 | 7f | 00 |
| 00 | 01 | 12 | b5 | 12 | b5 | 00 | 4b | 00 | d7 | 0c | 00 | 00 | 03 | 00 | 00 |
| 10 | 00 | a0 | 88 | c2 | 75 | 6b | 90 | 40 | b0 | 76 | 0d | b6 | e6 | 81 | 00 |
| 00 | 00 | 00 | 00 | 45 | 00 | 00 | 00 | 00 | 01 | 00 | 00 | 40 | 11 | 7f | 00 |

Options

The **File** → **Preferences** menu opens the **Preferences** window, where users can configure various debugger and debug session settings.



DPL Admin

Introduction

DPL Admin is a CLI tool used for managing and inspecting the DPL Runtime daemon. It provides the following functionalities:

- Viewing the current configuration of the DPL Runtime daemon
- Reading the DPL Runtime Service logs
- Setting the log level for the DPL Runtime Service
- Retrieving the hardware steering dump, capabilities, and NVConfig of the DPL Runtime Service

Options and Commands

Usage:

```
./dpl_admin.sh -a <host:port> [options] <command> [command-  
options]
```

Options:

- `-h`, `--help` – Displays help message and exits
- `-v`, `--version` – Displays `dpl_admin`'s version and exits
- `-a`, `--address` – Required; specifies `dpl_admin`'s server address (refer to agent config file) and port of the template ipv4/6: `[address]:[port]`. Examples:
 - `192.168.0.1:9600`
 - `ipv6:[2607:f8b0:400e:c00::ef]:443`

- `ipv6:[::]:1234`

config

Displays the DPL Runtime Service configuration information, including servers, logging, devices and HAL configurations.

- Usage:

```
./dpl_admin.sh -a <host:port> config
```

- Options:

-

- `-h`, `- help` – Displays command's help message and exit

- Example:

```
./dpl_admin.sh -a 10.1.1.1:9600 config
```

- Output:

```
CONFIGURATION

servers
server_address  tcp_port

[::]            9559
[::]            9560

dpl_rt
dpl_nspect
```

```

admin
[::]          9600

    log                               name:
/var/log/doca/dpl_rtd/agent.log

                                level:
INFO

    devices

controller attr          id    counter cache timeout [msec]
                                interfaces

port_id: 9876          1000          0
name port_id mtu          mac

eth7      4      1514  00:00:00:00:00:00
eth6      3      1514  00:00:00:00:00:00
eth5      2      1514  00:00:00:00:00:00
eth4      1      1514  00:00:00:00:00:00
eth2      0      1514  00:00:00:00:00:00

```

```
hal
number of queues:  1

queue size:       1024

burst size:       32
```

set-counter-cache-timeout

Sets counter cache timeout for device.

- Usage:

```
./dpl_admin.sh -a <host:port> set-counter-cache-timeout --
device_id/-did <DEVICE_ID> --timeout_value/-nto
<TIMEOUT_VALUE>
```

- Options:

- `-h`, `- help` – Displays command's help message and exit
- `--device_id`, `-did` – Required; specifies device ID.
- `--timeout_value`, `-nto` – Required; specifies new counter cache time-out value

- Example:

```
./dpl_admin.sh -a 10.1.1.1:9600 set-counter-cache-timeout --
device_id 1000 --timeout_value 3
```

- Output:

```
Device ID: 1000, new timeout: 3 msec
```

set-log-level

Sets DPL Runtime Service's log level.

- Usage:

```
dpl_admin.sh -a <host:port> set-log-level --level <LOG_LEVEL>
```

- Options:


- `-h`, `- help` – Displays command's help message and exit
- `--level` – Required; new log level out of `{DISABLE, CRITICAL, ERROR, WARNING, INFO, DEBUG, TRACE}`

- Example:

```
./dpl_admin.sh -a 10.1.1.1:9600 set-log-level --level DEBUG
```

- Output:

```
dpl_rtd log level set to: 'DEBUG'
```



Info

Setting log level is not persistent. DPL Admin restarts with log level according to config file.

get-log-full

Reads DPL Runtime Service's log into a file.

- Usage:

```
./dpl_admin.sh -a <host:port> get-log-full
```

- Options:

- `-h`, `- help` – Displays command's help message and exit
- `--outpath` – Optional; specifies output log file path. Defaults to current directory.

