



RoCE Driver Version 10.2 for Linux

Technical Preview Guide

Copyright © 2014 Emulex. All rights reserved worldwide. No part of this document may be reproduced by any means or translated to any electronic medium without the prior written consent of Emulex.

Information furnished by Emulex is believed to be accurate and reliable. However, no responsibility is assumed by Emulex for its use; or for any infringements of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright or related rights of Emulex.

Emulex, the Emulex logo, AutoPilot Installer, AutoPilot Manager, BlockGuard, Connectivity Continuum, Convergenomics, Emulex Connect, Emulex Secure, EZPilot, FibreSpy, HBAnyware, InSpeed, LightPulse, MultiPulse, OneCommand, OneConnect, One Network. One Company., SBOD, SLI, and VEngine are trademarks of Emulex. All other brand or product names referenced herein are trademarks or registered trademarks of their respective companies or organizations.

Emulex provides this manual "as is" without any warranty of any kind, either expressed or implied, including but not limited to the implied warranties of merchantability or fitness for a particular purpose. Emulex may make improvements and changes to the product described in this manual at any time and without any notice. Emulex assumes no responsibility for its use, nor for any infringements of patents or other rights of third parties that may result. Periodic changes are made to information contained herein; although these changes will be incorporated into new editions of this manual, Emulex disclaims any undertaking to give notice of such changes.

Emulex, 3333 Susan Street
Costa Mesa, CA 92626

Note: References to OCe11100 series products also apply to OCe11100R series products.

Table of Contents

List of Tables	5
1. Introduction	6
Overview	6
Supported Adapters	6
OFED Overview	6
Abbreviations	6
2. Installing and Uninstalling	8
Installing the RoCE Driver for the OCe14000-Series Adapters	8
Installing OFED	8
RoCE Driver Installation.....	9
Uninstalling the RoCE Driver	10
3. RoCE Configuration for the OCe14000-Series Adapters.....	11
Basic Configuration	11
Interface Configuration	11
Sample Applications	12
Setting the Profile	12
Configuring VLANs.....	13
MTU Configuration	13
Advanced Applications	14
NFS over RDMA	14
Server Configuration	14
Client Configuration	15
Enabling RDMA on all Four Ports on an OCe14000-Series Adapter	16
Advanced Configuration	18
QoS Behavior	18
OCe14000-Series Defaults	18
QoS Configuration Guidelines	19
Priority Groups	19
L2 Flow Control	19
DCBX Enabled Switch.....	19
DCBX Disabled Switch Connection (generic pause mode)	21
Updating the Adapter Firmware for RoCE.....	21
Determine Firmware Version.....	21
Updating the Firmware Manually.....	22
Updating the Firmware using OneCommand Manger	22

4. Troubleshooting 23
 Log Messages 23
 RoCE Error Log Messages and their Descriptions 23

List of Tables

Table 2-1	Compatible OFED Versions.....	8
Table 3-1	Setting sunrpc.rdma_slot_table_entries.....	16
Table 4-1	RoCE error log messages and their descriptions.....	23

1. Introduction

Overview

This technical preview guide provides requirements and steps to install and configure the Emulex® RDMA over converged Ethernet (RoCE) device driver in the Linux environment.

This guide is applicable for the following operating systems:

- Red Hat Enterprise Linux (RHEL) version 6.4
- SUSE Linux Enterprise Server (SLES) version 11 SP2

Supported Adapters

The OCe14000 family of adapters is supported.

OFED Overview

The OpenFabrics Enterprise Distribution (OFED™) is open-source software for Remote Direct Memory Access (RDMA) and kernel bypass applications. OFED is used in business, research and scientific environments that require highly efficient networks, storage connectivity and parallel computing. The software provides high performance computing sites and enterprise data centers with flexibility and investment protection as computing evolves towards applications that require extreme speeds, massive scalability and utility-class reliability.

Some operating system distributions ship OFED in-box. It is recommended to only install the OFED version for the supported operating system. Refer to Table 2-1 on page 8 for a list of compatible OFED versions.

