



Mellanox SwitchX® Firmware Release Notes

Rev 9.4.2000

Last Modified: November 29, 2016

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 NON-INFRINGEMENT 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 2016 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®, InfiniScale®, Kotura®, Kotura logo, Mellanox Federal Systems®, Mellanox Open Ethernet®, Mellanox ScalableHPC®, Mellanox TuneX®, Mellanox Connect Accelerate Outperform logo, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, NP-1c®, NP-2®, NP-3®, Open Ethernet logo, PhyX®, PSIPHY®, SwitchX®, Tiler®, Tiler logo, TestX®, TuneX®, The Generation of Open Ethernet logo, UFM®, 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

Table of Contents	3
Chapter 1 Overview	4
1.1 Supported Systems in 9.4.2000	4
1.2 Supported Cables and Modules	4
1.3 Firmware Upgrade	4
1.4 PRM Revision Compatibility	5
Chapter 2 Changes and Major New Features	6
2.1 Changes in Rev 9.4.2000	6
2.2 Changes in Rev 9.4.1000	6
2.3 Changes in Rev 9.3.8000	6
2.4 Changes in Rev 9.3.6000	6
2.5 Changes in Rev 9.3.4000	6
2.6 Changes in Rev 9.3.2000	6
2.7 Changes in Rev 9.3.1200	6
2.8 Changes in Rev 9.3.0000	6
2.9 Changes in Rev 9.2.8000	7
2.10 Changes in Rev 9.2.6100	7
2.11 Changes in Rev 9.2.4002	7
2.12 Changes in Rev 9.2.3000	7
2.13 Changes in Rev 9.2.0000	8
2.14 Changes in Rev 9.1.7000	8
2.15 Changes in Rev 9.1.5000	8
2.16 Changes in Rev 9.1.3000	8
2.17 Changes in Rev 9.1.2000	8
2.18 New Features and Changes in Rev 9.1.0000	8
Chapter 3 Known Issues	9
Chapter 4 Bug Fixes History	10

1 Overview

These are the release notes for the SwitchX® and SwitchX-2 firmware, 9.4.2000. This firmware complements the SwitchX silicon architecture with a set of advanced features, allowing easy and remote management of the switch.



Firmware 9.4.2000 is compatible with MFT tools v2.7.1 and above. Using older versions of MFT tools causes failure during image creation.

1.1 Supported Systems in 9.4.2000

This firmware supports the devices and protocols listed in Table 1. For the most updated list of switches supported, visit the firmware download pages via <http://www.mellanox.com>.

Table 1 - Supported PSIDs

Device Part Number	PSID	Device Description
MSX6015F-xxxS	MT_1260110020	SwitchX®-2 based FDR InfiniBand Switch, 18 QSFP ports, externally managed
MSX6015T-xxxS	MT_1260110029	SwitchX®-2 based FDR10 InfiniBand Switch, 18 QSFP ports, externally managed
MSX6025F-xxxR	MT_1010210021	SwitchX® based FDR InfiniBand Switch; 36 QSFP; externally managed
MSX6025F-xxxS	MT_1010310021	SwitchX®-2 based FDR InfiniBand Switch; 36 QSFP; externally managed
MSX6025T-xxxR	MT_1010210026	SwitchX® based FDR10 InfiniBand Switch; 36 QSFP; externally managed
MSX6025T-xxxS	MT_1010310026	SwitchX®-2 based FDR10 InfiniBand Switch; 36 QSFP; externally managed
MSX6005F-xxxS	MT_1260110021	SwitchX®-2 based FDR InfiniBand Switch; 12 QSFP ports.
MSX6005T-xxxS	MT_1260110026	SwitchX®-2 based FDR10 InfiniBand Switch; 12 QSFP ports.

