



NVIDIA DOCA Rivermax

Programming Guide

Table of Contents

Chapter 1. Introduction.....	1
Chapter 2. Prerequisites.....	2
Chapter 3. Architecture.....	3
Chapter 4. API.....	4
4.1. Objects.....	4
4.1.1. struct doca_rmax_in_stream.....	4
4.1.2. struct doca_rmax_flow.....	4
4.2. Library Initialization.....	4
4.2.1. doca_rmax_set_cpu_affinity_mask().....	4
4.2.2. doca_rmax_init().....	5
4.2.3. doca_rmax_release().....	5
4.3. Checking Device Capability.....	5
4.3.1. doca_rmax_get_ptp_clock_supported().....	5
4.4. Stream Life Cycle.....	6
4.4.1. Create and Configure Input Stream.....	6
4.4.1.1. doca_rmax_in_stream_create().....	6
4.4.1.2. doca_rmax_in_stream_set_*(()) and doca_rmax_in_stream_get_*(()).....	6
4.4.2. Attach Network Device.....	9
4.4.3. Query Buffer and Stride Size.....	9
4.4.3.1. doca_rmax_in_stream_get_memblok_size().....	10
4.4.3.2. doca_rmax_in_stream_get_memblok_stride_size().....	10
4.4.4. Create and Set Memory Buffer (doca_buf).....	10
4.4.4.1. doca_rmax_in_stream_set_memblok().....	10
4.4.5. Start Stream Context.....	11
4.4.6. Attach to WorkQ.....	11
4.4.7. Attach Flows.....	11
4.4.7.1. doca_rmax_flow_attach().....	11
4.4.8. Receive Data.....	11
4.4.8.1. doca_rmax_in_stream_completion.....	11
4.4.8.2. doca_rmax_stream_error.....	12
4.4.9. Teardown.....	12
4.4.9.1. Detach Flows.....	12
4.4.9.2. Detach from WorkQ.....	13
4.4.9.3. Stop Stream Context.....	13
4.4.9.4. Detach Network Device.....	13

4.4.9.5. Destroy Stream.....	13
4.5. Flow Functions.....	13
4.5.1. Create Flow.....	13
4.5.1.1. <code>doca_rmax_flow_create()</code>	13
4.5.2. Configure Flow Properties.....	13
4.5.2.1. <code>doca_rmax_flow_set_src_ip()</code>	14
4.5.2.2. <code>doca_rmax_flow_set_dst_ip()</code>	14
4.5.2.3. <code>doca_rmax_flow_set_src_port()</code>	14
4.5.2.4. <code>doca_rmax_flow_set_dst_port()</code>	14
4.5.2.5. <code>doca_rmax_flow_set_tag()</code>	15
4.5.3. Destroy Flow.....	15
4.5.3.1. <code>doca_rmax_flow_destroy()</code>	15
Chapter 5. Packaging.....	16
Chapter 6. Samples.....	17

Chapter 1. Introduction

DOCA Rivermax (RMAX) is a DOCA API for NVIDIA Rivermax, an optimized networking SDK for media and data streaming applications. Rivermax leverages NVIDIA® BlueField® DPU hardware streaming acceleration technology which enables direct data transfers to and from the GPU, delivering best-in-class throughput and latency with minimal CPU utilization for streaming workloads.

This document is intended for software developers wishing to accelerate their networking operations.

Chapter 2. Prerequisites

DOCA Rivermax-based applications can run on the target DPU.

It is recommended to have at least 800 huge pages enabled to achieve maximum performance:

```
dpu> echo 1000000000 > /proc/sys/kernel/shmmax  
dpu> echo 800 > /proc/sys/vm/nr_hugepages
```

Chapter 3. Architecture

DOCA Rivermax (RMAX) library contains two objects:

- ▶ `doca_rmax_flow` – is a flow object that represents an IP/port tuple
- ▶ `doca_rmax_in_stream` – represents the input stream and can be thought of as a receive queue which scatters the received data into memory. Each stream can receive one or more flows.

Chapter 4. API

For the library API reference, refer to DOCA Rivermax (RMAX) API documentation in [NVIDIA DOCA Libraries API Reference Manual](#).



