



Mellanox Messaging Accelerator (VMA)

Release Notes

Rev 8.5.7

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 2018. Mellanox Technologies Ltd. All Rights Reserved.

Mellanox®, Mellanox logo, Accelio®, BridgeX®, CloudX logo, CompustorX®, Connect-IB®, ConnectX®, CoolBox®, CORE-Direct®, EZchip®, EZchip logo, EZappliance®, EZdesign®, EZdriver®, EZsystem®, GPUDirect®, InfiniHost®, InfiniBridge®, InfiniScale®, Kotura®, Kotura logo, Mellanox CloudRack®, Mellanox CloudXMellanox®, Mellanox Federal Systems®, Mellanox HostDirect®, Mellanox Multi-Host®, Mellanox Open Ethernet®, Mellanox OpenCloud®, Mellanox OpenCloud Logo®, Mellanox PeerDirect®, Mellanox ScalableHPC®, Mellanox StorageX®, Mellanox TuneX®, Mellanox Connect Accelerate Outperform logo, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, NP-1c®, NP-2®, NP-3®, NPS®, Open Ethernet logo, PhyX®, PlatformX®, PSIPHY®, SiPhy®, StoreX®, SwitchX®, Tiler®, Tiler logo, TestX®, TuneX®, The Generation of Open Ethernet logo, UFM®, Unbreakable Link®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

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

Table of Contents

Document Revision History	5
1 Introduction	6
1.1 System Requirements for VMA 8.5.7	6
1.2 VMA Release Contents	6
1.3 Related Documentation	7
1.4 Certified Applications.....	7
2 Changes and New Features in Rev 8.5.7	8
2.1 Deprecated Features and Support	8
2.2 Important Notes	8
3 Known Issues	9
4 Bug Fixes History	18
5 Change Log History	32
Appendix A: VMA Licensing Model	36

List of Tables

Table 1: Document Revision History	5
Table 2: System Requirements.....	6
Table 3: Release Contents	6
Table 4: VMA Certified Applications	7
Table 5: VMA Known Issues	9
Table 6: Bug Fixes History.....	18
Table 7: Change Log History	32

Document Revision History

Table 1: Document Revision History

Revision	Date	Description
Rev 8.5.7	March 1, 2018	Initial release of this VMA version.

1 Introduction

These release notes pertain to the Mellanox Messaging Accelerator (VMA) library for Linux, software version 8.5.7. The VMA library accelerates TCP and UDP socket applications, by offloading traffic from the user-space directly to the network interface card (NIC) or Host Channel Adapter (HCA), without going through the kernel and the standard IP stack (kernel-bypass). VMA increases overall traffic packet rate, reduces latency, and improves CPU utilization.

VMA v8.5.7 can work on top of both MLNX_OFED driver stack that supports Ethernet and InfiniBand and on a lighter driver stack, MLNX_EN that support only Ethernet.

1.1 System Requirements for VMA 8.5.7

The following table presents the currently certified combinations of stacks and platforms, and supported CPU architectures for VMA 8.5.7.

Table 2: System Requirements

Specification	Value
Network Adapter Cards	<ul style="list-style-type: none"> ConnectX®-3/ConnectX®-3 Pro ConnectX®-4/ConnectX®-4 Lx ConnectX®-5/ConnectX®-5 Ex
Firmware	<ul style="list-style-type: none"> ConnectX®-3/ConnectX®-3 Pro v2.42.5000 ConnectX®-4 v12.22.1002 ConnectX®-4 Lx v14.22.1002 ConnectX®-5/ConnectX®-5 Ex v16.22.1002
Driver Stack	MLNX-OFED: v4.3-x.x.x.x [Ethernet and InfiniBand]
Supported Operating Systems and Kernels	All Linux 64 bit distributions supported by: <ul style="list-style-type: none"> MLNX_OFED 4.3-x.x.x.x MLNX_EN v4.2-x.x.x.x and above
CPU Architecture	x86_64 (Intel Xeon), arm64, ppc64
Minimum memory requirements	1 GB of free memory for installation 800 MB per process running with VMA
Minimum disk space requirements	1 GB
Transport	Ethernet / InfiniBand / VPI

1.2 VMA Release Contents

Table 3: Release Contents

Item	Description
Binary RPM and DEB packages for 64-bit architecture for Linux distribution	libvma_8.5.7_amd64.deb libvma-8.5.7.x86_64.rpm libvma-dev_8.5.7_amd64.deb libvma-devel-8.5.7.x86_64.rpm

Item	Description
	libvma-utils_8.5.7_amd64.deb libvma-utils-8.5.7.x86_64.rpm
Documentation	VMA Release Notes VMA Installation and Quick Start Guide VMA User Manual

1.3 Related Documentation

- Mellanox Messaging Accelerator (VMA) Library for Linux User Manual (DOC-00393)
- Mellanox VMA Installation Guide (DOC-10055)
- Performance Tuning Guidelines:
 - For Mellanox Network Adaptors: <https://community.mellanox.com/docs/DOC-2489>
 - For VMA: <https://community.mellanox.com/docs/DOC-2797>

1.4 Certified Applications

The VMA library version 8.5.7 was successfully tested, and is certified to work with the applications listed in the following table.

Table 4: VMA Certified Applications

Application	Company / Source	Application Type	Notes
Redis	Open Source	Advanced key-value store	http://redis.io/
sockperf	Mellanox (Open Source)	Bandwidth and Latency Benchmarking	Version 3.3 Included in VMA package. https://github.com/mellanox/sockperf
iperf	NLANR	Bandwidth Benchmarking	Version 2.0.5
netperf	Open Source	Bandwidth and Latency Benchmarking	Version 2.6.0
NetPIPE	Open Source	Network Protocol Independent Performance Evaluator	Version 3.7.2
UMS (formerly LBM)	Informatica	Message Middleware Infrastructures	Version 6.7
Opra FeedHandler	NYSE Technologies (WombatFS)	Market Data Infrastructures	Running with WDF/LBM/UMS/RV middleware

2 Changes and New Features in Rev 8.5.7

Following are the main changes and new features in VMA v8.5.7:

Feature/Category	Description
Static Address Resolution Protocol (ARP)	When Static ARP Entries are manually added on a locally cached ARP table, sending out ARP requests is no longer necessary.
Reproducible Build	<p>Added the option to override build date in order to make builds reproducible.</p> <p>For further information on reproducible builds, visit: https://reproducible-builds.org/</p> <p>For further information on build systems, visit: https://reproducible-builds.org/specs/source-date-epoch/</p>
TCP Congestion Control	<p>Added the option to disable TCP Congestion Control algorithm by setting VMA_TCP_CC_ALGO parameter to a new value of 2. VMA will now send data as long as there is room in the advertised window.</p> <p>For more details about Congestion Control Algorithms, see RFC-2581.</p>
Network Virtual Service Client (NetVSC)	Added support for VMA to detect NetVSC in the guest operating system.
	Added support for manage flows in daemon for NetVSC interfaces.
socketXtreme	Renamed vmapoll (Explicit Ring Polling) interfaces to socketXtreme.
Daemon	Added support for systemd service management system on hosts whose operating system has systemd enabled.

For further information on the features described in the table above, please refer to the VMA User Manual.

2.1 Deprecated Features and Support



NOTE: Starting VMA 8.6.x (future VMA versions), VMA will stop supporting MLNX_OFED versions earlier than 4.3-x.x.x.x.

2.2 Important Notes

- Starting VMA 8.5.x, the VMA_POLL interfaces are renamed to SocketXtreme.
- We recommend using libnl3 as this is the latest version and it fixes issues related to libnl1

3 Known Issues

The following table describes known issues in VMA Rev 8.5.7, and existing workarounds.

