



Emulex Drivers for Windows User Manual

FC and FCoE version 2.72.012.001

NIC version 4.2.390.6

iSCSI version 4.2.281.0

Copyright © 2012 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

Table of Contents

List of Figures	8
List of Tables	9
1. Installation.....	10
Driver Information	10
Compatibility	10
Operating System Requirements	10
Driver Installation Options	11
OneInstall Installer	11
Loading the OneInstall Package	12
AutoPilot Installer	12
Special Consideration When Updating the Emulex Storport Miniport FC/FCoE Driver	13
Driver Kit Installer Overview	14
Installing the Driver Kit	14
AutoPilot Installer Overview	15
Starting Installers from a Command Prompt or Script	15
Running a Software Installation Interactively	15
Option 1: Automatically Run the AutoPilot Installer	16
Option 2: Run the AutoPilot Installer Separately	16
Hardware-First Installation or Driver Update.....	17
Software-First Installation.....	18
Text-Only Driver Installation.....	18
Unattended Driver Installation	19
Option 1: Install the Driver Silently	19
Option 2: Run the Driver Kit Installer Separately	19
Installation Failure	20
Manually Installing or Updating the Emulex Protocol Drivers	20
Removing the Adjunct Driver Registry Key	20
Installing the Emulex PLUS (EIXPlus) Driver for the First Time	21
Updating the Emulex PLUS (EIXPlus) Driver.....	21
Installing or Updating the FC/FCoE Storport Miniport Driver	22
Installing or Updating the iSCSI Driver.....	23
Installing or Updating the NIC Driver	24
Windows Server 2008.....	24
Windows Server 2012.....	25
Removing Emulex Driver Kits and Drivers.....	26

Uninstalling Emulex Driver Kits	26
Windows Server 2008	26
Windows Server 2012	27
Uninstalling the Emulex Drivers	27
Windows Server 2008	27
Windows Server 2012	28
2. Configuration	30
FC/FCoE Driver Configuration	30
Configuring FC Driver Parameters	30
Server Performance with FC Drivers	38
I/O Coalescing	38
Performance Testing	38
NIC/TOE Driver Configuration	40
Configuring NIC/TOE Driver Options	40
Configuring Windows Server NIC/TOE Driver Parameters	49
Modifying Advanced Properties	49
Statistics Property Page	50
Using OCCFG for Windows NIC/TOE Driver Options	54
Displaying OCCFG Help	54
Selecting an Adapter	55
Configuring Device Parameters	56
Viewing Device Parameters	57
Resetting All Parameters	57
Displaying All Parameters	57
Using Interactive Mode	58
Parameter Help	58
Using SR-IOV with Emulex Devices	59
Server BIOS Configuration	59
Emulex PXESelect Configuration for SR-IOV	60
SR-IOV Server Validation	60
Verifying the Driver Version	61
Enabling SR-IOV in the Emulex Device	63
Hyper-V	64
Verifying SR-IOV	65
Network Driver Performance Tuning	66
Optimizing Server Hardware and BIOS Configuration	66
Windows Server Network Driver	67
Checksum Offloading and Large Send Offloading (LSO)	70
Receive Side Scaling (RSS) for Non-Offloaded IP/TCP Network Traffic	71
TCP Offloading (TOE)	72

Receive Window Auto Tuning and Compound TCP	75
Interrupt Coalescing	75
CPU Binding Considerations	76
Single TCP Connection Performance Settings	76
iSCSI Driver Configuration	77
Configuring iSCSI Driver Options	77
Backing Up and Editing the Registry	77
Interrupt Moderation Policy Settings	79
Creating Non-Bootable Targets	80
Using the Microsoft iSCSI Initiator Service	80
Logging into a Target Using the Microsoft Software Initiator	80
Windows Multipath I/O Support	81
Multipath Support on Windows Server 2008, 2008 R2, and 2012	81
Logging into Targets for Multipath Support	82
Maximum Transmission Unit (MTU) for iSCSI Connections	82
iSCSI Error Handling	82
Configuring LDTO and ETO on Windows Server	83
Error Handling Under MultiPath (MPIO) and Cluster Configurations	84
3. Troubleshooting	85
General Troubleshooting	85
Troubleshooting the FC/FCoE Driver	85
Troubleshooting the Cisco Nexus Switch Configuration	85
Event Trace Messages	86
ELS Log Messages (0100-0130)	86
Discovery Log Messages (0202-0262)	89
Mailbox Log Messages (0310-0326)	92
INIT Log Messages (0400-0463)	93
FCP Log Messages (0701-0749)	95
Link Log Messages (1302-1306)	98
Tag Messages (1400-1401)	99
NPIV Messages (1800-1899)	100
ELS Messages (1900-1999)	101
Troubleshooting the NIC/TOE Drivers	102
Monitoring TCP Offloads	103
TCP Offload Failure	104
Troubleshooting the iSCSI Driver	105
Appendix A. Error and Event Log Information	108
FC/FCoE Error and Event Logs	108
Viewing the FC/FCoE Error Log	108

Severity Scheme	109
Related Driver Parameter: LogError	109
Format of an Error Log Entry.....	109
Error Codes Tables.....	110
Viewing the FC/FCoE Event Log	115
Event Log Interpretation.....	115
Additional Event Log Information	115
ASC/ASCQ	117
Additional Notes on Selected Error Codes	118
NIC/TOE Error and Event Logs	119
Viewing the NIC/TOE Error Log.....	119
NIC/TOE Event Log	119
iSCSI Error and Event Log.....	123
Viewing the iSCSI Error and Event Log on Windows Server 2008	123
iSCSI Error Log on Windows Server 2008	124
Viewing the iSCSI Error Log on Windows Server 2008 R2 and 2012.....	128
Appendix B. Configuring iSCSI through DHCP	133
Dynamic Host Configuration Protocol (DHCP) Recommendations	133
Vendor-Specific Option 43.....	133
Format of Vendor-Specific Option 43.....	133
Description of Mandatory and Optional Parameters.....	134
Examples	134
Appendix C. Port Speed Specifications	136
Negotiating Speed on a Mezzanine Card.....	136
Appendix D. AutoPilot Installer Command Line and Configuration File Parameters..	137
AParg Driver Kit Parameter and Appending to the APInstall.exe File	137
AutoPilot Installer Syntax.....	138
Path Specifiers	139
Configuration File Location	139
Software Configuration Parameters.....	139
DiagEnable (Running Diagnostics)	139
ForceDriverTypeChange (Forcing a Driver Type Change)	139
ForceDriverUpdate (Forcing a Driver Version Update)	140
ForceRegUpdate (Forcing an Update of an Existing Driver Parameter Value).....	140
LocalDriverLocation (Specifying Location to Search for Drivers)	140
NoSoftwareFirstInstalls (Prohibiting Software First Installations)	141
ReportLocation (Setting Up an Installation Report Title and Location)	141

SilentInstallEnable (Enabling Unattended Installation)	141
SilentRebootEnable (Enabling Silent Reboot)	141
InstallWithoutQFE (Enabling Installation if a QFE Check Fails)	142
AutoPilot Configuration File	142
Using the Windows Environment Variable (%ProgramFiles%)	142
Configuration Identification [AUTOPILOT.ID]	143
Software Configuration [AUTOPILOT.CONFIG]	143
Configuration Prompts/Vendor-Specific Questions [STORPORT.CONFIGURATION] ..	143
QFE Checks [STORPORT.QFES]	144
Setting Up FC Driver Parameters [STORPORT.PARAMS]	145
Setting Up System Parameters [SYSTEM.PARAMS]	145
AutoPilot Installer Exit Codes	146
AutoPilot Installer Installation Reports	147
Command Script Example	147

List of Figures

Figure 1-1	AutoPilot Installer Warning (Software-First Installation)	18
Figure 2-1	Partial View of Windows Device Manager	49
Figure 2-2	NIC/TOE Advanced Properties in Windows Server 2008	50
Figure 2-3	NIC/TOE Statistics Properties in Windows Server 2008	51
Figure 2-4	Device Manager for Windows Server 2012	62
Figure 2-5	Emulex OneConnect NIC/TOE Driver Properties page	63
Figure 2-6	Emulex OneConnect NIC/TOE Advanced Properties page	64
Figure 2-7	Emulex OneConnect NIC Statistics Properties page	66
Figure A-1	Event Properties	108
Figure A-2	iSCSI Error	124

List of Tables

Table 2-1	Storport Miniport Driver Parameters	31
Table 2-2	Recommended Settings for I/O Coalescing	38
Table 2-3	Windows Server 2008, 2008 R2, and 2012 NIC/TOE Driver Options	41
Table 2-4	NIC/TOE Driver Properties Statistics	51
Table 2-5	Windows Server Performance Tuning Situations	67
Table 2-6	Statistics and Fine Tuning	68
Table 2-7	iSCSI Driver Options.....	78
Table 2-8	im_policy Settings.....	79
Table 2-9	LDTO and ETO Information on Windows Server	83
Table 3-1	General Troubleshooting.....	85
Table 3-2	Cisco Nexus Switch Situations.....	85
Table 3-3	Troubleshooting the NIC/TOE Drivers.....	102
Table 3-4	Troubleshooting TCP Offload Failures	104
Table 3-5	Troubleshooting the iSCSI Driver	105
Table A-1	Severe Errors.....	110
Table A-2	Malfunction Errors	112
Table A-3	Command Errors.....	113
Table A-4	Event Indications	114
Table A-5	ELS/FCP Command Error Status Codes.....	115
Table A-6	CT Command Response Codes.....	115
Table A-7	FC-CT Reject Reason Codes	116
Table A-8	ELS Command Codes.....	116
Table A-9	SCSI Status Codes	116
Table A-10	Local Reject Status Codes	117
Table A-11	SRB Status Codes	117
Table A-12	NIC/TOE Event Log Entries	119
Table A-13	iSCSI Error Log Entries on Windows Server 2008	124
Table A-14	iSCSI Error Log Entries on Windows Server 2008 R2 and 2012.....	128
Table B-1	Data String Parameters for Option 43.....	134
Table C-1	Negotiated Speed Specification per Adapter Port Connection.....	136
Table D-1	Unattended Installation Error Codes	146

1. Installation

Driver Information

This product supports the Emulex® OneConnect™ families of universal converged network adapters (UCNAs) and Emulex LightPulse® family of host bus adapters (HBAs).

The Windows drivers support the Fibre Channel (FC), FC over Ethernet (FCoE), Ethernet (NIC), including TCP Offload Engine (TOE), and Internet Small Computer System Interface (iSCSI) protocols.

This document explains how to install the Windows drivers on your system and configure the drivers' capabilities based on the supported networking protocols:

- FC and FCoE
 - Configuring the FC/FCoE driver parameters
 - Improving server performance with FC/FCoE drivers
- Ethernet and TOE
 - Configuring NIC/TOE driver options
 - Tuning network driver performance
- iSCSI
 - Configuring iSCSI driver options
 - Creating non-bootable targets
 - Configuring Multipath I/O

Compatibility

For a list of adapters that are compatible with this driver, see the driver's Downloads page on the Emulex website. For compatible firmware versions, see the Downloads page for the specific adapter.

Operating System Requirements

One of the following operating systems must be installed:

- Windows Server 2008 running on an x86 or x64 server
- Server Core installation option of Windows Server 2008 running on an x86 or x64 server
- Windows Server 2008 R2 running on an x64 server
- Windows Server 2008 R2 SP1
- Windows Server 2012
- Windows 7 Ultimate, Enterprise, or Professional edition (x64 only)

Note: Windows 7 is only supported on OCe11102 UCNAs.

- Windows Pre-installation Environment (PE) 2.x, 3.x, or 4.x

Note: AutoPilot Installer is not supported for use with Windows PE. See the Windows Automated Installation Kit (WAIK) documentation for instructions on including drivers with your Windows PE images.

- Minimum 512 MB memory

Note: Check the Emulex website for required updates to the Windows operating system or the Emulex drivers.

Driver Installation Options

There are two ways that you can install the Windows drivers: the OneInstall Installer (OneInstall) package or AutoPilot Installer®. You must choose one of the options.

- OneInstall contains all of the Emulex® Windows drivers (Emulex Storport Miniport and Network Driver Interface Specification [NDIS] Miniport drivers) and the OneCommand Manager application utility in a single download package.
- AutoPilot Installer for Emulex® Storport Miniport and NDIS Miniport drivers provides installation options from simple installations with a few mouse clicks to unattended installations that use predefined script files and text-only installations.

OneInstall Installer

The OneInstall Package is an executable file that self-extracts and copies the following software onto your system:

- All compatible protocol drivers to manage communication and data transfer between applications and I/O devices, using Emulex adapters as agents. There are four separate drivers for Emulex adapters:
 - FC
 - FCoE
 - iSCSI
 - Ethernet
- ElxPlus driver – Provides remote management support in the OneCommand Manager application.
- OneCommand Manager application for the Emulex adapters

OneInstall is available for download from the Emulex website. The limitation of the single package is that you cannot perform pre-installation tasks, unattended installations, or text-only installations. For these tasks, use the AutoPilot Installer kits.

Loading the OneInstall Package

To install the drivers using the OneInstall Package:

1. Navigate to the location of the OneInstall package that you downloaded to your hard drive.
2. Double-click the OneInstall package. The Welcome screen appears.
3. Click **Next**. The Installation options screen appears.
4. Select the drivers and application that you want to install. Click **Next**.
5. A progress screen appears while the OneInstall installer loads the selected drivers and applications.
6. When the drivers and application software are loaded, an Installation completed screen appears. Click **Finish**.

AutoPilot Installer

AutoPilot Installer is included with the Emulex drivers and utilities in Windows executable files (kit installers) that you can download from the Emulex website.

Each driver kit contains x86, x64, and IA64 FC/FCoE drivers and installs the correct files on a per-system basis. Run one of the following driver kit installers to extract the software needed for the driver installation.

- Fibre Channel (FC) driver package (elxdrv-fc-<version>.exe)
- Fibre Channel over Ethernet (FCoE) driver package (elxdrv-fcoe-<version>.exe)
- Internet Small Computer System Interface (iSCSI) driver package (elxdrv-iscsi-<version>.exe)
- Network Interface Card (NIC) driver package (elxdrv-nic-<version>.exe)

Note: Updating the NIC protocol driver may temporarily disrupt operation of any NIC teams configured on the system.

You can install a driver by any of the following methods:

Note: These installation methods are not mutually exclusive.

- **Hardware-first installation** – At least one Emulex adapter must be installed before you can install the Emulex drivers and utilities.
- **Software-first installation** – You can install drivers and utilities using AutoPilot Installer prior to the installation of any adapters. You do not need to specify the adapter models to be installed later. The appropriate drivers and utilities automatically load when you install the adapters.
- **Utility-Only installation** – If the drivers in the driver kit share the same version with those already installed on the system, you can reinstall or update the previously installed utility without reinstalling the drivers.
- **Text-Only installation** – Text-based installation mode is used automatically when AutoPilot Installer is run on a Server Core system.

- **Network installation** – You can place the driver kit installers on a shared network drive and install them across your local area network (LAN). Network-based installation is often used in conjunction with unattended installation and scripting. This allows you to configure and install the same driver version on all the hosts in a storage area network (SAN).
- **Unattended installation** – You can run the driver kit installers and AutoPilot installer with no user interaction from a command line or script. Unattended installation works for both hardware-first and software-first installations and all driver kits. An unattended installation operates in silent mode (sometimes referred to as quiet mode) and creates an extensive report file with installation status.

Note: Complete driver and utilities documentation may be downloaded from the Emulex website (www.emulex.com).

Special Consideration When Updating the Emulex Storport Miniport FC/FCoE Driver

Any pre-existing SCSIport Miniport or FC/FCoE Port drivers are replaced (overwritten) when you install the Emulex Storport Miniport FC/FCoE driver. The Emulex Storport Miniport FC/FCoE driver is installed for every adapter on your server. Installing the Emulex Storport Miniport FC/FCoE driver also deletes all customized driver parameters, persistent bindings, logical unit number (LUN) masking, and LUN mapping. Emulex recommends that you make a note of your current settings before you install the Emulex Storport Miniport FC/FCoE driver so that you can update your customized driver parameters. Or, you can use the AutoPilot Installer default parameters.

Note: This update process can leave files associated with older driver kits on your system. After running AutoPilot Installer and restarting your system, you may be required to uninstall old driver or utility program entries. Follow the procedures in “Uninstalling Emulex Driver Kits” on page 26. Ensure you delete only previous driver versions.

Note: If you are installing this driver as an update to the WS2012 driver, some parameter defaults are different from the inbox driver. Emulex recommends that, after you install the Emulex driver, you select “reset to default” on the Device Manager property page. This returns all adapter and driver settings to the default values listed in this manual.

Driver Kit Installer Overview

The driver kit installer is an executable file that self-extracts and copies the following software onto your system:

- FC/FCoE protocol drivers – Manage communication and data transfer between applications and I/O devices, using Emulex adapters as agents. There are four separate drivers for Emulex adapters:
 - FC
 - FCoE
 - iSCSI
 - Ethernet
- ElxPlus driver – Provides remote management support in the OneCommand Manager application.

If you are creating Windows PE disk images or performing boot from SAN installation, you must select the **Unpack All Drivers** option instead of running AutoPilot Installer. See step 2 of Installing the Driver Kit in the following procedure.

After the driver kit is installed, you have two options:

- Automatically run AutoPilot Installer immediately. This is the default option. See Option 1 on page 16 for more information.
- Run AutoPilot Installer later on your system. See Option 2 on page 16 for more information.

Installing the Driver Kit

To install the driver kit:

1. Do one of the following:
 - Download and install the single package from the Emulex website to your system.
 - Download the driver kit installer from the Emulex website to your system.

Note: The name of the driver kit installer depends on the current version identifier.
2. If you plan to create Windows Pre-installation Environment (PE) disk images or perform boot from SAN installations, unpack all drivers. On the driver kit installer Welcome page, select the **Unpack all drivers** checkbox. Checking this checkbox unpacks all the drivers into the 'My Documents' directory.
3. Choose one of three options:
 - For an automatic interactive installation, see "Option 1: Automatically Run the AutoPilot Installer" on page 16.
 - For control of all interactive installation settings, see "Option 2: Run the AutoPilot Installer Separately" on page 16.
 - For an unattended installation with AutoPilot Installer, see "Unattended Driver Installation" on page 19.

AutoPilot Installer Overview

AutoPilot Installer runs after the driver kit and the OneCommand Manager application are installed. AutoPilot Installer can be installed:

- Immediately after the driver kit installation has completed
- At a later time using an interactive installation
- Through an unattended installation

AutoPilot Installer provides

- Command line functionality – Initiates an installation from a command prompt or script. Configuration settings can be specified in the command line.
- Compatibility verification – Verifies that the driver to be installed is compatible with the operating system and platform.
- Driver installation and update – Installs and updates drivers.
- Multiple adapter installation capability – Installs drivers on multiple adapters, alleviating the need to manually install the same driver on all adapters in the system.
- Driver diagnostics – Determines whether the driver is operating properly.
- Silent installation mode – Suppresses all screen output (necessary for unattended installation).

Note: AutoPilot Installer does not allow you to install the driver if the minimum Windows service pack or Microsoft Storport driver update is not installed.

Starting Installers from a Command Prompt or Script

When a driver kit installer or AutoPilot Installer is run from a command prompt or command script (batch file), the Windows command processor does not wait for the installer to run to completion. As a result, you cannot check the exit code of the installer before the next command is executed. Emulex recommends that for command line invocation, always use the “start” command with the “/wait” option. This causes the command processor to wait for the installer to finish before it continues.

For more information on command line installation and configuration parameters, see appendix D, “AutoPilot Installer Command Line and Configuration File Parameters,” on page 137.

Running a Software Installation Interactively

There are two options when performing an installation interactively.

- Option 1 allows you to automatically run the AutoPilot Installer, which completes a driver kit installation and AutoPilot installation with a few mouse clicks.
- Option 2 allows you to run the AutoPilot Installer separately. This option is recommended when:

- Changing installation settings for a limited number of systems.
- Familiarizing yourself with AutoPilot Installer configuration options.

These options assume you have already downloaded the driver kit installer from the Emulex website.

Option 1: Automatically Run the AutoPilot Installer

Use this option unless you have specific configuration needs.

1. Double-click the driver kit installer or run it from a command line. See appendix D, “AutoPilot Installer Command Line and Configuration File Parameters,” on page 137 for information on the command line options. The command line parameter APargs allows you to specify arguments that are automatically passed to the AutoPilot Installer command. A Welcome page is displayed with driver kit version information and Emulex contact information.
2. Click **Next** to proceed to the Installation Options page.
For each installation option, the default installation location for that option is displayed. Browse to a different location, if desired.
3. Click **Install** to continue the installation.
The Progress dialog box is displayed. After all tasks are completed, the Finish dialog box is displayed. The Start AutoPilot Installer box is automatically selected.
4. Click **Finish**. AutoPilot Installer runs automatically and completes one of the following installations:
 - Hardware-First Installation or Driver and Utility Update (page 17).
 - Software-First Installation (page 18).

Option 2: Run the AutoPilot Installer Separately

To access these options, run AutoPilot Installer after driver kit installation is complete. This allows you to change the configuration options supplied to the AutoPilot Installer (see below).

1. Perform steps 1 through 3 for “Option 1: Automatically Run the AutoPilot Installer”.
2. Clear the **Run AutoPilot Installer** check box on the Finish dialog box.
3. Click **Finish**. The driver kit installer exits.

After the driver kit installation is complete, change the configuration in one of two ways:

- Change the configuration file. See “Software Configuration Parameters” on page 139 for details.
- Supply parameters on the command line. See appendix D, “AutoPilot Installer Command Line and Configuration File Parameters,” on page 137 for details.

Once you have finished this step, you can run AutoPilot Installer at a later time, using either of the following methods: (If you are supplying options via the command line, you must run AutoPilot Installer from the command line.)

- Select **Programs>Emulex>AutoPilot Installer** in the Start menu.
- Run AutoPilot Installer from the command line. Type

`C:\Program Files\Emulex\AutoPilot Installer\APInstall.exe`

Note: The location of APInstaller.exe may differ on your system, depending on your system's Program Files location. You may also specify a different location when you install the driver package.

Hardware-First Installation or Driver Update

The driver kit installer must be downloaded from the Emulex website and installed.

Note: Updating the NIC protocol driver may temporarily disrupt operation of any NIC teams configured on the system.

Note: To update the Emulex protocol drivers, begin the procedure at step 2.

To perform a hardware-first installation:

1. Install a new Emulex adapter and power-on the system. If the Windows Found New Hardware wizard is displayed, click **Cancel** to exit. AutoPilot Installer performs this function.

Note: If there are multiple adapters in the system, the Windows Found New Hardware wizard appears multiple times. Click **Cancel** to exit the wizard each time it appears.

2. Run AutoPilot Installer using one of the two options listed in “Running a Software Installation Interactively” on page 15.
3. When the AutoPilot Installer Welcome page appears, select an adapter in the list and click **Next**. The installation continues.

Consider the following:

- If you are replacing a SCSIport Miniport or FC/FCoE Port driver, an Available Drivers list is displayed. Click **OK** in this window.
 - If you are updating the driver, the existing port settings are used, unless otherwise specified in the configuration file. These settings are pre-selected but can be changed. Set or change settings, then click **Next**.
 - If you are initially installing a vendor-specific version of the Emulex driver installation program, a Driver Configuration page may be displayed. This page includes one or more windows with questions that you must answer before continuing the installation process. In this case, answer each question and click **Next** on each window to continue.
4. Click **Next**. The installation automatically completes. A dialog box is displayed if Windows requires a reboot. Once the installation is successful, the Finish dialog box appears.
 5. View or print a report, if desired.
 - View Installation Report – The installation report is a text file with current Emulex adapter inventory, configuration information and task results.

- Print Installation Report – The Windows print dialog box is displayed to select options for printing the installation report.
6. Click **Finish** to exit AutoPilot Installer. If the system must be rebooted, you are prompted to do so as indicated in step 4; you must reboot before using the drivers or utilities.

Software-First Installation

The driver kit installer must be downloaded from the Emulex website and installed. Either the full or core driver package may be installed; only one can be installed on a system.

To perform a software-first installation:

1. Run AutoPilot Installer using one of the two options listed in “Running a Software Installation Interactively” on page 15. The following message appears:




Figure 1-1 AutoPilot Installer Warning (Software-First Installation)

2. Click **OK**. A Welcome page appears.
3. Click **Next**. The installation automatically progresses. Once the installation is completed successfully, the Finish dialog box appears.
4. View or print a report, if desired.
 - View Installation Report – The installation report is a text file with current Emulex adapter inventory, configuration information and task results.
 - Print Installation Report – The Windows print dialog box is displayed to select options for printing the installation report.
5. Click **Finish** to exit the AutoPilot Installer.

Text-Only Driver Installation

Text-based installation mode is used automatically when the driver kit installer runs on a server with the Server Core installation option of Windows Server. During text-based installations, AutoPilot Installer uses a command prompt window. The driver kit installer notifies you when the driver is installed and also gives you a chance to stop the installation.

Whether it is launched from the command line or started programmatically, Windows always starts AutoPilot Installer as a separate stand-alone task. This means that AutoPilot Installer has its own command prompt window and cannot access others.

Unattended Driver Installation

An unattended driver installation, sometimes referred to as a quiet or silent installation, requires no user input. This is useful for performing an installation remotely from a command script, or when you want to make sure a custom configuration is not changed by a user during installation.

When in unattended installation mode, AutoPilot Installer does the following:

- Reads the configuration file.
- Reads any options that may be specified on the command line, overriding the configuration file settings as appropriate.
- Opens the installation report file.
- Validates the operating system.
- Discovers adapters and records the adapter inventory in the report file.
- Verifies mandatory configuration file parameters.
- Searches for drivers to install based on the LocalDriverLocation setting in the configuration file.
- If appropriate, verifies that the selected driver is a different type than the currently installed driver, or a more recent version of the currently installed driver.
- Copies the driver parameters from the configuration file into the registry for the driver's coinstaller (FC and FCoE drivers only).
- Installs or updates the driver.
- Rediscovered adapters and records the updated adapter inventory in the report file.
- Records the final results and closes the report file.

There are two ways to perform an unattended installation.

- Install the driver silently.
- Run the driver kit installer separately.

Option 1: Install the Driver Silently

Run the driver kit installer from a command prompt or script. Specify the `"/q"` (quiet) command line option. For example:

```
elxdrv-fc-fcoe<version>.exe /q
```

Note: The name of the driver kit installer depends on the current version identifier. For other command line options, see "AutoPilot Installer Command Line and Configuration File Parameters" on page 137.

Option 2: Run the Driver Kit Installer Separately

1. Follow steps 1–3 for "Running a Software Installation Interactively" on page 15.
2. Clear the **Run AutoPilot Installer** check box on the Finish dialog box.
3. Choose one of the following options:

- Run the AutoPilot Installer from a command prompt or script with the silent option:

```
APInstall.exe /silent
```

- Edit the AutoPilot Installer configuration file before running the AutoPilot Installer. The configuration file is typically located in:

```
C:\Program Files\Emulex\AutoPilot Installer\APInstall.cfg
```

Uncomment the line that sets `SilentInstallEnable` to "True". There are other settings in the same section of the configuration file related to unattended installations that you may also want to edit. See "Software Configuration Parameters" on page 139 for more information. After editing the file, you can run the AutoPilot Installer from the Start menu, a command prompt or a script.

Installation Failure

If the installation fails, the Diagnostics window opens showing that the adapter failed.

If the adapter fails:

1. Select the adapter to view the reason for the failure. The reason and suggested corrective action are displayed.
2. Perform the suggested corrective action and run AutoPilot Installer again.

Note: You can run AutoPilot Installer again from the Start menu (**Programs>Emulex>AutoPilot Installer**) or you can run APInstall.exe from a command prompt.

Manually Installing or Updating the Emulex Protocol Drivers

You can install or update the Emulex protocol drivers and utilities manually without using AutoPilot Installer.

The Emulex PLUS (ElxPlus) driver provides remote management support in the OneCommand Manager application.

Note: The ElxPlus driver must be installed before you install the Emulex protocol drivers.

Removing the Adjunct Driver Registry Key

The ElxPlus driver replaces the adjunct driver that was used with Storport Miniport 1.11a3 or earlier. If it is currently installed, the adjunct driver registry key must be removed using the `deladjct.reg` file before you install the ElxPlus driver.

Installing the Emulex PLUS (ElxPlus) Driver for the First Time

Note: Only one instance of the ElxPlus driver should be installed, even if you have multiple adapter ports installed in your system.

To install the ElxPlus driver from the desktop:

1. Run the driver kit installer, but do not run AutoPilot Installer. See “Running a Software Installation Interactively” on page 15 for instructions.
2. Select **Start>Settings>Control Panel>Add Hardware**. The Add Hardware Wizard window appears. Click **Next**.
3. Select **Yes, I have already connected the hardware** and click **Next**.
4. Select, **Add a new hardware device** and click **Next**.
5. Select **Install the hardware that I manually select from a list (Advanced)** and click **Next**.
6. Select **Show All Devices** and click **Next**.
7. Click **Have Disk...** Direct the Device Wizard to the location of elxplus.inf. If you have installed the driver installer kit in the default folder and C:\ is your Windows system drive, the path is:
 - C:\Program Files\Emulex\AutoPilot Installer\Drivers\Storport\x86\HBA for the 32-bit driver version
 - or-
 - C:\Program Files\Emulex\AutoPilot Installer\Drivers\Storport\x64\HBA for the x64 driver version
 - or-
 - C:\Program Files\Emulex\AutoPilot Installer\Drivers\Storport\IA64\HBA for the Itanium driver version
8. Click **OK**.
9. Select **Emulex PLUS**. Click **Next** and click **Next** again to install the driver.
10. Click **Finish**. The initial ElxPlus driver installation is complete. Continue with manual installation of the Storport Miniport Driver. See “Installing or Updating the FC/FCoE Storport Miniport Driver” on page 22 for this procedure.

Updating the Emulex PLUS (ElxPlus) Driver

Note: Only one instance of the ElxPlus driver should be installed, even if you have multiple adapter ports installed in your system.

To update an existing ElxPlus driver from the desktop:

1. Run the driver kit installer, but do not run AutoPilot Installer. See “Running a Software Installation Interactively” on page 15 for instructions on how to do this.
2. Select **Start>Settings>Control Panel>Administrative Tools>Computer Management**.

3. Click **Device Manager** (left pane).
4. Click the plus sign (+) next to the Emulex PLUS class (right pane) to show the ElxPlus driver entry.
5. Right-click the ElxPlus driver entry and select **Update Driver...** from the menu.
6. Select **No, not this time**. Click **Next** on the Welcome to the Hardware Update Wizard window. Click **Next**.
7. Select **Install from a list or specific location (Advanced)** and click **Next**.
8. Select **Don't Search. I will choose the driver to install**.
9. Click **Have Disk...** Direct the Device Wizard to the location of driver's distribution kit. If you have installed the driver installer kit in the default folder, the path is:
 - C:\Program Files\Emulex\AutoPilot Installer\Drivers\Storport\x86 for the 32-bit driver version
 - or-
 - C:\Program Files\Emulex\AutoPilot Installer\Drivers\Storport\x64 for the x64 driver version
 - or-
 - C:\Program Files\Emulex\AutoPilot Installer\Drivers\Storport\IA64 for the Itanium driver version
10. Click **OK**. Select **Emulex PLUS**.
11. Click **Next** to install the driver.
12. Click **Finish**. The ElxPlus driver update is finished. Continue with manual installation of the Storport Miniport Driver.

