



Release Notes

MT25208 InfiniHost III Ex Firmware

(MemFree)

FW-25218 Rev 5.1.0

© Copyright 2005. Mellanox Technologies, Inc. All Rights Reserved.

MT25208 InfiniHost III Ex Firmware FW-25218 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 MT25208 InfiniHost III Ex firmware, FW-25218 Rev 5.1.0. It is appropriate for the MT25208 device functioning in MemFree mode (that is, without local memory). This firmware supports 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	Code Name	Description
MHEL-CFXXX-T ¹ (previously MTLP25208)	Lion Cub SDR	InfiniHost III Ex PCI Express HCA Adapter Card. Note: IB ports support operation at single data rate (SDR) only.
MHEA28-XT	Lion Mini SDR	InfiniHost III Ex PCI Express MemFree HCA Adapter Card. Note: IB ports support operation at single data rate (SDR) only.
MHGA28-XT	Lion Mini DDR	InfiniHost III Ex PCI Express MemFree HCA Adapter Card. Note: IB ports support operation at double data rate (DDR).

1. XXX reflects the size of on-board memory (in MB): 128, or 256.

The document consists of the following sections:

- “Major New Features” (page 4)
 - “Future Unsupported Features” (page 4)
- “Known Issues” (page 5)
- “Bug Fixes” (page 7)
- “Creating a Device Configuration (.ini) File” (page 9)

2 Major New Features

- Auto Negotiation for IB ports Double Data Rate (DDR) operation.

Note: After burning new firmware to an HCA board, reboot the machine so that the new firmware can take effect. However, if the new firmware changes PLL frequencies, a complete shutdown of the machine is needed.

- PCIe BARs can be configured using the .ini file to be 32 bits wide rather than 64 bits. Each BAR can be configured separately.
- Client ReRegister flow is supported. An Event is generated according to the InfiniBand Architecture Specification, Vol.1, Release 1.2.
- Event generation upon APM State moving from ReARM to ARMED.
- The MT25208 device can be configured to push DIMMs with weak buffers (SSTL1).
- GPIO can be configured to disable an IB link.
- The InfiniHost III Ex device can be configured to generate events upon GPIO pin changes.
- POST_DOORBELL command is supported. This command is equal to ringing a DoorBell on the Privileged UAR.
- The new outbox format for the QUERY_DEBUG_MESSAGE command is supported.
- Mellanox firmware burning tools can now be provided with configuration (.ini) files instead of board (.brd) files to change configuration variable default values assumed by firmware. See “Creating a Device Configuration (.ini) File” on page 9.

2.1 Future Unsupported Features

- This is the last release to support board (.brd) files. Starting with the next release, only configuration (.ini) files will be supported.

3 Known Issues

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

Table 2 - Known Issues

Index	Issue	Description	Current Implemented Workaround in FW	Possible Workaround	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.	QPC.Flight_LIM	QPC field – no HW limit, infinite WQEs on send.	NA	NA	NA	NA
3.	QUERY_DDR	Query does not return JEDEC vendor ID yet. Scope of status is limited to active / not active.	NA	NA	NA	NA
4.	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 will not take effect, and no error indication is provided.	Mask these optional fields.	NA	NA
5.	PCI 2.3 control and status - for interrupts	InfiniHost III Ex does not support PCI2.3 control and status bits for interrupts.	NA	NA	NA	NA
6.	Change of memory bars on a disabled system	Changing memory bars size / addresses between SYS_DIS and SYS_EN may cause the InfiniHost III Ex to hang (ID: 24206)	NA	NA	NA	NA
7.	BAR resizing on an enabled system	Changing bar sizes when a system is enabled may cause the InfiniHost III Ex to hang (ID: 24208).	NA	NA	NA	NA
8.	SW reset via configuration cycles	SW reset via config cycles may create double PCI- Express completions for the configuration transaction.	NA	If InfiniHost III Ex boots in memory controller mode, perform power cycle / hot reset after restoring the flash.	NA	NA
9.	SW reset is performed during a configuration transaction	If SW reset is performed while a configuration transaction is outstanding, it may create double PCI- Express completions for the configuration transaction.	NA	Do not perform SW reset during configuration cycles.	NA	NA

