Mellanox MLX4_EN Driver for VMware README

Rev. 1.6.1-2

Last Modified: January 17, 2012

www.mellanox.com

Mellanox Technologies
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.
Beit Mellanox
PO Box 586 Yokneam 20692
Israel
www.mellanox.com
Tel: +972 (0)4 909 7200 ; +972 (0)74 723 7200
Fax: +972 (0)4 959 3245

Mellanox MLX4_EN Driver for VMware README

Table of Contents

List of Tables

1 Overview

This document provides instructions for installing the MLX4_EN drivers for Mellanox Technologies ConnectX®, ConnectX®-2 and ConnectX®-3 based network adapter cards in a VMware ESXi-5.x Server environment. The ConnectX®, ConnectX®-2 and ConnectX®-3 adapters identify on the PCI bus as 25408, 25418, 25448, 26418, 26428, 26438, 26448, 26468, 26478, 26488, 25400, 27500, 27510, 27511, 27520, 27521, 27530, 27531, 27540, 27541, 27550, 27551, 27560, 27561.

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
- Multiple physical functions support
- Rx/Tx traffic Rate Limiting per function
- Flex10 support
- WOL (only on supported hardware)
- NCSI

2 Revision Compatibility

This driver is compatible with:

- Release build ESXi-5.x
- ConnectX®/ConnectX®-2 firmware, version 2.9.1000 and higher
- ConnectX®-3 firmware, version 2.10.0000 and higher

Multifunction mode requires firmware version 2.9.1000 with special configuration. For further information, please contact your Mellanox representative.

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 VBI Driver on ESXi-5.x

1. Log into the VMware ESXi server machine as root.
2. You can either:
   a. Remove any earlier version of the driver from your VMware ESXi server machine prior to installing the new version. Run:

```
#> esxcli software vib list
#> esxcli software vib remove –n net-mlx4-en
```

   b. Install the mlx4_en driver VIB package. Run:

```
#> esxcli software vib install –v <vib_url>
```

   c. Reboot ESXi server (The driver will be loaded automatically).

   OR

   a. Update the driver. Run:

```
#> esxcli software vib update –v <vib_url>
```

   b. Reboot ESXi server (The driver will be loaded automatically).

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

```
#> vmkload_mod -l | grep mlx4_en
```

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

```
#> esxcli network nic list
```

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

In Non Multifunction Mode, port 2 is identified as a pseudo device. Therefore devices are not seen by vSphere when added as uplink.

For further information on how to manipulate the uplink, please refer to Section 5.3, “Adding the Device as an uplink to an Existing Vswitch using the CLI,” on page 10.

3.1.2 Installing and Running the offline_bundle Driver on ESXi-5.x

1. Copy the offline_bundle zip file to ESXi 5.0 machine and extract its contents.
2. You can install the driver in one of the following ways:
   a. Remove any earlier version of the driver from your VMware ESXi server machine prior to installing the new version. Run:

```
#> esxcli software vib list
#> esxcli software vib remove –n net-mlx4-en
```

   b. . Install the mlx4_en driver offline_bundle package. Run:

```
#> esxcli software vib install –d
<path>/mlx4_en-mlnx-1.6.1.2-offline_bundle-471530.zip
```

   c. Reboot ESXi server. (The driver will be loaded automatically).

   OR

   a. Update the driver. Run:

```
#> esxcli software vib update –n net-mlx4-en –d
<path>/mlx4_en-mlnx-1.6.1.2-offline_bundle-471530.zip
```

   b. Reboot ESXi server. (The driver will be loaded automatically).

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

```
#> vmkload_mod -l | grep mlx4_en
```

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

```
#> esxcli network nic list
```

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

In Non Multifunction Mode, port 2 is identified as a pseudo device. Therefore devices are not seen by vSphere when added as uplink.

For further information on how to manipulate the uplink, please refer to Section 5.3, “Adding the Device as an uplink to an Existing Vswitch using the CLI,” on page 10.

3.2 Removing the VIB/offline_bundle Driver

» ***To remove the VIB/offline_bundle driver package from the ESXi server machine, run:***

```
#> esxcli software vib  remove –n net-mlx4-en
```

3.3 Driver Default Values

The driver's default values for the following parameters:

Table 1 - Driver Default Values

Parameters	Default Value
Both ports	Activated
MSI-X	Enabled
NetQueue	Enabled
RX rings	• Multifunction Mode - Default 4, maximum 4 • Non Multifunction Mode - Default 8, maximum 8
TX rings	Default 8, maximum 8
RX rings size	• Multifunction Mode - 256 • Non Multifunction Mode - 512
TX ring size	512
Low Memory	Disabled If enabled, the value of RX/TX ring size and RX ring are changed as follow: • TX ring size: 256 • RX ring size 128 • RX rings 4
Number of LRO Sessions	16
Number of LRO Packets to Aggregate	16
Multifunction Mode	Disabled
Tx rate limiting	• Multifunction Mode - 2500MB/s • Non Multifunction Mode - disabled
Rx rate limiting	Disabled
Port1 Default Function	• Multifunction Mode - 0 • Non Multifunction Mode - disabled
Port2 Default Function	• Multifunction Mode - 1 • Non Multifunction Mode - disabled

`Tx rate limiting`, `Rx rate limiting`, `Port1 Default Function`, and `Port2 Default Function` parameters are enabled only if the Multifunction Mode is enabled.