- Example:

```
./dpl_admin.sh -a 10.1.1.1:9600 get-log-full --outpath  
'/tmp/dpl_agent.log'
```

- Output:

```
Find Log file at: '/tmp/dpl_agent.log'
```

get-log-recent

Displays recent DPL Runtime Service's log lines.

- Usage:

```
./dpl_admin.sh -a <host:port> get-log-recent
```

- Options:

- `-h`, `- help` – Displays command's help message and exit
- `--number_of_lines` – Optional; specifies number of recent log lines to read. Defaults to 10.

- Example:

```
./dpl_admin.sh -a 10.1.1.1:9600 get-log-recent --  
number_of_lines 6
```

- Output:

```
[12:45:09:480451][73079][DOCA][INF][OnReadDone] New client  
connected to device 1000  
[12:45:09:481295][73078][DOCA][INF] A program was previously  
loaded, clearing it before applying new one...  
[12:45:09:481314][73078][DOCA][INF] [DPL Device 1000]  
Removing DPL program ...  
[12:45:09:498476][73078][DOCA][INF] [DPL Device 1000] DPL  
program was removed.  
[12:45:09:499432][73078][DOCA][INF] [DPL Device 1000] Loading  
DPL program ...  
[12:45:09:522419][73078][DOCA][INF] [DPL Device 1000] DPL  
program was loaded successfully.
```


hws-dump

Dumps DPL Runtime Service's hardware steering rules into a file.

- Usage:

```
./dpl_admin.sh -a<host:port> hws-dump --device_id/-did  
<DEVICE_ID>
```

- Options:


- `-h`, `- help` – Displays command's help message and exit
- `--device_id`, `-did` – Required; specifies device ID
- `--parser_args` – Optional; specifies arguments to pass to HWS dump tool
- `--outpath` – Optional; specifies output file path. Defaults to current directory.

- Example:

```
./dpl_admin.sh -a 10.1.1.1:9600 hws-dump --device_id 1000 --  
parser_args '\-vvv' --outpath '/tmp/dpl_hws_dump.txt'
```

- Output:

```
Find HWS dump file at: '/tmp/dpl_hws_dump.txt'
```



Info

Intended as a debugging aid for developers.

hca-capabilities

Dumps NIC HCA capabilities into a file.

- Usage:

```
./dpl_admin.sh -a<host:port> hca-capabilities --device_id/-  
did <DEVICE_ID>
```

- Options:

- `-h`, `- help` – Displays command's help message and exit
- `--device_id`, `-did` – Required; specifies device ID
- `--outputpath` – Optional; specifies output file path. Defaults to current directory.

- Example:

```
./dpl_admin.sh -a 10.1.1.1:9600 hca-capabilities --device_id  
1000 --outputpath '/tmp/dpl_hca_capabilities.json'
```

- Output:

```
Find HCA capabilities file at:  
'/tmp/dpl_hca_capabilities.json'
```

nv-config

Dumps DPL Runtime Service's NV config into a file.

- Usage:

```
./dpl_admin.sh -a<host:port> nv-config --device_id/-did  
<DEVICE_ID>
```

- Options:

- `-h`, `- help` – Displays command's help message and exit
- `--device_id`, `-did` – Required; specifies device ID
- `--outpath` – Optional; specifies output file path. Defaults to current directory.

- Example:

```
./dpl_admin.sh -a 10.1.1.1:9600 nv-config --device_id 1000 --  
outpath '/tmp/dpl_nv_config.txt'
```

- Output:

```
Find NV config file at: '/tmp/dpl_nv_config.txt'
```

P4 Runtime Controller

DPL applications are deployed to the NVIDIA® BlueField® networking platform (DPU or SuperNIC) via the P4Runtime API.

Since DPL is derived from the P4-16 language, it is compatible with the P4Runtime specification, enabling standard runtime interaction with the compiled DPL pipeline.

Introduction

The P4 Runtime shell (`p4runtime_sh`) is an open-source CLI tool that provides an interface to the P4Runtime API. It is especially useful for:

- Loading simple DPL programs
- Testing match-action tables
- Debugging pipeline behavior

The shell can be invoked using the launch script provided in the DPL Development container.

Info

For detailed instructions, refer to [Loading DPL Applications](#).

p4runtime_sh Usage

The following are example commands for using the p4runtime_sh P4 Controller after loading a program.

