



NVIDIA DOCA URL Filter

Reference App

Table of Contents

Chapter 1. Introduction.....	1
Chapter 2. System Design.....	2
Chapter 3. Application Architecture.....	5
Chapter 4. Configuration Flow.....	7
Chapter 5. Running Application on BlueField.....	8
Chapter 6. Running Application on Host.....	10
Chapter 7. References.....	11

Chapter 1. Introduction

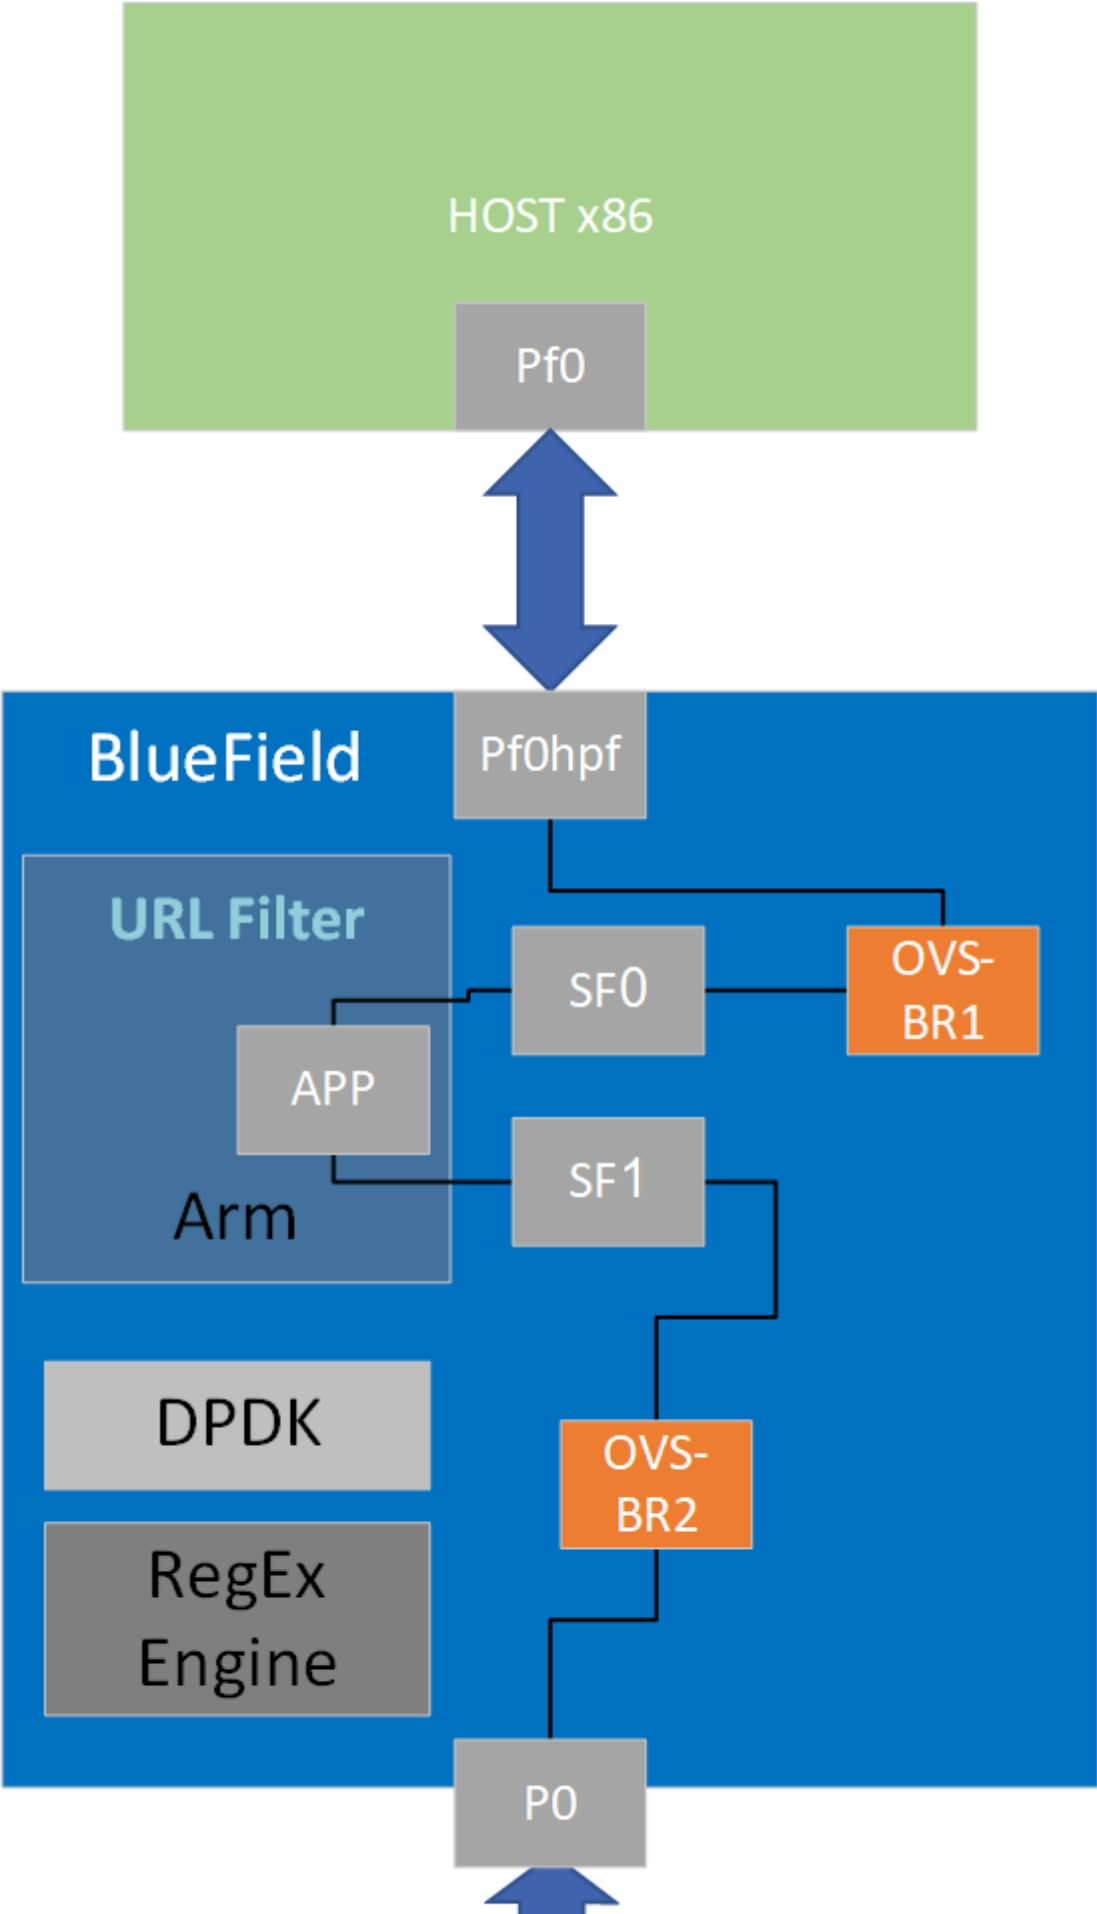
URL filtering limits access by comparing web traffic against a database to prevent user from different threats, malware and accessing harmful sites such as phishing pages.

