

Emulex Driver for Linux Release Notes

Versions: FC/FCoE Version 10.4.299.0 (RHEL, SLES)
NIC Version 10.4.301.0 (RHEL, SLES)
Open-iSCSI Version 10.4.299.0 (RHEL 6, RHEL 7, SLES 11, SLES 12)
Proprietary iSCSI Version 10.0.719.1009 (RHEL 5)
RoCE Version 10.4.301.0 (RHEL 6, RHEL 7, SLES 11)

Date: May 2015

Purpose and Contact Information

These release notes describe the recommended firmware versions, new features, resolved known issues, current known issues, and technical tips associated with these Emulex[®] driver for Linux releases.

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

Recommended Firmware Versions

- Emulex OneConnect[®] OCe14000-series: 10.4.303.14
- Emulex LightPulse[®] LPe16000-series: 10.4.303.3
- Emulex LightPulse LPe12000-series: Firmware 2.02x11, Boot Code 5.13x4

New Features

- Discontinued support for the LPe11000-series and OCe10100-series adapters.
- Fully supports RHEL 5.11, 6.6 inbox, 7.1 inbox; UEK R3 U4; SLES 11 SP4; Citrix XenServer 6.5; Ubuntu 14.10 (NIC and iSCSI)
- Enhanced performance of FCoE, NIC, and iSCSI drivers on OCe14000-series adapters.
- Supports RDMA over Converged Ethernet (RoCE). Table 1 on page 1 lists the operating system support for RoCE.

Table 1 Operating system support for RoCE

OS	OFED or Native Stack	NFS over RDMA	iSER Initiator	iSER Target	Required Emulex Software
RHEL 6.4	3.12-1	Technical preview	No	No	Download 10.4 NIC and RoCE driver.
RHEL 6.5	3.12-1	Technical preview	No	No	Download 10.4 NIC and RoCE driver.
RHEL 6.6	Native stack	Technical preview	No	No	Inbox drivers with 10.4 firmware

Table 1 Operating system support for RoCE (Continued)

OS	OFED or Native Stack	NFS over RDMA	iSER Initiator	iSER Target	Required Emulex Software
RHEL 7.0	Native stack	Not available	No	No	Download 10.4 NIC and RoCE driver.
RHEL 7.1	Native stack	Client: Production Server: Technical preview	Technical preview	Technical preview	Inbox drivers with 10.4 firmware
SLES 11 SP2	3.5	Technical preview	No	No	Download 10.4 NIC and RoCE driver.
SLES 11 SP3	3.12-1	Technical preview	No	No	Download 10.4 NIC and RoCE driver.

Resolved Issues

FC/FCoE

There are no resolved issues for this release.

NIC

There are no resolved issues for this release.

iSCSI

There are no resolved issues for this release.

RoCE

There are no resolved issues for this release.

Known Issues

The following notes apply to updating firmware on OCe14000-series adapters for this release:

- Once an adapter has been flashed to firmware version 10.2.261.36 or later, do not attempt to flash down to an older version without first contacting Lenovo support. Lenovo support will provide a special required downgrade flash procedure. If this procedure is not followed, there is risk in making the adapter permanently unusable.
- If the adapter in use is currently running firmware 10.0.803.2202 or earlier and iSCSI boot firmware table (iBFT) functionality is required, special steps must be followed when upgrading to this release. The recommended flash method is to use the Emulex OneConnect Offline Flash ISO, which will allow flashing in a single step.
- To upgrade and enable iBFT functionality with online tools, the most recent version of the NIC driver, the OneCommand Manager application, and the CIM Provider must first be installed. Additionally, the firmware must be flashed twice with a reboot after each flash.

Some online flash utilities, such as the OneCommand Manager application, may instruct you to reboot and flash the firmware a second time. If iBFT functionality is not required,

this message can be safely ignored. No additional procedures are necessary when iBFT functionality is not required.

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 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 4 Gbps adapters in Direct I/O virtualized environments.

