

Emulex Drivers for Linux Release Notes

| | |
|---|---|
| Versions: Driver for Linux, FC/FCoE Version 11.0.240.0 | (RHEL 7, RHEL 6, SLES 12, SLES 11 SP, OL 7, OL 6) |
| Driver for Linux, NIC Version 11.0.232.0 | (RHEL 7, RHEL 6, SLES 12, SLES 11 SP, OL 7, OL 6, Debian, Ubuntu) |
| Driver for Linux, iSCSI Version 11.0.230.0 | (RHEL 7, RHEL 6, SLES 12, SLES 11 SP, OL 7, OL 6) |
| Driver for Linux, RoCE Version 11.0.232.0 | (See Table 1, Operating System Support for RoCE) |

Date: March 1, 2016

Purpose and Contact Information

These release notes describe the new features, resolved known issues, current known issues, and technical tips associated with this Broadcom[®] Emulex[®]-supported drivers for Linux release.

For the latest product documentation, go to www.broadcom.com. If you have questions or require additional information, contact an authorized Broadcom technical support representative at tech.support-elx@broadcom.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 Emulex LPe31000-series 16Gb/s and LPe32000-series 32Gb/s adapters.
- On Emulex LPe31000-series adapters and LPe32000-series adapters, only Broadcom-approved optics are supported. Refer to the *LPe31000-Series and LPe32000-Series Adapters Installation Guide* for a list of approved optics.
- Supports the Red Hat Enterprise Linux 7.2 (RHEL 7.2) operating system.
- Supports SUSE Linux Enterprise Server 12 SP1 (SLES 12 SP1) operating system.
- Supports Ubuntu 16.04, inbox only.
- For details on operating system support for RoCE, see Table 1.

Note: All listed new features are not supported on RHEL 5.x operating systems.

Caution: Installing OFED replaces any Infiniband stack packages included.

Table 1 Operating System Support for RoCE

| OS | OFED or Native Stack | NFS over RDMA | iSER Initiator | iSER Target | Required Emulex Software |
|-------------|----------------------|---|-------------------|-------------------|---|
| RHEL 6.5 | 3.12-1 | Technical preview | No | No | Download 11.0 firmware, NIC driver and RoCE driver. |
| RHEL 6.6 | Native stack | Technical preview | No | No | Download 11.0 firmware, NIC driver and RoCE driver. |
| RHEL 6.7 | Native stack | Technical preview | No | No | Download 11.0 firmware, NIC driver and RoCE driver. |
| RHEL 7.0 | Native stack | Technical preview | No | No | Download 11.0 firmware, NIC driver and RoCE driver. |
| RHEL 7.1 | Native stack | Client: Production Server: Technical preview | Technical preview | Technical preview | Download 11.0 firmware, NIC driver and RoCE driver. |
| RHEL 7.2 | Native stack | Client: Production Server: Technical preview | Technical preview | Technical preview | Download 11.0 firmware, NIC driver and RoCE driver |
| SLES 11 SP3 | 3.18 | Technical preview | No | No | Download 11.0 firmware, NIC driver and RoCE driver. |
| SLES 11 SP4 | 3.18-1 | Technical preview | No | No | Download 11.0 firmware, NIC driver and RoCE driver. |
| SLES 12 | 3.18 | Technical preview | No | No | Download 11.0 firmware, NIC driver and RoCE driver. |

Resolved Issues

FC/FCoE

1. Issue of domain controller rport structure memory leak is fixed.

Known Issues

1. On Emulex OCE14000-series adapters, firmware versions 11.x and later include new features that require new flash regions to support them.

Firmware versions earlier than 10.0.803.37 did not have the ability to configure the flash regions to support these new features.

Note: Do not use Linux inbox drivers when performing the firmware update.

If you are updating from a firmware version earlier than 10.0.803.37, use one of these methods to update the firmware to 11.x:

- Use the ISO flash tool.
- Use the released 11.x version of the `ethtool` utility. You must perform the firmware update procedure twice to ensure that the flash regions are properly configured, and you must reboot the system after each firmware update.

If you are updating from firmware version 10.0.803.37 or later, use one of these methods to update the firmware to 11.x.:

- Use the ISO flash tool.
- Use the released 11.x version of the `ethtool` utility, and reboot the system after the firmware update. (It is not necessary to perform the firmware update process twice.)

Note: After you have updated the firmware, you must not downgrade the firmware to a version earlier than 10.0.803.37.

2. The host may fail to boot on certain vendor BIOS versions with a single port adapter.

This is most commonly observed when a 1-port adapter is using the RoCE protocol, but it can also occur when a 1-port adapter uses the NIC protocol with SR-IOV disabled.

Affected vendor systems include:

- SuperMicro X9DR3 -F/i
- Cisco USC-240-m3

Workaround

None.

FC/FCoE

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

Workaround

None.

2. PCI Hot Plug may cause applications such as the Emulex OneCommand Manager application or third party applications that use the Emulex libraries (for example, HBAAPI), to misbehave or malfunction.

