



Mellanox MLNX-OS® Release Notes for VPI

Software Ver. 3.4.1110

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

Mellanox Technologies, Ltd.
Hakidma 26
Ofer Industrial Park
Yokneam 2069200
Israel
www.mellanox.com
Tel: +972 (0)74 723 7200
Fax: +972 (0)4 959 3245

© Copyright 2015. Mellanox Technologies. All Rights Reserved.

Mellanox®, Mellanox logo, BridgeX®, ConnectX®, Connect-IB®, CoolBox®, CORE-Direct®, GPUDirect®, InfiniBridge®, InfiniHost®, InfiniScale®, Kotura®, Kotura logo, MetroX®, MLNX-OS®, PhyX®, ScalableHPC®, SwitchX®, TestX®, UFM®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

ExtendX™, FabricIT™, FPGADirect™, HPC-X™, Mellanox CloudX™, Mellanox Open Ethernet™, Mellanox PeerDirect™, Mellanox Virtual Modular Switch™, MetroDX™, NVMeDirect™, Switch-IB™, Unbreakable-Link™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

Table of Contents

Chapter 1	Introduction	4
Chapter 2	Supported Platforms, Firmware, Cables and Licenses	4
2.1	Supported Switch Systems	4
2.2	Supported CPU Architecture	5
2.3	Supported Firmware	5
2.4	Supported UFM® Release	5
2.5	Supported Mellanox OFED Release	5
2.6	Supported Software Licenses	5
2.7	Upgrade From Previous Releases	6
2.8	Supported Cables	6
Chapter 3	Changes and New Features	7
Chapter 4	Known Issues	13
4.1	General Known Issues	13
4.2	Ethernet Known Issues	18
4.3	IP Routing Known Issues	24
4.4	InfiniBand Known Issues	28
Chapter 5	Bug Fixes	33
Chapter 6	Submitting a Service Request	35

1 Introduction

This document is the Mellanox MLNX-OS® Release Notes for VPI.

MLNX-OS is a comprehensive management software solution that provides optimal performance for cluster computing, enterprise data centers, and cloud computing over Mellanox SwitchX® family. The fabric management capabilities ensure the highest fabric performance while the chassis management ensures the longest switch up time.

The MLNX-OS documentation package includes the following documents:

- User Manual – provides general information about the scope, organization and command line interface of MLNX-OS as well as basic configuration examples
- Release Notes – provides information on the supported platforms, changes and new features, and reports on software known issues as well as bug fixes

2 Supported Platforms, Firmware, Cables and Licenses

2.1 Supported Switch Systems

Table 1 - Supported Switch Systems

Model Number	Description
MSX6036F	36-Port FDR (56Gb/s) 1U InfiniBand switch system
MSX6036T	36-Port FDR10 (40Gb/s) 1U InfiniBand switch system
MSX6018F	18-Port FDR (56Gb/s) 1U InfiniBand switch system
MSX6018T	18-Port FDR10 (40Gb/s) 1U InfiniBand switch system
MSX6012F	12-Port FDR (56Gb/s) 1U InfiniBand switch system
MSX6012T	12-Port FDR10 (40Gb/s) 1U InfiniBand switch system
MSX6536	648-Port FDR (56Gb/s) / FDR10 (40Gb/s) InfiniBand switch system
MSX6518	324-Port FDR (56Gb/s) / FDR10 (40Gb/s) InfiniBand switch system
MSX6512	216-Port FDR (56Gb/s) / FDR10 (40Gb/s) InfiniBand switch system
MSX6506	108-Port FDR (56Gb/s) / FDR10 (40Gb/s) InfiniBand switch system
MSX6710	36-Port FDR (56Gb/s) 1U InfiniBand dual core x86 switch system
MSX6720	36-Port FDR (56Gb/s) 1U InfiniBand dual core x86 switch system with battery backup unit (BBU) support
MSB7700	36-Port EDR (100Gb/s) 1U InfiniBand dual core x86 switch system
MSX1036	36-Port QSFP 40GbE 1U Ethernet switch system
MSX1035	36-Port QSFP 40GbE 1U Ethernet Lite Layer 2 switch system
MSX1024	12-Port QSFP 40GbE 48-Port SFP+ 10GbE 1U Ethernet switch system
MSX1016	64-Port SFP+ 10GbE 1U Ethernet switch system

Table 1 - Supported Switch Systems

Model Number	Description
MSX1012	12-Port QSFP (40GbE) 1U Ethernet switch system
MSX1400	48-Port SFP+ 10GbE, 12-Port QSFP 40GbE 1U x86 Ethernet switch system
MSX1710	36-Port QSFP 40GbE 1U Ethernet dual core x86 switch system
MTX6000	16 long-haul FDR10 (40Gb/s) and 16 FDR (56Gb/s) down link long-haul switch system (MetroDX™)
MTX6100	6 long-haul FDR10 (40Gb/s) and 6 FDR (56Gb/s) down link long-haul switch system (MetroX®)
MTX6240	2 long-haul FDR10 (40Gb/s) and 2 FDR (56Gb/s) down link long-haul switch system (MetroX®)
MTX6280	1 long-haul FDR10 (40Gb/s) and 1 FDR (56Gb/s) down link long-haul switch system (MetroX®)

2.2 Supported CPU Architecture

- PPC 460
- x86

2.3 Supported Firmware

- SwitchX® firmware version 9.3.1750
- SwitchX®-2 firmware version 9.3.1750
- Switch-IB™ firmware version 11.0.1250
- ConnectX®-2 firmware version 2.9.1200 and higher
- ConnectX®-3 firmware version 2.30.8000 and higher
- Connect-IB® firmware version 10.10.4020 and higher

2.4 Supported UFM® Release

The Unified Fabric Manager® software version supported in this MLNX-OS® release is UFM® 4.9.

2.5 Supported Mellanox OFED Release

The Mellanox OpenFabrics Enterprise Distribution for Linux driver version supported in this MLNX-OS® release is OFED 2.3.

2.6 Supported Software Licenses

For the software licenses supported with MLNX-OS® software please refer to the “Licenses” section of the “Getting Started” chapter of the *Mellanox MLNX-OS User Manual*.

2.7 Upgrade From Previous Releases

Older versions of MLNX-OS may require upgrading to one or more intermediate versions prior to upgrading to the latest. Missing an intermediate step may lead to errors. Please refer to Table 2 and Table 3 to identify the correct upgrade order.

Table 2 - Supported Software Upgrades for PPC Systems

Target Version	Verified Versions From Which to Upgrade
3.4.1110	3.4.1100 (for Ethernet only)3.4.0012
3.4.0012	3.3.5006; 3.3.4402
3.3.5006	3.3.4402; 3.3.4304
3.3.4402	3.3.4304; 3.3.4100
3.3.4304	3.3.4100; 3.3.3500 (1U only); 3.2.0330-100 (director switches)
3.3.4100	3.3.3500 (1U only); 3.3.3000 (1U only); 3.2.0330-100 (director switches)

Table 3 - Supported Software Upgrades for x86 Systems

Target Version	Verified Versions From Which to Upgrade
3.4.1110	3.4.0012; 3.3.5006
3.4.0012	3.3.5006
3.3.5006	3.3.4406; 3.3.4304

For upgrade instructions refer to the section “Upgrading MLNX-OS Software” in *Mellanox MLNX-OS User Manual*.



Downgrading from MLNX-OS version 3.4.0012 or later on x86 switch systems is not supported.

2.8 Supported Cables

For a list of the Mellanox supported cables please visit the LinkX™ Cables and Transceivers page of the Mellanox Website at http://www.mellanox.com/page/cables?mtag=cable_overview.



