



Mellanox Rivermax™

Frequently Asked Questions (FAQ)

Rev 1.6

NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT ("PRODUCT(S)") AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES "AS-IS" WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER'S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



Mellanox Technologies
350 Oakmead Parkway Suite 100
Sunnyvale, CA 94085
U.S.A.
www.mellanox.com
Tel: (408) 970-3400
Fax: (408) 970-3403

© Copyright 2019. Mellanox Technologies Ltd. All Rights Reserved.

Mellanox®, Mellanox logo, Mellanox Open Ethernet®, LinkX®, Mellanox Spectrum®, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, ONE SWITCH. A WORLD OF OPTIONS®, Open Ethernet logo, Spectrum logo, Switch-IB®, SwitchX®, UFM®, and Virtual Protocol Interconnect® are registered trademarks of Mellanox Technologies, Ltd.

For the complete and most updated list of Mellanox trademarks, visit <http://www.mellanox.com/page/trademarks>.

All other trademarks are property of their respective owners.

Table of Contents

Document Revision History	4
1 Mellanox Rivermax™	5
2 FAQ: Mellanox Rivermax™	5
2.1 Supported Operating Systems/Network Interface Cards (NICs)	5
2.2 Virtual Machine (VM) and Containers.....	6
2.3 SDP, RTP and Network Headers	6
2.4 Supported Video Formats.....	8
2.5 Miscellaneous	8

Document Revision History

Table 1: Document Revision History

Revision	Date	Description
1.6	May 20, 2019	Updated list of currently supported adapters.
1.5	April 01, 2019	Updated various section throughout the document.
1.4	Sept 06, 2018	Updated various section throughout the document.
1.3	April 12, 2018	Official release of this document.
1.2	March 26, 2018	Final Remarks/Edits.
1.1	March 15, 2018	Added additional FAQ.
1.0	March 14, 2018	Initial version of the document.

1 Mellanox Rivermax™

Rivermax running on Mellanox network adapters is designed to provide a dedicated solution for IP-based video streaming, enabling use of common off-the-shelf (COTS) servers for HD to Ultra HD flows by reducing CPU usage, removing bottlenecks for high throughput and providing specification compliance of SMPTE 2110-21 with Mellanox’s network adapters. Rivermax implements an optimized software library API for the media video application while making use of special hardware offloads available in the Mellanox ConnectX®-5 network adapter, allowing the application to remain in frame/line(s) level while hardware strips packet-level headers and assembles frames/line(s) directly to application memory.

Rivermax is “JT-NM Tested”. For more details on the JT-NM Tested program at NAB 2019 and its test results please see https://jt-nm.org/jt-nm_tested.

2 FAQ: Mellanox Rivermax™

The sections below address FAQ related to the Mellanox Rivermax™.

2.1 Supported Operating Systems/Network Interface Cards (NICs)

Question	What Operating Systems (OS) does Rivermax support? Linux/Windows?
Answer	Rivermax supports both Linux and Windows <ul style="list-style-type: none"> • Windows Server 2012 R2 and 2016 • Windows 10 • Linux distros: Ubuntu 16.04 & 18.04, Centos 7.5 and RH7.5
Question	Which Mellanox Network Adapters (NICs) does Rivermax supports?
Answer	Rivermax supports Mellanox ConnectX-5 adapter cards only <ul style="list-style-type: none"> • ConnectX-5 EN - 10/25 GbE SFP28 x 2, Part number: MCX512A-ACAT (Supporting 10/25 GbE speeds) • ConnectX-5 Ex - 100 GbE QSFP28 x 2, Part number: MCX516A-CDAT (Supporting 10/25/40/50/100 GbE speeds)
Question	How can I ensure that Rivermax is using the supported NIC type? (i.e. I have both ConnectX-4 and ConnectX-5 NICs installed in my machine)
Answer	The device IP is used when creating Rivermax resources either directly or inside the SDP file. Rivermax uses the IP address to detect if the associated Mellanox NIC is supported by the Rivermax library. If the NIC is not supported, the API will return an error code.

2.2 Virtual Machine (VM) and Containers

Question	Does Rivermax support a VM/SR-IOV?
Answer	Currently Rivermax supports virtualization on top of VMware ESXi.
Question	Does Rivermax support Containers?
Answer	Current Rivermax release is able to run on Docker Containers at alpha level as it was not fully tested yet. Full Container support is planned for future releases.