Default driver configuration fails to initialize 4 Gbps 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 Storage Fixed Pass through (SFPT) and 4 Gbps adapters, you must load the LPFC driver with the following driver parameter set:

```
lpfc_sli_mode=2
```

For example:

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

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

6. Order of LPFC module in the initrd module list.

On SLES 11 systems, if another SCSI driver, such as `aic79xx`, is loaded right behind the LPFC driver through the `initrd` image, the loading of the SCSI driver right after the LPFC driver might interrupt the SCSI mid-layer discovery process on the LUNs connected to the LPFC driver. This causes the SCSI discovery reference count put on the LPFC driver to not release and the LPFC driver cannot unload.

Workaround

The Emulex driver kit installation script always puts the LPFC module at the end of the `INITRD_MODULES` list before building the `initrd` image. However, if this problem is observed make sure no SCSI drivers are added right after the LPFC module in the `INITRD_MODULES` list.

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

When using virtual local area network (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 not work if the scripts 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.

8. When using bonding on top of VLANs on a RHEL-based system, networking can 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 hangs 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 Media Access Control (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.

9. 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.)

10. 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.

11. 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.

12. 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 (driver version part of the Linux distribution), which was saved as part of the current driver kit install process, is restored and becomes the active driver. However, the driver kit un-installation 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 possibly 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, you see an error message 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
```

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

An issue was discovered with Cisco Nexus 5000-series FCoE 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 CNAs. 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 CNA 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 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
```


14. **On OCe11100-series adapters, changing the personality to iSCSI using the UEFI utility on RHEL 6.5 and RHEL 6.6 64 bit systems may cause the server not to boot the local operating system.**

Workaround

None.

15. **Low throughput may result when using VLAN bonded interfaces with LPe16000-series adapters on SLES 12 systems.**

Workaround

None.

NIC

1. **ifconfig displays "RUNNING" flag when the cable is not connected.**

This is a known issue in the Linux operating system stack.

Workaround

Use "ip link show" to properly display the "Oper status" value.

2. **PING is not working when attempting to bridge the 1 Gb/s or 10 Gb/s ports to the virtual machines when single root I/O virtualization (SR-IOV) is enabled for 10 Gb/s 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 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 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.

Note: Another option may be to use the MacVTap interface to the guest instead of the bridge interface. The MacVTap interface is only available in the RHEL 6.3 or newer kernels.

3. **Unable to ping the RHEL 6.4 kernel-based virtual machine (KVM) guest virtual machine (VM) from a remote host when using the MacVTap driver and IPv6.**

When an IPV6 unicast address is configured, the Linux kernel also configures an Ethernet multicast MAC address in the format 33-33-xx-xx-xx-xx. In the context of a guest operating system (VM), the mcast address that needs to be configured comes into existence after the VM is booted up and the unicast IPv6 address is configured in the guest operating system.

When using a MacVTap interface to bridge the NIC interface to the VM's interface, the previously-mentioned Ethernet MAC address configuration in the VM does not reach the MacVTap interface automatically. Since the MacVTap interface does not put the NIC

interface into promiscuous mode, this results in an IPv6 ping failure from the remote host to the VM's interface.

Workaround

Select one of the following workaround options:

- Use the following command:

```
ifconfig eth<x> allmulti
```

where “eth<x>” is the interface in the hypervisor to which the MacVTap interface is bridged. However, this causes all multicast traffic to be received.

- Alternatively, you can configure the unicast IPv6 address (assigned to the VM interface) on the MacVTap interface in the hypervisor. This results in the Ethernet mcast MAC address being configured in the NIC interface.

Use the following command:

```
ifconfig macvtap<x> inet6 add <ipv6 address>
```

where “macvtap<x>” is the MacVTap interface in the hypervisor and “<ipv6 address>” is the unicast IPv6 address assigned to the VM interface.

4. **When SR-IOV is enabled, NIC priority group (PG) and priority flow control (PFC) are not supported.**

Workaround

None.

5. **For OCe11100-series adapters, when the driver is loaded with num_vfs=32, the initialization of two of the virtual functions (VFs) fails and only 30 VF interfaces are created.**

Workaround

None.

6. **In certain configurations, timeout errors may occur during maximum performance socket testing.**

Workaround

The default Linux driver settings may not be appropriate for optimal performance in all scenarios. If performance appears to be lower than expected, there are several driver and system settings that can be modified to improve performance. See the server documentation to determine the correct system settings and the optimal memory and processor configuration.

