



NVIDIA DOCA DPI Compiler

User Guide

Table of Contents

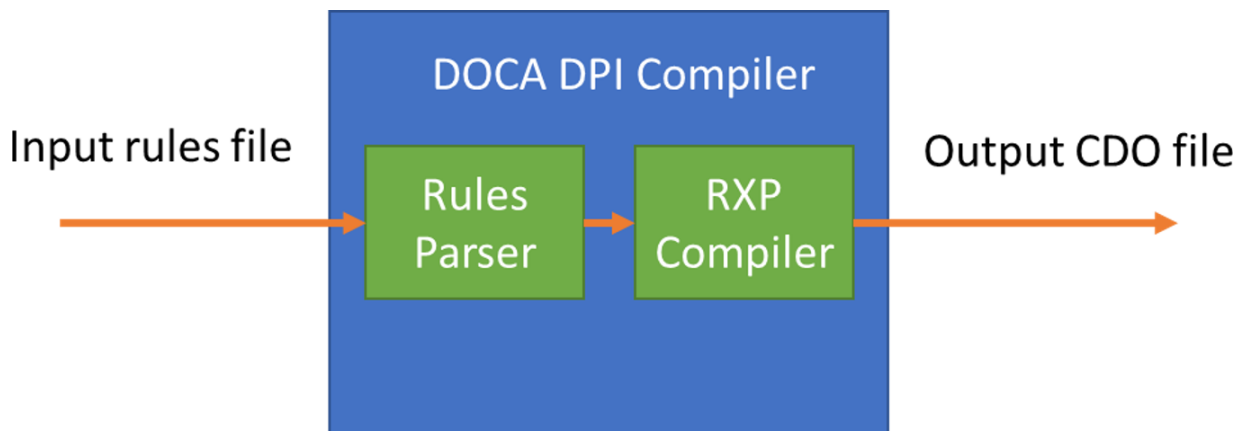
Chapter 1. Introduction.....	1
Chapter 2. Description.....	2
Chapter 3. Execution.....	4
Chapter 4. Arg Parser DOCA Flags.....	5

Chapter 1. Introduction



Important: No updates were made to the DOCA DPI Compiler tool in DOCA 2.2. Please refer to DOCA 2.5 for a note regarding future updates.

DPI compiler is used to compile a signature file which will be loaded into the BlueField RegEx HW accelerator using a dedicated API (`doca_dpi_load_signatures(cdo_file)`). The output for the DPI compiler is a JSON-based CDO file.



This diagram contains the following elements:

- ▶ Input rules file – file containing a list of rules
Supported formats:
 - ▶ Suricata (limited support—see section "[Description](#)" for more information)
 - ▶ Third party signature files (for more information, please contact NVIDIA Support at: networking-support@nvidia.com).
- ▶ RXP compiler – regular expression compiler which is part of the DPI compiler
- ▶ Output CDO file – compiled signature output from compiler. File containing list of rules to be loaded to the RegEx engine.

Chapter 2. Description

The following is a Suricata syntax example:

```
alert udp any any -> any 53 (msg:"ET DOS DNS BIND 9 Dynamic Update DoS attempt";  
content:"|00 00 06|"; offset:8; content:"|c0 0c 00 ff|"; sid:2009701; )
```

Suricata rules are made of three parts – action, header, and options. Take the following rule for example:

```
drop tcp 1.1.1.1/11 8080 -> 2.2.2.2/22 9090 (msg:"ET TROJAN Likely Bot Nick in IRC  
(USA +..)"; flow: to_server; content:"NICK "; nocase; pcre:"/NICK .*USA.*[0-9]{3,}/  
iI"; sid:2008124;)
```

The following table breaks down the rule syntax:

Part	Rule Element
Action	drop
Header	tcp any any -> any any 1.1.1.1/11 8080 -> 2.2.2.2/22 9090
Options	(msg:"ET TROJAN Likely Bot Nick in IRC (USA +..)"; flow: to_server; content:"NICK "; nocase; pcre:"/NICK .*USA.*[0-9]{3,}/i"; sid:2008124;) Detailed explanation can be found in Suricata's official documentation .

