



**Mellanox MLX4_EN Driver for
VMware
README**

Rev. 1.5.6

www.mellanox.com

NOTE:

MELLANOX TECHNOLOGIES, INC. AND ITS AFFILIATES ("MELLANOX") FURNISH THIS DOCUMENT "AS IS," WITHOUT WARRANTY OF ANY KIND. MELLANOX DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT AND THOSE ARISING FROM A COURSE OF PERFORMANCE, A COURSE OF DEALING, OR TRADE USAGE. MELLANOX SHALL NOT BE LIABLE FOR ANY ERROR, OMISSION, DEFECT, DEFICIENCY OR NONCONFORMITY IN THIS DOCUMENT AND DISCLAIMS ALL LIABILITY, INCLUDING LIABILITY FOR INFRINGEMENT OF ANY INTELLECTUAL PROPERTY RIGHTS RELATED TO THE INFORMATION CONTAINED IN THIS DOCUMENT.

No license, expressed or implied, to any intellectual property rights is granted under this document. This document, as well as the software described in it, are furnished under a separate license and shall only be used or copied in accordance with the terms of the applicable license. The information in this document is furnished for informational use only, is subject to change without notice, and should not be construed as any commitment by Mellanox. Except as permitted by the applicable license, no part of this document may be reproduced, stored in a retrieval system, or transmitted in any form or by any means without the express written consent of Mellanox.

Names and logos identifying products of Mellanox in this document are registered trademarks or trademarks of Mellanox. Voltaire is a registered trademark of Mellanox Technologies, Ltd. All other trademarks mentioned in this document are the property of their respective owners.

Copyright © 2011 Mellanox Technologies, Inc. All rights reserved.



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

© Copyright 2011. Mellanox Technologies. All rights reserved.

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

FabricIT and SwitchX are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

Contents

1 Overview	4
1.1 Main Features Overview.....	4
2 Revision Compatibility	5
3 Driver Software.....	6
3.1 Installing and Running the Driver	6
3.1.1 Installing and Running the Driver on ESX-4.x	6
3.1.2 Installing and Running the Driver on ESXi-4.x	6
3.2 Remove Driver.....	7
3.3 Driver Default Values.....	7
4 Firmware Programming.....	9
5 HOW TOs	10
5.1 How to Increase PacketHeapSize	10
5.2 How to Enable MSI-X	10
5.3 How to Modify Jumbo Frames.....	11

1 Overview

This document provides instructions for installing the MLX4_EN drivers for Mellanox Technologies ConnectX® and ConnectX®-2 based network adapter cards in a VMware ESX-4.x Server environment. The ConnectX® and ConnectX®-2 adapters identify on the PCI bus as 25408, 25418, 25428, 25448, 25458, 26418, 26428, 26438, 26448, 26468, 26478, 26488 and 26458.

1.1 Main Features Overview

The following are the main features for MLX4_EN:

- Single/Dual port
- MSI-X / Int-X
- NetQueue support
- Multiple Tx/Rx rings
- HW Tx/Rx checksum offload
- Large Send Offload (TCP Segmentation Offload)
- VLAN Tx/Rx acceleration (HW VLAN stripping/insertion)
- Ethtool support
- NAPI support
- Rx/Tx traffic Rate Limiting per function
- Flex10 support
- WOL
- NCSI

2 Revision Compatibility

This driver is compatible with:

- Release build ESX-4.x
- ConnectX firmware fw25408, version 2.8.0600 and above

3 Driver Software

The driver is a single kernel module and has no software dependencies. The Ethernet MLX4_EN driver installation on VMware ESX Server is done using a Red Hat package manager (RPM).

The following steps describe how to download, install, and run the driver.

3.1 Installing and Running the Driver

3.1.1 Installing and Running the Driver on ESX-4.x

1. Log into the VMWare ESX Server machine as root.
2. Remove any earlier version of the driver from your VMWare ESX Server machine prior to installing the new version. Run:

```
#> esxupdate query
#> esxupdate --bundle <bundle to remove> remove
```

3. Install the mlx4_en driver package. Run:

```
#> esxupdate --bundle <mlx4_en full path> --maintenancemode update
```

4. Reboot ESX server.

3.1.2 Installing and Running the Driver on ESXi-4.x

1. Open a terminal on the vMA virtual machine.
2. Place the driver ISO CD in the drive of the desktop machine.
3. Right click on the vMA virtual machine and select **Edit Settings**.
4. Select **CD ROM** and choose the **Client Device** option in **Device Type**.
5. Select **Connect CD/DVD** in the vMA virtual machine console toolbar to connect the desktop CD Drive to the vMA virtual machine, and choose the CD drive in which the driver CD is inserted.
6. Mount the driver ISO into the local directory in the terminal on the vMA virtual machine:

```
# mkdir /ISO
# mount /dev/cdrom /ISO
```

7. Enter the host server in the maintenance mode.
 1. Right click **ESXi 4** in the vSphere Client
 2. Choose the option **Enter Maintenance Mode**.
8. In the vMA terminal, execute the following commands:

- For ESXi-4.0

```
# vifpinit <ESXi host IP>
```

- For ESXi-4.1

```
# vifptarget -s <ESXi host IP>
```

9. Install the driver CD bulletin IDs. Run:

```
# vihostupdate -i --bundle=<path of driver.zip in mount location>
```

10. After host updates successfully, exit the maintenance mode.

1. Right click **ESXi-4.x** in the vSphere Client
2. Choose the option **Exit Maintenance Mode**.

11. Reboot ESXi-4.x.

1. Right-click the ESXi-4.x server
2. Select **Reboot**.



The driver is loaded automatically after the reboot.