Installing or Updating the FC/FCoE Storport Miniport Driver

To update or install the FC/FCoE Storport Miniport driver from the desktop:

1. Select **Start>Settings>Control Panel>System**.
2. Select the **Hardware** tab.
3. Click **Device Manager**.
4. Open the **SCSI and RAID Controllers** item.
5. Double-click the desired Emulex adapter.

Note: The driver affects only the selected adapter. If there are other adapters in the system, you must repeat this process for each adapter. All dual-channel adapter models are displayed in Device Manager as two adapters, and each adapter must be updated.

6. Select the **Driver** tab.
7. Click **Update Driver**. The Update Driver wizard starts.
8. Select **No, not this time**. Click **Next** on the Welcome to the Hardware Update Wizard window.
9. Select **Install from a list or specific location (Advanced)** and click **Next**.
10. Select **Don't search. I will choose the driver to install** and click **Next**.

Note: Using the OEMSETUP.INF file to update Emulex's FC/FCoE Storport Miniport driver overwrites customized driver settings. If you are updating from a previous installation, write down the settings. Following installation, use the OneCommand Manager application to restore the previous settings.

11. Click **Have Disk...** Direct the Device Wizard to the location of oemsetup.inf. If you have installed the driver installer kit to the default folder, the path is:
 - C:\Program Files\Emulex\AutoPilot Installer\Drivers\Storport\x86\HBA for the 32-bit driver version
 - or-
 - C:\Program Files\Emulex\AutoPilot Installer\Drivers\Storport\x64\HBA for the x64 driver version
 - or-
 - C:\Program Files\Emulex\AutoPilot Installer\Drivers\Storport\IA64\HBA for the Itanium driver version
12. Click **OK**. Select **Emulex LightPulse LPX000, PCI Slot X, Storport Miniport Driver** (your adapter model is displayed here).
13. Click **Next**.
14. Click **Finish**.

The driver installation is complete. The driver should start automatically. If the adapter is connected to a SAN or data storage device, a blinking yellow light on the back of the adapter indicates a link up condition.

Installing or Updating the iSCSI Driver

To update or install the iSCSI driver from the desktop:

1. Select **Start>Settings>Control Panel>System**.
2. Select the **Hardware** tab.
3. Click **Device Manager**.
4. Open the "SCSI and RAID Controllers" item.
5. Double-click the desired Emulex adapter.
6. Select the **Driver** tab.
7. Click **Update Driver**. The Update Driver wizard starts.
8. Select **No, not this time**. Click **Next** on the Welcome to the Hardware Update Wizard window

Note: The driver affects only the selected adapter. If there are other adapters in the system, you must repeat this process for each adapter. All DC adapter models are displayed in Device Manager as two adapters; therefore, you must update each adapter.

9. Select **Install from a list or specific location (Advanced)** and click **Next**.
10. Select **Don't search. I will choose the driver to install** and click **Next**.

11. Click **Have Disk...** Direct the Device Wizard to the location of oemsetup.inf. If you have installed the driver installer kit to the default folder, the path is:
 - C:\Program Files\Emulex\AutoPilot
Installer\Drivers\Storport\x86\iSCSI\[Windows Version] for the 32-bit driver version
 - or-
 - C:\Program Files\Emulex\AutoPilot
Installer\Drivers\Storport\x64\iSCSI\[Windows Version] for the x64 driver version
 - or-
 - C:\Program Files\Emulex\AutoPilot
Installer\Drivers\Storport\IA64\iSCSI\[Windows Version] for the Itanium driver version
12. Click **OK**. Select **Emulex OneConnect OCmXXXXX, iSCSI Initiator** (your adapter model is displayed here).
13. Click **Next**.
14. Click **Finish**.

The driver installation is complete. The driver should start automatically.

Installing or Updating the NIC Driver

Windows Server 2008

1. Select **Start>Settings>Control Panel>Device Manager**.
2. Open the **Network Adapters** item.
3. Double-click the desired Emulex adapter.
4. Select the **Driver** tab.
5. Click **Update Driver**. The Update Driver wizard starts.
6. Click **Browse my computer for driver software**.

Note: The driver affects only the selected adapter. If there are other adapters in the system, you must repeat this process for each adapter. All DC adapter models are displayed in the Device Manager as two adapters, therefore you must update each adapter.

7. Click **Let me pick from a list of device drivers on my computer** and click **Next**.
8. Select the network adapter that matches your hardware and click **Have Disk...**
9. Direct the Device Wizard to the location of oemsetup.inf. If you have installed the driver installer kit to the default folder, the path is:
 - C:\Program Files\Emulex\AutoPilot
Installer\Drivers\NDIS\x86\NIC\[Windows Version] for the 32-bit driver version

-or-

- C:\Program Files\Emulex\AutoPilot
Installer\Drivers\NDIS\x64\NIC\[Windows Version] for the x64 driver
version
- or-
- C:\Program Files\Emulex\AutoPilot
Installer\Drivers\NDIS\IA64\NIC\[Windows Version] for the Itanium driver
version

10. Click **OK**. The Windows Security dialog box opens.

11. Click **Install**.

12. When the device driver finishes installing, click **Close**.

The driver installation is complete. The driver should start automatically.

Windows Server 2012

1. Select **Server Manager>Dashboard>Tools>Computer Management>Device Manager**.

Note: Server Manager is set to open by default when booting Windows Server 2012. If not, you can open it with the Service Manager icon at the bottom left of the screen.

2. Open the **Network Adapters** item.
3. Double-click the desired Emulex adapter.
4. Select the **Driver** tab.
5. Click **Update Driver**. The Update Driver wizard starts.
6. Click **Browse my computer for driver software**.

Note: The driver affects only the selected adapter. If there are other adapters in the system, you must repeat this process for each adapter. All DC adapter models are displayed in the Device Manager as two adapters, therefore you must update each adapter.

7. Click **Let me pick from a list of device drivers on my computer**.
8. Select the network adapter that matches your hardware and click **Have Disk...**
9. Direct the Device Wizard to the location of oemsetup.inf. Select the desired oemsetup.inf file and click **Open**.

If you have installed the driver installer kit to the default folder, the path is:

- C:\Program Files\Emulex\AutoPilot
Installer\Drivers\NDIS\x86\NIC\[Windows Version] for the 32-bit driver
version
- or-
- C:\Program Files\Emulex\AutoPilot
Installer\Drivers\NDIS\x64\NIC\[Windows Version] for the x64 driver
version

-or-

- C:\Program Files\Emulex\AutoPilot Installer\Drivers\NDIS\IA64\NIC\[Windows Version] for the Itanium driver version
10. Click **Next**.
 11. When the device driver finishes installing, click **Close**.
- The driver installation is complete. The driver should start automatically.

Removing Emulex Driver Kits and Drivers

Uninstalling Emulex Driver Kits

Note: When you uninstall the Emulex driver kit, AutoPilot Installer is automatically uninstalled.

Windows Server 2008

To uninstall a driver kit on a Windows Server 2008 system:

1. Open the **Programs and Features** control panel.
2. Select one of the following in the program list and click the **Uninstall** icon in the tool bar above the program list. If you have User Access Control enabled, click **Continue** when asked for permission.
 - Emulex FC kit-2.xx.xxx
 - Emulex/FCoE kit-2.xx.xxx
 - Emulex/NIC 4xx.xxx
 - Emulex/iSCSI kit-4.xx.xxx
3. Click **Yes** when prompted to remove the kit. Once the kit is removed from the system, click **OK**.

To uninstall a driver kit on a Server Core system:

1. From the system prompt, navigate to the **Program Files** folder on an x86 system or the **Program Files (x86)** folder on a 64-bit system.
2. Navigate to **Emulex\AutoPilot Installer**.
3. Run one or more of the following batch files:
 - uninstall_fc_kit.bat (uninstall fc kit)
 - uninstall_cna_kit.bat (uninstall fcoe kit)
 - uninstall_iscsi_kit.bat (uninstall iscsi kit)
 - uninstall_nic_kit.bat (uninstall nic kit)
4. The proper driver kits are removed from the system.

On all platforms, the reports folder in the “Emulex\AutoPilot Installer” folder is not removed, so you can still view installation history and the drivers that have been installed on the system. You can delete the reports folder at any time if you want.

Windows Server 2012

To uninstall a driver kit on a Windows Server 2012 system:

1. Select **Start>Control Panel**.
2. From the Control Panel, select **Programs>Uninstall a Program**.
3. Select one of the following in the program list and click the **Uninstall** icon in the tool bar above the program list. If you have User Access Control enabled, click **Continue** when asked for permission.
 - Emulex FC kit-2.xx.xxx
 - Emulex/FCoE kit-2.xx.xxx
 - Emulex/NIC 4xx.xxx
 - Emulex/iSCSI kit-4.xx.xxx
4. Click **Yes** when prompted to remove the kit. When the kit is removed from the system, click **OK**.

To uninstall a driver kit on a Server Core system:

1. From the system prompt, navigate to the **Program Files (x86)** folder on a 64-bit system.
2. Navigate to **Emulex\AutoPilot Installer**.
3. Run one or more of the following batch files:
 - `uninstall_fc_kit.bat` (uninstall fc kit)
 - `uninstall_cna_kit.bat` (uninstall fcoe kit)
 - `uninstall_iscsi_kit.bat` (uninstall iscsi kit)
 - `uninstall_nic_kit.bat` (uninstall nic kit)
4. The proper driver kits are removed from the system.

On all platforms, the reports folder in the “Emulex\AutoPilot Installer” folder is not removed, so you can still view installation history and the drivers that have been installed on the system. You can delete the reports folder at any time if you want.

Uninstalling the Emulex Drivers

The Emulex Storport Miniport and PLUS (ElxPlus) drivers are uninstalled using the Device Manager.

Windows Server 2008

Note: On Windows 2008, after the message: “Warning – you are about to uninstall this device from your system”, is displayed. You must select the checkbox **Delete the software for this device** in order to uninstall the driver.

Emulex Storport Miniport Driver

To uninstall the Emulex Storport Miniport driver:

1. Select **Start>All Programs>Administrative Tools>Computer Management**.
2. Click **Device Manager**.

3. Double-click the adapter from which you want to remove the Storport Miniport driver. A device-specific console window is displayed. Select the **Driver** tab.
4. Click **Uninstall** and then click **OK** to uninstall.

ElxPlus Driver

To uninstall the ElxPlus driver (uninstall the ElxPlus driver only if all adapters and installations of Emulex miniport drivers are uninstalled):

1. Run the Device Manager (steps 1 and 2 above).
2. Click the plus sign (+) next to the Emulex PLUS driver class.
3. Right-click the Emulex driver and click **Uninstall**.
4. Click **OK** in the Confirm Device Removal window.

Older Versions of the Emulex Storport Miniport Driver

To uninstall or update an earlier version of the Storport Miniport driver (prior to version 1.20), you must remove the registry settings for the adjunct driver prior to manually installing a new driver.

To remove the adjunct driver registry settings:

1. Browse to the Storport Miniport driver version 1.20 (or later) driver kit that you downloaded and extracted.
2. Double-click on the **deladjct.reg** file. A Registry Editor window appears to confirm that you want to execute deladjct.reg.
3. Click **Yes**. The elxadjct key is removed from the registry.

Windows Server 2012

The Emulex Storport Miniport and PLUS (ElxPlus) drivers are uninstalled using the device manager.

Note: On Windows 2012, after the message: "Warning – you are about to uninstall this device from your system" is displayed, you must select the checkbox **Delete the software for this device** in order to uninstall the driver.

Emulex Storport Miniport Driver

To uninstall the Emulex Storport Miniport driver in Windows Server 2012:

1. Select **Server Manager>Dashboard>Tools>Computer Management>Device Manager**.
2. Double-click the adapter from which you want to remove the Storport Miniport driver. A device-specific console window is displayed. Select the Driver tab.
3. Click **Uninstall** and click **OK** to uninstall.

ElxPlus Driver

To uninstall the ElxPlus driver (uninstall the ElxPlus driver only if all adapters and installations of Emulex miniport drivers are uninstalled):

1. Run the Device Manager (step 1 above).
2. Click the plus sign (+) next to the Emulex PLUS driver class.
3. Right-click the Emulex driver and click **Uninstall**.
4. Click **OK** in the Confirm Device Removal window.

2. Configuration

FC/FCoE Driver Configuration

The Emulex Storport Miniport driver has many options that you can modify to provide different behavior. You can set Storport Miniport driver parameters using the OneCommand Manager application. Refer to the *OneCommand Manager Application User Manual* for information on using this utility to configure the driver.

Configuring FC Driver Parameters

The parameters listed in Table 2-1, Storport Miniport Driver Parameters, are applicable to the Storport Miniport driver versions 2.70.xxx or later.

The Driver Parameter table provides information such as the allowable range of values and factory defaults. Parameters can be entered in decimal or hexadecimal format.

A parameter has one of the following activation requirements:

- Dynamic – The change takes effect while the system is running.
- Reset – Requires an adapter reset from the utility before the change takes effect.
- Reboot – Requires a reboot of the entire machine before the change takes effect. In this case, you are prompted to perform a reboot when you exit the utility.

Note: If you are creating custom unattended installation scripts, any driver parameter can be modified and included in the script.

Note: If the Adapter/Protocol is blank it indicates that the parameter is supported on both LightPulse® as well as OneConnect adapters. “LightPulse only” indicates parameters supported only on LightPulse and not on OneConnect adapters. “FC only” indicates parameters supported on non-LightPulse FC adapters.

Most parameters default to a setting that optimizes adapter performance.

Table 2-1 Storport Miniport Driver Parameters

Parameter	Definitions	Activation Requirement	Adapter/Protocol
AutoMap=n	<p>AutoMap controls the way targets are assigned SCSI IDs. Discovered targets are assigned persistent SCSI IDs according to the selected binding method. Persistent bindings do not take effect with the driver in stand-alone mode.</p> <p>If set to 0 = automap is disabled. The OneCommand Manager application persistently sets the SCSI address of a discovered FCP capable FC node (target).</p> <p>If set to 1 = automap by WWNN.</p> <p>If set to 2 = automap by WWPN.</p> <p>If set to 3 = automap by DID</p> <p>Value: 0-3</p> <p>Default = 2</p>	Reboot	
Class=n	<p>Class selects the class of service on FCP commands.</p> <p>If set to 2, class = 2.</p> <p>If set to 3, class = 3.</p> <p>Value: 2-3</p> <p>Default = 3</p>	Dynamic	FC Only
CoalesceMsCnt=n	<p>CoalesceMsCn specifies wait time in milliseconds to generate an interrupt response if CoalesceRspCnt has not been satisfied. Zero specifies an immediate interrupt response notification. A non-zero value enables response coalescing at the specified interval in milliseconds.</p> <p>Value: 0-63 (decimal) or 0x0-0x3F (hex)</p> <p>Default = 0 (0x0)</p>	Reset	LightPulse Only
CoalesceRspCnt=n	<p>CoalesceRspCn specifies the number of response entries that trigger an Interrupt response.</p> <p>Value: 0-255 (decimal) or 0x1-0xFF (hex)</p> <p>Default = 8 (0x8)</p>	Reset	LightPulse Only

Table 2-1 Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement	Adapter/Protocol
ConfigScale	<p>ConfigScale sets the memory footprint profile in accord with the anticipated use case on a per port basis. The default value of 1 is considered to be the typical use case. The ConfigScale parameter supersedes the ExtTransferSize parameter for OneConnect adapters and Emulex 16-Gb adapters.</p> <p>For all values except 0, up to 1024 targets can be discovered and mapped.</p> <p>If set to 0, only 128 targets can be discovered and mapped. On a OneConnect adapter or an Emulex 16-Gb adapter, a value of 0 limits max XRLs to 512. Use ConfigScale = 0 to minimize the driver's per-port memory foot print on any adapter model.</p> <p>If set to 0 on a OneConnect adapter or an Emulex 16-Gb adapter, the max transfer size is limited to 500 KB.</p> <p>If set to 1 on a OneConnect adapter or an Emulex 16-Gb adapter, the max transfer size is limited to 1012 KB.</p> <p>If set to 2, the max transfer size is limited to 2036KB. Use ConfigScale = 2 if connecting to tape devices on a OneConnect or an Emulex 16-Gb port.</p> <p>If set to 3, which is the best setting if you are running performance benchmarks in a non-production environment, the max transfer size is limited to 2036KB.</p> <p>Values: 0, 1, 2, and 3 Default = 1</p>		OneConnect and Emulex 16-Gb adapters
DiscoveryDelay=n	<p>DiscoveryDelay controls whether the driver waits for 'n' seconds to start port discovery after link up.</p> <p>If set to 0 = immediate discovery after link up.</p> <p>If set to 1 or 2 = the number of seconds to wait after link-up before starting port discovery.</p> <p>Value: 0-2 seconds (decimal) Default = 0.</p>	Dynamic	

Table 2-1 Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement	Adapter/Protocol
DriverTraceMask	<p>The DriverTraceMask parameter is only available on operating systems that support extended system event logging.</p> <p>If set to 0 = the parameter is disabled.</p> <p>If set to 1 = error events logging is enabled.</p> <p>If set to 4 = warning events logging is enabled.</p> <p>If set to 8 = informational events logging is enabled.</p> <p>The values can be masked to generate multi-levels of events logging.</p> <p>Values: 0, 1, 4, and 8.</p> <p>Default =0.</p>	Dynamic	
EnableAck0=n	<p>Set to 1 to force sequence rather than frame level acknowledgement for class 2 traffic over an exchange. This applies to FCP data exchanges on IREAD and IWRITE commands.</p> <p>Value: 0-1 (decimal)</p> <p>Default = 1</p>	Reset	FC only
EnableAUTH	<p>EnableAUTH enables fabric authentication. This parameter requires the authentication to be supported by the fabric. Authentication is enabled when this value is set to 1.</p> <p>Value: 0-1</p> <p>Default = 0</p>	Reboot	
EnableFDMI=n	<p>If set to 1, enables management server login on fabric discovery. This allows Fabric-Device Management Interface (FDMI) to operate on switches that have FDMI-capable firmware.</p> <p>If set to 2, FDMI operates and uses the host name feature of FDMI.</p> <p>Value: 0-2 (decimal)</p> <p>Default = 0</p>	Reset	
EnableNPIV=n	<p>If set to 1, enables N_Port_ID virtualization (NPIV). Requires NPIV supported firmware for the adapter.</p> <p>Value: 0-1</p> <p>Default = 0 (disabled)</p> <p>Note: To run the driver using NPIV or SLI-3 optimization, the firmware must be version 2.72a0 or later. If an earlier version is used, the driver runs in SLI-2 mode and does not support NPIV.</p> <p>Note: NPIV is not available on 1Gb/s and 2 Gb/s adapters.</p>	Reboot	

Table 2-1 Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement	Adapter/Protocol
ExtTransferSize	ExtTransferSize is an initialization-time parameter that affects the maximum SGL that the driver can handle, which determines the maximum I/O size that a port will support. If set to 0 = the maximum default transfer size is 512KB for all controller models. If set to 1 = the maximum transfer size is 1MB. If set to 2 = the maximum transfer size is 2MB. If set to 3 = the maximum transfer size is 4MB. Value: 0-3 Default = 0 (disabled)		LightPulse adapters only including the LPe16000 HBAs.
FrameSizeMSB=n	FrameSizeMSB controls the upper byte of receive FrameSize if issued in PLOGI. This allows the FrameSize to be constrained on 256-byte increments from 256 (1) to 2048 (8). Value: 0-8 Default = 0	Reset	
InitTimeout=n	Determines the number of time-out seconds during driver initialization for the link to come up. If the link fails to come up by InitTimeout, driver initialization exits but is still successful. If the link comes up before InitTimeout, the driver sets double the amount for discovery to complete. Value: 5-30 seconds or 0x5-0x1E (hex) Default = 15 seconds (0xF)	Reboot	
LimTransferSize	Limits maximum transfer size when non-zero to selectable values. Values: 0 = Port Default 1 = 64Kb 2 = 128 Kb 3 = 256Kb	Reboot	
LinkSpeed=n	LinkSpeed has significance only if the adapter supports speeds other than one Gb/s. Value: Auto-select, 1 Gb/s, 2 Gb/s, 4 Gb/s, 8 Gb/s Default = Auto-select Note: Setting this option incorrectly can cause the adapter to fail to initialize.	Reset	FC Only

Table 2-1 Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement	Adapter/Protocol
LinkTimeOut=n	<p>LinkTimeOut applies to a private loop only. A timer is started on all mapped targets using the link timeout value. If the timer expires before discovery is re-resolved, commands issued to timed out devices returns a SELECTION_TIMEOUT. The Storport driver is notified of a bus change event which leads to the removal of all LUNs on the timed out devices.</p> <p>Value: 1-500 seconds or 0x0-0xFE (hex) Default = 30 (0x1E)</p>	Dynamic	
LogErrors=n	<p>LogErrors determine the minimum severity level required to enable entry of a logged error into the system event log. Errors are classified as severe, malfunction or command level.</p> <p>A severe error requires user intervention to correct a firmware or adapter problem. An invalid link speed selection is an example of a severe error.</p> <p>A malfunction error indicates that the system has problems, but user intervention is not required. An invalid fabric command type is an example of a malfunction error.</p> <p>An object allocation failure is an example of a command error.</p> <p>If set to 0 = all errors are logged. If set to 1 = command level errors are logged. If set to 2 = malfunction errors are logged. If set to 3 = severe errors are logged.</p> <p>Value: 0-3 Default = 3</p>	Dynamic	
NodeTimeout=n	<p>The node timer starts when a node (that is, a discovered target or UCNA) becomes unavailable. If the node fails to become available before the NodeTimeout interval expires, the operating system is notified so that any associated devices (if the node is a target) can be removed. If the node becomes available before NodeTimeout expires the timer is canceled and no notification is made.</p> <p>Value: 1-255 seconds or 0x0-0xFF (hex) Default = 30 (0x1E)</p>	Dynamic	

Table 2-1 Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement	Adapter/Protocol
QueueDepth=n	<p>QueueDepth requests per LUN/target (see QueueTarget parameter). If you expect the number of outstanding I/Os per device to exceed 32, then you must increase to a value greater than the number of expected I/Os per device (up to a value of 254). If the QueueDepth value is set too low, a performance degradation can occur due to driver throttling of its device queue. QueueDepth supports more than 1000 outstanding commands per port.</p> <p>Value: 1-254 or 0x1-0xFE (hex)</p> <p>Default = 32 (0x20)</p>	Dynamic	
QueueTarget=n	<p>QueueTarget controls I/O depth limiting on a per target or per LUN basis.</p> <p>If set to 0 = depth limitation is applied to individual LUNs.</p> <p>If set to 1 = depth limitation is applied across the entire target.</p> <p>Value: 0-1 or 0x0-0x1 (hex)</p> <p>Default = 0 (0x0)</p>	Dynamic	
RmaDepth=n	<p>RmaDepth sets the remote management buffer queue depth. The greater the depth, the more concurrent management controls can be handled by the local node.</p> <p>Value: 8-64, or 0x8-0x40 (hex)</p> <p>Default = 16 (0x10)</p> <p>Note: The RmaDepth driver parameter pertains to the functionality of the OneCommand Manager application.</p>	Reboot	
ScanDown=n	<p>If set to 0 = lowest AL_PA = lowest physical disk (ascending AL_PA order).</p> <p>If set to 1 = highest AL_PA = lowest physical disk (ascending SEL_ID order).</p> <p>Value: 0-1</p> <p>Default = 1</p> <p>Note: This option applies to private loop only in DID mode.</p>	Reboot	FC Only

Table 2-1 Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement	Adapter/Protocol
SLIMode=n	<p>If set to 0 = autoselect firmware, use the latest firmware installed.</p> <p>If set to 2 = implies running the adapter firmware in SLI-2 mode.</p> <p>If set to 3 = implies running the adapter firmware in SLI-3 mode.</p> <p>Value: 0, 2, and 3</p> <p>Default = 0</p>	Reboot	LightPulse Only
SrbTimeout	<p>SrbTimeout limits the SRB timeout value to 60 seconds when set to 1 or enabled. This is a non-displayed parameter where it has to be set manually into the registry. This option alters the I/O timeout behavior where an I/O will be returned in a max timeout of 60 seconds on some long I/O timeout.</p> <p>If set to 1 = enabled</p> <p>If set to 0 = disabled</p> <p>Values: 0, 1</p> <p>Default = 0</p>		
Topology=n	<p>Topology values can be 0 to 3.</p> <p>If set to 0 (0x0) = FC Arbitrated Loop (FC-AL).</p> <p>If set to 1 (0x1) = PT-PT fabric.</p> <p>If set to 2 (0x2) = *FC-AL first, then attempt PT-PT.</p> <p>If set to 3 (0x3) = *PT-PT fabric first, then attempt FC-AL.</p> <p>* Topology fail-over requires v3.20 firmware or higher. If firmware does not support topology fail-over, options 0,2 and 1,3 are analogous.</p> <p>Value: 0-3</p> <p>Default = 2 (0x2)</p>	Reset	FC Only
TraceBufSiz=n	<p>TraceBufSiz sets the size in bytes for the internal driver trace buffer. The internal driver trace buffer acts as an internal log of the driver's activity.</p> <p>Value: 250,000-2,000,000 or 0x3D090-0x1E8480 (hex).</p> <p>Default = 250,000 (0x3D090)</p>	Reboot	

Server Performance with FC Drivers

I/O Coalescing

I/O Coalescing is enabled and controlled by two driver parameters: `CoalesceMsCnt` and `CoalesceRspCnt`. The effect of I/O Coalescing will depend on the CPU resources available on the server. With I/O Coalescing turned on, interrupts are batched, reducing the number of interrupts and maximizing the number of commands processed with each interrupt. For heavily loaded systems, this will provide better throughput.

With I/O Coalescing turned off (the default), each I/O processes immediately, one CPU interrupt per I/O. For systems not heavily loaded, the default will provide better throughput. The following table shows recommendations based upon the number of I/Os per adapter.

Table 2-2 Recommended Settings for I/O Coalescing

I/Os per Second	Suggested CoalesceMsCnt	Suggested CoalesceRspCnt
I/Os < 10000	0	8
10000 < I/Os < 18000	1	8
18000 < I/Os < 26000	1	16
I/Os > 26000	1	24

CoalesceMsCnt

The `CoalesceMsCnt` parameter controls the maximum elapsed time in milliseconds that the adapter waits before it generates a CPU interrupt. The value range is 0–63 (decimal) or 0x0–0x3F (hex). The default is 0 and disables I/O Coalescing.

CoalesceRspCnt

The `CoalesceRspCnt` parameter controls the maximum number of responses to batch before an interrupt generates. If `CoalesceRspCnt` expires, an interrupt generates for all responses collected up to that point. With `CoalesceRspCnt` set to less than 2, response coalescing is disabled and an interrupt triggers for each response. The value range for `CoalesceRspCnt` is 1–255 (decimal) or 0x1–0xFF (hex). The default value is 8.

Note: A system restart is required to make changes to `CoalesceMsCnt` and `CoalesceRspCnt`.

Performance Testing

There are three driver parameters that need to be considered (and perhaps changed from the default) for better performance testing: `QueueDepth`, `CoalesceMsCnt`, and `CoalesceRspCnt`.

QueueDepth

If the number of outstanding I/Os per device is expected to exceed 32, increase this parameter to a value greater than the number of expected I/Os per device, up to a maximum of 254. The QueueDepth parameter defaults to 32. If 32 is set and not a high enough value, performance degradation may occur due to Storport throttling its device queue.

CoalesceMsCnt

CoalesceMsCnt defaults to zero. If you are using a performance evaluation tool such as IOMETER and if you expect the I/O activity to be greater than 8000 I/Os per second, set CoalesceMsCnt to 1 and reinitialized with an adapter reset or system reboot.

CoalesceRspCnt

CoalesceRspCnt defaults to 8. For all other values up to the maximum of 63, the adapter does not interrupt the host with a completion until either CoalesceMsCnt milliseconds has elapsed or CoalesceRspCnt responses are pending. The value of these two driver parameters reduces the number of interrupts per second which improves overall CPU utilization. However, there is a point where the number of I/Os per second is small relative to CoalesceMsCnt and this will slow down the completion process, causing performance degradation.

Examples

Test Scenario One:

- You execute IOMETER with an I/O depth of 1 I/O per device in a small-scale configuration (16 devices). In this case, the test does not exceed the adapter's performance limits and the number of I/Os per second are in the low thousands.
- Recommendation: set CoalesceMsCnt to 0 (or leave the default value).

Test Scenario Two:

- You execute IOMETER with an I/O depth of 48 I/Os per device in a small-scale configuration (16 devices).
- Recommendation: set QueueDepth to be greater than 48 (for example, 64).

NIC/TOE Driver Configuration

Note: Although TCP Offload Engine (TOE) is supported, it is enabled by default.

Configuring NIC/TOE Driver Options

The Windows Server NIC/TOE driver supports the following OneConnect configurable driver options through the Advanced Property page in Windows Device Manager. For information on how to configure the options through the Advanced Property page, see “Modifying Advanced Properties” on page 49.

Considerations for Using Universal Multichannel (UMC) and NIC/TOE

- No more than one VLAN can be used with each universal multichannel (UMC) virtual channel.
- You cannot run Link Aggregation Control Protocol (LACP) when UMC is enabled.
- UMC is not supported by TOE.

For additional information on UMC, refer to the *Boot for NIC, iSCSI, and FCoE Protocols User Manual*.

For more information on the options, see “Network Driver Performance Tuning” on page 66.

