EMULEX®

OneConnect™ OCm10102-N-X Virtual Fabric Adapter (CFFh) for IBM BladeCenter

Release Notes for OneCommand Manager Version 5.0.31

Introduction

This document contains important information about the Emulex® OneConnect™ OCm10102-N-X Virtual Fabric Adapter (VFA).

This document includes the following sections:

- Overview of the product.
- New features that were added to the product.
- Switches that support the VFA.
- The various operating systems that support the VFA.
- Pointer to the instructions for installing the OneConnect VFA software.
- Instructions for updating the Emulex OneConnect firmware.
- Information on how vNIC function numbers correlate to vNIC IDs on the switch.
- Resolved issues in this release.
- Known issues in this release.

OCm10102-N-X VFA Overview

The OCm10102-N-X VFA provides traditional network interface card (NIC) functionality for a 10-gigabit (Gb) Enhanced Ethernet (EE) fabric. Highlights of this initial release include:

- Support for the Emulex OneCommand Manager application.
- Dual port 10-Gb adapter supporting multiple virtual I/O functions.
- Support for 1-Gb and 10-Gb auto-negotiation.
- Virtual NIC: Physical ports can be divided into a maximum of eight virtual NICs per adapter with bandwidth allocation in 100-MB increments. Each virtual NIC appears as an individual adapter to the operating system.
- Support for Boot from Network via iSCSI software initiator.
- User documentation includes:
    - *SM-CLP to OneCommand Manager Migration Guide:* Describes how to update the driver and management application; you must follow these instructions if you previously installed the OneConnect UCNA Server Software package.
    - *OneCommand Manager User Manual:* Describes how to use the OneCommand Manager management application.
    - *Firmware Update Manual:* Provides a reference to Emulex applications that include firmware update capability.
    - Operating system—specific documentation: May include driver user manuals, quick installation guides, and release notes.

    Note: The driver versions listed in the operating system—specific documentation may not match the driver versions provided for the VFA. If this is the case, any differences in functionality are listed in the New Features section of the *Emulex OneConnect OCm10102-N-X Virtual Fabric Adapter (CFFh) for IBM BladeCenter Release Notes* (this document).

    - *Emulex OneConnect OCm10102-N-X Virtual Fabric Adapter (CFFh) for IBM BladeCenter Release Notes* (this document)*.*
    - *OneCommand NIC Teaming and VLAN Manager User Manual (for Windows only)*: Provides instructions for configuring NIC teaming functionality.

- *Boot Manual*: Provides instructions for UEFI boot.

New Features

- Supports VMware ESX 4.0.
- Supports SLES 10 SP 3.
- Supports Wake on LAN.

Supported Switches

- vNIC mode
  - BNT Virtual Fabric 10G Switch Module
- Two-port NIC mode
  - 10Gb Ethernet Pass-Thru Module for IBM BladeCenter
  - BNT 6-port 10Gb Ethernet Switch Module for IBM BladeCenter
  - Cisco Nexus 4001I Switch Module for IBM BladeCenter
  - BNT Virtual Fabric 10G Switch Module

Note: The Blade Network Technologies switch documentation is available at http://www.bladenetwork.net/. The switch *Application Guide and Command Reference* contains information on configuring the adapter.

Supported Operating Systems and Versions

- Microsoft Windows Server 2003
- Microsoft Windows Server 2008
- Microsoft Windows Server 2008 R2
- Red Hat Enterprise Linux 5 Update 3
- Red Hat Enterprise Linux 5 Update 4
- SUSE Linux Enterprise Server 10 SP2
- SUSE Linux Enterprise Server 10 SP3
- SUSE Linux Enterprise Server 11
- VMware ESX 3.5
- VMware ESX 4.0

Note: Windows and Linux operating systems are supported for x86 and x64 system architectures.

Updating the Firmware

- The *Firmware Update Manual* provides a reference to Emulex applications that include firmware update capability.

- The *Elxflash User Manual* provides information about the Linux ElxflashOffline NIC Kit, which you must use to perform firmware download to the inbox NIC drivers on RedHat operating systems.

