Emulex® Drivers for Linux Release Notes

Versions:	FC/FCoE Version 11.1.218.0-1	Red Hat Enterprise Linux (RHEL) 7, RHEL 6, SUSE Enterprise Linux (SLES) 12, SLES 11 SP, Oracle Linux (OL) 7, OL 6
	NIC Version1 1.1.207.0-2	RHEL 7, RHEL 6, SLES 12, SLES 11 SP, OL 7, OL 6, Debian, Ubuntu
	iSCSI Version 11.1.207.0-2	RHEL 7, RHEL 6, SLES 12, SLES 11 SP, OL 7, OL 6
	RoCE Version11.1.207.0-2	See Table 1, Operating System Support for RoCE
Date:	October 14, 2016	

Purpose and Contact Information

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

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 ecd-tech.support@broadcom.com, 800-854-7112 (US/Canada toll free), +1 714-885-3402 (International), or +49 8941352 0244 (Europe, Middle East, and Africa; UK business hours only 8:30 a.m. to 5:00 p.m. UTC Monday through Friday).

New Features

- This release adds support for the following operating systems:
  - RHEL 7.3 (out-of-box)
  - RHEL 6.8 (out-of-box)
  - SLES 12 SP2
- For details on operating system support for RDMA over converged Ethernet (RoCE), see Table 1.

Note: Installing OpenFabrics Enterprise Distribution (OFED) replaces any Infiniband stack packages included.

Table 1 Operating System Support for RoCE

Operating SystemS	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 11.1 firmware, NIC driver, and RoCE driver.
RHEL 6.6	Native stack	Technical preview	No	No	Download 11.1 firmware, NIC driver, and RoCE driver.
RHEL 6.7	Native stack	Technical preview	No	No	Download 11.1 firmware, NIC driver, and RoCE driver.
RHEL 6.7	Native stack	Technical preview	No	No	Download 11.1 firmware, NIC driver, and RoCE driver.
RHEL 7.0	Native stack	Technical preview	No	No	Download 11.1 firmware, NIC driver, and RoCE driver.
RHEL 7.1	Native stack	Client: Production Server: Technical preview	Technical preview	Technical preview	Download 11.1 firmware, NIC driver, and RoCE driver.
RHEL 7.2	Native stack	Client: Production Server: Technical preview	Technical preview	Technical preview	Download 11.1 firmware, NIC driver, and RoCE driver
SLES 11 SP3	3.18	Technical preview	No	No	Download 11.1 firmware, NIC driver, and RoCE driver.
SLES 11 SP4	3.18-1	Technical preview	No	No	Download 11.1 firmware, NIC driver, and RoCE driver.
SLES 12	3.18	Technical preview	No	No	Download 11.1 firmware, NIC driver, and RoCE driver.
SLES 12 SP2	3.18	Technical preview	No	No	Download 11.1 firmware, NIC driver, and RoCE driver.

iSCSI Extensions for RDMA (iSER)

Resolved Issues

1. **On RHEL 6.5 systems with OFED-3.12-1, idle Network File System (NFS) remote direct memory access (RDMA) connections no longer cause the system to stop responding.**

Known Issues

1. **Certain inbox drivers do not support LPe31000-series and LPe32000-series adapters, as described in the following table.**

Table 2 Operating System and Support for LPe31000-Series and LPe32000-Series Adapters

Operating System	Support for LPe31000-Series and LPe32000-Series Adapters
RHEL 7.2	No
RHEL 7.3	Yes
RHEL 6.8	Yes
RHEL 6.9	Yes
SLES 11 SP3	No; use the driver disk
SLES 11 SP4	No; use the driver disk
SLES 12 SP1	No; use the driver disk
SLES 12 SP2	Yes
UEK 3 R7	Yes
UEK 4 (OL 6.8 and 7.3)	Yes
Ubuntu 15.04	No
Ubuntu 16.04	Yes
XenServer 6.5	No; use the driver disk
XenServer 7.0	Yes

Workaround

Follow the guidelines in the table.

2. **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.

3. **The host might fail to boot on certain vendor basic input/output system (BIOS) versions with a single port adapter.**

   This issue is most commonly observed when a one-port adapter is using the RoCE protocol, but it can also occur when a one-port adapter uses the Network Interface Controller (NIC) protocol with single root I/O virtualization (SR-IOV) disabled.

   The following vendor systems are affected:

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

   Workaround

   None.