When using Mellanox AOC cables longer than 50m use one VL to achieve full wire speed.

3 Changes and New Features

Table 4 - VPI Changes and New Features

Category	Description
Release 3.4.1110	
General	Improved average CPU consumption in SX65xx systems
General	Improved CLI login performance in SX65xx systems
LACP	Added support for specific port in the command “show lacp neighbor”
Ethernet Switching	Added GA support for RPVST See section “RPVST” in the User Manual
Ethernet Switching	Added support for LACP individual mode See the command “suspend-individual” in the User Manual
Ethernet Switching	OpenFlow bitwise IP match
Ethernet Switching	Added support for BPDU guard See the command “spanning-tree port type” in the User Manual
IP Routing	Added GA support for IPv6 See section “IPv6 Routing” in the User Manual
IP Routing	Added support for VRRP counters
IP Routing	Added prefix filter in the command “show ip route <ip-address>”
MAGP	Added “MAGP disabled” indication in the command “show magp” when applicable
SNMP	Added support for SNMP SET upgrade See section “Upgrading MLNX-OS Software with SNMP” in the User Manual
WebUI	Added popup Welcome screen when connecting via WebUI See section “Starting the Web User Interface” in the User Manual
Systems	Added engineering sample support for MSX6720 switch systems See section “Back-Up Battery Units” in the User Manual
Systems	Added beta support for MSB7700 switch system
Security	Added default passwords to the XML default users See section “User Accounts” in the User Manual
Release 3.4.0012	
Security	Changed the HTTPS default ciphers to TLS.
IP Routing	GA support for PIM Bi-Dir. See section 6.3 “Multicast (IGMP and PIM)” in the User Manual.
IP Routing	Added support for ICMP redirect. See command “ip icmp redirect” in the User Manual.

Table 4 - VPI Changes and New Features

Category	Description
Configuration Management	Added support for embedded SM rev 4.0.5.
Ethernet Switching	Added support for disabling MAC-address learning.
Gateway	GA support for Proxy-ARP multicast in standalone mode. See chapter 8 “Gateway” in the User Manual.
General	Added support for Mellanox OFED 2.3 integration.
Interconnect	Added support for LR4 modules on Ethernet switch systems.
Interconnect	Added support for MC220731V FDR optic cables.
Interconnect	Added support for new EDR MCP1600-E0 copper cables.
IP Routing	GA support for Router Port. See Chapter 6 “IP Routing” in the User Manual.
IP Routing	GA support for BGP. See section 6.3 “BGP” in the User Manual.
IP Routing	GA support for route maps. See section 6.4 “Route Map” in the User Manual.
IGMP Querier	Beta support for IGMP Querier. See section 5.8.3 “IGMP Snooping Querier” in the User Manual.
MLAG	Added support for MLAG LACP.
Security	Added support for security strict mode.
SNMP	Added support for setting hostname through SNMP. See section 4.17.1.7 “SNMP SET Operations” in the User Manual.
SNMP	Added support for power cycling through SNMP. See section 4.17.1.7 “SNMP SET Operations” in the User Manual.
SNMP	Added support for changing configuration through SNMP. See section 4.17.1.7 “SNMP SET Operations” in the User Manual.
SNMP	Added SNMP traps for VLAN Interface, MLAG port channel and Port channel link up/down events.
SNMP	Added MSTP MIBs/traps support.
Systems	GA support for SX6710 and SX1710 switch systems.
Virtual Machine	GA support for Virtual Machine for x86-based switch systems. See section 4.19 “Virtual Machine” in the User Manual.
WebUI	Added support for Internet Explorer 11 web browser.
Release 3.3.5006	
General	Improved cable info read response time. See the command “show interfaces ib transceiver”.

Table 4 - VPI Changes and New Features

Category	Description
Systems	GA support for TX6000.
Ethernet Switching	GA support for MLAG.
ACLs	Added MAC-based VLAN ACL.
IP Routing	GA support for MAGP.
IP Routing	GA support for VRRP.
IP Routing	GA support for IGMP.
IP Routing	Alpha support for PIM.
OSPF	Loopback interface on OSPF.
OSPF	Added new OSPF MIB/Traps.
SNMP	Added cable info entries to entPhysicalTable.
SNMP	Added support for SNMP to trigger SNMP test trap via SNMP set command. See section “MLNX-EFM MIB”.
SNMP	Added system identifier (MAC address) to test trap.
Gateway	GA support for High Availability Gateway. See section “Proxy-ARP HA” in the User Manual.
Release 3.3.4402	
General	Added new certificate hashing algorithm (sha256). See section “Cryptographic (X.509, IPSec)” in the User Manual.
General	Added new Fiber Channel over Ethernet License. See section “Licenses” in the User Manual.
General	The command “show configuration full” is no longer supported.
CLI	Added support for command output filtering. See section “Command Output Filtering” in the User Manual.
Gateway	Added GA support for multiple Proxy-ARP. See chapter “Gateway” in the User Manual.
Gateway	Added beta-level support for High Availability Gateway. See section “High Availability Proxy-ARP” in the User Manual.
Gateway	Added support for 512 Ethernet hosts.
DCBX	Added support for LLDP-DCBX CEE TLVs.
SNMP	Added power supply of SX6012 to Entity MIB.
SNMP	SX65xx support for notification on internal link up/down state.
SNMP	SX65xx support for notification on internal link sub-optimal speed.

Table 4 - VPI Changes and New Features

Category	Description
Release 3.3.4304	
General	Added End-User License Agreement. See section “Getting Started” in the User Manual.
General	Improved configuration file format.
General	3rd party FDR cables are allowed to raise link at FDR (56Gb/s)
Power Management	Added support for link width reduction.
Systems	New CPU temperature for MSX10xx systems.
Systems	Beta support for x86 CPU family-based systems (MSX1400 and MSX1700).
Systems	GA support for MTX62x0 MetroX® family
Puppet Agent	GA support for Puppet. See section “Puppet Agent” in the User Manual.
Ethernet Switching	GA support for debug per protocol.
OpenFlow	GA support for OpenFlow. See section “OpenFlow” in the User Manual.
DCBX	GA support for DCBX. See section “DCBX” in the User Manual.
MSTP	GA support for MSTP protocol. See section “MSTP” in the User Manual.
DHCP Relay	GA support for DHCP relay. See section “DHCP Relay” in the User Manual.
IGMP Snooping	GA support for IGMP Querier relay. See section “IGMP Snooping Querier” in the User Manual.
SNMP	Added ability to display leaf/spine power state via SNMP.
Gateway	Beta level support for multiple Proxy-ARP. See chapter “Gateway” in the User Manual.
Release 3.3.4102	
General	Bug fix for MTX6100 port configuration.
Release 3.3.4100	
General	Improved debug file upload mechanism. Refer to “file debug-dump” command in the CLI reference guide.
General	Added support for displaying system hardware revision. Refer to “show inventory” command in the CLI reference guide.
General	Improve initialization time after upgrade procedure on SX65xx systems.