DPI compiler has limited support for Suricata format. Field support is detailed in the following table:

Part	Keyword	Support
Action	any	Accept any value. The application decides what to do with each action value.
Header	Protocol	Possible values: tcp, udp
Header	Port	Single port number or range (ex. 100:150)
Header	Address	Single IP address and netmask (ex. 1.2.3.4/5)
Header	Direction	<ul style="list-style-type: none">▶ -> - unidirectional: client-to-server OR server-to-client (default is client-to-server)▶ <> - bidirectional: client to server and server to client

Part	Keyword	Support
Option	msg	Any string – use to identify the rule in case of match on the DPI
Option	http. (must come before content option)	uri, cookie, user_agent, host, referer, content_type, request_body, response_body
Option	http_ (must come after content option)	uri, cookie, user_agent, host, referer, content_type, request_body, response_body
Option	content modifiers	nocase, startswith, endswith, offset Negation (!) for content is not supported
Option	PCRE modifiers	<ul style="list-style-type: none"> ▶ i – PCRE is case insensitive ▶ s – PCRE checks new line characters ▶ A – A pattern must match at the beginning of a buffer ▶ I – Makes PCRE match on the HTTP-raw-uri ▶ C – Makes PCRE match on the HTTP-cookie ▶ V – Makes PCRE match on the HTTP-User-Agent ▶ W – Makes PCRE match on the HTTP-Host ▶ P – Makes PCRE match on the HTTP-Content for CTS ▶ Q – Makes PCRE match on the HTTP-Content for STC
Option	ip_proto	Supported values: 1 (ICMP), 6 (TCP), 17 (UDP)
Option	flow	from_client, to_client, from_server, to_server
Option	tls.	sni – server name indication
Option	dns.	query
Option	sid	Any number – signature's ID
Option	gid	Any number – group ID that can further differentiate signatures according to Suricata rules



Note: If content or PCRE are provided without a field (e.g., `http_uri`), content/PCRE is searched on TCP/UDP payload according to the header protocol.

Chapter 3. Execution

To execute the compiler:

```
Usage: doca_dpi_compiler [DOCA Flags] [Program Flags]
DOCA Flags:
  -h, --help           Print a help synopsis
  -v, --version        Print program version information
  -l, --log-level      <CRITICAL=20, ERROR=30, WARNING=40, INFO=50, DEBUG=60>

Program Flags:
  -i, --input          Full path to signature file
  -f, --format         Signature file format
  -o, --output         Full path to CDO file
  -c, --classtype     Path to the classtype file for priority in suricata
                      file
  -F, --force_compile Ignore any rules containing regex not compatible with
                      hardware. Skipped rules and reasons are output to a
                      .uncompiled.log file.
```

For example:

```
doca_dpi_compiler -i ./suricata_rules.rules -o compiler_output.cdo -f suricata
```

To run doca_dpi_compiler using a JSON file:


```
doca_dpi_compiler --json [json_file]
```

For example:

```
doca_dpi_compiler --json /tmp/doca_dpi_compiler.json
```

Chapter 4. Arg Parser DOCA Flags

Refer to [NVIDIA DOCA Arg Parser Programming Guide](#) for more information.

Flag	Description	JSON Content
<code>-h / --help</code>	Print all information about the available options	N/A
<code>-i / --input</code>	Path to input files	<code>"input" : "/tmp/signatures.rules"</code>
<code>-o / --output</code>	Path to output file	<code>"output" : "/tmp/signatures.cdo"</code>
<code>-f / --format</code>	Signatures format	<code>"format" : "suricata"</code>
<code>-p / --print_rules-c / --classtype</code>	<p>Add the RegEx rules being passed to the RXPC to the CDO as well (useful for debug). Path to a classtype file containing information for priority in suricata format.</p> <p> Note: Needed only for input files that contain a <code>classtype</code> field in the rules.</p>	<code>"classtype" : "/tmp/classtype_config.txt"</code>
<code>-v / --version</code>	Print the current DOCA version in use	N/A

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

© 2023 NVIDIA Corporation & affiliates. All rights reserved.