Workaround

- a) Stop all applications that are accessing LPFC HBA API interface (Emulex OneCommand Manager application or third party applications) before performing PCI Hot Plug of an LPFC adapter.
- b) Use the following command to stop the Emulex OneCommand Manager application:

```
#/usr/sbin/ocmanager/stop_ocmanager
```
- c) After performing PCI Hot Plug of the adapter, you can restart the applications.

3. Deletion of Vports/PCI Hot Unplug.

On occasion the kernel might report SCSI errors when deleting Vports through the sysfs interface or performing a PCI Hot Unplug of a Broadcom Emulex adapter:

```
kernel: Synchronizing SCSI cache for disk
kernel: FAILED
```

Or

```
SCSI error: return code = 0x00010000
```

Workaround

None. These messages do not indicate a functional failure and can be ignored.

4. Deleting Vports while devices are in use.

Broadcom provides management utilities that allow you to delete Vports. However, there is no mechanism for the LPFC driver to detect whether devices accessed through that Vport are in use. This means you can delete a Vport when devices accessible through the Vport are mounted or when I/O is outstanding to the device. When file systems are mounted on Vports and Vports are deleted, the file systems still appear to be mounted; however, they will be inaccessible.

Workaround

Before deleting Vports, you must prepare the system affected by the Vport deletion accordingly, by unmounting all the devices accessible through the Vports, and ensuring there is no outstanding I/O.

5. Support of 4Gb/s adapters in Direct I/O virtualized environments.

Default driver configuration fails to initialize 4Gb/s adapters in virtualized environments that use Direct I/O or SFPT. This may result in the system not responding or an uninitialized LPFC adapter in Intel VT-d and AMD-V IOMMU systems.

Workaround

In these virtualized environments that use Direct I/O or SFPT and 4Gb/s adapters, you must load the LPFC driver with the following driver parameter:

- `lpfc_sli_mode=2`

For example:

```
# modprobe lpfc lpfc_sli_mode=2
```

Note: A consequence of this workaround is that virtual ports are no longer supported by the LPFC driver when the `lpfc_sli_mode` parameter is set to 2.

6. Devloss timeout after swapping ports.

The driver may not finish discovery when two initiator ports are swapped. This causes all devices accessible through one or both of these initiator ports to time out and all I/O to fail.

Workaround

Do one of the following:

- When swapping cables, replace each cable, one at a time, and allow discovery to finish before replacing the next cable. To determine if discovery is finished, read the `state sysfs` parameter.

- When swapping cables, allow devloss timeout to fire before replacing the cables (this fails all outstanding I/O).

7. LILO Boot Loader is not supported on i386 and x86_64 architectures.

The LILO Boot Loader on i386 and x86_64 architectures is not supported for this driver. If the LILO boot loader is used, after the LPFC driver package is installed and upon reboot an incorrect initial ramdisk is used, and the system might not boot correctly.

Workaround

The boot loader supported with this driver is GRUB, which is the default boot loader for most of the Linux distributions. LILO is an older boot loader used on i386 and x86_64 architectures only. GRUB works correctly with the driver package installation script.

8. Suspend to disk command results in a kernel Oops.

If you attempt to suspend to disk using the command:

```
#echo disk > /sys/power/state
```

the LPFC driver encounters a kernel Oops. The sysfs parameter `/sys/power/state` is used to suspend and resume the system. The LPFC driver does not support the suspend to disk and resume command. Do not attempt to use this sysfs parameter when the LPFC driver is loaded.

Workaround

None.

9. Potential error messages during the driver kit removal process.

As part of the driver kit removal process initiated using the `lpfc-install -u` command, the previous in-box LPFC driver version which was saved as part of the current driver kit install process is restored and becomes the active driver. However, the driver kit un-install process does not remove any entries in the Linux distribution configuration file (`modprobe.conf`). As such, parameters that would have been valid for the just-removed driver versions and entered in the `modprobe.conf` file are also used to load the just-restored in-box driver version. This can potentially create problems when

- The just-removed driver version might include module parameters that did not exist in the older just-restored driver version.
- One or more of these module parameters are included in the configuration file (`modprobe.conf`).

If the preceding criteria are met, an error message may be displayed during the uninstallation process of the driver kit, such as:

```
Loading LPFC Driver .FATAL: Error inserting lpfc
(/lib/modules/<kernel_revision>/kernel/drivers/scsi/lpfc/lpfc.ko):
Unknown symbol in module, or unknown parameter (see dmesg)
```

For example, this issue can be observed when uninstalling an 8.2.0.x driver kit, which had DH-CHAP functionality enabled, on a Linux distribution with an older 8.1.10.x in-box driver version.

Workaround

If such an error is seen during the kit removal process, edit the Linux configuration file (`modprobe.conf`) and remove all entries that list LPFC driver module parameters; that is, entries that start with:

```
options lpfc ...
```

Then attempt to uninstall the driver kit again.

