



Release Notes

MT23108 InfiniHost Firmware

fw-23108 Rev 3.5.000

© Copyright 2006. Mellanox Technologies, Inc. All Rights Reserved.

MT23108 InfiniHost Firmware Release Notes

Document Number:

Mellanox Technologies, Inc.
2900 Stender Way
Santa Clara, CA 95054
U.S.A.
www.mellanox.com

Tel: (408) 970-3400
Fax: (408) 970-3403

Mellanox Technologies Ltd
PO Box 586 Hermon Building
Yokneam 20692
Israel

Tel: +972-4-909-7200
Fax: +972-4-959-3245

Mellanox Technologies

1 Overview

These are the release notes for the MT23108 InfiniHost firmware, fw-23108 Rev 3.5.000. This firmware supports the MT23108 device (A1 silicon version only), and the Mellanox HCA Adapter Cards listed in Table 1.

Note: After burning new firmware to an HCA board, reboot the machine so that the new firmware can take effect.

Table 1 - Supported HCA Adapter Cards

HCA Card OPN	Description
MHET2X-1SC / MHET2X-ITC	InfiniHost™ HCA card, dual-port, SDR w/ media adapter support, PCI-X, 128MB memory, short/tall bracket, RoHS R5
MHET2X-2SC / MHET2X-2TC	InfiniHost™ HCA card, dual-port, SDR w/ media adapter support, PCI-X, 256MB memory, short/tall bracket, RoHS R5
MHXL-CFXXX / MHXL-CFXXX-T ¹ (previously MTLP23108)	Legacy InfiniHost™ HCA card, dual-port, SDR, PCI-X, 128/256/512MB memory, low profile short/tall bracket
MHX-CEXXX-T ¹ (previously MTPB23108)	Legacy InfiniHost™ HCA card, dual-port, SDR, PCI-X, 128/256/512MB memory, short/tall bracket

1. XXX reflects the size of on-board memory (in MB): 128, 256, or 512. OPN with -T indicates a tall bracket card; OPN without -T indicates a short bracket card.

This document consists of the following sections:

- “Major Changes and New Features” (page 3)
- “Bug Fixes” (page 4)
- “Invariant Sector (IS) Changes / Fixes” (page 4)
- “Known Issues” (page 5)
- “History” (page 6)
- “Creating a Device Configuration (.ini) File” (page 8)

2 Major Changes and New Features

- Indexing to the translation cache can now be programmed via the firmware INI file. See the parameter *tpt_map* in the file `fw-25204-defaults.ref`.
- Support for *log_ge_per_port* parameter which is used for assigning gather engines to queues. See Section 2.4, “HCA QoS Control” in the *InfiniHost Programmer’s Reference Manual, Rev. 1.19, Document no. 2111PM*.

3 Bug Fixes

The following table describes known issues from previous releases of InfiniHost firmware which were fixed in this firmware release.

Table 2 - Bug Fixes

Index	Issue	Description	Discovered in	Fixed in
1.	Wrong Client ReRegister bit	(ID:35356)	3.4.00	3.5.000
2.	Long PCIx retries	HCA may cause PCIx back-pressure upon PCIx DoorBells stress	3.4.00	3.5.000
3.	QP context corruption	May occur under stress (ID:36209)	3.4.00	3.5.000
4.	Corrupted Read/Atomic response	Possible corruption of Read or Atomic responses (ID:36353,37014,37026)	3.4.00	3.5.000
5.	Vendor Specific MADs	- Wrong Vendor Key (V_KEY) authentication and mismatch behavior (ID:36127) - Wrong Revision ID in General Info MAD (ID: 38025)	3.4.00	3.5.000
6.	RDMA_READ and ATOMIC Dead-lock	Internal leak may occur when doing RDMA_READs and ATOMIC on Send Queues (ID: 27761)	3.4.00	3.5.000
7.	Send WQ corruption	May occur for POST_LIST with RDMA_READs (ID:35649)	3.4.00	3.5.000
8.	Duplicated CQE	Multiple UD CQEs for same Recv WQE (ID:36075)	3.4.00	3.5.000

4 Invariant Sector (IS) Changes / Fixes

None in this firmware release.

5 Known Issues

The following table describes known issues in this firmware release and possible workarounds.

