Emulex Drivers for VMware Release Notes

Versions: ESXi 5.0 driver
FC/FCoE: 10.2.459.0
NIC: 10.2.453.0
iSCSI: 10.2.250.0

ESXi 5.1 driver
FC/FCoE: 10.2.459.0
NIC: 10.2.453.0
iSCSI: 10.2.250.0

ESXi 5.5 driver
FC/FCoE: 10.2.455.0
NIC: 10.2.445.0
iSCSI: 10.2.250.0

Systems: ESXi 5.0, 5.1, and 5.5

Date: March 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@emulex.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

- Supports up to four functions per port on 4-port OCe14000-series adapters (when Alternate routing ID interpretation (ARI) is available)
- Supports ExpressLane (also known as FOF priority queueing) for LPe16000-series adapters

Resolved Issues

ESXi 5.0/5.1 Resolved Issues

1. **Port speed and link status show correct information after disabling and re-enabling a port on OCe14000-series adapters with Universal Multi-channel (UMC) enabled.**
2. **For the ESXi 5.0 driver, the port state is correct after a cable pull.**
3. **Added 20G support for the ESXi 5.0 and ESXi 5.1 NIC drivers.**
4. **For the NIC driver, in the case of POST failure, the POST stage functionality is correct and per the required specification.**

5. **For the NIC driver, if asynchronous input/output control (ioctl) fails, status logs and additional status are reported.**
6. **For LPe16000-series adapters, the NIC resource descriptor VFT bit reflects the correct value.**
7. **The NIC driver correctly inserts the VLAN in the following conditions:**
   1. OCe10102 or OCe11102 chip
   2. Not TSO
   3. QnQ or UMC or VST/VGT
   4. Kernel does not request for Checksum.
   Previously the hardware inserted the VLAN and recalculated the packet even when the kernel or driver did not request it.

ESXi 5.1 Resolved Issues

1. **When a virtual machine (VM) is powered off (and to which a virtual function (VF) or physical function (PF) is assigned), network traffic in the VM does not cause the ESXi host to stop responding.**
2. **For the FC/FCoE driver with an ESX5.1 HP v2 CPU Setup, MSIX allocation no longer fails and LUNs no longer disconnect at random times.**
3. **On LPe16000-series adapters, rapidly powering on multiple VMs no longer causes the system to hang when VF are assigned to VMs.**
4. **Implemented the Set Mac List ioctl in embedded format.**
5. **The NIC driver successfully resumes traffic on Windows VM using netperf, while running a load/unload test.**
6. **The NIC driver in multi-channel mode returns a valid number of queues for both the OCe11102 adapter and the OCe14000-series adapter.**
7. **Packet replication does not happen when traffic is unicast and destined to a single, active VM. Packet replication was previously reported on the OCe11102 adapter and the OCe14000-series adapter.**
8. **For the OCe14000-series adapter, the profile number is logged during a driver load.**

ESXi 5.5 Resolved Issues

1. **Resolves an issue with abort processing that could result in a PSOD or slow I/O on LPe16000-series adapters running ESXi 5.5.**
2. **External loopbacks are supported on LPe16202 and OCe15100 adapters.**
3. **The link speed can be changed on OCe10100-series adapters using the esxcli software.**
4. **VFs on LPe16202 and OCe15100 adapters can be added to a VM using the vSphere Web Client.**
   This was a VMware Web Client UI issue that VMware has addressed (PR 1139259).
5. **Excessive "__lpfc_findnode_did" messages no longer appear in the vmkernel.log.**

6. **Tx Ethernet traffic unfairness is no longer observed on OCe14000-series adapters.**
7. **Erroneous "elxnet_mccComplProcess" failure messages no longer appear in the vmkernel.log for OCe14000-series adapters.**
8. **OCe14000-series adapter key value pages no longer show physical port IDs on a VPort page display.**
9. **Using the equivalent command with the Ethtool to modify flow control settings is working properly.**
10. **OneConnect OCe11000-series adapters report correct cable type.**
11. **Enabled global promiscuous mode for virtual functions.**
12. **For an LPe12002 adapter, an EMC Regression driver correctly reestablishes a link after a cable pull.**
13. **LUN paths return correctly.**
14. **For the OCe11102 adapter, the ESX OP Native NIC driver correctly reports the cable type.**
15. **The description of module parameters in the elxnet driver includes the defaults.**
16. **The cable type in the elxnet driver is displayed correctly. The elxnet driver correctly arms the TX CQs in the uplink startio callback**
17. **OPCODE_COMMON_FUNCTION_RESET is supported for the LPe16000-series adapters.**
18. **For the OCe14000-series adapters, an unrecoverable error is no longer observed when running NIC traffic on multiple VMs in ESXi5.5 U2.**
19. **On OCe1010x-series and OCe1110x-series adapters, when the vmkernel requests a VLAN tag insertion and there is not a checksum recalculation for Tx packets, the checksum and the Total Length fields in the IP Header are not recalculated.**
20. **Packets that are looped back in the ESX network stack are trimmed according to their receive packet size.**

Known Issues

Known Issues, All Versions

1. **Link Aggregation Control Protocol (LACP) cannot be used on the same port as FCoE or iSCSI.**

   Workaround

   None.