➤ *To verify that the driver is loaded, run:*

```
#> vmkload_mod -l
```

➤ *To query network uplinks installed on your machine, run:*

```
#> esxcfg-nics -l
```

The number of uplinks claimed by MLX4_EN driver should be displayed.

The resulting binary modules will be placed under: `/usr/lib/vmware/vmkmod`.

3.2 Remove Driver

➤ *To remove the driver package from the VMware ESX Server machine, run:*

```
#> esxupdate -b <bundle name to remove> remove
```



The `mlx4_en` driver is mandatory for the IB package. If you are using the IB drivers please remove them first.

3.3 Driver Default Values

The driver's default values for the following parameters:

Parameters	Default Value
Both ports	Activated
MSI-X	Enabled
NetQueue	Enabled
RX rings	Default 8, maximum 16
TX rings	8
RX ring size	512
TX ring size	512

Parameters	Default Value
Multifunction Mode	Disabled
TX rate limiting	<ul style="list-style-type: none">• Multifunction Mode – 2500MB/s• Non Multifunction Mode - disabled
RX rate limiting	Disabled

Some of these values can be changed by using module parameters, which can be obtained by running:

```
#> vmkload_mod -s <module name>
```


4 Firmware Programming

Use the mstflint tool to burn a firmware image onto the flash of the NIC device.

➤ *To install the mst tools packet use:*

```
COS# rpm -i mstflint-1.3-1.ofed1.3.1.i386.rpm
```

After installation, the mstflint binary file will be located in the user's path.

1. On the ESX server, log into the service console as root and perform the following steps:

3. Identify the PCI ID of your NIC. Run:

```
COS# lspci
```

4. Examine the output of the NIC you are going to burn. The first field of its entry is the PCI ID of the device.

2. Identify the board ID (PSID). Run:

```
COS# mstflint -d <NIC PCI ID> q
```

3. Retrieve the latest firmware image:

For standard NIC cards:

1. Go to [Firmware Downloads](#) and follow the firmware downloads link.
2. Locate and download the firmware image for the device according to the PSID identified in Step 2.



This release supports firmware rev. fw-25408-2.7.0 and above.

4. Burn the downloaded firmware. Run:

```
COS# mstflint -d <NIC PCI address> -i <FW image> b
```

Example:

```
mstflint -d 0e:00.0 -i fw-25408-2_7_000-MNEH28-XTC_A1.bin b
```

5. Reboot the ESX server

5 HOW TOs

5.1 How to Increase PacketHeapSize

1. Log into the VI Client.
2. Click the **Configuration** tab for the ESX Server host.
3. Select **Advanced Settings**.
4. Select **VMkernel**.
5. Find the corresponding value field for **VMkernel.Boot.netPktHeapMaxSize**.
6. If the value you wish to set is lower than the value displayed in the dialog, set the new value in the MB field (for example: 512)
 1. Click OK to save the changes.
 2. Reboot the ESX server host.



Increasing the packet heap size might reduce overall system performance if the ESX Server is low on physical memory.

5.2 How to Enable MSI-X

MSI-X can be enabled on platforms which support it, resulting in noticeable performance improvement. When MSI-X is enabled, the driver works in the RSS Steering mode (default).

By default, when working in the MSI-X mode, the driver uses one Tx ring and N Rx rings per port (N = number of physical cores).

➤ *To enable MSI-X in the driver, perform the following steps:*

1. Log into the VI Client.
2. Select the server from the inventory panel.

The hardware configuration page for this server should appear.
3. Click on the **Configuration** tab.
4. Choose **Advanced Settings**.
5. Select **VMkernel**.
6. Select **VMkernel.Boot.netNetQueueEnable**
7. Click OK.
8. Log in directly to your ESX server 3 host console.
9. Enable the mtNIC module from the COS. Run:

```
COS# esxcfg-module -e mtNIC
```

10. Enable the MSI-X on the MTNIC module. Run:

```
COS# esxcfg-module -s "msi_x=1" mtNIC
```

11. Reboot the ESX Server 3 host.

➤ *To check if the driver currently uses MSI-X or INTx, perform the following steps:*

1. Run `COS# grep "vmnic<x>-" /proc/vmware/interrupts`
 - If the output is empty, then the INTx is used.
 - If the MSI-X is used, see N+1 lines, where N is the number of cores in the machine.



If your hardware does not support MSI-X, or if the MSI-X is disabled on your ESX server, you will be limited to a single Rx ring (even if the driver is set to use multiple rings). Running with a single Rx ring can reduce performance considerably.

5.3 How to Modify Jumbo Frames

➤ *To modify the MTU of the virtual switch, perform the following steps:*

1. Change the size of `netPktHeapMaxSize` to 512MB.

For further information see the [How to Increase PacketHeapSize](#) section.

2. Change the MTU for the appropriate virtual switch. Run:

```
COS# esxcfg-vswitch -m <new-mtu> <vSwitch-Name>
```

For example if you wish to set the MTU to 9000 on vSwitch1, enter:

```
COS# esxcfg-vswitch -m 9000 vSwitch1
```

The last command sets the MTU for all uplinks on that vSwitch.



It is recommended to set the MTU to the largest size among all the virtual network adapters connected to the vSwitch.