Note: To find the module parameters supported by an LPFC driver module, type:

```
# modinfo <driver_dir>/lpfc.ko
```

10. Potential connection loss due to an FCF Failover issue with a Cisco FCoE switch.

An issue was discovered with Cisco Nexus 5000-series FCoE-capable Ethernet switch firmware 4.1(3)N2(1) or earlier, in NPV mode, that may cause the loss or interruption of SCSI connections when used with the Emulex OneConnect UCNAs. The switch incorrectly sends out a Discovery Advertisement to All-ENode-MACs from the FCF MAC with which the FC uplink was down. The end result is that sometimes the UCNA hangs on to an offline FCF or experiences back-to-back FCF failover and it may potentially lead to a Linux SCSI mid-layer devloss timeout.

Workaround

It has been verified that Cisco's IOS 4.2(1)N1(1) release has corrected the issue. You must upgrade your Cisco Nexus 5000-series FCoE switch firmware to 4.2(1)N1(1) or later to avoid this issue. If you decide to use 4.1(3)N2(1) or earlier firmware with your Cisco Nexus 5000-series FCoE switch and this issue is encountered, increase the FC transport `dev_loss_tmo` parameter to 60 seconds. This can be accomplished in one of two ways:

- Update the FC transport `dev_loss_tmo` parameter. For example:

```
# echo 60 > /sys/class/fc_remote_ports/rport-3:0-1/dev_loss_tmo
```

Or

- Update the LPFC driver `lpfc_nodev_tmo` parameter. For example:

```
# echo 60 > /sys/class/scsi_host/host3/lpfc_nodev_tmo
```

11. Potential multipath path loss on the SLES 11 SP3 operating system.

During the controller failback test, multipath cannot recover all paths. For SLES 11 SP3, there is a new `uid_attribute` method to obtain the LUN serial number which is defined in the `multipath.conf` file. This method will call `scsi_id` to send the `sg_io`. If any IOs are returned with a busy status, multipath will not reinstate that path.

Workaround

Edit the `multipath.conf` file and replace

```
uid_attribute 'ID_SERIAL'
```

with

```
getuid_callout '/lib/dev/scsi_id --whitelisted --device=/dev/%n'
```

12. Locked optics support for Emulex LPe31000-series and LPe32000-series adapters.

The adapters perform the following operations:

- Detect and enable both Broadcom or Emulex certified SFP and QSFP optics.
- For firmware revision 11.x unqualified optics are disabled, the link is down and an error message is written to the log file.
- The lpfc out-of-box driver revision 11.x shows this message and the link will not come up.

```
"3176 Port Name [wwpn] Unqualified optics - Replace with Avago optics for Warranty and Technical support"
```

When a 32Gb/s optic is installed in an Emulex LPe31000-series or LPe32000-series adapter, the link supports 32/16/8Gb/s speeds.

When a 16Gb/s optic is installed in an Emulex LPe31000-series or LPe32000-series adapter, the link supports 16/8/4Gb/s speeds.

13. Changed behavior of the static FC and FCoE driver parameter `lpfc_fdmi_on`.

NIC

1. **The physical interface can improperly name `eth0.123` when the `/etc/sysconfig/network-scripts/ifcfg-eth0.123` file contains a HWADDR declaration.**

When using VLAN on RHEL 6.x, the main interface is not created, and the VLAN interface does not actually work correctly as a VLAN.

This is an issue with the udev script in RHEL 6.x. In RHEL 6.x, the udev scripts interpret the HWADDR field in an `ifcfg-ethX.Y` configuration file to mean that the configuration belongs to a real interface. It does not check for the existence of the VLAN field to exclude the field as a real interface. Therefore, you may see unpredictable behavior when including the HWADDR field in the `ifcfg-ethX.Y` file for a VLAN. The setting may, or may not, work. When this setting does not work, it is because the scripts have created or renamed the main interface as `ethX.Y` instead of just `ethX`. Later, when adding the VLAN interface, it fails because `ethX` does not exist.

Workaround

Remove the HWADDR line in `ifcfg-ethX.Y` files that refer to VLAN configurations. You must also clean up the `/etc/udev/rules.d/70-persistent-net.rules` file to remove all the incorrect references to `ethX.Y`.

For more information, refer

https://bugzilla.redhat.com/show_bug.cgi?id=723936.

2. **In certain SKUs of controllers, when SR-IOV is enabled, the NIC priority group (PG) and priority flow control (PFC) are not supported.**

Workaround

None.

- 3. In certain SKUs of controllers, PING is not working when attempting to bridge the 1Gb Ethernet or 10GbE ports to the virtual machines when SR-IOV is enabled for 10GbE ports in the BIOS.**

This issue occurs due to limitations of the virtual Ethernet bridge. All transmitted broadcast packets are looped back by the controller. This affects the functionality of the Linux bridge, as it appears as if the same ARP broadcast packets are received on two different interfaces.

Workaround

- a) Set the aging of the bridge to 0 using the following command:

```
"brctl setageing <bridge> 0"
```

This causes the bridge to behave like a hub and flood the packet to all the ports (except the one on which the packet arrived) every time. This may impact performance. If you have only two interfaces on this bridge (one NIC interface and `virbr0-nic`), there is no performance impact.

