NVIDIA DOCA File Compression

Application Guide

MLNX-15-060553 _v1.5.2 | July 2023

Table of Contents

Chapter 1. Introduction..... 1
Chapter 2. System Design..... 2
Chapter 3. Application Architecture..... 4
Chapter 4. DOCA Libraries..... 7
Chapter 5. Configuration Flow..... 8
Chapter 6. Running Application..... 9
Chapter 7. Arg Parser DOCA Flags..... 11
Chapter 8. References..... 13

Chapter 1. Introduction

The file compression application exhibits how to use the DOCA Compress API to compress and decompress data using hardware acceleration as well as sending and receiving it using DOCA Comm Channel API.

The application's logic includes both a client and a server:

- Client side – the application opens a file, compresses it, and sends to the server the checksum of the source file alongside the compressed data
- Server side – the application saves the received file in a buffer, decompresses it, and compares the received checksum with the calculated one

Chapter 2. System Design

The file compression application client runs on the host and the server runs on the DPU.

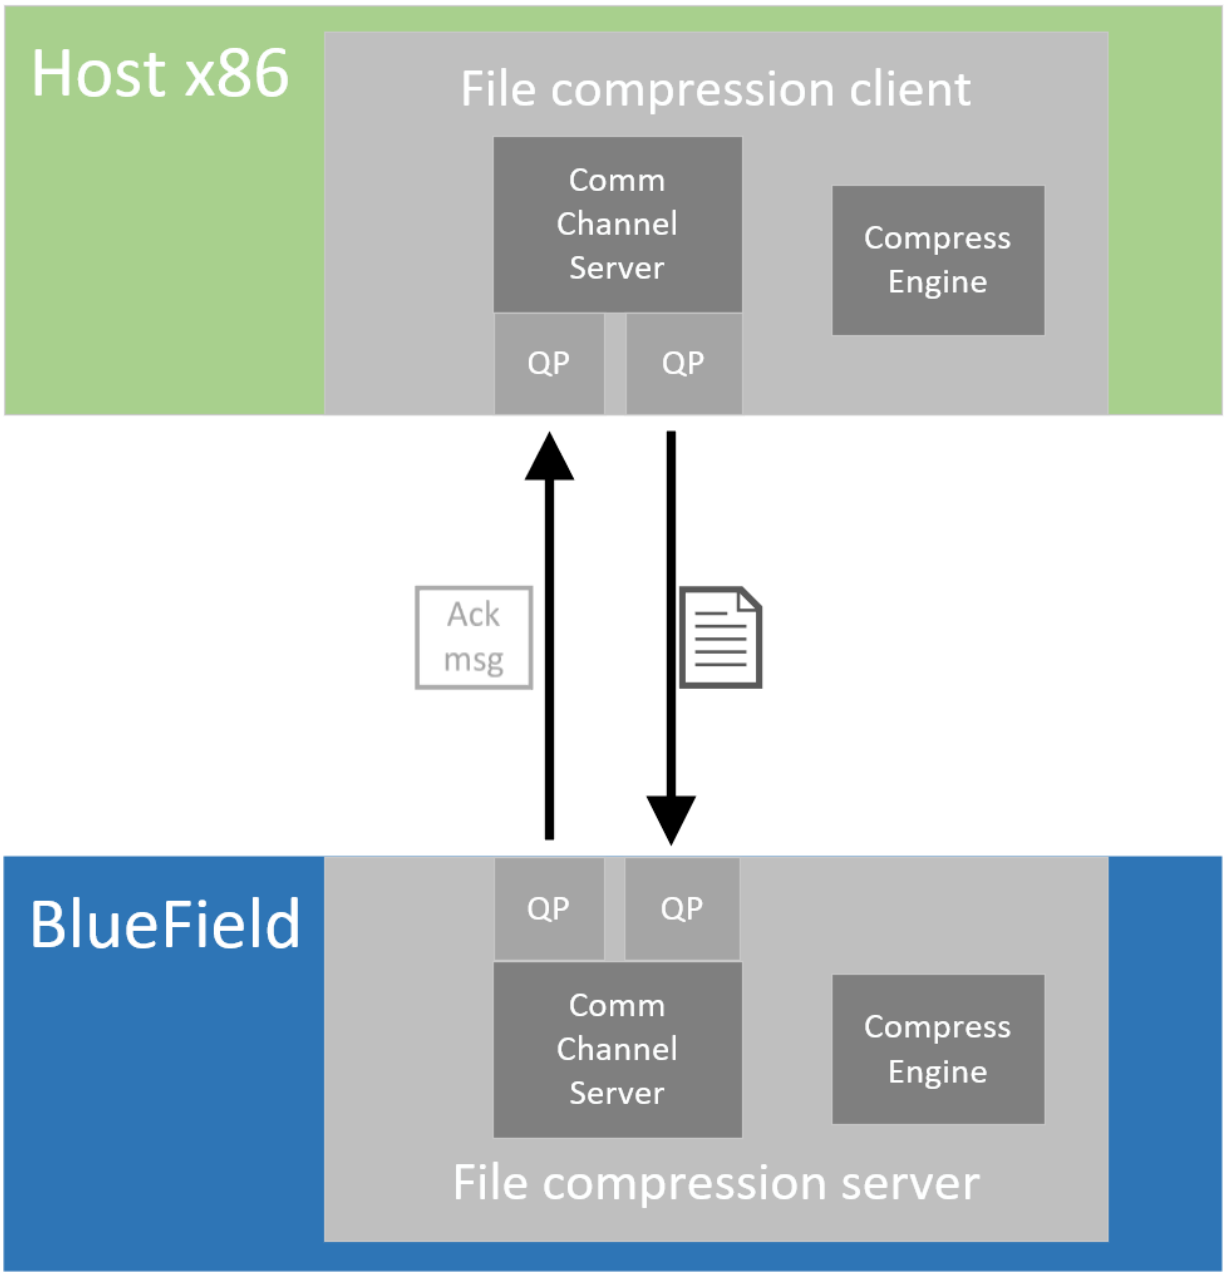
Host x86
File compression client
Comm Channel Server
Compress Engine
QP
QP
Ack msg
BlueField
QP
QP
Comm Channel Server
Compress Engine
File compression server