Note: The pkg-config (*.pc file) for the Rivermax library is named `doca-rmax`.

The following sections provide additional details about the library API.

4.1. Objects

4.1.1. `struct doca_rmax_in_stream`

Represents Rivermax's input stream. This is a main object in the RX part.

4.1.2. `struct doca_rmax_flow`

Represents the steering flow for the input stream to filter incoming data flow by match criteria.

4.2. Library Initialization

Users must explicitly call library initialization and deinitialization functions. Global parameters such as internal thread affinity mask must be set before initialization.

4.2.1. `doca_rmax_set_cpu_affinity_mask()`

Set the CPU affinity mask for the Rivermax (RMAX) internal thread. This must be called before library initialization.

The function `doca_rmax_get_cpu_affinity_mask()` can be used to query the CPU affinity mask for the Rivermax internal thread at any time.

```
doca_error_t doca_rmax_set_cpu_affinity_mask(const struct
doca_rmax_cpu_affinity_mask * mask);
```


mask [in]

Affinity mask. The CPU is included in the affinity mask if the corresponding bit is set. By default, affinity mask is not set so the internal thread can run on any CPU core.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.2.2. `doca_rmax_init()`

This function initializes the DOCA Rivermax (RMAX) global resources. This function must be called after `doca_rmax_set_cpu_affinity_mask()` and before any other DOCA RMAX library call.

```
doca_error_t doca_rmax_init(void);
```

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.2.3. `doca_rmax_release()`

This function cleans up the DOCA Rivermax (RMAX) resources. No DOCA Rivermax function may be called after calling this function.

```
doca_error_t doca_rmax_release(void);
```

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.3. Checking Device Capability

DOCA Rivermax (RMAX) can query additional device capabilities related to the library.

4.3.1. `doca_rmax_get_ptp_clock_supported()`

Query PTP clock capability for device.

```
doca_error_t doca_rmax_get_ptp_clock_supported(const struct doca_devinfo * devinfo);
```

devinfo [in]

Device to query

Returns

`doca_error_t` value

- ▶ ▶ `DOCA_SUCCESS` – PTP clock is supported
- ▶ `DOCA_ERROR_NOT_SUPPORTED` – PTP clock is not supported
- ▶ Other possible error values are documented in the header file

4.4. Stream Life Cycle

The following sections detail the stages of the stream life cycle from creation to teardown.

4.4.1. Create and Configure Input Stream

4.4.1.1. `doca_rmax_in_stream_create()`

Create a DOCA Rivermax (RMAX) input stream context. The stream can be downcasted to the DOCA context using the function `doca_rmax_in_stream_as_ctx()`.

```
doca_error_t doca_rmax_in_stream_create(struct doca_rmax_in_stream ** stream);
```

stream [out]

The input stream context created for DOCA Rivermax. Non-NULL upon success, NULL otherwise.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.4.1.2. `doca_rmax_in_stream_set_*()` and `doca_rmax_in_stream_get_*()`

Use `doca_rmax_in_stream_set_*()` functions to set the properties of the input stream and the corresponding `doca_rmax_in_stream_get_*()` functions to retrieve the current properties of the input stream.

4.4.1.2.1. Header-data Split Mode

By default, the DOCA Rivermax (RMAX) input stream places the whole packet data into a buffer (`memblk_count = 1`).

To enable header-data split mode when the headers and payload are placed to a different buffers, set the property `memblk_count` (`doca_rmax_in_stream_set_memblks_count`) to 2. In this case you must also set a non-zero value for the `min_size/max_size` property of memory descriptor (`doca_rmax_in_stream_memblk_desc_set_min_size/doca_rmax_in_stream_memblk_desc_set_max_size`).