Table 4 - VPI Changes and New Features

Category	Description
SNMP	Added a MELLANOX-ENTITY-MIB as an extension to the standard ENTITY-MIB to represent system GUID. Refer to SNMP chapter in the User Manual. The new MIB can be found in Mellanox support website.
SNMP	Added support for leaf/spine fatal state on the standard ENTITY-STATE-MIB.
Logging	Added support for event notification to monitor. Refer to “Event Notification” section in the User Manual.
Chassis Management	Improved temperature control algorithm.
User Interfaces	Improved login timeout mechanism.
User Interfaces	Added ability to display SM log from the CLI. Refer to “show ib sm log” command in the CLI guide.
SM HA	Added ability to control HA traffic. Refer to “ha dns enable” command in the CLI guide.
VPI	VPI flexible ports on SX6018, SX1036, SX6036, and SX6036G switch systems. Refer to VPI section in the User Manual.
Event Notifications	Port up/down event notification to log or terminal. Refer to Event Notification chapter in the User Manual.
OSPF	NSSA area type support. Refer to “area nssa” command in the OSPF chapter of the CLI guide.
Gateway	New user interface (CLI and WebUI) for the gateway. Refer to User Manual and CLI guides and webUI help pages.
Gateway	Support for 100 Ethernet hosts and 3500 InfiniBand hosts.
Systems	SX6036G and SX6012 support at GA level.
Systems	Gateway and VPI support on SX6018.
Systems	Gateway and VPI support on SX6012.
Release 3.3.4000	
Systems	SX1012 supported at GA level.
General	Bug fixes.
Release 3.3.3500	
IP Routing	OSPF support (GA level).
Gateway	Gateway support (GA level).
Management Interfaces	In-band management.
Ethernet Switching	Port mirroring.

Table 4 - VPI Changes and New Features

Category	Description
Ethernet Switching	sFlow.
IP Routing	OSPF support (beta level).

4 Known Issues

The following sections describe MLNX-OS® known issues in this software release and possible workarounds.



For hardware issues, please refer to the switch support product page.

4.1 General Known Issues

Table 5 - General Known Issues (Sheet 1 of 6)

Index	Category	Description	Workaround
1.	Gateway	Gateway is not supported in MLNX-OS release 3.4.1110.	N/A
1.	Chassis Management	BBU support only extends to the following CLI commands: <ul style="list-style-type: none"> battery-backup-unit charge battery-backup-unit discharge show battery-backup-unit show battery-backup-unit details show module show power show power consumers show temperature show voltage 	N/A
2.	Management Interfaces	DHCPv4/v6, VLAN, Zeroconf are not supported on IPoIB.	N/A
3.	Management Interfaces	When re-enabling interface ib0, MTU settings are not saved.	Manually configure MTU settings after re-enabling interface ib0.
4.	Management Interface	If using eth-single-swid system profile, static ARP cannot be configured on the mgmt0 interface without an L3 license.	N/A
5.	Management Interfaces	The CLI command <code>ip default-gateway <interface></code> sets the gateway address to 0.0.0.0 and prevents the user from adding other gateways.	Delete the entry by using the <code>no ip default-gateway</code> command.
6.	Management Interfaces	The “no” negation prefix is not supported prefix in the following CLI commands: <ul style="list-style-type: none"> <code>interface ethernet <s/p> ipv4 port access-group no dhcrelay</code> <code>interface port-channel <num> switchport mode</code> 	Use the “no” keyword as an infix (e.g. <code>interface port-channel <num> no switchport</code>).
1.	Management Interfaces	The switch may have an expired HTTPS certification.	Generate a new certificate by changing the hostname.

Table 5 - General Known Issues (Sheet 2 of 6)

Index	Category	Description	Workaround
2.	Management Interfaces	Consecutive hostname modification is not supported.	Wait 25 seconds before reattempting to modify the hostname.
3.	Management Interfaces	The command <code>reset factory keep-basic</code> removes the management IP configuration.	N/A
4.	Puppet Agent	InfiniBand port speed is updated only when the port is active (“up”) and the given speed is supported.	Ensure that the InfiniBand speed is supported before configuring puppet InfiniBand interface speed and setting the admin to “up” (in the puppet manifest).
5.	Puppet Agent	Only the resource ID “default” is supported for Puppet router.	N/A
6.	Puppet Agent	Duplex parameter is not supported for Ethernet, InfiniBand or LAG interfaces.	N/A
7.	In-Band Management	In-band management mode does not support IPv6.	N/A
8.	In-Band Management	In-band management ARPs are not taken into account in the command <code>show ip arp count</code> .	N/A
9.	NTP	The command <code>show ntp</code> always lists the last configured NTP server even if it has been deleted. This output can be safely ignored.	N/A
10.	Software Management	Only one image is allowed to be copied into the system (using the <code>image fetch</code> command). The user must remove old image files prior to fetching a new one.	N/A
11.	Configuration Management	Loading the wrong system configuration file may hang the system. For example: an attempt to use an SX1016 configuration file on an SX6036/SX6536/SX1035 system may cause the system to hang.	N/A
12.	Configuration Management	After loading a new configuration file, please reboot the system. Otherwise, configuration may not be properly applied and errors may appear in the log.	N/A

Table 5 - General Known Issues (Sheet 3 of 6)

Index	Category	Description	Workaround
13.	Configuration Management	Factory settings can be restored only if the system profile is the factory profile (for example, Ethernet profile on SX1036 or InfiniBand profile on SX6036).	To revert to factory defaults, change the system profile to the native factory mode (e.g., change the profile to InfiniBand on SX6036, or to Ethernet on SX1036) using the system profile command. This will delete all settings, except the basic settings (license, SSH keys and mgmt interfaces settings).
14.	Configuration Management	Prior to license deletion make sure to be in an allowed profile. Failing to do so may result in errors.	N/A
15.	Configuration Management	When using a large set of configuration files, configuration apply can take more time than usual due to parallel activity of statistics data collecting.	N/A
16.	Configuration Management	Applying a configuration file of one system profile to another is not supported.	N/A
17.	Configuration Management	If the last box leaving an HA (MLAG/SM HA/ Proxy-ARP HA) cluster is not the box on which the cluster was created, bogus configuration will result which does not allow recreating the same cluster name.	N/A
18.	Configuration Management	On TX6240/TX6280 systems only VL0 may be configured on the long-haul ports.	N/A
19.	Configuration Management	The system needs to be rebooted after running the command <code>configuration revert factory keep-basic keep-connect</code> .	N/A
20.	Configuration Management	When working in vpi-single-switch profile, the number of supported multicast groups is confined to 3072.	N/A
21.	Logging	“DROPPED MSG” errors may appear during reload (shutdown phase). These errors can be safely ignored.	N/A
22.	Logging	The warning “pgm_set_timeout” may appear in the log. This warning can be safely ignored.	N/A

Table 5 - General Known Issues (Sheet 4 of 6)

Index	Category	Description	Workaround
23.	Logging	During system de-init, the error “[mdreq.ERR]: init(), mdr_main.c:634, build 1: Error code 14014” may appear in the log. This error can be safely ignored.	N/A
24.	Logging	The warning “[mgmtd.WARNING]: Upgrade could not find node to delete: /iss/config/stp/switch/ethernet-default/spanning-tree/mode” may appear in the log. This warning can be safely ignored.	N/A
25.	Logging	When using a regular expression containing (OR) with the command <code>show log [not] matching <reg-exp></code> , the expression should be surrounded by quotes (“<expression>”), otherwise it is parsed as filter (PIPE) command.	N/A
26.	Logging	Port up/down events on a port quickly toggling states may be displayed in wrong order in the monitoring terminal.	For actual port stats, use the command <code>show interface</code> .
27.	Logging	The following error may appear in the log: “sn_mdc_msg_handler(), ../../tms/sn_mgmt.c:910, build 1: Error code 14002 (assertion failed) returned”. This error may be safely ignored.	N/A
28.	Logging	The message “[mgmtd.ERR]: md_cert_validate_new_cert_value(), md_cert.c:3388, build 1: Return status 512 from openssl verify!” may appear in the log. This message can be safely ignored.	N/A
29.	User Management	Some RADIUS and TACACS+ configurations keep the user locked out of the machine due to timeout limitation.	Press the reset button for 15 seconds, and then log in using your local authentication. Additionally, fix the configuration to avoid any future timeout issues.
30.	WebUI	Reversing the time clock can result in WebUI graphs’ corrupted data.	Clear the graphs data after setting the clock.
31.	WebUI	WebUI does not support arbitrary setting of all SM routing engines. If the CLI does not support the configuration, the routing engines option is disabled in the WebUI.	N/A
32.	WebUI	Enabling/disabling HTTPS while connected via HTTP to the WebUI may result in temporary loss of connection to the webpage.	Refresh the page or navigate back using the browser’s back button.