- b) Another option may be to use the MacVTap interface to the guest instead of the bridge interface.

- 4. In certain SKUs of controllers, when the driver is loaded with `num_vfs=32`, the initialization of two of the VFs fails and only 30 VF interfaces are created.**

Workaround

None.

- 5. In certain configurations, timeout errors may occur during maximum performance (MPx) socket testing.**

Workaround

For Linux Xen:

- Only configure interfaces created with the *netfront* source model and ignore the second set of interfaces created with *8139cp*

For Linux KVM:

- Use the *virtio* device model instead of *Hypervisor Default*

- 6. When using bonding on top of VLANs on a RHEL-based system, networking appears to stop responding during system boot or when restarting the network.**

On a RHEL-based system, having a bond on top of a VLAN while using `ONBOOT=yes`, can cause the system to stop responding during boot or when restarting the network. The system will appear to not respond while trying to start one of the slave interfaces.

For example, if you have a configuration similar to the following:

```
DEVICE=bond0
BROADCAST=10.255.255.255
GATEWAY=10.0.0.254
IPADDR=10.0.0.14
NETMASK=255.0.0.0
ONBOOT=yes
```



```
BONDING_OPTS="mode=0 miimon=100"  
DEVICE=eth2.2  
BOOTPROTO=none  
ONBOOT=yes  
MASTER=bond0  
SLAVE=yes  
VLAN=yes
```

```
DEVICE=eth3.2  
BOOTPROTO=none  
ONBOOT=yes  
MASTER=bond0  
SLAVE=yes  
VLAN=yes
```

This system hangs because the RHEL networking scripts try to bring up the eth2.2 and eth3.2 interfaces twice. It will bring them up once because the bonding interface specifies them as slave interfaces, and then it tries to bring them up a second time because it sees the `ONBOOT=yes` parameter on a VLAN device.

The issue occurs because after the devices are first brought up, the MAC addresses of the interfaces changes. Usually the first interface MAC address is set on the rest of the interfaces. When the device is brought up a second time, the MAC address of the interface no longer matches the `HWADDR` field. The network scripts will wait for an interface with a matching MAC address to appear, which does not occur.

Workaround

Set the `ONBOOT` parameter to `no` on the slave interfaces. This prevents the system network scripts from trying to bring up the interface a second time. This will not negatively affect the operation of the bonded interface because the `ifcfg-bond0` file still has `ONBOOT` set to `yes`. The networking scripts will attempt to bring up the `bond0` interface on boot, which will cause the slave interfaces to be brought up and configured correctly.

If you are only using bonding without VLAN, the network scripts do not bring up slave interfaces, so the issue is avoided. The issue occurs only when bonding is on top of VLAN interfaces.

7. **Firmware dump using `ethtool -W` is not supported for Emulex OCE14000-series adapters.**

Workaround

None.

8. **When a CNA is configured for a NIC + iSCSI profile, the NIC and iSCSI traffic can be configured to share the total bandwidth. In such a scenario, the NIC traffic will be assured a minimum bandwidth.**

In such a configuration, a VF interface will inherit the minimum bandwidth of the PF, while the Emulex OneCommand Manager application and BIOS components will

display the minimum and maximum bandwidths configured, and the PF will display the total bandwidth available.

9. Low throughput is observed for the VLAN bonded interface on SLES 12.

This issue does not occur in systems with SLES 12 SP1.

Workaround

Obtain the latest kernel patch from SUSE.

10. A call trace in dmesg is observed when running Tx traffic on Emulex LPe16000-series adapters with bonded VLAN interface with SLES 12.

This issue does not occur in systems with SLES 12 SP1.

Workaround

Obtain the latest kernel patch from SUSE.

11. Low throughput is observed with Emulex LPe16000-series adapters, when only one Tx queue is created on the interface.

Workaround

None.

12. In multichannel configurations with 4-port Emulex OCe14000-series adapters, the minimum bandwidth is not working as expected with default number of (7) queues.

Workaround

Reduce the number of queues to 4 using the `ethtool -L eth<X> combined 4` command for minimum bandwidth configured to work as expected.

13. When VxLAN offload is enabled in an Emulex OCe14000-series adapter, the TCP connections through other tunnel interfaces like GRE or IP/IP in the adapter may not work.

This is because checksum offload is not supported for tunnels other than VxLAN. When VxLAN offload is enabled, the driver message `Enabled VxLAN offloads for UDP port <port-num>` is displayed as the output of `dmesg` command.

Workaround

Disable Tx checksum offload in the adapter port on which other tunnel interfaces like GRE/IP/IP are configured.

```
# ethtool -K eth<X> tx off
```

14. In some of the supported operating systems, when SR-IOV is enabled in the adapter, the value of Number of VFs in the PCI config space is same as the Total VFs, although VFs have not been enabled in the driver.

This issue is fixed in RHEL 7.0, SLES 12 and later kernels.

Workaround

None.

15. In Emulex OCe14000-series adapters, throughput does not adhere to the VF Transmit (Tx) rate set in the hypervisor when multiple Tx queues are created on the VF.

