



Mellanox MLNX-OS® Release Notes for IBM SX90Y3474

Software Ver. 3.3.3706

www.mellanox.com

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.
Beit Mellanox
PO Box 586 Yokneam 20692
Israel
www.mellanox.com
Tel: +972 (0)74 723 7200
Fax: +972 (0)4 959 3245

© Copyright 2013. Mellanox Technologies. All Rights Reserved.

Mellanox®, Mellanox logo, BridgeX®, ConnectX®, CORE-Direct®, InfiniBridge®, InfiniHost®, InfiniScale®, MLNX-OS®, PhyX®, SwitchX®, UFM®, Virtual Protocol Interconnect® and Voltaire® are registered trademarks of Mellanox Technologies, Ltd.

Connect-IB™, FabricIT™, Mellanox Open Ethernet™, Mellanox Virtual Modular Switch™, MetroX™, MetroDX™, ScalableHPC™, 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	4
2.3 Supported Firmware	4
2.4 Supported Mezzanine	4
2.5 Supported CPLD	4
2.6 Supported System Director Plug-In	5
2.7 Supported Software Licenses	5
2.8 Supported Cables	5
Chapter 3 Changes and New Features	6
Chapter 4 Known Issues	8
4.1 General Known Issues	8
4.2 Ethernet Known Issues	10
Chapter 5 Bug Fixes	13
Chapter 6 Submitting a Service Request	13

1 Introduction

This document is the Mellanox MLNX-OS® Release Notes for IBM 90Y3474/6.

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 3 documents:

- A command reference guide – provides general information concerning the MLNX-OS command line interface
- A user manual – provides general information concerning the scope and organization 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
SX90Y3474	Sierra, 32-Port 40Gb/s Ethernet Blade Switch System

2.2 Supported CPU Architecture

- PPC 460 CPU families

2.3 Supported Firmware

- SwitchX® firmware version 9.2.3002
- SwitchX-2 firmware version 9.2.3002
- ConnectX®-2 firmware version 2.9.1314 and higher
- ConnectX-3 firmware version 2.11.1250 and higher

2.4 Supported Mezzanine

- ConnectX-2, Mezzanine P/N 90Y3460 (MalayaP), 2.9.1314 and higher
- ConnectX-2, Mezzanine P/N 90Y3480 (MalayaP-Net), 2.9.1314 and higher
- ConnectX-3, Mezzanine P/N 90Y3488 (Merlin), 2.30.2010 and higher
- ConnectX-3, Mezzanine P/N 90Y3484 (Nevada), 2.11.1450 and higher
- ConnectX-3, Mezzanine P/N 90Y3456 (MalayaX), 2.11.1450 and higher
- ConnectX-3, Mezzanine P/N 90Y3468 (MalayaX-Net), 2.11.1450 and higher

2.5 Supported CPLD

- 1.0.18

2.6 Supported System Director Plug-In

- Plug-in software version 003d

2.7 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*.

2.8 Supported Cables

Please refer to the *Mellanox Products Approved Cable Lists* document for the list of supported cables.

http://www.mellanox.com/related-docs/user_manuals/Mellanox_approved_cables.pdf



When using AOC cables longer than 100m use one VL to achieve full wire speed.

3 Changes and New Features

Table 2 - IBM 90Y3476 Changes and New Features

Category	Description
Release 3.3.3706	
General	Bug fixes.
Release 3.3.3704	
VPD	Changes to VPD block 1 capability bits.
Release 3.3.3702	
EHCM	Added detailed reasons for failure of CMM upgrade feature.
EHCM	Added fingerprint support for CMM update feature.
EHCM	Image bank 1 represents the active image and image bank 2 represents the non-active image.
VPD	Boot Rom will be reported in image segment 1.
WebUI	Security enhancements.
Release 3.3.3500	
Management Interfaces	In-band management.
Ethernet Switching	Port mirroring.
Ethernet Switching	sFlow.
EHCM	Enhancements to software update using CMM feature.
Release 3.3.3400	
EHCM	Added fwImageProtocols OIDs support.
Release 3.3.3000	
General	New Linux kernel 2.6.32.
WebUI	Added VLAN configuration page to the WebUI. Added temperature critical and warning thresholds to temperature graph.
U-boot	Updated u-boot – memory access optimization.
Network Interfaces	Added interface range support. By using the interface range configuration mode, a range of ports can be easily configured with the same parameters.
Ethernet Interfaces	Added Ethernet interfaces ingress and egress rate counters. The sampling interval can be tuned using the command load-interval in the interface configuration mode. The output is given in the “show interface ethernet” command. Refer to the CLI reference guide for additional details.
Release 3.2.1050	
WebUI	Added support for internal ports in WebUI.
WebUI	Applied new Apache version.