Table 5: VMA Known Issues

Internal Ref.	Issues
1310135	<p>Description: When working with Linux guest over Windows Hypervisor, and more than 256 flow steering rules are set in the VM, the following error message will appear: VMA ERROR: rfs[0x32ea410]:273:create_ibv_flow() Create of QP flow ID (tag: 0) failed with flow dst:5.5.5.77:6891, src:0.0.0.0:0, proto:UDP (errno=22 - Invalid argument) VMA ERROR: ring_simple[0x3292d50]:555:attach_flow() attach_flow=0 failed! As a result, TCP Receive Flow will not be supported, and UDP Receive Flow will not be offloaded.</p> <p>Workaround: N/A</p> <p>Keywords: Windows Hypervisor, flow steering limit</p> <p>Discovered in Version: 8.5.7</p>
-	<p>Description: When working with Linux guest over Windows Hypervisor, and exceeding the maximum amount of flow steering rules supported by the VM, the following error message will appear: VMA ERROR: rfs[0x32ea410]:273:create_ibv_flow() Create of QP flow ID (tag: 0) failed with flow dst:5.5.5.77:6891, src:0.0.0.0:0, proto:UDP (errno=12 - Cannot allocate memory) VMA ERROR: ring_simple[0x3292d50]:555:attach_flow() attach_flow=0 failed! As a result, TCP Receive Flow will not be supported, and UDP Receive Flow will not be offloaded.</p> <p>Workaround: Reduce the amount of supported VMs for the device in the Hypervisor to increase the total flow steering rules amount for each VM.</p> <p>Keywords: Windows Hypervisor, flow steering limit</p> <p>Discovered in Version: 8.5.7</p>
1309781	<p>Description: When working with Linux guest over Windows Hypervisor, and two VFs of the same physical device are mapped, VMA will manage to offload only one VF.</p> <p>Workaround: N/A</p> <p>Keywords: Windows Hypervisor, two VFs</p> <p>Discovered in Version: 8.5.7</p>
1161458	<p>Description: In kernel versions 3.10.0-693 and above, when creating a bond active-backup consisting of precisely two ports, the following error will be received: “VMA_ERROR - prepare_ibv_qp() failed to modify QP from ERR to INIT state (ret = -3)”</p> <p>Workaround: Starting from MLNX_OFED v4.3-x.x.x.x, you can enable RoCE LAG using sysfs: /sys/bus/pci/drivers/mlx5_core/<bdf>/roce_lag_enable (1 will enable RoCE LAG (default value) and 0 will disable it).</p>

Internal Ref.	Issues
	<p>Note that enablement and disablement through sysfs is non-persistent after driver restart.</p> <p>Keywords: RoCE LAG, bonding, bond</p> <p>Discovered in Version: 8.4.10</p>
1271529	<p>Description: On RedHat 7.2 OS, when using permanent file configuration for bonding type active_active, traffic is not offloaded.</p> <p>Workaround:</p> <ul style="list-style-type: none"> • Upon system start, configure the bond using the following manual commands: modprobe bonding, and ifenslave or • Use RedHat v7.3 and above <p>Keywords: Bonding, active_active, permanent</p> <p>Discovered in Version: 8.4.10</p>
1201675	<p>Description: When a non-privileged user uses VMA with RHEL inbox to perform networking operations (i.e. allocate IB resources) VMA crashes with a segmentation fault.</p> <p>Workaround: Use VMA with root privileges on the RHEL inbox driver.</p> <p>Keywords: RHEL inbox driver, segmentation fault</p> <p>Discovered in Version: 8.4.10</p>
1092906	<p>Description: Starting kernel version 3.10.0-693, when creating a bond LAG consisting of precisely two ports, the bond is not offloaded if both ports belong to a single device.</p> <p>Workaround: Make sure there are at least two ports belonging to different devices enslaved under the bond or add a dummy interface slave when creating a bond LAG,</p> <p>Keywords: Two port bond LAG, RoCE over LAG, mlx4, mlx5</p> <p>Discovered in Version: 8.4.8</p>
1161458	<p>Description: Starting kernel version 3.10.0-693, when creating a bond active-backup consisting of precisely two ports the following error will be received: VMA_ERROR - prepare_ibv_qp() failed to modify QP from ERR to INIT state (ret = -3)</p> <p>Workaround: Make sure there are be at least two ports belonging to different devices enslaved under the bond or add a dummy interface slave when creating a bond active-backup.</p> <p>Keywords: Two port bond active-backup, RoCE over LAG, mlx4</p> <p>Discovered in Version: 8.4.8</p>
-	<p>Description: The following VMA_ERRORS will be displayed when running ping with root permissions:</p> <ul style="list-style-type: none"> • VMA ERROR: ring_simple[0x7f257d18d720]:256:create_resources() ibv_create_comp_channel for tx failed. m_p_tx_comp_event_channel = (nil) (errno=13 Permission denied) • VMA ERROR: ib_ctx_handler213:mem_dereg() failed de-registering a memory region (errno=13 Permission denied)

Internal Ref.	Issues
	<p>Workaround: N/A</p> <p>Keywords: VMA_ERROR while running ping with root permissions</p> <p>Discovered in Version: 8.4.8</p>
965237	<p>Description: The following sockets APIs are directed to the OS and are not offloaded by VMA:</p> <ul style="list-style-type: none"> • int socketpair(int domain, int type, int protocol, int sv[2]); • int dup(int oldfd); • int dup2(int oldfd, int newfd); <p>Workaround: N/A</p> <p>Keywords: sockets, socketpair, dup, dup2</p>
965227	<p>Description: Multicast (MC) loopback within a process is not supported by VMA:</p> <ul style="list-style-type: none"> • If an application process opens 2 (or more) sockets on the same MC group they will not get each other's traffic. <p>Note: MC loopback between different VMA processes always work.</p> <ul style="list-style-type: none"> • Both sockets will receive all ingress traffic coming from the wire <p>Workaround: N/A</p> <p>Keywords: Multicast, Loopback</p>
965227	<p>Description: MC loopback between VMA and the OS limitation.</p> <ul style="list-style-type: none"> • The OS will reject loopback traffic coming from the NIC • MC traffic from the OS to VMA is functional <p>Workaround: N/A</p> <p>Keywords: Multicast, Loopback</p>
965227	<p>Description: MC loopback Tx is currently disabled and setsockopt (IP_MULTICAST_LOOP) is not supported.</p> <p>Workaround: N/A</p> <p>Keywords: Multicast, Loopback</p>
919301	<p>Description: VMA supports bonding in the following modes:</p> <ul style="list-style-type: none"> • For active-passive (mode=1), use either fail_over_mac=0 or fail_over_mac=1. • For active-active (mode=4), use fail_over_mac=0. • For VLAN over bond, use fail_over_mac=0 for traffic to be offloaded <p>Workaround: N/A</p> <p>Keywords: High Availability (HA)</p>
1011005	<p>Description: VMA SELECT option supports up to 200 sockets in TCP.</p> <p>Workaround: Use ePoll that supports up to 6000 sockets.</p> <p>Keywords: SockPerf</p>
977899	<p>Description: An unsuccessful trial to connect to a local interface, is reported by VMA as Connection timeout rather that Connection refused.</p> <p>Workaround: N/A</p> <p>Keywords: Verification</p>