Table 3 - Known Issues

Index	Issue	Description	Current Implemented Workaround in FW	Possible Workaround in Driver	Patch Release (fix)	Scheduled Release (fix)
1.	MSIx vectors	Writing to MSIX vectors (Address/Data/Mask) does not take immediate effect. There may be MSIX messages that leave the device according to the old vector.	NA	Commit a PCI configuration cycle after the MSIX modification	NA	NA
2.	QUERY_DDR	Query does not return JEDEC vendor ID yet. Scope of status is limited to active / not active.	NA	NA	NA	NA
3.	RTR2RTS_QPEE; SQD2RTS_QPEE: changing optional fields rra_max and ra_buf_index is not supported.	The optional fields rra_max and ra_buf_index are not supported in the RTR2RTS_QPEE and SQD2TRS_QPEE commands.	Change requests for these fields do not take effect, and no error indication is provided.	Mask these optional fields.	NA	NA
4.	PCI 2.3 control and status for interrupts	InfiniHost does not support PCI2.3 control and status bits for interrupts.	NA	NA	NA	NA
5.	SW reset in memory controller mode	On PCIX systems with the bus downgraded to PCI: When the Flash image is corrupted and InfiniHost comes up as a memory controller and the I2C connector is attached, a SW reset may hang the system.	NA	1. Reboot the system. 2. Disconnect I2C when issuing a SW reset.	NA	NA
6.	Config cycles during sys_en	PCIX cfg cycles issued while the system-enable command is in progress may take a long time to complete. This causes some ServerWork chipsets to time-out and hang the system.	NA	Do not issue config cycles during a sys_en command.	NA	NA
7.	Client ReRegister bit set	The Client ReRegister bit may be set in a response to SubnGet(). (ID:33958)	NA	NA	NA	NA

6 History

Table 4 - History of Bug Fixes

Issue	Description	Discovered in	Fixed in
Unreliable QP context corruption under stress	Fixed (ID: 31927)	3.3.5	3.4.000
RNR timeout is always set to 0x1F	RNR timeout is always set to 0x1F regardless of the QP set value (ID: 31959)	3.3.5	3.4.000
QP Flush during an SQ Drain may result in corruption	Flushing the QP during an SQ Drain may cause a schedule queue corruption. (ID: 34879)	3.3.5	3.4.000
Reset during SYS_EN execution may cause corruption	Fixed (ID: 34196)	3.3.5	3.4.000
False PCI-X Parity Error assertion	A PCIx Parity Error may be asserted with no reason (ID: 32665)	3.3.5	3.4.000
APM with changed SL may result in corruption	A schedule queue corruption may occur upon an APM with changed SL. (ID: 32672)	3.3.5	3.4.000
Wrong M_Key check	Given MKeyProtectBits<2, if a SubnGet(NodeInfo) with a wrong M_Key is received while the M_Key Lease Period counter is already active (due to a previous M_Key violation), the counter gets stopped instead of continuing with the original count. (ID: 33388)	3.3.5	3.4.000
Exit Self-Refresh is not followed by an INIT operation	Fixed	3.3.5	3.4.000
ACK is delayed when the sender is back-pressured	This caused a transport timeout to the remote peer. Fixed.	3.3.5	3.4.000
Incorrect ECC error handling	ECC errors may not be reported at all, may cause a wrong SERR# assertion, or may be reported with a wrong column address. (ID: 31614)	3.3.5	3.4.000
MAD status	MAD status may be "incorrect Method and Attribute combination" instead of "unsupported attribute". Fixed. (ID: 35226)	3.3.5	3.4.000
MADs:PortInfo Get()	When querying for information about an InfiniHost III Ex IB port via its other IB port, the wrong Local port number is returned. Instead of the number of the second port, the one which received the MAD packets, the number of the first port is being returned. (ID: 24177)	3.3.2	3.3.3
Requester ScatterList corruption upon CQ error	A CQ error can cause corruption in the Requester ScatterList Database. As a result QPs may move to error, and the device may stop sending packets (30670)	3.3.2	3.3.3
Big UAR pages	Support for big UAR pages is not complete (ID: 29496)	3.3.2	3.3.3
Multicast Index mis-calculation	Multicast Index mis-calculation may cause dropping of multicast packets instead of inserting them. (ID: 29469)	3.3.2	3.3.3
CQ error or QP error together with 2ERR_QPEE may cause CommandIF to hang	(ID: 29431,29737)	3.3.2	3.3.3
FW deadlock when flushing a QP	(ID: 29277)	3.3.2	3.3.3
After a Catastrophic Error, HCA start may fail	(ID: 29066)	3.3.2	3.3.3
Port state ACTIVE_DIFFER should be reported as ACTIVE	(ID: 28811)	3.3.2	3.3.3
DIMM Unrecoverable Error not detected	(ID: 28902)	3.3.2	3.3.3