Table 5 - General Known Issues (Sheet 5 of 6)

Index	Category	Description	Workaround
33.	WebUI	The following commands cannot be executed from “Execute CLI commands” section on WebUI. <ul style="list-style-type: none"> interface ethernet <interface> module-type dcb priority-flow-control enable 	To run those commands use the CLI.
34.	WebUI	Importing a configuration text file with a port-type change configuration in it is not supported via the WebUI.	Import the configuration text file through the CLI.
35.	WebUI	If the configured ciphers in versions prior to 3.4.0012 were SSL and TLS ciphers, upgrading to this version will override that. The new default is to allow TLS ciphers only. To enable SSL, please run the command <code>web https ssl ciphers all</code> .	N/A
36.	WebUI	When strict mode is activated with TLS 1.2, Firefox does not work properly.	N/A
37.	CLI	MLNX-OS support up to 50 CLI session open in parallel.	N/A
38.	CLI	Command output filtering does not support the following commands: <ul style="list-style-type: none"> show log show puppet-agent log show ib sm log show configuration text files <file> 	N/A
39.	SNMP	The error “Cannot find module (MELLANOX-MIB)” may appear in the log when performing rollback to a MLNX-OS version older than 3.3.3000. This error can be safely ignored.	N/A
40.	SNMP	Upon system shutdown, the following error may appear: “[mibd.ERR]: mdc_foreach_binding_prequeried_parsed(), mdc_main.c”. This error can be safely ignored.	N/A
41.	SNMP	The ifNumbers MIB (OID: 1.3.6.1.2.1.2.1.0) on x86 switch systems displays 42 interfaces while the ifTable displays 40 due to VM management interfaces that are not shown in the ifTable.	N/A
42.	Chassis Management	Supported USB flash drives are USBv2 protocol and above.	N/A
43.	Chassis Management	Duplicate notification is sent upon high temperature alerts, both specific module alerts and general “too high” alerts.	N/A

Table 5 - General Known Issues (Sheet 6 of 6)

Index	Category	Description	Workaround
44.	Chassis Management	When plugging in the power supply without connecting it to the power source, the system status and power supply LEDs become RED.	Plug in the power source.
45.	Chassis Management	Fans in a fatal state may appear as “not present” using the command <code>show fans</code> .	Use the command <code>show health report</code> for actual fan state.
46.	Chassis Management	Director switch systems boot time is slightly slower than previous releases.	N/A
47.	Chassis HA	Setting an IPv6 address for BIP and VIP is currently not supported.	N/A
48.	Unbreakable Links	LLR is not supported on MetroX® platforms’ long haul ports.	N/A
49.	Virtual Machine	For volume fetch, using a USB drive formatted with VFAT causes errors in the log and may require additional reboot for the USB to be registered for virtual machine volume usage.	Use EXT3 USB format.
50.	Virtual Machine	For virtual machine volume fetch, only ext3 USB format is supported. Using vfat formatted usb disk on key will cause errors in log and may require additional reboot for USB to be registered for virtual machine volume usage	N/A
51.	Systems	On SX6720, when running on battery, the system’s power may serve for up to 1 minute more than the reported time shown in the command “show battery-backup-unit”.	N/A
52.	Systems	On SX6720, 30 seconds before the battery runs out, the command “show battery-backup-unit details” may not display accurate data.	N/A
53.	Puppet Agent	On MSB7700, puppet agent is not supported in this release.	N/A

4.2 Ethernet Known Issues

Table 6 - Ethernet Known Issues (Sheet 1 of 7)

Index	Category	Description	Possible Workaround
1.	Ethernet Interfaces	The system allows tolerance of 4 bytes on top of the set MTU.	N/A
2.	Ethernet Interfaces	In case of faulty cable, rise time of other interfaces in the system may be delayed.	N/A

Table 6 - Ethernet Known Issues (Sheet 2 of 7)

Index	Category	Description	Possible Workaround
3.	Ethernet Interfaces	40GbE passive copper cable of 5m length might experience link rise issues when connected to some 3rd party 40GbE interfaces.	Replace copper cable with fiber cable
4.	Ethernet Interfaces	Default MTU is modified as 1500 instead of 1522.	N/A
5.	Ethernet Interfaces	There are no port counters for packet sizes 1518-1522.	N/A
6.	Ethernet Interfaces	Traffic loss may occur when enabling flow control on a packet with an MTU larger than 8K on SX1012.	N/A
7.	Ethernet Interfaces	LAG cost is not calculated correctly.	N/A
8.	User Interfaces	SX1012 hybrid cable information may not appear on WebUI/CLI.	N/A
9.	LAG	Convergence time of port removal from a LAG exceeds the specification by up to 30 seconds.	Shut down the port prior of removing it from the LAG.
10.	LAG	The remote LACP admin-key #1 cannot be established with some third party vendors.	Use a different admin-key number.
11.	LAG	Traffic running over LAG may not be evenly distributed when testing small number of SMACs.	N/A
12.	LAG	When converting static LAG to dynamic LAG (LACP), some of the ports remain in DOWN state.	Disable and enable the LAG.
13.	LAG	LAG Hash function lacks configuration of a randomization seed parameter. Traffic passed over LAG in hierarchy topology via multiple SwitchX based systems is not distributed evenly.	Change the distribution function between layers performing LAG using the command <code>port-channel load-balance</code> . The work-around is applicable in case there are multiple mixed flows.
14.	LAG	When removing all ports of a port-channel and then re-adding them to the port-channel in a different type, port-channel may not converge.	Disable and enable the port-channel.
15.	LAG	The following LACP port-states are not available in MLNX-OS®: <ul style="list-style-type: none"> • I – Individual H – hot-standby (LACP only) • s – Suspended R – module-removed 	N/A
16.	LAG	Configuring multiple LAGs as a range causes the switch to hang for up to 10 minutes.	N/A

Table 6 - Ethernet Known Issues (Sheet 3 of 7)

