Emulex Drivers for VMware Release Notes

Versions: ESXi 5.1 driver
FC/FCoE: 10.6.87.0
NIC: 10.6.118.0
iSCSI: 10.6.150.3

ESXi 5.5 driver
FC/FCoE: 10.6.126.0
NIC: 10.6.163.0
iSCSI: 10.6.150.3

ESXi 6.0 driver
FC/FCoE: 10.6.126.0
NIC: 10.6.163.0
iSCSI: 10.6.150.3

Systems: ESXi 5.1, 5.5, and 6.0

Date: July 2015

Purpose and Contact Information

These release notes describe the new features, resolved issues, known issues, and technical tips associated with these Emulex® drivers for VMware releases.

For the latest product documentation, go to www.Emulex.com. If you have questions or require additional information, contact an authorized Emulex technical support representative at tech.support-elx@avagotech.com, 800-854-7112 (US/Canada toll free), +1 714-885-3402 (US/International), or +44 1189-772929 (Europe, Middle East, and Africa).

New Features

1. **Supports OCe14000B-series adapters.**
2. **Discontinued support for the LPe11000-series adapters.**

Resolved Issues

All Versions, Resolved Issues

There are no resolved issues.

ESXi 5.1 Resolved Issues

There are no resolved issues.

ESXi 5.5 Resolved Issues

There are no resolved issues.

Known Issues

All Versions, Known Issues

1. **The VMware driver nodelist depletes when attempting logins to unresponsive nodes.**

   When there are remote nodes on the fabric that are not responsive, the driver correctly fails the login attempt, but does not update the internal state. If the fabric zone is sufficiently large and there are enough unresponsive nodes, this condition can exhaust the node table or cause undesired and constant node swapping in the driver.

   Workaround

   None.

2. **The driver node table for static vPort instances has initiator entries in a logged out state, but with valid World Wide Name (WWN) entries. The driver completes a port login (PLOGI) and process login (PRLI) and then logs out of the remote node.**

   Workaround

   None.

3. **Link Aggregation Control Protocol (LACP) cannot be used on the same port as Fibre Channel over Ethernet (FCoE) or Internet Small Computer System Interface (iSCSI).**

   Workaround

   None.

4. **If there is more than one port group (PG) configured with the same virtual local area network identifier (VLAN ID) on a vSwitch with a 1Gbps port as an uplink, and a virtual network interface card (vNIC) that is connected to one of these PGs is removed (by rebooting the VM to which this vNIC is connected or by other means), connectivity through other vNICs in the PGs with same VLAN IDs is lost.**

   Workaround

   None.

5. **The adapter model name and description do not match the vendor brand name field.**

   The VMware vSphere client only shows adapter family model names with a single port. For example, it shows the OneConnect OCe11102 universal converged network adapter (UCNA) as the OneConnect OCe11100 10 Gigabit Ethernet (GbE), Fibre Channel over Ethernet (FCoE) UCNA. This issue is seen with the native tool lspci as well. The vSphere client gets the adapter model name and description from xml package files installed with the operating system or with an out-of-box kit rather than from the adapter's VPD data. Other native tools work in a similar manner.

   Workaround

   For ESXi 5.1:

   To see the correct model name and description, read the driver's procfs node with this command:

   ```
   cat /proc/scsi/lpfc820/<Instance_Number>
   ```

   or use the Emulex OneCommand® Manager application.

   For ESXi 5.5 and 6.0:

   To show all KV instances available, use the following command:

```
/usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -d
```

For storage adapters, you should look for vmhbaX/Emulex.

6. **On only OCe14000-series adapters, when UMC is enabled and the maximum bandwidth is set using the PXESelect utility, iSCSISelect utility or the OneCommand Manager application, the I/O does not actually adhere to the set value.**

   Workaround:

   None

7. **Possible issues with boot from SAN support on LPe12000-series adapters.**

   Boot from storage area network (SAN) can experience issues on the LPe12000-series adapters. Issues include not finding the boot logical unit number (LUN) or not successfully booting from the boot LUN.