Internal Ref.	Issues
1019085	Description: Poll is limited with the amount of sockets.
	Workaround: Use ePoll for large amount of sockets (tested up to 6000)
	Keywords: Poll, ePoll
-	Description: Some performance degradation can be seen in low latency use case of IPoIB (not related to the ETH mode). Note: This will be fixed in the next MLNX_OFED version.
	Workaround: N/A
	Keywords: Performance
-	Description: VLAN on the bond interface does not function properly when bonding is configured with <code>fail_over_mac=1</code> due to a kernel bug.
	Workaround: Set the <code>fail_over_mac=0</code>
	Keywords: VLAN and High Availability (HA)
-	Description: Occasionally, traffic is not offloaded when manually configuring bonding the bond interface IP is given before enslaving the slave interfaces
	Workaround: Configure the active-active bond on the host according to one of the following ways: <ul style="list-style-type: none"> • Configure it automatically using the <code>ifcfg</code> files • Configure it manually in the following order: <ol style="list-style-type: none"> a. <code>modprobe bonding mode=4 miimon=100 fail_over_mac=0</code> b. <code>ifconfig bond0 up</code> c. <code>ifenslave bond0 <list_of_slaves></code> d. <code>ifconfig bond0 <bond_ip></code> Note: The IP is assigned to the bond interface after the slaves are added.
	Keywords: High Availability (HA) with <code>fail_over_mac=0</code>
-	Description: If LACP is not configured properly on the switch, received multicast traffic might be duplicated
	Workaround: Make sure LACP is configured properly, or move to kernel > 3.14 or RH7.2 and higher as they already include a fix
	Keywords: LACP High Availability (HA) – multiple MC traffic
-	Description: RX UDP UC and MC traffic in Ethernet and RX UDP UC in InfiniBand with fragmented packages (message size is larger than MTU) is not offloaded by VMA and will pass through the Kernel network stack. There might be performance degradation.
	Workaround: N/A
	Keywords: Issues with UDP fragmented traffic reassembly
-	Description: <code>VMA_TRACELEVEL=4</code> debug mode prints more info, which causes higher latency.
	Workaround: For best performance, run VMA with a lower than 4 <code>VMA_TRACELEVEL</code> value.
	Keywords: <code>VMA_TRACELEVEL=4</code> causes performance degradation
-	Description: The system runs out of memory due to huge-page reserved resources.

Internal Ref.	Issues
	<p>Workaround: Use Contiguous Pages instead of Huge Pages to gain performance improvements.</p> <p>The following parameter should be set as follow: <code>VMA_MEM_ALLOC_TYPE=1</code> (this is the default mode).</p>
	<p>Keywords: Huge-page reserved resources</p>
-	<p>Description: The following VMA_PANIC will be displayed when there are not enough open files defined on the server:</p> <pre>VMA PANIC : si[fd=1023]:51:sockinfo() failed to create internal epoll (ret=-1 Too many open files)</pre> <p>Workaround: Verify that the number of max open FDs (File Descriptors) in the system (<code>ulimit -n</code>) is twice as number of needed sockets. VMA internal logic requires one additional FD per offloaded socket.</p>
	<p>Keywords: VMA_PANIC while opening large number of sockets</p>
-	<p>Description: In MLNX_OFED versions earlier than 4.0-2.0.0.0 and VMA 8.2.10 socket API is only supported in the child process if the parent process has not called any socket routines prior to calling fork.</p> <p>In MLNX_OFED versions 4.0-1.6.1.0, VMA 8.2.10 and later the above restriction no longer exists however the child process cannot use any of the parent's socket resources.</p> <p>Workaround: VMA supports <code>fork()</code> if <code>VMA_FORK=1</code> (is enabled) and the Mellanox-supported stack OFED 4.0-1.6.1.0 or later is used. MLNX_OFED support for fork is for kernels supporting the <code>MADV_DONTFORK</code> flag for <code>madvise()</code> (2.6.17 and later), provided that the application does not use threads.</p> <p>The Posix <code>system()</code> call is supported.</p>
	<p>Keywords: There is limited support for <code>fork()</code>.</p>
-	<p>Description: Applications written in Java use IPv6 by default which is not supported by VMA and may lead to VMA not offloading packets</p> <p>Workaround: To change Java to work with IPv4, instruct the application to use <code>"Java -Djava.net.preferIPv4Stack=true"</code></p>
	<p>Keywords: Java applications using IPv6 stack</p>
-	<p>Description: When a VMA-enabled application is running, there are several cases when it does not exit as expected pressing CTRL-C.</p> <p>Workaround: Enable SIGINT handling in VMA, by using:</p> <pre>#export VMA_HANDLE_SIGINTR=1</pre>
	<p>Keywords: The VMA application does not exit when you press CTRL-C.</p>
-	<p>Description: VMA does not support network interface or route changes during runtime</p> <p>Workaround: N/A</p>
	<p>Keywords: Dynamic route or IP changes</p>
-	<p>Description: The send rate is higher than the receive rate. Therefore when running one sockperf server with one sockperf client there will be packets loss.</p> <p>Workaround: Limit the sender max PPS per receiver capacity.</p> <p>Example below with the following configuration :</p> <ul style="list-style-type: none"> • O.S: Red Hat Enterprise

Internal Ref.	Issues
	<ul style="list-style-type: none"> • Linux Server release 6.2 (Santiago) Kernel \r on an \m • Kernel: 2.6.32-220.el6.x86_64 • link layer: InfiniBand 56G Ethernet 10G • GEN type: GEN3 • Architecture: x86_64 • CPU: 16 • Core(s) per socket: 8 • CPU socket(s): 2 • NUMA node(s): 2 • Vendor ID: GenuineIntel • CPU family: 6 • Model: 45 • Stepping: 7 • CPU MHz: 2599.926 <p>MC 1 socket max pps 3M MC 10 sockets (select) max pps 1.5M MC 20 sockets (select) max pps 1.5M MC 50 sockets (select) max pps 1M UC 1 socket max pps 2.8M UC 10 sockets max pps 1.5M UC 20 sockets max pps 1.5M UC 50 sockets max pps</p>
	<p>Keywords: Packets loss occurs when running sockperf with max pps rate</p>
-	<p>Description: VMA behavior of epoll EPOLLET (Edge Triggered) and EPOLLOUT flags with TCP sockets differs between OS and VMA.</p> <ul style="list-style-type: none"> • VMA - triggers EPOLLOUT event every received ACK (only data, not syn/fin) • OS - triggers EPOLLOUT event only after buffer was full.
	<p>Workaround: N/A</p>
	<p>Keywords: VMA behavior of epoll EPOLLET (Edge Triggered) and EPOLLOUT flags with TCP sockets</p>
-	<p>Description: VMA does not close connections located on the same node (sends FIN to peers) when its own process is terminated.</p>
	<p>Workaround: N/A</p>
	<p>Keywords: VMA connection</p>
-	<p>Description: VMA is not closed (sends FIN to peers) when its own process is terminated when the /etc/init.d/vma is stopped.</p>
	<p>Workaround: Launch /etc/init.d/vma start</p>
	<p>Keywords: VMA connection</p>
-	<p>Description: When a non-offloaded process joins the same MC address as another VMA process on the same machine, the non-offloaded process does not get any traffic.</p>
	<p>Workaround: Run both processes with VMA</p>
	<p>Keywords: MC traffic with VMA process and non VMA process on the same machine</p>