1.2 Supported Cables and Modules

Please refer to the interconnect overview page on the Mellanox website:

http://www.mellanox.com/page/cables?mtag=cable_overview

1.3 Firmware Upgrade

Firmware upgrade may be performed directly from any previous version to this version. To upgrade firmware, please refer to the Mellanox Firmware Tools (MFT) package at:

http://www.mellanox.com/page/management_tools

1.4 PRM Revision Compatibility

Firmware 9.4.2000 complies with the following programmer's reference manual:

- *Switches Programmer's Reference Manual (PRM), Rev 1.40* or later

2 Changes and Major New Features

2.1 Changes in Rev 9.4.2000

- N/A

2.2 Changes in Rev 9.4.1000

- Added system MKey support

2.3 Changes in Rev 9.3.8000

- Support FORE EMAD query operation according to section 2.17.5 of the SwitchX® PRM rev 1.05
- Added support to read UCD sensor and V_{in} sensing according to section 2.8.5 of the SwitchX® PRM rev 1.05
- Added AccessRegisterGMP support for MSPS and fix MSPS firmware configuration

2.4 Changes in Rev 9.3.6000

- Added support for bandwidth monitoring according to Performance histogram in Vendor Specific MAD spec

2.5 Changes in Rev 9.3.4000

- Added support for fan-out-of-range event (FORE) according to PRM section “FORE – Fan Out of Range Event”
- Added LLR max retransmission rate as specified in Vendor Specific MAD V1.1, Table 110 – PortLLRStatistics MAD Description

2.6 Changes in Rev 9.3.2000

- Lowered the externally managed switches default fan speed from 80% to 60%
- Added congestion control attribute “PORT_VL_XMIT_TIME_CONG” according to 3.23 CounterGroupInfo in Vendor Specific MADs version 1.1
- Added support for adaptive routing per L4 transport layer

2.7 Changes in Rev 9.3.1200

- Added GA support for MFA1A00-Exxx EDR 100Gb/s AOCs over SwitchX® based systems

2.8 Changes in Rev 9.3.0000

- Added support for SX6710 and SX1710 switch systems
- Added support for timestamp in mirrored packages
- Added support for MCP1600-E001/MCP1600-E002 EDR cables
- Added support for FDR 2.5m (MC2207128-0A2)

2.9 Changes in Rev 9.2.8000

- Added power optimizations when ports are disabled
- Added TuneX support
- Bug fixes

2.10 Changes in Rev 9.2.6100

- Added LLR capability indication as part of extended port info MAD according to Vendor Specific MADs document version 1.1
- 3rd party FDR cables are allowed to raise link at FDR (56Gb/s)
- Added support for fan and power supply monitor MAD according to section 2.12 in Vendor Specific MADs document version 1.1 and section 2.17 in Mellanox PRM version 1.0
- Improved FDR systems link-up time after reboot

2.11 Changes in Rev 9.2.4002

- Added support for 12-port SX6005 switch system
- Added ability to Set Node Description according to errata proposal #9109
- Updated SMP and GMP CapabilityMask as appear in VS Mad Specification document version 0.56
- Added support for temperature sensor MAD according to section 2.12 in Vendor Specific MADs document version 1.1 and section 2.17 in Mellanox PRM version 1.0
- Added support for voltage sensor MAD
- Added a port disable power-saving feature
- Added support for reading VPD from flash (supported with systems coming out of factory with FW 9.2.4002 or higher only)
- Added support for alternate source of the following fiber cable P/Ns: MC2207310-XXX-F-A2, MC2210310-XXX-F-A2, MC2206310-XXX-F-A2 (See PCN MLNX-15-838)
- Added support for the MC2207312-XXX 56Gb/s active fiber cables.
- Added support for MC2207310-XXX-F-A2, MC2207310-XXX-F-A2, MC2206310-XXX-F-A2 active fiber cables
- Added support for MC2207411-SR4 56Gb/s active optical module

2.12 Changes in Rev 9.2.3000

- Improved FDR and FDR10 link stability
- Port mirroring bug fixes
- Added best-effort mirroring capability
- Fixing FDR10 switch link issue discovered in 9.2.0000

2.13 Changes in Rev 9.2.0000

- Added support for SwitchX®-2 SX6015
- Added support for Congestion Control Class
- Added support for InfiniBand multi-swid
- Enabled MultiCastFDBTop=0xBFFF to discard MC traffic
- Added a new INI field for the NodeInfo and DeviceID parameters