This kind of content filtering can increase network security and enforce policies on different network resources.

Chapter 2. System Design


URL filtering is designed to run as "bump-on-the-wire" on the BlueField-2 instance. It intercepts traffic coming from the wire and passes it to the physical function (PF) representor connected to the host.

URL filter utilizes the SFT and RegEx engines which are HW accelerators on BlueField.



Chapter 3. Application Architecture

- ▶ User adds a URL using the URL command from the CLI. Available commands:
 - ▶ Create database – create DB of URLs to be filtered (also removes old signatures)
 - ▶ Filter http [MSG] [RegEx] – add specific URL/URL RegEx pattern to be filtered
 - ▶ MSG – user information for the URL filter
 - ▶ RegEx – can hold URL or RegEx pattern
 - ▶ Commit database [PATH] – compile and load signature database
- ▶ Ingress traffic is identified using the connection tracking module
- ▶ Traffic is scanned against compiled signature DB
- ▶ Post-processing is performed for match decision
- ▶ Matched flows are filtered and traffic is blocked



1. Signatures are compiled by the DPI compiler and are then loaded to the DPI engine.
2. Ingress traffic is identified and classified using the stateful table module in the DPDK libs which utilizes the connection tracking hardware offloads. This allows flow classifications

to be done at the hardware level and forwarded to a hairpin queue without being processed by the software, which increases performance dramatically.

3. Traffic is scanned against the DPI engine compiled signature DB.
4. Post-processing is performed for match decision.
5. Matched flows are identified and can be offloaded to the hardware to increase performance as not further inspection is needed.
6. Flow termination is done by the aging timer (set in the SFT to 60 seconds). When a flow is offloaded it cannot be tracked and destroyed.
7. Is it important to note that only sites that support non-encrypted HTTP traffic can be matched against signatures created by the URL filtering as it specifically targets the URI field in the HTTP request.

