



NVIDIA DOCA File Scan

Application Guide

Table of Contents

- Chapter 1. Introduction..... 1
- Chapter 2. System Design..... 2
- Chapter 3. Application Architecture..... 5
- Chapter 4. DOCA Libraries..... 6
- Chapter 5. Configuration Flow..... 7
- Chapter 6. Running Application..... 8
- Chapter 7. Arg Parser DOCA Flags..... 10
- Chapter 8. References..... 12

Chapter 1. Introduction

BlueField DPU supports high-speed hardware RegEx acceleration using the RegEx engine. This allows accelerating different security and networking applications which require RegEx processing.

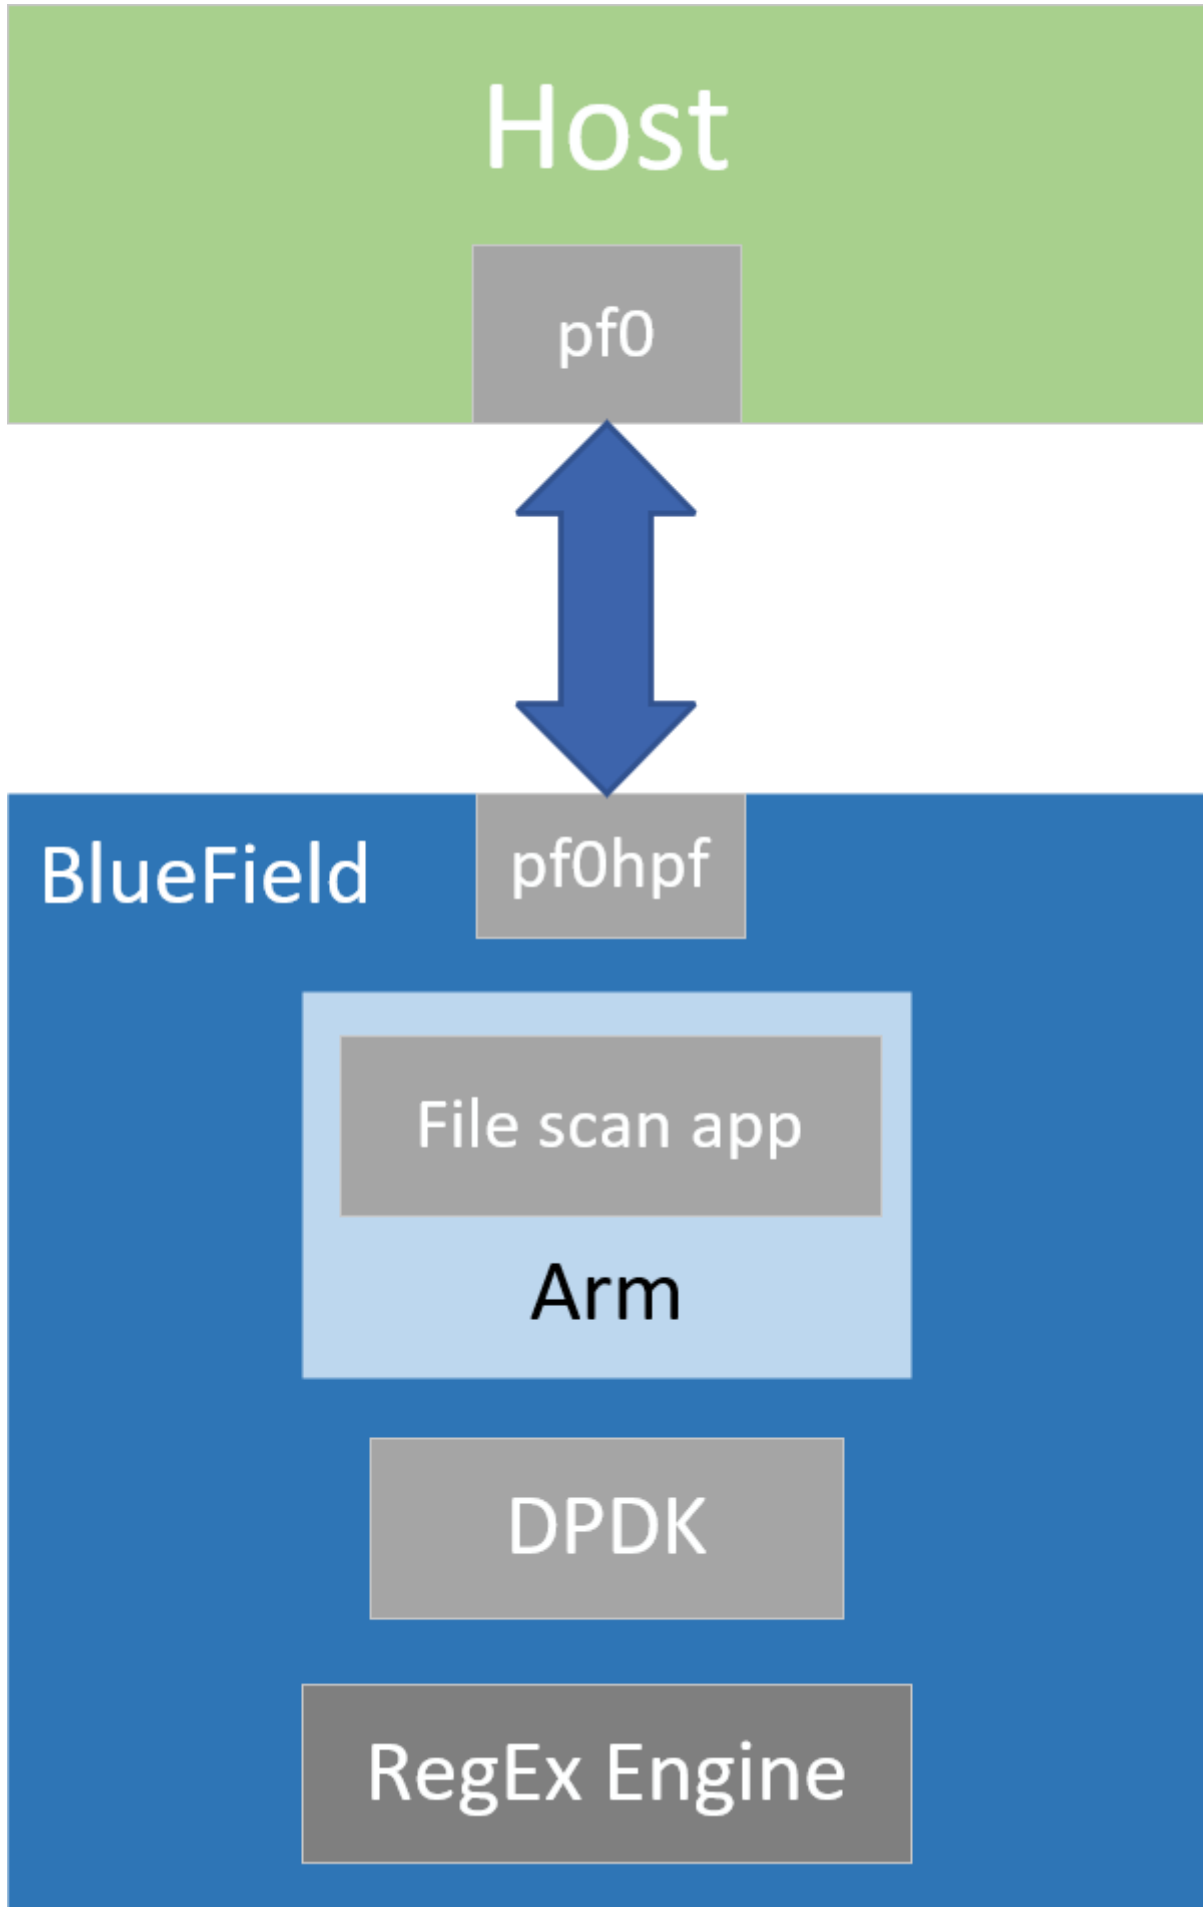
The file scan application uses the hardware RegEx engine to scan a file and returns whether there is a match or not. This application can be a base for many security verifications for sensitive data and personal information.