2. **If there is more than one port group (PG) configured with the same VLAN ID on a vSwitch with a 1Gb/s 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.

3. **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 OCe10102 universal converged network adapter (UCNA) as the OneConnect OCe10100 10GbE, 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 5.0 and 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:

   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.

4. **Storage I/O with VMotion fails because of lost vmkernel 10GB network.**

   Rarely, while running I/O intensive applications on Windows Server 2003 with a 32-bit VM, CPU usage can spike to 100% and Ethernet performance can be slow. This is due to the Transmission Control Protocol/Internet Protocol (TCP/IP) offload-enabled network adapter consuming a great deal of nonpaged pool memory and requesting large blocks of contiguous memory.

   Workaround

   Disabling the TCP Checksum Offload (IPv4) feature improves system and network performance because the large amounts of memory are no longer in use.

   To disable TCP Checksum Offload from your Windows desktop:

   a) Click on **Start > Control Panel > Network Connections > Local Area Connection**.
   b) Click **Properties**.
   c) Click **Configure** (next to vmxnet3 Ethernet Adapter).
   d) Click **Advanced**.
   e) Disable **TCP Checksum Offload (IPv4)**.

5. **Possible issues with boot from SAN support on the LPe12000 family of 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.

6. **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.

7. **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.

8. **An OCe11102-series adapter may freeze if a VLAN header is inserted by the hardware in a tx packet with large IP options, and the IPv4 option size causes the Ethernet + IP + TCP/UDP header sizes to exceed 80 bytes. This typically requires a 26 byte IPv4 option.**

   Workaround

   Use an unused gso field in the wrb header.

9. **On LPe16000-series adapter, power cycling may fail.**

   If multiple VFs are assigned to multiple VMs, when the system is powered on and off the system may not power back on.

   Workaround

   Follow these steps:

   a) Do not enable the

   PROMISCUOUS/MULTICAST_PROMISCUOUS/VLAN_PROMISCUOUS flags during iface creation.

   b) Create all the RQs.

   c) Post one round of buffers on all RQs and DB ring.

   d) Enable PROMISCUOUS/MULTICAST_PROMISCUOUS/VLAN_PROMISCUOUS flags on the iface.

   During the RQ deletion sequence, remove the

   PROMISCUOUS/MULTICAST_PROMISCUOUS/VLAN_PROMISCUOUS flags from the iface and then proceed with RQ deletion.

10. **The NIC driver only logs the firmware version during a driver load.**

    Workaround

    None.

ESXi 5.0 Known Issues

1. **When using the vSphere 5.0 NIC driver with Emulex UCNAs in an HP Flex-10 or IBM Virtual Fabric Adapter (VFA) environment, connectivity may not work properly on Windows virtual machines or on the server when VLANs are configured.**

   Workaround

   Do not use the NIC driver bundled with vSphere 5.0. Obtain an updated driver from Emulex, HP, or IBM that supports HP Flex-10 or IBM VFA systems.

2. **Network timeouts and system lockups have been reported when updating HP firmware and installing a new NIC driver.**

   Issues may occur if you are using the ESXi 5.0 NIC driver with Emulex UCNAs in an HP FlexFabric 10 Gb 2-port 554FLB adapter. If you update firmware from HP via HP SPP 2014.09 and then install a new driver via update manager, network timeouts and system lockups may occur.

   Workaround

   None.

ESXi 5.1 Known Issues

1. **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.

2. **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.

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

   Workaround

   None.

4. **Issues with SR-IOV support on the LPe16000 family of adapters.**

   On the LPe16000 family of 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.

5. **Loopback tests take a longer time to complete on ESXi 5.1 compared with ESXi 5.5.**

   Workaround

   None.

6. **You may encounter low performance on an OCe14000-series adapter with UMC if there is a Windows guest operating system.**

   Workaround

   None.

ESXi 5.5 Known Issues

1. **Correction for the *Emulex Driver for VMware ESXi User Manual*, Installing the NIC Driver heading, Installing the Native Mode Ethernet Driver esxcli Plug–In for ESXi 5.5 topic.**

   In the procedure for Installing the Native Mode Ethernet Driver esxcli Plug-In for ESXi 5.5, remove step 4 and replace it with the following text:

   4. Once the driver is restarted, the elxnet namespace appears under elxnet.

   Supported commands

   - Stats – The stats command is invoked as follows:

     ```
     esxcli elxnet stats get –p <pci dev name>
     ```

     For example:

     ```
     esxcli elxnet stats get –p 0000:003:00.1
     ```

   - Regdump (fatdump) – The regdump command is invoked as follows:

     ```
     esxcli elxnet regdump get –f <absolute file pathname> –p <pci device name>
     ```

     For example:

     ```
     esxcli elxnet regdump get –f /fat.txt –p 0000:02:00.0
     ```

2. **Firmware netdump does not work properly on OCe11000-series adapters.**

   Workaround

   None.