Index	Category	Description	Possible Workaround
17.	MLAG	Each MLAG VIP group must be configured with a different unicast IP address. If not, MLAG behavior is not anticipated.	N/A
18.	MLAG	MLAG is not supported in WebUI.	N/A
19.	MLAG	MLAG port-channels (Mpo) appear as regular port-channels (Po) in WebUI.	N/A
20.	MLAG	MLAG configuration including (mlag-port-channel configuration) must be identical in both switches.	N/A
21.	MLAG	MLAG VIP is limited to 2 switches. If more than 2 switches connected to the same MLAG VIP, MLAG behavior is not anticipated.	N/A
22.	MLAG	When the MPO is connected to the host, spanning tree should be disabled on the MLAG port-channel before it is enabled (port type set to edge and BPDU filter enabled). When connected to a switch, spanning tree should be disabled on the switch (<code>no spanning-tree</code>).	N/A
23.	MLAG	Deleting an IPL port-channel interface while MLAG is enabled (using the command <code>no interface port-channel <id></code>) might cause errors. These errors may be safely ignored.	Shutdown the port before deleting it.
24.	MLAG	The IPL VLAN interface must be used only for MLAG protocol and must not be used by any other interfaces (e.g. port-channel, Ethernet).	N/A
25.	MLAG	An MLAG-port-channel must be administratively disabled (<code>interface mpo <id> shutdown</code>) before deleting it.	N/A
26.	MLAG	IGMP snooping is automatically enabled on MLAGs and may not be disabled.	N/A
27.	MLAG	Keepalive interval should be set to 3 seconds when working with over 1k VLANs on PPC.	N/A
28.	MLAG	Configuring a LAG to become an IPL when another LAG is already configured as IPL leads to MLAG going down and an error being printed in the log.	Have only one LAG set as IPL. Remove any previous configuration if a new one is to be applied.
29.	MLAG	Changing STP mode affects traffic and may cause MLAG to stop working for an interval of 21 seconds until the new STP topology is built.	User should disable MLAG prior to changing the STP mode.

Table 6 - Ethernet Known Issues (Sheet 4 of 7)

Index	Category	Description	Possible Workaround
30.	VLAN	It might take approximately 20 seconds to create/delete 1000 VLANs, or to change the switchport mode of an interface.	N/A
31.	VLAN	Deleting VLAN ranges on trunk/hybrid ports is correlated to the number of the trunk/hybrid ports in the system.	Restrict the range of VLANs deleted to groups of up to a 100 VLANs.
32.	VLAN	When working with more than 2000 VLANs and trying to change a port's VLAN membership in modes hybrid/trunk using the command <code>interface 1/x switchport trunk allowed-vlan</code> , the operation may timeout on account of it being too heavy. Thus, the port does not become properly configured as a member of those VLANs. Assuming port 1 is in trunk mode and there are 3000 VLANs created, a good practice would be to achieve membership of port 1 in VLANs 100-3000, use the command <code>interface ethernet 1/x switchport trunk allowed-vlan remove 1-100</code> instead of <code>interface ethernet 1/x switchport trunk allowed-vlan 100-3000</code> .	To fix the problem, run the commands <code>interface ethernet 1/x switchport trunk allowed-vlan all</code> and <code>interface ethernet 1/x switchport trunk allowed-vlan remove 1-100</code> in succession.
33.	Spanning Tree	Non-standard behavior may occur on third party switch STP functionality when the port moves from Root to Alternate state. The standard Proposal-Agreement sent by MLNX-OS is unanswered which may result in traffic loss of up to 30 seconds.	N/A
34.	Spanning Tree	STP convergence time is slower than expected.	N/A
35.	Spanning Tree	BPDU's are sent from both the superior and the inferior switches.	N/A
36.	MSTP	The command <code>no vlan</code> on a VLAN interface associated with an MSTP instance is not supported.	Dissociate the VLAN from the MSTP instance before running the command.
37.	MSTP	Under heavy MSTP configuration, the system may exhibit slowness after adding a new MSTP instance. In the ensuing minute, if a show command is issued an error message may appear. That message can be safely ignored.	N/A
38.	IGMP Snooping	Setting the <code>ip igmp snooping unregistered multicast option forward-to-mrouter-ports</code> option reduces the amount of supported VLANs to 250.	N/A

Table 6 - Ethernet Known Issues (Sheet 5 of 7)

Index	Category	Description	Possible Workaround
39.	IGMP Snooping	There is no possibility to have more than one mrouter port on a single VLAN if IGMP snooping is disabled for that VLAN.	After enabling IGMP snooping on that VLAN, the user needs to re-add the mrouter-ports on that VLAN.
40.	IGMP Snooping	When no IP interfaces exist on the switch, the default IGMP querier address is 10.0.0.1.	Create an interface VLAN with the required IP address and reconfigure the IGMP querier on the required VLAN.
41.	IGMP Snooping	IGMP snooping is limited to a single MC IP per MC MAC.	N/A
42.	IGMP Snooping	IGMP snooping querier is not supported over more than one VLAN when IP routing is enabled.	Use PIM IGMP querier.
43.	LLDP	LLDP notifications are not supported.	N/A
44.	LLDP	The mgmt0, mgmt1 IPv6 addresses and the management address are not advertised on LLDP TLV(s).	N/A
45.	LLDP	When configuring LLDP with PFC on a LAG, after switch reload, the LLDP does not show the PFC configuration anymore.	After switch reload, reconfigure the PFC LLDP on the LAG using the command <code>interface ethernet <x/y/z> lldp tlv-select dcbx.</code>
46.	LLDP	LLDP ifIndexes do not match the ones published in the ifTable in IEEE 802.1.	Use the ifDescr and fields.
47.	ACLs	Packets dropped by the switch due to congestion or ACL rules are added to the bad type counter of the matching VLAN interface.	N/A
48.	Modules	When using 5m splitter cable (P/N MC2609125-005), wrong data is displayed under transceiver information.	N/A
49.	Modules	Using 1GbE copper modules leads to the port status being constantly reported as active regardless from the actual link state.	N/A
50.	Modules	The command <code>show interface</code> does not display 1GbE modules correctly.	N/A
51.	SNMP	LLDP MIB OID <code>lldpRemPortDesc</code> is not supported.	N/A

Table 6 - Ethernet Known Issues (Sheet 6 of 7)

Index	Category	Description	Possible Workaround
52.	In-Band Management	In-band management may not be enabled with IP Routing or Proxy-ARP.	N/A
53.	Puppet Agent	Speed parameter is not supported for LAG interfaces.	N/A
54.	Puppet Agent	Description parameter is not supported on L2 interfaces.	
55.	Port Mirroring	When port mirroring congestion is configured to “drop-excessive-frames” (best effort), while using the same analyzer port for more than a single mirroring session, that port could drop packets from all flows not according to their configured priorities.	N/A
56.	sFlow	The discarded packets counter in sFlow samples may not be accurate and may expose a skew of one second.	N/A
57.	sFlow	Discard counters are not reported to the collector.	
58.	OpenFlow	OpenFlow ARP packet-matching rule does not include IP source and destination addresses.	N/A
59.	OpenFlow	OpenFlow is not supported over LAG, member ports of LAG, or split-port interfaces.	N/A
60.	OpenFlow	OpenFlow can support up to 1000 flows, supporting 12 tuples with wildcards according to the standard.	N/A
61.	OpenFlow	Standard ACL cannot be configured while OpenFlow is in use.	N/A
62.	OpenFlow	OpenFlow counters can support counting of packets.	N/A
63.	OpenFlow	OpenFlow packet modify action supports VLAN (setting VLAN ID, setting VLAN priority and stripping VLAN header – for QinQ packets).	N/A
64.	OpenFlow	Output can only be set to flood, physical port normal, controller and drop.	N/A
65.	OpenFlow	OpenFlow EtherType-matching rule can be set to one of: IPv4, IPv6, ARP/RARP, FCoE and 2 user defined Ethertypes.	N/A
66.	OpenFlow	OpenFlow ICMP packet-matching rule does not include type and code.	N/A

Table 6 - Ethernet Known Issues (Sheet 7 of 7)