Table 2 - IBM 90Y3476 Changes and New Features

Category	Description
SNMP	Added support for SNMPv3 notifications.
Configuration Management	Removed “jump-start configuration wizard” feature.
Network Interfaces	Quality enhancement to link initialization.
Software Management	Added support for fetching image from TFTP server using IPv6.

4 Known Issues



For any possible errata due to hardware issues, please refer to the switch support product page.

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

4.1 General Known Issues

Table 3 - General Known Issues (Sheet 1 of 3)

Index	Category	Description	Possible Workaround
1.	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.
2.	Management Interfaces	The switch might have an expired HTTPS certification.	Generate a new certificate by changing the hostname.
3.	In-Band Management	In-band management mode does not support IPv6.	
4.	In-Band Management	In-band management ARPs are not taken into account in the “show ip are count” command.	
5.	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 ignore.	
6.	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.	
7.	Software Management	Upgrading to v3.2.1050 (or later) forces all ConnectX®-3 adapters that are connected to this switch to upgrade their firmware to adapter firmware v2.11.0960 (or later). For switch-to-switch connections, it is highly recommended that all SwitchX® switches be upgraded to MLNX-OS™ version v3.2.1050 (or later) to guarantee robust interoperability.	
8.	Software Management	In case of an upgrade, ACL sequence-number “0” is not valid any more. In case there are two ACL entries with sequence number of “0” and “1” (sequential entries), the ACL with sequence-number 0 is deleted. In other cases where there is an ACL entry with sequence-number “0”, it is moved to be sequence-number 1 within the system upgrade procedure.	In case you have an ACL entry with sequence number “0”, recreate it with different sequence number, prior to the upgrade procedure.

Table 3 - General Known Issues (Sheet 2 of 3)

Index	Category	Description	Possible Workaround
9.	Software Management	Upon upgrade if there are more than 4 ACLs, all of the ACL configuration is deleted. If there are 4 ACLs or less, rules with a sequence number larger than 495 are deleted.	
10.	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.	
11.	Configuration Management	After loading a new configuration file, please reboot the system. Otherwise, configuration might not be properly applied.	
12.	Configuration Management	The command set <code>revert {factory [keep-basic keep-connect] saved}</code> is removed.	Use the equivalent CMM command instead
13.	Configuration Management	Merging two binary configuration files using the command <code>configuration merge</code> is currently not supported.	Use the configuration text file “Apply” option instead.
14.	Logging	“DROPPED MSG” errors may appear during reload (shutdown phase). These errors can be safely ignored.	
15.	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.
16.	User Management	Logging into the system as USERID from the Serial Connection results in login failure the first attempt.	Log in again. The second attempt will result is successful login.
17.	WebUI	Only a single WebUI login per host is possible.	To have multi-sessions, open several browser simultaneously.
18.	WebUI	Reversing the time clock can result in WebUI graphs’ corrupted data.	Clear the graphs data after setting the clock.
19.	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.
20.	WebUI	Switching between binary configuration files when connected to the WebUI using HTTPS might result in the following message being displayed: “Switched configuration to ‘***’, which was already the active database.” This message is incorrect and can be safely ignored.	

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

Index	Category	Description	Possible Workaround
21.	SNMP	The error “Cannot find module (MELLANOX-MIB)” can appear in the log when performing roll-back to an older MLNX-OS version. This error can be safely ignored.	
22.	Chassis Management	Upon reaching critical thermal threshold, SR bit 2 is not set although the system is shut down and SR bit 3 is set instead.	
23.	IP Routing	The current release does not support L3.	