Note: The entire set of OpenFabrics Software – from which modules and patches are selected to form OFED releases – resides on the OpenFabrics servers and is available for download.

Abbreviations

API	application programming interface
BIOS	basic input/output system
CPU	central processing unit
DCBX	data center bridging exchange
FCoE	Fibre Channel over Ethernet
HCA	host channel adapter
IOV	I/O virtualization
IP	Internet Protocol
IPL	initial program load

iSCSI	internet Small Computer System Interface
MR	memory region
MPI	parallel message passing
MTU	maximum transmission unit
NAS	network attached storage
NFS-RDMA	network file system over RDMA
NIC	network interface card (or controller)
OFED	OpenFabrics Enterprise Distribution
PCI	Peripheral Component Interconnect
PCIe	Peripheral Component Interconnect Express
QoS	quality of service
RDS	remote desktop service
RDMA	remote direct memory access
RHEL	Red Hat Enterprise Linux
RoCE	RDMA over converged Ethernet
Rx	receive mode
SAN	storage area network
SCSI	Small Computer System Interface
SDP	Sockets Direct Protocol
SLES	SUSE Linux Enterprise Server
SLI	service level interface
SR-IOV	single-root I/O virtualization
SRP	SCSI RDMA Protocol
TCP	Transmission Control Protocol
Tx	transmit mode
VLAN	virtual local area network

2. Installing and Uninstalling

Installing the RoCE Driver for the OCe14000-Series Adapters

RoCE is a network protocol that allows remote direct memory access over an Ethernet network. RoCE is a link layer protocol which allows communication between any two hosts in the same Ethernet broadcast domain.

Network-intensive applications like networked storage or cluster computing require a network infrastructure with a high bandwidth and low latency. The advantages of RDMA over other network application programming interfaces are lower latency, lower CPU load and higher bandwidth.

The following operating systems are supported for RoCE:

- SLES 11 SP2
- RHEL 6.4

Installing OFED

The following table lists the available OFED versions that are compatible with the supported operating systems.

Table 2-1 Compatible OFED Versions

OFED Version	Operating System	
	SLES 11 SP2	RHEL 6.4
OFED 3.5	YES	NO
OFED 3.5-1	NO	YES

Supported OFED packages can be downloaded from the following website:

<https://www.openfabrics.org/downloads/OFED/>

Note: When installing OFED-3.5-1 packages on the RHEL 6.4 operating system, perform the following additional step to enable NFS over RDMA:

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]/) {
```


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

1. Open a web browser and navigate to
`www.openfabrics.org/downloads/OFED`
2. Download the appropriate .tgz tarball file.
3. Extract the downloaded OFED-x.tgz tarball to the /tmp directory:

```
#tar xvzf/tmp/OFED-3.5.tgz
```

4. Install OFED:

```
#cd /tmp/OFED-3.5  
#install.pl -all
```

Note: Package dependencies must be resolved for a successful installation.

5. When installation is complete, do one of the following:
 - reboot the system
 - restart OFED by typing

```
service openibd restart
```

RoCE Driver Installation

To install the RoCE driver:

1. Download the following package from the Emulex website:
`elx-ocrdma-dd-<release>-<version>.tar.gz`
2. Copy the package to /tmp and run

```
tar xvzf elx-ocrdma-dd-<release>-<version>.tar.gz
```
3. Change directory to the RoCE packages directory:

```
cd /tmp/X.X.X.X/Linux/RoCE
```
4. Select the required package based on the host on which the driver is being installed.
5. The current Linux distribution on the host may be determined by typing

```
lsb_release-i
```
6. Extract the selected package using tar.
7. Change directory to the resulting directory from the tar extraction command to find the installer script 'elx_roce_install.sh'.
8. Run the installer script to install the required RoCE driver and user library RPMs.

Uninstalling the RoCE Driver

To uninstall the RoCE driver, type