Table 2-3 Windows Server 2008, 2008 R2, and 2012 NIC/TOE Driver Options

Option Name	Acceptable Values	Supported Operating Systems	Definition
Class of Service (802.1p)	Automatic Priority (default) Filtered Priority User Priority Disable Priority	Windows 2008 Windows 2008 R2 Windows 2012	<p>The following modes are supported for selecting 802.1p priority tags:</p> <ul style="list-style-type: none"> Automatic Priority - The DCBX standard allows the network adapter to negotiate priority class usage with DCBX aware endpoints such as switches or network cards. If the peer indicates that priority pause is supported for a non-zero priority, the NIC automatically inserts the default priority in all transmitted packets. This is the default mode, allowing priority pause to operate for both storage and network traffic. If the peer indicates a zero default priority (such as when the peer does not support priority pause), the device uses the "Non-Storage Priority" mode discussed below. Filtered Priority - This mode coerces the user priorities in each packet to avoid sending packets on the network function that may disrupt the converged adapter's storage traffic. The network device uses the next lower priority if a conflict exists. This mode is useful if multiple network priorities are necessary. Only a limited number of classes are supported for priority pause, so typically it does not function optimally in this mode. User Priority - This mode allows any user specified priority value and should be limited to cases where storage functions are not used. Disable Priority - The adapter always transmits either untagged packets, or VLAN ID (802.1q) tagged packets with a priority value (802.1p) of zero.
Flow Control	Disabled RX and TX Enabled (default) Rx Enable/Tx Disable Tx Enable/Rx Disable	Windows 2008 Windows 2008 R2 Windows 2012	<p>Flow control is almost always advantageous to avoid packet drops on the network. The switch or network peer must also have flow control enabled.</p> <p>The IEEE 802.3x Ethernet specification defines a control frame between peers that can request a pause in packet transmissions. This allows one system to request a temporary halt of all incoming traffic when receive buffer space is exhausted.</p> <p>The network device may be configured to respond to pause frames (Rx Enable) and/or to send pause frames (Tx Enable).</p>

Table 2-3 Windows Server 2008, 2008 R2, and 2012 NIC/TOE Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
Interrupt Moderation	None Static (90k) Int/sec Static (70k) Int/sec Static (50k) Int/sec Static (40k) Int/sec Static (30k) Int/sec Static (25k) Int/sec Static (20k) Int/sec Static (15k) Int/sec Static (10k) Int/sec Static (5k) Int/sec Adaptive (70k) Int/sec Adaptive (50k) Int/sec Adaptive (40k) Int/sec Adaptive (30k) Int/sec (default) Adaptive (20k) Int/sec Adaptive (15k) Int/sec Adaptive (12k) Int/sec Adaptive (10k) Int/sec Adaptive (7k) Int/sec Adaptive (5k) Int/sec	Windows 2008 Windows 2008 R2 Windows 2012	<p>The network device uses interrupt moderation algorithms to reduce the total amount of CPU cycles spent processing interrupts which increases efficiency for the system. However, interrupt moderation increases the latency of each send and receive. It should only be disabled when short latencies are more important than efficient CPU utilization.</p> <p>The “None” setting disables all delays to minimize latency.</p> <p>The static settings use a constant interrupt delay to avoid any spikes in interrupt rate.</p> <p>The adaptive settings cause the driver to dynamically maintain a target interrupt rate. If CPU usage is too high, choose a higher adaptive value to increase the interrupt delays (and reduce the rate of interrupts).</p>
Receive Buffers	64-16384, inclusive The default value is 896.	Windows 2008 Windows 2008 R2 Windows 2012	<p>This determines the number of Ethernet receive buffers allocated per receive queue. If RSS is enabled, 4 or more receive queues are used. Otherwise, a single queue is enabled.</p> <p>Decreasing this value will reduce the required system memory, but performance may decrease. Each buffer is the size of the “Packet Size” parameter.</p>
IP Checksum Offload (IPv4)	Disabled RX and TX Enabled (default) RX Enabled TX Enabled	Windows 2008 Windows 2008 R2 Windows 2012	<p>This offloads the transmit and/or receive IPv4 checksum computation.</p> <p>Offloading checksums increases system efficiency.</p>
IP Checksum Offload (IPv6)	Disabled RX and TX Enabled (default) RX Enabled TX Enabled	Windows 2008 Windows 2008 R2 Windows 2012	<p>This offloads the transmit and/or receive IPv6 checksum computation.</p> <p>Offloading checksums increases system efficiency.</p>

Table 2-3 Windows Server 2008, 2008 R2, and 2012 NIC/TOE Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
Large Send Offload (LSO) v1 (IPv4)	Disabled Enabled (default)	Windows 2008 Windows 2008 R2 Windows 2012	Large Send Offload allows the NIC hardware to segment large TCP packets (up to 64kB) into smaller packets (\leq "Packet Size") that may be transmitted. This segmentation increases transmit efficiency for TCP applications that send large buffers. During segmentation, the hardware computes the IPv4 and TCP checksums for each individual packet. The Windows Version 1 LSO supports only IPv4.
Large Send Offload (LSO) v2 (IPv4)	Disable Enable (default)	Windows 2008 Windows 2008 R2 Windows 2012	Large Send Offload allows the NIC hardware to segment large TCP packets (up to 64kB) into smaller packets (\leq "Packet Size") that may be transmitted. This segmentation increases transmit efficiency for TCP applications that send large buffers. During segmentation, the hardware computes the IPv4 and TCP checksums for each individual packet. The Windows Version2 LSO supports larger offload sizes.
Large Send Offload (LSO) v2 (IPv6)	Disabled Enabled (default)	Windows 2008 Windows 2008 R2 Windows 2012	Large Send Offload allows the NIC hardware to segment large TCP packets (up to 64kB) into smaller packets (less than the MTU) that may be transmitted. This segmentation increases transmit efficiency for TCP applications that send large buffers. During segmentation, the hardware computes the TCP checksums for each individual packet. IPv6 support requires LSO Version 2, included in Windows 2008 and later.
Network Address	Valid MAC Address The default setting is None.	Windows 2008 Windows 2008 R2 Windows 2012	This overrides the permanent MAC address for the interface. The MAC address should follow this format XX:XX:XX:XX:XX:XX, where X is a hex digit (0-9 or A-F). <ul style="list-style-type: none"> The address cannot be a multicast address, which has the lowest bit in the first byte set. The address cannot be all zeros. For example, 01:00:00:00:00:00 is not valid, while 02:00:00:00:00:00 is valid.

Table 2-3 Windows Server 2008, 2008 R2, and 2012 NIC/TOE Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
Packet Size	1514 (default) 9014 4088 8222	Windows 2008 Windows 2008 R2 Windows 2012	<p>Configuration packet size for OneConnect NIC only. Limit normal sized frames and one value for jumbo packet.</p> <p>This parameter determines the maximum packet size transmitted and received on the interface. A 1514 byte frame size is standard, while larger packets are called jumbo frames.</p> <p>Using a higher frame size is generally more efficient, but it uses more system memory. A larger frame size also requires support on the network switch.</p>
RSS (Receive Side Scaling)	Disabled Enabled (default)	Windows 2008 Windows 2008 R2 Windows 2012	<p>Support for multiple RSS queues if enabled.</p> <p>Receive Side Scaling (RSS) scales receive processing over multiple CPUs in parallel. This scaling typically improves application performance; however, it tends to increase CPU usage on low end machines.</p> <p>RSS is only supported on two primary adapters per device.</p> <p>Note: For additional PCI functions, RSS does not appear in the Properties List.</p>
TCP Offload Optimization	Optimize Latency Optimize Throughput (default)	Windows 2008 Windows 2008 R2 Windows 2012	<p>This parameter only applies to TCP connection offload, which must be enabled in the "Protocol Offloads" section.</p> <p>Most applications perform better with TCP Offload Optimization set to "Optimize Throughput" which handles large data transfers with minimal CPU impact.</p> <p>Setting this parameter to "Optimize Latency" causes receive data to be delivered to the application without waiting for a TCP PSH. This causes additional receive indications that typically decrease total throughput.</p>
Transmit Buffers	64-256, inclusive The default setting is 256.	Windows 2008 Windows 2008 R2 Windows 2012	<p>Transmit Buffers sets the number of Ethernet transmits that may be posted to the hardware at any given time.</p> <p>The default value is sufficient to achieve maximum performance. Reducing this value conserves system memory.</p>

Table 2-3 Windows Server 2008, 2008 R2, and 2012 NIC/TOE Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
Preferred NUMA Node	Not present or a value from 0-65535. Optional. No default setting is set.	Windows 2008 Windows 2008 R2 Windows 2012	<p>Most modern multi-socket servers have separate memory controllers for each CPU socket. These systems have non-uniform memory access (NUMA) latencies for a given CPU core to access the local versus remote memory node.</p> <p>By setting this property, the driver attempts to use both memory and CPU cores from the given NUMA node.</p> <p>If the Preferred NUMA node is not set, the driver uses the preferred NUMA node as specified by the computer's BIOS.</p> <p>For best performance, the network applications should try to use memory and CPU affinity from the same NUMA node. This level of tuning is primarily noticeable when multiple adapters are running.</p>
Transmit CPU	Not present or a value from 0 through (number of CPUs -1). Optional. No default setting is set.	Windows 2008 Windows 2008 R2 Windows 2012	Sets the CPU to be used to process transmit completions. By default, the driver intelligently chooses a CPU in the system, so this parameter should only be set for advanced performance tuning.
TCP Checksum Offload (IPv4)	Disable RX and TX Enabled (default) RX Enabled TX Enabled	Windows 2008 Windows 2008 R2 Windows 2012	TCP Checksum Offload (IPv4) offloads the transmit and/or receive IPv4 TCP checksum computation. Offloading checksums increases system efficiency.
TCP Checksum Offload (IPv6)	Disable RX and TX Enabled (default) RX Enabled TX Enabled	Windows 2008 Windows 2008 R2 Windows 2012	TCP Checksum Offload (IPv6) offloads the transmit and/or receive IPv6 TCP checksum computation. Offloading checksums increases system efficiency.

Table 2-3 Windows Server 2008, 2008 R2, and 2012 NIC/TOE Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
TCP Connection Offload (IPv4)	Enabled Disabled (default)	Windows 2008 Windows 2008 R2 Windows 2012	<p>If TCP offload is enabled, the device offloads the entire TCP protocol, including ACK processing, retransmits, and timers. Applications that prepost receive buffers (before the data arrives) may avoid data copies in the receive path, which substantially increases the system efficiency and data rates.</p> <p>Windows does not offload TCP connections if any of the following are enabled:</p> <ul style="list-style-type: none"> • Network Load Balancing • IPSEC • Network Address Translation • NDIS 5.1 Intermediate Drivers <p>TCP offload must be enabled in the Windows operating system with the shell command:</p> <pre>netsh int tcp set global chimney=enabled</pre> <p>This parameter appears disabled if the firmware installed on your device does not support TCP connection offload. Upgrading the firmware may resolve this issue.</p> <p>View the "Statistics" property page to ensure that TCP connection offload is working.</p>
UDP Checksum Offload (IPv4)	Disabled RX and TX Enabled (default) RX Enabled TX Enabled	Windows 2008 Windows 2008 R2 Windows 2012	<p>User Datagram Protocol (UDP) offload checksum settings offload the transmit and/or receive IPv4 UDP checksum computation.</p> <p>Offloading checksums increases system efficiency.</p>
UDP Checksum Offload (IPv6)	Disable RX and TX Enabled (default) RX Enabled TX Enabled	Windows 2008 Windows 2008 R2 Windows 2012	<p>User Datagram Protocol (UDP) offload checksum settings offload the transmit and/or receive IPv6 UDP checksum computation.</p> <p>Offloading checksums increases system efficiency.</p>
VLAN Identifier	Not Present (default) 1 to 4094	Windows 2008 Windows 2008 R2 Windows 2012	<p>If selected, the adapter adds a VLAN tag to all transmitted packets, and only receives packets with the matching VLAN tag.</p> <p>Note: This property should not be used when the Emulex Teaming Driver is enabled. In that case, VLAN configuration should be performed in the Teaming Driver application.</p> <p>Note: This property should not be used with Hyper-V. In that case, the Microsoft Hyper-V Manager should be used to configure VLANs on each virtual machine.</p>

Table 2-3 Windows Server 2008, 2008 R2, and 2012 NIC/TOE Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
Wake on LAN	Disabled (default) Enabled	Windows 2008 Windows 2008 R2 Windows 2012	Enabling "Wake on LAN" allows the network device to wake up the computer when a magic packet is received during standby. In Blade server configurations, "Wake On Lan" is only supported on two primary adapters per device. Additional PCI functions appear disabled.
Recv Segment Coalescing (IPv4)	Disabled (default on Windows 2008, 2008 R2) Enabled (default on Windows 2012)	Windows 2008 Windows 2008 R2 Windows 2012	Receive Segment Coalescing merges multiple TCP segments and identifies them as a single coalesced unit to the operating system's TCP/IP stack. This reduces the per-packet receive processing overhead and CPU usage when standard 1514 byte sized frames are in use.
Recv Segment Coalescing (IPv6)	Disabled (default on Windows 2008, 2008 R2) Enabled (default on Windows 2012)	Windows 2008 Windows 2008 R2 Windows 2012	Receive Segment Coalescing merges multiple TCP segments and identifies them as a single coalesced unit to the operating system's TCP/IP stack. This reduces the per-packet receive processing overhead and CPU usage when standard 1514 byte sized frames are in use.
Maximum Number of RSS Queues	4-16 8 (default)	Note: For OCe11102 UCNA's only. Windows 2008 R2 Windows 2012	When RSS is enabled, this parameter controls the number of receive queues. Typically, this is left at the maximum value. Windows reduces the number of queues as necessary based on the number of installed CPU cores. This value may be reduced during performance tuning for a particular application. It is possible that system performance may improve by limiting the number of RSS queues. Greater than 4 RSS queues requires Advanced Mode Support be enabled in the BIOS controller configuration.
VMQ	Enabled (default) Disabled	Note: For OCe11102 UCNA's only. VMQs require Windows Server 2008 R2 or later with Hyper-V.	Virtual Machine Queues are dedicated hardware receive queues for virtual machines that filter receive packets based on the destination MAC address and/or VLAN. Receive buffers can be allocated for each queue from VM memory. This improves network throughput by distributing processing of network traffic for multiple VMs among multiple processors. It reduces CPU utilization by offloading receive packet filtering to NIC hardware. VMQs prove beneficial when 4 or more VMs are in use.
VMQ Lookahead Split	Enabled (default) Disabled	Note: For OCe11102 UCNA's only. Windows 2008 R2 Lookahead split is not supported for jumbo frames.	If Lookahead split is enabled with VMQs, the NIC can directly DMA received packets to the virtual machine memory, eliminating a packet copy from the host to VM. Lookahead split requires Advanced Mode Support is enabled in the BIOS controller configuration.

Table 2-3 Windows Server 2008, 2008 R2, and 2012 NIC/TOE Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
VMQ Transmit	Enabled (default) Disabled	Note: For OCe11102 UCNAs only. Windows 2008 R2 Windows 2012	If this option is enabled with VMQs, separate transmit queues are created for each VM network interface. Send and receive interrupts for a VM network interface are processed on the same CPU(s).
SRIOV	Enabled (default) Disabled	Note: For OCe11102 UCNAs only. Windows 2012	<p>Single Root I/O Virtualization (SR-IOV) enables the adapter to allocate virtual PCI functions for each virtual machine in Hyper-V. Note that the virtual switch and virtual network adapter must have SR-IOV enabled in the Hyper-V Manager.</p> <p>When using SR-IOV, the Emulex NIC driver must be installed on each virtual function within the virtual machine. SR-IOV provides a direct hardware interface from the virtual machine to the networking adapter, which reduces latency and improves performance.</p> <p>The Windows 2012 Server SR-IOV architecture establishes each Emulex virtual NIC with a corresponding emulated NIC. This allows the virtual machine to seamlessly failover to the emulated NIC if SR-IOV is disabled. It also allows Live Migration to another system, regardless of the installed NIC hardware.</p>
ETS	Disabled (default) Enabled	Note: For OCe11102 UCNAs only. ETS is not supported in conjunction with VMQ technology. Windows 2008 Windows 2008 R2 Windows 2012	<p>If ETS is enabled, the driver filters transmit packets based on the 802.1p priority tag into multiple separate transmit rings. The network switch should be configured for ETS to group priorities into a priority group (or traffic class). Each priority group may be assigned an QoS bandwidth limit. For example, one network priority may to support priority flow control to achieve loss-less network traffic. Using separate hardware interfaces in the driver allows each priority to progress at a different rate, or pause temporarily without affecting the other priorities.</p> <p>When ETS is enabled, all configuration regarding bandwidth and priority flow control should be performed on the network switch. The adapter will learn the configuration using the DCBx protocol.</p>

Configuring Windows Server NIC/TOE Driver Parameters

The Windows Server NIC/ TOE drivers support driver options through the Advanced Property page in Windows Device Manager.

Modifying Advanced Properties

Modify the advanced properties for the driver for Windows with the Windows Device Manager. For more information on advanced properties, see “Network Driver Performance Tuning” on page 66.

To modify the advanced properties:

1. Enter the Windows Device Manager.
 - Click **Start> Control Panel>System** and click the **Device Manager** hyperlink
 - or-
 - Click **Start>Run**, then type `devmgmt.msc` and click **OK**.

The Windows Device Manager is displayed.

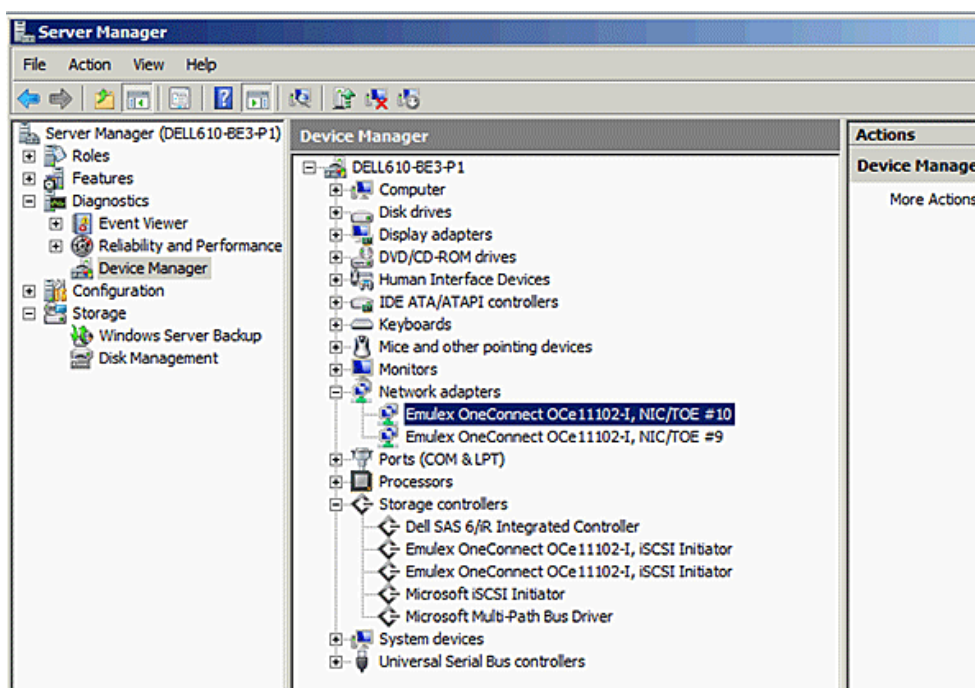


Figure 2-1 Partial View of Windows Device Manager

2. Right-click the network adapter for which you wish to modify advanced properties.
3. Click **Properties**, then click the **Advanced** tab (Figure 2-2 on page 50).
4. From the list of properties, click the property (parameter) you want to modify, then select the new value of the property by clicking the down arrow on the Value pull-down field.
5. Click **OK**.

Note: Modifying properties causes the network driver to reload and some Transmission Control Protocol (TCP) connections may be temporarily dropped.

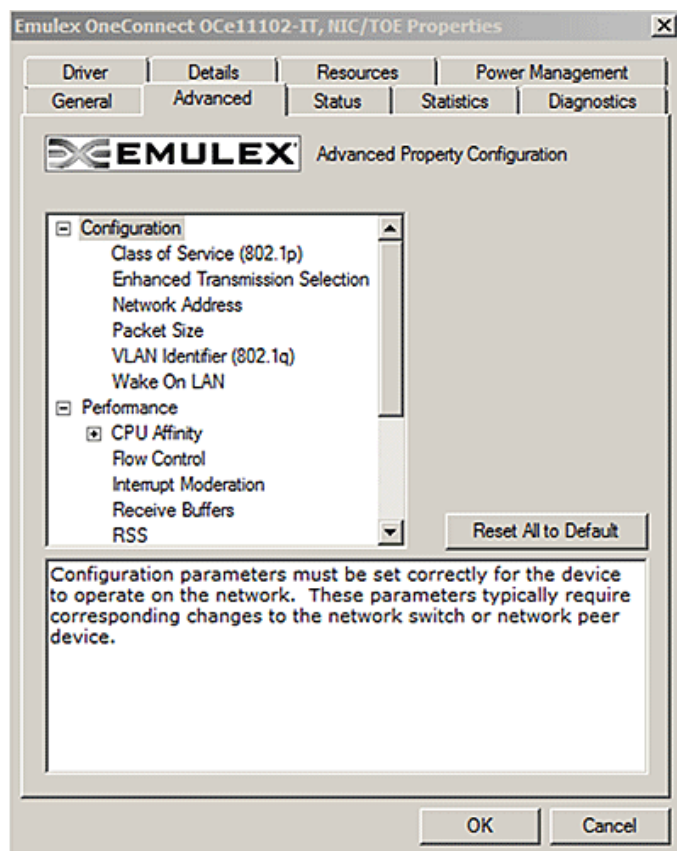


Figure 2-2 NIC/TOE Advanced Properties in Windows Server 2008

Statistics Property Page

Use the Statistics Properties tab to view the performance of the device and network. By viewing the statistics properties, you can troubleshoot problems and performance tune the system, for example you can assess how different device properties change the performance of the system.

To view the statistics properties:

1. Enter the Windows Device Manager.
 - Click **Start> Control Panel>System** and click the **Device Manager** hyperlink
- or-
- Click **Start>Run**, then type `devmgmt.msc` and click **OK**.

The Windows Device Manager is displayed (Figure 2-1).

2. Right-click the network adapter for which you wish to view the statistics properties.
3. Click **Properties**, then click the **Statistics** tab (Figure 2-3 on page 51).

4. From the list of properties, select the property (parameter) you want to view.

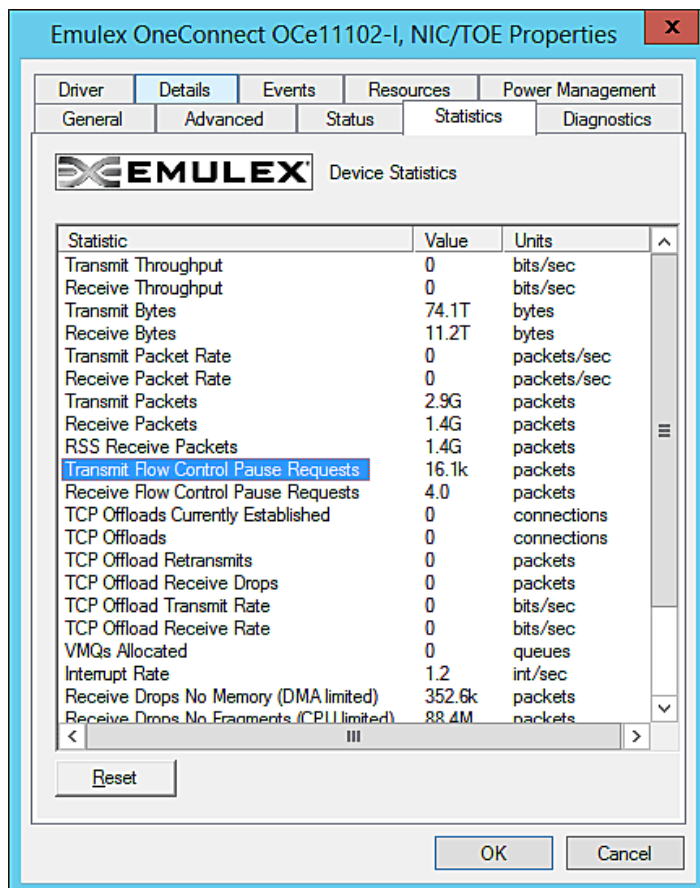


Figure 2-3 NIC/TOE Statistics Properties in Windows Server 2008

Table 2-4 NIC/TOE Driver Properties Statistics

Statistic Name	Description
Transmit Throughput	The data rate for this adapter on the network, including all packet headers. It is expressed in terms of bits/sec, where 1 byte = 8 bits. This is computed as the average over approximately 3 seconds.
Receive Throughput	The receive rate for this adapter.
Transmit Bytes	The total number of bytes transmitted by this adapter, since the last statistics reset or the last driver reload.
Receive Bytes	The total number of bytes received by this adapter.
Transmit Packet Rate	The rate of transmit packets for the adapter.
Receive Packet Rate	The rate of receive packets for the adapter.
Transmit Packets	The total number of packets transmitted by the adapter since the last statistics reset, or the driver was reloaded.
Receive Packets	The total number of packets received. This includes both RSS and non-RSS packets.

Table 2-4 NIC/TOE Driver Properties Statistics (Continued)

Statistic Name	Description
RSS Receive Packets	The number of receive packets that were suitable for RSS.
Transmit Flow Control Pause Requests	The number of times the network adapter sent a PAUSE frame to request that the peer stop sending data temporarily. This indicates a potential bottleneck in the system. Typically, this bottleneck is the result of the DMA of packets from the adapter to host memory.
Receive Flow Control Pause Requests	The number of times the network adapter received a PAUSE frame from the peer. This indicates a potential bottleneck in the attached switch or network peer device. This statistic only increments when the switch is correctly configured for flow control.
TCP Offloads Current Established	The current number of TCP connections offloaded to the adapter's TCP Offload Engine (TOE).
TCP Offloads	The total number of TCP connections that have been offloaded since the last statistics reset, or the driver was reloaded.
TCP Offload Retransmits	The number of packets retransmitted for TCP offloaded connections.
TCP Offload Receive Drops	The number of packets dropped by in the offloaded TCP stack. These drops may be the result of TCP protocol errors or bottlenecks in the system for consuming receive data.
TCP Offload Transmit Rate	The transmit data rate of the offloaded TCP connections. This is the portion of the total "Transmit Throughput" contributed by offloaded TCP connections.
TCP Offload Receive Rate	The receive data rate of the offloaded TCP connections.
VMQs Allocated	The current number of Virtual Machine Queues allocated.
Interrupt Rate	The number of interrupts per second generated by the adapter. The interrupt rate may be tuned by modifying the Interrupt Moderation parameter.

Table 2-4 NIC/TOE Driver Properties Statistics (Continued)

Statistic Name	Description
Receive Drops No Memory (DMA Limited)	<p>The number of packets dropped as a result of insufficient buffers posted by the driver. This is generally the result of the CPU core used for any receive queue reaching 100%. The system may lack sufficient CPU cycles to post receive buffers at the necessary rate. A lot of small packets lead to this behavior on almost any CPU, since the processing time for small packets is very high in the networking stack. Using a teaming driver may also lead to this, since it increases the CPU load during receive.</p> <p>Increasing the number of "Receive Buffers" in the advanced property page may alleviate some of these drops, in particular if the drops are the result of bursts of small receive packets on the network. However, if the CPU is the limit, increasing the buffer resources does not help because the driver cannot post them fast enough.</p> <p>Enabling RSS is another strategy to reduce drops since it allows the NIC driver to use additional CPU cores. The number of RSS queues may be increased to increase the total number of posted buffers available to the adapter.</p> <p>Enabling RSC can also reduce CPU consumption in the networking stack by combining multiple TCP packets into one larger packet.</p> <p>For best performance, the system BIOS should be set to "Maximum Performance" or manually disable C-states. The transitions to low power, C-states may cause a steady trickle of drops due to increased latencies from packet reception until the driver's interrupt processing code is invoked.</p>
Receive Drops No Fragments (CPU Limited)	<p>The number of receive packets dropped because of a DMA bottleneck from the network adapter to host memory. This may be caused by bottlenecks in either the PCI-Express bus or main memory.</p> <p>In the Status tab of the Custom property page, the Emulex NIC reports the PCI-Express link parameters and the maximum supported parameters. For example, installing a 8x device in a 4x PCI-Express slot cuts the available PCI-Express bandwidth in half. The PCI-Express MTU and Read Request size are also reported, and these may be configurable in the system BIOS for the computer.</p> <p>The performance of the main memory is the other major concern for networking throughput. The ideal situation is using high speed memory with all memory channels populated per CPU - typically 3 or 4 DIMMs per CPU socket. For the ideal performance, the same DIMM size should be used in each memory channel to allow perfect memory channel interleaving. Features such as memory sparing or memory mirroring dramatically decrease the memory bandwidth of the system and cause drops.</p> <p>TCP connection offload may lead to increased drops as a result of "no memory". If TCP connection offload is used, enabling flow control may reduce the drops. Alternatively, disabling TCP connection offload may improve performance.</p>

Table 2-4 NIC/TOE Driver Properties Statistics (Continued)

Statistic Name	Description
CRC Errors	The number of packets dropped as the result of CRC errors on the layer 2 Ethernet packet. In products that expose multiple PCIE functions per Ethernet port, this statistic is only incremented for the lowest PCI function per port since the packet cannot be further classified because of the error.
Receive IP Checksum Errors	The number of receive packets with an incorrect IPv4 checksum. These packets are provided to the TCP/IP stack for disposal in the operating system.
Receive UDP Checksum Errors	The number of receive packets with an incorrect UDP checksum. These packets are provided to the TCP/IP stack for disposal in the operating system.
Receive TCP Errors	The number of receive packets with an incorrect TCP checksum. These packets are provided to the TCP/IP stack for disposal in the operating system.

Using OCCFG for Windows NIC/TOE Driver Options

The occfg.exe program supports configuring parameters for the network functions on Emulex OneConnect adapters either through interactive mode with a set of menus, or command line mode that is scriptable.

If you did a standard driver installation, the occfg.exe is located in the following directory:

```
Directory of C:\Program Files\Emulex\AutoPilot  
Installer\NIC\Drivers\NDIS\<platform>\<OS>
```

The following section describes how to use the occfg.exe program to configure the Windows OneConnect device driver from the command line.

Displaying OCCFG Help

To display help, use the -? option by typing `occfg -?` on the command line. The following text will be displayed:

```
OneConnect Network Config (0.0.9999.0)
```

```
Copyright 2011 Emulex
```

```
Usage:   occfg.exe [-options]
```

Running with no arguments will display a menu to select the adapter and parameters to modify. Using the command line arguments allow scripting this process.

Options:

<code>-a str[,str]</code>	Selects all adapters with any of the given strings in the connection or device name. If omitted, <code>occfg</code> prompts for an adapter from a list.
<code>-s name=v,[name=v]</code>	Sets the parameter's value and reloads the devices.
<code>-g name[,name]</code>	Gets parameter value.
<code>-r</code>	Skips reloading the driver when setting a parameter.
<code>-f</code>	Force reloading the driver.
<code>-l</code>	List available adapters and exit.
<code>-x</code>	Reset all parameters to the default value.
<code>-p</code>	Show all registry parameter values.
<code>-q</code>	Show all driver parameter values.
<code>-h</code>	Show help text for all parameters.
<code>-?</code>	Show this help.

Examples:

```
Run interactively with menus:          occfg.exe
Set a parameter on all Emulex adapters:  occfg.exe -a Emulex -s
                                         rss=1
Set multiple parameters on one adapter:
    occfg.exe -a "Local Area Connection 23" -s "Flow=3,rss=0"
```

Selecting an Adapter

In batch mode, the “-a” parameter should be followed by a substring that is contained within the adapter name. The name is a combination of the device manager name (for example, Emulex OneConnect OCe11102) and the network connection name (for example, Local Area Connection). The later may be modified by using the Window's Network Connections applet (`ncpa.cpl`).