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

All Linux:

- Run the perf_tune_benet.sh script (included with Linux NIC driver)

Linux Xen:

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

Linux KVM:

- Use the “virtio” device model instead of “Hypervisor Default”

7. The description for OCe14102-U3-D shows as ‘unknown device 0720’ in the Xen Center for Xen 6.1.

Workaround

Run the update-pciid command to display the correct adapter name.

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

Workaround

None.

9. For OCe14000-series adapters, if the port is not linked up to the switch at the time of driver loading with a non-zero num_vfs parameter, ethtool / ifconfig will report that the physical link is detected, even though there is no physical link.

Workaround

None.

10. When a CNA is configured in a NIC + iSCSI profile, the NIC and iSCSI traffic could be configured to share the total bandwidth and NIC traffic will be assured a minimum bandwidth.

In such a configuration, a VF interface will inherit the minimum bandwidth of the PF.

Workaround

None.

11. In SLES 11 SP3, when a bond is created between fn0 and fn1 (or fn2 and fn3 for example), bandwidth that is dynamically changing is not being updated in the bonded interfaces (slaves) in Unified Fabric Port (UFP) mode.

Workaround

Disable and then enable the ports to update the bandwidth in the bond.

12. Assigning VFs to a VM on the SLES operating system.

To assign VFs to the VM in the SLES Xen kernel, the VF must be unbound from the NIC module and then bound to the pciback module.

Note: In the following steps, “0000:07:0b.5” is used as an example. To match those instances to the port that you want to use, you will need to select the entry which matches the PCI bus, device, or function which corresponds to the port that you want to assign. The ethtool utility can be used to determine this information, such as ethtool -i eth0 (where eth0 is the interface you want to assign).

- a) Load the pciback driver “modprobe pciback”.
- b) Navigate to the /sys/bus/pci/drivers/pciback directory and ensure that the following is displayed:

```
Panama-Sles11sp2:/sys/bus/pci/drivers/pciback # ls -lrt
```

```
total 0
--w----- 1 root root 4096 Sep  5 15:29 unbind
--w----- 1 root root 4096 Sep  5 15:29 uevent
-r----- 1 root root 4096 Sep  5 15:29 slots
--w----- 1 root root 4096 Sep  5 15:29 remove_slot
--w----- 1 root root 4096 Sep  5 15:29 remove_id
-rw----- 1 root root 4096 Sep  5 15:29 quirks
-rw----- 1 root root 4096 Sep  5 15:29 permissive
--w----- 1 root root 4096 Sep  5 15:29 new_id
lrwxrwxrwx 1 root root    0 Sep  5 15:29 module ->
../../../../module/pciback
--w----- 1 root root 4096 Sep  5 15:34 new_slot
--w----- 1 root root 4096 Sep  5 15:34 bind
```

- c) Navigate to the `/sys/bus/pci/drivers/be2net` directory and ensure that the following is displayed:

```
--w----- 1 root root 4096 Sep  5 15:32 uevent
--w----- 1 root root 4096 Sep  5 15:32 remove_id
--w----- 1 root root 4096 Sep  5 15:32 new_id
lrwxrwxrwx 1 root root    0 Sep  5 15:32 module ->
../../../../module/be2net
--w----- 1 root root 4096 Sep  5 15:32 bind
--w----- 1 root root 4096 Sep  5 15:33 unbind
lrwxrwxrwx 1 root root    0 Sep  5 15:32 0000:07:0b.5 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.5
lrwxrwxrwx 1 root root    0 Sep  5 15:32 0000:07:0b.4 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.4
lrwxrwxrwx 1 root root    0 Sep  5 15:32 0000:07:0b.3 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.3
lrwxrwxrwx 1 root root    0 Sep  5 15:32 0000:07:0b.2 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.2
lrwxrwxrwx 1 root root    0 Sep  5 15:32 0000:07:0b.1 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.1
lrwxrwxrwx 1 root root    0 Sep  5 15:32 0000:07:0b.0 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.0
```

- d) Unbind the VF from the NIC driver:

```
echo -n "0000:07:0b.5" > /sys/bus/pci/drivers/be2net/unbind
```

- e) Bind the driver to the pciback module:

```
echo -n "0000:07:0b.5" > /sys/bus/pci/drivers/pciback/new_slot
echo -n "0000:07:0b.5" > /sys/bus/pci/drivers/pciback/bind
```

- f) Navigate to the directory `/sys/bus/pci/drivers/pciback` and ensure that the device "0000:07:0b.5" is listed under it. In addition, check that `dmesg` logs report the same device.
- g) You can now launch `qemu-kvm` and attach VF "0000:07:0b.5" to any desired VM.