4.2 Ethernet Known Issues

Table 4 - Ethernet Known Issues (Sheet 1 of 3)

Index	Category	Description	Possible Workaround
1.	Ethernet Interfaces	The system allows tolerance of 4 bytes on top of the set MTU.	
2.	Ethernet Interfaces	Running RDMA traffic without enabling flow control might cause poor performance.	Enable flow control over the specific ports.
3.	Ethernet Interfaces	In case of faulty cable, rise time of other interfaces in the system may be delayed.	
4.	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
5.	Ethernet Interfaces	Default MTU is modified as 1500 instead of 1522.	
6.	Ethernet Interfaces	There are no port counters for packet sizes 1518-1522.	
7.	Ethernet Interfaces	Setting an interface switchport mode to “access” using the command <code>switchport mode</code> then deleting the access VLAN using the command <code>no vlan</code> results in said interface showing as a member of VLAN 0. If the running-config of that switch is exported to another, the process fails.	Reconfigure the interface to have a valid access VLAN.
8.	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.
9.	LAG	The remote LACP admin-key #1 cannot be established with some third party vendors.	Use a different admin-key number.
10.	LAG	Traffic running over LAG may not be evenly distributed when testing small number of SMACs.	
11.	LAG	When converting static LAG to dynamic LAG (LACP), some of the ports remain in DOWN state.	Disable and enable the LAG.

Table 4 - Ethernet Known Issues (Sheet 2 of 3)

Index	Category	Description	Possible Workaround
12.	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 workaround is applicable in case there are multiple mixed flows.
13.	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.
14.	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 	
15.	VLAN	It might take approximately 20 seconds to create/delete 1000 VLANs, or to change the switchport mode of an interface.	
16.	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.
17.	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.
18.	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.	
19.	IGMP Snooping	Setting the <code>ip igmp snooping unregistered multicast option to forward-to-mrouter-ports</code> option reduces the amount of supported VLANs to 250.	

Table 4 - Ethernet Known Issues (Sheet 3 of 3)

Index	Category	Description	Possible Workaround
20.	IGMP Snooping	Only one mrouter interface can be added on a certain VLAN when IGMP Snooping is disabled on that VLAN (or globally). If more are added, they overwrite their predecessors.	Before adding an mrouter interface, enable IGMP Snooping globally and per VLAN using the command "ip igmp snooping".
21.	LLDP	LLDP notifications are not supported.	
22.	LLDP	The mgmt0, mgmt1 IPv6 addresses and the in-band management address are not advertised on LLDP TLV(s).	
23.	ACLs	A mask of IPv4 ACL rule can contain either 255 or 0 in any of the 4 bytes assembling it.	
24.	ACLs	Packets dropped by the switch due to congestion or ACL rules are added to the bad type counter of the matching VLAN interface.	
25.	Modules	Using 1GbE copper modules leads to the port status being constantly reported as active regardless from the actual link state.	
26.	SNMP	LLDP MIB OID lldpRemPortDesc is not supported.	
27.	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 drops packets from all flows not according to their configured priorities.	
28.	sFlow	The discarded packets counter in sFlow samples may not be accurate and may expose a skew of one second.	

5 Bug Fixes

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

Table 5 - Bug Fixes

Index	Category	Description
1.	Configuration Management	Binary and text configuration files apply fix.
2.	Configuration Management	The output of the commands “show running-config” and “configuration text generate” are not displayed properly when using the VLAN range commands (e.g. “interface ethernet <port> switchport trunk/hybrid allowed-vlan <range>”). Applying a configuration text file generated with these commands on a switch will result with port configured as being a member only of the last VLAN in the range.
3.	Configuration Management	The output of the commands “show running-config” and “configuration text generate” is not displayed properly after the command “dcb priority-flow-control” enable is used.
4.	EHCM	Switch reload is added after “update -activate” CMM command run.
5.	EHCM	Image update may not complete successfully if performed using CMM.

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.