FC/FCoE

1. **On RHEL systems with LPe31000-series or LPe32000-series adapters installed, setting the rq_affinity value to 2 results in a central processing unit (CPU) lockup.**

   Workaround

   None.

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

   Workaround

   None.

3. **Peripheral Component Interconnect (PCI) Hot Plug might cause applications, such as the Emulex OneCommand® Manager application or third-party applications that use the Emulex libraries (for example, a host bus adapter [HBA] application programming interface [API], to malfunction.**

   Workaround

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

      ```
      #/usr/sbin/ocmanager/stop_ocmanager
      ```

   3) After performing PCI Hot Plug of the adapter, restart the applications.

4. **Small Computer System Interface (SCSI) errors might occur on deletion of virtual ports (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. Ignore these messages; they do not indicate a functional failure.

5. **An issue exists while deleting Vports when devices are in use.**

   Broadcom provides management utilities that allow you to delete Vports. However, there is no mechanism exists for the LPFC driver to detect whether devices accessed through that Vport are in use. This situation means you can delete a Vport when devices accessible through the Vport are mounted or when input/output (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.

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

   The default driver configuration fails to initialize 4Gb/s adapters in virtualized environments that use Direct I/O or Secure File Transfer Protocol (SFPT). This might 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 the LPFC driver does not support virtual ports when the `lpfc_sli_mode` parameter is set to 2.

7. **Devloss timeout occurs after swapping ports.**

   The driver might not finish discovery when two initiator ports are swapped. This situation 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 the devloss timeout to occur before replacing the cables (this action fails all outstanding I/O).

8. **The LInux LOader (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 GRand Unified Bootloader (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.

9. **The 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` suspends and resumes 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.

10. **Potential error messages are displayed during the driver kit removal process.**

    As part of the driver kit removal process initiated using the `lpfc-install -u` command, the previous inbox LPFC driver version that was saved as part of the current driver kit installation process is restored and becomes the active driver. However, the driver kit uninstallation 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 also are used to load the just-restored inbox driver version. This situation can potentially create problems when the following conditions occur:

    - 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 might 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 Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) functionality enabled, on a Linux distribution with an older 8.1.10.x inbox 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 the following:

```
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
```

11. **A potential connection loss occurs due to a Fibre Channel over Ethernet Forwarder (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 N_Port Virtualization (NPV) mode, that might cause the loss or interruption of SCSI connections when used with the Emulex OneConnect Universal Converged Network Adapter (UCNA). The switch incorrectly sends out a Discovery Advertisement to All-ENode-MACs from the FCF Media Access Control (MAC) with which the Fibre Cahnnel (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 Fibre Channel over Ethernet (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 action 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
      ```

12. **A potential multipath path loss occurs on the SLES 11 SP3 operating system.**

    During the controller failback test, multipath cannot recover all paths. For SLES 11 SP3, a new `uid_attribute` method obtains the LUN serial number that is defined in the `multipath.conf` file. This method calls `scsi_id` to send the `sg_io`. If any I/Os are returned with a busy status, multipath does 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'
    ```

13. **The behavior of the static FC and FCoE driver parameter `lpfc_fdmi_on` is changed.**

14. **When OCe14000-series and LPe31000-series or LPe32000-series adapters coexist, if you attempt to install the RHEL 7.2 out-of-box (OOB) LPFC driver on a boot from storage area network (SAN) device, the installation fails.**

    Workaround

    To install RHEL 7.2 OOB LPFC on a server on which OCe14000-series and LPe31000-series or LPe32000-series adapters coexist, perform the following steps:

    1) Boot the RHEL 7.2 DVD, select **Install Red Hat Linux Enterprise 7.2**, and press the **Tab** key.
    2) Add `modprobe.blacklist=lpfc linux dd` to the end of the line `>vmlinuz initrd=`, and press **Enter**.

       ```
       >vmlinuz initrd=modprobe.blacklist=lpfc linux dd
       ```

       Note: `modprobe.blacklist=lpfc` indicates to the operating system not to load the inbox LPFC driver, and `linux dd` prompts you to install the out-of-box LPFC driver.

    3) Complete the RHEL 7.2 operating system installation using the Red Hat 7.2 driver update disk (DUD).

NIC

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

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

   This issue is 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 an actual interface. It does not check for the existence of the VLAN field to exclude the field as an actual interface. Therefore, you might see unpredictable behavior when including the HWADDR field in the `ifcfg-ethX.Y` file for a VLAN. The setting might 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 to:
   `https://bugzilla.redhat.com/show_bug.cgi?id=723936`.