2.14 Changes in Rev 9.1.7000

- Enabled LLR by default
- Added Mellanox Proprietary Link negotiation (MLPN) support
- Added support for MC2207312-XXX FDR optical cables up to 300m
- Added support for SwitchX®-2 based FDR and FDR10 edge switch systems
- Improved Eye Opener Machine

2.15 Changes in Rev 9.1.5000

- Added InfiniBand Port Mirroring support
- Reduced port PLLs jitter
- Enabled Gradient Decent (GDS) in MSX6025T and MSX6025F systems
- Added optimize SL2VL support
- Improved link negotiation process

2.16 Changes in Rev 9.1.3000

- Link negotiation protocols improvements

2.17 Changes in Rev 9.1.2000

- Added PKEY enforcement on external ports

2.18 New Features and Changes in Rev 9.1.0000

- InfiniBand switch supported speed:
 - FDR v1.3
 - FDR10 - Mellanox Technologies propriety
 - QDR/DDR/SDR v1.2.1
- 36 FDR (56Gb/s) ports in a 1U switch
- FDR10 supports 20% more bandwidth over QDR using the same cables/connectors

3 Known Issues

Table 2 describes known issues in this firmware release and possible workarounds.

Table 2 - Known Issues

Index	Category	Description	Workaround
1.	Counters	The SampleCounter mechanism is not supported.	N/A
2.	Counters	Resetting the counters of port receive packets and port receive unicast packets using “Performance class MAD” or the VL15 received packets and data counters using Mellanox Vendor Specific MAD resets all of the above counters.	N/A
3.	Network Interfaces	FDR link may raise as FDR10 when plugged in near an FDR10 cable.	N/A
4.	Network Interfaces	Link width negotiation is not supported.	N/A
5.	Network Interfaces	Link up negotiation may take up to 70 seconds.	N/A
6.	Network Interfaces	DDR link may rise with errors or at SDR speed when connecting to ConnectX®-3 version 2.30.3000 or earlier, Connect-IB™ version 10.10.1000 or earlier, or previous releases of SwitchX®.	N/A
7.	Switch Capabilities	QDR capabilities at the INI must be enabled when working in FDR mode.	N/A
8.	General	Downgrading to firmware version 9.2.0000 requires power cycling the switch for the firmware to load.	N/A
9.	Cables	Loopback modules are not supported and do not link up.	N/A
10.	Link	56GbE ports may take up to 60 seconds to link up.	N/A

4 Bug Fixes History

Table 3 lists the bugs fixed in this release.

Table 3 - Bug Fixes History

Index	Category	Description	Discovered in Release	Fixed in Release
1.	LEDs	UID LED MAD is not responsive.	N/A	9.4.1000
2.	Chassis Management	PSU fans set to work with 60% max speed by default	9.3.8000	9.4.1000
3.	Chassis Management	The command “show interfaces ib * transceiver” shows no cable is connected while link is up.	9.3.8000	9.4.1000
4.	Power	AC/DC status info not available when reading address 0x900D (MSPS via in-band MAD).	9.3.6000	9.3.8000
5.	Network Interfaces	Cannot enable 40GbE on backplane ports	9.3.4000	9.3.6000
6.	System Management	Reset cannot be performed during I ² C operation	9.3.4000	9.3.6000
7.	Adaptive routing	Fixed case of traffic loss between node pairs	9.3.2000	9.3.4000
8.	pLFT	When activating private linear forwarding table (pLFT), last port was not mapped to 0 by default	9.3.2000	9.3.4000
9.	IB spec	MLID limit at 0xFDE8 instead of 0xFE00	9.3.2000	9.3.4000
10.	Latency	Reduced forwarding latency when Skip symbols were lost	9.3.2000	9.3.4000
11.	Links	Fixed fast FDR mechanism issue	9.3.2000	9.3.4000
12.	Latency	Fixed latency issues in SX65xx switch systems	9.3.2000	9.3.4000
13.	LEDs	Unexpected LED behavior once a LED Vendor Specific MAD configures any port LED to a blinking mode.	9.3.0000	9.3.2000
14.	Links	Fix for parts going down randomly.	9.3.0000	9.3.2000
15.	Chassis Management	Fix fan error indication when changing AC modules.	9.3.0000	9.3.2000
16.	Cables	Fix for supporting SFP+ cables in force FDR10 mode.	9.2.8000	9.3.0000
17.	VS MADs	Fix MSGI register to return data with little Endian instead of big Endian.	9.2.8000	9.3.0000
18.	Counters	Congestion control supported counters in the Sample counters option mask is missing.	9.2.8000	9.3.0000
19.	Counters	PortRevPkts counter is miscalculated when LLR is enabled.	9.2.8000	9.3.0000