Internal Ref.	Issues
-	<p>Description: Occasionally, epoll with EPOLLONESHOT does not function properly.</p> <p>Workaround: N/A</p> <p>Keywords: epoll with EPOLLONESHOT</p>
-	<p>Description: Occasionally, when running UDP SFNT-STREAM client with poll muxer flag, the client side ends with an expected error: ERROR: Sync messages at end of test lost ERROR: Test failed. This only occurs with poll flag</p> <p>Workaround: Set a higher acknowledgment waiting time value in the <code>sfnt-stream</code>.</p> <p>Keywords: SFNT-STREAM UDP with poll muxer flag ends with an error on client side</p>
-	<p>Description: Occasionally, SFNT-STREAM UDP client hangs when running multiple times.</p> <p>Workaround: Set a higher acknowledgment waiting time value in the <code>sfnt-stream</code>.</p> <p>Keywords: SFNT-STREAM UDP client hanging issue</p>
-	<p>Description: MC loopback in InfiniBand functions only between 2 different processes. It will not work between threads in the same process.</p> <p>Workaround: N/A</p> <p>Keywords: MC loopback in InfiniBand</p>
-	<p>Description: Ethernet loopback functions only if both sides are either off-loaded or not-offloaded.</p> <p>Workaround: N/A</p> <p>Keywords: Ethernet loopback is not functional between the VMA and the OS</p>
-	<p>Description: The following error may occur when running netperf TCP tests with VMA: remote error 107 'Transport endpoint is not connected'</p> <p>Workaround: Use netperf 2.6.0</p> <p>Keywords: Error when running netperf 2.4.4 with VMA</p>
-	<p>Description: Occasionally, a packet is not sent if the socket is closed immediately after <code>send()</code> (also for blocking socket)</p> <p>Workaround: Wait several seconds after <code>send()</code> before closing the socket</p> <p>Keywords: A packet is not sent if the socket is closed immediately after <code>send()</code></p>
-	<p>Description: It can take for VMA more time than the OS to return from an <code>iomux</code> call if all sockets in this <code>iomux</code> are empty sockets</p> <p>Workaround: N/A</p> <p>Keywords: <code>Iomux</code> call with empty sockets</p>
-	<p>Description: TCP throughput with maximum rate may suffer from traffic "hiccups".</p> <p>Workaround: Set the <code>mps = 1000000</code></p>

Internal Ref.	Issues
	Keywords: TCP throughput with maximum rate
-	Description: Netcat with VMA on SLES 11 SP1 does not function.
	Workaround: N/A
	Keywords: Netcat on SLES11 SP1
-	Description: Sharing of HW resources between the different working threads might cause lock contentions which can affect performance.
	Workaround: Use different ring allocation logics.
	Keywords: Issues with performance with some multi-threaded applications
-	Description: Known NetPIPE bug - Netpipe is trying to access read-only memory.
	Workaround: Upgrade to NetPIPE 3.7 or later.
	Keywords: Segmentation fault on NetPIPE exit.
-	Description: If VMA runs when VMA_HANDLE_SIGINTR is enabled, an error message might be written upon exiting.
	Workaround: Ignore the error message, or run VMA with VMA_HANDLE_SIGINTR disabled.
	Keywords: When exiting, VMA logs errors when the VMA_HANDLE_SIGINTR is enabled.
-	Description: VMA suffers from high latency in low message rates.
	Workaround: Use “Dummy Sent”.
	Keywords: VMA ping-pong latency degradation as PPS is lowered
-	Description: VMA does not support broadcast traffic.
	Workaround: Use libvma.conf to pass broadcast through OS
	Keywords: No support for direct broadcast
-	Description: Directing VMA to access non-valid memory area will cause a segmentation fault.
	Workaround: N/A
	Keywords: There is no non-valid pointer handling in VMA
-	Description: VMA allocates resources on the first connect/send operation, which might take up to several tens of milliseconds.
	Workaround: N/A
	Keywords: First connect/send operation might take more time than expected
-	Description: Calling select upon shutdown of socket will return “ready to write” instead of timeout.
	Workaround: N/A
	Keywords: Calling select() after shutdown (write) returns socket ready to write, while select() is expected to return timeout
-	Description: VMA does not raise sigpipe in connection shutdown.
	Workaround: N/A
	Keywords: VMA does not raise sigpipe

Internal Ref.	Issues
-	Description: VMA polls the CQ for packets; if no packets are available in the socket layer, it takes longer.
	Workaround: N/A
	Keywords: When there are no packets in the socket, it takes longer to return from the read call
-	Description: Select with more than 1024 sockets is not supported
	Workaround: Compile VMA with SELECT_BIG_SETSIZE defined.
	Keywords: 1024 sockets

4 Bug Fixes History

The following table describes the issues that have been resolved in VMA.

Table 6: Bug Fixes History

Internal Ref.	Description
1280902	Description: Fixed the following issues related to <code>tcp_split_segment()</code> <ul style="list-style-type: none"> VMA accumulating several buffers in a single segment VMA not updating the unsent queue as needed
	Keywords: TCP, TCP split segment
	Discovered in Release: 8.4.10
	Fixed in Release: 8.5.7
1149532	Description: Fixed an issue where deadlock may have occurred while the socket Send buffer was full.
	Keywords: TCP, TCP Send, TCP deadlock
	Discovered in Release: 8.4.10
	Fixed in Release: 8.5.7
1256248	Description: Fixed an issue which caused VMA daemon installation failures in some Debian packages.
	Keywords: Installation, systemd, daemon
	Discovered in Release: 8.4.10
	Fixed in Release: 8.5.7
1214453	Description: Fixed an issue in <code>inotify</code> (inode file monitoring system) event processing in VMA daemon.
	Keywords: <code>inotify</code> , VMA daemon
	Discovered in Release: 8.4.10
	Fixed in Release: 8.5.7
1262198	Description: Fixed an issue in <code>epoll_wait()</code> which might have caused VMA to return <code>EPOLLOUT</code> event twice for the same fd.
	Keywords: <code>EPOLLOUT</code> , <code>epoll_wait</code> , <code>epoll</code>
	Discovered in Release: 8.4.10
	Fixed in Release: 8.5.7
1247358	Description: Fixed an issue where VMA could crash when checking an invalid neighbor value.
	Keywords: ARP, Failed ARP entry
	Discovered in Release: 8.4.10
	Fixed in Release: 8.5.7
1258245	Description: Fixed an issue in <code>INCOMPLETE</code> neighbor state handling that caused VMA to send SYN messages to an invalid peer.
	Keywords: SYN, <code>INCOMPLETE</code> neighbor

Internal Ref.	Description
	Discovered in Release: 8.4.10
	Fixed in Release: 8.5.7
1263120	Description: Fixed a configuration issue that emerged when libibverbs library was not compiled with libnl library.
	Keywords: libibverbs, libnl
	Discovered in Release: 8.4.10
	Fixed in Release: 8.5.7
1235797	Description: Fixed an issue that caused VMA to handle instead of ignore SYN-FIN packets in Listen state.
	Keywords: TCP Listen
	Discovered in Release: 8.4.10
	Fixed in Release: 8.5.7
1229443	Description: Fixed an issue that caused VMA to handle instead of ignore SYN-RST packets in Listen state.
	Keywords: TCP Listen
	Discovered in Release: 8.4.10
	Fixed in Release: 8.5.7
1185372	Description: Fixed an issue that prevented VMA from opening a QP using high value of VMA_RX_WRE on each ETH device (verify_eth_qp_creation()).
	Keywords: VMA_RX_WRE parameter
	Discovered in Release: 8.4.8
	Fixed in Release: 8.4.10
1201040	Description: Fixed an issue that caused the VMA_BUFFER_BATCHING_MODE to be ignored when set from the spec.
	Keywords: VMA_BUFFER_BATCHING_MODE
	Discovered in Release: 8.4.8
	Fixed in Release: 8.4.10
1175657	Description: Fixed an issue that caused the ring statistics not to be updated when the flowtag was enabled.
	Keywords: VMA statistics, ring statistics
	Discovered in Release: 8.4.8
	Fixed in Release: 8.4.10
1192017	Description: Fixed an ibv_destroy_cq() failure scenario while VMA_TCP_3T_RULES=1. When an unexpected closing occurred while an established TCP connection occurred, VMA tried to destroy the ibv_flow from the RFS which did not create it. Now, during the RFS destruction, m_p_rule_filter's map is updated in order to destroy the correct flow.
	Keywords: VMA_TCP_3T_RULES, ibv_destroy_cq
	Discovered in Release: 8.4.8

