

Driver for Linux Release Notes

Versions: Driver for Linux, FC/FCoE Version 10.4.255.16 (RHEL, SLES)
Driver for Linux, NIC Version 10.4.255.20 (RHEL, SLES, Debian, Ubuntu)
Driver for Linux, Open-iSCSI Version 10.4.245.0 (RHEL 6, SLES 11)
Driver for Linux, Proprietary iSCSI Version 10.0.719.1009 (RHEL 5)
Driver for Linux, RoCE Version 10.4.255.20 (RHEL, SLES)

Date: March 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@emulex.com, 800-854-7112 (US/Canada toll free), +1 714-885-3402 (US/International), or +44 1189-772929 (Europe, Middle East, and Africa).

New Features

- 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.
- Adds RoCE, iSCSI, and FCoE support for OCe14102-NT and OCe14102-UT 10GBASE-T adapters.
- Dropped support for OCe10100-series adapters.
- Supports RDMA over Converged Ethernet (RoCE). Refer to Table 1 on page 2 for details.

Note: Technical Preview indicates RoCE is available for testing, but should not be used for production.

Note: Firmware version 10.4 is required.

Table 1 Supports RDMA over Converged Ethernet (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
RHEL 7.0	Native stack	Technical preview	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.

New NIC Features

- Adds VxLAN offload support for NIC.
- Supports RHEL inbox driver kdump testing.
- Adds Vendor ID check for 'Emulex Qualified' SFP modules and logs vendor information for unsupported modules.
- Adds Checksum, TSO and RSS steering offload support for VXLAN (QC 3415, BZ 161751)
- Enables support to configure source MAC address spoofing (in ip-link(8) command) for LPe16000-series adapters.
- Adds support to configure RSS hash key via ethtool(8) command.
- Supports to retrieve EEPROM information via ethtool(8) command.
- Exports board temperature via the sysfs interface.
- Support added for IRQ pinning to different CPU cores.
- Enabled PCIe error reporting on virtual functions.

Resolved Issues

FC/FCoE

1. Enhanced performance of FCoE drivers on OCe14000-series adapters.

NIC

1. The *Emulex Linux Drivers version 10.4 User Manual* has instructions for assigning VFs to a VM on the SLES operating system.
2. Enhanced performance of NIC drivers on OCe14000-series adapters.
3. Supports Firmware dump using "ethtool -W" command for LPe16000-series adapters.
4. RSS support is enabled for PFs in OCe14000-series adapters, when SRIOV is enabled and requested number of VFs are less than maximum supported.

iSCSI

1. The open-iSCSI driver be2iscsi dependency on some operating systems utilities and packages regarding session establishment after driver installation is resolved.
2. Enhanced performance of iSCSI drivers on OCe14000-series adapters.

RoCE

1. Enhanced performance of RoCE drivers on OCe14000-series adapters.

Known Issues

Firmware version 10.4.xxx.xx includes new features that required 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 10.4.xxx.xx:

- Use the ISO flash tool.
- Use the released 10.4 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 10.4.xxx.xx:

- Use the ISO flash tool.
- Use the released 10.4 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.

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

Default driver configuration fails to initialize 4 Gb/s 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 4 Gb/s adapters, you must load the LPFC driver with the following driver parameter:

- `lpfc_sli_mode=2`

For example:

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

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

6. 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 SCSI mid-layer discovery process may be interrupted 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 places 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 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.

8. The 16Gb LPe16000B and LPe16002B adapters are not supported with the RHEL 5.9 and RHEL 6.4 inbox FC/FCoE drivers.

The RHEL 5.9 inbox FC/FCoE (LPFC) driver version 8.2.0.128.3p and the RHEL 6.4 inbox FC/FCoE (LPFC) driver version 8.3.5.86.1p do not support the 16Gb LPe16000B and LPe16002B adapters (PCI Express 3.0 Bus Type). However, these driver versions do support the 16Gb LPe16000 and LPe16002 adapters (PCI Express 2.0 Bus Type).