File scan logic includes:

- ▶ Compiling RegEx rules using Regex compiler
- ▶ Loading a compiled set of RegEx rules to the BlueField RegEx engine
- ▶ Sending file to the RegEx for scanning
- ▶ HW scanning file and returning a list of matches (e.g., matching rule, offset, length)

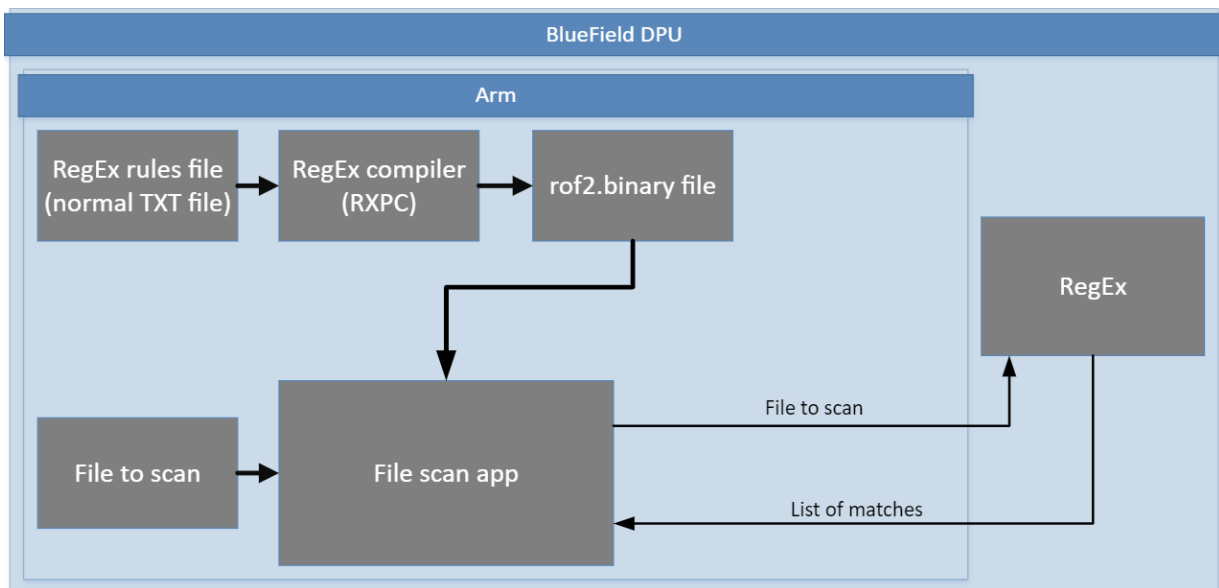
Chapter 2. System Design

File scan utilizes the RegEx engine which is a hardware accelerator on the BlueField.



Chapter 3. Application Architecture

The file scan application runs on top of DOCA RegEx API to configure, send, and receive scans to and from the BlueField RegEx engine.



1. RegEx rules file is compiled into a `rof2.binary` file by the user.
2. The RegEx binary rules file is loaded into the RegEx engine.
3. The file to scan is sent to the RegEx triggering the engine to start matching
4. The results are sent back to the application.



Note: It is the application's responsibility to check for finished jobs. In this example, you Busy-Wait until the RegEx finishes scanning.

5. The matches are printed on the screen, and optionally to a CSV file (which should input to the application, see [Running Application](#)).

Chapter 4. DOCA Libraries

This application leverages the [DOCA RegEx Library](#).

Chapter 5. Configuration Flow

1. Parse application argument.

```
doca_arg_init();
```

- a). Initialize Arg Parser resources.
- b). Register DOCA general flags.

```
register_file_scan_params();
```
- c). Register file scan application flags.

```
doca_arg_start();
```
- d). Parsing APP flags.

2. Initialize file scan resources.

```
file_scan_init();
```

- a). Initialize DOCA RegEx.
- b). Allocate memory pool and QP.
- c). Start the RegEx engine.

3. Run the application:

```
file_scan_run();
```

- a). Initialize jobs.
- b). Enqueue jobs to the RegEx engine.
- c). Dequeue jobs from the RegEx engine and print the results if found.

4. Clean up the File Scan app.

```
file_scan_cleanup();
```

- a). Stop and free RegEx resources.
- b). Free application resources.

Chapter 6. Running Application

1. Refer to the following documents:
 - ▶ [NVIDIA DOCA Installation Guide for Linux](#) for details on how to install BlueField-related software.
 - ▶ [NVIDIA DOCA Troubleshooting Guide](#) for any issue you may encounter with the installation, compilation, or execution of DOCA applications.
2. The application binary is located under `/opt/mellanox/doca/applications/file_scan/bin/doca_file_scan`. To build all the applications together, run:

```
cd /opt/mellanox/doca/applications/  
meson build  
ninja -C build
```
3. To build the file scan application only:
 - a). Edit the following flags in `/opt/mellanox/doca/applications/meson_options.txt`:
 - ▶ Set `enable_all_applications` to `false`
 - ▶ Set `enable_file_scan` to `true`
 - b). Run the commands in step 2.



Note: `doca_file_scan` is created under `./build/file_scan/src/`.