2. **In certain stock keeping unit s (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 does not work 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 situation affects the functionality of the Linux bridge, as it appears as if the same address resolution protocol (ARP) broadcast packets are received on two different interfaces.

   Workaround

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

     `"brctl setageing <bridge> 0"`

     This command 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 situation might impact performance. If you have only two interfaces on this bridge (one NIC interface and virbr0-nic), there is no performance impact.
   - Another option might be to use the MacVTap interface for the guest instead of the bridge interface.

4. **In certain configurations, timeout errors might occur during maximum performance (MPx) socket testing.**

Workaround

For Linux XenServer:

- Configure interfaces only created with the *netfront* source model, and ignore the second set of interfaces created with *8139cp*.

For Linux K Virtual Machine (KVM):

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

5. **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 an 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 appears unresponsive 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 brings 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 change. 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 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 action prevents the system network scripts from trying to bring up the interface a second time. This action does not negatively affect the operation of the bonded interface because the `ifcfg-bond0` file still has `ONBOOT` set to `yes`. The networking scripts attempt to bring up the bond0 interface on boot, which causes the slave interfaces to be brought up and configured correctly.

If you only use 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.

6. **Firmware dump using `ethtool -W` is not supported for Emulex OCe14000-series adapters.**

   Workaround

   None.

7. **When a converged network adapter (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 is assured a minimum bandwidth.**

   In such a configuration, a VF interface inherits the minimum bandwidth of the physical function (PF). The Emulex OneCommand Manager application shows the minimum and maximum bandwidths configured, and the PF displays the total bandwidth available.

8. **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.

9. **A call trace in dmesg is observed when running transmit (Tx) traffic on Emulex LPe16000-series adapters with a 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.

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

    Workaround

    None.

11. **In multichannel configurations with four-port Emulex OCe14000-series adapters, the minimum bandwidth is not working as expected with default number of queues (seven).**

    Workaround

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

12. **When Virtual extensible LAN (VXLAN) offload is enabled in an Emulex OCe14000-series adapter, the Transmission Control Protocol (TCP) connections through other tunnel interfaces, such as Generic Routing Encapsulation (GRE) or IP-in-IP (IPIP), in the adapter might not work.**

    This situation occurs 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, such as GRE or IPIP, are configured.

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

13. **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.

14. **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.

15. **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 situation occurs because the kernel does not support an 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. Obtain the maximum number of VLANs per VF 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 might further need to be reduced by 1 because 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.

16. **On OCe14000-series 1Gb/s adapters, traffic flows only on single connections when traffic is run using the iperf utility with multiple connections.**

    Workaround

    Create a 1TX queue using the `ethtool -L` command, and traffic will flow on all TCP connections.

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 is invalid. MAC addresses can be changed using firmware updates, or by changing the profile or personality of the adapter.**

   Nonboot persistent sessions no longer work. Therefore, all previously mounted partitions are not found. When `iscsid` attempts to open sessions through the iSCSI port with a changed MAC address, the following message appears:

   ```
   beiscsi_ep_connect shost is NULL
   ```

   Workaround

   Clean up the saved configuration and then re-create a new interface, discovery, and login.

2. **When RHEL 6.x systems are used with Emulex OCe14102-UM adapters with iSCSI/NIC enabled, iSCSI BIOS might fail to boot from the iSCSI LUN and kernel panic might occur.**

   Workaround

   1) Boot the system with the `rdshell` option in the kernel parameter. When the system cannot find or mount the boot partition, it exits to the basic command shell.
   2) Manually run `iscsistart -b`. This action adds the boot target to the system.
   3) Type `exit`, and the system continues to boot normally.

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

4. **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 parameters are not available in the new Open-iSCSI driver.

RoCE

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

   In such cases, the installer might fail to load the modules even though the Red Hat Package Managers (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) Converged Enhanced Ethernet (CEE) Data Center Bridging (DCB) on priority 5 cannot be enabled on a switch for RoCE, the connection might break due to packet drops.**

   Workaround

   Enable generic pause on the switch.

3. **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.

4. **Unloading the `ocrdma` module using the `rmmod` command while there are active NFS RDMA mounts might 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.

5. **A CPU lockup might occur on some systems using multiple queue pairs during high traffic tests. This situation 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.

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

Workaround

None.