2.3 SDP, RTP and Network Headers

Question	Does Rivermax generate RTP packets?
Answer	No. Rivermax receives from the application RTP packets and based on that it generates network packets (adds MAC/IP/UDP). Mellanox is considering adding this feature to future versions.
Question	Does Rivermax support receiving or transmitting RTP payload/header with variable length?
Answer	<p>Yes this is supported by Rivermax.</p> <p>For send requests, the sizes must be indicated upon the output stream creation via <code>rmax_mem_block.app_hdr_size_arr/data_size_arr</code>.</p> <p>For receive requests, the min and max sizes must be indicated upon the input stream creation via <code>rmax_in_buffer_attr.min_data/hdr_size</code> and <code>rmax_in_buffer_attr.max_data/hdr_size</code>.</p> <p>Note: Rivermax input stream is capable of receiving traffic that consists packets with variant sizes that can dynamically change during runtime but must be in the range: <code>[min_data/hdr_size - max_data/hdr_size]</code></p> <p>Note: It is strongly recommended to work with constant length in order to fully utilize hardware offloads and achieve the optimal performance.</p>
Question	On application TX, can the data be separated from the RTP headers?
Answer	Yes. The RTP headers can be separated from the data or placed with the data. Rivermax supports both options.
Question	Is the usage of SDP mandatory?
Answer	<p>Currently yes. SDP provides important information used by Mellanox NICs to support Packet Pacing and other configurations.</p> <p>In future releases Rivermax will also support a more generic API with no need for SDP.</p>

Question	Media attributes are provided to Rivermax via string, why?
Answer	Rivermax is using a standard SDP format which is textual. SDP string stores a lot of information. Some of this information is being used today by Rivermax and some of it will be used in future versions.
Question	Application holds the data in a binary format, is it complicated to convert it to an SDP format?
Answer	The application can easily convert the binary format to an SDP textual format, e.g. via <code>sprintf()</code> .
Question	How can I know which SDP attributes are required by Rivermax?
Answer	Documentation of <code>rmax_create_output_stream()</code> routine explains what is the minimal set of required SDP attributes. Please refer to the <code>rivermax_api.h</code> file for an example of an SDP.
Question	Does Rivermax create the Ethernet & IP layer?
Answer	Yes. The Rivermax layer will create all the Ethernet/IP/UDP headers (based on the information extracted from the SDP).
Question	Does Rivermax support offload of localhost addresses?
Answer	No. Using addresses in the 127.0.0.0/8 subnet is not supported for loopback offload. To offload packets that must be looped back one of the following should be used: - NIC local IP address - Multicast address (which the host has joined)
Question	Does Rivermax support PTP? (timing synchronization)
Answer	Yes. Using Rivermax, each frame must be committed with the time parameter set to the system time (in nanoseconds) at which it must be sent. The system time is being synchronized using a 3rd party software to lock the system time to a PTP GM. The timing synchronization is supported for both Linux and Windows. Note: Currently only gapped mode (TP=2110TPN) is supported with PTP. The PTP client is provided by a 3 rd party

2.4 Supported Video Formats

Question	Which SMPTE streams are supported by Rivermax?
Answer	Rivermax supports 2110-20, 2110-30, 2110-40 2022-7, 2022-6 and 2110-31
Question	Does Rivermax support 2022-7 stream? If yes how?
Answer	<p>Yes.</p> <p>Transmit:</p> <p>Rivermax reads from the SDP string if a redundant transmission is requested and what are the properties of the duplicated streams.</p> <p>When redundancy is requested, Rivermax internally replicates chunks to multiple streams. This mode is currently not validated.</p> <p>Receive:</p> <p>Rivermax supports 2022-7 on the receive side just by allowing the application to receive 2 media streams. The redundancy is made in the application layer.</p> <p>Fully supported by the hardware will be available on the ConnectX-6 adapter cards.</p>
Question	Does Rivermax support 8K UHD streams?
Answer	<p>Using ConnectX-5 Rivermax Linux version can support receiving of 8K UHD video format of 20-80Gb/s. On the transmit side, the Packet Pacing will support 8K UHD video format of up to ~40Gbps.</p> <p>Further bandwidths will be supported in future Rivermax versions.</p>
Question	Does Rivermax support both interlaced & progressive (i.e. 1080i & 1080p) modes?
Answer	Yes. Rivermax supports both interlaced and progressive modes.