```
elx_roce_install.sh --uninstall
```

Note: SLES 11 platforms may prevent loading of unsupported modules by default. In such cases, the installer may fail to load the modules even though the RPMs are installed. If this occurs, try manually loading the ocrdma module after doing one of the following:

- set 'allow_unsupported_modules' to '1' in /etc/modprobe.d/unsupported-modules
- specify '--allow-unsupported' on the command line

3. RoCE Configuration for the OCe14000-Series Adapters

Notes:

- RoCE + multichannel is not supported.
- SR-IOV is not supported when RoCE is enabled.
- It is highly recommended that PFC is enabled as the default mode when possible while using RoCE. See “Advanced Configuration” on page 18 to enable QoS for RoCE. In addition, VLAN interfaces must be configured and used for RoCE traffic in order for PFC to work correctly (see “Configuring VLANs” on page 13). VLANs greater than 1 should be used for best interoperability.

RoCE profiles can be specified in the OneCommand Manager GUI application, the OneCommand Manager CLI application, and the PXE Boot utility. Regardless of the utility you use, follow these guidelines to select the appropriate RoCE profile:

- Choose the RoCE-2 profile for NFS on Linux.

Note: Check the Implementer's Lab on the Emulex website for any updated information on additional use cases for the RoCE-2 profile.

- For the RoCE-1 profile, check the Implementer's Lab on the Emulex website for any updated information on use cases for the RoCE-1 profile.

For specific information on selecting RoCE profiles, see the applicable manual:

- *OneCommand Manager Application User Manual*
- *OneCommand Manager Command Line Interface User Manual*
- *Boot for NIC, iSCSI, FCoE, and RoCE Protocols User Manual*

Basic Configuration

Interface Configuration

Identify the RoCE interfaces and their corresponding NIC interfaces by using the following commands:

To list the ocrdma interfaces.:

```
ibv_devinfo -l
```

To list the corresponding NIC interfaces:

```
ibdev2netdev
```

Follow the standard procedure to assign a valid IP address to the desired Ethernet interface that corresponds to the RoCE port. You can assign an IP address to eth4 to use 'ocrdma0' for RoCE. Example output:

```
ibv_devinfo -l
2 HBAs found:
```

```
ocrdma1
ocrdma2
ibdev2netdev
ocrdma0 port 1 ==> eth0 (Up)
ocrdma1 port 1 ==> eth1 (Up)
ocrdma2 port 1 ==> eth2 (Up)
ocrdma3 port 1 ==> eth3 (Up)
```

Sample Applications

The RoCE connectivity can be tested using the following OFED i-built tools/applications:

```
ibv_rc_pingpong
ib_send_bw
ib_read_bw
ib_write_bw
```

Note: The use of these commands assume that the client (ocrdma0) and server (ocrdma0) interfaces are configured with IPs 11.192.168.x and 11.192.168.x respectively.

Examples:

```
ibv_rc_pingpong
Server: ibv_rc_pingpong -g 0 -d ocrdma0
Client: ibv_rc_pingpong -g 0 -d ocrdma0 11.192.168.x

ib_send_bw
Server: ib_send_bw -d ocrdma0
Client: ib_send_bw -d ocrdma0 11.192.168.x

ib_read_bw
Server: ib_read_bw -d ocrdma0
Client: ib_read_bw -d ocrdma0 11.192.168.x

ib_write_bw
Server: ib_write_bw -d ocrdma0
Client: ib_write_bw -d ocrdma0 11.192.168.x
```

Setting the Profile

The RoCE profile can be enabled by using PXESelect BIOS or the OneCommand Manager application.

- To configure the adapter using PXESelect BIOS, refer to the *Boot for NIC, iSCSI, FCoE, and RoCE Protocols User Manual* for more information on the PXESelect BIOS utility.

- To configure the adapter using the OneCommand Manager application, refer to the *OneCommand Manager Application User Manual*, or the *OneCommand Manager Command Line Interface User Manual*.

The profile can be confirmed from the Ethernet driver load messages in `/var/log/messages`. For example:

```
Active profile ID ROCE-2
```

Note: If the correct profile is not reported, update the profile ID using the OneCommand Manager or PXESelect BIOS.

Configuring VLANs

To configure a VLAN interface, perform the following steps:

1. Load the 8021q module (if necessary) by typing