7. **The Multiple Queue Pair feature must be enabled only while using NFS RDMA. Enabling this feature on upper layer protocols other than NFS RDMA might result in unpredictable behavior.**

Workaround

None.

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

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

```
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 the `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, perform these steps:

   1) Make sure the Open-iSCSI package is installed. If it is not, mount the installation CD and install it.
   2) Edit the `/etc/iscsi/initiatorname.iscsi` file to add the initiator iSCSI Qualified Name (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.

   3) (Optional step): Change the directory to `/boot`, copy the default initrd to a new name, and create a boot entry in the `/boot/grub/menu.lst` file for this new initrd name using the same kernel. This step ensures that the system can still boot if something goes wrong.
   4) Install the Broadcom Emulex Open-iSCSI driver RPM file.
   5) 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.
   6) Change the directory to `/boot`, and 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`
   7) 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 exits into a command shell.

8) Run the following command from the command shell to log into the boot target:

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

9) Exit the shell using `exit` for the system to continue to boot with the iSCSI target.

10) Regenerate the `/boot/initrd` to include `iscsi.sh`. `iscsi.sh`, which is required for the initialization script to log into the iSCSI boot target. This step is necessary because `mkinitrd` includes `iscsi.sh` in the initrd only if an entry, such as `/sys/firmware/iscsi_boot*`, exists; which is only created under Open-iSCSI boot.

11) Reboot the system again. It should boot normally.

12) If the system has other non-boot iSCSI targets, follow the Open-iSCSI instruction to log into those targets.

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

1) 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.

2) Copy the `initramfs` image to the `initrd-kdump` image, or build a new `initramfs` image using Dracut, and rename it to `initrd-kdump`.

3) Restart the `kdump` service.

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

1) 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
```

2) Edit the `/etc/iSCSI/initiatorname.iscsi` file 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 must be the same name used in iSCSIselect. One method to obtain the current adapter IQN when running the proprietary be2iscsi is as follows:

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

For example:

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

3) 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.

4) 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
```

5) 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, and continue to the next step.

6) Change the directory to `/boot`, and rebuild the initrd to include the iSCSI feature. The output of this command lists the features included in the initrd.

   Note: Make sure iSCSI is listed. This feature includes `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`

7) To automatically include the iSCSI script to log in 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.

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

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

9) Type `exit` to exit the shell, and the system will continue to boot with the iSCSI target.

10) After the system has booted up with the open be2iscsi driver using the workaround, regenerate the `/boot/initrd` to include `iscsi.sh`.

    Note: `iscsi.sh` is required for the init script to log in to the iSCSI boot target. This workaround is required because `mkinitrd` only includes `iscsi.sh` in the initrd if there is an entry, such as `/sys/firmware/iscsi_boot*`, which can only be created only if Open-iSCSI boot is enabled. In addition, add `-f multipath` 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`

11) Reboot the system again; it should boot normally.

12) If the system has other non-boot iSCSI targets, follow the Open-iSCSI instructions to log in to those targets.

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

   Workaround

   The following link has instructions for adding an additional physical NIC to the XenServer.

   `http://support.citrix.com/article/CTX121615/`

5. **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.

6. **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 might not work.

7. **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 for PFC to work correctly. Use VLANs greater than 1 for best interoperability.

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

8. **On NFS servers where `fsid` must be specified while exporting file systems, ensure that the `fsid` is unique for each exported file system.**

9. **When the RoCE maximum transmission unit (MTU) on Emulex OCe1410x-series adapters and Emulex OCe1440x-series adapters 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 reported.**

10. **When using the RoCE profile on Emulex OCe14000-series adapters with PFC settings enabled, the `8021q.ko kernel` module must be loaded, if it is not already loaded by default.**

    Load this module before the Broadcom Emulex drivers are loaded using the following command.

```
# modprobe 8021q
```

11. **On Emulex OCe1440x-series adapters for better performance of NFS over RDMA, enable the multiple queue pair feature on the RoCE device using the following command:**

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

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

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

    For information on setting the profile, refer to the Basic Configuration section 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`.

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

15. **Refer to the OFED-3.18-1 release notes for OFED issues.**

16. **Locked optics are supported on Emulex LPe31000-series and LPe32000-series adapters.**

    The adapters perform the following operations:

    - Detect and enable both Broadcom or Emulex certified small form-factor pluggable (SFP) and quad small form-factor pluggable (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 32Gb/s, 16Gb/s, and 8Gb/s speeds.

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

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 (*the NFS over RDMA section).

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.