4.4.1.2.1.1. `doca_rmax_in_stream_set_memblks_count()`

Set the number of configured memory blocks.

```
doca_error_t doca_rmax_in_stream_set_memblks_count(struct doca_rmax_in_stream * stream, uint32_t value);
```

stream [in]

The input stream to write the property.

value [in]

Property value.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.4.1.2.2. Mandatory Properties

4.4.1.2.2.1. `doca_rmax_in_stream_set_elements_count()`

Set number of elements in the stream buffer. The actual number of elements in a buffer can be increased by the library, so the actual value can be queried using `doca_rmax_in_stream_get_elements_count()`.

```
doca_error_t doca_rmax_in_stream_set_elements_count(struct doca_rmax_in_stream *
stream, uint32_t value);
```

stream [in]

The input stream to write the property.

value [in]

Property value.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.4.1.2.2.2. `doca_rmax_in_stream_membblk_desc_set_min_size()`

Set minimal packet segment size(s). The value is an array of the minimal packet segment sizes received by the input stream.

```
doca_error_t doca_rmax_in_stream_membblk_desc_set_min_size(struct doca_rmax_in_stream
* stream, uint16_t * value);
```

stream [in]

The input stream to write the property.

value [in]

- ▶ When `membblk_count=1`, pointer to a variable that contains the packet size
- ▶ When `membblk_count>1`, pointer to an array of packet sizes

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.4.1.2.2.3. `doca_rmax_in_stream_membblk_desc_set_max_size()`

Set maximal packet segment size(s). The value is an array of the maximal packet segment sizes received by the input stream.

```
doca_error_t doca_rmax_in_stream_membblk_desc_set_min_size(struct doca_rmax_in_stream
* stream, uint16_t * value);
```

stream [in]

The input stream to write the property.

value [in]

- ▶ When `membblk_count=1`, pointer to a variable that contains the packet size
- ▶ When `membblk_count>1`, pointer to an array of packet sizes

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.4.1.2.3. Optional Properties

The following properties have a default value and may be set as long as the input stream is not yet active.

4.4.1.2.3.1. `doca_rmax_in_stream_set_type()`

Set the input stream type.

```
doca_error_t doca_rmax_in_stream_set_type(struct doca_rmax_in_stream * stream, enum doca_rmax_in_stream_type value);
```

stream [in]

The input stream to write the property.

value [in]

Property value.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.4.1.2.3.2. `doca_rmax_in_stream_set_scatter_type()`

Set the type of packet's data scatter:

- ▶ All packet data including network headers
- ▶ Only User-Level Protocol data (discard network header up to L4)
- ▶ Payload data only (all headers will be discarded)

```
doca_error_t doca_rmax_in_stream_set_scatter_type(struct doca_rmax_in_stream * stream, enum doca_rmax_in_stream_scatter_type value);
```

stream [in]

The input stream to write the property.

value [in]

Property value.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.4.1.2.3.3. `doca_rmax_in_stream_set_timestamp_format()`

Set stream timestamp format.

```
doca_error_t doca_rmax_in_stream_set_scatter_type(struct doca_rmax_in_stream * stream, enum doca_rmax_in_stream_scatter_type value);
```

stream [in]

The input stream to write the property.

value [in]

Property value.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.4.1.2.4. Run-time Properties

The following properties have a default value and may be set at any time.

4.4.1.2.4.1. `doca_rmax_in_stream_set_min_packets()`

Set minimal number of packets that the input stream must return in a read event.

```
doca_error_t doca_rmax_in_stream_set_min_packets(struct doca_rmax_in_stream *
stream, uint32_t value);
```

stream [in]

The input stream to write the property.

value [in]

Property value.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.4.1.2.4.2. `doca_rmax_in_stream_set_max_packets()`

Set the maximal number of packets that the input stream must return in read event.

```
doca_error_t doca_rmax_in_stream_set_max_packets(struct doca_rmax_in_stream *
stream, uint32_t value);
```

stream [in]

The input stream to write the property.

value [in]

Property value.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.4.1.2.4.3. `doca_rmax_in_stream_set_timeout_us()`

Set receive timeout. The number of μ secs that library would do busy wait (polling) for reception of at least `min_packets` number of packets.

```
doca_error_t doca_rmax_in_stream_set_timeout_us(struct doca_rmax_in_stream * stream,
int value);
```

stream [in]

The input stream to write the property.

value [in]

Property value.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.4.2. Attach Network Device

Attach a network device to the stream context using the `doca_ctx_dev_add()` function. Only one device per stream is supported.