Workaround

Decrease the Tx queue count on the VF interface to 1 using the command:

```
ethtool -L eth<X> combined 1
where eth<X> is the VF interface.
```

- 16. On Red Hat 6.5 and 7.1 operating systems, using OCe14000-series adapter with NPar/NPar-EP enabled, you will be able to attach a maximum of two partitions (coming from different bus devices) per adapter to different guests (through PCI-passthru) out of the possible 8/16 partitions.**

This issue arises as vfio only allows functions in different IOMMU groups to be assigned to different guests.

Workaround

Use a RHEL 6.7 or RHEL 7.2 operating system.

- 17. In RHEL 6.5, the command to configure a VLAN interface (vconfig or ip link) on a VF succeeds, although it exceeds the maximum number of VLANs supported on the VF.**

This is because the kernel does not support error return from the driver even when the maximum number of VLANs supported on the VF is exceeded. This issue does not occur in RHEL 6.6 and later kernels.

Workaround

Do not configure more than the maximum number of VLANs supported per VF. The maximum number of VLANs per VF can be obtained by running the following command in the host (or Hypervisor).

```
# dmesg | grep be2net | grep vlans
```

For example:

```
# dmesg | grep be2net | grep vlans
be2net 0000:86:00.0: Max: uc-macs 42, mc-macs 16, vlans 22
be2net 0000:86:04.0: Max: uc-macs 42, mc-macs 16, vlans 21
be2net 0000:86:04.1: Max: uc-macs 42, mc-macs 16, vlans 21
```

In this example the maximum number of VLANs allowed per VF is 21.

Note: The VF count in the example may further need to be reduced by 1 as the Guest operating system could add VLAN 0 by itself without any user command.

The following example shows the steps to confirm if VLAN 0 has been added by the operating system in the VM.

```
# dmesg | grep "VLAN 0" | grep eth0
kernel: 8021q: adding VLAN 0 to HW filter on device eth0"
```

where eth0 is the VF interface.

18. In SUSE Xen kernel with an NPar-EP or SR-IOV configuration, the NIC driver might fail to load.

In such cases, the following messages are displayed in the output of the `dmesg` command.

```
map irq failed
MSIx enable failed
```

This is due to a limitation in the default SUSE Xen kernel to allocate MSIx interrupt vectors.

Workaround

Boot the system with the following kernel boot parameters:

```
extra_guest_irqs=64,2048 nr_irqs=2048
```

iSCSI

- 1. The Open-iSCSI administration utility binds the MAC address of each iSCSI port as an identity to create a configuration database. If a MAC address for an iSCSI port changes, the configuration data will be invalid. MAC addresses can be changed using firmware updates, or by changing the profile or personality of the adapter.**

Non-boot persistent sessions will no longer work. Therefore, all previously mounted partitions will not be found. When `iscsid` attempts to open sessions through the iSCSI port with a changed MAC address, expect to see the following message:

```
beiscsi_ep_connect shost is NULL
```

Workaround

Clean up the saved configuration and then recreate a new interface, discovery, and login.

- 2. When Oracle UEK 6.5 x64 and RHEL 6.x and 5.x systems are used with the Emulex OCE14102-UM adapter with iSCSI/NIC enabled, iSCSI BIOS can fail to boot from the iSCSI LUN and kernel panic can occur.**

Workaround

- a) Boot the system with the `rdshell` option in the kernel parameter. When the system cannot find or mount the boot partition, it will exit to the basic command shell.
- b) Manually run `iscsistart -b`. This will add the boot target to the system.
- c) Type `exit`, and the system will continue to boot normally.

- 3. The system may crash when logging into a target after injecting a PDU error.**

Workaround

None.

- 4. When the IPv6 address is programmed on adapters using the `iscsiadm` tool, it is reflected in the Emulex One Command Manager application as auto-configure and**

greyed out. You cannot configure the routable IP address using the Emulex OneCommand Manager application.

Workaround

Use only one tool either the `iscsiadm` or the Emulex OneCommand Manager application for network configuration of iSCSI adapters.

5. **Appendix A: iSCSI driver in the *Emulex Drivers for Linux User Manual* must be appended to include the following note in the introduction.**

Note: The iSCSI driver parameters are only available in the Emulex provided proprietary iSCSI driver. These are not available in the new Open-iSCSI driver.

RoCE

1. **SLES 11 platforms may prevent loading of unsupported modules by default.**

In such cases, the installer may fail to load the modules even though the RPMs are installed.

Workaround

Manually load the `ocrdma` module after doing one of the following:

- Set `allow_unsupported_modules` to 1 in `/etc/modprobe.d/unsupported-modules`
Or
- Specify `--allow-unsupported` on the command line using the `modprobe` command.

```
# modprobe ocrdma --allow-unsupported
```

2. **If the priority flow control (PFC) (CEE/DCB) on priority 5 cannot be enabled on a switch for RoCE, the connection may break due to packet drops.**

Workaround

Enable generic pause on the switch.