In Boot from Storage Area Network (SAN) configurations over the unsupported LPe16000B and LPe16002B adapters, the process of booting from SAN will fail. In

standard local disk boot configurations, although the FC/FCoE driver will load upon system boot, the target storage LUNs will not be discovered.

Note: You can use the Linux “lspci” utility to determine whether an adapter is LPe1600xB or LPe1600x, and therefore whether it is supported by the RHEL 5.9 inbox LPFC driver version 8.2.0.128.3p or the RHEL 6.4 inbox LPFC driver version 8.3.5.86.1p.

Executing the “lspci -x” Linux utility displays the PCI configuration space for each device in the form of:

```
<bus/slot/function info> <strings matching vendor id, etc>
00: xx xx xx xx ...
10: xx xx xx xx ...
...
```

For example, for the Emulex 16Gb adapter:

```
0a:00.1 Fibre Channel: Emulex Corporation Lancer-X: LightPulse
Fibre Channel Host Adapter
00: df 10 00 e2 46 00 10 00 00 00 04 0c 10 00 80 00
10: 0c 00 cd f1 00 00 00 00 00 00 00 00 00 00 00 00
20: 00 00 00 00 00 00 00 00 00 00 00 00 00 df 10 00 e2
30: 00 00 00 00 54 00 00 00 00 00 00 00 00 0b 02 00 00
```

The 9th byte (offset 0x8) of the PCI configuration space of the LPe1600x/LPe1600xB adapters indicates the Revision ID register value of the adapter.

- If the Revision ID register value is 0x00, this represents an LPe1600x adapter.
- If the Revision ID register value is 0x10, this represents an LPe1600xB adapter.

In the example above, the value of the byte with offset 0x8 is 0x00; therefore, this represents an LPe1600x adapter, so the RHEL 5.9 inbox FC/FCoE (LPFC) driver version 8.2.0.128.3p and RHEL 6.4 inbox FC/FCoE (LPFC) driver version 8.3.5.86.1p do support this adapter.

Workaround

Support for the 16Gb LPe16000B and LPe16002B adapters will be provided with the FC/FCoE (LPFC) driver included in the RHEL 6.5 release. In addition, support is provided for the RHEL 5.9 and later releases through independent out-of-box FC/FCoE (LPFC) driver packages hosted on the Emulex website.

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. Driver package installation failure may occur when installing on systems with SLES 11 SP2 maintenance kernel version 3.0.80-0.7.

The SLES11 SP2 maintenance kernel release with version 3.0.80-0.7 has a kernel version higher than the latest SLES11 SP3 GM release (which was 3.0.76-0.11.1). This is a non-standard kernel versioning scheme that is used by Linux operating system vendors. See the related SuSE KB article at

<https://www.suse.com/support/kb/doc.php?id=7012717>.

This non-standard versioning scheme causes the Emulex driver package installation scripts to fail when an attempt is made to install the packages on this specific SLES11 SP2 maintenance release.

Workaround

To install the driver on this specific SLES11 SP2 maintenance kernel (version 3.0.80-0.7), you should not use the driver package's installation script. Instead, you should manually install the driver binary RPM built for the SLES11 SP2 GM kernel (version 3.0.13-0.27) and corresponds to the architecture of the target system (for example, x86_64, PPC, and so on). These driver binary RPMs include the kernel version (3.0.13_0.27) in their name.

For example:

```
# rpm -Uhv
elx-lpfc-kmp-default-8.3.7.27_3.0.13_0.27-1.sles11sp2.x86_64.rpm
```

11. 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.

12. 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.

13. 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
```

14. 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
```

15. 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.

Work around

Edit the multipath.conf file and replace

```
uid_attribute 'ID_SERIAL'
```

with

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

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

Work around

None.

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

Work around

None.

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

Work around

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 display the "Oper status" value.

2. 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.

3. Unable to ping the RHEL 6.4 KVM guest (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. In certain SKEWs of controllers, when SR-IOV is enabled, NIC priority group (PG) and priority flow control (PFC) are not supported.

