

Emulex Driver for Linux Release Notes

Versions: Driver for Linux, FC/FCoE Version 10.6.144.21 (RHEL 6, RHEL 7, SLES, OL 6, OL 7)
Driver for Linux, FC/FCoE Version 10.4.255.16 (RHEL 5, OL 5)
Driver for Linux, NIC Version 10.6.144.21 (RHEL 6, RHEL 7, SLES, OL 6, OL 7,
Debian, Ubuntu)
Driver for Linux, NIC Version 10.4.255.20 (RHEL 5, OL 5)
Driver for Linux, Open-iSCSI Version 10.6.144.21 (RHEL 7, RHEL 6, SLES 12, SLES
11, OL)
Driver for Linux, Proprietary iSCSI Version 10.0.719.1009 (RHEL 5, OL 5)
Driver for Linux, RoCE Version 10.6.144.21 (See Table 1 on page 2)

Date: July 2015

Purpose and Contact Information

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

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

New Features

- Supports OCe14000B-series adapters.
- Supports the Red Hat Enterprise Linux 7.1 (RHEL 7.1) operating system.
- Enhanced performance of FCoE, NIC, iSCSI drivers on OCe14000-series adapters.
- Discontinued support for LPe11000-series and OCe10100-series adapters.
- Discontinued support for RHEL 6.4, SLES 11 SP1, SLES 11 SP1e, and SLES 11 SP2.
- Support for receive-side scaling (RSS) is enabled for virtual functions (VFs) in OCe14000-series adapters, when SR-IOV is enabled and requested number of VFs are less than maximum supported.
- Supports enabling or disabling number of VFs using sysfs.
- Fully supports RDMA over Converged Ethernet (RoCE) on SUSE Linux Enterprise Server 12 (SLES 12) operating system with OFED-3.12-1.
- Enhanced NFS RDMA performance with multiple queue pairs.
- Refer to Table 1 on page 2 for details on operating system support for RoCE.

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

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 10.6 NIC and RoCE driver.
RHEL 6.6	Native stack	Technical preview	No	No	Download 10.6 NIC driver. Use inbox RoCE driver.
RHEL 6.7	Native stack	Technical preview	No	No	Download 10.6 NIC driver. Use inbox RoCE driver.
RHEL 7.0	Native stack	Technical preview	No	No	Download 10.6 NIC and RoCE driver.
RHEL 7.1	Native stack	Client: Production Server: Technical preview	Technical preview	Technical preview	Download 10.6 NIC and RoCE driver.
SLES 11 SP3	3.12-1	Technical preview	No	No	Download 10.6 NIC and RoCE driver.
SLES 12	3.12-1	Technical preview	No	No	Download 10.6 NIC and RoCE driver.

Note: Currently, RoCE is not supported on SLES 11 SP4.

Resolved Issues

FC/FCoE

There are no resolved issues.

NIC

1. The VF interface, when in VLAN promiscuous mode, no longer sees multicast and broadcast traffic intended for other VMs on other VLANs.
2. The status of the Virtual Ethernet Bridge (VEB) is exposed correctly to the operating system.
3. Fixed external loop-back test "ethtool -t <eth> external_lb" on OCe14000-series adapters.
4. The VF interface displays link status correctly when the state is changed from "disable" to "auto".

iSCSI

There are no resolved issues.

RoCE

There are no resolved issues.

Known Issues

1. **On OCe14000-series adapters, firmware versions 10.6.xxx.xx 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.

If you are updating from a firmware version earlier than 10.0.803.37, use one of these methods to update the firmware to 10.6.xxx.xx:

- Use the ISO flash tool.
- Use the released 10.6 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.

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

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

If you are updating from firmware version 10.0.803.37 or later, use one of these methods to update the firmware to 10.6.xxx.xx:

- Use the ISO flash tool.
- Use the released 10.6 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: Do not use Linux inbox drivers when performing the firmware update.

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's 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 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's HBAAPI interface (OneCommand Manager application or third party applications) before performing PCI Hot Plug of an LPFC adapter.
- b) Use the following command to stop the 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 via the sysfs interface or performing a PCI Hot Unplug of an 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.**

Emulex 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 4Gbps adapters in Direct I/O virtualized environments.**

Default driver configuration fails to initialize 4Gbps adapters in virtualized environments that use Direct I/O or SFPT. This may result in a system hang 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 4Gbps 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. 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, see https://bugzilla.redhat.com/show_bug.cgi?id=723936.