The most typical scenario involves setting parameters the same for all ports of a network adapter. This is accomplished by specifying “-a emulex”.

Often it is convenient to rename the connections to have a common name to easily operate on a group. For example, naming the network connections “dot1, dot2, dot3” allows operating on all adapters using the substring “dot”, or on any individual adapter by specifying the exact name such as “dot1”.

Configuring Device Parameters

OneConnect Config is used to query and modify registry parameters for Emulex network devices. The registry keys are stored at:

```
HKLM/System/CurrentControlSet/Control/Class/{4D36E972-E325-11CE-BFC108002bE10318}/####
```

where “####” is the device instance number.

The occfg program allows you to modify registry keys on a set of network devices. Once modified, the driver must be restarted to apply these parameters. In batch mode, occfg will automatically restart the driver when changing a parameter, and in interactive mode there is a menu item to select to restart the driver.

In batch mode the commands to modify parameters will look like the following examples:

```
occfg -a emulex -s rss=0
occfg -a emulex -s "Interrupt Moderation=4,Flow Control=3"
```

The parameter name must uniquely specify one parameter to modify, but it may be only a substring on the full parameter name. For example, the following are all equivalent:

```
occfg -a emulex -s "Flow Control=3"
occfg -a emulex -s flow=3
occfg -a emulex -s control=3
```

Note that the parameter name is generally the text readable parameter description name, but you may specify the exact registry key name as well. Microsoft has defined many documented standard registry key names that start with a '*' character. The '*' is not a wildcard — it is part of the registry key name. The following examples are equivalent:

```
occfg -a emulex -s "Flow Control=3"
occfg -a emulex -s "*FlowControl=3"
```

Note: Quotes are required if the parameter name contains a space character

To modify a parameter without a driver reload, use “-r”. This is useful to modify several parameters in sequence, then force a reload of the driver at the end. To force a driver reload use the “-f” parameter.

The following is an example of such a sequence:

```
occfg -a emulex -r -s rss=0
occfg -a emulex -r -s "interrupt moderation=0"
occfg -a emulex -f
```

Registry keys may be set to two special values:

- The “delete” value will cause the key to be entirely deleted and the driver will use the default value. This is appropriate for keys that are optional, such as the “Network Address”.

- The “default” value will set the key to the driver’s default value. If the key is optional, the default value may be equivalent to deleting the key.

For example:

```
occfg -a emulex -s vlan=delete
occfg -a emulex -s rss=default
```

Viewing Device Parameters

The occfg.exe program can query device parameters from either the registry or the device driver (if running driver version >= 2.103.x.x).

The registry and driver values may differ until the driver is reloaded. If the driver reload fails for any reason (such as another application has an open handle to the device driver), it may be necessary to reboot the system to apply the registry changes.

Note: If the driver has been disabled or if the driver failed to load due to any error, the driver query will return the error, “Failed to query driver for the parameter”.

The following are batch mode examples:

```
occfg -a emulex -g "Interrupt Moderation"
occfg -a "(Local Area Connection)" -g interrupt,rss
Emulex OneConnect OCell102-I, NIC/TOE (Local Area Connection):
    [Registry] Interrupt Moderation = 4 (Adaptive [Default])
    [Driver] Interrupt Moderation = 4 (Adaptive [Default])
Emulex OneConnect OCell102-I, NIC/TOE (Local Area Connection):
    [Registry] RSS = 0 (Disable)
    [Driver] RSS = 0 (Disable)
```

Resetting All Parameters

Resetting all parameters will restore the default values for each adapter. This is accomplished by using the command:

```
occfg -a emulex -x
```

Displaying All Parameters

To display the current value of all parameters, use either “-p” or “-q” command line options. This shows the registry value or driver value of the parameter, or both when using “-pq” together.

For example:

```
occfg -a emulex -p
Emulex OneConnect OCell102-I, NIC/TOE #2 (Local Area Connection 2)
Display all properties.
[Registry] Class of Service (802.1p) = 1 (Auto Priority Pause)
[Registry] Flow Control = 3 (Tx Enable/Rx Enable)
[Registry] IPv4 IP Checksum Offload = 3 (Tx Enable/Rx Enable)
```

```
[Registry] IPv4 Large Send Offload v1 = 1 (Enable)
[Registry] IPv4 Large Send Offload v2 = 1 (Enable)
[Registry] IPv4 TCP Checksum Offload = 3 (Tx Enable/Rx Enable)
[Registry] IPv4 TCP Connection Offload = 0 (Disable)
[Registry] IPv4 UDP Checksum Offload = 3 (Tx Enable/Rx Enable)
[Registry] IPv6 Large Send Offload v2 = 1 (Enable)
[Registry] IPv6 TCP Checksum Offload = 3 (Tx Enable/Rx Enable)
[Registry] IPv6 UDP Checksum Offload = 3 (Tx Enable/Rx Enable)
[Registry] Interrupt Moderation = 4 (3 Adaptive [Default])
[Registry] Network Address = <not set>
[Registry] Packet Size = 9014 (9014)
[Registry] Preferred NUMA Node = <not set>
[Registry] RSS = 1 (Enable)
[Registry] Receive Buffers = 896
[Registry] Receive CPU = <not set>
[Registry] TCP Offload Optimization = 0 (Optimize Throughput)
[Registry] Transmit Buffers = 256
[Registry] Transmit CPU = <not set>
[Registry] VLAN Identifier (802.1q) = <not set>
[Registry] Wake On LAN = 0 (Disabled)
```

Using Interactive Mode

The occfg.exe program also supports the interactive mode with a set of menus. To start this utility in interactive mode, run occfg.exe from a command console and do the following:

1. A list of adapters displays on which to operate. Type either a number of the list or a substring from any part of the name (for more information, see “Selecting an Adapter” on page 55.).
2. The program prompts for an operation, such as modifying or querying a parameter value. Follow the prompt.
3. The program provides a list of available registry parameters to modify or query. Type either the number of the corresponding option or a substring in the parameter name. The substring must uniquely identify the parameter or occfg will display all potential options.
4. To apply the parameters, select the menu item to exit and reload the drivers. Pressing control -c at any point may leave modifications in the registry, but the driver does not use the new parameters until is reloaded.

Parameter Help

In interactive mode, setting a parameter will display help text and information regarding the legal values for each parameter. This information can be dumped for all parameters by specifying the -h option.

The following is an example help text for the RSS parameter:

RSS:

Receive Side Scaling (RSS) scales receive processing over multiple CPUs in parallel. This scaling typically improves application performance; however, it tends to increase CPU usage on low end machines.

RSS is only supported on two primary adapters per device. It will appear disabled for additional PCI functions in blade server configurations.

RSS requires Windows 2008 and later.

Registry Key: *RSS

Default Value : 1 (Enable)

Valid Values :

0 = Disable

1 = Enable

Using SR-IOV with Emulex Devices

Advisory

OCe11100 series adapters may have a problem recovering from corrupted use of SR-IOV. Assigning an SR-IOV device to a virtual machine could leave the system vulnerable and lead to instability. It is strongly recommended that you assign SR-IOV devices only to virtual machines that run trusted workloads, or consider disabling the use of SR-IOV.

This advisory is highlighting a use case where a “rogue” [non-Emulex] digitally signed driver is installed by the system administrator in a virtual machine. It is then possible for that rogue driver to crash an OCe11100 series networking adapter. While there are many benefits to using SR-IOV with virtualized workloads, these benefits should be weighed against the potential risks in doing so. As an example, see the Microsoft TechNet Blog cited in the link below where the benefits and usage of the Windows Server 2012 Hyper-V switch versus NIC SR-IOV are noted.

<http://www.emulex.com/downloads/sr-iov.html>

Note: SR-IOV is only supported on OCe11000-series NIC adapters installed in Windows 2012 servers. It is not supported on OCe10102 adapters.

Note: SR-IOV support requires the Emulex NIC driver for Windows 2012. SR-IOV is supported in-box with Windows Server 2012 and in device drivers from Emulex with driver versions 4.2.390.6 and higher.

Server BIOS Configuration

SR-IOV requires support in the server chipset beyond standard virtualization technologies, including operating system control of PCI-Express and interrupt remapping. The server may have BIOS options to control SR-IOV, and typically these are disabled by default. The following may need modification in your system BIOS during boot:

- Enable “Virtualization”, such as Intel VT-x or AMD-V. This is required for any virtual machine.
- Explicitly enable SR-IOV in the system BIOS. This may be called Intel VT-d (Virtualization Technology for Direct I/O), AMD-Vi (AMD I/O Virtualization Technology), or IOMMU.

Emulex PXESelect Configuration for SR-IOV

The Emulex OCe11000 family of adapters requires enabling firmware support for SR-IOV within the Emulex PXESelect BIOS. See the *Boot for NIC, iSCSI, and FCoE Protocols User Manual* for more information.

To enable firmware support in the PXESelect BIOS:

1. Press <Ctrl+ P> during the Emulex PXE Select splash screen as the server boots. A screen appears showing global options.
2. Set the following options to use SR-IOV:
 - Advanced mode = Enable
 - Universal Multichannel (UMC) = Disable
3. Save the settings and enable SR-IOV for each PCI function. The server reboots after this modification.

SR-IOV Server Validation

Use the following Microsoft PowerShell commands to determine if your server is capable of SR-IOV. These commands may be run prior to installing the Hyper-V role.

- Get-VmHost
- Get-NetAdapterSriov
- Get-VmNetworkAdapter
- Get-VmSwitch

See Microsoft documentation for more information.

Note: Early SR-IOV capable chipsets had errors that may prevent SR-IOV from operating in Windows 2012. The PowerShell command “Get-VmHost | fl *” includes ‘IovSupportReasons’ that indicates if the chipset suffers from this issue.

Enabling SR-IOV on Unqualified Servers

If Windows 2012 detects a problem with the system I/O remapping hardware, you may still be able to use SR-IOV by explicitly enabling SR-IOV in the registry using ‘IovEnableOverride’.

Note: Emulex recommends this procedure for trusted virtual machines only.

Note: Emulex recommends that you make a back up of your registry before you make any changes.

Caution: Using registry editor can cause serious issues that may require you to reinstall the computer's operating system. Emulex cannot guarantee that issues resulting from changes you make to the registry can be repaired. Use the registry editor at your own risk.

Backing Up and Editing the Registry

1. Create a system restore point.
2. Open the registry editor by running regedit.exe at the command prompt.
3. Select the hive (the top level key) and export it to a .reg file.
4. Save the .reg file to a location off of the server as a precaution.
5. Navigate to:
HKLM\Software\Microsoft\Windows NT\CurrentVersion\Virtualization
6. Create a DWORD type entry named "IovEnableOverride".
7. Set the value of "IovEnableOverride" to 1.
8. Reboot the system.
9. If the system does not boot, press <F8> and select **Previous Known Good**, or use the system restore function while booting from an operating system install disc or recovery disc.
10. If the system boots but does not work properly, restore from a previous restore point, or import the saved .reg file and reboot.

Verifying the Driver Version

To verify the Emulex device driver meets the minimum requirements:

1. Select **Server Manager>Dashboard>Tools>Computer Management**.

2. Click **Device Manager**. The Device Manager opens.

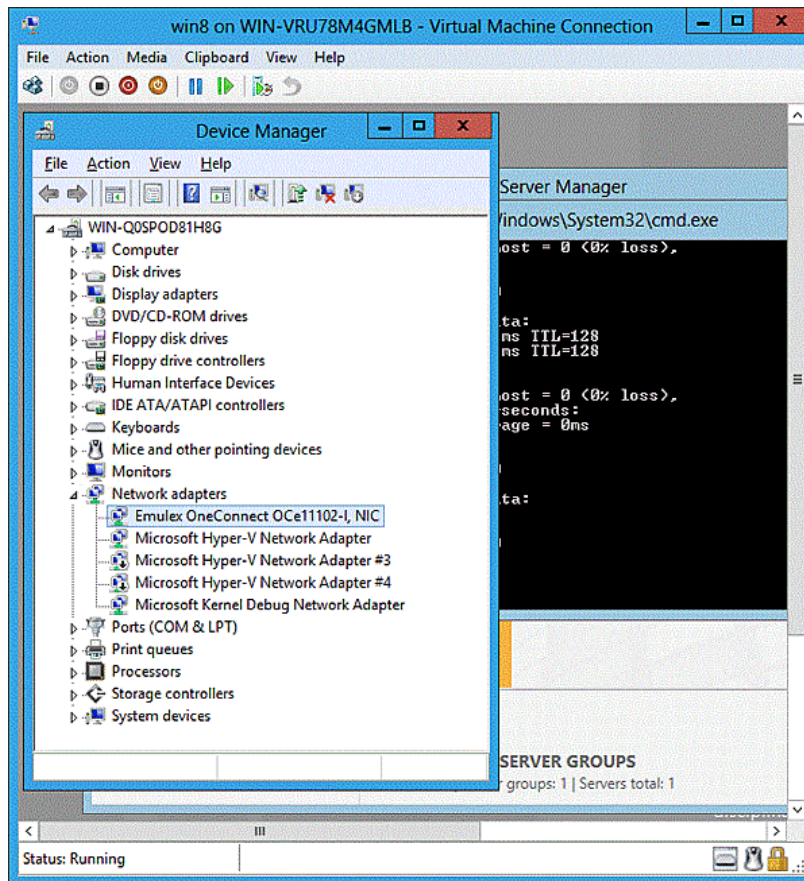


Figure 2-4 Device Manager for Windows Server 2012

3. Open the **Network Adapters** item, find the Emulex OneConnect device and right-click. Select **Properties** from the context-menu. The OneConnect Properties dialog box opens showing the Driver page. The Driver page contains the driver version number.

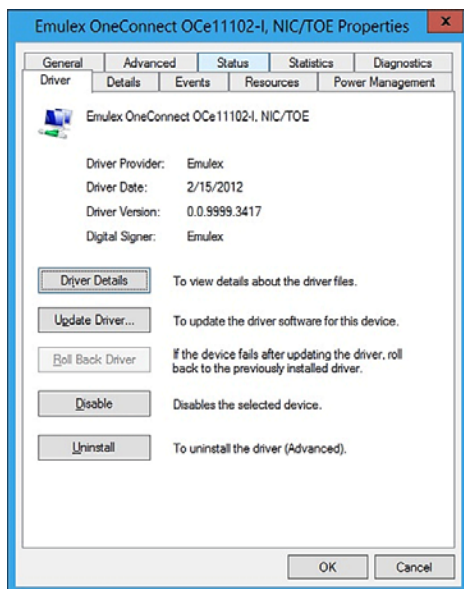


Figure 2-5 Emulex OneConnect NIC/TOE Driver Properties page

4. Click **Driver Details**. A window opens displaying the driver name.

Enabling SR-IOV in the Emulex Device

To enable SR-IOV in the Emulex device :

1. Select **Server Manager>Dashboard>Tools>Computer Management**.
2. Click **Device Manager**. The Device Manager opens. See Figure 2-4 on page 62.
3. Open the **Network Adapters** item, find the Emulex OneConnect device and right-click. Select **Properties** from the context-menu. The OneConnect Properties dialog box opens. See Figure 2-5.

4. Click the **Advanced** tab. The Advanced Property Configuration page opens.

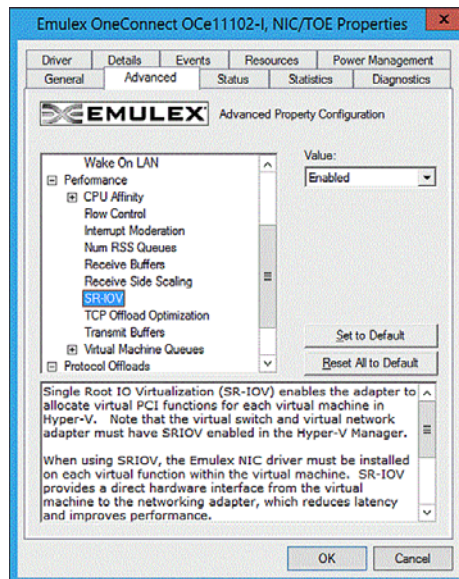


Figure 2-6 Emulex OneConnect NIC/TOE Advanced Properties page

5. Select **SR-IOV** from the list and select **Enabled** from the Value list.

Note: You must configure Hyper-V to create an SR-IOV enabled virtual machine. See Microsoft Hyper-V documentation for more information.

Hyper-V

The Hyper-V role must be added using the Server Manager. Once the Hyper-V role is added, you can enable SR-IOV in the Hyper-V Manager by:

- Creating the virtual switch
- Creating each virtual NIC

See Microsoft documentation for more information.

Note: Make sure SR-IOV is enabled on the server and on the Emulex adapter prior to configuring the Hyper-V virtual switch.

The Windows 2012 Server treats SR-IOV as an offload - an active-active team with virtual function and an emulated adapter. That means each Emulex SR-IOV adapter is accompanied by a fully functional, emulated NIC. The emulated NIC is named "Microsoft Virtual Network Adapter," and the TCP/IP stack is only bound to this device.

Once the Emulex driver is loaded, the Emulex SR-IOV virtual function is used for all unicast receive and transmit traffic. The emulated NIC handles multicast and broadcast traffic. If SR-IOV is disabled, the Emulex adapter is removed from the virtual machine, and all traffic automatically uses the emulated NIC. This technology allows Live Migration of Virtual Machines when using SR-IOV.

Note: If multiple adapters are added to the virtual machine, Emulex recommends using MAC addresses to map the Emulex Network adapter to the corresponding Microsoft Virtual Network adapters.

Verifying SR-IOV

When SR-IOV is working, it can be verified by opening the Device Manager within the virtual machine and examining the information about the transmit and receive packets that are using the SR-IOV virtual function. This is the final verification that SR-IOV is working correctly. SR-IOV can also be verified from the host Hyper-V server.

Note: Because current versions of Windows 2012 require that SR-IOV be enabled in different locations prior to creating the virtual switch, if SR-IOV is not working, delete the virtual switch and create it again. The SR-IOV option is always available during switch creation.

Verifying SR-IOV from the Virtual Machine

To verify the SR-IOV from within the virtual machine:

1. From within the virtual machine, select:
`Server Manager>Dashboard>Tools>Computer Management`
2. Click **Device Manager**. The Device Driver opens. See Figure 2-4 on page 62.
3. Open the **Network Adapters** item, click the Emulex OneConnect device and right-click. Select **Properties** from the context-menu. The OneConnect Properties dialog box opens showing the Driver page. See Figure 2-5 on page 63.

Note: The Emulex adapter may initially appear as a “Network Adapter” before the driver is loaded.

4. Select the **Statistics** tab. Information about the transmit and receive packets that are using the SR-IOV virtual function are displayed; specifically, the number of “Transmit Bytes” and “Receive Bytes” that are transmitted directly to hardware from the virtual function. If this number is greater than zero, the device is successfully using the SR-IOV direct hardware access.

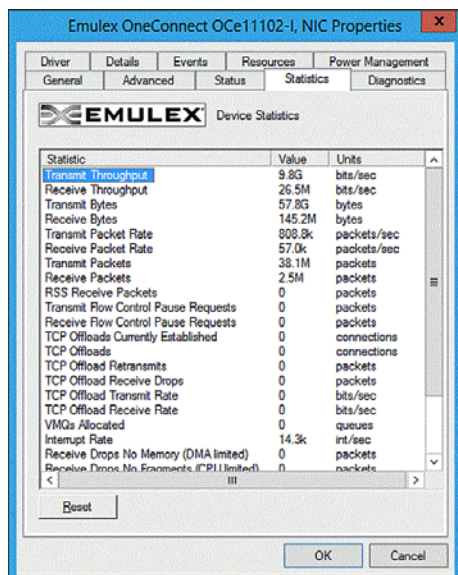


Figure 2-7 Emulex OneConnect NIC Statistics Properties page

Verifying SR-IOV from the Host Hyper-V Server

1. From the Device Manager, open the **Network Adapters** item, click the Microsoft Hyper-V Network adapter and right-click. Select **Properties** from the context-menu. The Hyper-V Network adapter Properties dialog box opens showing the Driver page.
2. Select the **Statistics** tab.
3. From the Statistics tab, locate the “Established Vf” item. “Established Vf” shows the count of currently enabled virtual functions.

Note: The Microsoft Powershell command “Get-NetAdapterSriovVf” lists each SR-IOV virtual function. See Microsoft documentation for more information.

Network Driver Performance Tuning

Network driver performance tuning improves performance of the network and TOE driver for the Windows Server operating system. The OneConnect UCNA (OCe111xx and OCe101xx) is a x8, Generation 2 (Gen 2, or Gen2) PCI-Express device and requires substantial memory bandwidth in a system to support 10 Gb/s data streams.

Optimizing Server Hardware and BIOS Configuration

UCNA performance can be improved by selecting a more efficient PCI-Express packet payload size. If the system BIOS allows selection of a larger PCI-Express packet size, selecting at least a 512-byte PCIe packet payload size provides the best efficiency for PCIe data transfers. This may be an option in the server's system BIOS. The current value is displayed in Device Manager on the Status property page for the adapter.

Most computers offer multiple distinct memory channels, which should be configured for channel interleaving for optimal performance. Optimal interleaving is achieved by using the exact same DIMM configuration for each memory channel. Check the manufacturer's documentation and BIOS parameters for details about optimizing memory bandwidth. Typically, all the dual in-line memory module (DIMM) slots must be populated to make use of all the memory channels. As a general rule, more DIMMs provide better performance by allowing a higher degree of memory-access interleaving to occur. However, some servers decrease the memory speed if using more than 2 DIMMs per memory channel – so it is important to understand the trade-off for a particular server platform.

Some servers may allow memory mirroring or memory sparing, where the total memory is divided in half and each location is stored twice. This allows fault recovery if one memory location detects an error, but it greatly reduces the perceived memory bandwidth of the system.

Nearly any desktop or low-end server has enough memory bandwidth for the UCNA to support Direct Memory Access (DMA) at 20 Gb/s of data (10 Gb/s read, 10 Gb/s write). However, most of the memory demands come from the processor accessing the data for either packet copies in the non-offloaded networking stack or application. Increasing the clock speed of the memory interface to the processor can be critical for achieving the best networking performance. This interface may be the front side bus (FSB), Intel's QPI, or AMD's HyperTransport.

Windows Server Network Driver

Table 2-5 describes ways to use various OneConnect driver properties and Microsoft Windows properties to performance tune a system.

Table 2-5 Windows Server Performance Tuning Situations

Situation	Answer/Solution
There are a large number of short-lived TCP connections such as web server or e-mail server.	Enable RSS, increase number of RSS queues, and disable TCP offload.
There are large data transfers such as to a file server, web server with file downloads, or an FTP server.	Use TCP connection offload.
There are large data transfers such as to a backup server.	Enable jumbo packets, and use TCP offload.
There is a small server struggling to keep up with larger servers on the network.	Disable RSS, enable TCP offload, enable jumbo packets, and increase the interrupt moderation to allow fewer interrupts per second.
There is a general purpose server such as Active Directory server, DHCP server, or a DNS server.	Use TCP offload, and enable RSS.

Analyzing Performance Problems

You can use the Windows Performance Monitor (perfmon) to view statistics for each network device.

To view statistics for each network device:

1. Click **Start > Run > perfmon** to launch the Windows Performance Monitor.
2. Right-click and select **Add Counters** to add additional statistics.

Table 2-6 lists a few statistics to use for troubleshooting performance issues. For network performance, all the counters from the table are useful: Network Interface, TCPv4, IPv4, and Processor.

Table 2-6 Statistics and Fine Tuning

Situation	Answer/Solution
Network Interface > Packets Received Errors.	If this is incrementing even a small amount, a physical problem may exist on the network, such as a loose connection or bad cable, causing CRC errors in Ethernet packets. Find and eliminate the physical problem.
Network Interface > Packets Received Discarded.	If this is incrementing dramatically, the computer system may be receiving a lot of unsolicited traffic using network resources.
IPv4 > Fragmented Datagrams / sec.	If this is greater than 0, the computer system is sending or receiving IP fragments. This is a serious performance problem. See "Jumbo Packet" on page 68.
TCPv4 > Segments Retransmitted / sec.	TCP retransmits indicate that packets are being dropped by the receiving system (or in a network switch). Ideally, reduce retransmits to 0.
Processor > % Processor Time.	If CPU usage is high, try to enable all available offloads, such as TCP offload, checksum offloads and use jumbo packets.

Jumbo Packet

The jumbo packet setting in the registry determines the maximum Ethernet packet size. It includes the Ethernet frame header (typically 14 bytes) but excludes the trailing Cyclic Redundancy Check (CRC). The standard packet size is 1514 bytes plus a 4 byte trailing CRC.

Vendors use many terms that refer to this same quantity, such as packet size, frame size, or Maximum Transmission Unit (MTU). The MTU is the Ethernet packet payload size. This does not include the Ethernet frame header or the trailing CRC. The standard MTU is 1500 bytes, corresponding to a 1514-byte packet size plus a 4-byte trailing CRC. Historically, any 1514-byte frame is a standard packet, while any frame larger than 1514 bytes is called a jumbo packet. Windows Server attempts to standardize the terminology across vendors so the jumbo packet parameter refers to the byte size of the packet.

The Windows Server driver supports several jumbo packet values. The larger packet size provides better throughput and CPU usage. Typically, all devices on the network, including switches, must be configured for the larger size. The drawbacks of using jumbo packets are interoperability and increased memory usage on the server.

To set a jumbo packet value, go to the Advanced Properties page in Windows Device Manager. For information on how to configure the options through the Advanced Property page, see “Modifying Advanced Properties” on page 49.

The path MTU is the maximum MTU that can be used before IP fragmentation occurs, taking into account the MTU for the endpoints and all routers between the endpoints. To verify the path MTU, ping a remote target with an increasing payload size. Eventually, the IP packet length exceeds the path MTU, and the packet fragments. This can be seen by using a packet sniffing application, such as Ethereal, Wireshark, or Microsoft Network Monitor.

IP fragmentation degrades performance dramatically, because all fragments must be received and reassembled before delivering the network packet to the upper layer protocol. In many cases, IP fragmentation may lead to a 10x performance degradation. The MTU parameter should be modified on all systems to avoid IP fragmentation for optimal network throughput.

Typical cases for using the MTU:

- Server interconnects are typically deployed using jumbo frames. This is the most efficient configuration for high bandwidth server-to-server communication, such as Network Attached Storage, iSCSI and database transactions.
- Servers connected to client systems that run desktop operating systems typically use standard 1500-byte frames. Most desktop systems do not support jumbo packets.
- Servers that need both high performance server-to-server communication and client access can be configured with jumbo frames with Path MTU Discovery enabled. Path MTU Discovery is enabled by default in Windows Server, and it allows TCP connections to negotiate the optimal packet size that avoids IP fragmentation.

Flow Control

The UCNA supports IEEE 802.3x standard flow control, which uses control packets to temporarily pause the transmission of packets between two endpoints. These control messages are point-to-point, they are not forwarded by switches or routers. You must configure both endpoints for flow control. The UCNA can either respond to flow control packets (by temporarily pausing transmits) or send flow control PAUSE packets when the transmitter is overwhelming the system's receive bandwidth. For best performance, flow control must be enabled on the switches as well as on UCNAs. Receive and transmit flow control are on by default. Flow control is not available if using FCoE on a converged network adapter. In this situation, priority pause is negotiated with the network switch and used only for the FCoE protocol packets.

The NIC function can also use priority pause if supported by the switch. This requires tagging packets in the operating system with the correct priority value, and enabling ETS in the driver properties.

Configurations that support multiple PCI functions per port generally configure flow control from the switch or blade configuration application. Since flow control is an Ethernet port property, it must be the same for all PCI functions using the same port.

If multiple PCI functions are exposed for a single 10-Gb/s Ethernet port, such as in a blade configuration, the flow control parameter must be set the same on all adapters for the port. The results are unpredictable if the setting differs among PCI functions, because this is a shared property of the 10-Gb/s port.

Examples:

Flow control greatly improves the following situations:

- The UCNA is installed in 4x PCI-Express slot or an underpowered server system.

If the PCI-Express bus does not provide 10 Gb/s of throughput due to chipset limitations or the bus width, the UCNA cannot maintain 10 Gb/s of incoming receive data. It starts dropping packets quickly. In this situation it may be beneficial to enable receive flow control in the UCNA, and enable flow control in the attached switch for all devices. This helps to slow down the transmitters.

- The UCNA transmits to 1-Gb devices, especially non-TCP protocol.

If the UCNA transmits to a 10-Gb/s switch with attached 1-Gb clients, the UCNA may overwhelm the switch. The switch is then forced to start dropping packets because, although it may receive a 10-Gb/s stream, the client can only sink a 1-Gb stream. In this situation, it may be beneficial to enable transmit flow control in the UCNA, and enable flow control for the 10-Gb/s switch port.

Note: If multiple PCI functions are exposed for a single 10-Gb/s Ethernet port, such as in a blade configuration, the flow control parameter must be set the same on all adapters for the port. The results are unpredictable if the setting differs among PCI functions, because this is a shared property of the 10-Gb/s port.

For information on modifying the Flow Control parameter, see “Configuring NIC/TOE Driver Options” on page 40.

Checksum Offloading and Large Send Offloading (LSO)

The UCNA supports IP, TCP, and UDP checksum offloading. All these protocols are enabled by default. You can disable offloading through the Windows Device Manager Advanced Properties. Disabling checksum offloading is only useful for packet sniffing applications, such as Ethereal or Microsoft Network Monitor, on the local system where the UCNA is installed and monitored. When packets are sniffed, transmit packets may appear to have incorrect checksums because the hardware has not yet calculated them.

The UCNA supports transmit LSO, which allows the TCP stack to send one large block of data, and the hardware segments it into multiple TCP packets. This is recommended for performance, but it can be disabled for packet sniffing applications. LSO sends appear as giant packets in the packet sniffer, because the hardware has not yet segmented them.

For information on modifying the CheckSum Offload or Large Send Offload parameter, see “Configuring NIC/TOE Driver Options” on page 40.

Receive Side Scaling (RSS) for Non-Offloaded IP/TCP Network Traffic

The UCNA can process TCP receive packets on multiple processors in parallel. This is ideal for applications that are CPU limited. Typically, these applications have numerous client TCP connections that may be short-lived. Web servers and database servers are prime examples. RSS typically increases the number of transactions per second for these applications.

Understanding RSS

To better understand RSS, it helps to understand the interrupt mechanism used in the network driver. Without RSS, a network driver receives an interrupt when a network packet arrives. This interrupt may occur on any CPU, or it may be limited to a set of CPUs for a given device, depending on the server architecture. The network driver launches one deferred procedure call (DPC) that runs on the same CPU as the interrupt. Only one DPC ever runs at a time. In contrast, with RSS enabled, the network driver launches multiple parallel DPCs on different CPUs.

