



InfiniScale IV fw-IS4 Release Notes

For fw-IS4 Rev 7.4.3000

NOTE:

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



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

© Copyright 2013. Mellanox Technologies. All Rights Reserved.

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

Connect-IB™, FabricIT™, MLNX-OS™, MetroX™, ScalableHPC™, Unbreakable-Link™, UFM™ and Unified Fabric Manager™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

Table of Contents

Table of Contents	3
Chapter 1 Introduction	4
1.1 Requirements	4
Chapter 2 Changes an Major New Features	5
2.1 Changes From Version 7.4.2070 to 7.4.3000	5
2.2 Changes From Version 7.4.2040 to 7.4.2070	5
2.3 Changes From Version 7.4.0000 to 7.4.2040	5
2.4 Changes From Version 7.3.200 to 7.4.0000	5
2.5 Changes From Version 7.3.100 to 7.3.200	5
2.6 Changes From Version 7.2.400 to 7.3.100	6
2.7 Changes From Version 7.2.000 to 7.2.400	6
Chapter 3 Unsupported Features	7
Chapter 4 Known Issues	8
Chapter 5 Bug Fixes History	9
Appendix A SMA/GSA Attributes	11

1 Introduction

These are the release notes for the MT48436 InfiniScale® IV firmware fw-IS4, Rev 7.4.3000. This firmware (FW) complements the InfiniScale IV MT48436 silicon architecture with a set of advanced features, allowing easy and remote management of the switch.

Important Notes:

- This firmware version does not support Avago active optical cables with cable firmware version 3.0. See issue #1 in [Table 1, “Known Issues,” on page 8](#) for more details.
- This firmware release complies with the InfiniScale IV MT48436 Programmer’s Reference Manual (PRM), Doc. #2885, Rev 0.40.
- After burning new firmware to the InfiniScale IV switch device, reboot the switch so that the new firmware can take effect.

1.1 Requirements

- One of the following burning tools:
 - **flint** or **mlxburn** applications (part of MFT ver. 2.7.0 or later)
- Burning firmware to the switch can be performed In-Band or via the I2C-compatible interface. Burning via the I2C-compatible interface requires one of the following I2C cards:
 - MTUSB-1 (USB to I2C adapter)
 - ISA Calibre

2 Changes and Major New Features

2.1 Changes From Version 7.4.2070 to 7.4.3000

- Improved MC join / leave processing time
- Added support for InterProcessPacket (IPC) in INIT_MAD_DEMUX
- Bug Fixes - see [Section 5, “Bug Fixes History”, on page 9](#)

2.2 Changes From Version 7.4.2040 to 7.4.2070

- Updated IS5023 INI to resolve wrong lane polarity

2.3 Changes From Version 7.4.0000 to 7.4.2040

- Bug Fixes - see [Section 5, “Bug Fixes History”, on page 9](#).

2.4 Changes From Version 7.3.200 to 7.4.0000

- New MAD: VendSpecLedControl – enables controlling port LEDs, regardless of hardware port state. This MAD also enables controlling other system status LEDs such as the Unit Identifier (UID) LED and the Bad Port LED.
- New MAD: LossyLinkConfig – is used for configuring “lossy” InfiniBand links to ignore credits on one or more VLs of a port’s TX lanes, and to ignore buffer overrun errors on one or more VLs of a port’s RX lanes
- New Bad port LED indication - The Bad Port LED is turned on if one or more ports do not operate properly
- New UID LED indication - The UID LED can be turned on/off to help in physically locating the switch system within a rack or a row of racks
- Added a secondary HCR
- ABA - During IBTA spec 1.2.1 speed negotiation, the serdes parameters are based on the cable attention as read from the QSFP eeprom
- Added fallback from Adaptive Routing mode to Normal Mode when a new Subnet Manager is activate

