



Release Notes

HP IB-Mezz 2 Image

hpmezz Rev 002_03_02

Mellanox Technologies

© Copyright 2009. Mellanox Technologies, Inc. All Rights Reserved.

hpmezz Image Release Notes

Mellanox Technologies, Inc.
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
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 *hpmezz Image, Rev 002_03_02*. This image supports the HCA Adapter cards listed in Table 1.

Table 1 - Supported HCA Adapter Cards

HCA Card OPN	Code Name	Description
448262-B21	HP IB Mezz 2	ConnectX IB HCA card, dual-port, DDR w/ media adapter support, PCIe x8, mem-free, tall bracket, RoHS R5

The image consists of the following components:

- CLP image -- Expansion ROM version 0x000200ab
- ConnectX[®] IB DDR fw-25408 firmware image, Rev 2.6.000

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

The remainder of this document consists of the following sections:

- “CLP Image” (page 4)
 - “CLP Command Strings” (page 4)
 - “CLP Image Fixed Bugs” (page 6)
- “ConnectX IB Firmware fw-25408” (page 7)
 - “Revision Compatibility” (page 7)
 - “Major New Features from v2.5.000” (page 7)
 - “Bug Fixes” (page 8)
 - “Known Issues” (page 9)
 - “Creating a Device Configuration (.ini) File” (page 10)

2 CLP Image

CLP is a library that enables configuring the HCA device before booting the OS. This allows management software to configure parameters that must **not** change while the device driver is running.

The CLP code resides in a Flash memory attached to the HCA device, and is presented to the software as an expansion ROM. The expansion ROM contains two binary images: x86 for a BIOS environment, and ia64 for a UEFI environment.

This CLP image version supports the following configuration capabilities:

1. Enable/disable an InfiniBand port
2. Set a port GUID
3. Set SerDes parameters
4. Restore to default

2.1 CLP Command Strings

This following describes the supported CLP command strings appropriate for use with HP IB-Mezz 2 Image Rev 002_03_02.

2.1.1 Command Strings

Note: X refers to a hexadecimal digit.

1. "set netport1 enabledstate=enabled"

Action: Enable the IB port

2. "set netport1 enabledstate=disabled"

Action: Disable the IB port

3. "set netport1 PermanentAddresses[0]=XXXXXXXXXXXXXXXXXXXX"

Action: Set the IB port GUID

4. "set netport1 OEMHP_hss=XXXX"

Action: Set the HSS parameters for lanes 0-1 of the IB port

5. "set netport2 OEMHP_hss=XXXX"

Action: Set the HSS parameters for lanes 2-3 of the IB port

6. "exit"

Action: Stop accepting CLP commands until the next power cycle. This implementation will complete every command before returning, thus it leaves no pending commands.

7. "set netport1 default"

Action: Restore device defaults to similar values after power on.

2.1.1.1 Notes

- Each 4X InfiniBand (IB) port has four physical lanes. In this implementation, each such two physical lanes are considered a "physical port". Thus, each IB port is comprised of two physical ports, and in the commands syntax the ports are referred to asnetports.
- Commands that operate on the IB port use "netport1" to represent the IB port.
- Commands that operate on the physical port use netport1 and netport2 to specify the target physical port.
- If a command that operates on an IB port specifies netport2, then the command will be completed with error.
- Example: The following command will end with error since it is supposed to disable the IB port, but netport2 does not denotean IB port.

```
set netport2 enabledstate=disabled
```

2.1.2 Returned Output Values

1. Successful execution of the command

```
[DS:ESI] "status=0\nstatus_tag=COMMAND COMPLETED\nendoutput\n"  
[AL] 0  
[AH] 2  
[EAX 30-16] 0  
[EAX 31] 0
```

2. The command string does not match exactly any of the command strings described above

```
[DS:ESI] "status=2\nstatus_tag=COMMAND PROCESSING FAILED\nerror=254\nerror_tag=COM-  
MAND NOT SUPPORTED\n"  
[AL] 2  
[AH] 254  
[EAX 30-16] 0  
[EAX 31] 0
```