```
modprobe 8021q
```

2. Create a VLAN interface by typing

```
vconfig add eth<x><vlan id>
```

For example:

```
vconfig add eth4 100
```

3. Configure an IP address on the VLAN interface by typing

```
ifconfig eth<x>.<vlan id> x.x.x.x up
```

For example:

```
ifconfig eth4.100 11.192.168.2 netmask 255.255.255.0 up
```

4. Verify the configuration.

From the server, type

```
ibv_rc_pingpong -g 1 -d ocrdma0
```

Note: `-g 1` corresponds to the GID index for the first VLAN.

From the client, type

```
ibv_rc_pingpong -g 1 -d ocrdma0 11.192.168.1
```

Notes:

- If QoS is set to use PFC, the interfaces must be configured with VLANs.
- RoCE PFC works well if VLANs other than 0 and 1 are used.

MTU Configuration

A maximum MTU of 4200 is allowed for RoCE; therefore, the corresponding Ethernet interface for a given RoCE interface should be configured for a MTU of greater than 4200. See “Interface Configuration” on page 11 to find the Ethernet interface for a given RoCE interface.

For example:

```
ifconfig eth4 mtu 4200
```

The following procedure must be performed when changing the MTU on a network interface or a switch port:

1. Unmount all the existing NFS RDMA mounts.
2. Change MTU (on host interface and the switch).
3. Remount the NFS RDMA mounts.

Advanced Applications

NFS over RDMA

Notes:

- When enabling NFS over RDMA on RHEL 6.4 systems, refer to the note on page 8 in the “Installing OFED” section.
- On NFS servers where ‘fsid’ needs to be specified while exporting file systems, ensure that they are unique for each exported file system.
- The usage of memory regions (MRs) by the NFS over RDMA service is high by default (32). It has been found that there is no performance loss by reducing the MR usage to a lower level. This gives flexibility to the number of mounts that can be allowed over a single port.

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
```

See Table 3-1 on page 16 for additional values that can be used.

- The NFS over RDMA service available with OFED 3.5-1 and OFED 3.5 has been known to fail in certain error conditions. It is recommended to use Emulex packaged NFS RDMA modules until future OFED releases address the issue.

Server Configuration

1. Load and configure the ocrdma driver/library.
2. Load the RDMA transport module by typing

```
#modprobe svcrdma
```

3. Start the NFS server by typing
#/etc/init.d/nfs start
(for RHEL)

-or-

```
#service nfsserver start  
(for SLES)
```

4. Configure the server Listen port number for RDMA transport:

```
#echo "rdma 20050">/proc/fs/nfsd/portlist
```

5. Configure the /etc/exports file by adding the required entries:

```
#echo "<path>* (rw,fsid=0,insecure,no_subtree_check,async,  
no_root_squash)">>/etc/exports
```

For example:

```
#echo"/export*(rw,insecure,no_subtree_check,async,  
no_root_squash)">>/etc/exports  
cat/etc/exports  
/export*(rw,insecure,no_subtree_check,async,no_root_squash)
```

6. Export the file system configured in /etc/exports:

```
#exportfs -a
```

Client Configuration

1. Load and configure the ocrdma driver/library.
2. Load the RDMA client module:

```
#modprobe xprtrdma
```

3. List the file system exported by the NFS server:

```
#showmount -e <server_roce_ip>
```

Note: <server_roce_ip> is the NIC IP address of the corresponding RoCE interface on the NFS Server.

For example:

```
#showmount -e 11.192.168.1
```

Export list for 11.192.168.1:

```
/export (everyone)
```

4. Mount the file system:

```
#mount -t nfs4 <server_roce_ip>:<path> -o rdma,port=20050 <mount  
point>
```

For example:

```
#mount -t nfs4 11.192.168.1:/ -o rdma,port=20050 /mnt
```

5. Verify the NFS mount using RDMA:

```
#cat /proc/mounts | grep <mount point>
```

Note: The "Proto" field should be rdma.