Internal Ref.	Description
	Fixed in Release: 8.4.10
1164732	Description: Fixed a deadlock issue during routing.
	Keywords: Routing
	Discovered in Release: 8.4.8
	Fixed in Release: 8.4.10
1185978	Description: Fixed a compilation error for "--enable-opt-log=high" configuration option.
	Keywords:
	Discovered in Release: 8.4.8
	Fixed in Release: 8.4.10
1181379	Description: Fixed the TCP zero-window probe when the data is in-flight. Now VMA advances the sender's next sequence number by one as needed when sending a zero window probe.
	Keywords: TCP Zero Window Probe
	Discovered in Release: 8.4.8
	Fixed in Release: 8.4.10
1182981	Description: Fixed the TCP zero-window probe message sequence. Now VMA advances the sender's next sequence number by one as needed when sending a zero window probe.
	Keywords: TCP Zero Window Probe
	Discovered in Release: 8.4.8
	Fixed in Release: 8.4.10
1073223	Description: Enable Debian package timestamp.
	Keywords: Timestamp, Debian
	Discovered in Release: 8.3.7
	Fixed in Release: 8.4.8
863457	Description: Fixed the logic of the interface selection in the attach receiver.
	Keywords: Interface, receiver
	Discovered in Release: 8.3.7
	Fixed in Release: 8.4.8
1031172	Description: Fixed issues reported by Checkmarx security application.
	Keywords: Checkmarx security application
	Discovered in Release: 8.3.5
	Fixed in Release: 8.3.7
1043382	Description: Fixed the VMA_PROGRESS_ENGINE_WCE_MAX incorrect value in the README.txt file.
	Keywords: VMA_PROGRESS_ENGINE_WCE_MAX
	Discovered in Release: 8.3.5

Internal Ref.	Description
	Fixed in Release: 8.3.7
1037215	Description: Improved rdma_lib_reset detection.
	Keywords: rdma_lib_reset
	Discovered in Release: 8.2.10
	Fixed in Release: 8.3.5
1027871	Description: Fixed various valgrind issues.
	Keywords: valgrind
	Discovered in Release: 8.2.10
	Fixed in Release: 8.3.5
962481	Description: Fixed a completion queue error during initialization in vmapoll mode.
	Keywords: vmapoll mode
	Discovered in Release: 8.2.10
	Fixed in Release: 8.3.5
1030299	Description: Updated verbs consumer index before calling the <code>ibv_req_notify_cq</code> function in the <code>cq_mgr_mlx5</code> class.
	Keywords: IPoIB, ConnectX-4
	Discovered in Release: 8.2.10
	Fixed in Release: 8.3.5
1008712	Description: Fixed the mutual exclusion access in the ring bonding flow.
	Keywords: Mutual exclusion access
	Discovered in Release: 8.2.10
	Fixed in Release: 8.3.5
1003524	Description: Fixed a <code>fanotify_init()</code> failure in ARM 64bit architecture.
	Keywords: ARM
	Discovered in Release: 8.2.10
	Fixed in Release: 8.3.5
946914	Description: Fixed <code>epoll_create()</code> libvma incompatibility with OS. As the OS does not support creating <code>epoll</code> fds using <code>epoll_create()</code> with zero size, libvma should handle a non-positive size like OS does: return -1 and set the <code>errno</code> to <code>EINVAL</code> .
	Keywords: <code>epoll_create()</code> libvma incompatibility with OS
	Discovered in Release: 8.2.10
	Fixed in Release: 8.3.5
922411	Description: Fixed the <code>fork()</code> behavior in the child process.
	Keywords: <code>fork()</code> behavior
	Discovered in Release: 8.2.10

Internal Ref.	Description
	Fixed in Release: 8.3.5
957729	Description: Fixed a potential race condition in vma_stats.
	Keywords: vma_stats
	Discovered in Release: 8.2.8
	Fixed in Release: 8.2.10
906042	Description: Fixed a VMA errno when reaching TCP max backlog to match the OS.
	Keywords: VMA errno, TCP
	Discovered in Release: 8.2.8
	Fixed in Release: 8.2.10
973194	Description: Fixed various parameters' configuration that caused compilation errors.
	Keywords: Compilation errors
	Discovered in Release: 8.2.8
	Fixed in Release: 8.2.10
972524	Description: Optimized the data-path by removing unnecessary "if" statements and protected the pointer access to the OS API.
	Keywords:
	Discovered in Release: 8.2.8
	Fixed in Release: 8.2.10
910917	Description: Added a return error on socket API if an exception occurred
	Keywords: Socket API
	Discovered in Release: 8.2.8
	Fixed in Release: 8.2.10
1003524	Description: Fixed a fanotify_init() failure on ARM 64bit architecture.
	Keywords:
	Discovered in Release: 8.2.8
	Fixed in Release: 8.2.10
922411	Description: Fixed the fork() behavior in the child process.
	Keywords: fork() behavior
	Discovered in Release: 8.2.8
	Fixed in Release: 8.2.10
924683	Description: Fixed an src address issue upon sending MCAST with IP_MULTICAST_IF option.
	Keywords: multicast, IP_MULTICAST_IF
	Discovered in Release: 8.1.7
	Fixed in Release: 8.2.8

Internal Ref.	Description
865172	Description: Fixed an illegal memory access by netlink wrapper.
	Keywords: netlink
	Discovered in Release: 8.1.7
	Fixed in Release: 8.2.8
911076	Description: Fixed an rcvmsg failure with MSG_VMA_ZCOPY_FORCE on OS socket.
	Keywords: zero copy
	Discovered in Release: 8.1.7
	Fixed in Release: 8.2.8
888475	Description: Fixed the core dump for IB pkeys under bond.
	Keywords: IB, pkeys, bond
	Discovered in Release: 8.1.7
	Fixed in Release: 8.2.8
898505	Description: Fixed an issue which caused seg fault when connecting twice to a destination IP through a non-existing gateway.
	Keywords: Routing, gateway
	Discovered in Release: 8.1.7
	Fixed in Release: 8.2.8
888475	Description: Fixed an issue which did not allow the IB to be offloaded in case of two bonds over pkeys.
	Keywords: IB, bond, pkeys
	Discovered in Release: 8.1.7
	Fixed in Release: 8.2.8
945914	Description: Fixed epoll_create() - libvma incompatibility with OS when the number of descriptor was 0.
	Keywords: epoll_create()
	Discovered in Release: 8.1.7
	Fixed in Release: 8.2.8
-	Description: Fixed an issue which caused the machine to await for all Rx flushed WQE's to CQE's be completed when closing the QP.
	Keywords: Rx flushed WQE
	Discovered in Release: 8.0.4
	Fixed in Release: 8.1.7
-	Description: Fixed an issue which prevented VMA from receiving data on <code>epoll_wait()</code> on some cases.
	Keywords: <code>epoll_wait()</code>
	Discovered in Release: 8.0.4
	Fixed in Release: 8.1.7