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

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

For further information, please refer to Section 5.1, “Changing Driver Default Values,” on page 9.

4 Firmware Programming

1. Download the bootable binary image (md5sum: e7b3e9357ca4045fabe2e8a95d951343) from the Mellanox Firmware Tools (MFT) site.
2. Install the image according to the steps described in the README file.

The following procedure requires custom boot image downloading, mounting and booting from a USB device.

5 HOW TOs

5.1 Changing Driver Default Values

The driver's default values can be changed in one of the following methods:

» ***To remove the VIB/offline_bundle driver package from the ESXi server machine, run:***

1. Unload the driver. Run:\

```
#> vmkload_mod -u mlx4_en
```

2. Query the available module parameters. Run:

```
#> vmkload_mod -s mlx4_en
```

3. Load the driver with the required parameters. Run:

```
#> vmkload_mod mod mlx4_en param1=<param value>,param2=<param value>
```

The parameters must be redefined upon every boot.

» ***To permanently change the driver default values :***

1. Query the available module parameters. Run:

```
#> esxcli system module get –m mlx4_en
#> esxcli system module parameters list -m mlx4_en
```

2. Set the driver with the required parameters. Run:

```
#> esxcli system module parameters set –p “param1=<param value>,param2=<param
value>” –m mlx4_en
```

3. Verify that the parameters are set correctly. Run:

```
#> esxcli system module parameters list –m mlx4_en
```

The parameters must be redefined upon every boot.

5.2 Disabling/Enabling Automatic Load of the Driver upon Boot

1. Query the driver auto load status. Run:

```
#> esxcli system module list
```

2. Disable auto load on boot. Run:

```
#> esxcli system module set –enabled=false –m mlx4_en
```

3. Enable auto load on boot. Run:

```
#> esxcli system module set –enabled=true –m mlx4_en
```

The parameters must be redefined upon every boot.

5.3 Adding the Device as an uplink to an Existing Vswitch using the CLI

1. Log into the ESXi server with root permissions.
2. Add an uplink to a vswitch. Run:

```
#> esxcli network vswitch standard uplink add –u <uplink_name> -v
<vswitch_name>
```

Once you add a device via the CLI, it is visible in the vSphere client console, thus removing it can be performed via the UI.

3. Check that the uplink was added successfully. Run:

```
#> esxcli network vswitch standard list –v <vswitch_name>
```

» ***To remove the device locally:***

1. Log into the ESXi server with root permissions.
2. Remove an uplink from a vswitch. Run:

```
#> esxcli network vswitch standard uplink remove –u <uplink_name> -v
<vswitch_name>
```

For additional documents, please refer to the VMware site:

http://pubs.vmware.com/vsphere-50/topic/com.vmware.vcli.migration.doc_50/cos_upgrade_technote.1.9.html#1024629

5.3.1 Locally

1. Log into the ESXi server with root permissions.

2. Add an uplink from a vswitch. Run:

```
#> esxcli network vswitch standard uplink add –u <uplink_name> -v
<vswitch_name>
```

Once you add a device via the CLI, it is visible in the vSphere client console, thus removing it can be performed via the UI.

3. Check that the uplink was added successfully. Run:

```
#> esxcli network vswitch standard list –v <vswitch_name>
```

» ***To remove the device locally:***

1. Log into the ESXi server with root permissions.
2. Remove an uplink from a vswitch. Run:

```
#> esxcli network vswitch standard uplink remove –u <uplink_name> -v
<vswitch_name>
```

For additional documents, please refer to the VMware site:

http://pubs.vmware.com/vsphere-50/topic/com.vmware.vcli.migration.doc_50/cos_upgrade_technote.1.9.html#1024629

5.3.2 Remotely

1. Download and install VMware vSphere Management Assistant (vMA) from:

   http://downloads.vmware.com/d/details/vma50/dHRAYnQld3RiZHAlJQ

2. Use the command “vicfg-vswitch” from the vMA environment.

For additional documents, please refer to the VMware site:

http://pubs.vmware.com/vsphere-50/index.jsp?topic=/com.vmware.vcli.ref.doc_50/vicfg-vswitch.html

5.3.3 Renaming the uplink Name

The uplink naming format is in an increasing order, e.g. when working in either Multifunction Mode or Flex10 "vmnic0" to "vmnic7".

If the order is disrupted/inconsistent and you wish to correct, please follow the procedure below:

1. Log into the ESXi server with root permissions.
2. Open the “vi /etc/vmware/esx.conf” file.
3. Locate the /device/<PCi device>/vmkname = "vmnicX".
4. Change the vmnic numbers to the desired order.

5. Save the file.
6. Reboot the server.

Two vmnics with the same number cannot exist in the same ESXi server.