Chapter 3. Application Architecture

The file compression application runs on top of the DOCA Comm Channel API to send and receive the file from the host and to the DPU.

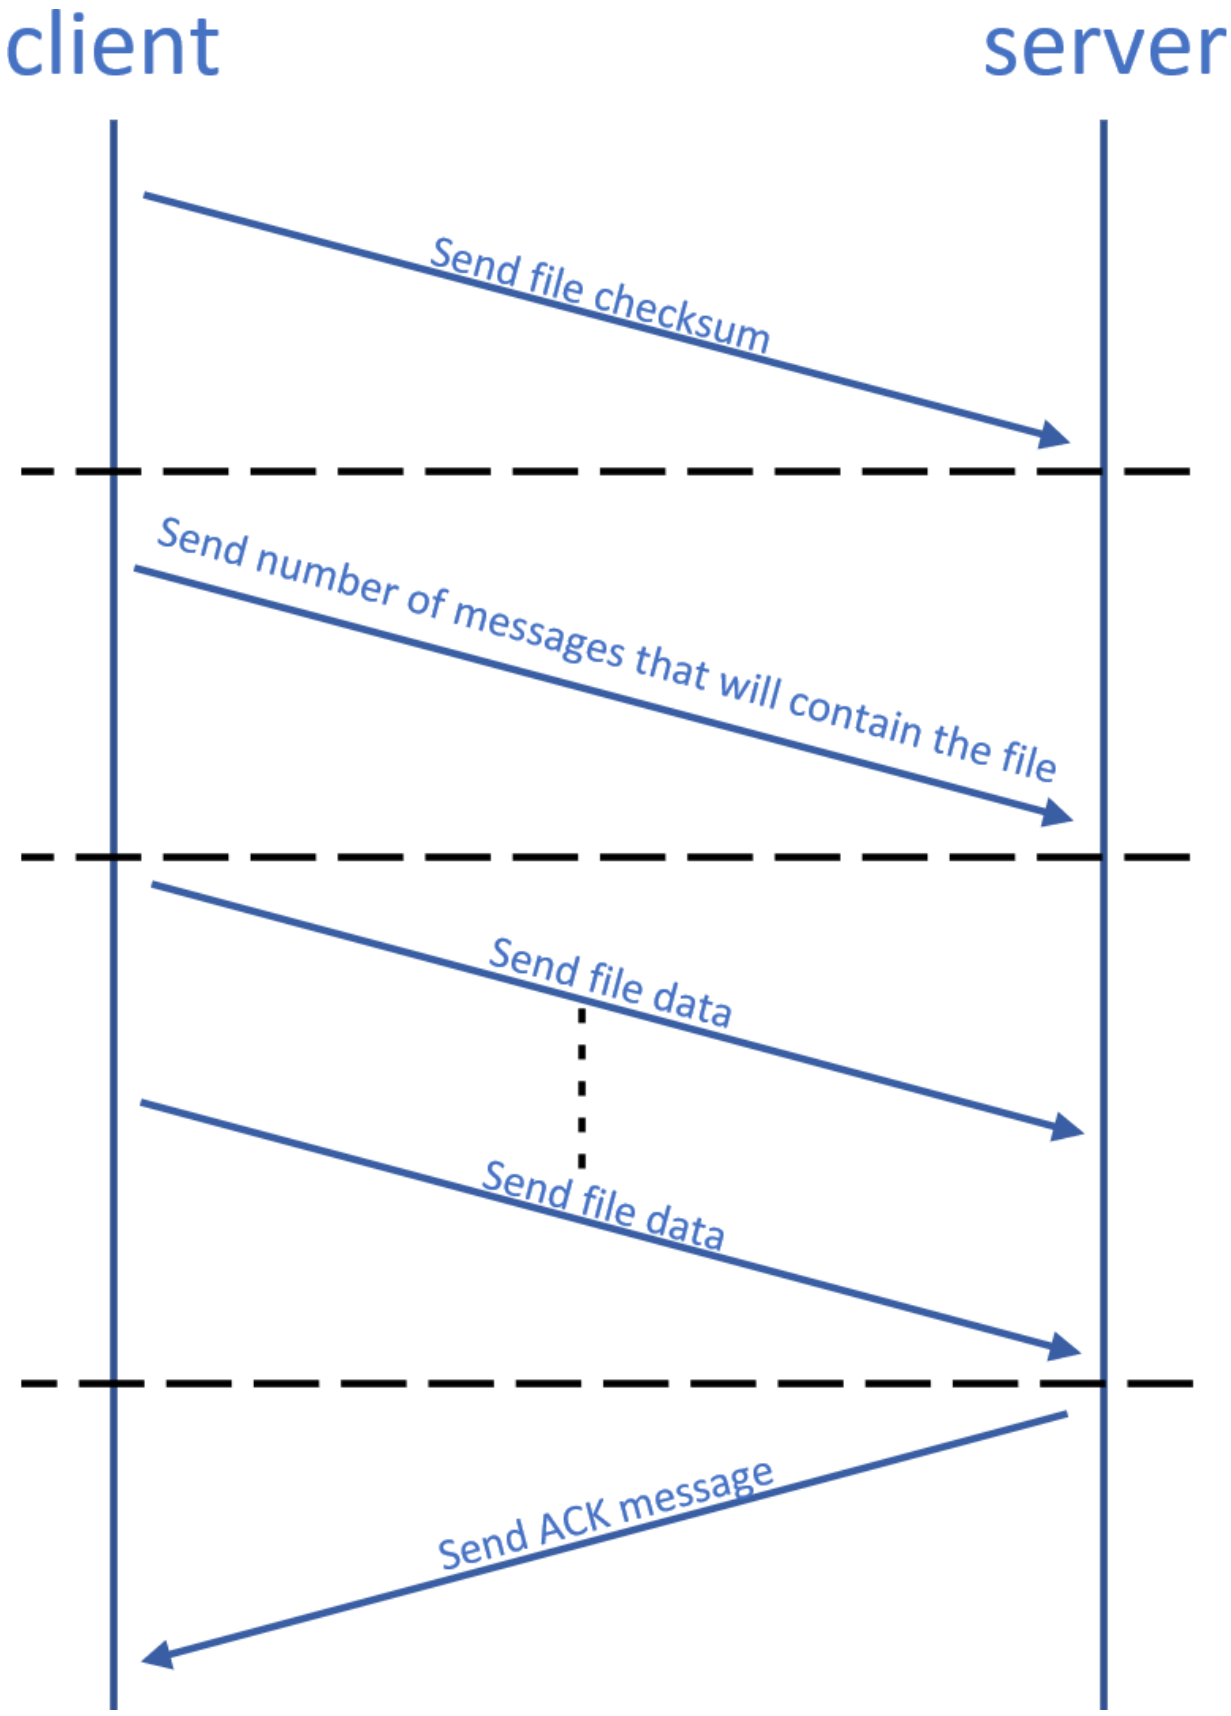
client
server
Send file checksum
Send number of messages that will contain the file
Send file data
Send file data
Send ACK message

1. Connection is established on both sides by DOCA Comm Channel API.
2. Client submits compress job with DOCA Compress API and sends the result to the server.
3. Client sends the number of messages needed to send the compressed content of the file.
4. Client sends data segments in size of up to 4080 bytes.
5. Server saves the received data in a buffer and submits a decompress job.
6. Server sends an ACK message to the client when all parts of the file are received successfully.
7. Server compares the received checksum to the calculated checksum.
8. Server writes the decompressed data to an output file.