4.4.3. Query Buffer and Stride Size

Query the memory block size to determine the required size of the memory buffers.

4.4.3.1. `doca_rmax_in_stream_get_memblk_size()`

Get the size of the memory blocks.

```
doca_error_t doca_rmax_in_stream_get_memblk_size(struct doca_rmax_in_stream *
stream, size_t * value);
```

stream [in]

The input stream to write the property.

value [out]

Size of the memory block (array of sizes for multiple memblks, the number of memory blocks in stream is more than one).

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.4.3.2. `doca_rmax_in_stream_get_memblk_stride_size()`

Get stride sizes.

```
doca_error_t doca_rmax_in_stream_get_memblk_stride_size(struct doca_rmax_in_stream *
stream, uint16_t * value);
```

stream [in]

The input stream to write the property.

value [out]

Stride size of the memory block (array of stride sizes for multiple memory blocks)

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.4.4. Create and Set Memory Buffer (`doca_buf`)

Allocate the memory of at least one previously queried size and then configure the input stream to use this buffer.

For header-data split mode, users must use two buffers and the buffers for the header and the payload data must be chained. See `doca_buf_list_chain()` for more information.

4.4.4.1. `doca_rmax_in_stream_set_memblk()`

Set memory buffer(s) to be used as received data storage.



Note: Must be set before starting the stream context.

```
doca_error_t doca_rmax_in_stream_set_memblk(struct doca_rmax_in_stream * stream,
struct doca_buf * buf);
```

stream [in]

The input stream to write the property.

buf [in]

Memory buffer (or head of linked list of memory buffers) to store received data. The length of the linked list must be the same as the number of memory blocks configured.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.4.5. Start Stream Context

Cast a stream object to the DOCA context and start the context.

4.4.6. Attach to WorkQ

Use the `doca_ctx_workq_add()` function.

4.4.7. Attach Flows

You can attach one or more flows to a stream after start.

4.4.7.1. `doca_rmax_flow_attach()`

Attach a flow to a stream.

```
doca_error_t doca_rmax_flow_attach (const struct doca_rmax_flow * flow, const struct
doca_rmax_in_stream * stream);
```

flow [in]

Flow to operate on.

buf [in]

The context for attaching a flow.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.4.8. Receive Data

DOCA Rivermax (RMAX) does not accept jobs posted to the WorkQ. Instead, you can retrieve available input packets by polling WorkQ progress. Use `doca_workq_progress_retrieve()` to query available data.



Note: Event-driven mode is not yet supported.

If progress status is `DOCA_SUCCESS`, then event of type `DOCA_RMAX_ACTION_TYPE_RX_DATA` contains a pointer to the `doca_rmax_in_stream_completion` structure in the field of the `result.ptr` event.

4.4.8.1. `doca_rmax_in_stream_completion`

Completion is returned by the input stream describing the incoming packets.

```
struct doca_rmax_in_stream_completion {
    uint32_t elements_count;
    uint64_t ts_first;
    uint64_t ts_last;
    uint32_t seqn_first;
    uint32_t memblk_ptr_arr_len;
```

```
void **memblk_ptr_arr;
};
```

elements_count

Number of packets received.

ts_first

Time of arrival of the first packet.

ts_last

Time of arrival of the last packet.

seqn_first

Sequence number of the first packet.

memblk_ptr_arr_len

Number of memory blocks placed in `memblk_ptr_arr`.

memblk_ptr_arr

Array of pointers to the beginning of the memory block as configured by the input stream create step. The offset between packets inside the memory block is a stride size that can be queried using `doca_rmax_in_stream_get_memblk_stride_size`.

If the progress status is `DOCA_ERROR_IO_FAILED`, then the event of type `DOCA_RMAX_ACTION_TYPE_RX_DATA` would contain a pointer to the `doca_rmax_stream_error` structure in the event's `result.ptr` field.

4.4.8.2. `doca_rmax_stream_error`

Detailed completion error information.

```
struct doca_rmax_stream_error {
    int code;
    const char *message;
};
```

code

Raw Rivermax error code.

message

Human-readable error message.