P4 info

| Operation | Command |
|--|-------------------|
| Retrieves the content of <code>p4info.txt</code> of the currently loaded DPL program | <pre>p4info</pre> |

P4 Table

| Operation | Command |
|--|--|
| Lists all P4 tables | <pre>tables</pre> |
| Displays information about a specific P4 table | <pre>tables["<P4_TABLE_NAME>"]</pre> |

Working with P4 Table Entries

| Operation | Command |
|--|--|
| Default Entry | |
| Reads P4 table's default entry without counter's value | <pre>te = table_entry[" <P4_TABLE_NAME>"] te.is_default = True te.read(lambda te: print(te))</pre> |

| Operation | Command |
|--|---|
| <p>Reads P4 table's default entry with counter's value</p> <div data-bbox="159 359 993 579" style="background-color: #ffffcc; padding: 10px;"> <p>Note Supported only if a direct counter is enabled on the P4 table.</p> </div> | <pre data-bbox="1015 226 1463 674">te = table_entry[" <P4_TABLE_NAME>"] te.is_default = True te.counter_data.byte count = 0 te.read(lambda te: print(te))</pre> |
| <p>Modifies P4 table's default entry action</p> <div data-bbox="159 905 993 1163" style="background-color: #ffffcc; padding: 10px;"> <p>Note A default entry cannot be removed or inserted. It can only be modified to perform a different P4 action.</p> </div> | <pre data-bbox="1015 703 1463 1325">te = table_entry[" <P4_TABLE_NAME>"] (action=" <P4_ACTION_NAME>") te.is_default = True # Set value for all parameters required by desired action. te.action["<PARAMETER NAME>"] = "<PARAMETER VALUE>" te.modify()</pre> |
| Regular Entries | |
| <p>Note: In the following examples, some commands require specifying the match key of the desired regular P4 table key. The syntax for specifying a match key varies according to the defined match method for each key (per the DPL source code where the keys are defined on the P4 table).</p> | |
| Match Method | Syntax for Specifying Match Key Value |
| exact | <pre data-bbox="316 1732 1451 1877">te.match["<MATCH_KEY_NAME>"] = "<MATCH_VALUE>"</pre> |

| Operation | | Command |
|---|--|---------|
| Match Method | Syntax for Specifying Match Key Value | |
| ternary | <pre>te.match["<MATCH_KEY_NAME>"] = "<MATCH_VALUE>&&&<MASK_VALUE>" te.priority = <PRIORITY VALUE></pre> <p>Note that mask is provided in the match line, separated by &&&. If mask is not specified, a full match mask will be used.</p> | |
| lpm | <pre>te.match["<MATCH_KEY_NAME>"] = "<MATCH_VALUE>/<PREFIX_LENGTH_VALUE>"</pre> <p>Note that LPM prefix_len is provided in the match line, separated by /. If LPM prefix_len is not specified, a prefix_len with full field bitwidth is used.</p> | |
| For simplicity, the following examples are written using <code>exact</code> match syntax. | | |
| Reading Entries | | |
| Reads a specific regular P4 table entry | <pre>te = table_entry[" <P4_TABLE_NAME>"] try: pass # Comment out the next line to disable reading counters te.counter_data.byte e_count = 0 except Exception as e: # Table does not have a Direct Counter pass # Set value for all keys required by the P4 table.</pre> | |

| Operation | Command |
|--|--|
| | <pre>te.match[" <MATCH_KEY_NAME>"] = " <MATCH_VALUE>" te.read(lambda te: print(te))</pre> |
| <p>Reads all regular entries from a P4 table</p> | <pre>num = 1 def hndlr(te): global num print(f">> Entry number {num}:") print(te) print("----- -----") num += 1 te = table_entry[" <P4_TABLE_NAME>"] try: pass # Comment out the next line to disable reading counters te.counter_data.byte_count = 0 except Exception as e: # Table does not have a Direct Counter pass # Read regular entries te.is_default = False</pre> |

| Operation | Command |
|---|---|
| | <pre data-bbox="1073 222 1393 310">te.read(lambda te: hndlr(te))</pre> |
| <p data-bbox="159 394 922 468">Reads all regular entries from all P4 tables in the P4 program</p> | <pre data-bbox="1073 453 1446 1944">for tbl in tables: num = 1 def hndlr(te): global num print(f">> Entry number {num}:") print(te) print("----- -----") num += 1 print(f"===== {tbl.name} =====") te = table_entry[tbl.name] try: pass # Comment out the next line to disable reading counters te.counter_data.byte e_count = 0 except Exception as e: # Table does not have a Direct Counter pass</pre> |

