Connect. Accelerate. Outperform.™

Mellanox FlexBoot for ConnectX-3

Release Notes

Rev 3.4.648

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

Table of Contents

List of Tables

Document Revision History

Table 1: Document Revision History

Revision	Date	Description
3.4.648	September, 2015	Initial release of this FlexBoot version

1 Overview

These are the release notes for "Mellanox FlexBoot", the software for Boot over Mellanox Technologies adapter devices supporting VPI. FlexBoot enables booting kernels or operating systems (OS) from remote servers in compliance with the PXE specification.

This package is based on the open source project iPXE available at http://www.ipxe.org.

1.1 Supported Adapter Cards and Firmware

Table 2: Supported Adapter Cards and Firmware

Adapter Cards	Firmware Version
ConnectX-3 Pro	2.35.5000
ConnectX-3	2.35.5000

1.2 Supported Tools (MFT)

Table 3: Supported Tools (MFT)

Tools	Version
MFT (Mellanox Firmware Tools)	4.1.0

1.3 Supported Operating Systems

Table 4: Operating Systems

OSs	Version
ESXi	5.5
RHEL	6.4
RHEL	6.6
RHEL	7.0
RHEL	7.1
SLES	11.3
SLES	11.4 - Ethernet only
SLES	12
Ubuntu	14.04 - Ethernet only
Ubuntu	15.4
Windows Server	2012
Windows Server	2012R2

1.4 Tested Platforms

Table 5: Tested Platforms

Server Vendor & Model	BIOS Manuf. & Version
DELL C6220	Dell Inc. 2.5.3
DELL C6220 II	Dell Inc. 2.4.2
DELL C7220N	2.1.11
DELL M630	DELL 1.0.3
DELL R720	DELL 2.5.2
DELL R730	DELL 1.3.6
DELL Stryper	3A06
HP BL460c Gen9	I36 v1.50
HP DL160 Gen9	U20
HP DL360 Gen9	P89 v1.50
HP DL380 Gen9	P89 v1.50
HP Moonshot Anders M710	HP systems – H03
HP ProLiant SE2140s Gen8 (MiniPooch)	P87
HP ProLiant XL230b Gen9	U12
HP SE2140s Gen8	P87
HP SE2160w Gen9	U16
HP SL230s Gen8	P75
HP XL230a Gen9	U13 v1.40
HP XL230b Gen9	U12 v1.30
IBM Flex system x220	KSE142AUS-1.50
IBM Flex system x240	B2E142AUS-1.50
IBM System x3650 M4	VVE142EUS-1.80
IBM System x3850 X6 M4 (WillyCat)	A8E11AUS-1.00
Quanta T6MG	American Megatrends Inc. C1031.BS.3A09
Quanta Winterfell	American Megatrends Inc. F03_3B07
Supermicro X9DRW	American Megatrends Inc. 3.0c
Wiwynn WCS	American Megatrends Inc. C1032.BS.1C03.GN1
ZT	American Megatrends Inc. 2.59

2 Changes and New Features in Rev 3.4.648

Table 6: Changes and Fixes in Rev 3.4.648

Category	Description
FlexBoot Code Size	Added support for .mrom images larger than 128kB
Boot over IB	Added boot over IB with non-default PKey for ConnectX®-3, ConnectX®-3 Pro cards
Upstream sync	Synced the source with iPXE (upstream sync)
Spanning Tree Protocol	Added support for detecting Spanning Tree Protocol non-forwarding ports (RSTP/MSTP)
BIOS Compatibility	Moved to flat real mode when calling INT 1a,b101 to avoid BIOSes issues
Bug Fixes	See section Bug Fixes History.

3 Known Issues

The following is a list of general limitations and known issues of the various components of this FlexBoot release.