Internal Ref.	Description
-	<p>Description: Fixed an issue which prevented <code>epoll</code> from being cleared from the <code>vma_stats</code>.</p> <p>Keywords: <code>epoll</code></p> <p>Discovered in Release: 8.0.4</p> <p>Fixed in Release: 8.1.7</p>
-	<p>Description: Fixed an issue which resulted in failure to receive fragmented multicast packets over ConnectX-3 InfiniBand (updated CSUM computation).</p> <p>Keywords: Fragmented multicast packets</p> <p>Discovered in Release: 8.0.4</p> <p>Fixed in Release: 8.1.7</p>
-	<p>Description: Fixed the log WARNING to DEBUG in case an <code>mlx4</code> driver did not exist on the IPoIB interface.</p> <p>Keywords: WARNING to DEBUG log</p> <p>Discovered in Release: 8.0.4</p> <p>Fixed in Release: 8.1.7</p>
-	<p>Description: Fixed error handling while registering a socket with plural <code>epoll</code> instances.</p> <p>Keywords: Error handling</p> <p>Discovered in Release: 8.0.4</p> <p>Fixed in Release: 8.1.7</p>
-	<p>Description: Fixed <code>exp-cq</code> help message when executing "<code>./configure -help</code>".</p> <p>Keywords: <code>exp-cq</code> help message</p> <p>Discovered in Release: 8.0.4</p> <p>Fixed in Release: 8.1.7</p>
-	<p>Description: Fixed segmentation fault in <code>vma_stats</code> when exceeding <code>fds</code> string length.</p> <p>Keywords: <code>vma_stats</code></p> <p>Discovered in Release: 8.0.4</p> <p>Fixed in Release: 8.1.7</p>
-	<p>Description: Modified the way partially-built QPs are destroyed.</p> <p>Keywords: QPs</p> <p>Discovered in Release: 8.0.4</p> <p>Fixed in Release: 8.1.7</p>
-	<p>Description: Fixed segmentation fault issue when running with <code>CTL_THREAD_NO_WAKEUP</code>.</p> <p>Keywords: Segmentation fault</p> <p>Discovered in Release: 8.0.4</p>

Internal Ref.	Description
	Fixed in Release: 8.1.7
-	<p>Description: Fixed an issue which prevented vma_stats from using the right shared memory for cleaning.</p> <p>Keywords: vma_stats</p> <p>Discovered in Release: 7.0.14</p> <p>Fixed in Release: 8.0.4</p>
-	<p>Description: Fixed an issue which prevented ib1 (second IB interface, port 2) from being recognized.</p> <p>Keywords: IB interface, port 2</p> <p>Discovered in Release: 7.0.14</p> <p>Fixed in Release: 8.0.4</p>
-	<p>Description: Fixed an issue which caused the machine to crash when running in a DEBUG mode and when the interface dis not have an IP address.</p> <p>Keywords: DEBUG mode</p> <p>Discovered in Release: 7.0.14</p> <p>Fixed in Release: 8.0.4</p>
-	<p>Description: Fixed the zero_copy_rx () handling of iov scattered data.</p> <p>Keywords: iov scattered data</p> <p>Discovered in Release: 7.0.14</p> <p>Fixed in Release: 8.0.4</p>
-	<p>Description: Fixed a TCP crash case when executed setsockopt () SO_RCVBUF on a listening socket.</p> <p>Keywords: setsockopt () SO_RCVBUF</p> <p>Discovered in Release: 7.0.14</p> <p>Fixed in Release: 8.0.4</p>
-	<p>Description: Fixed misleading user message about the libnl package.</p> <p>Keywords: libnl package</p> <p>Discovered in Release: 7.0.14</p> <p>Fixed in Release: 8.0.4</p>
-	<p>Description: Fixed wrong package name librdma_cm-devel in the error message.</p> <p>Keywords: librdma_cm-devel</p> <p>Discovered in Release: 7.0.14</p> <p>Fixed in Release: 8.0.4</p>
-	<p>Description: Fixed GCC 6.0.0-0.13 compilation errors.</p> <p>Keywords: GCC 6.0.0-0.13</p> <p>Discovered in Release: 7.0.14</p> <p>Fixed in Release: 8.0.4</p>

Internal Ref.	Description
-	<p>Description: Fixed TCP WIN SCALE incompatibility with RFC-1323, section 2.2/2.3.</p> <p>Keywords: TCP WIN SCALE</p> <p>Discovered in Release: 7.0.14</p> <p>Fixed in Release: 8.0.4</p>
-	<p>Description: Fixed realloc size parameter in event_handler_manager class.</p> <p>Keywords: event_handler_manager</p> <p>Discovered in Release: 7.0.14</p> <p>Fixed in Release: 8.0.4</p>
-	<p>Description: Enabled the usage of safe_mce_sys() calls instead of unsafe mce_sys calls.</p> <p>Keywords: safe_mce_sys() calls</p> <p>Discovered in Release: 7.0.14</p> <p>Fixed in Release: 8.0.4</p>
-	<p>Description: Fixed panic on TX LAG bond connections/disconnections .</p> <p>Keywords: TX LAG</p> <p>Discovered in Release: 7.0.7</p> <p>Fixed in Release: 7.0.14</p>
-	<p>Description: Fixed a segfault on VMA Server with IPerf UDP on Server closing.</p> <p>Keywords: IPerf UDP</p> <p>Discovered in Release: 7.0.7</p> <p>Fixed in Release: 7.0.14</p>
-	<p>Description: Fixed traffic over ConnectX-4 when configuring VLAN over LAG.</p> <p>Keywords: VLAN over LAG</p> <p>Discovered in Release: 7.0.7</p> <p>Fixed in Release: 7.0.14</p>
-	<p>Description: Eliminated LEX/YACC compilation warnings.</p> <p>Keywords: LEX/YACC compilation</p> <p>Discovered in Release: 7.0.7</p> <p>Fixed in Release: 7.0.14</p>
-	<p>Description: Fixed a segfault in vma_stats -v 4.</p> <p>Keywords: segfault in vma_stats</p> <p>Discovered in Release: 7.0.7</p> <p>Fixed in Release: 7.0.14</p>

Internal Ref.	Description
-	<p>Description: Fixed a crashing incident when loading VMA due to unsafe global variable values.</p> <p>Keywords: Global variable values</p> <p>Discovered in Release: 7.0.7</p> <p>Fixed in Release: 7.0.14</p>
-	<p>Description: Fixed a VMA error that occurred when running getsockopt with an invalid level and optlen.</p> <p>Keywords: getsockopt</p> <p>Discovered in Release: 7.0.7</p> <p>Fixed in Release: 7.0.14</p>
-	<p>Description: Improved msg_flag returned by recvmmsg() flow.</p> <p>Keywords: msg_flag</p> <p>Discovered in Release: 7.0.7</p> <p>Fixed in Release: 7.0.14</p>
-	<p>Description: Fixed a race condition leading to connection refuse.</p> <p>Keywords: race condition</p> <p>Discovered in Release: 7.0.7</p> <p>Fixed in Release: 7.0.14</p>
-	<p>Description: Prevented TCP accept() call from exiting when the socket closed.</p> <p>Keywords: TCP accept() call</p> <p>Discovered in Release: 7.0.7</p> <p>Fixed in Release: 7.0.14</p>
-	<p>Description: Fixed a crashing incident when FD limit was less than the required for initialization.</p> <p>Keywords: FD limit</p> <p>Discovered in Release: 7.0.7</p> <p>Fixed in Release: 7.0.14</p>
-	<p>Description: Fixed a scenario in which VMA kept spinning (burning CPU) even when it should not have to.</p> <p>Keywords: Burning CPU</p> <p>Discovered in Release: 7.0.7</p> <p>Fixed in Release: 7.0.14</p>
-	<p>Description: Fixed wrong MTU calculation.</p> <p>Keywords: MTU calculation</p> <p>Discovered in Release: 7.0.7</p> <p>Fixed in Release: 7.0.14</p>
-	<p>Description: TCP connection occasionally fails binding to INPORT_ANY, while REUSEADDR is on.</p>