Chapter 4. Configuration Flow

1. DPDK initialization.

```
dpdk_init(&argc, &argv, &nb_queues);
```

2. DPDK port initialization.

```
dpdk_ports_init(nb_queues, nb_ports);
```

- a). Mempool allocation
- b). Port initialization

3. Stateful Flow Table (SFT).

```
dpdk_sft_init(ar_config.ct, nb_queues, nb_ports);
```

- a). SFT initialization with the configurable parameters (CT on/off)
- b). Configure RTE flow to forward non-L4 traffic to hairpin queue and L4 traffic to SFT
- c). Configure RTE flow to forward matched traffic to hairpin queue

4. DPI initialization.

```
doca_dpi_init(&doca_dpi_config);
```

- a). Configure RegEx engine
- b). Configure DPI queues

5. Configure DPI packet processing.

```
dpi_worker_lcores_run(nb_queues, CLIENT_ID, url_filter_worker_attr);
```

- a). Configure DPI enqueue packets
- b). Send jobs to RegEx engine
- c). Configure DPI dequeue packets

6. Add URL shell command.

- a). Create database
- b). Filter http
- c). Commit database

7. DPI destroy

```
doca_dpi_destroy(dpi_ctx);
```

- ▶ Free DPI resources

Chapter 5. Running Application on BlueField

1. Please refer to the [DOCA Installation Guide](#) for details on how to install BlueField related software.
2. To build the application

- a). The URL filtering example is installed as part of the `doca-dpi-lib` package. The binary is located under `/opt/mellanox/doca/examples/url_filter/bin/doca_url_filter`. To re-build the URL filter sample, run the following command:

```
cd /opt/mellanox/doca/examples/url_filter/src
meson /tmp/build
ninja -C /tmp/build
```

`doca_url_filter` will be created under `tmp/build`.

- b). The build process depends on the `PKG_CONFIG_PATH` environment variable to locate the DPDK libraries. If the variable was accidentally corrupted, and the build fails, run the following command:

- ▶ For Ubuntu:

```
export PKG_CONFIG_PATH=$PKG_CONFIG_PATH:/opt/mellanox/dpdk/lib/aarch64-linux-gnu/pkgconfig
```

- ▶ For CentOS:

```
export PKG_CONFIG_PATH=$PKG_CONFIG_PATH:/opt/mellanox/dpdk/lib64/pkgconfig
```

- c). The URL filter example is based on DPDK libraries. Therefore, the user is required to provide DPDK flags, and allocate huge pages. Run:

```
echo 2048 > /sys/kernel/mm/hugepages/hugepages-2048kB/nr_hugepages
sudo mkdir /mnt/huge
sudo mount -t hugetlbfs nodev /mnt/huge
systemctl start mlx-regex
```

3. To run the application:

```
doca_url_filter [dpdk flags] -- --print_match
```

Sub-functions need to be enabled according to the [Scalable Function Setup Guide](#).

For example:

```
/opt/mellanox/doca/examples/url_filter/bin/doca_url_filter -a
0000:03:00.0,class=regex -a auxiliary:mlx5_core.sf.4,sft_en=1 -a
auxiliary:mlx5_core.sf.5,sft_en=1 -c3 -- -p
```



Note: The flag `-a 0000:03:00.0,class=regex -a auxiliary:mlx5_core.sf.4,sft_en=1 -a auxiliary:mlx5_core.sf.5,sft_en=1`

is necessary for proper usage of the application. Modifying this flag will result in unexpected behavior as only 2 ports are supported. The subfunction number is arbitrary and configurable. The RegEx device, however, is not and must be initiated on port 0.



Note: URL filter will only work when using a single worker thread on Ubuntu systems.

To print the output when DPI engine finds a match, use `--print_match` or `-p` flag.

For additional information on available flags for DPDK, use `-h` before the `--` separator and for the app, use `-h` after `--`.

Application flags:

- ▶ `-t` or `--connection_tracking` – enable SFT connection tracking (may decrease performance)
- ▶ `-l` or `--log_level` – set the log level for the app (ERR=0, DEBUG=3)

If the `print_match` flag is used, a print will be posted to the screen whenever a match is found.

The URL filter is based on user interaction with shell commands. Using the TAB key allows autocompletion while the `quit` command terminates the application. Other available commands are as follows:

- ▶ `create database` – removes and creates a new signature database at `/tmp/signature.txt` so it can be used in the filter command
- ▶ `filter http [msg] [regex]` – a signature containing the regular expression written is created in the database. When a match is found, a message is printed.
- ▶ `commit database [path]` – compiles and loads the signatures created by the filter command from the file path provided. The default database is `/tmp/signature.txt`.



Note: To load a signatures file that was created beforehand, simply run `commit database` with the desired path to load the file.



Note: For instructions on running the application on the host, please refer to [Running Reference Applications Over Host Guide](#).

Chapter 6. Running Application on Host

Please refer to [Running Reference Applications Over Host Guide](#).

Chapter 7. References

- ▶ `/opt/mellanox/doca/examples/url_filter/src/url_filter.c`

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation nor any of its direct or indirect subsidiaries and affiliates (collectively: "NVIDIA") make no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assume no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF 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, the NVIDIA logo, and Mellanox are trademarks and/or registered trademarks of Mellanox Technologies Ltd. and/or NVIDIA Corporation in the U.S. and in other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2021 NVIDIA Corporation & affiliates. All rights reserved.