Chapter 4. DOCA Libraries

This application leverages the following DOCA libraries:

- DOCA Compress Library
- DOCA Comm Channel Library

Chapter 5. Configuration Flow

1. Parse application argument.
   a). Initialize arg parser resources and register DOCA general parameters.
   ```
   doca_argp_init();
   ```
   b). Register file compression application parameters.
   ```
   register_file_compression_params();
   ```
   c). Parse the arguments.
   ```
   doca_argp_start();
   ```
      i. Parse app flags.
2. Set queue-pair attributes.
   ```
   set_endpoint_properties();
   ```
   a). Set maximum message size of 4080 bytes.
   b). Set maximum messages allowed on queue pair.
3. Create Comm Channel endpoint.
   ```
   doca_comm_channel_ep_create();
   ```
   a). Create endpoint for client/server.
4. Create compress context.
   ```
   doca_compress_create();
   ```
   a). Create compress context for submitting compress jobs for client/server.
5. Run client/server main logic.
   ```
   file_compression_client/server();
   ```
6. Clean up the file compression application.
   ```
   file_compression_cleanup();
   ```
   a). Free all application resources.
7. Arg parser destroy.
   ```
   doca_argp_destroy()
   ```

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.
   - NVIDIA DOCA Applications Overview for additional compilation instructions and development tips regarding the DOCA applications.
2. The file compression binary is located under `/opt/mellanox/doca/applications/file_compression/bin/doca_file_compression`. To build all the applications together, run:

```
cd /opt/mellanox/doca/applications/
meson build
ninja -C build
```

3. To build only the file compression application:
   a). Edit the following flags in `/opt/mellanox/doca/applications/meson_options.txt`:
      - Set `enable_all_applications` to `false`
      - Set `enable_file_compression` to `true`
   b). Run the commands in step 2.

> **Note:** `doca_file_compression` is created under `./build/file_compression/src/`.

Application usage:

```
Usage: doca_file_compression [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:
  -d, --dev-pci             Comm Channel DOCA device PCI address
  -r, --rep-pci             Comm Channel DOCA device representor PCI address
```

```
 -f, --file                  File to send by the client / File to write by the
server
```

Note: For additional information on the application, use `-h` option:

```
/opt/mellanox/doca/applications/file_compression/bin/doca_file_compression
 -h
```

4. CLI example for running the application on BlueField:

```
/opt/mellanox/doca/applications/file_compression/bin/doca_file_compression -d
 03:00.0 -r 3b:00.0 -f received.txt
```

5. CLI example for running the application on the host:

```
/opt/mellanox/doca/applications/file_compression/bin/doca_file_compression -d
 3b:00.0 -f send.txt
```

Note: Refer to section "Running DOCA Application on Host" in NVIDIA DOCA Virtual Functions User Guide.

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

```
doca_file_compression --json [json_file]
```

For example:

```
cd /opt/mellanox/doca/applications/file_compression/bin
./file_compression --json ./file_compression_params.json
```

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: ▶ CRITICAL=20 ▶ ERROR=30 ▶ WARNING=40 ▶ INFO=50 ▶ DEBUG=60	`"log-level": 60`
	`v`	`version`	Print program version information	N/A
	`h`	`help`	Print a help synopsis	N/A
Program Flags	`f`	`file`	For client – path to the file to be sent For server – path to write the file into **Note:** This is mandatory flag.	`"file": "/tmp/data.txt"`
	`d`	`dev-pci`	Comm Channel DOCA device PCIe address **Note:** This is	`"dev-pci": 03:00.1`

Flag Type	Short Flag	Long Flag/JSON Key	Description	JSON Content
			mandatory flag.	
	r	rep-pci	Comm Channel DOCA device representor PCIe address **Note:** This flag is a mandatory only on the DPU.	"rep-pci": b1:00.1

Chapter 8. References

- `/opt/mellanox/doca/applications/file_compression/src/file_compression.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. 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.

NVIDIA Corporation | 2788 San Tomas Expressway, Santa Clara, CA 95051
http://www.nvidia.com