7. 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).

8. 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's installation script.

9. 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.

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

As part of the driver kit removal process initiated via 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 above 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
```

11. 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. Emulex highly recommends that you 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's lpfc_nodev_tmo parameter. For example:

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

12. 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'
```

13. On SLES 12 systems, a cable pull before tmo of the guest IO causes IO errors.

Workaround

None.

14. LP16000-series adapters connected to an "I+T" target with LUNs via a fabric may not see the expected number of targets.

Workaround

None.

15. After unloading the driver, an rport dump shows it is deleted, but the counter is not freed.

Workaround

None.

NIC

1. **In certain SKEWs of controllers, PING is not working when attempting to bridge the 1G or 10G ports to the virtual machines when SR-IOV is enabled for 10G 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.

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

Workaround

None.

3. **In certain SKEWs 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.

4. **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"

5. **When using bonding on top of VLANs on a RHEL-based system, networking appears to hang 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 hang during boot or when restarting the network. The system will appear to hang 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's 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.

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

Workaround

None.

7. **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 OneCommand Manager and BIOS components will display the minimum and maximum bandwidths configured, and the PF will display the total bandwidth available.

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

Workaround

Obtain the latest kernel patch from SUSE.

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

Workaround

Obtain the latest kernel patch from SUSE.

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

Workaround

None.

11. **In UMC configuration with 4-port OCe14000-series adapters, minimum bandwidth is not working as expected with default number of (7) queues.**

Workaround

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

12. **When VxLAN offload is enabled in a OCe14000-series adapter, TCP connections through other tunnel interfaces like GRE or IPIP 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/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 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.
```

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 via 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. There is a known Dracut issue that occurs on some systems with specific network settings, where 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. In systems with an open-iSCSI be2iscsi driver installed but not used as bootable, any iSCSI LUN configured in iSCSISelect will not get detected automatically. When the open-iSCSI be2iscsi driver is installed as bootable, any non-bootable iSCSI LUNs configured in iSCSISelect will not get detected automatically.
Configure the iSCSI LUNs using iscsiadm or the OneCommand Manager application.
 4. When Oracle UEK 6.5 x64 and RHEL 6.x and 5.x systems are used with the 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.
5. The system may crash when logging into a target after injecting a PDU error.

Workaround

None.

6. **The storage path may go down while running converged I/O on bond interface and multipath devices.**

Workaround

None.

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

Try manually loading the ocrdma module after doing one of the following:

- set allow_unsupported_modules to '1' in /etc/modprobe.d/unsupported-modules
 - specify --allow-unsupported on the command line
2. **RHEL 6.6, RHEL 6.7, and RHEL 7.x do not have "ibdev2netdev" included in the ibutils package.**

Workaround

As a root user run the following commands on the shell:

```
#git clone git://flatbed.openfabrics.org/compat-rdma/compat-rdma.git
/tmp/compat-rdma
#cp /tmp/compat-rdma/ofed_scripts/ibdev2netdev /usr/local/bin
#rm -rf /tmp/compat-rdma
```

3. **Current implementations of NFS RDMA and KVM VM migration over RDMA do not support IPv6.**

Workaround

None.

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 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 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 Emulex Ethernet driver for 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's 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 iqname. This name can be the same name as in iSCSIselect or a new unique one.

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

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

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.**
- Recommendation**
- 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 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. Setting PFC as the default mode is recommended 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 the 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 PFC settings are enabled on Emulex OCe14000 series adapters, the 8021q.ko kernel module must be loaded if not already loaded by default.**
- This module must be loaded before the 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, it is recommended that you terminate the application(s) that may be using ocrdma and try again.**
- 14. It is recommended that you 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 Emulex website). When you successfully update the profile, the Active-Profile-ID must display 0x15.**
- 15. The recommended RoCE profile for Linux iSER, NFS, or VM Migration is now RoCE-2.**
- 16. Refer to ofed-3.12-1 release notes for ofed issues.**

Copyright ©2012- 2015 Emulex. All rights reserved worldwide. This document refers to various companies and products by their trade names. In most, if not all cases, their respective companies claim these designations as trademarks or registered trademarks. This information is provided for reference only. Although this information is believed to be accurate and reliable at the time of publication, Emulex assumes no responsibility for errors or omissions. Emulex reserves the right to make changes or corrections without notice. This report is the property of Emulex and may not be duplicated without permission from the Company.