Table 4 - History of Bug Fixes (Continued)

Issue	Description	Discovered in	Fixed in
EQC.intr for the Catastrophic Error EQ is hard wired to 0x0	It now can be any legal value (including MSIx) (28815,28377)	3.3.2	3.3.3
SRQ performance is too low	(ID: 28702)	3.3.2	3.3.3
MSIx vector race when updating MSIx Table	(ID: 26599)	3.3.2	3.3.3
UD starvation	UD messages are not sent because RC ACKs are not arriving (ID: 28374,28427)	3.3.2	3.3.3
DIMM timing parameters	DIMM timing parameters are not configured correctly.	3.3.2	3.3.3
Configuration Cycles get many RETRIES when FW does SysEn/SysDis	(ID: 21332)	3.3.2	3.3.3
Long ScatterLists may cause data corruption	(ID: 28888)	3.3.2	3.3.3
Enforcement of burnt Max_Read_Byte_count	Burnt Max_Read_Byte_count was not enforced. (ID: 29807)	3.3.2	3.3.3
Single Symbol Errors are breaking the link instead of being ignored	Cannot raise link with IB spec 1.2 compatible device. (ID: 28465)	3.3.2	3.3.3
Consumer Index corruption in a Completion Queue	When using Increment_CI doorbells, to increment the CI in more than 1, CI may advance wrongly, casing a false CQ overrun, or not detecting a real overrun (ID: 27893)	3.3.1	3.3.2

The following table summarizes the history of Invariant Section bug fixes:

Table 5 - IS Changes/Fixes

Issue	Description	Discovered in	Fixed in
FailSafe failure report	Corruption of FW image in NVRAM does not result in a correct PCI cfg syndrome in register 20 as required. (ID:15912)	1.16.0000	1.18.0000
Header corruption on boot2 section	If one of the first two words in the image is corrupted (0), FailSafe code cannot recover nor respond to configuration cycles. (ID:16181)	1.16.0000	1.18.0000
IS memory controller non-PCIX compliant	On PCIX systems, when InfiniHost comes up as a memory controller, a SW reset may hang the system.(ID:16181)	1.16.0000	1.18.0000
FailSafe when two images are damaged	In Failsafe mode, when both the primary and secondary images are damaged, InfiniHost does not come up as a memory controller.	1.16.0000	1.16.0008

7 Creating a Device Configuration (.ini) File

Mellanox firmware burning tools enable setting and/or changing configuration variables by the use of an *optional* configuration (.ini) file. This is needed in case the default values of some variables do not suit a user's specific system requirements. This section describes how to create this configuration file.

To begin with, the .ini file is a text file is composed of one or several configuration *sections* (see [Section 7.1](#) for the format and/or an example). It is recommended to include, under the appropriate sections, only those variables that need to be changed.

A firmware release includes a reference file called fw-23108-defaults.ref. This file contains the list of all variables which can be configured by a configuration (.ini) file. For each variable the reference file includes a short explanation, the [<section>] it should be under, the range of possible values, and a line with the *default* setting of the variable which is assumed by the firmware release.

To create the .ini file, simply copy the lines with the variables you wish to set, paste them under their appropriate [<section>] headings, and change the setting values as desired.

7.1 Configuration (.ini) File Format

The .ini file is composed of one or more *sections* with variable settings. Each section in the file starts with its name between square brackets, e.g. [ADAPTER], [HCA], [IB], etc. The section name is followed by one or more lines of configuration settings and comments, as in the .ini file example shown below. Note that comment lines start with a semicolon.

Excerpt from fw-23108-defaults.ref:

```
;;;; VPD support can be Disabled/Enabled
;;;; Under [ADAPTER] section
;;;; Boolean parameter. Possible values: true, false .
```

```
vpd_enable = true
```

Example of a .ini file:

```
;Begin of .ini file

[ADAPTER]
vpd_enable = false

;This is a comment line
;End of .ini file
```