   Workaround

   Boot from SAN support on the LPe12000-series adapters requires the Emulex Universal Boot Code 5.12a2 or later. Emulex recommends using firmware version 2.01a4 or later.

8. **Restriction in assigning data center bridging (DCB) priorities to priority groups.**

   Although there are eight priority groups to which priorities can be assigned, you are able to assign priorities and bandwidths to only two of the priority groups. To one priority group you must assign the FCoE or iSCSI priority and to the other priority group you must assign the other seven (NIC) priorities.

   Note: If you are using a data center bridging exchange (DCBX)-enabled switch to configure the priority groups, configure it for only two priority groups to work correctly with the OneConnect adapter.

   Workaround

   None.

9. **Concurrent firmware upgrades performed on the same ESXi host may not be fully copied.**

   If a firmware upgrade is attempted using the HTTP protocol, the firmware file is copied to the / filesystem of the ESXi system. This file is deleted after the firmware upgrade is complete. Therefore, if concurrent firmware upgrades are performed on the same ESXi host, due to space constraints, the firmware file may not be copied fully, and the firmware upgrades may fail with an error # 40197 (the firmware image is corrupt).

   Workaround

   None.

ESXi 5.1 Known Issues

1. **During link initialization or bootup, the vmkernel.log file may show a series of failed Mailbox commands with a status of 0x44 while the system is running.**

   Workaround

   None required, these messages can be ignored.

2. **If single root I/O virtualization (SR-IOV) is enabled on an OCe14000-series adapter, you cannot create more than 16 VFs per port.**

Workaround

None.

3. **When a transparent VLAN is set as 1, the behavior of the VF attached to the VM is different than expected.**

   When VLAN1 is set as the default VLAN for the VF in the .vmx file, it cannot connect to the interface tagged with vlan-1 on the peer on a Linux system. However, it is pinging with the base interface (untagged interface) on the peer.

   Workaround

   None.

4. **On LPe12000-series adapters, running a long term port shutdown test may result in an operating system error.**

   Workaround

   None.

5. **Issues with SR-IOV support on LPe16000-series adapters.**

   On LPe16000-series adapters, using the VMDirectPath feature on ESXi 5.1 to pass-through the PCI functions to a VM on a Dell R710 or an IBM x3650 M4 server is not supported with SR-IOV.

   Workaround

   This issue is resolved in ESXi 5.5. If you are experiencing this issue, Emulex recommends that you upgrade to ESXi 5.5.

ESXi 5.5 Known Issues

1. **In INT-X mode with UMC enabled, an ESXi NIC native driver reload may take up to 20 minutes to load.**

   Workaround

   None.

2. **On an OCe14000-series adapter, VXLAN performance may require several minutes to reach the line rate with 30 VMs.**

   Workaround

   None.

3. **Firmware netdump does not work properly on OCe11100-series adapters.**

   Workaround

   None.

4. **The ESXi 5.5 NIC driver has limited support for extended SR-IOV support options.**

   ESXi 5.5 has limited support for extended configuration networking options of a virtual function that is assigned to a virtual machine with compatibility for ESXi 5.5 or later.

   ### Supported extended SR-IOV options

   - Enabling Virtual Guest Tagging (VGT)
   - Enabling VLAN Switch Tagging (VST) mode

Extended SR-IOV option limitations

- Cannot change the size of the maximum transmission unit (MTU) (cannot enable jumbo frames)
- Cannot accept or drop incoming frames for a new address with the MAC address change option
- Cannot enable global promiscuous mode for virtual machine network adapters

For more information regarding supported networking configurations see the *Emulex Drivers for VMware ESXi User Manual* and the support section of the VMware website.