Progress status `DOCA_ERROR_AGAIN` signifies that no data is available yet. All other codes should be handled as a DOCA error.

4.4.9. Teardown

4.4.9.1. Detach Flows

All flows must be detached from an input stream before stopping it.

4.4.9.1.1. `doca_rmax_flow_detach()`

Detach a flow from a stream.

```
doca_error_t doca_rmax_flow_detach (const struct doca_rmax_flow * flow, const struct
doca_rmax_in_stream * stream);
```

flow [in]

Flow to operate on.

buf [in]

The context for detaching a flow.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.4.9.2. Detach from WorkQ

Use the `doca_ctx_workq_rm()` function.

4.4.9.3. Stop Stream Context

Cast the stream object to the DOCA context and stop the context.

4.4.9.4. Detach Network Device

Use the `doca_ctx_dev_rm()` function.

4.4.9.5. Destroy Stream

4.4.9.5.1. `doca_rmax_in_stream_destroy()`

Destroy a DOCA input stream context. Free all allocated resources associated with a DOCA Rivermax (RMAX) input stream.

```
doca_error_t doca_rmax_in_stream_destroy(struct doca_rmax_in_stream * stream);
```

stream [in]

The context to be destroyed.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.5. Flow Functions

4.5.1. Create Flow

4.5.1.1. `doca_rmax_flow_create()`

Create a steering flow for the input stream to filter the incoming data flow by match criteria.

```
doca_error_t doca_rmax_flow_create(struct doca_rmax_flow ** flow);
```

flow [out]

The context to be destroyed.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.5.2. Configure Flow Properties

4.5.2.1. `doca_rmax_flow_set_src_ip()`

Set the source IP filter for the flow.

```
doca_error_t doca_rmax_flow_set_src_ip(struct doca_rmax_flow * flow, const struct in_addr * ip);
```

flow [in]

Flow to operate on.

ip [in]

Source IPv4 address.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.5.2.2. `doca_rmax_flow_set_dst_ip()`

Set the destination IP filter for the flow.

```
doca_error_t doca_rmax_flow_set_dst_ip(struct doca_rmax_flow * flow, const struct in_addr * ip);
```

flow [in]

Flow to operate on.

ip [in]

Destination IPv4 address.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.5.2.3. `doca_rmax_flow_set_src_port()`

Set the source port filter for the flow.

```
doca_error_t doca_rmax_flow_set_src_port(struct doca_rmax_flow * flow, uint16_t port);
```

flow [in]

Flow to operate on.

port [in]

Source port number. If zero, then any source port is accepted.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.5.2.4. `doca_rmax_flow_set_dst_port()`

Set the destination port filter for the flow.

```
doca_error_t doca_rmax_flow_set_dst_port(struct doca_rmax_flow * flow, uint16_t port);
```

flow [in]

Flow to operate on.

port [in]

Destination port number (non-zero)

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.5.2.5. `doca_rmax_flow_set_tag()`

Set the tag for the flow.

```
doca_error_t doca_rmax_flow_set_tag(struct doca_rmax_flow * flow, uint32_t tag);
```

flow [in]

Flow to operate on.

tag [in]

Non-zero tag

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

4.5.3. Destroy Flow

4.5.3.1. `doca_rmax_flow_destroy()`

Set the tag for the flow.

```
doca_error_t doca_rmax_flow_destroy(struct doca_rmax_flow * flow);
```

flow [in]

Flow to destroy.

Returns

`doca_error_t` value. `DOCA_SUCCESS` if successful, or an error value upon failure. Possible error values are documented in the header file.

Chapter 5. Packaging

DOCA Rivermax (RMAX) is distributed separately from other DOCA libraries:

- ▶ `doca-rmax-libs` – library binary
- ▶ `libdoca-rmax-libs-dev` – header files and samples

DOCA Rivermax binaries depend on:

- ▶ `doca-runtime`
- ▶ `rivermax`

DOCA Rivermax headers and samples:

- ▶ `doca-sdk`
- ▶ `rivermax`
- ▶ `rivermax-ext`

Chapter 6. Samples

Please refer to the [NVIDIA DOCA Rivermax Sample Guide](#) for more information about the API of this DOCA library.

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.