Table 7: Known Issues

Category	Description	Workaround
Firmware	Firmware 2.30.3000 and below does not support booting over Ethernet with VPI adapter cards when using QSFP cables (40GE, hybrid cables)	N/A
Booting from WDS	Booting from WDS running on Windows 2008 R2 in InfiniBand mode, is currently not supported due to compatibility issues between FlexBoot 3.4.142 and Mellanox WinOF v4.40 and below	N/A
BIOS	Several BIOS vendors have limited boot-vector space and may not display FlexBoot in their boot menu	Disable the embedded NIC boot agent in BIOS
BIOS	In several BIOS, the server might hang during FlexBoot booting due to wrong configuration of the PMM	N/A
Large Receive Offload (LRO)	Large Receive Offload (LRO) and iSCSI may not interoperate due to a bug in current Linux kernel distributions.	Disable LRO in the IPoIB module when using iSCSI. See the Mellanox FlexBoot user's manual for details under the Diskless Machines chapter (InfiniBand Ports).
iSCSI over IB	iSCSI over IB is not tested	N/A
iSCSI over DCB	iSCSI over DCB is not supported	N/A
MTU	Flexboot supports only 2K MTU	N/A
BIOS Compatibility	Only EBX,ESI,DS,ES registers can be saved in Boot Entry	N/A
Protocols	SLAM,FTP,HTTPS and SRP are currently not supported	N/A
FCoE, BCV	FCoE, BCV are not supported	N/A
Boot over VLAN	Boot over VLAN with IB port is currently not supported	N/A
FlexBoot Boot Menu	Boot menu is displayed as READ ONLY if the HCA card does not support flash configuration	N/A
PXE Boot after iSCSI Boot	PXE boot after iSCSI boot with static configuration is currently not supported	N/A

Category	Description	Workaround
iSCSI Boot	Boot retries is currently not functional when booting from iSCSI	N/A
Blink LEDs	Blink LEDs are currently not functional	N/A
iSCSI Connection	FlexBoot supports only a single active iSCSI connection. Thus, when iSCSI-boot via Port 1 succeeds to connect but fails to boot, it will fail to connect via Port 2.	N/A
FlexBoot Boot Menu	FlexBoot Boot Menu will not be visible in serial output	N/A
Link Speed	56Gb/s is currently not supported	N/A
Boot Loaders	Some faulty boot loaders do not close the underlying UNDI device which may result in unexpected behavior and possible system crash after the OS starts to load	N/A
BIOS	If a client returned control to the BIOS after a successful connection to an iSCSI target (but did not boot from it), then, unexpected behavior may occur.	Follow the instructions described in the UM for the proper iSCSI boot/install
IPv6	IPv6 is currently not supported	
Spanning Tree Protocol (STP)	Occasionally, using the Spanning Tree Protocol (STP) in the switches may cause packet drops and boot failure in the system.	Enable the "edgemode" if disabled on the switch, or use either portfast or edgemode functionality on the switch ports connected to the NICs.
SR-IOV	Setting the number of Virtual Functions higher than the machine's memory capability may cause memory issues and system instability	N/A
PXE Boot	In some cases, PXE boot will not work if the client was given only the filename without next-server (siaddr)	N/A
Chain-loading gPXE	Chain-loading gPXE stack may result in undesirable behavior	N/A

4 Bug Fixes History

Table 8: Bug Fixes History

#	Description	Fixed in Release
1.	Fixed HTTP boot over IPoIB	3.4.648
2.	Fixed BEV posting issue in hybrid BIOS	3.4.648
3.	Fixed transmission of IPoIB multicast packets as broadcasts	3.4.648
4.	Fixed HTTP download over IPoIB	3.4.521
5.	Fixed completion with error handling process	3.4.521
6.	Fixed an issue for TLV with length 0	3.4.306
7.	Fixed an issue related to sync VLAN IRQ operation with trunk IRQ operation	3.4.306
8.	Fixed an issue which enabled a netdevice (VLAN) to open/close twice	3.4.306
9.	Fixed an issue which prevented the iSCSI initiator's name from being received from HII	3.4.306
10.	Fixed an issue related to dual port adapters; occasionally, booting from the second port resulted in TFTP download failure when the first port was already linked up with DHCP, and has received a TFTP address	3.4.306
11.	Fixed an issue which caused PXE boot failure when using a filename if iSCSI rootpath is set	3.4.306
12.	Fixed an issue which prevented the device to PXE boot from the 2nd port if first port was already downloaded	3.4.306
13.	Fixed compilation issue	3.4.306
14.	Fixed a broken VLAN issue	3.4.306
15.	Fixed a retry issue when the value is infinite	3.4.306
16.	Fixed "Impossible to PXE boot from 2nd port if first port already downloaded." issue	3.4.225
17.	Fixed "no more network devices" issues	3.4.000

5 Change Log History

Table 9: Change Log History

Release	Changes
3.4.521	• Added iSCSI CHAP and mutual CHAP configuration • Added the GRH size when allocating receive buffer for IPoIB • Updated VLAN netdevice's settings with all the trunk's iSCSI required settings • Updated the port event handling process • Enabled console output in Debug mode • Disabled the serial output • Disabled the banner in BEV execution • Disabled function 0x04 (in int21) when serial console is disabled • Preserved COM port settings