For example, on a four-processor server that interrupts all processors, without RSS the DPC jumps from CPU to CPU, but it only runs on one CPU at a time. Each processor is busy only 25 percent of the time. The total reported CPU usage of the system is about 25 percent (perhaps more if other applications are also using the CPU). This is a sign that RSS may help performance. If the same four-processor server uses RSS, there are four parallel executing DPCs, one on each processor. The total CPU usage that is available for networking processing is increased from 25 percent to 100 percent.

Some server machines and some network traffic profiles do not benefit from RSS. Because the non-offloaded TCP stack includes a data copy during receive processing, it is possible that memory bandwidth will limit performance before the CPU. In this situation, the CPU usage is very high while all processors wait for memory accesses. To overcome this problem, you can reduce the number of RSS CPUs, or disable RSS entirely.

Poor RSS behavior is typical only in network performance testing applications that receive data, but perform no other processing. For other applications, RSS allows the application to scale other processing tasks across all CPUs, thereby improving overall performance. RSS offers the most benefit for applications that create numerous, short-lived connections. These applications are typically CPU limited instead of network bandwidth limited.

For information on modifying the RSS Queues parameter, see “Configuring NIC/TOE Driver Options” on page 40.

Note: Microsoft currently does not schedule RSS processing on all hyper-threaded CPUs. For example, only CPU 1 and 3 have RSS queues on a dual-core, hyperthreaded CPU.

Enabling Windows to Use Up to Eight Processors

Windows Server 2008 uses only four processors by default. It is possible for UCNAs to use up to eight processors. In order for the OneConnect driver to use up to eight processors, the registry must be changed and the system restarted.

For Windows Server 2008, set the registry keyword `MaxNumRssCpus` (a `DWORD` type) to 8 at the location:

```
HKEY_LOCAL_MACHINE\\SYSTEM\\CurrentControlSet\\Services\\Ndis\\Parameters
```

Note: Do not set the registry keyword to a value greater than the number of processors in the system or 16, whichever is smaller.

For Windows Server 2008 R2 and Windows Server 2012, the operating system uses all available CPU cores for RSS without manual configuration.

TCP Offloading (TOE)

The UCNA and drivers support TCP offload, which provides significant performance improvements. The performance improvements are:

- A zero-copy receive data path exists. In contrast, all non-offloaded TCP packets are copied in the network stack. This copy dramatically increases the memory bandwidth and CPU requirements for receive data.
- Sending and receiving of ACK packets is handled entirely in hardware, reducing PCI-Express bus usage and interrupts.
- TCP timers, including delayed ACK, push, retransmit and keep alive, are implemented in hardware. This reduces host CPU usage.
- Retransmits are handled entirely in hardware.
- Packetizing data, including segmenting, checksums, and CRC, is supported. The network driver can receive send and receive buffers that are larger than 1 MB for maximum efficiency.
- The driver provides efficient parallel processing of multiple connections TCP on multiple CPU systems.

The UCNA receive path is zero-copy for applications that prepost receive buffers or that issue a socket read before the data arrives. Ideal applications use Microsoft's Winsock2 Asynchronous Sockets API, which allows posting multiple receive buffers with asynchronous completions, and posting multiple send operations with asynchronous completions. Applications that do not prepost receive buffers may incur the penalty of the data copy, and the performance improvement is significantly less noticeable.

Applications that transmit large amounts of data show excellent CPU efficiency using TCP offload. TCP offload allows the network driver to accept large buffers of data to transmit. Each buffer is roughly the same amount of processing work as a single TCP packet for non-offloaded traffic. The entire process of packetizing the data, processing the incoming data acknowledgements, and potentially retransmitting any lost data is handled by the hardware.

TCP Offload Exclusions

Microsoft provides a method to exclude certain applications from being offloaded to the UCNA. There are certain types of applications that do not benefit effectively from TCP offload. These include TCP connections that are short-lived, transfer small amounts of data at a time, exhibit fragmentation from end to end, or make use of IP options.

If an application sends less data than the Max Segment Size (MSS), the driver, like most TCP stacks, uses a Nagling algorithm. Nagling reduces the number of TCP packets on the network by combining small application sends into one larger TCP packet. Nagling typically reduces the performance of a single connection to allow greater overall performance for a large group of connections. During Nagling, a single connection may have long pauses (200 ms) between sending subsequent packets, as the driver waits for more data from the application to append to the packet. An application can disable Nagling using the TCP_NO_DELAY parameter. TCP offload does not improve the performance for connections that Nagle, because the performance is intentionally limited by the Nagling algorithm. Telnet and SSH consoles are examples of connections that typically use Nagling.

Windows Server has not optimized the connection offload path. Some applications that use numerous short-lived TCP connections do not show a performance improvement using TCP offload.

Windows Server provides control over the applications and TCP ports that are eligible for TCP offload using the netsh tool. Refer to the Microsoft documentation for these netsh commands:

```
netsh interface tcp add chimneyapplication state=disabled  
application=<path>  
netsh interface tcp add chimneyport state=disabled remoteport=23  
localport=*
```

Note: The netsh commands require the Windows firewall to be running. If the firewall is disabled, all applications and ports added with the netsh commands may fail to connect.

TCP Offload Optimization Settings

The UCNA supports an option for optimizing TCP connection offload characteristics for throughput or latency. This option is available through the Advanced Property Page. See “Configuring NIC/TOE Driver Options” on page 40 for the TCP Offload Optimization settings.

The default option is Optimize Throughput, which produces the best throughput characteristics for certain types of traffic flows. This configuration setting has produced the best results on benchmarks such as Chariot, ntttcp, and iperf.

The other available option, Optimize Latency, improves the latency characteristics for the class of traffic flows not ideally suited for offloaded by sacrificing throughput. These are applications that typically do not pre-post receive buffers at a rate fast enough to keep up with the traffic flow, causing the received data to be buffered until the application has pre-posted a receive buffer. Some applications intentionally are written this way to “peek” at incoming data to determine how large of a receive buffer to post. The timings of such a usage semantic in some cases (depending on factors such as CPU-Memory performance, line rates, the sizes of the receive buffers, and system loading at the time) will result in no observable improvement.

It is recommend that you leave this parameter set to the default of Optimize Throughput.

Windows Networking and TOE

If certain Windows Server 2008 and Windows Server 2008 R2 networking features are enabled, TOE does not operate as expected, and connections are not offloaded.

Installing or activating firewall applications causes no connections to be offloaded by the Windows Server 2008 and Windows Server 2008 R2 network stack. By default, Windows Firewall Services are enabled at operating system installation time, and they must be explicitly disabled in order to use TOE. Firewall services can be disabled through the Service Control panel, or the following commands at the command line prompt:

To set firewall services to load on demand `sc config MpsSvc start= demand`

To stop firewall services `Net stop MpsSvc`

To temporarily disable firewall services `netsh advfirewall set all state off`

Enabling certain Windows networking features, such as network bridging, VPN, and routing, may cause the operating system to enable IP NAT services and the IPSEC policy agent. These services, if enabled, disallow connections from being offloaded to the UCNA. To disable these functions, use the Services Control panel, or the following commands at the command line prompt:

```
net stop accesspolicy
net stop sharedaccess
net stop ipnat
```

Windows TCP Parameters

Emulex does not recommend modifying the TCP registry parameters, such as TcpAckFrequency, provided by Microsoft. The default parameters are suitable for a wide variety of situations, with or without using TCP offloading.

Receive Window Auto Tuning and Compound TCP

Windows Server adds several features to the host TCP stack, such as receive window auto-tuning and Compound TCP (CTCP). These features affect only non-offloaded TCP traffic.

Performance of some 10 Gb/s stress applications may suffer with these features enabled. In particular, Emulex has seen some bi-directional data stream test performance degradation when receive window auto-tuning is enabled. This is due to increased receive performance that adversely affects the same TCP connection's transmit performance.

To disable these features, type these commands at the command line:

```
netsh interface tcp set global autotuning=disabled  
netsh interface tcp set global congestionprovider=none
```

Interrupt Coalescing

The Windows Server network driver automatically performs adaptive interrupt coalescing. During periods of low network usage, the interrupt delay is set to a minimum for lower latency. As the interrupt rate increases, the delay is increased. This allows the driver to perform more work in a single interrupt, which reduces the amount of wasted cycles from additional interrupts.

The interrupt coalescing algorithm automatically tunes the system to maintain responsiveness and performance in a wide variety of situations, including RSS and TOE traffic.

On slower machines, excessive interrupts cause user input to become non-responsive, and they may not allow sufficient CPU cycles for higher level drivers (such as Microsoft iSCSI Initiator) and applications. This may result in timeouts in upper layer applications, because they are never scheduled to run. Increasing the level of interrupt coalescing can alleviate these problems. Increasing interrupt coalescing may improve total bandwidth for applications that transfer large data buffers. Additionally, servers running numerous parallel TCP connections may benefit from higher interrupt coalescing.

Some applications run slower with interrupt coalescing enabled, such as applications that depend on the completion of the current network transfer before they post additional work. If an application sends and receives one network message before posting the next message, it is considered latency bound. For latency bound applications, an interrupt is required to proceed to the next work item, so reducing the number of interrupts directly reduces the network throughput. The Microsoft iSCSI Initiator is generally considered a latency bound application unless the I/O sizes are very large.

When tuning the system, you must balance the extra CPU usage caused by interrupts with the potential decrease in total throughput for latency bound applications.

CPU Binding Considerations

Windows applications may set a processor affinity, which binds a program to a particular CPU in a multiple processor computer. However, with the recent additions to the Windows networking stack, manually configuring CPU affinity is not recommended.

The advantage of application affinity for network applications is based on choosing the ideal relationship between the DPC and application affinity to reduce processor-cache coherency cycles. The ideal mapping may require that both the DPC and application run on the same processor, different processors, or different cores of a dual-core processor that share a common memory cache. Even when the best affinity relationship is determined, it is impossible to enforce this relationship because RSS or TCP offloading choose the DPC processor.

The driver uses multiple parallel DPCs that are explicitly assigned to particular CPUs for processing both RSS and TCP offloading tasks. Each TCP connection is assigned to a particular CPU for processing. This provides the advantage of assigning CPU affinities by reducing CPU cache misses, without any user configuration. Explicit processor affinity assignments are not necessary for the driver because the advantages of assigning processor affinities are realized by using RSS.

The only reason to experiment with application and interrupt CPU affinity is when performing isolated networking benchmarks.

Single TCP Connection Performance Settings

One common benchmark is to run a single TCP connection between two computers as fast as possible. The following are a few suggestions to deliver the best possible performance:

- Use TCP window scaling with a 256 Kb or 512 Kb window. This may be controlled with show socket applications, such as ntttcp from Microsoft.
- Use send and receive buffers that are larger than 128 Kb with an efficient application such as ntttcp.
- Disable RSS and use an interrupt filter driver. Experiment with all relative CPU affinities to find the best combination.
- Disable timestamps and SACK, because the test should run without dropping any packets.
- Unbind unused network protocols in the Network Connections property page.
- Disable any firewall services, IPSEC, or NAT.

iSCSI Driver Configuration

Table 2-7 lists the user-configurable iSCSI driver options available on Windows Server. It includes a description of the parameters, their default values, and the limits within which they can be configured.

Note: If the value given for a parameter is outside the supported range, the driver logs an error in the Event Log and continues to load by using the parameter's default value.

Configuring iSCSI Driver Options

To change the iSCSI driver's configuration parameters, you must edit the registry. Because there is the risk of corrupting your system and rendering it unusable, Emulex recommends that you make a back up of your registry before you make any changes.

Caution: Using registry editor can cause serious problems that may require you to reinstall the computer's operating system. Emulex cannot guarantee that problems resulting from changes you make to the registry can be repaired. Use the registry editor at your own risk.

Backing Up and Editing the Registry

1. Create a system restore point.
2. Open the registry editor by running regedit.exe at the command prompt.
3. Select the hive (the top level key) and export it to a .reg file.
4. Save the .reg file to a location off of the server as a precaution.
5. Navigate to:
`HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\be2iscsi\Parameters\Device`
6. Double-click on **DriverParameters** to edit.
7. Enter the configuration parameter. See Table 2-7, iSCSI Driver Options, on page 78 for iSCSI configuration parameters.

For example, `eto=60;ldto=30;im_policy=3;`

Note: A semicolon is required after every parameter.

8. Reboot the system.
9. If the system does not boot, press <F8> and select **Previous Known Good**, or use the system restore function while booting from an operating system install disc or recovery disc.
10. If the system boots but does not work properly, restore from a previous restore point, or import the saved .reg file and reboot.

Table 2-7 iSCSI Driver Options

Parameter	Default Value	Minimum Value	Maximum Value	Description
LDTO	20 seconds	0 seconds	3600 seconds	Link-down timeout (LDTO), in seconds. This parameter determines the amount of time the universal converged network adapter (UCNA) driver waits for the controller's physical link to be available before reporting that the logical unit numbers (LUNs) are unavailable to the operating system. Note: LDTO values are configurable via the Windows registry.
ETO	90 seconds	0 seconds	3600 seconds	Extended timeout (ETO) in seconds. This parameter determines the amount of time the UCNA driver waits for the target to be available after it has lost connection to the target during an I/O operation. Note: ETO values are configurable via the Windows registry.
im_policy	2	0	4	The Interrupt Moderation policy parameter controls the rate of interrupts for the UCNA. For more information, see "Interrupt Moderation Policy Settings" on page 79.
lqd	128	1	255	The LUN queue depth parameter configures the number of concurrent commands to a logical unit via Storport API StorPortSetDeviceQueueDepth. The lqd parameter also sets the maximum number of concurrent commands allowed per LUN.

Table 2-7 iSCSI Driver Options (Continued)

Parameter	Default Value	Minimum Value	Maximum Value	Description
large_io	64	64	512	<p>Maximum transfer size in a single I/O request, in KB. By default, the iSCSI driver supports a maximum of 64 KB of data and 16 scatter/gather entries in a single I/O request. This option enables support for 512 KB of data in a single I/O request. If an application issues an I/O request that is larger than 64 KB or that needs more than 16 scatter/gather entries, the request is split into multiple requests by the Storport driver.</p> <p>Note: If the large_io parameter is set to 512, the amount of physical memory consumed by the driver increases. Also, although intermediate values between 64 and 512 are accepted, the memory used by the driver is the same as is used if large_io is set to 512.</p>

Interrupt Moderation Policy Settings

The Interrupt Moderation policy settings control the rate of interrupts for UCNA hardware. By default, the driver implements an interrupt moderation scheme that is based on the I/O load and the interrupt rate. The default setting for im_policy tries to vary the interrupt rate between 3500 to 10000 interrupts per second. In addition, the iSCSI driver allows other configuration settings, as shown in Table 2-8.

Table 2-8 im_policy Settings

Parameter Value	Setting	Description
im_policy=0	Disabled	The Interrupt rate algorithm is turned off in the driver.
im_policy=1	Aggressive	The highest interrupt rate among all available settings.
im_policy=2	Moderate	The default value.
im_policy=3	Conservative	A lower interrupt rate than moderate.
im_policy=4	Very conservative	The lowest interrupt rate among all available settings.

While the default setting may work for most configurations, there are instances when the setting may need to be altered. The im_policy parameter setting should be based on the UCNA system configuration, the number of iSCSI targets to be connected, the I/O load, and the throughput and latency offered by these iSCSI targets.

On systems that are capable of sustaining a higher interrupt rate and on which the number of connected targets is low (eight or fewer), setting the `im_policy` to 1 results in lower latency and higher values of I/O operations per second (IOPs). But this aggressive interrupt rate can also result in system stalls and freezes, especially if queue depth values are high and I/O requests are small.

In a configuration that involves a large number of iSCSI targets (more than 32 or 64) and higher values of queue depth, the default setting may prove to be too aggressive. In such a case, you may need to change the `im_policy` parameter setting to 3 or 4. Although this increases latency of an I/O request, the lower interrupt rate may allow the system to be functional under a high load.

Creating Non-Bootable Targets

To set up non-bootable targets, proceed with the driver and operating system installation, then download and use the Microsoft iSCSI Initiator Service to configure and manage the adapter.

Using the Microsoft iSCSI Initiator Service

You can use the Microsoft iSCSI Initiator Service to configure and manage the UCNA. The Microsoft Initiator Service is available as a free download from www.microsoft.com. See the documentation that accompanies it for detailed information.

Note: When you install the Microsoft iSCSI Initiator Service, you need to select only the Initiator Service check box and not the Software Initiator check box.

The Microsoft iSCSI Initiator Service sets its own initiator name. Once you have installed it, you must replace this with your chosen initiator name. To do this:

1. In the Microsoft iSCSI Initiator Service, under the general tab, click **Change**.
2. Type your initiator name and click **OK**.

Logging into a Target Using the Microsoft Software Initiator

If you install the Software Initiator, you must select the UCNA initiator when logging into the target. To do this:

1. From the Targets tab, select the target and click **LogOn**.
2. Click **Advanced**. Under the General Tab, everything appears as default.
3. Select the UCNA initiator as the local adapter, select your source IP, and click **OK**.

Windows Multipath I/O Support

This section describes the installation and login processes for multipath I/O support on Windows Server operating systems.

Multipath Support on Windows Server 2008, 2008 R2, and 2012

On Windows Server 2008, 2008 R2, and 2012, MPIO must be installed from the Server Manager. After installing the MPIO feature, you must launch and configure the MPIO GUI to enable multipath support for iSCSI devices.

The following steps describe the installation process for setting up Microsoft iSCSI DSM and enabling multipath I/O for all iSCSI devices irrespective of their vendor and device IDs. You can use the MPIO GUI to configure DSMs other than Microsoft iSCSI DSM. Also, you can use the GUI to enable multipath support for a specific vendor ID and device ID. For details on these topics, refer to the Microsoft TechNet Library on the Microsoft website.

To configure the iSCSI drivers to fail over to an alternate path as quickly as possible, edit the Driver Parameters for the iSCSI service in the registry:

1. Locate the key and set `eto=0;ldto=0`;
`HKLM\System\CurrentControlSet\Services\be2iscsi\Parameters\Device\DriverParameters`
2. Reboot the system for the registry changes to take effect.
3. Log into the iSCSI target using WMI. For more information, see “Logging into Targets for Multipath Support” on page 82.
4. Enable MPIO.
 - a. Select **Start>Administrative Tools>Server Manager**.
 - b. In the Server Manager tree, click **Features**.
 - c. In the Features area, click **Add Features**.
 - d. In the Add Features wizard on the Select Features page, select the **Multipath I/O** check box and click **Next**.
 - e. On the Confirm Installation Selections page, click **Install**.
 - f. When the installation is completed, click **Close** on the Installation Results page.
 - g. When prompted to restart the computer, click **Yes**.
 - h. Click **Close**.
5. Discover all possible paths to all devices on the system.
 - a. Open the MPIO control pane: select **Start>Administrative Tools>MPIO**.
 - b. On the User Account Control page, click **Continue**. The Properties dialog box is displayed.
 - c. Select the **Discover Multi-Paths** tab.
 - d. Select **Add support for iSCSI Devices** and click **Add**.

6. Reboot the system when prompted to do so.

After rebooting, the Microsoft iSCSI DSM claims all iSCSI discovered disks. The MPIO GUI shows device id MSFT2005iSCSIBusType_0x9 under the MPIO Devices tab. The Disk Manager does not show duplicate disks.

You can configure load balancing policies on the LUN from the Device Manager after you click on the disk and select the MPIO tab.

Logging into Targets for Multipath Support

After you have successfully installed and enabled MPIO support on a Windows Server, you must log in to the target. This section describes the steps to log into iSCSI targets through the WMI GUI. For information on using the iSCSISelect utility to log into an iSCSI target, see the *Emulex Boot Manual*.

To log in to a target using WMI:

1. Select the **Enable multi-path in the Log On to Target** window. This must be selected for every target to be logged in for MPIO. Use the Advanced tab to select the specific UCNA port to use for login.
2. After the target login is complete, select the target and click on the **Details** or **Properties** button (depending on the Windows operating system) to see the multiple sessions connected.

Maximum Transmission Unit (MTU) for iSCSI Connections

Because the Emulex OneConnect UCNA is a multi-function adapter, the MTU settings for iSCSI functions are different than the ones for NIC functions.

For iSCSI, there is no explicit way to configure MTU from the OneCommand Manager application. Instead this value is auto-negotiated by the firmware. Before establishing a TCP connection for an iSCSI Login, the iSCSI firmware issues an ICMP Echo with a large payload to the iSCSI target. If Jumbo Frames has been enabled on all the switches leading to the target, as well as, on the target interface and if there is a successful ICMP Echo reply, the iSCSI firmware uses Jumbo Frames for that connection. The MTU used in this case is 8342 bytes.

If the large ping request is unsuccessful, the firmware defaults to non-jumbo mode and uses an MTU size of 1514 bytes.

The Max MTU value is displayed in the OneCommand Manager application for the iSCSI controller under the Port Information Tab on the Max MTU field. The TCP MSS used for an active iSCSI connection is displayed in the OneCommand Manager application in the 'TargetSessions' screen on the TCPMSS field.

iSCSI Error Handling

The goal of iSCSI error handling is to be tolerant of link level and target level failures up to configured timeout values, so that I/O errors are not seen by the application or operating system. The error handling is triggered under the following conditions:

- Loss of immediate link to the UCNA (such as a cable disconnect or port failure). The UCNA firmware detects the loss of link and notifies the driver. When this happens, the driver queues the I/O requests internally, up to a configured timeout period, so that the operating system does not see I/O errors. This timeout period is known as link-down timeout (LDTO).
- Loss of connection to the target because of target or network disconnection at the target. If the driver has I/O requests pending with the target and the target becomes unavailable (because the target is down, has failed over, or network issues are detected at the target), the driver queues the I/O request internally up to a configured timeout period. This timeout period is known as extended timeout (ETO).

If the configured threshold for LDTO and ETO is reached and the UCNA is still unable to connect to the target, the driver fails all I/O requests. I/O errors are seen by the application and operating system.

Note: Following a link up, switch ports can take a long time to initialize and go to a forwarding state. Because of this, add additional time to the ETO and LDTO settings to eliminate I/O disruption or target unavailability. If the switch port is connected to a single host, then PortFast mode can be enabled on the switch port to eliminate delays in transitioning to a forwarding state.

Configuring LDTO and ETO on Windows Server

LDTO and ETO values are configured via the Windows registry. Table 2-9 lists the default values of LDTO and ETO on Windows Server and the limits within which they can be configured.

Note: If the ETO is set to a number between 0 and 19, the driver assumes the value to 20 seconds internally. You will not see any modification to the registry.

Table 2-9 LDTO and ETO Information on Windows Server

Value	Default	Minimum	Maximum
LDTO	20 sec	0 sec	3600 sec
ETO	90 sec	0 sec	3600 sec

To modify LDTO and ETO values, edit the Driver Parameters for the iscsi service in the registry. Locate the key:

```
HKLM\System\CurrentControlSet\Services\be2iscsi\Parameters\Device\
DriverParameters
```

and edit the REG_SZ parameter to `ldto=25;eto=50;`

This sets the default value of LDTO to 25 seconds and the default value of ETO to 50 seconds. The settings are applied the next time the driver is loaded. You must reboot the system (boot drivers) or disable the iSCSI driver and enable it again (non-boot drivers) in Device Manager for the settings to take effect.

Error Handling Under MultiPath (MPIO) and Cluster Configurations

In an MPIO or cluster configuration, fault tolerant software is present on the system in addition to the iSCSI driver's default error handling scheme. Depending on the type of failover configuration, the iSCSI driver's error handling parameter needs to be configured appropriately.

If the iSCSI target is in Active-Active failover mode, the iSCSI driver must report I/O errors as soon as they are detected. This allows the failover software to trigger a path failover to an active path or active node as quickly as possible. This configuration requires that the iSCSI driver's LDTO and ETO parameters be set to 0.

If the iSCSI target is in Active-Standby failover mode, then the iSCSI driver must wait for the target side failover operation to complete before reporting device unavailability to the operating system. For such configurations, the driver's ETO must be set to the amount of time the iSCSI target needs to complete its failover operation.

3. Troubleshooting

There are several circumstances in which your system may operate in an unexpected manner. The Troubleshooting section contains reference tables on event codes and error messages and provides information regarding unusual situations.

General Troubleshooting

Table 3-1 General Troubleshooting

Problem	Answer/Solution
The operating system fails to install or does not successfully install the driver.	Verify that the operating system is supported by the driver.
Windows Device Manager shows a code 10 or code 39 with a yellow or red exclamation point on the device.	The firmware image does not match the installed device drivers, or the firmware is corrupt. Using the OneCommand Manager application or one of the WinPE offline or online utilities, install a version of firmware that is compatible with the driver.
The firmware is corrupt or non-responsive.	Using the OneCommand Manager application or one of the WinPE offline or online utilities, install a version of firmware that is compatible with the driver.
The Emulex iSCSI BIOS banner is not displayed during system POST.	Configure the motherboard BIOS to enable the Option ROM for the PCI-E slot in which the UCNA is installed.

Troubleshooting the FC/FCoE Driver

Troubleshooting the Cisco Nexus Switch Configuration

Table 3-2 Cisco Nexus Switch Situations

Problem	Solution
<ol style="list-style-type: none"> 1) Windows creates the NTFS partition ok, but then reports that "The hard disk containing the partition or free space you chose has a LUN greater than 0. Setup cannot continue". (Dell 1850 server). 2) Windows reboots successfully, but then gets stuck during the GUI portion of the installation right from the beginning. (HP DL385G2 server). 	Set up the FCoE switch ports as follows: <ul style="list-style-type: none"> • no priority-flow-control mode on • untagged cos 0 • flowcontrol receive on • flowcontrol send on • spanning-tree port type edge

Event Trace Messages

ELS Log Messages (0100-0130)

lpfc_mes0100: FLOGI failure - ulpStatus: x%x, ulpWord[4]:x%x

Description	An ELS fabric login (FLOGI) command that was sent to the fabric failed.
Severity	Error
Log	LOG_ELS verbose
Action	Check the fabric connection.

lpfc_mes0101: FLOGI completes successfully - NPortId: x%x, RaTov: x%x, EdTov: x%x

Description	An ELS FLOGI command that was sent to the fabric succeeded.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

lpfc_mes0102: PLOGI completes to NPortId: x%x

Description	The adapter performed an N_port_login (PLOGI) into a remote NPort.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

lpfc_mes0103: PRLI completes to NPortId: x%x, TypeMask: x%x, Fcp2Recovery: x%x

Description	The adapter performed a process login (PRLI) into a remote NPort.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

lpfc_mes0104: ADISC completes to NPortId x%x

Description	The adapter performed a discover address (ADISC) into the remote NPort.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

lpfc_mes0105: LOGO completes to NPortId: x%x

Description	The adapter performed a LOGO into a remote NPort.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

lpfc_mes0112: ELS command: x%x, received from NPortId: x%x

Description	Received the specific ELS command from a remote NPort.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.
Remarks	lpfc_mes0114 and lpfc_mes0115 are also recorded for more details if the corresponding severity level is set. You can use the XRI to match the messages.

lpfc_mes0114: PLOGI chkparm OK

Description	Received a PLOGI from a remote NPORT and its FC service parameters match this adapter. Request can be accepted.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.
See Also	lpfc_mes0112

lpfc_mes0115: Unknown ELS command: x%x, received from NPortId: x%x\n

Description	Received an unsupported ELS command from a remote NPORT.
Severity	Error
Log	LOG_ELS verbose
Action	Check remote NPORT for potential problem.
See Also	lpfc_mes0112

lpfc_mes0128: Accepted ELS command: OpCode: x%x

Description	Accepted an ELS command from a remote NPORT.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

lpfc_mes0129: Rejected ELS command: OpCode: x%x

Description	Rejected ELS command from a remote NPORT.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

lpfc_mes0130: ELS command error: ulpStatus: x%x, ulpWord[4]: x%x

Description	ELS command failure.
Severity	Error
Log	LOG_ELS verbose
Action	Check remote NPORT for potential problem.

Discovery Log Messages (0202-0262)

lpfc_mes0202: Start Discovery: Link Down Timeout: x%x, initial PLOGICount:%d

Description	Device discovery/rediscovery after FLOGI, FAN or RSCN has started. TMO is the current value of the soft link time. It is used for link discovery against the LinkDownTime set in parameters. DISC CNT is number of nodes being discovered for link discovery. RSCN CNT is number of nodes being discovered for RSCN discovery. There will be value in either DISC CNT or RSCN CNT depending on which discovery is being performed.
Severity	Information
Log	LOG_DISCOVERY verbose
Action	No action needed, informational.

lpfc_mes0204: Discovered SCSI Target: WWN word 0: x%x, WWN word 1: x%x, DID: x%x:, RPI: x%x

Description	Device discovery found SCSI target.
Severity	Information
Log	LOG_DISCOVERY verbose
Action	No action needed, informational.

lpfc_mes0214: RSCN received: Word count:%d

Description	Received RSCN from fabric.
Severity	Information
Log	LOG_DISCOVERY verbose
Action	No action needed, informational.

lpfc_mes0215: RSCN processed: DID: x%x

Description	Processed RSCN from fabric.
Severity	Information
Log	LOG_DISCOVERY verbose
Action	No action needed, informational.

lpfc_mes0225: Device Discovery completes

Description	This indicates successful completion of device (re)discovery after a link up.
Severity	Information
Log	LOG_DISCOVERY verbose
Action	No action needed, informational.

lpfc_mes0229: Assign SCSIId x%x to WWN word 0: x%x, WWN word 1: x%x, NPortId x%x

Description	The driver assigned a SCSI ID to a discovered mapped FCP target. BindType - 0: DID 1:WWNN 2:WWPN
Severity	Information
Log	LOG_DISCOVERY verbose
Action	No action needed, informational.

lpfc_mes0230: Cannot assign SCSIId to WWN word 0: x%x, WWN word 1: x%x, NPortId x%x

Description	SCSI ID assignment failed for discovered target.
Severity	Warning
Log	LOG_ELS verbose
Action	Review system configuration.

lpfc_mes0232: Continue discovery at sequence number%d, PLOGIs remaining:%d

Description	NPort discovery sequence continuation.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

lpfc_mes0235: New RSCN being deferred due to RSCN in process

Description	An RSCN was received while processing a previous RSCN.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

lpfc_mes0236: Issuing command to name server" type: x%x

Description	The driver is issuing a nameserver request to the fabric. Also recorded if a GID_FT is sent.
Severity	Information
Log	LOG_DISCOVERY verbose
Action	No action needed, informational.
See Also	lpfc_mes0239 or lpfc_mes0240

lpfc_mes0238: NameServer response DID count:%d

Description	Received a response from fabric name server with N DIDs.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

lpfc_mes0239: NameServer Response: next DID value: x%x

Description	The driver received a nameserver response. And, this message is recorded for each DID included in the response data.
Severity	Information
Log	LOG_DISCOVERY verbose
Action	No action needed, informational.
See Also	lpfc_mes0236

lpfc_mes0240: NameServer Response Error - CmdRsp:x%x, ReasonCode: x%x, Explanation x%x

Description	The driver received a nameserver response containing a status error.
Severity	Error
Log	LOG_DISCOVERY verbose
Action	Check Fabric configuration. The driver recovers from this and continues with device discovery.
See Also	lpfc_mes0236

lpfc_mes0256: Start node timer on NPortId: x%x, timeout value:%d

Description	Starting timer for disconnected target with NPort ID and timeout value.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

lpfc_mes0260: Stop node timer on NPortId: x%x, SCSIId: x%x

Description	Discontinuing timer for reconnected target with NPort ID and SCSI ID.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

lpfc_mes0262: Node timeout on NPortId: x%x, SCSIId: x%x

Description	Disconnected NPort ID, SCSI ID has failed to reconnect within timeout limit.
Severity	Error
Log	LOG_ELS verbose
Action	Review system configuration.