Index	Category	Description	Possible Workaround
67.	DCBX	Enabling LLDP triggers a faulty notation of PFC oper-state. This status should be ignored.	N/A
68.	DCBX	There is no automatic fallback between IEEE and CEE.	N/A
69.	DCBX	When moving from DCBX CEE to DCBX IEEE, TLVs may be sent inappropriately.	After mode change, the TLVs to be sent must be reset.
70.	DCBX	DCBX auto-select type (IEEE/CEE) is not supported.	DCBX type should be selected manually.
71.	Chassis Management	Problems with the power supply cannot be monitored on SX1016 switch systems using the command <code>show health-report</code> or by observing the PS LED on the switch system itself.	N/A

4.3 IP Routing Known Issues

Table 7 - IP Routing Known Issues (Sheet 1 of 5)

Index	Category	Description	Possible Workaround
1.	IP Interfaces	ICMP echo packets are not counted as part of the VLAN interface counters.	ICMP packets are counted in the interface L2 counters.
2.	IP Interfaces	A VLAN bound with an ACL action must not be bound to a router port. Else, router port behavior is unpredictable.	N/A
3.	IP Interfaces	Router ports are not supported on SX10xx-xxxR and SX60xx-xxxR systems.	N/A
4.	IP Interfaces	It is not possible to add or remove member interfaces from a LAG, or delete that LAG when it is configured as a router port.	Physical port must explicitly be changed from “no switchport” to “switchport” so that router capabilities are reset from physical port.
5.	IP Interfaces	ICMP redirect is not supported over router port.	N/A
6.	LLDP	A port connected to a router-port receives from it an LLDP packet containing IPv4 and IPv6 management addresses. However, the IPv6 address is not included in the LLDP MIB as a remote management address. Only the IPv4 remote management address is published.	N/A

Table 7 - IP Routing Known Issues (Sheet 2 of 5)

Index	Category	Description	Possible Workaround
7.	IP Diagnostic Tools	While using invoked shell commands such as ping, traceroute, tcpdump, for egressing from a specific router interface use <code>-I/i vlan<vlan id></code> . For example, to ping 10.10.10.10 via interface VLAN 30 run the command: <code>ping 10.10.10.10 -I vlan30</code> .	N/A
8.	IP Routing	Prior to upgrading to this software release the user must align the number of configured multicast routes to the number of routes defined in section “IPv4 Routing Mode” in the User Manual.	N/A
9.	IP Routing	LAG and L3 ECMP Hash functions lack the configuration of a randomization seed parameter. Traffic passed over LAG or L3 ECMP in hierarchy topology via multiple SwitchX based systems is not distributed evenly.	Change the distribution function between layers performing LAG or L3 ECMP using the command <code>port-channel load-balance / ip load-sharing</code> . The workaround is applicable in case there are multiple mixed flows.
10.	IP Routing	In IP Routing mode, multicast or broadcast packets such as ARP may be counted as “RX discard” in case only one interface is a member of the VLAN. These kind of packets do reach the CPU for processing. The progressing counter can be ignored.	N/A
11.	IP Routing	Errors in the log may appear when using ECMP routes if next-hop belongs to a VLAN interface which lies on port-channel whose links operational state is toggling.	N/A
12.	IP Routing	If there is a static ARP on top of an interface VLAN and ECMP route with its next-hop equal to the static ARP entry, and if the static ARP owner is not responding to ARP requests, then performing <code>shutdown</code> and <code>no shutdown</code> to the interface VLAN, within a minute, causes errors to appear on the log. These errors may be safely ignored.	N/A
13.	IP Routing	ARP aging timeout minimum value is 240. After upgrade to release 3.3.3500, any value lower than this is clamped to 240.	N/A
14.	IP Routing	The configuration of <code>ip arp timeout</code> does not affect entries related to gateway, out-of-band, or management interfaces.	N/A

Table 7 - IP Routing Known Issues (Sheet 3 of 5)

Index	Category	Description	Possible Workaround
15.	IP Routing	If the number ARP entries configured on a system is over the permitted limit, the machine becomes overloaded and hangs.	Power cycle the system.
16.	IP Routing	If static route next-hop of a VLAN interface is configured before configuring its subnet, that configured route is not added to the interface and cannot be deleted later.	Make sure to configure the VLAN interface subnet before configuring the next-hop itself.
17.	OSPF	Changing the IP address on an OSPF Interface VLAN deletes the interface-OSPF attachment.	Reconfigure all OSPF settings after changing the IP.
18.	OSPF	Interface configuration cannot be performed while OSPF is admin-disabled (shutdown).	Enable OSPF admin state using the command <code>router ospf no shutdown</code> .
19.	OSPF	Removing a static ID using the command <code>no router-id</code> requires at list one VLAN interface to be configured.	N/A
20.	OSPF	OSPF supports learning of up to 4K routes.	N/A
21.	OSPF	Default route advertise by ABR to NSSA area not installed by NSSA area router.	Install the default route manually on the NSSA area using the command <code>ip route</code> .
22.	OSPF	Changing administrative distance (the command <code>router ospf distance</code>) is not supported.	N/A
23.	OSPF	Reverse packet filtering is not supported. When misconfigured, traffic may pass on an unintended subnets.	N/A
24.	OSPF	If the MTU is reconfigured while the interface is up, disruptions may be experienced in L3 traffic.	Shut down the interface before reconfiguring MTU.
25.	OSPF	A loopback interface assigned to area 0 does not get redistributed as direct route after removal from the OSPF.	Keep the loopback assigned to the OSPF area 0
26.	DHCP Relay	DHCP Relay is supported on SwitchX-2 platforms only.	N/A
27.	DHCP Relay	DHCP Relay is not supported on SX10xx-xxxR and SX60xx-xxxR systems.	N/A
28.	VRRP	Using the same Virtual IP (VIP) as the VLAN interface is not supported.	N/A
29.	MAGP	Designated traffic, such as ping, to the MAGP interface is not supported.	N/A

Table 7 - IP Routing Known Issues (Sheet 4 of 5)

Index	Category	Description	Possible Workaround
30.	PIM	PIM BIDIR routers must have point-to-point connection.	N/A
31.	PIM	The command <code>show ip pim rp</code> does not distinguish elected RP from other active RPs.	Use the command <code>show ip pim rp-hash</code> to see the elected RP, and <code>show ip mroute</code> to see active RPF-interface.
32.	PIM	Per-interface multicast TTL threshold, and static OIF cannot be configured without enabling PIM in the interface.	Enable PIM-SM on the interface.
33.	PIM	Local setting of candidate-BSR interval affects the BSR hold-time of other winner BSR router.	Use the default BSR interval on the BSR router (60 seconds), or configure all routers with the same interval. (Note that this does not affect per-RP hold-time within BSR).
34.	PIM	At most 300 IGMP packets can be processed per second.	Do not use PIM in environments where there are more than 300 joins/leaves per second.
35.	PIM	RP candidate can be enabled only on one interface at a time.	N/A
36.	PIM	BSR candidate can be enabled only on one interface at a time.	N/A
37.	PIM	In an L3 network with two or more PIM-BIDIR routers, IGMP snooping should be disabled.	N/A
38.	Multicast	Multicast protocol (PIM) must be enabled in order to configure static multicast route.	Enable PIM.
39.	Multicast	Router port does not filter outgoing multicast packets according to configured multicast time-to-live (TTL) value.	N/A
40.	BGP	Route aggregation configuration cannot be modified.	Delete route aggregation and create a new configuration.
41.	BGP	Aggregate overlapping not supported	N/A
42.	BGP	Using the command <code>clear ip bgp <ip> soft</code> in mandates enabling route-fresh on the peer.	If route-refresh is not supported on the peer device, use the command <code>clear ip bgp <ip> in</code> instead.