Enabling RDMA on all Four Ports on an OCe14000-Series Adapter

RDMA cannot run on all four ports on the OCe14000-series 4-port adapter due to insufficient MRs. To enable RDMA on all four ports, decrease the `sunrpc.rdma_slot_table_entries` parameter from the default value of 32 by performing the following steps:

1. Determine the total number of client-server connections desired per adapter. A connection is defined as a link between a single client and a server, and is not considered an NSF mount. There can be any number of mounts on a single connection.
2. Using Table 3-1, look up the appropriate value for `sunrpc.rdma_slot_table_entries`.

Table 3-1 Setting `sunrpc.rdma_slot_table_entries`

Max Connections Per Adapter	<code>sunrpc.rdma_slot_table_entries</code>
3	32
4	31
5	24
6	20
7	17
8	15
9	13
10	12
11	11
12	10
13	9
15	8
17	7
20	6
24	5
31	4
41	3
62	2

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

3. To change the value of `sunrpc.rdma_slot_table_entries`, add the following line to the `/etc/rc.local` file:

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

(where N is the desired value)

4. Reboot the system for the new settings to take effect.

Note: The total number of connections is dependent on the number of MRs used per connection, and the number of MRs used per connection is dependent on the value of `sunrpc.rdma_slot_table_entries` within the Linux kernel.

The parameter `sunrpc.rdma_slot_table_entries` can have a value between 2 and 32 with the default being 32. If the default of 32 is used, the total number of client to server connections is only three for the entire adapter, which is why all four ports of a 4-port adapter cannot be connected by default.

Advanced Configuration

This section describes the configuration and behavior aspects of RoCE QoS on the OCe14000-series adapters.

QoS Behavior

- Supported:
 - Limited QOS configuration via the OneCommand Manager
 - A single traffic class group for RoCE per port
 - A single RoCE priority in PFC mode
 - Bandwidth allocation for priority groups
- Not supported:
 - RoCE + Multichannel is not supported

OCe14000-Series Defaults

Note: If generic pause is used, ensure that switches have the proper support for this feature. Use PFC with priority 5 if the switch does not support generic pause.

- Adapter boot time
 - PFC is disabled on all the ports at adapter boot time in the NIC+RoCE profile.
 - Generic pause is enabled on all the ports at adapter boot time in the NIC + RoCE profile.
- Back-to-back connection (OCe14000 to OCe14000):
 - PFC is disabled by default.
 - Generic pause is enabled on that port.
- DCBX enabled switch connection
 - When the OCe14000-series adapter is connected to a DCBX-enabled switch, it will shift the mode from generic pause to PFC mode.
 - The OCe14000-series adapter configures RoCE traffic for priority 5.
 - Manually enables priority 5 on a switch under a priority group other than a FCoE/ISCSI/NIC priority group.
 - In the absence of priority 5 at the switch side, the OCe14000-series adapter continues to be configured for PFC mode for priority 5. This can result in packet losses, unrecoverable errors, or infinite retries for RoCE traffic.
- DCBX disabled switch connection
 - When the OCe14000-series adapter is connected to a DCBX-disabled switch, it will be in generic pause mode.

QoS Configuration Guidelines

Priority Groups

It is advisable to split traffic into two or more priority groups:

- one priority group for RoCE
- other groups for non-RoCE traffic

Many RoCE applications use TCP/IP for out-of-band connection establishment. Therefore, it is advisable to allocate sufficient bandwidth to non-RoCE priority groups.

L2 Flow Control

When a port is running in generic pause mode, RoCE latencies can be adversely affected. In this situation, it is advisable to configure RoCE to use PFC for better results.

For switches and adapters that do not support PFC, RoCE can continue to operate in generic pause mode. Bandwidth allocation can be still done for RoCE versus NIC traffic. However, this allocation cannot be guaranteed, since all of the outgoing traffic can be paused in case of congestion.

DCBX Enabled Switch

Switch Configuration

At this time, none of the known switch vendors (for example, Arista, Brocade, Cisco, and Juniper) allow configuring priority for RoCE specific traffic. Priority 5 must be manually enabled on the switch under a priority group other than the FCoE/iSCSI/NIC priority group.