Mailbox Log Messages (0310-0326)

lpfc_mes0310: Mailbox command timeout - HBA unresponsive

Description	A Mailbox command was posted to the adapter and did not complete within 30 seconds. sync - 0: asynchronous mailbox command is issued 1: synchronous mailbox command is issued.
Severity	Error
Log	LOG_MBOX verbose
Action	This error could indicate a software driver or firmware problem. If no I/O is going through the adapter, reboot the system. If these problems persist, report these errors to Technical Support.

lpfc_mes0326: Reset HBA - HostStatus: x%x

Description	The adapter has been reset.
Severity	Information
Log	LOG_MBOX verbose
Action	No action needed, informational.

INIT Log Messages (0400-0463)

lpfc_mes0400: Initializing discovery module: OptionFlags: x%x

Description	Driver discovery process is being initialized with internal flags as shown.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

lpfc_mes0401: Initializing SLI module: DeviceId: x%x, NumMSI:%d

Description	PCI function with device id and MSI count as shown is being initialized for service level interface.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

lpfc_mes0405: Service Level Interface (SLI) 2 selected\n");

Description	Service Level Interface level 2 is selected.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

lpfc_mes0406: Service Level Interface (SLI) 3 selected\n");

Description	Service Level Interface level 3 is selected.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

lpfc_mes0436: Adapter not ready: hostStatus: x%x

Description	The adapter failed during powerup diagnostics after it was reset.
Severity	Error
Log	LOG_INIT verbose
Action	This error could indicate a hardware or firmware problem. If problems persist report these errors to Technical Support.

lpfc_mes0442: Adapter failed to init, CONFIG_PORT, mbxStatus x%x

Description	Adapter initialization failed when issuing CONFIG_PORT mailbox command.
Severity	Error
Log	LOG_INIT verbose
Action	This error could indicate a hardware or firmware problem. If problems persist report these errors to Technical Support.

lpfc_mes0446: Adapter failed to init, CONFIG_RING, mbxStatus x%x

Description	Adapter initialization failed when issuing CFG_RING mailbox command.
Severity	Error
Log	LOG_INIT verbose
Action	This error could indicate a hardware or firmware problem. If problems persist report these errors to Technical Support.

lpfc_mes0454: Adapter failed to init, INIT_LINK, mbxStatus x%x

Description	Adapter initialization failed when issuing INIT_LINK mailbox command.
Severity	Error
Log	LOG_INIT verbose
Action	This error could indicate a hardware or firmware problem. If problems persist report these errors to Technical Support.

lpfc_mes0458: Bring Adapter online

Description	The FC driver has received a request to bring the adapter online. This may occur when running HBAnyware.
Severity	Warning
Log	LOG_INIT verbose
Action	None required.

lpfc_mes0460: Bring Adapter offline

Description	The FC driver has received a request to bring the adapter offline. This may occur when running HBAnyware.
Severity	Warning
Log	LOG_INIT verbose
Action	None required.

lpfc_mes0463: Adapter firmware error: hostStatus: x%x, Info1(0xA8): x%x, Info2 (0xAC): x%x

Description	The firmware has interrupted the host with a firmware trap error.
Severity	Error
Log	LOG_INIT verbose
Action	Review HBAnyware diagnostic dump information.

FCP Log Messages (0701-0749)**lpfc_mes0701: Issue Abort Task Set to PathId: x%x, TargetId: x%x, Lun: x%x**

Description	The driver has issued a task management command for the indicated SCSI device address.
Severity	Warning
Log	LOG_INIT verbose
Action	Review system configuration.

lpfc_mes0703: Issue LUN reset to PathId: x%x, TargetId: x%x, Lun: x%x, Did: x%x

Description	Storport is requesting a reset of the indicated LUN.
Severity	Warning
Log	LOG_INIT verbose
Action	Review system configuration. Possible side-effect of cluster operations.

lpfc_mes0713: Issued Target Reset to PathId:%d, TargetId:%d, Did: x%x

Description	Storport detected that it needs to abort all I/O to a specific target. This results in login reset to the target in question.
Severity	Warning
Log	LOG_FCP verbose
Action	Review system configuration. Possible side-effect of cluster operations.
See Also	lpfc_mes0714

lpfc_mes0714: Issued Bus Reset for PathId:%d

Description	Storport is requesting the driver to reset all targets on this adapter.
Severity	Warning
Log	LOG_FCP verbose
Action	Review system configuration. Possible side-effect of cluster operations.
See Also	lpfc_mes0713

lpfc_mes0716: FCP Read Underrun, expected%d, residual%d

Description	FCP device provided less data than was requested.
Severity	Supplement Information
Log	LOG_FCP verbose
Action	No action needed, informational.
See Also	lpfc_mes0730

lpfc_mes0729: FCP command error: ulpStatus: x%x, ulpWord[4]: x%x, XRI: x%x, ulpWord[7]: x%x

Description	The specified device failed an I/O FCP command.
Severity	Warning
Log	LOG_FCP verbose
Action	Check the state of the target in question.
Remarks	lpfc_mes0730 is also recorded if it is a FCP Rsp error.

lpfc_mes0730: FCP response error: Flags: x%x, SCSI status: x%x, Residual:%d

Description	The FCP command failed with a response error.
Severity	Warning
Log	LOG_FCP verbose
Action	Check the state of the target in question.
Remark	lpfc_mes0716, lpfc_mes0734, lpfc_mes0736 or lpfc_mes0737 is also recorded for more details if the corresponding SEVERITY level is set.
See Also	lpfc_mes0729

lpfc_mes0734: Read Check: fcp_parm: x%x, Residual x%x

Description	The issued FCP command returned a Read Check Error.
Severity	Warning
Log	LOG_FCP verbose
Action	Check the state of the target in question.
See Also	lpfc_mes0730

lpfc_mes0737: SCSI check condition, SenseKey x%x, ASC x%x, ASCQ x%x, SrbStatus: x%x

Description	The issued FCP command resulted in a Check Condition.
Severity	Warning
Log	LOG_FCP verbose
Action	Review SCSI error code values.
See Also	lpfc_mes0730

lpfc_mes0747: Target reset complete: PathId: x%x, TargetId: x%x, Did: x%x

Description	A target reset operation has completed.
Severity	Warning
Log	LOG_FCP verbose
Action	Review system configuration. Possible side-effect of cluster operations.
Remark	See also Message 0713.

lpfc_mes0748: Lun reset complete: PathId: x%x, TargetId: x%x, Lun: x%x

Description	A LUN reset operation has completed.
Severity	Warning
Log	LOG_FCP verbose
Action	Review system configuration. Possible side-effect of cluster operations.
Remark	See also Message 0703.

lpfc_mes0749: Abort task set complete: Did: x%x, SCSId: x%x

Description	A task management has completed.
Severity	Warning
Log	LOG_FCP verbose
Action	Review system configuration. Possible side-effect of cluster operations.
Remark	See also Message 0701.

Link Log Messages (1302-1306)

lpfc_mes1302: Invalid speed for this board:%d, forced link speed to auto

Description	The driver is re-initializing the link speed to auto-detect.
Severity	Warning
Log	LOG_LINK_EVENT verbose
Action	None required.

lpfc_mes1303: Link Up event: tag: x%x, link speed:%dG, topology (0 = Pt2Pt, 1 = AL):%d

Description	A link up event was received. It is also possible for multiple link events to be received together.
Severity	Error
Log	LOG_LINK_EVENT verbose
Action	If numerous link events are occurring, check physical connections to the FC network.
Remarks	lpfc_mes1304 is recorded if Map Entries > 0 and the corresponding mode and SEVERITY level is set.

lpfc_mes1305: Link down even: tag x%x

Description	A link down event was received.
Severity	Error
Log	LOG_LINK_EVENT verbose
Action	If numerous link events are occurring, check physical connections to the FC network.

lpfc_mes1306: Link Down timeout

Description	The link was down for greater than the configuration parameter (HLinkTimeOut) seconds. All I/O associated with the devices on this link will be failed.
Severity	Warning
Log	LOG_LINK_EVENT verbose
Action	Check adapter cable/connection to SAN.

Tag Messages (1400-1401)

lpfc_mes1400: Tag out of range: ContextIndex: x%x, MaxIndex: x%x, ulpCommand: x%x

Description	Firmware has generated an invalid response.
Severity	Error
Log	LOG_LINK_EVENT verbose
Action	Review hardware configuration. Contact Emulex Technical Support.

lpfc_mes1401: Invalid tag: ContextIndex: x%x, ulpCommand: x%x

Description	Firmware has generated an invalid response.
Severity	Error
Log	LOG_LINK_EVENT verbose
Action	Review hardware configuration. Contact Emulex Technical Support.

NPIV Messages (1800-1899)

lpfc_mes1800: NPIV FDISC failure VPI: x%x Error x%x Reason x%x

Description	Virtual Port fails on a FDISC to the switch with the error and reason listed.
Severity	Error
Log	LOG_NPIV verbose
Action	Check to ensure the switch supports NPIV.

lpfc_mes1801: Memory allocation failure for NPIV port: x%x

Description	Fails to allocated the block of memory for the Virtual Port.
Severity	Error
Log	LOG_NPIV verbose
Action	Check to ensure system has sufficient kernel memory.

lpfc_mes1802: Exceeded the MAX NPIV port: x%x

Description	Exceeded the number of Virtual Port allows on the adapter.
Severity	Error
Log	LOG_NPIV verbose
Action	Reduce the number of Virtual Ports.

lpfc_mes1803: Virtual Port: x%x VPI:x%x successfully created.

Description	Virtual Port ID is successfully created.
Severity	Information
Log	LOG_NPIV verbose
Action	No action needed, informational.

lpfc_mes1804: Removing Virtual Port: x%x VPI:x%x

Description	Removing Virtual Port ID.
Severity	Information
Log	LOG_NPIV verbose
Action	No action needed, informational.

ELS Messages (1900-1999)

lpfc_mes1900: x%x sends ELS_AUTH_CMD x%x with TID x%x

Description	An ELS_AUTH_CMD is sent.
Severity	Information
Log	LOG_FCSP verbose
Action	No action needed, informational.

lpfc_mes1901: x%x sends ELS_AUTH_REJECT x%x x%x to x%x

Description	An ELS_AUTH_REJECT is sent.
Severity	Information
Log	LOG_FCSP verbose
Action	No action needed, informational.

lpfc_mes1902: Receives x%x from x%x in state x%x

Description	Receives an ELS_AUTH_CMD.
Severity	Information
Log	LOG_FCSP verbose
Action	No action needed, informational.

lpfc_mes1903: Receives ELS_AUTH_RJT x%x x%x

Description	Receives an ELS_AUTH_REJECT.
Severity	Information
Log	LOG_FCSP verbose
Action	No action needed, informational.

lpfc_mes1904: Authentication ends for x%x with status x%x (%d %d)

Description	Authentication is done.
Severity	Information
Log	LOG_FCSP verbose
Action	No action needed, informational.

lpfc_mes1905: Authentication policy change for local x%08x x%08x remote x%08x%08x

Description	Authentication policy has been changed.
Severity	Information
Log	LOG_FCSP verbose
Action	No action needed, informational.

Troubleshooting the NIC/TOE Drivers

The following table provides troubleshooting information for the NIC/TOE drivers.

Table 3-3 Troubleshooting the NIC/TOE Drivers

Problem	Answer/Solution
Performance is not as expected.	<p>The UCNA may be installed in a PCIe slot other than an 8x PCIe slot or a Gen2 slot. Verify that the UCNA has been properly installed and that the system is properly set for 8x Gen2 PCIe operation.</p> <p>If TOE is enabled and performance is not as high as expected, the operating system may not offload TOE connections. For more information, see "Improving Performance with TCP Offload" on page 71.</p>
There are frequent event log entries for link changes, or statistics that show more than expected CRC errors.	<p>Unload and reload the driver to reset available target IDs. Ensure that the SAN configuration is correct prior to reloading the driver. This clears the driver's consistent binding table and frees target IDs for new target nodes.</p>
The driver fails to load, and an event log entry states that the driver failed to load due to memory constraints.	<p>There may not be enough memory installed in the system to provide sufficient memory for all devices installed in the system. Try installing more memory if possible.</p>
Unpredictable results occur when flow control setting differs among PCI functions.	<p>If multiple PCI functions are exposed for a single 10-Gb/s Ethernet port, such as in blade configurations, the flow control parameter must be set the same on all adapters for the port.</p> <p>Results are unpredictable if the setting differs among PCI functions because this is a shared property of the 10-Gb/s port.</p>
On servers that support PCIe hot unplug, the system may hang or produce a bugcheck if a PCIe hot unplug or replace is attempted.	<p>Hot unplug is not supported in this release.</p>
When Wake on LAN is set through the advanced properties page, the system does not wake when receiving a magic packet or a packet that would normally wake the system.	<p>The system may not support Wake on LAN on the PCIe slot in which the UCNA is installed. Check the system documentation to determine whether the system is capable of Wake on LAN operation.</p> <p>A system BIOS setting may not be correct for Wake on LAN to work as expected. Check the system documentation to determine whether Wake on LAN must be enabled in the system BIOS.</p> <p>The system may not go to a lower power state because another software component, device, or driver is preventing it from doing so.</p> <p>Microsoft provides several useful references for troubleshooting Wake on LAN configuration issues in the Microsoft TechNet Library on the Microsoft website.</p>

Table 3-3 Troubleshooting the NIC/TOE Drivers (Continued)

Problem	Answer/Solution
When running Windows Server 2008, the computer restarts and shows various Stop error codes when performing one of the following operations: <ul style="list-style-type: none"> Enabling or disabling TCP Chimney Offload Changing the network adapter settings Upgrading the NIC drivers 	Apply the 979614 hotfix as described on the Microsoft website.
When an NDIS driver is being installed manually on a Windows 2008 system, the installer installs the first driver it finds, even if it is not the latest version of the driver.	Windows 2008 picks up the first available driver it finds when an NDIS driver is being installed manually. Thus, an NDIS5 driver will be installed even if a Windows NDIS6 driver is available. An event log message advises you to update to the latest driver for best performance.
The system crashes or appears to hang. In the case of a hang, there could be a message indicating that the driver experienced a hardware malfunction.	<p>There are several possible causes for this problem.</p> <ul style="list-style-type: none"> Certain systems require an updated BIOS to properly manage the power states of newer Intel and AMD processors. Check with your OEM for information regarding BIOS and firmware updates that may be required to run well with the latest releases of the Windows operating systems. Also, certain BIOS settings may be required. For example, it is recommended that you disable any low power processor states and low power settings for PCI-Express. On certain AMD systems, it is possible the intelppm.sys driver is enabled, and should not be. To query this system driver's run state, log in as administrator and at the command line type: <pre>sc query intelppm</pre> <p>If the results indicate that the intelppm driver is running, you must disable it. At the command line type:</p> <pre>sc config intelppm start= disabled</pre> <p>On all systems, it may be necessary to set the power options to High Performance. See the operating system documentation for details.</p>

Monitoring TCP Offloads

To monitor TCP offloads, in a command window type:

```
netstat -t
```

This command indicates the offload state for each TCP connection of the system.

Windows Server 2008 (and later versions) allows TCP offloads in more scenarios than previous versions of Windows Server. In particular, TCP offloads may occur with the Windows firewall enabled.

TCP Offload Failure

The following table lists common reasons why TCP offloads do not occur and their suggested fixes.

Table 3-4 Troubleshooting TCP Offload Failures

Reasons for No TCP Offload	Solutions
Chimney offload is disabled on the system.	<p>For Windows Server 2008 and Windows Server 2008 R2 To determine whether Chimney offload is enabled or disabled, at the command line type:</p> <pre>netsh interface tcp show global</pre> <p>To enable Chimney offload, at the command line type:</p> <pre>netsh interface tcp set global chimney=enabled</pre> <p>To disable Chimney offload, at the command line type:</p> <pre>netsh interface tcp set global chimney=disabled</pre> <p>For Windows Server 2008 and Windows Server 2008 R2 To verify whether offloading is enabled type:</p> <pre>netstat -nt</pre> <p>This command displays a list of connections and their offloading state.</p>
Offloads are disabled for specific ports or applications.	<p>To view any TCP ports or applications that may be configured to disable TCP offload, at the command line type:</p> <pre>netsh interface tcp show chimneyports netsh interface tcp show chimneyapplications</pre>
A third-party firewall is running.	The Windows firewall does not affect TCP offload, but third-party firewalls may prevent TCP offloads. Uninstall third-party firewall software to allow TCP offloads.
In the network properties, some intermediate drivers prevent offloading.	Go to Network Connections > Properties and clear check boxes for unused drivers. In particular, Network Load Balancing and some third-party drivers prevent offloads.
IPSec is enabled.	Disable IPSec.
IP Network Address Translation (NAT) is enabled.	Disable IP NAT.
The driver supports an Advanced Property to disable TCP offloading.	Make sure TCP offloading is enabled.
The TCP connection uses using IPv6.	The driver supports offloading TCP connections only with IPv4.

Note: Packet sniffing applications such as Ethereal or Microsoft Network Monitor, do not see TCP offloaded packets.

Troubleshooting the iSCSI Driver

The following table provides troubleshooting information for the iSCSI driver.

Table 3-5 Troubleshooting the iSCSI Driver

Problem	Answer/Solution
Overall failure.	Use the iSCSISelect utility to clear the Adapter Configuration. See the <i>Emulex Boot Manual</i> for instructions.
The driver loads successfully, but there are event 11 entries in the event log for the iSCSI driver.	The most common cause is that the UCNA link is down. See “iSCSI Error Log on Windows Server 2008” on page 124 and look for specific event codes to confirm.
Unable to create a memory dump file on a system booted over iSCSI.	Make sure the disk has enough free disk space to create the dump file. If a full memory dump is selected, the disk must have free space at least equivalent to the amount of physical memory in the system.
Unable to log in to target from WMI.	<ul style="list-style-type: none"> • Ensure that the IP address on the UCNA is valid and the network connection has been set up to reach the target. • If login is attempted after discovering the target, ensure that the correct UCNA port has been selected for the login.
The iSCSI WMI GUI shows the target state as connected, but no LUNs are seen from the disk manager.	Verify that the UCNA name used to connect to the target matches the UCNA name configured on the iSCSI target.
Multipath configuration shows duplicate LUNs on the disk manager.	Ensure that MPIO software is installed and the login options have selected the MPIO flag. On Windows Server 2008 and Windows Server 2008 R2 Operating Systems, the server role must be set up for Multipath. See the <i>Emulex Boot Manual</i> for more information on MPIO.
Multipath configuration takes a long time to fail over or failover does not occur.	Ensure that link down timeout (LDTO) settings and extended timeout (ETO) settings have been configured for MPIO. These values must be set to 0. For more information, see “Configuring LDTO and ETO on Windows Server” on page 83.
Sendtargets to an IET target fails because it violates the iSCSI specification.	If you still want to add an IET target, you must add the target manually. This issue affects Sendtargets only.
<p>The following POST error message appears:</p> <p>Initiator iSCSI Name mismatch, Please use iSCSISelect to set a single name for all controllers.</p> <p>Press <Ctrl><S> to enter iSCSISelect. (Reboot required)</p>	In the iSCSI BIOS, the Emulex iSCSI initiator name may be different if more than one OneConnect UCNAs are in the system. This message appears if the iSCSI initiator name is different on multiple controllers. You must enter iSCSISelect and save a new initiator name on the first iSCSISelect menu window so that the iSCSI initiator name on all controllers match. All logins from the multiple controllers will use the new name. See the <i>Emulex Boot Manual</i> for more information.

Table 3-5 Troubleshooting the iSCSI Driver (Continued)

Problem	Answer/Solution
<p>When an <code>iscsicli logouttarget</code> command is issued back-to-back in a script, event 12 errors from the PlugPlayManager are seen in the Windows Event Viewer. The error message is similar to this string:</p> <p>The device 'SE ISCSI 00 SCSI Disk Device' (SCSI\Disk&Ven_SE_ISCSI&Prod_00&Rev_3.64\5&17659873&2&020000) disappeared from the system without first being prepared for removal.</p>	<p>This behavior is not specific to the OneConnect UCNA.</p>
<p>On a system running Windows 2008, or Windows 2008 R2, the iSCSI driver fails to load after many iterations of enable/disable from Device Manager. Because the system failed to allocate contiguous uncached extension memory, the iSCSI driver failed to load, and an attention icon is displayed next to the OneConnect iSCSI device. The Device Status shows "This device cannot start. (Code 10)", and an Event 11 error is logged in the Windows system event log for the iSCSI driver with 0x31840006 in the 5th DWORD.</p>	<p>There is no workaround for this issue.</p>
<p>When an iSCSI UCNA is used to log in to an iSCSI target and the LUN configuration on the target is changed, neither the UCNA nor the WMI GUI see the updated LUN configuration.</p>	<p>If an iSCSI target provides an asynchronous event notification to the UCNA when its logical unit inventory has changed, the iSCSI driver initiates a bus rescan and the LUNs are updated dynamically. However, if an iSCSI target does not provide an asynchronous event notification, the LUN list is not updated dynamically. Perform a manual rescan in Disk Management.</p>
<p>A login to new target fails after Microsoft iSCSI Initiator Service is installed.</p>	<p>When Microsoft iSCSI software is installed, the service chooses a default IQN name for the UCNA. The Microsoft iSCSI service issues the request to the iSCSI driver via the WMI interface to set this new IQN name. Therefore, any IQN name that was configured earlier (such as by using iSCSISelect) will be overridden and the new IQN name will be in effect.</p> <p>Although this will not affect existing boot sessions and persistent sessions, new target logins could fail because the new IQN name does not match the incoming IQN name configured on the target.</p> <p>After the Microsoft iSCSI Initiator Service is installed, the initiator name must be renamed to the previous name configured from the WMI GUI.</p>

Table 3-5 Troubleshooting the iSCSI Driver (Continued)

Problem	Answer/Solution
<p>When software-based iSCSI targets are logged into the UCNA, Event ID 56 (Driver SCSI (000000)). Appears in the Windows event viewer. This issue has been observed on Windows 2008 R2 under the following conditions:</p> <ul style="list-style-type: none">• The iSCSI target is a software-based target (MSiSCSI, IET, StarWind) that uses a local hard drive or a RAM disk for its backend LUN.• Different UCNA ports are involved in the login.• A SAS controller is present on the system.	<p>This occurs caused because of an issue with the data reported by the iSCSI target in the Product Identification field in response to the standard inquiry from the UCNA. This field should be unique among different targets' LUNs, but software-based targets report the same pre-formatted data for all the LUNs across all targets. When Windows encounters the same Product Identification field for different LUNs with the same Bus Target Lun field, it records error in the event log. No other effect has been found as a result of this behavior.</p> <p>The workaround for this error is to use non-overlapping LUN numbers for the various LUNs across the various iSCSI targets. On the iSCSI target system, LUNs can be numbered sequentially; they do not have to start at zero.</p>

Appendix A. Error and Event Log Information

FC/FCoE Error and Event Logs

Viewing the FC/FCoE Error Log

The system event log is a standard feature of Windows Server software. All events logged by the Emulex Storport Miniport will be Event ID 11 with source "elxstor".

To view the error LOG:

1. Open the Event Viewer window:
 - Click **Start>Programs>Administrative Tools>Event Viewer**
 - or-
 - Right-click **My Computer**, **Manage** and **Event Viewer** in **Computer Management**.

The Event Viewer window is displayed.

2. Double-click any event with the source name ELXSTOR.
3. Examine the entry at offset 0x10 and Event ID 11. The Emulex event code is found in byte 0010 and supplementary data is in the byte offsets 0011 through 0013 (in example Figure A-1, byte 0010 = 9b, byte 0011 = 00, byte 0012 = 29 and byte 0013 = 00).




Figure A-1 Event Properties

Severity Scheme

When the Event Viewer is launched, there are three branches: Application, Security and System. All ELXSTOR error log entries are found under the System branch and all ELXSTOR error log entries have the Event Viewer severity level of “error”.

- A severe error code indicates that the driver, firmware or adapter is behaving abnormally and your intervention is required to correct the problem.
- A malfunction error code indicates that there is a problem with the system, but your intervention is not required.
- A command error code indicates that an event has transpired, but does not require your intervention. An event may be problem-oriented, such as an invalid fabric command sub-type. An event may not be problem-oriented, such as exhausted retries on PLOGI or PDISC.

Related Driver Parameter: LogError

The LogError driver parameter determines the minimum severity level to enable entry of a logged error into the system. See the Configuration Section for instructions on how to set driver parameters.

- If set to 0 = all errors regardless of severity are logged.
- If set to 1 = severe, malfunction and command level errors are logged.
- If set to 2 = both severe and malfunction errors are logged.
- If set to 3 = only severe errors are logged.

Note: Set LogError to 1 if you are troubleshooting SAN connectivity or device discovery issues.

Format of an Error Log Entry

An error log entry will take the form of an event. This event is described by:

- Date (date entry was logged)
- Source (elxstor)
- Time (time entry was logged)
- Category (none)
- Type (error)
- Event id (0)
- User (N/A)
- Computer (name of computer)

Error Codes Tables

Severe Errors

Table A-1 Severe Errors

Bits 0 - 7	Interpretation
0x00	Invalid link speed selection (SLI2-3 mode)
0x01	READ_REV failed (SLI2-3 mode)
0x02	Invalid adapter type (LightPulse)
0x03	Invalid adapter type (LightPulse)
0x04	CONFIG_PORT failed
0x06	READ_CONFIG_failed
0x07	CONFIG_RING 0 failed
0x09	CONFIG_RING 1 failed
0x08	CONFIG_RING 2 failed
0x0A	CONFIG_RING 3 failed
0x0B	INIT_LINK failed (SLI2-3 mode)
0x0C	INIT_LINK failed (SLI2-3 mode)
0x0D	READ_REV failed (SLI2-3 mode)
0x0E	Invalid adapter type (LightPulse)
0x0F	Invalid adapter type (LightPulse)
0x10	CONFIG_PORT failed (reinitialization)
0x12	READ_CONFIG command failed (reinitialization)
0x13	CONFIG_RING 0 failed (reinitialization)
0x14	CONFIG_RING 1 failed (reinitialization)
0x15	CONFIG_RING 2 failed (reinitialization)
0x16	CONFIG_RING 3 failed (reinitialization)
0x17	Unresponsive adapter port (SLI2-3 mode)
0x1C	Firmware trap: info1 (SLI2-3 mode)
0x1D	Firmware trap: info2 (SLI2-3 mode)
0x1E	Over-temperature error condition (LightPulse)
0x1F	Firmware-initiated adapter port reset (LightPulse)
0x20	Adapter port error attention (LightPulse)
0x22	Over-temperature warning (LightPulse)
0x23	Returned to safe temperature (LightPulse)
0x24	Invalid response tag (SLI2-3 mode)

Table A-1 Severe Errors (Continued)

Bits 0 - 7	Interpretation
0x25	Invalid response tag (SLI2-3 mode)
0x26	Invalid response tag (SLI2-3 mode)
0x27	Invalid response sequence (SLI2-3 mode)
0x28	Failure on REG_LOGIN mailbox command
0x29	Unable to initiate fabric binding operation
0x2A	Attempted ADISC to non-existent node
0x2B	Failure on iocb context allocation
0x2C	Unable to initiate nport unbinding operation
0x2D	Unable to initiate nport binding operation
0x30	Failure on mailbox context allocation
0x7C	Menlo initialization error
0x7D	Menlo initialization error
0x7E	Menlo initialization error
0xA0	Failed to initialize adapter port (OneConnect)
0xA1	Failed to initialize adapter port (LightPulse)
0xC0	Insufficient revision level for STORPORT.SYS
0xC1	Failed to allocate miniport un-cached extension
0xC2	Insufficient un-cached extension space
0xC3	Port initialization failure (OneConnect)
0xC4	Port initialization failure (LightPulse)
0xC5	Utility mailbox command error
0xC6	SLI4 Pre-initialization failure
0xD3	NPIV memory allocation failure
0xF0	Unresponsive adapter port (SLI4 mode)
0xF4	ULP Unrecoverable Error: low part (SLI4 mode)
0xF5	ULP Unrecoverable Error: high part (SLI4 mode)
0xF6	ARM Unrecoverable Error (SLI4 mode)
0xF7	READ_NV failed (SLI4 mode)
0xF8	READ_NV failed (SLI4 mode)
0xF9	READ_REV failed (SLI4 mode)
0xFA	READ_CONFIG failed (SLI4 mode)
0xFB	Failed to post header templates (SLI4 mode)
0xFC	Invalid Completion Queue Entry (SLI4 mode)

Table A-1 Severe Errors (Continued)

Bits 0 - 7	Interpretation
0xFD	Invalid Completion Queue Entry (SLI4 mode)
0xFE	Invalid Completion Queue Entry (SLI4 mode)

Malfunction Errors

Table A-2 Malfunction Errors

Bits 0 - 7	Interpretation
0x05	SET_VAR command failed
0x11	SET_VAR command failed (reinitialization)
0x21	Spurious mailbox command interrupt
0x31	Unrecognized mailbox command completion
0x32	Duplicate link attention: event tag unchanged
0x33	Invalid link attention: no link state indicated
0x34	Duplicate link attention: link state unchanged
0x35	Error reading common service parameters for port
0x36	Error reading common service parameters for fabric
0x37	Error reading common service parameters for nport
0x3B	Failed to create node object
0x3C	PRLI initiation failure
0x42	Exhausted retries on FLOGI
0x45	ELS command rejected
0x49	Exhausted retries on PLOGI
0x4E	World Wide Port Name mismatch on ADISC
0x4F	World Wide Node Name mismatch on ADISC
0x50	ADISC response failure
0x55	LOGO response failure
0x57	PRLI to non-existent node
0x5A	PRLI response error
0x5F	CT command error
0x62	Name server response error
0x66	State Change Notification registration failure
0x6A	Unrecognized ELS command received
0x6F	Received PRLI from un-typed source
0x73	Failed to pend PRLI for authentication

Table A-2 Malfunction Errors (Continued)

Bits 0 - 7	Interpretation
0x77	Failed to allocate Node object
0x7A	REG_VPI failed
0xA3	Command context allocation failure
0xAB	SCSI command error
0xAC	Read check error
0xB0	Node timeout: device removal signaled to Storport

Command Errors

Table A-3 Command Errors

Bits 0 - 7	Interpretation
0x43	Fabric login succeeded
0x46	ELS command failed
0x47	Exhausted retries on ELS command
0x4A	PLOGI accepted
0x56	LOGO accepted
0x59	PRLI accepted
0x63	Fabric name server response
0x6B	ELS RSCN processed
0x71	LOGO received from fabric
0x79	FDISC accepted
0xA2	SCSI address assigned to discovered target
0xA4	Report LUNs error (initial I/O to discovered target)
0xA5	Local error indication on FCP command
0xA8	Data overrun
0xA9	FCP command error
0xAA	SCSI check condition
0xAD	Local reject indication on FCP command
0xAE	Error on SCSI pass-through command
0xAF	Error on Menlo CT command

Event Indicators

Table A-4 Event Indications

Bits 0 - 7	Interpretation
0x18	Port shutdown event (LightPulse)
0x19	Port in off-line state (LightPulse)
0x1A	Port in on-line state (LightPulse)
0x1B	Port in off-line state (LightPulse)
0xA7	Data underrun
0xD0	NPIV Virtual Port creation success (Virtual Port Did in bits 8-31)
0xD1	NPIV Virtual Port creation failed (Virtual Port index in bits 8-31)
0xD2	NPIV Virtual Port FDISC failed (Virtual Port index in bits 8-31)
0xD4	Exceeded max Virtual Port supported (Virtual Port index in bits 8-31)
0xD5	NPIV Virtual Port removal (Virtual Port Did in bits 8-31)
0xE0	Authenticated successfully (remote Did in bits 8-31)
0xE1	Failed to authenticate (remote Did in bits 8-31)
0xE2	Authentication not support (remote Did in bits 8-31)
0xE3	Authentication ELS command timeout (remote Did in bits 8-31)
0xE4	Authentication transaction timeout (remote Did in bits 8-31)
0xE5	LS_RJT other than Logical Busy received for Authentication transaction (remote Did in bits 8-31)
0xE6	LS_RJT Logical Busy received for Authentication Transaction (remote Did in bits 8-31)
0xE7	Received Authentication Reject other than Restart (remote Did in bits 8-31)
0xE8	Received Authentication Reject Restart (remote Did in bits 8-31)
0xE9	Received Authentication Negotiate (remote Did in bits 8-31)
0xEA	Authentication spurious traffic (remote Did in bits 8-31)
0xEB	Authentication policy has been changed (remote Did in bits 8-31)
0xED	Same passed were set for both local and remote entities (remote Did in bits 8-31)
0xF1	Port shutdown event (OneConnect)
0xF2	Port in off-line state (OneConnect)
0xF3	Port in on-line state (OneConnect)

Viewing the FC/FCoE Event Log

Event Log Interpretation

- All events logged by Emulex Storport Miniport are in Event ID 11 with source "elxstor".
- The Storport Miniport driver parameter LogErrors determines what type of events are logged by the driver; the default setting is "3" which logs only events of a SEVERE nature; the optional setting of "2" logs events of both SEVERE and MALFUNCTION type; the optional setting of "1" logs events of SEVERE, MALFUNCTION and COMMAND type.