Table 7 - IP Routing Known Issues (Sheet 5 of 5)

Index	Category	Description	Possible Workaround
43.	BGP	The command <code>bgp aggregate-address summary-only</code> does not silence aggregate-route re-advertisements.	N/A
44.	BGP	AS-path request/set are limited to 128. When more are received, TCP FIN is sent and connection to the peer is lost.	N/A
45.	BGP	When running both BGP and OSPF, changing the configuration of the command <code>router bgp <AS> <external distance> <internal distance> <local distance></code> requires running <code>clear ip bgp all</code> afterwards.	N/A
46.	BGP	Four-byte ASN is not supported.	N/A
47.	BGP	The command <code>no neighbor <ip-address> local-as</code> is not functional.	Delete the neighbor with <code>no neighbor {<ip-address>} remote-as</code> and restore the neighbor without the command <code>neighbor local-as</code> .
48.	ACL	When upgrading to this release, it is advised to reduce the number of ACL configured to 3.	N/A

4.4 InfiniBand Known Issues

Table 8 - InfiniBand Known Issues

Index	Category	Description	Possible Workaround
1.	InfiniBand Interfaces	Port hardware speed and width capabilities settings affect port speed and width admin capabilities.	N/A
2.	InfiniBand Interfaces	Setting the port width to 1x in the WebUI and/or CLI is currently not supported.	N/A
3.	InfiniBand Interfaces	Clearing the port counters when no SM is running in the subnet, may result in error in the log, it can be safely ignored.	N/A
4.	SM	DOR and LASH OpenSM routing algorithms can be configured only via the CLI.	N/A
5.	SM	The SM must be disabled prior to and during the software upgrade.	N/A
6.	SM	Chassis HA takeover might terminate open SM on new master.	Run open SM manually after takeover.

Table 8 - InfiniBand Known Issues

Index	Category	Description	Possible Workaround
7.	SM	The SM and InfiniBand diagnostics cannot run in case the system profile is vpi-single-switch.	N/A
8.	SM	The “Subnet Manager started” email may not be sent on system initialization.	N/A
9.	SM	Two notifications are sent via SNMP trap or email when the SM is disabled.	N/A
10.	WebUI	SX65xx systems WebUI: Recent manual changes made to the port configuration of the InfiniBand port are applied only after pressing the Apply button. Refreshing the page without saving the changes, cause the loss of the changes that were not saved, and the Port 1 of Leaf 1 (if present) is shown.	N/A
11.	WebUI	On SX65xx systems, removed PS appears as powered off through the WebUI.	N/A
12.	WebUI	On SX65xx systems, management box IP (BIP) cannot be removed from WebUI.	Remove BIP from CLI.
13.	SNMP	The SNMP timeout in an external SNMP manager should be set to 60 seconds when accessing SX65xx.	N/A
14.	SNMP	ifPhysAddress OID returns the prefix of the Node GUID of the ib0 management.	N/A
15.	Chassis Management	On SX65xx systems, hot-plugging Spine 1 or Spine 2, which are connected to the management modules, is not supported. Hotplugging them may cause the system to freeze.	N/A
16.	Chassis Management	Unresponsive module notifications on SX65xx are sent for each module that is powered off along with an insufficient power notification.	N/A
17.	Chassis Management	When taking out a slave management module on dual management system, the following error may appear in the log: “[health.WARN-ING]: get_voltage_limits : 2531 - got null while performing sweep scan”. This error can be safely ignored.	N/A
18.	Chassis Management	Upgrading an SX65xx system can take up to 40 minutes during which time the system is indisposed.	N/A
19.	Chassis Management	The status of the power supply fan is not available for SX65xx family via CLI/GUI/SNMP.	N/A

Table 8 - InfiniBand Known Issues

Index	Category	Description	Possible Workaround
20.	Chassis Management	Errors appear in the log after plug-out of a fatal management. These errors may be safely ignored	N/A
21.	Chassis Management	Plugging in spines or leafs during system initialization may result in the modules not activating.	Turn the modules off and on again.
22.	Chassis Management	Internal link notification requires a 2-second delay between link states changes on SX65xx systems.	N/A
23.	Chassis Management	Retrieving chassis information (such as temperature, power etc.) by running show commands can only be performed from chassis master.	N/A
24.	Configuration Management	In InfiniBand profile, the system needs to be rebooted after running the command <code>configuration revert factory keep-basic keep-connect</code> .	N/A
25.	Chassis HA	A management module that has been reset 5 times by the other management module on a SX65xx system, is treated as malfunctioning and gets disabled even after plugging out and plugging in the module.	To resolve the issue, perform the following: <ul style="list-style-type: none"> • Remove the powered off management module. • Run <code>chassis ha power enable other</code>. • Remove the powered off management local spine and wait 1 minute • Plug in the removed spine (this step will raise the spines internal links). • Plug in the removed/new management (spine's link will fall and raise again once it is done).
26.	Chassis HA	On dual management systems using the command <code>no chassis ha power enable other</code> may result in other management to be appear in "fatal" state although it is powered off. This state can be safely ignored.	N/A
27.	Chassis HA	The following error may appear in the log in dual management systems: "[clusterd.ERR]: Heartbeat: failed to send heartbeat message to node 1". This error can be safely ignored.	N/A

Table 8 - InfiniBand Known Issues

Index	Category	Description	Possible Workaround
28.	Chassis HA	When taking out a master management on a dual management system, its locally attached spine may appear in the CLI as being off after a take-over.	Turn the spine off and then on using the CLI or WebUI.
29.	SM HA	The SM HA works only over mgmt0.	N/A
30.	SM HA	The SM master may not be determined according to the GUID.	N/A
31.	SM HA	On a network with SM HA enabled, the following errors may appear on SM HA master when members are leaving the subnet: [clusterd.NOTICE]: Heartbeat: potential failure of node [clusterd.NOTICE]: Heartbeat: suspect node XXX of failure These errors can be safely ignored.	N/A
32.	SM HA	Change hostname to members of SM HA network may result in these member to be appear as offline.	Configure hostname to a member prior adding it to the cluster.
33.	SM HA	After upgrade SX65xx system with both management modules that are part of an SM HA cluster, the SM HA slave may see the SM HA master as offline.	Reboot the SM HA slave.
34.	SM HA	SM HA between systems with different CPU types (i.e. PPC and x86) is not supported.	N/A
35.	SM HA	The following error may appear in the log “smm_ha_ready(), smm_main.c:2445, build 1: Error code 14004 (item not found) returned”. This error may be safely ignored.	N/A
36.	SM HA	If an SM HA node fails, the VIP might not be reachable, but functionality is not affected.	Reconnect via the VIP.
37.	Fabric Inspector	Wrong Port 0 physical link state reported in Fabric Inspector.	N/A

Table 8 - InfiniBand Known Issues

Index	Category	Description	Possible Workaround
38.	Logging	<p>Upon system init the following errors may appear in the log:</p> <ul style="list-style-type: none"> • [ibd.ERR]: sa_is_master_sm_in_swid(), swidapi.c:1563, build 1: Sysapi function called but madrpc was not initiated • [ibd.WARNING]: get_is_sm_in_fabric: sa_is_master_sm_in_fabric() err=1(ERROR) • [ibd.ERR]: ibd_mon_handle_get(), ibd_main.c:1347, build 1: Error code 14000 (generic error) returned • [ibd.ERR]: mdc_dispatch_mon_request_nosend(), mdc_misc.c:1535, build 1: Get handler for '/ib/state/swid/1/is_sm_in_fabric' <p>These errors may be safely ignored.</p>	N/A
39.	Logging	<p>The following error may appear in the log when inserting a new leaf/spine “ibd_mon_handle_get(), ibd_main.c:1333, build 1: Required condition was not met”. This error can be safely ignored.</p>	N/A
40.	Logging	<p>The error “[mibd.ERR]: remove_device_from_db(), mibd_entity.c” may appear in the log when plugging out an FRU module. This error can be safely ignored.</p>	N/A
41.	Logging	<p>Several “[cablemond.ERR]” errors may appear in the log when plugging out a leaf or spine module. These errors can be safely ignored.</p>	N/A
42.	Software Management	<p>The following error may appear in the log: “[clusterd.ERR] Failed to receive CCL response for CCL message handle.” This error may be safely ignored.</p>	N/A

5 Bug Fixes

The following table describes MLNX-OS® bug fixes in this software release.