Note: In the absence of priority 5 at the switch side, the OCe14000-series adapter will continue to be configured for PFC mode for priority 5. This can result in packet losses, unrecoverable errors, or infinite retries for RoCE traffic.

Perform the following steps to configure the switch:

1. Create a priority group 1 (PG 1) for RoCE traffic.
2. Assign priority 5 to PG 1.
3. Assign the appropriate bandwidth (for example, 90%) to PG 1.
4. Create PG 2 (or something different from PG 1).
5. Assign NIC traffic to PG 2.
6. Assign the remaining bandwidth to PG 2 (for example, 10%).
7. Enable PFC on the switch ports.
8. Set both switch ports to pass all VLAN traffic.

Note: Some switches have jumbo frame size support disabled by default on the port and/or global level. Enable jumbo frame support, or set MTU to at least 4200.

Host Configuration

1. Enable PFC using the OneCommand Manager (see the *OneCommand Manager Application User Manual*).
2. Create a VLAN.
3. Assign an appropriate IP address to the VLAN interface.

Example Switch PFC Configuration

Note: This example is for a Cisco switch that is connected to the OCe14000-series adapter.

By default, Cisco is configured with two priority groups enabled, which are fixed and cannot be deleted:

- Default group name: default-group
- FCoE group name: fcoe-group

Perform the following steps:

1. Use default-group as the non-RoCE priority group (see step 4 on page 19).
2. Create another group for priority 5, such as PG 5 (see step 1 on page 19).
3. Set 90% bandwidth to the PG 5 group and 10% to the default-group. No other changes are required to the default-group or fcoe-group.

The following are example switch PFC configurations:

Cisco Global QoS Configuration

Global QoS configuration on the Cisco Switch

```
class-map type qos roce
  match qos 5
class-map type queuing roce
  match qos-group 5
class-map type network-qos roce
  match qos-group 5
policy-map type qos roce
  class roce
    set qos-group 5
  class class-fcoe
    set qos-group 1
  class class-default
policy-map type queuing roce
  class type queuing roce
    bandwidth percent 90
  class type queuing class-fcoe
    bandwidth percent 0
  class type queuing class-default
    bandwidth percent 10
```

```
policy-map type network-qos roce
  class type network-qos roce
  pause no-drop
  mtu 4200
  class type network-qos class-default
  mtu 9216
class type network-qos class-fcoe
  pause no-drop
  mtu 2158
system qos
  service-policy type qos input roce
  service-policy type queuing input roce
  service-policy type queuing output roce
  service-policy type network-qos roce
```

DCBX Disabled Switch Connection (generic pause mode)

1. Host Configuration:

On the host and peer systems, ensure that Tx and Rx pause flow control is enabled using the operating system standard tools on all of the ports/interfaces which are RoCE enabled.

a. To verify status:

```
ethtool -a ethX
```

b. To configure:

```
ethtool -A ethX [ autoneg on|off ]
[ rx on|off ]
[ tx on|off ]
```

2. Switch configuration:

a. Enable Tx and Rx generic pause flow control on each port.

b. Some switches have jumbo frame size support disabled by default on the port and/or global level. Enable jumbo frame support, or set MTU to at least 4200.

Updating the Adapter Firmware for RoCE

Note: This section can be skipped if the adapter already has the required firmware version and RoCE profile.

Determine Firmware Version

The adapter firmware should be upgraded to version 10.2.x.x. The upgrade can be done manually or by using the OneCommand Manager application.

1. Determine if the firmware needs to be updated. List the ocrdma interfaces by typing

```
#ibdev2netdev
```

2. Each ocrdma interface maps to an Ethernet interface ethX. Use 'ethtool -i ethX' to determine the firmware version on the adapter.

If the reported firmware version does not match the version listed above, then the firmware for that adapter needs to be updated.

The latest firmware can be downloaded from the Emulex website. The name of the firmware file will have a format of

```
oc14-x.x.x.x.ufi
```

oc14 in the filename refers to the OCe14000-series NIC adapter, and x.x.x.x refers to the version.