Note: For troubleshooting SAN connectivity or device discovery issues, set the LogErrors to 1.

- The Emulex event code is found in byte 0010 and supplementary data is in byte offsets 0011 through 0013.

Additional Event Log Information

The following tables are not comprehensive but do include those codes, which through Emulex's experiences in our support and testing environments, we feel are most likely to show up in SAN environments where problems occur.

ELS/FCP Command Error Status Codes

Internal firmware codes posted by the adapter firmware that explain why a particular ELS or FCP command failed at the FC level.

Table A-5 ELS/FCP Command Error Status Codes

Explanation	Code
Remote Stop - Remote port sent an ABTS	0x2
Local Reject - Local Reject error detail	0x3
LS_RJT Received - Remote port sent LS_RJT	0x9
A_RJT Received - Remote port sent BA_RJT	0xA

CT Command Response Codes

Codes that indicate the response to a FC Common Transport protocol command.

Table A-6 CT Command Response Codes

Explanation	Code
FC Common Transport Reject	0x8001
FC Common Transport Accept	0x8002

FC-CT Reject Reason Codes

Codes that indicate the reason a CT command was rejected.

Table A-7 FC-CT Reject Reason Codes

Explanation	Code
Invalid command code	0x01
Invalid version level	0x02
Logical busy	0x05
Protocol error	0x07

ELS Command Codes

FC protocol codes that describe what particular Extended Link Services command was sent.

Table A-8 ELS Command Codes

Explanation	Code
Link Service Reject (LS_RJT)	0x01
Accept (ACC)	0x02
N_Port Login (PLOGI)	0x03
Fabric Login (FLOGI)	0x04
N_Port Logout (LOGO)	0x05
Process Login (PRLI)	0x20
Process Logout (PRLO)	0x21
Discover F_Port Service Params (FDISC)	0x51
Discover Address (ADISC)	0x52
Register State Change Notify (RSCN)	0x61

SCSI Status Codes

The SCSI status returned from a SCSI device which receives a SCSI command.

Table A-9 SCSI Status Codes

Explanation	Code
GOOD	0x00
CHECK CONDITION	0x02
BUSY	0x08
RESERVATION CONFLICT	0x18
QUEUE FULL	0x28

Local Reject Status Codes

Codes supplied by the Emulex adapter firmware which indicate why a command was failed by the adapter.

Table A-10 Local Reject Status Codes

Explanation	Code
SEQUENCE TIMEOUT - Possible bad cable/link noise	0x02
INVALID RPI - Occurs when link goes down	0x04
NO XRI - Possible host or SAN problem	0x05
TX_DMA FAILED - Possible host system problem	0x0D
RX_DMA FAILED - Possible host system problem	0x0E
ILLEGAL FRAME - Possible bad cable/link noise	0x0F
NO RESOURCES - Port out of exchanges or logins	0x11
LOOP OPEN FAILURE - FC_AL port not responding	0x18
LINK DOWN - Queued cmds returned at link down	0x51A
OUT OF ORDER DATA - Possible bad cable or noise	0x1D

SRB Status Codes

SCSI Request Block status provided by the driver to the operating system based upon response from SCSI device in the SAN.

Table A-11 SRB Status Codes

Explanation	Code
ERROR	0x04
BUSY	0x05
TIMEOUT	0x09
SELECTION TIMEOUT	0x0A
COMMAND TIMEOUT	0x0B
BUS RESET	0x0E
DATA OVERUN	0x12

ASC/ASCQ

Additional Sense Code/Additional Sense Code Qualifier information can be found in any SCSI specification document – these codes contain detailed information about the status/condition of the SCSI device in question.

Additional Notes on Selected Error Codes

These are error codes which may be seen more frequently than others or which indicate conditions that you might be able to solve by investigation and correction of problems in the SAN configuration.

Note: The nomenclature of "0x" is used as the prefix for the byte code fields because those byte codes are actually hex values.

Node Timeout (Code 0xAA)

This event code indicates that a particular device has not been found (if the message is logged during device discovery) or that a particular device has been removed from the fabric. If this message is seen, determine if there is something wrong with the connection of that device to the SAN (cables, switches or switch ports, status of the target device itself).

SCSI Command Error (0x9A) and SCSI Check Condition (code 0x9B)

Code 0x9A indicates that the SCSI command to a particular device was responded to with an error condition (the target and LUN information, along with the SCSI status, are provided).

In the specific case of code 0x9B, this code indicates that the device responded with the specific status of Check Condition – the ASC/ASCQ information provided in bytes 0x12 and 0x13 will allow you to find out what status is being reported by the target and determine if there is an action that can be performed to return the device to functional status.

Nameserver Response (Code 0x98)

This code is useful in determining if the expected number of targets in a SAN configuration are being presented by the nameserver to the requesting adapter. The number in byte 0x11 is the number of targets returned to the nameserver query made by the adapter – if the number of targets does not match expectations, examine the SAN configuration found in the switch tables and if that information shows targets or devices still missing, check connections between the switch ports and those devices.

Context Allocation Failures

There are a number of event codes for which the interpretation contains the phrase "context allocation failure" – these types of events are referring to the internal memory constructs of the Emulex Storport Miniport driver and as such are intended for Emulex design engineer's information. If a customer encounters such an event, they should contact Emulex support for analysis and determination if that particular event may be an indicator of a failed adapter or of some problem with interaction between the adapter, the Emulex Storport Miniport driver, the host operating system and the host memory.

Note: Context allocation failures are rare.

NIC/TOE Error and Event Logs

Viewing the NIC/TOE Error Log

For Windows Server operating systems, the OneConnect network driver generates error codes in the system event log. These error codes can be viewed by using the Event Viewer application.

To view the error codes:

1. Click the **Start** tab on the bottom of the screen.
2. Click **Run**.
3. Type **eventvwr** and click **OK**.
4. Click **Windows Log**.
5. Click **System**.
6. Click the be2net error under System Events to show the event details.

NIC/TOE Event Log

Windows Device Manager generates error log codes if any errors occur during the installation of the NIC driver. Each log contains a Message Id, Severity and Symbolic Link. The Message Id is unique and tracks the error message (if not displayed). Table A-12 shows the list of error codes, the severity of the error, the message displayed, the meaning of the error and recommended resolutions. When reporting a problem with the UCNA to Emulex, check the event log and report any of these entries that may be present.

Table A-12 NIC/TOE Event Log Entries

Message ID Hexadecimal/ Decimal	Severity	Message	Recommended Resolution
0x0002f/47	Warning	Virtual Machine Queue (VMQ) offload is disabled. Disable SR/IOV support in PXE BIOS to use VMQ.	Disable SR-IOV support in PXE BIOS to use VMQ.
0x0002e/46	Error	Device is not supported on Windows 7 Operating System.	
0x0002d/45	Error	Error recovery failed. The device is no longer operational. Update all drivers and firmware.	See the Downloads section of the Emulex website for compatible firmware and drivers.
0x0002c/44	Warning	Error recovery is disabled on the system. The device is no longer operational.	This message is informational.
0x0002b/43	Informational	The driver successfully recovered from an error.	This message is informational.
0x0002a/42	Warning	Legacy driver loaded. Move to the NDIS 6.20 driver for Windows 2008 R2 for best performance.	

Table A-12 NIC/TOE Event Log Entries (Continued)

Message ID Hexadecimal/ Decimal	Severity	Message	Recommended Resolution
0x0029/41	Warning	Legacy driver loaded. Move to the NDIS 6.x driver for Windows 2008 for best performance.	
0x0028/40	Warning	The firmware is outdated and does not support TOE offloads for this driver. Update the firmware.	The firmware and the driver are not compatible versions. See the Downloads section of the Emulex website for compatible firmware and drivers.
0x0026/38	Warning	The device firmware does not support RSS functionality for this network adapter.	The firmware and the driver are not compatible versions. See the Downloads section of the Emulex website for compatible firmware and drivers.
0x0025/37	Warning	The device firmware does not support TCP offload functionality.	The firmware and the driver are not compatible versions. See the Downloads section of the Emulex website for compatible firmware and drivers.
0x0024/36	Error	The device firmware does not support network functionality.	The firmware and the driver are not compatible versions. See the Downloads section of the Emulex website for compatible firmware and drivers.
0x0023/35	Warning	The Ethernet link is down due to a remote fault.	The Ethernet link is down due to the remote partner signaling a fault. Check the peer device for errors.
0x0022/34	Warning	The Ethernet link is down due to a local fault.	The Ethernet link is down due to a link-down event detected at the OneConnect driver.
0x0021/33	Informational	Network device is operating in Gen2 mode and installed in a 4x PCI-Express slot.	For best performance, install the UCNA in an 8x Gen2 PCIe slot. Note: A 16x slot will not provide any additional performance.
0x0020/32	Informational	The network device is operating in Gen2 mode and installed in a 1x PCI-Express slot.	For best performance, install the UCNA in an 8x Gen2 PCIe slot. Note: A 16x slot will not provide any additional performance.
0x001f/31	Informational	The network device is operating in Gen1 mode and installed in a 8x PCI-Express slot.	For best performance, install the UCNA in an 8x Gen2 PCIe slot. Note: A 16x slot will not provide any additional performance.
0x001e/30	Informational	The network device is operating in Gen1 mode and installed in a 4x PCI-Express slot.	For best performance, install the UCNA in an 8x Gen1 PCIe slot. Note: A 16x slot will not provide any additional performance.

Table A-12 NIC/TOE Event Log Entries (Continued)

Message ID Hexadecimal/ Decimal	Severity	Message	Recommended Resolution
0x001d/29	Informational	The network device is operating in Gen1 mode and installed in a 1x PCI-Express slot.	For best performance, install the UCNA in an 8x Gen1 PCIe slot. Note: A 16x slot will not provide any additional performance.
0x001c/28	Error	Vital product data is not initialized correctly.	Use the offline flash utility to reconfigure the device.
0x000c/12	Informational	The Ethernet link is down.	This message is informational.
0x000b/11	Informational	The Ethernet link is up.	This message is informational.
0x000a/10	Error	The network device detected an error.	A hardware error occurred. Verify that the firmware flash image is not corrupted. Contact Emulex Technical Support.
0x0015/21	Error	Firmware version does not match driver version.	The firmware version and driver must match. This is a warning message, but it is recommended that you reinstall matching versions of the firmware and driver.
0x0015/21	Warning	Firmware version does not match driver version.	The firmware version and driver must match. This is a warning message, but it is recommended that you reinstall matching versions of the firmware and driver.
0x0014/20	Error	Failed to read registry configuration.	The registry is corrupted. Reinstall the driver and/or operating system.
0x0013/19	Error	Resource conflict.	The operating system failed to allocate resources for the device. Check low memory conditions and operating system hardware resource conflicts.
0x0012/18	Error	Failed to enable bus mastering.	Verify that the BIOS allows bus mastering and that no resource conflicts exist.
0x0011/17	Error	The driver is incompatible with the device.	The driver is loaded on the incorrect hardware device. Verify that the correct driver is installed.
0x0010/16	Warning	The network driver was reset.	This may indicate a system hang or hardware problem. Verify other system devices are working properly.
0x0009/9	Error	Failed to register interrupt service routine.	This is an NDIS error. Verify that hardware resource conflicts do not exist.

Table A-12 NIC/TOE Event Log Entries (Continued)

Message ID Hexadecimal/ Decimal	Severity	Message	Recommended Resolution
0x0008/8	Error	Failed to get TCP offload handlers.	This is an NDIS error. Verify the NDIS version is valid for the driver.
0x0007/7	Warning	A memory allocation failure occurred during driver load. Performance may be reduced.	This warning occurred due to a failed memory allocation. Check low memory conditions. Use a smaller MTU or disable TCP offload to reduce driver memory requirements.
0x0006/6	Error	Driver load failed due to memory allocation failure	This failure occurred due to a failed memory allocation in the driver. Check low memory conditions.
0x0005/5	Error	Failed to register scatter gather DMA.	This failure occurred due to a failed memory allocation in the operating system. Check low memory conditions.
0x0004/4	Error	Failed to map device registers.	This failure occurred due to a failed memory allocation in the operating system. Check low memory conditions.
0x0003/3	Error	Unsupported medium.	This is an internal NDIS error. Check the operating system installation.
0x0002/2	Error	The network driver initialization failed.	This may be a firmware driver mismatch or corrupt installation. Check the firmware version, reinstall the firmware and try again. This may also indicate a hardware problem.
0x0001/1	Informational	The driver successfully loaded.	This message is informational and indicates successful loading of the device driver.

iSCSI Error and Event Log

Viewing the iSCSI Error and Event Log on Windows Server 2008

The iSCSI driver generates error codes in the system event log in the form of Event ID 11 errors. These error codes can be viewed by using the Event Viewer application.

To view the error codes:

1. Click the **Start** tab on the bottom of the screen.
2. Click **Run**.
3. Type **eventvwr** and click **OK**.
4. Click **Windows Log**.
5. Click **System**.
6. Click the be2iscsi error under System Events to show the details of the event.

The iSCSI driver logs errors with the port driver error code of SP_INTERNAL_ADAPTER_ERROR, which translates to an Event ID 11 entry in the system event log.

The following is an example of the iSCSI driver error code 0x11800003 viewed with the Event Viewer application. The window shows the driver generated error code in the fifth DWORD (offset 0x10) of the word dump.

Note: To improve the visibility of the error code in the Data field of the Event Properties window, select the Words option.

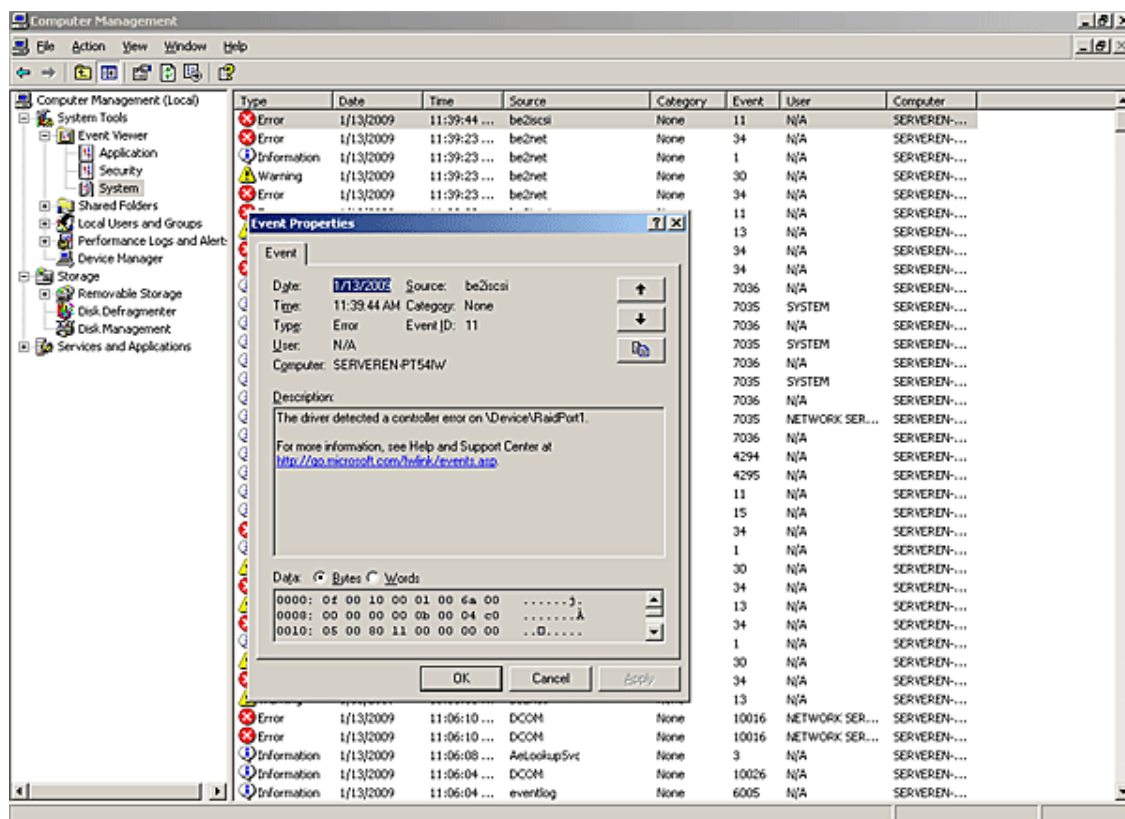


Figure A-2 iSCSI Error

Because the UCNAs are dual PCI-function adapters, the `\Device\RaidPort<n>` value changes depending on the device that observed the error.

iSCSI Error Log on Windows Server 2008

The following is a brief description of the error log codes generated by the iSCSI driver for Windows Server 2008. It includes the error code, the message displayed, and the meaning of the message and the recommended resolution.

Table A-13 iSCSI Error Log Entries on Windows Server 2008

Message ID	Message	Description/Recommended Resolution
0x348d0008	The iSCSI driver failed a WMI IOCTL request from the port driver because the request was failed by the ARM firmware. This error is immediately followed by another error code entry indicating the WMI request code in error.	This failure indicates that an operation attempted from the Microsoft WMI application resulted in an error. Check the operation being attempted for user or configuration errors.

Table A-13 iSCSI Error Log Entries on Windows Server 2008 (Continued)

Message ID	Message	Description/Recommended Resolution
0x348d0007	The iSCSI driver failed a WMI IOCTL request from the port driver. This error is immediately followed by another error code entry indicating the WMI request code in error.	This failure indicates that an operation attempted from the Microsoft WMI application resulted in an error. Check the operation being attempted for user or configuration errors.
0x33900002	The iSCSI driver failed an I/O request because it failed to retrieve a scatter gather list from the Storport driver.	This failure occurred due to a failed memory allocation in the operating system. Check low memory conditions.
0x31880001	The iSCSI driver failed to load because initialization failed during a power management bootup.	This failure may be due to the firmware not being present or currently running. This failure may also indicate a hardware problem.
0x3184000c	The iSCSI driver was unable to map one or more PCI Base Address Registers and failed to load.	This failure may indicate a low memory condition or a hardware error.
0x3184000b	The iSCSI driver ignored a configuration entry because the entry was invalid.	The invalid entry must be removed or corrected. Check the registry configuration for any new valid values added to the driver parameters. For more information on valid driver values, see "Driver Parameters for Windows" on page 90.
0x31840009	The iSCSI driver failed to load a configuration value specified in the registry because the value was out of range. The driver will use the default value for this configuration parameter instead.	The range specified for a configuration parameter is too large or too small and must be corrected. Check the registry configuration for any new valid values added to the driver parameters. For more information on valid driver values, see "Driver Parameters for Windows" on page 90.
0x31840006	The iSCSI driver failed to load due to memory allocation failure.	This failure occurred due to a failed memory allocation in the driver. Check low memory conditions.
0x31840001	The iSCSI driver failed to load because initialization failed during normal bootup.	This failure may be due to the firmware not being present or currently running. This failure may also indicate a hardware problem.
0x31640004	An internal API failed in the iSCSI driver during initialization.	This failure may indicate a low memory condition.
0x3164000D	The driver failed to allocate its complete memory requirement and will attempt to load with reduced capabilities. Total number of targets available will be reduced.	This message indicates a low memory condition.
0x14831000	There was an Unrecoverable Error detected by the iSCSI driver. Following this error log entry, the next 3 entries indicate the error codes.	This may be due to hardware errors or due to unhandled exceptions in the hardware or firmware.

Table A-13 iSCSI Error Log Entries on Windows Server 2008 (Continued)

Message ID	Message	Description/Recommended Resolution
0x138e0103	The iSCSI driver failed an IOCTL request because the number of scatter gather elements required for the IOCTL buffer exceeded the firmware limit. Following this error log entry, the next entry will indicate the IOCTL opcode and the payload length requested.	This error may indicate an incorrect configuration option for the iSCSI driver. It may also indicate a low memory condition.
0x138d0101	The iSCSI driver detected an error offloading the iSCSI connection. The operation will be retried again. Following this error log entry, the next entry will indicate the session handle and the firmware error code.	This may indicate a target is in error or may point to transient network connectivity issues. It may also indicate a firmware error.
0x12990013	The iscsi driver did not receive an iSCSI command window update within 25 seconds during I/O operations. Following this error log entry, the next entry will indicate the session handle where this error occurred. The iSCSI driver will trigger a session recovery on the session and continue.	<ul style="list-style-type: none"> Verify that the iSCSI target is certified by Microsoft. Check for errors reported at the target. The Emulex iSCSI UCNA is only supported with certified targets. Check for software updates at the target vendor's website. If applicable, update the software. Check for software updates at the Emulex website. If applicable, update the software.
0x127b0012	The iSCSI driver received an invalid iSCSI Command Sequence Number update from the target. Following this error log entry, the next three entries will indicate the session handle and the iSCSI parameters - MaxCmdSN and ExpCmdSN respectively.	<ul style="list-style-type: none"> Verify that the iSCSI target is certified by Microsoft. Check for errors reported at the target. The Emulex iSCSI UCNA is only supported with certified targets. Check for software updates at the target vendor's website. If applicable, update the software. Check for software updates at the Emulex website. If applicable, update the software.
0x12790006	A connection to the target was lost for a period exceeding the Extended Timeout (ETO). The error log entry immediately following this entry will indicate the session ID of the target that lost the connection. There will be event log entries from the disk subsystem indicating that the drives were lost. If any I/Os were in progress, the system may see I/O errors or failures.	Check the connection to the target or the state of the target device. If the target is made available, any sessions that existed previously are reestablished and the devices are available for I/O.
0x11990007	The iSCSI driver received a Task Management Function (TMF) that is not supported and rejected this request. The error log entry immediately following this entry will indicate the TMF function code that was rejected.	The operating system version is not supported.

Table A-13 iSCSI Error Log Entries on Windows Server 2008 (Continued)

Message ID	Message	Description/Recommended Resolution
0x11940008	The iSCSI driver received a TMF Abort request for an I/O request that is not present with the driver.	This may indicate a slow connection to the target. Check network connectivity to the target for any errors.
0x1184000B	Firmware returned invalid data in its configuration. iSCSI login and offload are disabled.	Reload the firmware.
0x11840002	The iSCSI driver encountered a mismatched version of the firmware running on the board. This error may be followed by error codes 0x31840001 or 0x31880001 indicating that the iSCSI driver failed to load.	This failure indicates that the driver version that is running on the system does not match the version of the firmware on the board. Correct this by running the installer from the desired version.
0x11840001	The iSCSI driver detected a failure in the hardware during initialization. This error may be followed by error codes 0x31840001 or 0x31880001 indicating that the iSCSI driver failed to load.	This failure indicates that the hardware has not been initialized or is malfunctioning. This may also indicate that the firmware is not running correctly.
0x11800005	Both Port 0 and Port 1 links were down for a period exceeding the Link Down Timeout (LDT0). If the UCNA has connection to the target, there will be event log entries from the disk subsystem indicating that the drives were lost. If any I/Os were in progress, the system may see I/O errors or failures.	Check the links to the UCNA. If the link is reestablished, any sessions that previously existed are reestablished and the devices is available for I/O.
0x11800003	Both Port 0 and Port 1 links are down.	Check the links to the UCNA.
0x31840005	Driver load failed because the PCI Vendor ID and Device ID are not supported.	Check the configuration on the UCNA.
0x1180000A	The logical link on the OneConnect Port is down, traffic is disallowed on this function.	The iSCSI function may have been disabled in the PXESelect application. If you disabled it intentionally, you can ignore this message.

Viewing the iSCSI Error Log on Windows Server 2008 R2 and 2012

The iSCSI driver on the Windows Server 2008 R2 and 2012 operating systems support the new event logging mechanism introduced by Storport. Custom event messages are logged for a variety of events with different severity, such as informational, warning or error. The source of the events indicates the service name and every event includes a unique ID and a symbolic name.

Table A-14 iSCSI Error Log Entries on Windows Server 2008 R2 and 2012

Message ID	Severity	Message	Recommended Resolution
0x02	Info	Driver loaded successfully.	N/A
0x04	Error	Firmware version does not match with driver version.	The driver version that is running on the system does not match the version of the firmware on the UCNA. Install a driver that is compatible with the firmware.
0x05	Warning	Port link is down, check connection to UCNA.	Check the links to the UCNA.
0x06	Info	Port link is up.	N/A
0x07	Error	Link down timeout expired on the port, all targets are lost.	The link on the UCNA is down for a period exceeding the link down timeout (LDTO) value. If the UCNA has a connection to the target, event log entries from the disk subsystem indicate that the drives were lost. If any I/O was in progress, the system may see I/O errors or failures. Check the links to the UCNA. If the link is reestablished, any sessions that previously existed are reestablished and the devices are available for I/O.
0x08	Error	Target with session id N failed to connect within the configured timeout.	A connection to the target was lost for a period exceeding the Extended Timeout (ETO). The error log entry includes the session ID of the target that lost the connection. Event log entries from the disk subsystem indicate that the drives were lost. If any I/O was in progress, the system may see I/O errors or failures. Check the connection to the target or the state of the target device. If the target is made available, any sessions that previously existed are reestablished and the devices are available for I/O.
0x09	Error	Task Management request N was unhandled.	The iSCSI driver received a Task Management Function that is not supported, and it rejected this request. An application or service that is installed on the system may not be compatible with the driver.