3. The argument for a command is not recognized. For example

```
"set netport1 enabledstate=on" would cause the following to be returned.  
[DS:ESI] "status=2\nstatus_tag=COMMAND PROCESSING FAILED\nerror=250\nerror_tag=INVALID ARGUMENT\n"  
[AL] 2  
[AH] 250  
[EAX 30-16] 0  
[EAX 31] 0
```

4. Command interpreted but execution failed by the device

```
[DS:ESI] "status=3\nstatus_tag=COMMAND EXECUTION FAILED\nprob-  
cause=34\nprobcause_desc=Output Device Error\n"  
[AL] 3  
[AH] 0  
[EAX 30-16] 34  
[EAX 31] 0
```

2.2 CLP Image Fixed Bugs

The following table describes known issues from previous releases of the CLP image which were fixed in this release.

Table 2 - Bug Fixes

		Description	Discovered in CLP Version	Fixed in CLP Version
1.	SerDes parameter change	Optimized InfiniBand ports SerDes parameters	0x0002009e	0x000200ab

Mellanox Technologies

3 ConnectX IB Firmware fw-25408

These are the release notes for the ConnectX IB DDR Adapter firmware, fw-25208 Rev 2.6.000. This firmware is appropriate for the PCI Device IDs listed in Table 3, “PCI Device ID”.

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

Table 3 - PCI Device ID

PCI Device ID (Decimal)	Device Part Number	Device Name
25418	MT25408A0-FCC-DI	ConnectX IB (4X 20Gb/s IB)

3.1 Revision Compatibility

- ConnectX Programmer's Reference Manual (PRM), Rev 0.39 or later, which has Command Interface Revision 0x3. The command interface revision can be retrieved by means of the QUERY_FW command and is indicated by the field cmd_interface_rev.
- Firmware fw-25408 Rev 2 works with the OpenFabrics Enterprise Distribution (OFED) version OFED-1.3.1. Earlier versions of OFED will not work. See <http://www.openfabrics.org/downloads/OFED/>.

3.2 Major New Features from v2.5.000

- VPI support is now at GA level
- Support for QDR interoperability with InfiniScale IV switch platforms
- Support for Reserved LKey
- Support for Fast Registration Work Request (FRWR)
- Support for congestion control
- Support for Raw QP
- Support for QUERY_ADAPTER_vsd_vendor_id parameter
- Disabled the PCI Express Extended Tag feature
- Changes to the SET_PORT command (as described in PRM rev 0.39):
 - Support for the Reset Capability Mask (rcm) field
 - Added the ability to configure VLCap and MTUCap using the command SET_PORT
 - Support for the multicast loopback disable bit field
- Support for the sniffer feature
- Support for PortXmitWait Port Counter
- Added a re-read WQE bit to the WQE Control segment

The new bit is located at bit offset 6 of the first Dword in the Control Segment. See Table 19, “General - Ctrl Segment Format” in *ConnectX Programmer's Reference Manual (PRM), Rev 0.38*.

Firmware that supports the new bit sets the device capability bit 15 at offset 0x44 to 1. See Table 145, “Query Device Capabilities Layout” in *ConnectX Programmer's Reference Manual (PRM), Rev 0.38*
- Added support for twinax cables used for connecting Single and Dual SR ConnectX[®] EN 10 Gigabit Ethernet Adapters

- Ethernet performance improvement
- Firmware configuration (.ini) file change: To reflect a new assignment scheme for the PCI Express SubSystem ID parameter of each Mellanox adapter card, .ini files of standard Mellanox adapter cards now list the parameter `hca_header_subsystem_id`

3.3 Bug Fixes

The following table describes known issues from previous releases of ConnectX[®] firmware which were fixed in this firmware release.