Updating the Firmware Manually

Note: These steps need to be done only once for each adapter.

Perform the following steps to update the firmware:

1. The firmware download command needs to be invoked once for each adapter by specifying any Ethernet interface (ethX) configured on the adapter.
2. List the ocrdma interfaces:

```
#ibdev2netdev
```

3. Update the firmware for ethX:

```
# copy oc14-x.x.x.x.ufi
```

to

```
/lib/firmware
```

```
# ethtool -f ethX oc14***.ufi
```

Note: Multiple ocrdma interfaces may be present on a adapter, so by using the mapping of ocrdma interfaces to Ethernet interfaces (through ibdev2netdev) and 'ethtool -i ethX', only one Ethernet interface per adapter may be selected.

Updating the Firmware using OneCommand Manger

If the OneCommand Manager application is used to the update the firmware, refer to the appropriate section of the *OneCommand Manager Application User Manual* for the procedure.

4. Troubleshooting

Log Messages

RoCE Error Log Messages and their Descriptions

Table 4-1 lists RoCE error log messages and their descriptions.

Table 4-1 RoCE error log messages and their descriptions.

RoCE Log Message	Description	Type
<pci bus info> <hca_name>: <speed> " <model_number> port <port_num>	Prints the information about the RoCE PCI function. For example: 0000:04:00.1 Emulex OneConnect RoCE HCA: 10Gbps "OneConnect OCe14000" port 1	Information
<pci bus info> ocrdma<D> driver loaded successfully	Driver loaded successfully on the device.	Information
ocrdma_add() leaving. ret=<D>	Adding device failed with Error = D.	Error
ocrdma_dispatch_ibevent () unknown type=0x<D>	Received unknown event from the hardware.	Error
ocrdma_dispatch_ibevent: Fatal event received	Device reported a fatal event.	Error
ocrdma_process_acqe (<D>) invalid evt code=0x<DD>	Invalid event code <DD> reported on the device <D>.	Error
ocrdma_process_mcqe () cqe for invalid tag0x<D> expected=0x<DD>	Invalid completion tag reported.	Error
ocrdma_wait_mqe_cmpl (<D>) mailbox timeout: fw not responding	Mailbox failed because of timeout.	Error
ocrdma_mbx_cmd() cqe_status=0x<D>, ext_status=0x<DD>	Completion and Extended status in case of mailbox errors.	Error
opcode=0x<D>, subsystem=0x<DD>	Opcode and subsystem ids of the failed mailbox commands.	Error
ocrdma_irq_handler(): Fatal Error, EQ full eq_id = 0x<D>, eqe = 0x<DD>	EQ Full detected.	Error
ocrdma_mbx_create_cq (<D>) max_cqe=0x<DD>, requester_cqe=0x<DDD>	Requesting more CQ entries than what the device supports.	Error
ocrdma_set_create_qp_sq_cmd: () req. max_send_wr=0x<D>	The total number of SQ entries requested is more than what the device supports.	Error
ocrdma_set_create_qp_rq_cmd() req. max_rcv_wr=0x<D>	The total number of RQ entries requested is more than what the device supports.	Error

Table 4-1 RoCE error log messages and their descriptions. (Continued)