5. **The ESXi 5.5 NIC driver can be configured using the vSphere Client.**

   ESXi 5.5 has full support for configuration networking options of a virtual function that is assigned to a virtual machine with compatibility for ESXi 5.1 or later. The *Emulex Driver for VMware ESXi User Manual* provides instructions for default VLAN Tagging, VGT, and configuring a static MAC for a virtual function using the ESXi driver.

   You can also perform these configuration tasks using the vSphere Client. For more information regarding the vSphere Client, see the support section of the VMware website.

6. **1Gbps connectivity is lost after you disable the netqueue feature in the ESXi host.**

   The native driver framework of the vmkernel has an issue in the netqueues disabled path that causes transmit packets to be sent to the elxnet driver with an invalid Tx queue ID. This causes the elxnet driver to drop the Tx packets that are marked with an invalid TxQID. This issue is PR 1122401 tracked by VMware. The issue is resolved in vSphere 2015.

   Workaround

   Keep the netqueues feature of the ESXi host enabled, but load the elxnet driver in legacy interrupt mode. This results in creating a single net queue.

   ```
   esxcli system module parameter set -p "msix=0" -m elxnet
   ```

7. **Disabling MultiQueue (MQ) may seriously impact performance of the LPFC driver.**

   Workaround

   For this release, do not disable MQ.

8. **When the OneCommand Manager CIM Client or VI Client plug-in attempts to connect to an ESXi 5.5 system with an LPe16000-series adapter installed and running firmware version 1.1.35.0, the server never gets displayed in the application.**

   Workaround

   Update the adapter firmware to a newer version.

9. **VPort's NDLP state does not transition to unused after receiving unsolicited PLOGI from an initiator.**

   Workaround

   None.

10. **Errant NULL print string write to console log may occur.**

Workaround

None.

11. **A system hang may occur on servers with more than eight CPU sockets.**

    Workaround

    None. The driver and adapter must be installed on servers with eight CPU sockets or less.

12. **Direct-attached configurations in a point-to-point topology do not connect.**

    Workaround

    Use the driver's default topology (loop).

13. **When using the same CIM Provider versions in customized ESXi images, the VMware upgrade process from version 5.1 to 5.5 or 6.0 fails with the error: "cannot merge VIBs Emulex_bootbank_emulex-cim-provider_<version> with unequal payloads attributes."**

    Workaround

    Use different versions of the CIM Provider in customized images of ESXi 5.1, 5.5, and 6.0.

14. **Firmware netdump fails when NIC I/O is running on the same uplink that is configured for the netdump feature.**

    Workaround

    None.

15. **The physical function (PF) interface receives packets when network traffic is running through the SR-IOV VF interface on LPe16202 and OCe15100 adapters.**

    Workaround

    None.

16. **Elxnet network stats for SR-IOV VFs are not working properly.**

    PFs on Emulex adapters do not support retrieval of VF network statistics from the PF interface.

    Workaround

    None.

17. **On OCe11100-series adapters if you update the driver and firmware, ESXi 5.5 hosts may report large numbers of packet loss and errors in the vmkernel logs.**

    Throughput is not affected, but errors may fill management software logs.

    Workaround

    None.

18. **For the OCe14000-series adapters, driver parameter changes may not take effect if the server is power cycled immediately after the change.**

    Workaround

    Use the "reboot" operation instead of "Power Cycle".

ESXi 6.0 Known Issues

1. **On an OCe14102 adapter with VxLAN, IPV6 Tx traffic goes down when TSO is disabled on both the hypervisor as well as on the VM.**

   Workaround

   Disable TSO on the hypervisor and the VM.