Workaround

None.

5. 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.

6. 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”

7. 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.

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

Workaround

None.

9. 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.

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

Workaround

Obtain the latest kernel patch from SUSE.

11. VF interface is able to see all multicast and broadcast traffic which is intended for other VMs on other VLANs, when put into promiscuous mode.

Workaround

None.

12. The status of the Virtual Ethernet Bridge (VEB) is not exposed to the operating system.

Workaround

None.

13. GRE tunnels do not work when a VxLAN tunnel is configured.

Workaround

None.

14. 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.

15. External Loopback test "ethtool -t <eth> external_lb" does not work on OCe14000-series adapters.

Workaround

None.

16. In certain SKEWs of controllers, VF does not show link up when the state is changed from disable to auto, although PF uplink is connected.

Workaround

None.

17. Low throughput is observed with LPe16000-series adapters, when only one TX queue is created on the interface.

Workaround

None.

18. 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.

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. **Adding a target profile fails with the iSCSI driver on RHEL 6.6 and RHEL 7.0.**

Workaround

None.

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

Workaround

None.

RoCE

1. **Reducing of CQ depth is not supported.**

Workaround

None.

2. **The minimum supported MTU RoCE interfaces is 512.**

Note: The MTU values on the RoCE interface can be only powers of 2 with the maximum MTU value being 4096.

Workaround

None.

3. **The host may fail to boot on certain vendor system BIOS versions when a single PF is configured with SR-IOV disabled.**

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 enabled.

Note: This issue can also be observed on 2 and 4-port adapters that support "link reconfigurability" and are reconfigured for a single port.

- ◆ Affected vendor systems include:
- ◆ Colfax/Intel system with S2600C0 motherboard
- ◆ SuperMicro X9SRL-F motherboard
- ◆ SuperMicro X9SCL/X9SCM

Workaround

None.

4. **When unloading the ocrdma driver, a "module ocrdma is in use" error message may appear. Linux does not allow unloading of modules when they are in use.**

Workaround

When trying to upgrade existing drivers please use `elx_roce_install.sh` and reboot the system.

Technical Tips

1. **On systems running RHEL 6.4 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 SP2 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 6.4 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 SP2 and SP3.

- a) Ensure that the Open-iSCSI package is installed. If it is not, mount the installation CD and install it.
SLES 11 SP3 example:

```
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
```

SLES 11 SP2 example:

```
linux-8vq8:/media/SLES-11-SP2-DVD-x86_6407551/SuSE/x86_64 # rpm
-ihv open-iscsi-2.0.872-0.35.1.x86_64.rpm
```

- b) Edit `/etc/iscsi/initiatorname.iscsi` to add the initiator iqname. This name can be the same name as in `iscsi`select 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 `iscsi`select. 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.

SLES 11 SP2 example:

```
linux-8vq8:/SLES11_64 # rpm
-Uhv be2iscsi-kmp-default-10.0.739.0_3.0.13_0.27-1.sles11sp2.x86
64.rpm
```

SLES 11 SP3 example:

```
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`.

SLES 11 SP3 example:

```
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

SLES 11 SP2 example:

```
linux-8vq8:/boot # mkinitrd -i initrd-3.0.13-0.27-default -k  
vmlinuz-3.0.13-0.27-default -f iscsi
```

Features:

- acpi
 - dm
 - multipathd
 - block
 - usb
 - iscsi
 - multipath
 - kpartx
 - resume.userspace
 - resume.kernel
- 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.

SLES 11 SP3 example:

```
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

SLES 11 SP2 example:

```
linux-8vq8:/boot # mkinitrd -i initrd-3.0.13-0.27-default -k
vmlinuz-3.0.13-0.27-default -f multipath
```

Features:

- acpi
  - dm
  - multipathd
  - block
  - usb
  - 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.

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.

Note: References to OCE11100 series products also apply to OCE11100R series products.

### 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.