3. **When the RoCE maximum transmission unit (MTU) on an OCe14102 adapter is set to 4096, the MTU on the switch port in use must be set to at least 5000; otherwise, cyclic redundancy check (CRC) errors are indicated in the statistics counter.**

Workaround

None.

4. **On RHEL 6.5 systems with `ofed-3.12-1`, idle NFS RDMA connections may cause the system to stop responding.**

Workaround

None.

5. **On RHEL 6.5, RHEL 6.6 and RHEL 6.7 systems, cable pulls or link down events can cause the system to stop responding.**

Workaround

None.

6. **Unloading the `ocrdma` module using the `rmmmod` command while there are active NFS RDMA mounts may cause the system to stop responding. This behavior is observed on the following platforms - RHEL 6.7, RHEL 7.0, RHEL 7.1, RHEL 7.2, SLES 11 SP3, and SLES 12.**

Workaround

Unmount the existing NFS RDMA mounts, before unloading the `ocrdma` module.

7. **CPU lockup may occur, on some systems using multiple queue pairs, during high traffic tests. This is due to interrupts from different traffic streams being directed to the same CPU core.**

Workaround

- 1) Stop the `irqbalance` service by typing:

```
# service irqbalance stop
```

- 2) Run the `ocrdma_irq_affinity.sh` script, distributed along with the `ocrdma` driver package, by typing:

```
# sh ocrdma_irq_affinity.sh ocrdma<x>
```

where `<x>` is the `ocrdma` device ID.

8. **Current implementations of NFS RDMA do not support IPv6.**

Workaround

None.

The Multiple Queue Pair feature must be enabled only while using NFS RDMA. Enabling this feature on upper layer protocols, other than NFS RDMA, may result in an unpredictable behavior. Workaround

None.

9. **On systems with OFED-3.12-1 or later, installation of out-of-box drivers (`be2net` and `ocrdma`) may fail.**

When an out-of-box driver installation fails, the following message is shown:

```
file /lib/modules/<kernel-version>/updates/be2net/be2net.ko from
install of be2net-kmp-<version_num> conflicts with file from package
compat-rdma-3.12-1.<version_num>
```

Workaround

OFED-3.12-1 includes the inbox `be2net` and inbox `ocrdma` drivers. If out-of-box drivers are desired, use the following procedure to install OFED without the inbox `be2net` and `ocrdma` drivers.

- 1) If OFED is already installed, type the following commands to uninstall OFED and remove the installation directory:

```
# cd OFED-3.12-1
# ./uninstall.pl
# cd ..
# rm -rf OFED-3.12-1
```

- 2) To extract the `OFED-3.12-1.tgz` file, type the following command:

```
#tar -xvzf OFED-3.12-1.tgz
```

- 3) Verify if `asm` symbolic link is present in the `/usr/include` directory. If the symbolic link is not present, create it. Type the following command:

```
#ln -s /usr/include/asm-<xxx>/ /usr/include/asm
```

where `<xxx>` represents the machine architecture. For example, for x86, use `asm-x86`.

- 4) To install OFED without the inbox drivers, type the following command:

```
#./install.pl --without-ocrdma --without-be2net
```

Continue installing the drivers using the instructions in the *Emulex Drivers for Linux User Manual*.

Technical Tips

1. On systems running RHEL 6.5 or later, when upgrading or downgrading from the proprietary `be2iscsi` driver to the open `be2iscsi` driver, you must install the `Dracut-network RPM` package, and pass the `netroot` command parameter into `Dracut` to build the `initramfs` image.

To update the proprietary iSCSI driver to the Open-iSCSI driver for SLES 11 SP3 and later, you must do the following:

- a) Make sure the Open-iSCSI package is installed. If it is not, mount the installation CD and install it.
- b) Edit `/etc/iscsi/initiatorname.iscsi` to add the initiator IQN name. This name can be the same name as in `iSCSIselect` or a new unique one.

Note: The name must conform to the standard IQN.

- c) (Optional step): Change directory to `/boot`, copy the default `initrd` to a new name, and create a boot entry in `/boot/grub/menu.lst` for this new `initrd` name using the same kernel. This step ensures that the system can still boot if something goes wrong.
- d) Install the Broadcom Emulex Open-iSCSI driver RPM file.
- e) If this is not an iSCSI boot system, reboot and follow the Open-iSCSI instructions to log into targets; you do not need to continue following these steps.
If this is an iSCSI boot system that previously used the proprietary iSCSI driver, do not reboot. Instead, continue with the next step.

- f) Change directory to `/boot`, rebuild the `initrd` to include the `iscsi` feature. The output of this command lists the capabilities included in the `initrd`. Make sure `iscsi` is listed. This option includes capabilities such as `iscsiadm` in the `initrd`.