Table 4 - Bug Fixes (Sheet 1 of 2)

	Issue	Description	Discovered in	Fixed in
1.	Possible live lock in QP upon retransmission stress	Fixed – (ID: 49870, 52066)	2.5.000	2.6.000
2.	Wrong link state reported during link speed negotiation	Fixed – (ID: 49951)	2.5.000	2.6.000
3.	CQs may be generated after CQ overrun error	Fixed – (ID: 49982)	2.5.000	2.6.000
4.	Slow handling of configuration cycles	Fixed – (ID: 49807)	2.5.000	2.6.000
5.	Wrong fields in CQE-w-Error on XRC QP	Fixed – (ID: 49742)	2.5.000	2.6.000
6.	Wrong handling of Remote Invalidate Error	Fixed	2.5.000	2.6.000
7.	Multiple RNR Nack may cause slowdown	Fixed – (ID: 49559)	2.5.000	2.6.000
8.	QUERY_DEV_CAP.apm bit was fixed at 0 even though APM was active	Fixed – (ID: 49548)	2.5.000	2.6.000
9.	PCIe physical errors upon entering L1 state	Fixed – (ID: 52025)	2.5.000	2.6.000
10.	PCI_CFG.interrupt_disable has no impact	Fixed – (ID: 53350)	2.5.000	2.6.000
11.	Non-default setting of VLCap via .ini does not take effect	Fixed	2.5.000	2.6.000
12.	SET_PORT may lead to non-optimal RX buffer reallocation if opvl was less than vlcap	Fixed	2.5.000	2.6.000
13.	Modified PLL parameter settings in .ini	Some PLL parameter settings were changed to allow a longer period for PLL stabilization	2.5.000	2.6.000
14.	Adapter may generate PCIe transactions with wrong function ID	Fixed	2.5.000	2.6.000
15.	Adapter may generate PCIe ERR_NON_FATAL in case of an unsupported request	Fixed	2.5.000	2.6.000
16.	<u>Ethernet only</u> : Different VLAN priorities for WQE and QP may cause wrong SchedQ allocation	Fixed	2.5.000	2.6.000
17.	Wrong VLAN Priority in PPP mode	The ConnectX device may send pause frames for the wrong priority in PPP mode (ID: 49533)	2.5.000	2.5.900
18.	Wrong aliasing in address steering mode	This issues applies only to the VMware [®] Infrastructure 3 v3.5 operating system in netq mode	2.5.000	2.5.900
19.	Possible packet dropping though the pause policy is set	Fixed	2.5.000	2.5.900
20.	Wrong PCI Class Code for Ethernet Network Controller	This issue applies to Ethernet devices only. Fixed.	2.5.000	2.5.900

Table 4 - Bug Fixes (Sheet 2 of 2)

	Issue	Description	Discovered in	Fixed in
21.	Bringing up/down of an adapter port may lead to PHY errors on the second adapter port	Applies only to dual-port adapter cards. Fixed (ID: 51356)	2.5.000	2.5.900
22.	PPP does not work for an MTU of 9600	Fixed	2.5.000	2.5.900

3.4 Known Issues

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

Table 5 - Known Issues

	Issue	Description	Current Implemented Workaround in FW	Possible Workaround	Scheduled Release (fix)
1.	UAR Bar is too small for 64k-page machines	The small BAR causes driver loading to fail	N/A	Change the "log2_uar_bar_megabytes".ini parameter under the [HCA] section as follows: log2_uar_bar_megabytes = 5	N/A
2.	Change of memory bars on a disabled system	Changing memory bars size / addresses between SYS_DIS and SYS_EN may cause the device to hang (ID: 24206)	N/A	N/A	N/A
3.	BAR resizing on an enabled system	Changing bar sizes when a system is enabled may cause the device to hang (ID: 24208)	N/A	N/A	N/A
4.	<u>Ethernet only</u> : Must query all capabilities upon boot	If not all capabilities are queried upon boot, then the query command may fail. See the QUERY_CAP command in <i>ConnectX EN Programmer's Reference Manual</i>	N/A	Query all capabilities upon boot	2.7.000
5.	Disrupting QDR negotiation may lead to port rising as SDR	Disconnecting an IB cable (or closing the port) during QDR negotiation and then reconnecting (or reopening) may cause the adapter to bring up the port at SDR	N/A	Disconnect the cable (or close the port) again and then reconnect (reopen). To avoid this scenario, wait for QDR negotiation to finish prior to disconnecting the cable (or closing the port) and reconnecting (or reopening). The following are two possible methods to verify QDR negotiation is complete: a. The physical (green) LED is on. b. A query of LinkPhyState using a GetPortInfo MAD indicates LinkUp.	N/A

3.5 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 3.5.1 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.

3.5.1 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
```