Internal Ref.	Description
	<p>Keywords: TCP connection</p> <p>Discovered in Release: 6.9.1</p> <p>Fixed in Release: 7.0.7</p>
-	<p>Description: Possible deadlock when using VMA_RING_ALLOCATION_LOGIC_RX=31.</p> <p>Keywords: VMA_RING_ALLOCATION_LOGIC_RX=31</p> <p>Discovered in Release: 6.9.1</p> <p>Fixed in Release: 7.0.7</p>
-	<p>Description: For VMA_RING_ALLOCATION_LOGIC_RX=31 – CPU, affinity fails on certain systems.</p> <p>Keywords: VMA_RING_ALLOCATION_LOGIC_RX=31 – CPU</p> <p>Discovered in Release: 6.9.1</p> <p>Fixed in Release: 7.0.7</p>
-	<p>Description: Fixed TCP backlog handling.</p> <p>Keywords: TCP backlog</p> <p>Discovered in Release: 6.9.1</p> <p>Fixed in Release: 7.0.7</p>
-	<p>Description: Fixed cubic Congestion Control algorithm.</p> <p>Keywords: Congestion Control algorithm</p> <p>Discovered in Release: 6.9.1</p> <p>Fixed in Release: 7.0.7</p>
-	<p>Description: Fixed crash on exec() system cal.</p> <p>Keywords: exec() system cal</p> <p>Discovered in Release: 6.9.1</p> <p>Fixed in Release: 7.0.7</p>
-	<p>Description: Fixed a crash that occurred when the TCP Header was bigger than the packet length.</p> <p>Keywords: TCP Header</p> <p>Discovered in Release: 6.9.1</p> <p>Fixed in Release: 7.0.7</p>
-	<p>Description: Fixed wrong TCP window scale option handling.</p> <p>Keywords: TCP window scale</p> <p>Discovered in Release: 6.9.1</p> <p>Fixed in Release: 7.0.7</p>
-	<p>Description: Fixed an issue which prevented TCP connection from being closed when the SEND window was full.</p> <p>Keywords: TCP connection</p>

Internal Ref.	Description
	Discovered in Release: 6.9.1
	Fixed in Release: 7.0.7
-	Description: Fixed a crash that occurred when receiving a multithread UD.
	Keywords: multithread UD
	Discovered in Release: 6.9.1
	Fixed in Release: 7.0.7
-	Description: Fixed SO_LINGER functionality.
	Keywords: SO_LINGER functionality
	Discovered in Release: 6.9.1
	Fixed in Release: 7.0.7
-	Description: Fixed device port selection in new kernels.
	Keywords: Device port
	Discovered in Release: 6.9.1
	Fixed in Release: 7.0.7
-	Description: Fixed a TCP connection hang in syn_recv state.
	Keywords: TCP connection
	Discovered in Release: 6.9.1
	Fixed in Release: 7.0.7
-	Description: Fixed the functionality of the SO_RCVBUF and SO_SNDBUF socket options for TCP.
	Keywords: TCP
	Discovered in Release: 6.9.1
	Fixed in Release: 7.0.7
-	Description: Fixed error when giving getsockopt an invalid level value.
	Keywords: getsockopt
	Discovered in Release: 6.9.1
	Fixed in Release: 7.0.7
-	Description: Adjusted TCP RECEIVE window size to the socket RECEIVE buffer size.
	Keywords: TCP RECEIVE
	Discovered in Release: 6.9.1
	Fixed in Release: 7.0.7
-	Description: Fixed crash when there is no route back to syn sender.
	Keywords: syn sender
	Discovered in Release: 6.8.3
	Fixed in Release: 6.9.1
-	Description: Fixed issues in listen socket shutdown.

Internal Ref.	Description
	<p>Keywords: listen socket</p> <p>Discovered in Release: 6.6.4</p> <p>Fixed in Release: 6.7.2</p>
-	<p>Description: Fixed issues that caused multithread deadlocks and races in the system.</p> <p>Keywords: multithread deadlocks</p> <p>Discovered in Release: 6.5.9</p> <p>Fixed in Release: 6.6.4</p>
-	<p>Description: Fixed wrong usage of route gateway information.</p> <p>Keywords: Route gateway</p> <p>Discovered in Release: 6.5.9</p> <p>Fixed in Release: 6.6.4</p>
-	<p>Description: Fixed buffer management issues and leaks.</p> <p>Keywords: buffer management</p> <p>Discovered in Release: 6.5.9</p> <p>Fixed in Release: 6.6.4</p>
-	<p>Description: Fixed multicast loopback filtering on RX flow.</p> <p>Keywords: RX flow</p> <p>Discovered in Release: 6.5.9</p> <p>Fixed in Release: 6.6.4</p>
-	<p>Description: Fixed issues that caused multithread deadlocks and races in the system.</p> <p>Keywords: multithread deadlocks</p> <p>Discovered in Release: 6.4.11</p> <p>Fixed in Release: 6.5.9</p>
-	<p>Description: Fixed wrong handling of IGMP packets in multithread environment.</p> <p>Keywords: IGMP packets</p> <p>Discovered in Release: 6.4.11</p> <p>Fixed in Release: 6.5.9</p>
-	<p>Description: Fixed wrong usage of route gateway information.</p> <p>Keywords: route gateway information</p> <p>Discovered in Release: 6.4.11</p> <p>Fixed in Release: 6.5.9</p>
-	<p>Description: TCP close socket (active and passive sides) – buffer leaks, segmentation faults, hangs.</p> <p>Keywords: TCP</p>

Internal Ref.	Description
	Discovered in Release: -
	Fixed in Release: 6.4.11
-	Description: IGMP handling – buffer leak when having IB MC over IPR to a router.
	Keywords: IGMP handling
	Discovered in Release: -
	Fixed in Release: 6.4.11
-	Description: VMA does not handle MSG_TRUNC correctly.
	Keywords: MSG_TRUNC
	Discovered in Release: -
	Fixed in Release: 6.4.11
-	Description: TCP EPOLL on non-offloaded listen socket does not deliver events (hangs) on new connection.
	Keywords: TCP EPOLL
	Discovered in Release: -
	Fixed in Release: 6.4.11
-	Description: Receive timeout (using SO_RCVTIMEO) set to zero should block.
	Keywords: Receive timeout
	Discovered in Release: -
	Fixed in Release: 6.4.11

5 Change Log History

Table 7: Change Log History

Release	Description
8.4.10	<ul style="list-style-type: none"> • VMA TCP QUICKACK: Added TCP_QUICKACK_THRESHOLD parameter to lwip/opt.h. TCP_QUICKACK_THRESHOLD value is now set to disable quickack for payloads larger than the threshold. By default, the quickack threshold is disabled and is effective only when TCP_QUICKACK is enabled. • Delayed ack Control: Added support for delayed ack control for small packets. • Socket properties inherence: Enabled the option to inherit socket properties SO_LINGER and SO_BINDTODEVICE from the listener socket during session establishment. • Source based routing rules: Added support for source based routing rules. • Bug fixes: (see Bug Fixes History)
8.4.8	<ul style="list-style-type: none"> • Performance: Reduced latency for the following cases: <ul style="list-style-type: none"> • packet > 190bytes when using ConnectX-5 adapter cards • multi connection using ePoll • General latency reduction for all use cases • Media: Added statistics to multi packet ring • Bug fixes: (see Bug Fixes History)
8.2.8	<ul style="list-style-type: none"> • OOB User Experience: Optimized performance by measuring it via the VMA predefined specification profile for latency. Latency profile spec - optimized latency on all use cases. System is tuned to keep balance between the kernel and VMA. Note: It may not reach the maximum bandwidth. Example: VMA_SPEC = latency. • Peer-notification daemon: Peer-notification daemon implementation for half-opened TCP connections. • "Dummy Send": "Dummy Send" for improving low msg rate latency (minimize cache hit) • Socket offload : Added support for socket offload even when the link is down on VMA bringup • Hardware and Software timestamps : Added hardware and software timestamps to user callback data • Bug fixes (see Bug Fixes History)
8.1.7	<ul style="list-style-type: none"> • IGMPv3: Added support for 'Source Specific Multicast' • sockperf-2.8-0: Added support for 'source specific multicast' via the '--mc-source-filter' flag in the command line and in the feed file • IGMPv2: Removed VMA's IGMPv2 limitations • RX SW Checksum: Enabled software checksum validation for ingress TCP/UDP IP packets where the HW checksum is not supported • Lightweight IP (LWIP): Split the head of TCP unsent segment when a larger than the SEND window size and NODELAY flag is set (better adjustment to peer TCP window size advertisements) • Logging levels: Enabled excluding logging levels via configure to improve performance