13. TCP Traffic stalls on a tunnel interface (GRE/IPIP) if a VxLAN device is created on a physical port.

This could happen if VxLAN offloads are enabled in the device. When VxLAN offloads are enabled, the following `be2net` driver message can be seen in the output of `dmesg`.

```
"Enabled VxLAN offloads for UDP port <port-num>"
```

This issue has been resolved in Linux kernel versions 3.19.4 and above.

Workaround

Disable TX checksum offload on the network interface on which GRE tunnel is configured.

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

14. When ARI is enabled on a CNA-based profile, only seven functions are properly initialized out of the 14 NIC functions in UMC.

Workaround

Add "realloc = on" to the kernel boot parameters.

`pci=realloc`

<https://www.kernel.org/doc/Documentation/kernel-parameters.txt>

`realloc` = Enable/disable reallocating PCI bridge resources if allocations done by the BIOS are too small to accommodate resources required by all child devices.

off: Turn realloc off

on: Turn realloc on

iSCSI

1. In Linux, the device driver no longer loads in the operating system after a kernel update.

During the installation of a Linux kernel update, the following message (or a similar message) may be seen:

```
Cannot determine dependencies of module be2iscsi. Is modules.dep up to date?
```

Workaround

Issue the following commands from the Command Line Interface (CLI) to resolve any driver conflicts. The [new-kernel] and [old-kernel] statements below must be replaced as appropriate.

- a) `cd /lib/modules/[new-kernel]/weak-updates/`
- b) `ln -s /lib/modules/[old-kernel]/updates/be2iscsi/be2iscsi.ko`
- c) `ln -s /lib/modules/[old-kernel]/updates/lpfc.ko`
- d) `ln -s /lib/modules/[old-kernel]/updates/be2net/be2net.ko`

Note: Steps b through d may not all be necessary depending on the personality of the installed CNA.

- e) `depmod 2.6.18.60-0.93.1-smp`
- f) `mkinitrd`
- g) `shutdown -r now`

After this workaround is applied, the server can boot to the new kernel.

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. **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.

4. **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 log in.

5. **The open-iSCSI driver be2iscsi depends on some operating system utilities and packages. If these packages are not present, the session establishment may not be possible after successful installation of the driver rpm.**

Workaround

Normally, a system previously installed with iSCSI hardware will have these packages installed automatically. In some corner cases, such as a new iSCSI adapter that was installed after the operating system installation, these rpm packages will need to be installed manually. These packages can be found on the installation CD.

For RHEL 6.x and RHEL 7, the dependencies will be:

- dracut-network
- iscsi-initiator-utils

For SLES 11 SPx, the dependency will be:

- open-iSCSI

6. **In RHEL 6.4 x86 systems with the Open-iSCSI driver installed, iSCSI bootable LUNs configured in iSCSISelect are not automatically detected.**

Workaround

Configure the iSCSI LUNs using iscsiadm or the OneCommand Manager application.

7. **When open-iSCSI is used in the SLES 11 SP3 operating system, all adapters will function as a single iSCSI initiator. For previous releases, the following behavior can be expected:**