| Operation | Command |
|--|--|
| | <pre data-bbox="1015 212 1463 575"> # Read regular entries te.is_default = False te.read(lambda te: hndlr(te)) </pre> |
| Adding Entries | |
| <p data-bbox="159 1031 589 1066">Adds a regular P4 table entry</p> | <pre data-bbox="1015 638 1463 1457"> te = table_entry[" <P4_TABLE_NAME>"] (action=" <P4_ACTION_NAME>") # Set value for all keys required by the P4 table. te.match[" <MATCH_KEY_NAME>"] = " <MATCH_VALUE>" # Set value for all parameters required by desired action. te.action["<PARAMETER NAME>"] = "<PARAMETER VALUE>" te.insert() </pre> |
| Deleting Entries | |
| <p data-bbox="159 1535 748 1570">Deletes a specific regular P4 table entry</p> | <pre data-bbox="1015 1520 1463 1944"> te = table_entry[" <P4_TABLE_NAME>"] # Set value for all keys required by the P4 table. te.match[" <MATCH_KEY_NAME>"] = " <MATCH_VALUE>" </pre> |

| Operation | Command |
|---|---|
| | <pre>te.delete()</pre> |
| Deletes all regular entries from a P4 table | <pre>te = table_entry[" <P4_TABLE_NAME>"] te.read(lambda te: te.delete())</pre> |

Working with Direct Counters

This allows working with Direct Counter directly without getting the whole Table Entry info.

| Operation | Command |
|--|---|
| Lists all defined Direct Counters | <pre>direct_counters</pre> |
| Default Entry | |
| Reads counter data of a default entry in a Direct Counter | <pre>DirectCounterEntry(" <P4_DIRECT_COUNTER_NAME>") ce.table_entry.is_default = True ce.read((lambda ce: print(ce)))</pre> |
| Clears counter data of a default entry in a Direct Counter | <pre>DirectCounterEntry(" <P4_DIRECT_COUNTER_NAME>")</pre> |

| Operation | Command |
|--|---|
| | <pre>ce.table_entry.is_default = True ce.packet_count = 0 ce.modify()</pre> |
| Regular Entries | |
| <p>Reads counter data of a specific P4 table entry in a Direct Counter</p> | <pre>ce = DirectCounterEntry(" <P4_DIRECT_COUNTER_NAME>") # Set value for all keys required by the P4 table. ce.table_entry.match[" <MATCH_KEY_NAME>"] = " <MATCH_VALUE>" ce.read((lambda ce: print(ce)))</pre> |
| <p>Reads all counter data of specific Direct Counter</p> <div data-bbox="159 1236 841 1499" style="background-color: #ffffcc; padding: 10px; margin-top: 10px;"> <p>i Info This will also read the default entry counter data.</p> </div> | <pre>ce = DirectCounterEntry(" <P4_DIRECT_COUNTER_NAME>") ce.read((lambda ce: print(ce)))</pre> |
| <p>Clears counter data of a a specific P4 table entry in a Direct Counter</p> | <pre>ce = DirectCounterEntry(" <P4_DIRECT_COUNTER_NAME>") # Set value for all keys required by the P4 table. ce.table_entry.match[" <MATCH_KEY_NAME>"] = "</pre> |

| Operation | Command |
|---|--|
| | <pre><MATCH_VALUE>" ce.packet_count = 0 ce.modify()</pre> |
| <p>Clears all counter data of specific Direct Counter</p> <p>Note This also clears the default entry counter data.</p> | <pre>ce = DirectCounterEntry(" <P4_DIRECT_COUNTER_NAME>") ce.packet_count = 0 ce.modify()</pre> |
| <p>Clears all direct counters from all table</p> | <pre>for dc in direct_counters: ce = DirectCounterEntry(dc.name) ce.packet_count = 0 ce.modify()</pre> |