For example: `linux-fiai:/boot # mkinitrd -i initrd-3.0.76-0.11-default -k vmlinuz-3.0.76-0.11-default -f iscsi`

- g) The previous step did not include the iSCSI script to log into the boot iSCSI target. To include it automatically, boot the system using the Broadcom Emulex Open-iSCSI driver. Reboot the system and wait until it fails to mount root and exit into a command shell.
- h) Run the following command from the command shell to log into the boot target:

```
iscsiadm -m fw -l
```

- i) Exit the shell using `exit` for the system to continue to boot with the iSCSI target.
- j) Regenerate the `/boot/initrd` to include `iscsi.sh`. `iscsi.sh` is needed for the initialization script to log into the iSCSI boot target. This is necessary because `mkinitrd` will include `iscsi.sh` in the `initrd` only if there is an entry such as `/sys/firmware/iscsi_boot*`, which can only be created under Open-iSCSI boot.
- k) Reboot the system again. It should boot normally.
- l) If the system has other non-boot iSCSI targets, follow the Open-iSCSI instruction to log into those targets.

2. Red Hat does not officially support saving kdump files to iSCSI targets. In RHEL 6.5 and later, to save kdump files to iSCSI targets:

- a) The system must reserve at least 256 MB for the crash dump. Issue the kernel command with the option `crashkernel=256M`. Reboot the system for this change to take effect.
- b) Copy the `initramfs` image to the `initrd-kdump` image, or build a new `initramfs` image using Dracut and rename it to `initrd-kdump`.
- c) Restart the `kdump` service.

3. Updating OneConnect UCNA firmware in Oracle Linux 5 Unbreakable Enterprise Kernel systems.

The Broadcom Emulex Ethernet driver for Emulex OneConnect adapters supports updating the firmware image in the UCNA flash through the `request_firmware` interface in Linux. For Oracle Linux 5 Unbreakable Enterprise Kernel systems, an updated version of the `ethtool` package `ethtool-6-4.0.1.el5.<arch>.rpm` is required. If this package is not already installed on your system, it is available from the Oracle Linux 5 LATEST channel on the Unbreakable Linux Network (Advisory ELBA-2013-2544 released Aug 29, 2013) and from Oracle public yum at

<http://public-yum.oracle.com/repo/OracleLinux/OL5/latest>

Note: Package updates are not required in Oracle Linux 6 Unbreakable Enterprise Kernel systems.

You can perform this update when the UCNA is online and passing network or storage traffic.

To update the Ethernet firmware image:

- a) Copy the latest firmware image to the `/lib/firmware` directory using the following command:

```
# cp be3flash.ufi /lib/firmware
```
- b) Start the update process using the following command:

```
# ethtool -f eth<X> be3flash.ufi 0
```
- c) Reboot the system to enable the new firmware image.

4. Upgrading the `be2iscsi` driver from proprietary to Open-iSCSI in SLES 11 SP3.

- a) Ensure that the Open-iSCSI package is installed. If it is not, mount the installation CD and install it.

```
linux-oe3a:/media/SLES-11-SP3-DVD-x86_6407031/SuSE/x86_64 # rpm
-ihv open-iscsi-2.0.873-0.21.1.x86_64.rpm
```


- b) Edit `/etc/iSCSI/initiatorname.iscsi` to add the initiator iqn name. This name can be the same name as in `iSCSIselect` or a new unique one.

Note: The name must conform to the standard iqn, and for iSCSI boot, it has to be the same name used in `iSCSIselect`. One method to obtain the current adapter iqn when running the proprietary `be2iscsi` is:

```
cat /sys/class/scsi_host/host4/adapter_iqn/etc/iscsi/initiator-  
name.iscsi
```

For example:

```
InitiatorName=iqn.1996-04.de.suse:01:619c53c296b2
```

- c) Optional step: Change the directory to `/boot`, copy the default `initrd` to a new name, and create a boot entry in `/boot/grub/menu.lst` for this new `initrd` name using the same kernel. This step ensures that the system can still be rebooted if something unexpected occurs.
- d) Install the `be2iscsi` Open-iSCSI driver rpm.

```
linux-oe3a:/tmp # rpm -Uhv  
be2iscsi-kmp-default-10.0.713.0_3.0.76_0.11-1.sles11sp3.x86_64.rpm
```

- e) If this is a non-iSCSI boot system, reboot and follow the Open-iSCSI instructions to log in to targets. If this is an iSCSI boot system that previously used the proprietary `be2iscsi`, do not reboot, continue to the next step.
- f) Change directory to `/boot`, and rebuild the `initrd` to include the iSCSI feature. The output of this command will list the features included in the `initrd`.

Note: Make sure iSCSI is listed. This feature will include `iscsiadm` in the `initrd`.

```
linux-oe3a:/boot # mkinitrd -i initrd-3.0.76-0.11-default -k  
vmlinuz-3.0.76-0.11-default -f iscsi
```

Features:

- `acpi`
 - `dm`
 - `multipathd`
 - `kms`
 - `block`
 - `usb`
 - `network`
 - `multipath`
 - `kpartx`
 - `resume.userspace`
 - `resume.kernel`
 - `kdump`
- g) To automatically include the iSCSI script to login to the boot iSCSI target, boot the system using `open be2iscsi`. Reboot the system and wait until it fails to mount root and exits into a command shell.