2.5 Changes From Version 7.3.100 to 7.3.200

- New MAD: CableInfo – enables reading and writing into/from QSFP devices
- New MAD: EyeOpeningResults – enables reading the results from InfiniScale® IV devices
- SwitchInfo MAD – enables setting multicast FDBTop
- SetPort Command – enables changing LinkWidthSupported and LinkSpeedSupported

Added on September 27, 2010

This firmware also enables the Status and Fan LEDs behavior to change as follows:

- Status LED: If one fan fails, then the Status LED will change to YELLOW. After the failing fan is replaced, the Status LED will return back to GREEN without the need for reset. Previously, the Status LED remained GREEN even after one fan failed.
- Fan LED: If one fan fails, then the Fan LED will change to RED. After the failing fan is replaced, the Fan LED will return back to GREEN without the need for reset. Previously, the Fan LED remained GREEN even after one fan failed.

2.6 Changes From Version 7.2.400 to 7.3.100

- Support for active cables, including:
 - Power class enforcement
 - Enhanced dynamic SerDes parameters selection based on the cable type indicated by the cable EEPROM
- Link-up flow per IBTA Release 1.2.1 is enabled by default
- Added a PrivateLFT – enables setting an FDB per port
- SwitchInfo MAD – added a new field for *multicast top*
- Added the following MADs:
 - LEDTest – a new vendor specific MAD that enables turning ON/OFF port LEDs regardless of port state
 - PortRcvDataVLExtended, PortXmitDataVLExtended, PortRcvPktsVLExtended, PortXmitPktsVLExtended – vendor specific 64-bit counters that provide statistics per port VL
- Added the following commands to the device's Command Interface:
 - PORT_REMAP – allows changing port mapping after boot prior to bringing up port links
 - SET_PORT – this command is now available for external ports as well (previously for Port 0 only). It allows modifying the PortInfo MTUCap and VLCap. Using this commands requires the port link to be in the INIT state
- New firmware configuration (.ini) parameters:
 - ignore_credits_port<#> – using this parameter enables the transmitting port to assume that the receiving port always has enough credits to receive the transmitted packet. Flow control packets are ignored. This feature can be enabled per port VL
 - MTU_CAP_IB_PORT<#> – allows setting the MTU to 8K

2.7 Changes From Version 7.2.000 to 7.2.400

- Mellanox auto-negotiation:
 - Enhanced the sweep procedure to yield better eye measurements
 - Added an optional crosstalk reduction algorithm. Enabling this algorithm multiplies the auto-negotiation duration by 4.
- Command interface:
 - Updated the SET_PORT command to allow for changing the node description

3 **Unsupported Features**

- Baseboard Management Agent (BMA)

4 Known Issues

Table 1 - Known Issues

Index	Issue	Description	Possible Workaround	Scheduled Release (fix)
1.	Switch links may not rise in systems connected with Avago AOCs	Avago active optical cables (AOCs) with cable firmware version 3 may prevent other cable modules from accessing the I2C bus, resulting in: wrong power class enforcement, wrong SerDes parameters selection, and the prevention of switch links from rising	Power cycle the entire switch system	Avago has released firmware version 9 which resolves this issue. Please contact your assigned Mellanox FAE for details.
2.	Wrong return values by the <i>Flow Monitoring Counters Info</i> or <i>Flow Monitoring Counters MADs</i>		NA	-
3.	AUTOCONF_1X4X8X12X setting is not supported	AUTOCONF_1X4X8X12X setting for the [LINK_WIDTH] parameter width_supported_clusterX is not supported	NA	-
4.	Wrong M_KeyViolations counter value is reported for Enhanced Port 0		NA	-
5.	Temperature Sensing	a. TRAP is not implemented. b. If the over-temperature warning GPIO mode is set (<i>wgm=1</i>), the GPIO cannot be cleared by MAD	NA	-

5 Bug Fixes History

Table 2 lists the bugs fixed in this release.