Release	Description
	<ul style="list-style-type: none"> • UDP socket: Un-offloaded the UDP socket that is bound to any non-offloaded interfaces e.g. loopback (127.0.0.1) to improve performance • Auto tuning: Added auto tuning for sockperf ping-pong latency test using <code>VMA_SPEC=10</code> • Global variables: Modified the global variables usage in fast path to improve performance • Read-time stamp counter (RDTSK): Modified the RDTSK logic to consider CPU speeds taken from <code>/proc/cpuinfo</code> • fd statistics: Added fd statistics' dumping into the VMA logger using <code>vma_stats</code> (see VMA wiki page for further details) • Acting gracefully: Added an error/warning before system crash when running with Mellanox upstream drivers for non-supported features • OS connections acceptance by the server: Enabled the server to accept OS' new connections even when they are bounded to <code>INADDR_ANY</code> • DNS server: Added to the VMA configuration file (<code>libvma.conf</code>) the option of handling the DNS server (UDP port 53) by the OS by default • Coverity: Cleaned Coverity error reports based on Coverity 8.5 version • Bug fixes (see Bug Fixes History)
8.0.4	<ul style="list-style-type: none"> • Changed VMA license from "GPLv2 or commercial" to "GPLv2 or BSD" (dual license) – see Appendix A: VMA Licensing Model for more details • Added support for Fedora 24 • Added the ability to return ENOMEM if select() uses 'ndfs' larger than the system's <code>FD_SETSIZE</code> • Updated the neigh table search using both <code>dst addr</code> and <code>ifindex</code> • Converted <code>m_ready_fds</code> (<code>epoll_info.h</code>) from <code>std::tr1::unordered_map</code> to <code>vma_list</code> • Converted <code>socket_fd_vec</code> (<code>epoll_wait_call.cpp</code>) from <code>std::vector</code> into <code>vma_list</code> • Added support for multiple node names in <code>vma_list</code> • Expanded <code>vma_stats</code> to include <code>Send-Q</code> and <code>Recv-Q</code> • Added internal thread TCP timer handling mode: <code>immediate/deferred</code> • Added <code>vma_stats -v 5 PID</code> column (netstat like) • Added <code>--disable-exp-cq</code> parameter to control CQ polling using experimental verbs API • Disabled VMA's CQ interrupt moderation logic in case of <code>MLNX_OFED's</code> experimental CQ moderiation API are missing • Added jenkins check support on github pull request • Added support for Intel compiler • Updated logging on <code>iomux</code> flows • Limited documentation of cased that use <code>VMA_RING_ALLOCATION_LOGIC_RX TX</code> with <code>select()/poll()</code> • Added the option to skip 'OFED Version' logging in VMA banner if <code>ofed_info</code> is not supported
7.0.14	<ul style="list-style-type: none"> • ConnectX@-4 EN and ConnectX-4 Lx EN support • Bonding LAG (mode=4) support • VLAN over bond support • UDP RX hardware Time Stamping support • UDP TX flow optimization

Release	Description
	<ul style="list-style-type: none"> • Socket API exception flows handling • TCP Time Stamping support • Software CSUM calculation for L3 (IP) & L4 (TCP) layers in environments where the TX hardware CSUM is not supported (L4/UDP CSUM=0) • RX/TX stddev and latency improvements for UDP and TCP sockets
7.0.7	<ul style="list-style-type: none"> • Unconventional VLAN names support • Added an option to avoid system calls on TCP socket • Connection establishment rate improvements • TCP receive window and socket receive buffers cohesion improvement • Optimized accept() call to consume less CPU time
6.9.1	<ul style="list-style-type: none"> • Reduced contention between VMA internal thread and the user threads • Moved control flow tasks to VMA internal thread to improve latency • Added syn/fin throttling support to improve latency during connection control tasks • Added support for creation of vma_stats shared memory files in a given directory • Added retransmission counters to vma_stats • Added RX CSUM verification offload support (depending on device support) • Fixed DEVICE_FATAL event to support hot-unplug
6.8.3	<ul style="list-style-type: none"> • Added support for all Linux OSs supported in MLNX_OFED v2.4-x.x.x • Added support for TCP zero copy in the extra API
6.7.2	<ul style="list-style-type: none"> • Added support for all Linux OSs supported in MLNX_OFED v2.3-1.0.X • Added support for routing rules and secondary route tables • Added support for ARM 64 bit architecture (at beta level) • Added support for PowerPC 64 bit architecture (at beta level)
6.6.4	<ul style="list-style-type: none"> • Added support to all Linux Operating Systems supported in MLNX_OFED 2.2-1.0.0 • Improved interrupt driven mode performance • Added interrupt moderation and adaptive interrupt moderation support • Added UDP software timestamps support • Added support for accept4 system call • Added support for SO_BINDTODEVICE socket option • Added support for SOCK_NONBLOCK and SOCK_CLOEXEC socket() flags • Added ring statistics to vma_stats
6.5.9	<ul style="list-style-type: none"> • Added support for all Linux OSs supported in MLNX_OFED 2.1-1.0.0 • Improved TCP latency performance • Improved VMA blacklist • Added EXTRA API to control off-load capabilities • Added TCP cubic congestion control algorithm • Increased the number of supported sockets to thousands of sockets • Added IP_PKTINFO support in UDP recvmsg • ETH loopback support
6.4.11	<ul style="list-style-type: none"> • Ubuntu 12.04 OS support • Improved TCP stability • Improved performance of TCP throughput

Release	Description
	<ul style="list-style-type: none">• Improved performance of applications with multiple epoll instances• Added TCP window scaling and Generic Receive Offload (GRO)• Support for Ethernet Unicast loopback via Kernel network stack (not offloaded)• Support sendmmsg• Support epoll_pwait, pselect, and ppoll• Support ip_mreqn in setsockopt ADD/DROP_MEMBERSHIP and MULTICAST_IF• Support TCP socket MSG_PEEK recv flag• Support TCP socket MSG_DONTWAIT send flag• Support getsockopt() SOL_SOCKET SO_ERROR

Appendix A: VMA Licensing Model

As of VMA v8.0.4, VMA is licensed under a dual license of either GPL version 2 or BSD. Earlier versions of VMA were released under GPL version 2 license which is a restrictive license with limitations and obligations that forced the customers to obtain a license for VMA as proprietary software, i.e. subject to payment.

VMA's new licensing model enables customers to use either GPL version or BSD. The BSD license is permissive, enabling the customers to distribute and modify VMA, without the obligation to contribute their modifications to the community, all subject to the BSD terms.