Table 3 - Bug Fixes History

Index	Category	Description	Discovered in Release	Fixed in Release
20.	Counters	Port receive packets and port receive unicast packets counters may show incorrect values when port is down.	9.2.4002	9.2.8000
21.	Health	Fixed temperature sensing MAD.	9.2.6000	9.2.8000
22.	Network Interfaces	Ports may stay at INIT state after power-cycle.	9.2.6000	9.2.6100
23.	IB Spec	Fixed Neighbor MTU default size to be 256, as defined in IB SPEC.	9.2.4002	9.2.6000
24.	IB Spec	Set operation in MADs PortXmitWaitVL and PortXmitWaitVLExtended does not function properly when setting VL15.	9.2.4002	9.2.6000
25.	IB Spec	Wrong behavior of trap 128 when physical link is moving to “down”.	9.2.3000	9.2.4002
26.	Links	Link may have stability issues after closing a port.	9.2.3000	9.2.4002
27.	Links	Cables limited only to QDR may not rise links at all.	9.2.3000	9.2.4002
28.	LLR	Rate counters are not reset upon request.	9.2.3000	9.2.4002
29.	Cables	Cable insertion internal interrupts may be lost in rare cases resulting in the link not rising.	9.2.3000	9.2.4002
30.	Remote IB Mirroring	Packet always sent with vl=0.	9.2.0000	9.2.3000
31.	Remote IB Mirroring	Mirrored packets sent from the switch are corrupted if their size is bigger than 64 bytes.	9.2.0000	9.2.3000
32.	Port Mirroring	Truncation does not properly crop the packet.	9.2.0000	9.2.3000
33.	Cables	Link issues when working with Finisar (non-Mellanox) non-FDR AOC cables.	9.2.0000	9.2.3000
34.	Links	FDR10 switch fails to raise links when several links are disconnected at the remote side.	9.2.0000	9.2.3000
35.	Traps	SwitchX® does not send Congestion Control traps.	9.2.0000	9.2.3000
36.	Routing	SL2VL mapping issue which may render the SwitchX® unresponsive.	9.1.7000	9.2.0000
37.	Traps	Trap257 and Trap258 are not supported.	9.1.7000	9.2.0000
38.	Counters	Performance counters option mask corrupted.	9.1.7000	9.2.0000
39.	Counters	Wrong behavior of PortRcvDataSL & PortXmit-DataSL counters.	9.1.7000	9.2.0000

Table 3 - Bug Fixes History

Index	Category	Description	Discovered in Release	Fixed in Release
40.	LLR	Wrong behavior of PortUnicastRcvPkts counter when LLR is active.	9.1.7000	9.2.0000
41.	Network Interfaces	The link up is raised as SDR or is not raised when connected to the Qlogic switch.	9.1.7000	9.2.0000
42.	FRU LEDs	On all externally managed 1U switches, FRU LEDs remain OFF after boot.	9.1.5000	9.1.7000
43.	Cables	Fixed support for MFS4R12CB-XXX FDR10 optical cables up to 100m.	9.1.5000	9.1.7000
44.	Cables	When repeatedly plugging out and plugging in cables, the cable might not be recognized, causing the links to remain down.	9.1.2000	9.1.3000
45.	Cables	FDR10 cable P/N MFS4R12CB-XX is supported in ports 9-36 only.	9.1.0000	9.1.2000
46.	Network Interfaces	The link is raised as SDR.	9.1.0000	9.1.2000
47.	Network Interfaces	InfiniBand link vs FDR device rises as DDR.	9.1.0000	9.1.2000
48.	Network Interfaces	When sending PortInfo MAD to an Active port using the POLLING command, the link might rise as SDR.	9.1.0000	9.1.2000