3. **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.

4. **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 5.0 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.

5. **1 Gbs 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
   ```

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

   Workaround

   For this release, do not disable MQ.

7. **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.

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

   Workaround

   None.

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

Workaround

None.

10. **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.

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

    Workaround

    Use the driver's default topology (loop).

12. **When using the same CIM Provider versions in customized ESXi images, the VMware upgrade process from version 5.0 or 5.1 to 5.5 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 5.0, 5.1, and 5.5.

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

    Workaround

    None.

14. **The PF interface receives packets when network traffic is running through the SR-IOV VF interface on LPe16202 and OCe15100 adapters.**

    Workaround

    None.

15. **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.

16. **On OCe11102-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 effected, but errors may fill management software logs.

    Workaround

    None.

17. **On OCe14000-series adapters, you cannot install a signed elxnet and be2iscsi driver using the VMware VUM update feature.**

    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".

19. **The elxnet native driver reports incorrect information after being configured for auto-negotiation.**

    esxcli shows Supports Auto Negotiation as false, even when the driver updates the vmkernel that supports auto-negotiation. exscli also shows Auto neg as false, even when the driver updates the supported modes to have auto-negotiation support.

20. **If jumbo frames are enabled and the NIC is reset by the ESX network stack, the MTU size will not be re-initialized to jumbo frames and all jumbo frames will be dropped.**

    Workaround

    Manually reset the MTU size on the port to 1500 and then back to the correct MTU size.

21. **If a port level VLAN tag insertion is enabled, the number of bytes "inflight" is not being tracked correctly for transmit data. The result is that the NIC could be reset regularly. Traffic will continue to flow after the port is reset.**

    Workaround

    Disable VLAN tag insertion at the port level.

Technical Tips

1. **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.

2. **Installing an offline bundle locally on ESXi 5.0, 5.1, or 5.5 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.

3. **Throughput performance degrades on block sizes greater than 64KB on Emulex UCNAs.**

   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 one of these values or higher:

   - ESXi 5.1 GA - build #799733
   - ESXi 5.0 Patch Release - build #821926

   To engage this fix, enter the follow commands:

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

4. **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.

5. **Using the Emulex esxcli vib add–on for the ESXi 5.5 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.**

6. **In ESXi 5.5 NIC teaming scenarios, when Failback is set to no, it is recommended to have at least one stand-by vmnic for it to work properly.**

7. **The elxnet driver informational logs are categorized into Groups. You can enable or suppress logs of a specific Group by setting or clearing the corresponding bit on DebugMask**.

   The elxnet driver information logs include categories such as DRIVER, UPLINK, QUEUE, SRIOV, MCC, and VxLAN. The DebugMask is a bit-vector (uint32) and each bit represents a Group.

   Log groups reduce informational log clutter at the default log level and aid in debugging by enabling or disabling group logs. Warning and error logs are always logged, regardless of the debug mask. The debug mask is not meant as a replacement of the log level adjustments supported by vsish. vmkernel log levels control the verbosity of logs at each group. For example, level 0 results in less logging and level 4 results in more verbose logs. Debug masks control which groups are enabled or disabled.

   There are two ways to modify the Debug mask, with the ModuleParam debugMask command or the esxcli-Plugin dbgmask command.

   ModuleParam debugMask

   You must reload (or reboot) the driver after modifying the debugMask module parameter. This is a Global debugMask, therefore the same value applies to all Emulex NIC PFs.

   Use either of these commands:

   ```
   esxcfg-module -s "debugMask=0x0120" elxnet
   ```

   ```
   esxcli system module parameters set -p "debugMask=0x0120" -m elxnet
   ```

   esxcli-Plugin

   A driver reload (or reboot) is not required after modifying the dbgmask parameter. This is per-PF debugMask, therefore the same value applies to specific Emulex NIC PF (PF with SBDF is specified with "-p").

   Command to Get:

   ```
   esxcli elxnet dbgmask get -p 0000:03:00.1
   ```

   Command to Set:

```
esxcli elxnet dbgmask set -p 0000:03:00.1 0x000000ff
```

8. **For the ESXi 5.0/5.1 NIC driver, the COMMON_NTWK_GET_LINK_SPEED command fails for an OCe10102 adapter. Use the GET_LINK_STATUS_QUERY instead.**

Note: References to OCe11100 series products also apply to OCe11100R series products.