Table 9 - Bug Fixes

Index	Category	Description
1.	Security	Memory leak in the <code>tls_decrypt_ticket</code> function in <code>t1_lib.c</code> in OpenSSL before 0.9.8zc, 1.0.0 before 1.0.0o, and 1.0.1 before 1.0.1j (CVE-2014-3567).
2.	Security	Memory leak in <code>d1_srtp.c</code> in the DTLS SRTP extension in OpenSSL 1.0.1 before 1.0.1j (CVE-2014-3513).
3.	Security	The SSL protocol 3.0, as used in OpenSSL through 1.0.1i and other products, uses nondeterministic CBC padding (CVE-2014-3566).
4.	Security	Double free vulnerability in <code>d1_both.c</code> in the DTLS implementation in OpenSSL 0.9.8 before 0.9.8zb, 1.0.0 before 1.0.0n, and 1.0.1 before 1.0.1i (CVE-2014-3505).
5.	Security	<code>d1_both.c</code> in the DTLS implementation in OpenSSL 0.9.8 before 0.9.8zb, 1.0.0 before 1.0.0n, and 1.0.1 before 1.0.1i (CVE-2014-3506).
6.	Security	Memory leak in <code>d1_both.c</code> in the DTLS implementation in OpenSSL 0.9.8 before 0.9.8zb, 1.0.0 before 1.0.0n, and 1.0.1 before 1.0.1i (CVE-2014-3507).
7.	Security	The <code>OBJ_obj2txt</code> function in <code>crypto/objects/obj_dat.c</code> in OpenSSL 0.9.8 before 0.9.8zb, 1.0.0 before 1.0.0n, and 1.0.1 before 1.0.1i, when pretty printing is used, does not ensure the presence of '\0' characters (CVE-2014-3508).
8.	Security	Race condition in the <code>ssl_parse_serverhello_tlsext</code> function in <code>t1_lib.c</code> in OpenSSL 1.0.0 before 1.0.0n and 1.0.1 before 1.0.1i, when multithreading and session resumption are used (CVE-2014-3509).
9.	Security	The <code>ssl3_send_client_key_exchange</code> function in <code>s3_clnt.c</code> in OpenSSL 0.9.8 before 0.9.8zb, 1.0.0 before 1.0.0n, and 1.0.1 before 1.0.1i (CVE-2014-3510).
10.	Security	The <code>ssl23_get_client_hello</code> function in <code>s23_srvt.c</code> in OpenSSL 1.0.1 before 1.0.1i (CVE-2014-3511).
11.	Security	OpenSSL before 0.9.8za, 1.0.0 before 1.0.0m, and 1.0.1 before 1.0.1h does not properly restrict processing of ChangeCipherSpec messages (CVE-2014-0224).
12.	Security	The <code>config_auth</code> function in <code>ntpd</code> in NTP before 4.2.7p11, when an auth key is not configured, improperly generates a key (CVE-2014-9293).
13.	Security	<code>util/ntp-keygen.c</code> in <code>ntp-keygen</code> in NTP before 4.2.7p230 uses a weak RNG seed (CVE-2014-9294).
14.	Security	Multiple stack-based buffer overflows in <code>ntpd</code> in NTP before 4.2.8 (CVE-2014-9295).

Table 9 - Bug Fixes

Index	Category	Description
15.	Security	The receive function in ntp_proto.c in ntpd in NTP before 4.2.8 continues to execute after detecting a certain authentication error (CVE-2014-9296).
16.	Security	Hidden directory found in the WebUI.
17.	Security	Cacheable SSL page found.
18.	SNMP	Updated SNMP fan and leaf health report trap verbiage.
19.	SNMP	Cannot map remote IP to remote host in lldpRemManAddrTable.
20.	SNMP	SNMP refresh configuration does not support value lower than 60 seconds.
21.	DCBx	Command “show dcb ets interface eth 1/x” output error.
22.	MLAG	The command “no ipl <ipl-id>” command fails.
23.	WebUI	The power graph does not show any data in the WebUI for and SX6710 switch systems.
24.	BGP	BGP import-localpref does not operate properly.
25.	BGP	BGP routes sent to neighbors do not get revoked if policy is placed on the direction “out”.
26.	BGP	Router-ID is not properly used when selecting the best route
27.	BGP	BGP sends withdraw traffic for routes for which it has replacement.
28.	BGP	The command show ip bgp summary always shows zero multi-path network entries and zero multi-path paths.
29.	BGP	BGP connected routes are not chosen over remote routes.
30.	BGP	Removing a neighbor from peer group is not added to running-config.
31.	BGP	Neighbors are unable to establish a session using local-as feature.
32.	BGP	Community received from one neighbor is sent to another neighbor even though the send community feature is not enabled.
33.	DHCP Relay	IP DHCP relay does not flood L2 DHCP.
34.	Spanning Tree	Loop guard cannot be enabled on a shared link although the topology is of single switch.
35.	Spanning Tree	Loop guard cannot be enabled.
36.	MSTP	The indicators for “Inconsistent” are shown only in show spanning-tree.
37.	MSTP	Upon TCN MAC is flushed for all instances and all VLANs instead of the specific instance.
38.	IP Routing	MAC is not updated when GARP is sent to the switch.
39.	IP Routing	ARPs of unreachable hosts are stuck in the ARP table and will be deleted only by interface shutdown.

Table 9 - Bug Fixes

Index	Category	Description
40.	IP Routing	Platform does not shift traffic after receiving gratuitous ARP.
41.	VRRP	Disabling preemption does not take effect.
42.	VRRP	VRRP node uses the wrong source MAC.
43.	VRRP	Master VR does not take over properly after reload with over 1000 instances.
44.	VRRP & MAGP	After reboot with multiple MAGP/VRRP instances, some might get stuck in init if the digits of the MAGP/VRRP ID of some instances are contained in another (e.g “12” and “122”).
45.	VRRP & MAGP	When changing configuration of MAGP/VRRP, the traffic is dropped on other MAGP/VRRP instances.
46.	Logging	When inserting a PS into SX6710/SX1710 switch systems an error log may be printed. These errors may be safely ignored.
47.	Chassis Management	In SX6710 and SX1710 switch systems, fan speeds lower than 4000rpm are detected as 4000rpm by the system and hence do not trigger any alarms.
48.	Chassis Management	Inconsistent power supply part numbers in SX65xx systems when running the command <code>show inventory</code> .
49.	Software Management	MLNX-OS management becomes unavailable after accessing XML tree at top level.
50.	Software Management	On SX1016 systems, management interface <code>mgmt1</code> appears in the CLI even though it does not exist on the switch box.

6 Submitting a Service Request

The Mellanox® Support Center is at your service for any issues. You may access the Warranty Service through the Web Request Form by using the following link:

http://www.mellanox.com/content/pages.php?pg=support_index.