Working with Indirect Counters

| Operation | Command |
|---|---|
| <p>Lists all defined Indirect Counters</p> | <pre>counters</pre> |
| <p>Shows info about a specific Indirect Counter</p> | <pre>counter_entry[" <P4_INDIRECT_COUNTER_NAME>"]</pre> |

| Operation | Command |
|--|--|
| Reads a specific value from a specific indirect counter | <pre> ce = counter_entry[" <P4_INDIRECT_COUNTER_NAME>"] ce.index = <COUNTER_CELL_INDEX> ce.read((lambda ce: print(ce))) </pre> |
| Reads all values from a specific indirect counter | <pre> ce = counter_entry[" <P4_INDIRECT_COUNTER_NAME>"] ce.read((lambda ce: print(ce))) </pre> |
| Clears a specific value from a specific indirect counter | <pre> ce = counter_entry[" <P4_INDIRECT_COUNTER_NAME>"] ce.index = <COUNTER_CELL_INDEX> ce.byte_count = 0 ce.packet_count = 0 ce.modify() </pre> |
| Clears all values from a specific indirect counter | <pre> ce = counter_entry[" <P4_INDIRECT_COUNTER_NAME>"] ce.byte_count = 0 ce.packet_count = 0 ce.modify() </pre> |

P4 Actions

| Operation | Command |
|--|---|
| Lists all defined P4 actions | <pre>actions</pre> |
| Shows info about a specific P4 actions | <pre>actions[" <P4_ACTION_NAME>"]</pre> |

i Info
This includes its parameters' names and sizes.

Packet IO

To use packet IO, it must be enabled in the P4 program source code.

| Operation | Command |
|--|--|
| Packet In – For receiving packets sent from the DPL Runtime daemon to the P4 Controller (according to defined rules in the DPL program) | |
| Captures packets for 10 second, then displays them | <pre>packet_in.sniff(lambda m: print(m), timeout=10)</pre> |
| Capturs packets for 10 second, then displays them parsed as well as their metadata info | <p>This example command uses the <code>impacket</code> package for parsing the packets, so make sure that it is installed on your system prior to running the <code>p4runtime_sh</code> P4 Controller:</p> <pre>pip install impacket</pre> <p>Then, from the <code>p4runtime_sh</code> P4 Controller, run:</p> |

| Operation | Command |
|---|---|
| | <pre> from impacket.ImpactDecoder import * for msg in packet_in.sniff(timeout=10): print("-----") print(msg) print("+++++") print("Raw packet (hex):\n") print(msg.packet.payload.hex()) print("+++++") print("Parsed packet:\n") print(EthDecoder().decode(msg.packet.payload)) print("+++++") print("Metadata:\n") for md in msg.packet.metadata: val = int.from_bytes(md.value, "big") print("metadata ID (" + md.metadata_id + "):", hex(val)) </pre> |
| <p>Packet Out – For sending a packet from the <code>p4runtime_sh</code> P4 Controller to the DPL Runtime daemon (which will process it according to defined rules in the DPL program)</p> | |
| <p>Syntax of sending a packet</p> | <pre> my_pkt = b'<PACKET_BYTE_STRING_MESSAGE_GOES_HERE>' </pre> |

| Operation | Command |
|---|---|
| | <pre>packet_out(payload=my_pkt, <METADATA_NAME>="METADATA_VALUE_GOES_HERE") . send() Repeat the <METADATA_NAME>= "METADATA_VALUE_GOES_HERE " parameter for each defined metadata in the P4 program.</pre> |
| <p>Example of building and sending a packet</p> | <p>Using <code>scapy</code> tool, build the desired packet:</p> <pre>>>> p = Ether(src='00:11:11:11:11:11', dst='00:22:22:22:22:22') / IP(src="1.1.1.1", dst="2.2.2.2") >>> print(p.build()) b'\x00\x00\x00\x00\x11\x11\x11\x11\x11\x11\x08\x00E\x00\x00\x14\x00\x01\x00\x00@\x00t\xe4\x01\x01\x01\x01\x02\x02\x02\x02' >>></pre> <p>Then send it using the <code>p4runtime_sh</code> P4 Controller:</p> <pre>my_pkt = b'\x00\x00\x00\x00\x11\x11\x11\x11\x11\x11\x08\x00E\x00\x00\x14\x00\x01\x00\x00@\x00t\xe4\x01\x01\x01\x01\x02\x02\x02\x02' packet_out(payload=my_pkt, controller_metadata="0x123").send()</pre> |

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