Table 2 - Fixed Bugs List

#	Issue	Description	Discovered in Release	Fixed in Release
1.	Symbol error counter	The symbol error counter was reset in link training	7.4.2070	7.4.3000
2.	Perf tool	Perf option mask was reported incorrectly	7.4.2070	7.4.3000
3.	MC groups	The switch could not support more than 2K MC groups	7.4.2070	7.4.3000
4.	PortMirror issues	PortMirror issues: <ul style="list-style-type: none"> PortMirrorRoute Get() always returned 0 When using Multicast LID, the out port is selected as the first port in the multicast group 	7.4.0000	7.4.2040
5.	TRAPS 256, 257, 258 data	Fixed minor bugs in TRAPS 256, 257, 258 data	7.4.0000	7.4.2040
6.	Credit allocation for single VL single port configuration,	Issue regarding the credit allocation for single VL single port configuration, 3Km link are now doable	7.4.0000	7.4.2040
7.	Device ID and Revision ID via MADS	Values of Device ID and Revision ID via MADS were wrongly reported	7.4.0000	7.4.2040
8.	Counter issues	Counter issues: <ul style="list-style-type: none"> LinkDownedCounter could not count beyond 0 ConfigGroupCounter when setting to all ports it failed Port0 counter per VL reported wrong values 	7.4.0000	7.4.2040
9.	VLCAP reporting of Port0	Issues with VLCAP reporting of Port0	7.4.0000	7.4.2040
10.	PhyConfig MAD	PhyConfig MAD, prevented access to Port0	7.4.0000	7.4.2040
11.	Traps	Trap 258 fields are empty	7.3.100	7.3.200
12.	Congestion Control	Minor Congestion Control bugs	7.3.100	7.3.200
13.	INI file issue	The user could not set a value bigger than 0xF in the preemp_post field	7.3.100	7.3.200
14.	Auto-negotiation	Mellanox auto-negotiation does not support 12X links	7.2.400	7.3.100
15.	LinkWidth in INI file	LinkWidth in INI file cannot be specified using IB port numbers	7.2.400	7.3.100
16.	Adaptive routing	Adaptive routing	7.2.400	7.3.100
17.	Speed auto-negotiation for 1X ports	Speed auto-negotiation for 1X ports always brings up the port at SDR	7.2.400	7.3.100
18.	M_Key check	No M_Key check is performed for packets sent by the external SMA	7.2.400	7.3.100

Table 2 - Fixed Bugs List

#	Issue	Description	Discovered in Release	Fixed in Release
19.	Overtemp warning GPIO cannot be set	The GPIO indicating the overtemp warning cannot be set because the support bit is missing in the ini file	7.2.100	7.2.200
20.	Command interface	command interface - using MAD_IFC causes unexpected behavior	7.2.000	7.2.100
21.		command interface - QUERY_FW does not support OpMod=0x2	7.2.000	7.2.100
22.		command interface - INIT_PORT and CLOSE_PORT are not supported	7.2.000	7.2.100
23.		command interface - bit BM in INIT_MAD_DEMUX is not supported	7.2.000	7.2.100
24.		HCA driver interface: HW2SW_DQ may post extra CQEs	7.2.000	7.2.100
25.		VL_CAP	Cannot set VL_CAP by ini	7.1.000
26.	Non- qp0/qp1 packets	Non- qp0/qp1 packets to the SMA of an unmanaged switch hang the FW	7.1.000	7.2.000
27.	PKEY	Inbound P_Key Enforcement partial bit is set by default instead of cleared	7.1.000	7.2.000
28.	Symbol error counter	Symbol error counter of IB Port1 is not cleared after boot	7.1.000	7.2.000
29.	Auto-negotiation	Mellanox auto-negotiation ini parameters	7.1.000	7.2.000
30.	PortCounters MAD	Cannot clear PortxMmitWait counter using set()	7.1.000	7.2.000
31.	ClassPortInfo	Wrong capability mask in ClassPortInfo of performance management class	7.1.000	7.2.000
32.	Auto-negotiation	If NOI/NRI of peer device is higher than InfiniScale IV NOI/NRI values, SerDes parameters may all be set to zero	7.1.000	7.2.000
33.	I2C address conflicts	Secondary I2C address conflicts with temperature sensor I2C address on Shark system	7.1.000	7.2.000
34.	MAD issues	Cannot use MAD to open a port disabled by INI file	7.1.000	7.2.000

Appendix A: SMA/GSA Attributes

The following tables summarize the attributes supported by the management agents provided in this release.