3.1 Unsupported *InfiniHost III Programmer's Reference Manual* (Device-ID 25218) Changes

The following features of the *InfiniHost III Programmer's Reference Manual*, Document no. 2248PM, are not scheduled to be supported:

1. Flight lim value in QPC may show a value other than 0'1111 even when set for unlimited usage.

Mellanox Technologies

4 Bug Fixes

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

Table 3 - Bug Fixes

Issue	Description	Discovered in	Fixed in
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)	5.0.1	5.1.0
A Concurrent Bind and Deallocate for the same Memory Window will prevent closing the Memory Region of this Window	Bind and Deallocate modify the same 'unprotected' variable of the Memory Region. If both operations are attempted simultaneously for the same Memory Window, the variable does not get updated correctly. This prevents closing the Memory Region as the corrupted variable value may indicate that a Memory Window is still bound to it.	5.0.1	5.1.0
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 (ID: 30670)	5.0.1	5.1.0
IB Tx phase detector should be opened when link state is Config Debounce	(ID: 24332)	5.0.1	5.1.0
MLX QP may be broken under heavy load of DoorBells	(ID: 30544)	5.0.1	5.1.0
FW deadlock in an environment of BIND and HW2SW_MPT	(ID: 29814)	5.0.1	5.1.0
Internal error when receiving RDMA_READ or ATOMIC requests	RDB access fails when MTT base != 0x0 for an RDMA_READ or ATOMIC request, and internal error occurs. (ID: 29701)	5.0.1	5.1.0
QP deadlocks when doing a RETRY	(ID: 29676)	5.0.1	5.1.0
Big UAR pages	Support for Big UAR pages is not complete. (ID: 29496)	5.0.1	5.1.0
QP.send_doorbell_record_index not returned in queryQP	(ID: 29488)	5.0.1	5.1.0
CQ error may cause FW to deadlock when doing HW2SW_CQ	(ID: 29439)	5.0.1	5.1.0
Multicast Index miscalculation	Multicast Index miscalculation may cause dropping of multicast packets instead of inserting them. (ID: 29469)	5.0.1	5.1.0
CQ error or QP error together with 2ERR_QPEE may cause CommandIF to hang	(ID: 29431,29737)	5.0.1	5.1.0
Schedule Queue corruption may cause QPs to deadlock	(ID: 29292)	5.0.1	5.1.0
FW deadlock when flushing a QP	(ID: 29277)	5.0.1	5.1.0
SRQ deadlocks when QP goes to error	(ID: 29174)	5.0.1	5.1.0
After a Catastrophic Error, HCA start may fail	(ID: 29066)	5.0.1	5.1.0
Port state ACTIVE_DIFFER should be reported as ACTIVE	(ID: 28811)	5.0.1	5.1.0

Table 3 - Bug Fixes

Issue	Description	Discovered in	Fixed in
DIMM Unrecoverable Error not detected	(ID: 28902)	5.0.1	5.1.0
EQC.intr for the Catastrophic Error EQ is hard wired to 0x0	It now can be any legal value (including MSIx) (ID: 28815,28377)	5.0.1	5.1.0
e2e credits may not be seen for a short period. Thus UC QP may drop packets, and RC QP may RNR nack	(ID: 28825)	5.0.1	5.1.0
SRQ performance is too low	(ID: 28702)	5.0.1	5.1.0
MSIx vector race when updating MSIx Table	(ID: 26599)	5.0.1	5.1.0
SQ DoorBells may be lost in rare cases	(ID: 28646)	5.0.1	5.1.0
UD starvation because RC ACKs are not arriving	UD messages are not sent because RC ACKs are not arriving (ID:28374,28427)	5.0.1	5.1.0
SerDes electrical stress may occur if VDDIO > 1.2V	(ID: 28385)	5.0.1	5.1.0
Request Notification for multiple CQEs may take the wrong "multiple" number, and wait for more CQEs in order to notify	(ID: 28096)	5.0.1	5.1.0
MPT window count is corrupted when BIND is used excessively	(ID: 27759)	5.0.1	5.1.0
Fast Self Refresh feature is not functional	(ID: 30014)	5.0.1	5.1.0
DIMM timing parameters	DIMM timing parameters are not configured correctly	5.0.1	5.1.0

5 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 5.1 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-25218-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.

5.1 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-25218-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
```

Mellanox Technologies