RoCE Log Message	Description	Type
ocrdma_mbx_create_srq() req. max_wr=0x<D>	The total number of SRQ entries requested is more than what the device supports	Error
ocrdma_mbx_create_qp(<D> rq_err	Failed to create qp.	Error
ocrdma_mbx_create_qp(<D> sq_err	Failed to create qp.	Error
ocrdma_parse_dcbxcfg_rsp(): DCBX state is disabled.	DCBX state is disabled in the adapter.	Information
ocrdma_parse_dcbxcfg_rsp(): pfc is disabled.	PFC is disabled in the adapter.	Information
ocrdma_init_service_level(): status=<D>	Failed to get the DCBX configuration from the adapter.	Error
ocrdma is using default service level	Using the default service level.	Information
ocrdma_init_hw() status=<D>	Initialization of the hardware failed.	Error
ocrdma_alloc_resources(<D> error	Failed to allocate driver resources.	Error
Unable to allocate ib device	ib_alloc_device failed.	Error
ocrdma_add_stat: No space in stats buff	Response data for debugfs request is more than the stats buff size.	Error
ocrdma_alloc_stats_mem: stats mbox allocation failed	Failed to allocate memory for statistics command.	Error
ocrdma_alloc_stats_mem: stats debugfs mem allocation failed	Failed to allocate memory for debugfs.	Error
ocrdma_update_stats: stats mbox failed with status = <D>	The Statistics command failed from the hardware.	Error
ocrdma_query_port(<D> invalid_port=0x<DD>	Invalid port specified in Query port.	Error
ocrdma_modify_port(D) invalid_port=0x<DD>	Invalid port specified in Modify port.	Error
ocrdma_dealloc_ucontext_pd(D) Freeing in use pdid=0x<DD>	Dealloc ucontext requested on a PD which is in use.	Information
ocrdma_get_dma_mr err, invalid access rights	Invalid access rights while allocating lkey.	Error
ocrdma_dereg_mr(<D>) fw not responding	De-registration of MR failed because firmware is not responding.	Information
ocrdma_copy_cq_uresp(<D>) copy error cqid=0x<DD>	Failed to copy cq create response.	Error

Table 4-1 RoCE error log messages and their descriptions. (Continued)

RoCE Log Message	Description	Type
ocrdma_check_qp_params(<D>) unsupported qp type=0x<DD> requested	Validation checks during QP create.	Error
ocrdma_check_qp_params(<D>) unsupported send_wr=0x<DD> requested ocrdma_check_qp_params(<D>) supported send_wr=0x<DD>	Validation checks during QP create.	Error
ocrdma_check_qp_params(<D>) unsupported recv_wr=0x<DD> requested ocrdma_check_qp_params(<D>) supported recv_wr=0x<DD>	Validation checks during QP create.	Error
ocrdma_check_qp_params(<D>) unsupported inline data size=0x<DD> requested ocrdma_check_qp_params(<D>) supported inline data size=0x<DD>	Validation checks during QP create.	Error
ocrdma_check_qp_params(<D>) unsupported send_sge=0x<DD> requested ocrdma_check_qp_params(<D>) supported send_sge=0x<DD>	Validation checks during QP create.	Error
ocrdma_check_qp_params(<D>) unsupported recv_sge=0x<DD> requested ocrdma_check_qp_params(<D>) supported recv_sge=0x<DD>	Validation checks during QP create.	Error
ocrdma_check_qp_params(<D>) Userspace can't create special QPs of type=0x<DD>	Validation checks during QP create.	Error
ocrdma_check_qp_params(<D>) GSI special QPs already created	Validation checks during QP create.	Error
ocrdma_check_qp_params(<D>) Consumer QP cannot use GSI CQs	Consumer QPs should not use the CQ of GSI QP.	Error
ocrdma_copy_qp_uresp(<D>) user copy error	Failed to copy the QP create response back to the user.	Error
ocrdma_create_qp(<D>) error=<DD>	Failed QP create command.	Error
ocrdma_modify_qp(<D>) invalid attribute mask=0x<M> specified for qpn=0x<QP> of type=0x<T> old_qps=0x<OS>, new_qps=0x<NS>	Parameter error while trying to modify the QP.	Error

Table 4-1 RoCE error log messages and their descriptions. (Continued)

RoCE Log Message	Description	Type
<code>ocrdma_build_inline_sges() supported_len=0x<L>,unspported len req=0x<LL></code>	Failed to build inline SGEs.	Error
<code>ocrdma_update_wc() invalid opcode received = 0x%x</code>	Invalid opcode received from the hardware completion.	Information
<code>ocrdma_reg_mr() status=<D></code>	Failed to register MR.	Error
<code>ocrdma_resolve_dmac () fail to resolve mac_addr</code>	Failed to resolve MAC address.	Error