Table 3 - SMA Attributes

Attribute	Support
Notice	-
NodeDescription	+
NodeInfo	+
SwitchInfo	+
GUIDInfo	+
PortInfo	+
Partition Key Table	+
SLtoVLMappingTable	+
VLArbtration	+
LinearForwardingTable	+
RandomForwardingTable	-
MulticastForwardingTable	+
SMInfo	-
VendorDiag	-
LedInfo	-
PrivateLFTMap (VendorSpecific)	+
PrivateLFTAccess (VendorSpecific)	+
PrivateLFTTops (VendorSpecific)	+
AdaptiveRoutingInfo (VendorSpecific)	+
AdaptiveRoutingGroupTable (VendorSpecific)	+
AdaptiveRoutingLinearForwardingTable (VendorSpecific)	+
RemotePortMirroring (VendorSpecific)	+
TemperatureSensing (VendorSpecific)	+
CableInfo (VendorSpecific)	+
EyeOpeningResult (VendorSpecific)	+
VendSpecLedControl (VendorSpecific)	+

Table 3 - SMA Attributes

Attribute	Support
LossyLinkConfig (VendorSpecific)	+

Table 4 - Mandatory Performance Management Attributes

Attribute	Support
ClassPortInfo	+
PortSamplesControl	+
PortSamplesResult	+
PortCounters	+
PortCountersExtended (for data counters only)	+

Table 5 - Optional Performance Management Attributes

Attribute	Support
PortRcvErrorDetails	+
PortXmitDiscardDetails	+
PortOpRcvCounters	-
PortFlowCtlCounters	+
PortVLOpPackets	-
PortVLOpData	-
PortVLXmitFlowCtlUpdateErrors	+
PortVLXmitWaitCounters	+
PortCountersExtended	+
PortSamplesResultExtended	+
SwPortVLCongestion	+
PortXmitConCtrl	+
PortVLXmitTimeCong	-
PortXmitDataSL	+
PortRcvDataSL	+

Table 6 - VendorSpecific MADs

InfiniBand Management Class	Management Attribute	Supported by Last Release	Supported by New Release
0x0A VendorSpecific class	0x0001 ClassPortInfo	+	+
	0x0011 PortPowerState	-	-
	0x0012 DeviceSoftReset	+	+
	0x0013 ExtPortAccess	+	+
	0x0014 PhyConfig	+	+
	0x0015 MFT	-	-
	0x0017 GeneralInfo	+	+
	0x0050 ConfigSpaceAccess	+	+
	0x0060 PortRcvDataVL	+	+
	0x0061 PortXmitDataVL	+	+
	0x0062 PortRcvPktsVL	+	+
	0x0063 PortXmitPktsVL	+	+
	0x0068 FlowMonitoringCountersInfo	-	-
	0x0069 FlowMonitoringCounter	-	-
	0x0070 CongestionNotificationInfo	-	-
	0x0073 PortRcvDataVLExtended	-	+
	0x0074 PortXmitDataVLExtended	-	+
	0x0075 PortRcvPktsVLExtended	-	+
	0x0076 PortXmitPktsVLExtended	-	+
	0x0080 Virtual switch Info	-	-
	0x0090 CounterGroupInfo	+	+
	0x0091 ConfigCounterGroup	+	+
	0x00A0 EnhancedConfigSpaceAccess	-	-