- Proprietary iSCSI driver:
 - Uses the initiator iSCSI qualified name (IQN) name which was set in the iSCSISelect utility per adapter. Multiple adapters can have different IQN names. All logins were done by arm firmware via saved sessions from flash, which were previously configured in the iSCSISelect utility.
- Open-iSCSI driver:
 - For the case of boot target, the IQN name for the initiator connected to the boot target is used in the operating system for all adapters, even if there are different IQN names for other adapters set in the iSCSISelect utility. In the operating system, iscsid will use the entry in /etc/iscsi/initiatorname.iscsi for all sessions. There is one initiator IQN name per host.
 - For the case of non-iSCSI boot, the new IQN in /etc/iscsi/initiatorname.iscsi can be manually set or left as-is. There is one initiator IQN name per host.
 - Driver sysfs entries for adapter initiator iqn names will still have the original IQN names on flash per adapter. iscsid does not use this information.
 - Unlike the proprietary iSCSI driver, sessions login by arm and sessions login by the driver/operating system are not the same. The driver/operating system is not aware of existing arm firmware sessions other than the boot session. It is possible to have multiple sessions logged in to the same target that has the same initiator IQN name.

8. **RHEL 6.x and SLES 11 SPx inbox drivers may not boot when the adapter is in the iSCSI personality.**

Note: The driver update disk or package (DUD or DUP) must be used if the operating system is being installed while the adapter is in the iSCSI personality.

Workaround

Install the latest out-of-box Linux iSCSI driver before using the iSCSI port.

-or-

For the inbox Linux iSCSI driver, edit the grub menu during boot or in /boot/grub/menu.lst or /boot/efi/EFI/redhat/grub.conf:

add "maxcpus=4" at the line starting with "kernel".

9. When IBM Fabric Manager (IFM) is enabled in RHEL 6.5 x86, the main system Ethernet (based on Dynamic Host Control Protocol [DHCP]) is not functioning. This prevents Dracut from logging into the boot session, and a kernel panic can occur.

Workaround

Remove the following two kernel parameters from the `/boot/grub/menu.lst` file:

```
ip=xyz
ifname=<MAC>
```

RoCE

1. When installing the OpenFabrics Enterprise Distribution (OFED) 3.5-1 package on the RHEL 6.4 operating system, the following additional steps must be performed to enable Network File System (NFS) over remote direct memory access (RDMA):

- a) Change the following lines in the `install.pl` file from

```
#NFSRDMA
    if ($kernel =~ m/^3\.5/ or $DISTRO =~ /SLES11.2|RHEL6.[23]/) {
to
#NFSRDMA
    if ($kernel =~ m/^3\.5/ or $DISTRO =~ /SLES11.2|RHEL6.[234]/) {
```

Note: If OFED has already been installed, it must be uninstalled, modified, and reinstalled.

- b) When installation of the OFED stack is complete, do one of the following:

- reboot the system
- restart OFED by typing

```
service openibd restart
```

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

In such cases, it is possible that the installer can fail to load the modules even though the RPM 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`
- specify `--allow-unsupported` on the command line

3. If generic pause is being used, ensure that switches support this feature.

If CEE/DCB mode is enabled on the switch, enable PFC for priority 5 (configure RoCE traffic on priority 5). If CEE/DCB is not enabled, generic pause should be enabled on the switch and the NIC port to prevent loss of packets.

4. 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.

Workaround

None.

Technical Tips

1. **Additional physical NICs added to the XenServer do not appear in XenCenter.**

Workaround

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

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

2. **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, SLES 11, and SLES 12 releases.

3. **When the SLES 11u1x32 or SLES 11u2x32 operating systems are used with the OCe11100-series adapters with the Open-iSCSI driver, the OEMSEAddSendTargetPortal 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.

4. **FC in-band management is no longer supported.**
5. **The OneCommand Manager application no longer installs OneCommand Vision components.**
6. **To enable flexibility to the number of mounts that can be allowed over a single port, decrease the number of memory regions (MRs) from the default high of 32.**

For example:

To lower the default MR usage to 16, add the following line to the etc/rc.local file:

```
sysctl -w sunrpc.rdma_slot_table_entries=16
```

Note: Setting an MR value to less than 8 is not recommended, because a significant performance drop will occur.

7. **Setting PFC as the default mode is recommended when possible while using RoCE.**

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

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

Copyright © 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.
Open-iSCSI GPLv2 Notice

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.
51 Franklin Street, Fifth Floor, Boston, MA 02110-1301, USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Lesser General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The “Program”, below, refers to any such program or work, and a “work based on the Program” means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term “modification”.) Each licensee is addressed as “you”.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange;

or,

b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.