- If you previously installed the OneConnect UCNA Server Software package on the VFA, when you update the driver and management application you must first update the firmware manually by using the flash utility as described in the *SM-CLP to OneCommand Manager Migration Guide*. For more information, see the *SM-CLP to OneCommand Manager Migration Guide*.

vNIC Interface Names

When in vNIC mode, the VFA presents eight vNICs to the OS or hypervisor (four for each of the two physical NIC ports). Each vNIC is identified in the OS or hypervisor with a different vNIC function number (0-7). vNIC function numbers correlate to vNIC IDs on the switch as follows:

Table 1: vNIC ID Correlation

Virtual Fabric Adapter		Virtual Fabric Switch Module		
PCIe Function	VFA Port	Chassis Bay	vNIC Pipe	vNIC ID
0	0	HSSM 7	1	INTx.1
2	0	HSSM 7	2	INTx.2
4	0	HSSM 7	3	INTx.3
6	0	HSSM 7	4	INTx.4
1	1	HSSM 9	1	INTx.1
3	1	HSSM 9	2	INTx.2
5	1	HSSM 9	3	INTx.3
7	1	HSSM 9	4	INTx.4

Resolved Issues

This section lists the resolved issues in this release. If you have any questions or require additional information, contact an authorized Emulex technical representative.

1. If the NIC teaming driver is installed, you do not need to reboot the system before running the teaming software.

2. You can now update the NIC Teaming manager to preserve the existing configuration by running the elx_octeamupdate.exe utility after you install the new version of the NIC Teaming manager.

3. In Windows Server 2008 R2, a problem resulting in slow startup when the teaming driver is configured and bound to the LWF driver has been fixed.

4. If teaming is used and an adapter is missing, you no longer need to uninstall the NIC teaming driver to delete the team.

Known Issues

This section lists the known issues associated with this release. If you have any questions or require additional information, contact an authorized Emulex technical representative.

1. Serial over LAN (SOL) currently does not work.

2. VFA MACs do not appear in the Networking menu in the uEFI utility.

3. In the 'System Settings > Network' menu, there is an option for 'PXE Configuration'. Currently, if you use that menu, there are two issues:

    - Only one port of the adapter is displayed instead of both of the physical ports.

    - If a BOFM configuration is applied to the blade in use, the PXE Configuration menu will list the adapter port with its default factory MAC address instead of the BOFM MAC address it is actually using. When PXE booting on this device, the BOFM MAC is actually used, so this is only a display issue.

4. The system will not boot to storage devices when booting from SAN to iSCSI targets over a LAN.

5. An unexpected network congestion can occur when running a Linux system with more than one network interface connected to the same physical network. This behavior is known as Address Resolution Protocol (ARP) flux. To prevent ARP flux, enable arp_filter on all affected interfaces by adding the line net.ipv4.conf.eth$x$.arp_filter=1, where "eth$x$" is a VFA interface number, to the /etc/sysctl.conf file.

    For example, if your VFA interfaces are eth2 through eth9, add the following lines to the /etc/sysctl.conf file:

    net.ipv4.conf.eth2.arp_filter = 1

    net.ipv4.conf.eth3.arp_filter = 1

    net.ipv4.conf.eth4.arp_filter = 1

    net.ipv4.conf.eth5.arp_filter = 1

    net.ipv4.conf.eth6.arp_filter = 1

    net.ipv4.conf.eth7.arp_filter = 1

    net.ipv4.conf.eth8.arp_filter = 1

    net.ipv4.conf.eth9.arp_filter = 1

    Refer to http://linux-ip.net/html/ether-arp.html for more information on how to deal with ARP flux.

6. If the "vNIC Enabled" setting is set to "disabled" and you update the firmware on the VFA, the setting returns to the default "enabled" setting.

7. When modifying the bandwidth allocated to a vNIC on a running VMware system, the VMware driver will improperly report the link speed as 0 until the host has been rebooted.

8. When modifying the bandwidth allocated to a vNIC on a running Linux system, the Linux driver will improperly report the original link speed until the interface is brought down then brought back up.