3.4.460	• Boot Menu support: Added new FlexBoot GUI. The device can now be configured in the POST stage
	• Non-volatile memory read/write support
	• Configurable URI boot retry and delay between retries
	• Configurable iSCSI settings using DHCP/NVM
	• Added new interface in order to update the registered devices on the PXE stage
	• Enabled ConnectX Ethernet adapter cards family to work with interrupts
	• Enabled PXE to work in promiscuous VLAN mode (configurable through the INI)
	• Synced version with ipxe.org: Now the latest code in iPXE is used
	• Added boot priority capability: iSCSI vs PXE and fallback in case one fails
	• Updated the Proxy DHCP request method for non-existing option 54.
	ProxyDHCP request is sent to port 67 with broadcast IP address if the server identifier in option 54 is zero.
	Packets with source port different than BOOTPS_PORT and PXE_PORT are filtered by the PROXY
	• SHELL CLI is currently supported on ConnectX-3 and ConnectX-3 Pro adapter cards only
	• Both the GUID and the MAC are printed on the screen when the port link layer is set as InfiniBand
	• PROXYDHCP and PXEBS settings are saved under netdevice settings
	• rootpath/filename/nextserver are now fetched from the netdevice settings
	• The cached DHCP packet are received only if working with the same net device. When pxelinux.0 receives the cached DHCP packet from the UNDI API, it constructs a new (fake) packet for the current net device.
	  • If the process is stopped and then restarted and booted from the next boot device which serves as the second port in the HCA, a new (fake) DHCP packet is not constructed. The previous packet which includes all the information of the first port (IP, MAC, Netmask, etc...) is used.
	  • If an old (fake) DHCP packet is discovered, its chaddr is compared to the chaddr in the pxe_netdev, if not similar, a new (fake) DHCP packet is created.
	• PXE shutdown is called if int22 with function 0x000C is called.
	• The server's IP address in DHCP server replies is now checked before checking the reply type. This will ignore NACK replies from servers which already were ignored by the

Release	Changes
	client. In case of 2 DHCP servers in the same subnet, the client will eventually choose one of them, by sending the DHCP REQUEST with 'DHCP Server Identifier' (option 54) filled with the requested server's IP address. • Changed DHCP discover timeouts to comply with PXE spec
3.4.306	• Added validation script for the released ROMs • Added the option to always keep SAN hook to enable WIN install on iSCSI target • Added compilation flag around the flash readout • Added URI Boot retry. Default retries = 0 • Added Unmap MPT command in teardown • Added support for HII ISCSI configuration • Added 64-bit PCI BAR support (Large bar) • Added the option added for running PXE with promiscuous VLAN • Re-added COMBOOT image support by default • Enabled pages-function handling in Connect-IB initialization stage to work according to the PRM • Applied additional patches from ipxe.org • Updated the window even if ACK does not acknowledge new data • Modified the error print to debug print • Modified the printed string when initializing devices • Modified the error print. Added additional information to make the output more user-friendly • Changed the size of the domain name array to 0xfd • Disabled the waiting period for link up on trunk-net-device when VLAN is enabled on port • Removed unsupported EQ event in Connect-IB
3.4.225	• Added additional information to the error print output • Added compilation flag around the flash readout • Added URI Boot retry. Default retries = 0 • Added Unmap MPT command in teardown • Added 64-bit PCI BAR support • Added an option for running PXE with promiscuous VLAN • Added support for HII iSCSI configuration • Enlarged the mailbox size to 4kb • Enlarged the number of WQE to 64 (from 4) • Enabled multiple DHCP offers to be received before proceeding to request state • Changed the size of the domain name array to 0xfd • Changed error print to debug print • Changed printed string when initializing devices • Kept the SAN connection permanently open to enable Windows install on iSCSI target even when the iSCSI target is empty • Readded COMBOOT image support by default • Prevented a netdevice (VLAN) from opening/closing twice • Removed unsupported EQ event in Connect-IB

Release	Changes
	• Disabled the waiting time for link up on trunk net device when VLAN is enabled on a port
3.4.142	• Enabled firmware to handle the link state with the Subnet Manager • Updated the DHCP class code to NONE • Added flash access capability for reading software-to-software configurations • Enabled DHCP validation of MAC address and XID for a unique tuple • Improved randomness algorithm for DHCP XID
3.4.000	• Increased the number of the supported systems • Added the ability to merge the PXE image with the UFI image
3.3.420	• Increased DHCP timeout
3.3.400	• Added ConnectX-3 support • Increased the number of the supported systems • All iPXE features are now part of the ROM image • Support multiple functions (Flex10) • Added support for Link Aggregation Control Protocol (LACP) • Added ifenable/ifdisable network interface command • Improved reliability
3.0.000	• Initial Release