- h) Run the following commands from the command shell to login to the boot target:

```
iscsid -n  
iscsiadm -m fw -l
```

- i) Type `exit` to exit the shell and the system will continue to boot with the iSCSI target.
j) After the system has booted up with the open `be2iscsi` driver using the workaround, the `/boot/initrd` must be regenerated to include `iscsi.sh`.

Note: `iscsi.sh` is needed for the init script to login to the iSCSI boot target. This workaround is required because `mkinitrd` will only include `iscsi.sh` in the `initrd` if there is an entry such as `/sys/firmware/iscsi_boot*`, which can only be created if Open-iSCSI boot is enabled. In addition, `-f multipath` must be added to the `mkinitrd` command if `multipath` is being used.

```
linux-oe3a:/boot # mkinitrd -i initrd-3.0.76-0.11-default -k  
vmlinuz-3.0.76-0.11-default -f multipath
```

Features:

- `acpi`
- `dm`
- `multipathd`
- `kms`
- `block`
- `usb`
- `network`
- `iSCSI`
- `multipath`
- `kpartx`
- `resume.userspace`
- `resume.kernel`

- k) Reboot the system again; it should boot normally.
l) If the system has other non-boot iSCSI targets, follow the Open-iSCSI instructions to log in to those targets.

5. Additional physical NICs added to the XenServer do not show up in the XenCenter.

Workaround

See the following link for instructions on adding an additional physical NIC to the XenServer.

<http://support.citrix.com/article/CTX121615/>

6. The Open-iSCSI driver persistent targets are maintained in the Host Nodes database. The driver recognizes only the boot-target that is persistent, and is not aware of any non-boot persistent targets on the adapter.

Workaround

Do not use `iSCSISelect` to configure persistent non-boot targets for all variants of RHEL 6, RHEL 7, and SLES 11 releases.

7. **When the SLES 11u1x32 or SLES 11u2x32 operating systems are used with the Emulex OCE11100-series adapter with the open-iSCSI driver, the `OEMSEAaddSendTargetPortal` command does not work.**

Workaround

When installing a new SLES 11u1x32 or SLES 11u2x32 operating system, ensure that the open-iSCSI rpm package that came on the installation CD is used during the installation. Other rpm packages may not work.

8. **Set PFC as the default mode when possible, while using RoCE.**

If CEE or DCB mode is enabled on the switch, enable PFC for priority 5 (configure RoCE traffic on priority 5).

In addition, VLAN interfaces must be configured and used for RoCE traffic in order for PFC to work correctly. VLANs greater than 1 must be used for best interoperability.

If CEE or DCB is not enabled, generic pause must be enabled on the switch and the NIC port to prevent loss of packets.

9. **On NFS servers where `fsid` needs to be specified while exporting file systems, ensure that the `fsid` is unique for each exported file system.**
10. **When the RoCE maximum transmission unit (MTU) on Emulex OCE14102 adapter is set to 4096, the MTU on the switch port in use must be set to at least 5000; otherwise, cyclic redundancy check (CRC) errors will be reported.**
11. **When using the RoCE profile on Emulex OCE14000-series adapters with PFC settings enabled, the `8021q.ko` kernel module must be loaded, if not already loaded by default.**

This module must be loaded before the Broadcom Emulex drivers are loaded using the following command.

```
# modprobe 8021q
```

12. **For better performance of NFS over RDMA, enable multiple queue pair feature on the RoCE device using the following command:**

```
#echo 1 > /sys/class/infiniband/ocrdmaX/qp_expansion
```

13. **When unloading the `ocrdma` driver, an error message `module ocrdma is in use` may appear. Linux does not allow unloading of modules when they are in use. In such cases, terminate the application(s) that may be using `ocrdma` and try again.**

14. **Use the RoCE-2 NIC profile with NFS-RDMA.**

For information on setting the profile, refer to Section *Basic Configuration* in the most recent *Emulex Drivers for Linux User Manual* (available on the Broadcom website). When you successfully update the profile, the `Active-Profile-ID` must display `0x15`.

15. **The RoCE profile for Linux iSER, NFS, or VM Migration is RoCE-2.**

16. **Refer to ofed-3.18-1 release notes for ofed issues.**

17. On Emulex OCe14000-series adapters, the maximum number of possible mounts for the default value of `rdma_slot_table_entries` on the host is shown in the following table.

| OS distributions | 1 port card OCe14101B | 2 port card OCe14102B | 4 port card OCe14104B |
|---|--------------------------|--------------------------|--------------------------|
| RHEL 6.5 | 5 | 2 | 1 |
| RHEL 6.6, RHEL 6.7 | 38 | 19 | 9 |
| RHEL 7.0, RHEL 7.1, RHEL 7.2, SLES 11 SP3, SLES 11 SP4, SLES 12 | 5 | 2 | 1 |

For information on how to further increase the maximum number of possible mounts, refer to the *Emulex Drivers for Linux User Manual* (Section, NFS over RDMA).

For each operating system in the table use the operating system with the RoCE stack as specified in Table 1, Operating System Support for RoCE.