Table A-14 iSCSI Error Log Entries on Windows Server 2008 R2 and 2012 (Continued)

Message ID	Severity	Message	Recommended Resolution
0x0a	Error	Task Management Function abort was received on a task that is not present.	<p>The iSCSI driver received a Task Management Function Abort request for an I/O request that is not present with the driver.</p> <p>This may indicate a slow connection to the target. Check network connectivity to the target for any errors.</p>
0x0b	Error	Error in determining firmware configuration.	<p>An error in determining the firmware configuration occurred. The firmware on the UCNA may not be functioning properly. Check the UCNA and reinstall the firmware if required.</p>
0x0e	Warning	iSCSI error was detected on session A, ExpCmdSn B, MaxCmdSn C.	<p>The iSCSI driver received an invalid iSCSI Command Sequence Number update from the target. The event log entry indicates the session handle, MaxCmdSN, and ExpCmdSN.</p> <ul style="list-style-type: none"> • Verify that the iSCSI target is certified by Microsoft. Check for errors reported at the target. The Emulex iSCSI UCNA is only supported with certified targets. • Check for software updates at the target vendor's website. If applicable, update the software. • Check for driver and firmware updates at the Emulex website. If applicable, update the driver and firmware.
0x0f	Warning	The iSCSI target on session id N failed to open the command window within configured timeout.	<p>The iSCSI driver did not receive an iSCSI command window update for up to 25 seconds during I/O operations. The event log entry indicates the session handle on which the error occurred. The iSCSI driver triggers a session recovery on the session and continues.</p> <ul style="list-style-type: none"> • Verify that the iSCSI target is certified by Microsoft. Check for errors reported at the target. The Emulex iSCSI UCNA is only supported with certified targets. • Check for software updates at the target vendor's website. If applicable, update the software. • Check for driver and firmware updates at the Emulex website. If applicable, update the driver and firmware.
0x10	Warning	Encountered an error offloading an iSCSI connection, error code N.	<p>The iSCSI driver detected an error while offloading the iSCSI connection. The operation is retried up to five times. The session handle and the UCNA firmware error code are included in the event log message.</p> <p>This may indicate a target is in error or it may point to transient network connectivity issues. It may also indicate a UCNA firmware error.</p>

Table A-14 iSCSI Error Log Entries on Windows Server 2008 R2 and 2012 (Continued)

Message ID	Severity	Message	Recommended Resolution
0x11	Warning	The IOCTL opcode A requires more scatter gather elements than allowed. Transfer length is B.	The iSCSI driver failed an IOCTL request because the number of scatter/gather elements required for the IOCTL buffer exceeded the UCNA firmware limit. The IOCTL opcode and the payload length requested are included in the event log entry. This error may indicate an incorrect configuration option for the iSCSI driver. It may also indicate a low memory condition.
0x12	Error	Unrecoverable error detected. UE Low: A UE High: B FW Line: C.	An unrecoverable error was detected by the iSCSI driver. This may be caused by hardware errors or by unhandled exceptions in the hardware or firmware.
0x13	Error	Hardware initialization failed, failing driver load	The iSCSI driver detected a failure in the hardware during initialization. This failure indicates that the hardware has not been initialized or is malfunctioning. This may also indicate that the firmware is not running correctly.
0x14	Warning	Failed to retrieve scatter gather list for an SRB, an IO has failed.	The iSCSI driver failed an I/O request because it failed to retrieve a scatter/gather list from the Storport driver. This failure occurred because of a failed memory allocation in the operating system. Check low memory conditions.
0x15	Error	ACIT library table initialization failed.	An internal API failed in the iSCSI driver during initialization. This failure may indicate a low memory condition.
0x16	Error	An ACIT API failed.	An internal API failed in the iSCSI driver during initialization. This failure may indicate a low memory condition.
0x17	Error	Unsupported hardware, failing driver load.	Driver loading failed because the PCI Vendor ID and Device ID are not supported. Check the UCNA configuration.
0x18	Error	Memory could not be allocated, failing driver load.	This failure occurred because of a failed memory allocation in the driver. This failure may indicate a low memory condition

Table A-14 iSCSI Error Log Entries on Windows Server 2008 R2 and 2012 (Continued)

Message ID	Severity	Message	Recommended Resolution
0x19	Warning	WMI driver error, code A.	<p>The iSCSI driver failed a WMI IOCTL request from the port driver. The event log entry includes the WMI request code in error.</p> <p>An operation attempted from the Microsoft WMI application resulted in an error. Check the operation being attempted for errors.</p>
0x1a	Warning	WMI IOCTL error, code A.	<p>The iSCSI driver failed a WMI IOCTL request from the port driver because the request was failed by the ARM firmware. The event log entry includes the request code in error.</p> <p>An operation attempted from the Microsoft WMI application resulted in an error. Check the operation being attempted for errors.</p>
0x1b	Warning	A configuration parameter is out of range.	<p>The iSCSI driver failed to load a configuration value specified in the registry because the value was out of range. The driver uses the default value for this configuration parameter.</p> <p>The range specified for a configuration parameter is either too large or too small, and it must be corrected. Check the registry configuration for new driver parameter entries. See "Driver Parameters for Windows" on page 90 for the correct range of values</p>
0x1d	Warning	A configuration parameter is invalid.	<p>The iSCSI driver ignored a configuration entry because the entry was invalid.</p> <p>Check the registry configuration for new driver parameter entries. The invalid entry must be removed or corrected.</p> <p>See "Driver Parameters for Windows" on page 90 for the correct range of values.</p>
0x1e	Error	Failed to map Base Address Register, failing driver load.	<p>The iSCSI driver was unable to load because it was unable to map one or more PCI Base Address registers.</p> <p>This failure may indicate a low memory condition or a hardware error.</p>
0x1f	Error	Hardware initialization has failed - error code A.	<p>The hardware initialization has failed. This error causes the driver load to fail. The error code included in the event log entry identifies the specific point of failure.</p> <p>This failure indicates that the hardware has not been initialized or is malfunctioning. This may also indicate that the firmware is not running correctly.</p>

Table A-14 iSCSI Error Log Entries on Windows Server 2008 R2 and 2012 (Continued)

Message ID	Severity	Message	Recommended Resolution
0x20	Warning	Initial memory allocation failed, driver is running with reduced capabilities.	The driver failed to allocate its complete memory requirement and attempts to load with reduced capabilities. The total number of targets available is reduced. This message indicates a low memory condition.
0x21	Info	Target Reconnected for Session id N.	N/A
0x22	Info	Interrupt Redirection capability is enabled.	N/A
0x23	Warning	Interrupt Redirection capability is not supported by this firmware. Update your firmware.	Update the firmware to the latest version.
0x24	Info	Interrupt Redirection capability is not supported by this hardware.	N/A
0x26	Warning	Logical link on the OneConnect Port is down, traffic is disallowed on this function.	The iSCSI function may have been disabled in the PXESelect application. If you disabled it intentionally, you can ignore this message.
0x27	Error	Firmware returned invalid data in its configuration. iSCSI login and offload are disabled.	Reload the firmware.
0x28	Warning	Error Recovery is not being attempted. Adapter is no longer functional.	An Unrecoverable Error (UE) has occurred, but UE recovery is not enabled. A system reboot is required to make the adapter operational again.

Appendix B. Configuring iSCSI through DHCP

Dynamic Host Configuration Protocol (DHCP) Recommendations

If you are using the DHCP server to obtain an IP address for the UCNA, Emulex recommends that you set up a reservation. A reservation assigns a specific IP address based on the MAC address of the UCNA. If you do not reserve an IP Address through DHCP, then you must set the lease length for the UCNA IP address to unlimited. This prevents the IP address lease from expiring.

Vendor-Specific Option 43

This section describes the format for the data returned in DHCP vendor-specific option 43. The method and format for specifying the Vendor ID is outside the scope of this document and is not included here. The UCNA offers this Vendor ID to the DHCP server to retrieve data in the format described in this section.

Format of Vendor-Specific Option 43

The following describes the format of option 43 and includes guidelines for creating the data string:

```
'iscsi':'<TargetIP>':'<TargetTCPPort>':'<LUN>':'<TargetName>':'<InitiatorName>':'<HeaderDigest>':'<DataDigest>':'<AuthenticationType>
```

- Strings shown in quotes are part of the syntax and is therefore mandatory
- Fields enclosed in angular brackets (including the angular brackets) should be replaced with their corresponding values. Some of these fields are optional and may be skipped.
- If an optional field is skipped, a colon must be used as a placeholder to indicate the default value for that field.
- When specified, the value of each parameter should be enclosed in double quotes. See “Examples” on page 134.
- All options are case sensitive.

Description of Mandatory and Optional Parameters

Table B-1 describes the parameters used in the data string for option 43.

Table B-1 Data String Parameters for Option 43

Parameter	Description	Field Type
<TargetIP>	A valid IPv4 address in dotted decimal notation	Mandatory
<TargetTCPPort>	A decimal number ranging from 1 to 65535 (inclusive). The default TCP port is 3260.	Optional
<LUN>	A hexadecimal representation of the LUN of the boot device. By default, LUN 0 is assumed to be the boot LUN. It is an eight-byte number which should be specified as a hexadecimal number consisting of 16 digits, with an appropriate number of zeroes padded to the left, if required.	Optional
<TargetName>	A valid iSCSI name of up to 223 characters.	Mandatory
<InitiatorName>	A valid iSCSI iqname of up to 223 characters. If it is not provided, the default initiator name (generated by the UCNA based on its MAC address) is used.	Optional
<HeaderDigest>	Either E (the header digest is enabled) or D (the header digest is disabled).	Optional
<DataDigest>	Either E (the data digest is enabled) or D (the data digest is disabled).	Optional
<AuthenticationType>	D (authentication is disabled), E (one-way CHAP is enabled; the user name and secret must be specified by non-DHCP means), or M (mutual-CHAP is enabled; the user name and passwords must be specified by non-DHCP means). D is the default setting.	Optional

Examples

The following is an example of default initiator name and data digest settings:

```
iscsi:"192.168.0.2":"3261":"0000000000000000E":"iqn.2009-4.com:1234567890"::"E"::"E"
```

- Target IP address: 192.168.0.2
- Target TCP port: 3261
- Target boot LUN: 0x0E
- Target iqname: iqn.2009-04.com:1234567890
- Initiator name: Not specified. Use the initiator name that is already configured, or use the default name if no initiator name is already configured.
- Header digest: Enabled
- Data digest: Not specified; assume disabled.
- Authentication type: One-way CHAP.

The following is an example of default TCP port and mutual-CHAP settings:

```
iscsi:"192.168.0.2"::"0000000000000000E":"iqn.2009-4.com:1234567890"  
::"E":"D":"M"
```

- Target IP address: 192.168.0.2
- Target TCP port: Use default from RFC 3720 (3260)
- Target boot LUN: 0x0E
- Target iqn name: iqn.2009-04.com:1234567890
- Initiator name: Not specified. Use the initiator name that is already configured, or use the default name if no initiator name is already configured.
- Header digest: Enabled
- Data digest: Disabled
- Authentication type: Mutual-CHAP

Appendix C. Port Speed Specifications

A UCNA can support only one Ethernet port speed at a time, and the preference is always for 10 Gb/s. The type of module used (copper or optical) does not make a difference. As soon as a 10-Gb module is plugged into one of the ports, the UCNA switches to a 10-Gb speed no matter what speed the other port is running, or even if I/O is running on that port. This behavior is a per-UCNA constraint; another UCNA can be running on a different speed.

Table C-1 lists negotiated speed specifications per an adapter's port connection:

Table C-1 Negotiated Speed Specification per Adapter Port Connection

Port 0	Port 1	Port Link	Status Speed
10 Gb/s	10 Gb/s	Both ports link up	10 Gb/s
10 Gb/s	1Gb/s	Only Port 0 links up	10 Gb/s
1Gb/s	10 Gb/s	Only Port 1 links up	10 Gb/s
1Gb/s	1Gb/s	Both ports link up	1 Gb/s
1Gb/s	-	Only Port 0 links up	1 Gb/s
-	1Gb/s	Only Port 1 links up	1 Gb/s
10 Gb/s	-	Only Port 0 links up	10 Gb/s
-	10 Gb/s	Only Port 1 links up	10 Gb/s

Negotiating Speed on a Mezzanine Card

A mezzanine card retains the first negotiated speed. This could be either 10 Gb/s or 1 Gb/s, depending on the switch connected. To change the speed on a mezzanine card:

1. Remove the switch from both the ports.
2. Insert the switch on one port and wait for the link to come up.
3. After the link is up, insert the switch on the other port.

The mezzanine card retains the speed of the first link until both links are down.

Appendix D. AutoPilot Installer Command Line and Configuration File Parameters

The AutoPilot Installer can initiate an installation from a command prompt or script. You can pass command line parameters to the AutoPilot Installer, yet still have it run automatically by the driver kit, if you run the driver kit installer from a command prompt or script.

If you specify the `/q` switch with the driver kit installer command, the driver kit installer runs in unattended mode and automatically invokes the APInstall.exe with its `/silent` switch. See “Unattended Driver Installation” on page 19 for more information.

AParg Driver Kit Parameter and Appending to the APInstall.exe File

If you specify a value for the “APargs” driver kit parameter, this value is appended to the APInstall.exe command line. For example, if you execute this installer file as:

```
elxdrv-fc-fcoe<version>.exe /q APargs=SilentRebootEnable=True
```

then after installing the AutoPilot Installer, the driver kit automatically executes it as:

```
APInstall.exe /silent SilentRebootEnable=True
```

To specify more than one parameter, separate the settings by one or more spaces and put quotes around the entire APargs expression. For example, the command line (all on one line):

```
elxdrv-fc-fcoe<version>.exe "APargs=SilentRebootEnable=True  
localDriverLocation = "d:\drivers\new\Storport"
```

results in the AutoPilot Installer being run as:

```
APInstall.exe SilentRebootEnable=True localDriverLocation =  
"d:\drivers\new\Storport"
```

Parameter values that contain spaces, such as path names, must be enclosed in quotes. To add such a setting to APargs, you must insert backslashes before the quotes around the value, and then add quotes around the entire APargs expression. For example, the command line (all on one line):

```
elxdrv-fc-fcoe<version>.exe "APargs=ReportLocation=\"C:\Documents  
and Settings\Administrator\My Documents\reports\""
```

results in AutoPilot Installer being run as:

```
APInstall.exe ReportLocation="C:\Documents and  
Settings\Administrator\My Documents\reports"
```

If you have many parameters to pass to the AutoPilot Installer, or if you want to do so repeatedly, then it may be less error prone to run the utility kit installer interactively, delay AutoPilot Installer execution, and then run the AutoPilot Installer command yourself. The procedure for doing so is described in “Option 2: Run the AutoPilot Installer Separately” on page 16 and “Unattended Driver Installation” on page 19.

You can specify a non-default directory for the driver kit by specifying an 'installation folder' on the command line. For example:

```
elxdrv-fc-fcoe<version>.exe install:"C:\Emulex"
```

This option can be used in conjunction with the “APArgs” directive.

AutoPilot Installer Syntax

The syntax used to run AutoPilot Installer silently from a command line or script is:

```
APInstall [/silent] [parameter setting][parameter setting...]
```

The “silent” switch and parameter settings can occur in any order. One or more spaces must separate the switch and each parameter setting.

The syntax of a parameter setting is:

```
parameter_name =["]value["]
```

Double quotes are required only around values that contain spaces. Spaces may separate parameters, equal signs and values. Parameter names and values are not case-sensitive.

The APInstall command may contain the settings listed below. Each setting, except ConfigFileLocation, may also be specified in the AutoPilot Configuration file. For descriptions of each parameter, see “Software Configuration Parameters” on page 139.

Settings specified in the APInstall command override those specified in the configuration file.

```
ConfigFileLocation = path-specifier
NoSoftwareFirstInstalls = { TRUE | FALSE }
SilentRebootEnable = { TRUE | FALSE }
ForceDriverUpdate = { TRUE | FALSE }
ForceDriverTypeChange = { TRUE | FALSE }
SkipDriverInstall = { TRUE | FALSE }
InstallWithoutQFE = { TRUE | FALSE }
ForceRegUpdate = { TRUE | FALSE }
LocalDriverLocation = path-specifier
ReportLocation = path-specifier
```

Path Specifiers

Paths may be specified as

- an explicit path:

```
ReportLocation="C:\Program Files\Emulex\AutoPilot  
Installer\Reports"
```

- a relative path:

```
LocalDriverLocation="Drivers\Storport Miniport\"
```

(assuming installation into "C:\Program Files\Emulex\AutoPilot Installer\", this path would logically become "C:\Program Files\Emulex\AutoPilot Installer\Drivers\Storport Miniport\")

- with the %ProgramFiles% environment variable:

```
LocalDriverLocation = "%ProgramFiles%\Emulex\AutoPilot  
Installer\Driver"
```

Configuration File Location

The optional setting ConfigFileLocation contains the path to the configuration file that should be used. If this parameter is not specified, AutoPilot Installer uses the file named APInstall.cfg in the same folder as APInstall.exe.

The format is the same as that of the other path settings.

Example:

```
APInstall /silent SkipDriverInstall=True  
configFileLocation=MyConfiguration.cfg
```

Software Configuration Parameters

DiagEnable (Running Diagnostics)

Note: The DiagEnable parameter cannot be specified on the command line; it must be specified within the configuration file

Default: True

By default, AutoPilot Installer runs its diagnostics after all driver installation tasks have been completed. To disable this function, set this parameter to false.

ForceDriverTypeChange (Forcing a Driver Type Change)

Default: False

When installing a driver, set this parameter to true to cause silent mode installations to update or install the Storport Miniport driver on each adapter in the system, without regard for the currently installed driver type (replacing any installation of the SCSIport Miniport or FC Port driver).

ForceDriverUpdate (Forcing a Driver Version Update)

Default: False

By default, if the same version of the driver is already installed, an unattended installation proceeds with installing only the utilities. To force a driver update even if the same version of the driver is installed, set this parameter to true.

Note: ForceDriverUpdate applies to unattended installations only; in interactive installations this parameter is ignored. Instead you are asked if the driver should be updated.

ForceRegUpdate (Forcing an Update of an Existing Driver Parameter Value)

Default: False

The ForceRegUpdate driver parameter setting determines whether existing driver parameters are retained or changed when you update the driver. By default, all existing driver parameter settings are retained. The ForceRegUpdate parameter does not affect any existing persistent bindings. To set up an installation to remove the existing driver parameters from the registry and replace them with parameters specified in the AutoPilot Configuration file, set this parameter to true.

Note: You can use this setting for attended installations with the AutoPilot Installer wizard if you modify the AutoPilot Configuration file in an AutoPilot Installer Kit.

LocalDriverLocation (Specifying Location to Search for Drivers)

Default: Drivers (The default "Drivers" folder is located in the same folder as AutoPilot Installer.)

You can specify a local location that is to be searched for drivers during unattended installations. The location may be a local hard drive or a network share. Removable media are not searched.

Example:

```
LocalDriverLocation = "d:\drivers\new\Storport"
```

Note: On x64 and 32-bit systems, the path specified by 'LocalDriverLocation' must contain at least one instance of an FC, FCoE, iSCSI, and NIC driver. AutoPilot Installer automatically selects the most recent revisions that it finds.

For IA64 systems, only FC and FCoE drivers are supported, no iSCSI or NIC drivers are included for installation.

NoSoftwareFirstInstalls (Prohibiting Software First Installations)

Default: False

When this parameter is set to true, AutoPilot Installer prevents unattended installations from performing software-first installations. This way you can execute an automated installation on multiple machines in your network, but only machines with Emulex adapters actually have Emulex drivers updated or installed.

If this parameter is omitted from the configuration file or explicitly set to true, the page is not displayed. AutoPilot Installer uses configuration file parameters to determine the appropriate management mode.

ReportLocation (Setting Up an Installation Report Title and Location)

The automatically generated file name for this report is

```
"report_mm-dd-yy.txt"
```

where 'mm' is the month number, 'dd' is the day, and 'yy' indicates the year.

You can change only the installation report folder; the file name is auto-generated. In the following example x could be any available drive:

```
ReportLocation = "x:\autopilot\reports\installs\"
```

SilentInstallEnable (Enabling Unattended Installation)

Note: Setting the SilentInstallEnable parameter to true in the configuration file is functionally equivalent to supplying the "/silent" switch on the command line. You cannot specify the SilentInstallEnable parameter on the command line.

Default: False

Setting this parameter to true causes AutoPilot Installer to operate with no user interaction.

SilentRebootEnable (Enabling Silent Reboot)

Default: False

AutoPilot Installer's default behavior in unattended installations is not to restart the system. AutoPilot Installer continues with the installation. Restarts often require you to log in as part of the Windows start up process. If there is no login, the installation process would hang if the system is restarted. However, Windows can be configured to start up without requiring you to log in. You must make sure it is safe to restart the system during unattended installations if you are going to set this parameter to true.

InstallWithoutQFE (Enabling Installation if a QFE Check Fails)

Default: False

AutoPilot Installer checks for Microsoft's Quick Fix Engineering updates (QFEs), also known as KB (Knowledge Base) updates, based on the checks you have specified in the [STORPORT.QFES] section. By default, the installation terminates if the QFE check fails. To enable a driver installation to proceed even if a check for QFEs fails, set this parameter to true.

AutoPilot Configuration File

The AutoPilot configuration file is organized into sections, grouped according to related commands. There are six main sections.

- [AUTOPILOT.ID] - Configuration Identification
- [AUTOPILOT.CONFIG] - Software Configuration
- [STORPORT.CONFIGURATION] - Configuration Prompts/Vendor-Specific Questions
- [STORPORT.QFES] - QFE Checks
- [STORPORT.PARAMS] - Setting Up FC Driver Parameters
- [SYSTEM.PARAMS] - Setting Up System Parameters

Each section begins with a heading. The heading is required even if there are no settings in the section. The only section not required is the Installation Prompts section, which has the heading [STORPORT.CONFIGURATION]. That section cannot exist if AutoPilot Installer runs in silent mode. You must delete or comment-out that entire section for unattended installation.

Lines that begin with a semicolon are comments. Some of the comments are sample settings. To use the setting, remove the semicolon.

Using the Windows Environment Variable (%ProgramFiles%)

You can use the Windows ProgramFiles environment variable in the LocalDriverLocation and ReportLocation strings within the configuration file. This allows you to specify strings in a driver-independent manner, allowing the same configuration file to be used on different systems where Windows may have been installed on different drives. To use this option, "%ProgramFiles%" must be the first component specified in the string. The portion of the string that follows is appended to the contents of the ProgramFiles environment variable. For example:

```
ReportLocation = "%ProgramFiles%\my company\reports"
```

Note: The contents of the ProgramFiles environment variable is not terminated with a slash, so you must provide one in the string. Windows environment variables are not case-sensitive.

Configuration Identification [AUTOPILOT.ID]

This section appears at the beginning of every AutoPilot configuration file and contains revision and label information. The revision entry identifies the file's version number and the date on which it was produced. The label entry is used to identify the configuration that the file supports. This section may appear only once in the APInstall.cfg file.

Software Configuration [AUTOPILOT.CONFIG]

This section can contain settings that control and configure AutoPilot Installer and the OneCommand Manager application operation. This section can appear only once in the AutoPilot configuration file. See “Software Configuration Parameters” on page 139 for information about settings that may be specified in this section.

Configuration Prompts/Vendor-Specific Questions [STORPORT.CONFIGURATION]

Note: You must remove or comment out the entire [STORPORT.CONFIGURATION] section for an unattended installation.

A [STORPORT.CONFIGURATION] section may exist in the AutoPilot configuration file. The first items in this section are the driver parameters to be used regardless of how the questions are answered. This is followed by a subsection that contains questions (these may be vendor-specific questions). A line containing '[QUESTIONS]' marks the start of the subsection, and the end of it is marked by a line containing '[ENDQUESTIONS]'. Within the question subsection there can be as many questions as needed. Each question uses the format:

```
question= "question?", "explanation", "answer0", "answer1",  
"answer2",.... , "answern"
```

Where:

- “question?” contains the text of the question to be asked.
- “explanation” contains brief text to help explain the question. The explanation appears below the question in a smaller font. If there is no explanatory text, empty quotes must be used in its place.
- “answer0” contains the 1st answer to be displayed in the drop down list.
- “answer1” contains the 2nd answer to be displayed in the drop down list.
- “answern” contains the nth answer to be displayed in the drop down list.

For each question there can be as many answers as needed. For each answer there must be a corresponding “answer =” section with its corresponding driver parameters listed beneath it. The answer uses the format:

```

answer = 0
DriverParameter="Param1=value; Param2=value;"
answer = 1
DriverParameter="Param1=value; Param2=value;"
....
answer = n
DriverParameter="Param1=value; Param2=value;"

```

Example of [STORPORT.CONFIGURATION] section:

```

[STORPORT.CONFIGURATION]

;The first section contains the driver parameters common to all
configurations, no matter what answers are given.
DriverParameter="EmulexOption=0;"

[QUESTIONS]
question = "What is your link speed?", "Note: select 'Auto-detect'
if you are unsure about the answer.", "4GB", "2GB", "1GB",
"Auto-detect"
ANSWER = 0
DriverParameter = "LinkSpeed=4;" ;4 GB
ANSWER = 1
DriverParameter = "LinkSpeed=2;" ;2 GB
ANSWER = 2
DriverParameter = "LinkSpeed=1;" ;1 GB
ANSWER = 3
DriverParameter = "LinkSpeed=0;" ;Auto-detect question = "Describe
the topology of your storage network.", "Note: Select 'Arbitrated
Loop' when directly connected to the array (no fibre switch). Select
'Point-to-Point' when connected to a SAN (fibre switch).",
"Arbitrated Loop", "Point-to-Point"
ANSWER = 0
DriverParameter = "Topology=2;"
ANSWER = 1
DriverParameter = "Topology=3;"
[ENDQUESTIONS]
[END.STORPORT.CONFIGURATION]

```

QFE Checks [STORPORT.QFES]

This section specifies an additional Quick-Fix Engineering (QFE) check, also known as KB (Knowledge Base) updates, during installation. To add a Windows QFE check to the configuration file, edit the [STORPORT.QFES] section in the AutoPilot configuration file. You may place this section anywhere within the file as long as it is not contained within another section. This section contains a single line for each QFE that is to be checked. Up to 10 lines are checked, more than that may exist but they are ignored. All parameters in each line must be specified. These lines have the format:

<code>qfe = "qfe name", "path and file name", "file version", "applicable OS"</code>	
<code>qfe name</code>	The name of the item being checked. For example, QFE 838896. The name should facilitate searching Microsoft's website for any required code updates.
<code>path and file name</code>	This string identifies the file to be checked and its location relative to the Windows home folder. In most cases, the file to check is the Microsoft Storport driver, for example, " <code>\system32\drivers\storport.sys</code> ". This string is also used in dialogs and log file messages.
<code>file version</code>	This is the minimum version that the file to be checked must have for the QFE to be considered installed. It is specified as a text string using the same format as is used when displaying the files property sheet. For example, " <code>5.2.1390.176</code> ".
<code>applicable OS</code>	This is used to determine if the QFE applies to the operating system platform present. The acceptable value is " <code>Win2008</code> ".

For example:

```
[STORPORT.QFES]
qfe = "QFE 83896", "\system32\drivers\storport.sys",
"5.2.1390.176", "Win2008"
```

Setting Up FC Driver Parameters [STORPORT.PARAMS]

This section specifies driver parameters. Parameters are read exactly as they are entered and are written to the registry. To change driver parameters, modify this section of the AutoPilot configuration file. Locate the [STORPORT.PARAMS] section in the AutoPilot configuration file. This section follows Optional Configuration File Changes. Under the [STORPORT.PARAMS] heading, list the driver parameters and new values for the driver to use.

For example:

```
Driver Parameter = "LinkTimeout = 45"
```

See "Driver Parameters for Windows" on page 90 for a listing of driver parameters, defaults and valid values.

Setting Up System Parameters [SYSTEM.PARAMS]

To change the system parameters, create a [SYSTEM.PARAMS] section in the APInstall.cfg file. Create this section under the Optional Configuration File Changes heading in the [AUTOPILOT.CONFIG] section.

For example, you can adjust the operating system's global disk timeout. The timeout is stored in the registry under the key HKML\CurrentControlSet\Services\disk and is specified with the following string:

```
TimeOutValue = 0x3C (where the number is the timeout value in seconds.)
```

AutoPilot Installer Exit Codes

AutoPilot Installer sets an exit code to indicate whether an installation was successful or an error occurred. These exit codes allow AutoPilot Installer to be used in scripts with error handling. In unattended installations, AutoPilot Installer sets the following exit codes:

Table D-1 Unattended Installation Error Codes

Error Code	Hex	Description
0	0x00000000	No errors.
2399141889	0x8F000001	Unsupported operating system detected.
2399141890	0x8F000002	The AutoPilot Configuration file is not found.
2399141891	0x8F000003	Disabled adapters detected in the system.
2399141892	0x8F000004	The selected driver is 64-bit and this system is 32-bit.
2399141893	0x8F000005	The selected driver is 32-bit and this system is 64-bit.
2399141894	0x8F000006	Installation activity is pending. AutoPilot Installer cannot run until it is resolved.
2399141895	0x8F000007	(GUI Mode only) You cancelled execution because you did not wish to perform a software-first install.
2399141896	0x8F000008	No drivers found.
2399141897	0x8F000009	One or more adapters failed diagnostics.
2399141904	0x8F000010	(GUI Mode only) You chose to install drivers even though a recommended QFE or Service Pack was not installed.
2399141920	0x8F000020	(GUI Mode only) You chose to stop installation because a recommended QFE or Service Pack was not installed.
2399141899	0x8F00000B	Unattended installation did not find any drivers of the type specified in the config file.
2399141900	0x8F00000C	A silent reboot was attempted, but according to the operating system a reboot is not possible.
2399141901	0x8F00000D	(GUI Mode only) A driver package download was cancelled.
2399141902	0x8F00000E	(Non-Enterprise) No adapters were found in the system.
2399141903	0x8F00000F	A required QFE or Service Pack was not detected on the system.
2399141836	0x8F000030	AutoPilot Installer was not invoked from an account with Administrator-level privileges.
2391419952	0x8F000040	AutoPilot Installer has detected unsupported adapters on the system.
2399141968	0x8F000050	Unattended software-first installations are disallowed.
2399141984	0x8F000060	You cancelled APInstall before any driver/utility installation occurred.
2399142000	0x8F000070	You cancelled APInstall after driver/utility installation occurred.

Table D-1 Unattended Installation Error Codes (Continued)

Error Code	Hex	Description
2399142032	0x8F000090	APIInstaller encountered an error while parsing the command line (Report file contains details).
2399142048	0x8F0000A0	NIC or iSCSI adapters are installed, but are not supported on IA64 systems.

AutoPilot Installer Installation Reports

During each installation, AutoPilot Installer produces a report describing events that occurred during the installation. This report has several sections.

- The first section provides basic information including the time and date of the installation, the name of the machine that the installation was performed on, the version number of AutoPilot Installer, and the identification of the configuration file that was used.
- The second section provides an inventory of the Emulex adapters as they were before AutoPilot Installer performed any actions.
- The third section lists the tasks that AutoPilot performs in the order they are done.
- The fourth section records the results of each task. When all driver installation tasks are completed, an updated adapter inventory is recorded.

Note: If you cancel AutoPilot Installer, that fact is recorded along with when you cancelled the installation. The contents of any error dialogs that are displayed are also recorded.

Command Script Example

Modify the configuration file to script the installation of a system's driver. The following example command script (batch file) assumes that you have made mandatory changes to the AutoPilot configuration file, as well as any desired optional changes. If your systems were set up with a service that supports remote execution, then you can create a command script to remotely update drivers for all of the systems on the storage net. If Microsoft's Remote Command Service (RCMD) service was installed, for example, a script similar to the following would run remote execution:

```
rcmd \\server1 g:\autopilot\x86_full_kit\apinstall
if errorlevel 1 goto serverlok
echo AutoPilot reported an error upgrading Server 1.
if not errorlevel 2147483650 goto unsupported
    echo Configuration file missing.
goto serverlok
:unsupported
if not errorlevel 2147483649 goto older
echo Unsupported operating system detected.
:older
```

```

if not errorlevel 2001 goto none
    echo The driver found is the same or older than the existing driver.
    goto serverlok
:none
if not errorlevel 1248 goto noreport
    echo No Emulex adapter found.
goto serverlok
:noreport
    if not errorlevel 110 goto nocfg
        echo Could not open installation report file.
        goto serverlok
    :nocfg
        if not errorlevel 87 goto badcfg
            echo Invalid configuration file parameters.
            goto serverlok
        :badcfg
            if not errorlevel 2 goto serverlok
            echo No appropriate driver found.
serverlok
rcmd \\server2 g:\autopilot\ApInstall
ConfigFileLocation=g:\autopilot\mysetup\apinstall.cfg
if errorlevel 1 goto server2ok
echo AutoPilot reported an error upgrading Server 2.
if not errorlevel 2147483650 goto unsupported
    echo Configuration file missing.
goto server2ok
:unsupported
    if not errorlevel 2147483649 goto older
        echo Unsupported operating system detected.
    :older2
    if not errorlevel 2001 goto none2
        echo The driver found is the same or older than the existing driver.
        goto server2ok
    :none2
        if not errorlevel 1248 goto noreport2
            echo No adapter found.
goto server2ok
:noreport2
    if not errorlevel 110 goto nocfg2
        echo Could not open installation report file.
        goto server2ok
    :nocfg2
        if not errorlevel 87 goto badcfg2
            echo Invalid configuration file parameters.
            goto server2ok

```

```
:badcfg2
    if not errorlevel 2 goto server2ok
    echo No appropriate driver found.
server2ok
```