Application usage:

```
Usage: doca_file_scan [DOCA Flags] [Program Flags]  
DOCA Flags:  
-h, --help                Print a help synopsis  
-v, --version             Print program version information  
-l, --log-level           Set the log level for the program <CRITICAL=20,  
ERROR=30, WARNING=40, INFO=50, DEBUG=60>  
Program Flags:  
-r, --rules <path>       Path to compiled rules file (rof2.binary)  
-d, --data <path>        Path to data file  
-p, --pci_addr <PCI-ADDRESS> RegEx device PCI address  
-o, --output-csv <path>  Path to the output of the CSV file  
-c, --chunk-size <Number of Bytes> Chunk size of each job sent to the regex,  
use 0 to send the file as 1 chunk.
```

```
-s, --overlap-size <Number of Bytes> Number of bytes of overlap to use for huge jobs
```

Note: The RegEx accelerator must explicitly be initialized on p0.

Note: For additional information on the application, use the `-h` option:
`/opt/mellanox/doca/applications/file_scan/bin/doca_file_scan -h`

4. Running application on BlueField:

- ▶ To run the application, the RegEx-compiled rule files must be supplied to the app. These files usually end with `*.rof2.binary`. To compile the example rules file, run:

```
cd /opt/mellanox/doca/applications/file_scan/bin/
rxpc -f regex_rules.txt -p 0.01 -o /tmp/regex_rules
```

The results of the `rxpc` are written to `/tmp/` directory, each file with the prefix `regex_rules`.

Note: For more information, refer to [NVIDIA RXP Compiler](#).

- ▶ Pre-run setup.

- The file scan example is based on the DOCA RegEx library. Therefore, the user is required to allocate huge pages to allow activation of the RegEx engine.

```
sudo echo 2048 > /sys/kernel/mm/hugepages/hugepages-2048kB/nr_hugepages
```

- Make sure the RegEx engine is active:

```
systemctl status mlx-regex
```

If the status is inactive (Active: failed), run:

```
systemctl start mlx-regex
```

- Copy the example data file to the `tmp` directory:

```
cp /opt/mellanox/doca/applications/file_scan/bin/data_to_scan.txt /tmp/
data_to_scan.txt
```

- ▶ CLI example for running the application on BlueField assuming the PCIe address of the DPI is `03:00`:

```
cd /opt/mellanox/doca/applications/file_scan/bin/
./doca_file_scan --pci-addr 03:00.0 --rules /tmp/regex_rules.rof2.binary --
data data_to_scan.txt
```

5. Running the application on the host, CLI example:

```
cd /opt/mellanox/doca/applications/file_scan/bin/
./doca_file_scan --pci-addr 21:00.0 --rules /tmp/regex_rules.rof2.binary --data
data_to_scan.txt
```

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

6. To run `doca_file_scan` using a JSON file:

```
doca_file_scan --json [json_file]
```

For example:

```
cd /opt/mellanox/doca/applications/file_scan/bin
./doca_file_scan --json file_scan_params.json
```

Note: The file scan application prints the results to the standard output.

Chapter 7. Arg Parser DOCA Flags

For more information, refer to [NVIDIA DOCA Arg Parser User Guide](#).

Flag Type	Short Flag	Long Flag/JSON Key	Description	JSON Content
General Flags	l	log-level	Set the log level for the application: <ul style="list-style-type: none"> ▶ CRITICAL=20 ▶ ERROR=30 ▶ WARNING=40 ▶ INFO=50 ▶ DEBUG=60 	<code>"log-level": 60</code>
	v	version	Print program version information	N/A
	h	help	Print a help synopsis	N/A
Program Flags	r	rules	Path to precompiled rules file (rof2.binary) <div style="border: 1px solid gray; padding: 5px; margin-top: 10px;"> <p>Note: rules is mandatory flag.</p> </div>	<code>"rules": "/tmp/regex_rules.rof2.binary"</code>
	d	data	Path to data file <div style="border: 1px solid gray; padding: 5px; margin-top: 10px;"> <p>Note: data is</p> </div>	<code>"data": "/tmp/data_to_scan.txt"</code>

Flag Type	Short Flag	Long Flag/JSON Key	Description	JSON Content
			mandatory flag.	
	p	pci-addr	RegEx device PCIe address	"pci-addr": "03:00.0"
	o	output-csv	Path to the output CSV file	"output-csv": "/tmp/ results.csv"
	s	overlap	Number of bytes of overlap to use for huge jobs	"overlap-size": 0
	c	chunk-size	Chunk size of each job sent to the RegEx	"chunk-size": 0

Chapter 8. References

- ▶ `/opt/mellanox/doca/applications/file_scan/src/file_scan.c`
- ▶ `/opt/mellanox/doca/applications/file_scan/src/regex_rules.txt`
- ▶ `/opt/mellanox/doca/applications/file_scan/src/data_to_scan.txt`

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 NON-INFRINGEMENT, 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. The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a world-wide basis. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2022 NVIDIA Corporation & affiliates. All rights reserved.