Technical Tips

1. **A max_multiq parameter has been added to the ESXi 5.5 and ESXi 6.0 native mode FC/FCoE drivers. This parameter determines how many IO channel (MSI-X vector/EQ/CQ tuples) are used to parallelize Fibre Channel Protocol (FCP) 10. The disable_mq and fcp_io_channels driver parameters have been removed from the ESXi 5.5 and ESXi 6.0 native mode FC/FCoE drivers. For more information, refer to the *Emulex Drivers Version 10.6 for VMware ESXi User Manual.***

2. **On the OCe14000-series adapters, teaming of physical functions from the same port in Multi-Channel mode may result in lost connectivity. Partitions belonging to the same port must be assigned to different vSwitches (vNetwork Standard Switch [vSwitch/vSS] or VNetwork Distributed Switch [dvSwitch/vDS]).**

   By design, the OCe14000-series adapters do not allow teaming from network interface card (NIC) partitions on the same port. On OCe14000-series adapters configured in Multi-Channel Mode (this includes UMC and Flex Modes), no more than one partition from a port can be assigned to a particular vSwitch/vSS or dvSwitch/vDS).

   The OCe14000-series adapters do not allow this because:

   - Replication of Traffic, results in excess usage of Peripheral Component Interconnect (PCI) Bandwidth and lowers throughput.
   - If a port on an adapter is in a bad state or the link is down, all the partitions of that port also go down. Teaming partitions on the same port do not provide the redundancy expected from teaming.

   Example of a vmkernel.log message that may be seen in ESXi 5.5 and 6.0 (elxnet driver):

   ```
   WARNING: elxnet: elxnet_applyQueueMACFilter:3438: vmnicX: Failed to add mac..

   qidVal:0, status (0x1)

   WARNING: elxnet: elxnet_mccComplProcess:1105: 0000:00N:00.M: Mailbox/MCC

   command opcode 59-1 failed:status 1-22

   Example of a vmkernel.log message that may be seen by in ESXi 5.1 (be2net driver):

   Mailbox command opcode = 0x3b on pci function number 0x4 in subsystem = 0x1

   Failed with status = 0x1, additional status=0x16

   Command to apply MAC address filter failed
   ```

3. **Safe method of changing firmware in flash memory.**

   For operations such as firmware download, Emulex recommends the server be brought into ESXi maintenance mode prior to flashing the firmware.

4. **Installing an offline bundle locally on ESXi 5.1, 5.5, or 6.0 with the esxcli command.**

   Install an offline bundle on a VMware ESXi server locally using the esxcli command.

```
[root@testmachine ~]# esxcli software vib install --maintenance-mode -d
<offline-bundle.zip>
```

Where <*offline-bundle.zip*> is the file name of the offline bundle to be installed.

5. **Throughput performance degrades on block sizes greater than 64KB on Emulex adapters.**

   This issue is fixed in this driver release, but you must administratively set this fix and reboot the system. There is one prerequisite, the ESXi operating system release must be ESXi 5.1 GA - build #799733 or higher:

   To engage this fix, enter the follow commands:

   ```
   esxcli system module parameter set -p "lpfc_dma_boundary" -m lpfc820
   reboot
   ```

6. **The default VMware driver settings may not be appropriate for optimal performance in all scenarios.**

   If performance appears to be lower than expected, there are several driver and system settings that can be modified to improve performance. See the server documentation to determine the correct system settings and the optimal memory and processor configuration.

   Along with proper hardware configuration, some driver settings can be modified to improve performance. The following are some recommended settings to examine while tuning for better performance. See the *Emulex Drivers for VMware ESXi User Manual f*or an explanation of the available settings.

   - Use the "vmxnet 3" adapter type for all guest operating system interfaces.
   - Install VMware Tools for all guest operating systems.

7. **Using the Emulex esxcli vib add–on for the ESXi 5.5 or 6.0 operating system will provide extended statistics and complete VXLAN statistics viewing. The vsish command is limited to 4k size output and is not useful for viewing statistics.**

8. **In ESXi 5.5 or 6.0 NIC teaming scenarios, when Failback is set to no, Emulex recommends having at least one stand-by vmnic for it to work properly.**

9. **For ESXi hosts, the Emulex OneCapture utility should be used to retrieve dumps instead of using the OneCommand Manager application remote dump.**

   For more information on the OneCapture utility, see the *OneCapture User Manual*.