2.5 Miscellaneous

Question	Which ST 2110-21 sender types are supported by Rivermax?
Answer	<p>Rivermax supports Narrow Sender (2110TPN), Wide Sender (2110TPW) and Narrow Linear Sender (2110TPNL).</p> <p>Note: Wide sender was not fully tested yet.</p> <p>Note: Currently PTP synchronization is supported only for the Narrow Sender.</p>
Question	Does Rivermax support the extended UDP Size?
Answer	Yes, please make sure to increase interface MTU and set Rivermax strides sizes correctly.
Question	Does Rivermax support the direct access to GPU memory?
Answer	<p>Rivermax supports direct access to GPU memory which allows allocating memory directly on the GPU.</p> <p>Note: This feature is supported only for Linux.</p>

Question	Memory allocation: Is it done by Rivermax or user/application?
Answer	Both modes are supported: memory can be allocated by application or by Rivermax. For input stream it is necessary to precede the memory allocation with a call to <code>rmax_in_query_buffer_size()</code> to determine the amount of memory required to be allocated by the application.
Question	Does Rivermax check the content of the data?
Answer	No
Question	Can an application dynamically, during runtime, change the output stream data/app headers sizes or can these sizes be set only once upon output stream creation?
Answer	The sizes must be set once upon output stream creation. Dynamic modification of data/app header sizes is supported.
Question	Does Rivermax support Receive and Transmit?
Answer	Yes. Both are supported and provide very high bandwidth with low CPU consumption
Question	Why use Rivermax and not DPDK?
Answer	Rivermax is a dedicate solution for M&E. As such, it provides specific features such as Frame/Line(s) read/write, Packet Pacing SMPTE ST 2110-21 spec compliance, and extremely low CPU load. Rivermax is supported on Microsoft Windows. Rivermax provides the same API for Windows and Linux. Rivermax integration time is fast and does not require any networking knowledge. Rivermax is being used by major industry leaders and proven to provide best performance with quick integration time. Mellanox Packet Pacing is not planned to be available under DPDK at this time.
Question	Why use Rivermax and not Mellanox VMA?
Answer	Rivermax is based on the Mellanox VMA Linux technology combined with a more dedicated solution for the M&E industry needs. Mellanox VMA does not provide a solution for Windows, whereas Rivermax works on both Linux and Windows OSs and provides the same API and functionality set. Dedicated features (SW & HW) for the M&E will be supported only by Rivermax.
Question	Do we still need VMA for Linux?
Answer	Yes VMA is needed. Rivermax is using VMA services in control path: e.g. HW resources allocation, address resolution, memory registration and more. However, data path is fully implemented in Rivermax in order to provide the optimal performance.

Question	Does Rivermax provide information such as how much buffers are available, committed or sent?
Answer	No. This capability will be available in future versions
Question	Does Rivermax expose advanced programmable capabilities of the Mellanox NIC for enhanced filtering?
Answer	No. Mellanox NICs have some advanced parser capabilities that are not fully exposed by Rivermax.
Question	Does Rivermax support audio or ancillary on top of video?
Answer	Yes - these are 2110-30 and 2110-40. An example code for 2110-30 and 2110-40 is available as part of our SDK (and Rivermax Player).
Question	Can I receive an application example that sends a real video using Rivermax API?
Answer	Yes - we can provide example code for playout of a real video using Rivermax
Question	What is the maximum transmit and receive bandwidth supported by Rivermax?
Answer	<p>Rivermax bandwidth depends on the link speed of the port which can be 100/50/40/25/10Gbps.</p> <p>The ConnectX-5 silicon ASIC (for all port speeds) is limited by the Gen3 PCIe max bandwidth.</p> <p>The 100GbE version has a X16 PCIe bus (max bandwidth is ~120Gbps). Hence dual port 100GbE NIC is limited to a total of ~120Gbps on total (2ports) in each direction (RX and TX) and single port 100GbE NIC is limited to 100Gbps in the single port in each direction (RX and TX).</p> <p>The 50GbE version (and below) are X8 PCIe bus. Hence dual port 50GbE NIC is limited to ~63Gbps on total on both ports in each direction and single port 50GbE NIC is limited to 50Gbps in the single port in each direction.</p> <p>On RX, NIC's full PCIe bandwidth capacity can be reached except for 50 and 100GbE NICs which have two ports of 50 or 100, adding up to more than the PCIe bandwidth.</p> <p>On TX side, the Packet Pacing mechanism has an overhead. The limit is 85Gbps on total per NIC (both ports for dual port NICs).</p> <p>Therefore, 10 and 25 GbE dual/single port NICs can have full line rate in both ports or in the single port, while the 50 and 100GbE dual/single port NICs can distribute the outbound traffic between the ports or transmit via the single port with the total limit of 85Gbps.</p>