



Emulex[®] OneCommand[®] Manager Application

User Guide
Release 12.4

Broadcom, the pulse logo, Connecting everything, Avago Technologies, Avago, the A logo, Brocade, ClearLink, Emulex, ExpressLane, LightPulse, and OneCommand are among the trademarks of Broadcom and/or its affiliates in the United States, certain other countries, and/or the EU.

Copyright © 2003–2019 Broadcom. All Rights Reserved.

The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, please visit www.broadcom.com.

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Table of Contents

Chapter 1: Introduction	7
1.1 Abbreviations	8
Chapter 2: Installing and Uninstalling OneCommand Manager Application Components	9
2.1 Installing the OneCommand Manager Application	9
2.1.1 In Windows	9
2.1.2 In Linux	11
2.1.3 In Solaris	13
2.1.4 In VMware	15
2.2 Uninstalling the OneCommand Manager Application	15
Chapter 3: Starting and Stopping the OneCommand Manager Application	16
3.1 In Windows	16
3.2 In Linux and Solaris	16
Chapter 4: Using the OneCommand Manager Application	18
4.1 The OneCommand Manager Application Window Element Definitions	18
4.1.1 Menu Bar	19
4.1.2 Toolbar	19
4.1.3 Discovery-Tree	21
4.1.4 Property Tabs	22
4.1.5 Status Bar	22
4.2 Using OneCommand Manager Secure Management	23
4.2.1 OneCommand Manager Secure Management Configuration Requirements	24
4.3 Changing Management and Read-Only Mode	24
4.3.1 Management Host	25
4.4 Using CIM (Windows Only)	26
Chapter 5: Configuring Discovery	27
5.1 Discovery Using the TCP/IP Access Protocol	27
5.1.1 Hosts File	28
5.1.2 Adding a Single Host	29
5.1.3 Adding a Range of Hosts (IPv4 Only)	30
5.1.4 Removing Hosts	31
5.2 Configuring Discovery and Default CIM Credentials	32
5.3 Viewing Discovery Information	33
Chapter 6: Managing Hosts	35
6.1 Viewing Host Information	35
6.2 Viewing Host Grouping Information	36

6.3 Grouping Hosts	37
6.3.1 Managing Host Groups	38
6.3.2 Creating a Host Group	39
6.3.3 Deleting a Host Group	40
6.3.4 Adding a Host to a Host Group	40
6.3.5 Removing a Host from a Host Group	41
6.3.6 Restoring a Host Group	41
6.3.7 Restoring All Host Groups	41
6.3.8 Exporting Host Grouping Configurations	41
6.4 Searching for Hosts in the Discovery-Tree	42
Chapter 7: Managing Adapters and Ports	43
7.1 Viewing Adapter Information	43
7.1.1 Configuring Trunking.....	44
7.2 Viewing Port Information	46
7.2.1 Enabling and Disabling a Port.....	47
7.3 Viewing Firmware Parameters	48
7.3.1 Configuring Link Speed.....	49
7.3.2 Enabling and Disabling FA-PWWN.....	50
7.3.3 Disabling and Enabling Dynamic D_Port Tests	51
7.3.4 Enabling and Disabling FEC	51
7.4 Viewing Firmware Information	52
7.5 Viewing and Clearing Statistics	53
7.6 Viewing Virtual Port Information	55
7.7 Creating and Deleting Virtual Ports	56
7.7.1 Creating Virtual Ports.....	56
7.7.2 Deleting Virtual Ports	58
7.8 Viewing Fabric Information	60
7.9 Viewing Port Transceiver Information	61
7.10 Viewing VPD Information	62
7.11 Viewing Maintenance Information	64
7.12 Viewing Target Information	65
7.13 Viewing LUN Information	66
7.14 Viewing Target Mapping	68
7.14.1 Using Automapping and Persistent Binding (Windows Only)	69
7.15 Masking and Unmasking LUNs (Windows)	71
7.15.1 LUN Masking Conventions and Guidelines	72
7.16 Changing the WWPN and WWNN	73
7.16.1 Changing Port Names.....	77
7.16.2 Resetting the FC Functions	77

7.17 Configuring the Driver Parameters	78
7.17.1 Activation Requirements	78
7.17.2 Host Driver Parameters Tab	78
7.17.3 Setting the Driver Parameters	80
7.17.4 Creating a Batch Mode Driver Parameters File	84
7.17.5 Configuring Boot from SAN	87
7.18 Configuring Advanced Settings (Boot from SAN)	91
7.18.1 x86 Boot Advanced Adapter Settings Dialog	92
7.18.2 EFIBoot Advanced Adapter Settings Dialog	93
7.18.3 OpenBoot Advanced Adapter Settings Dialog	94
7.19 Using FC-SP DHCHAP Authentication	95
7.19.1 Deleting Authentication for All Ports	97
7.19.2 Viewing Saved Authentication Configuration Entities	98
7.19.3 Setting or Changing Secrets	99
7.19.4 Changing the Authentication Configuration	100
7.20 Guest Operating System Discovery and Management from the Base Host Operating System	100
Chapter 8: Updating Adapter Firmware	102
8.1 Updating Firmware for a Single Adapter	102
8.2 Updating Firmware for Multiple Adapters	104
Chapter 9: Exporting SAN Information	107
9.1 Creating a SAN Report	107
Chapter 10: Diagnostics	108
10.1 Viewing Flash Contents, Overlay Details, PCI Registers, and Wakeup Information	109
10.1.1 Viewing Flash Contents	110
10.1.2 Viewing Overlay Details	110
10.1.3 Viewing the PCI Registers	110
10.2 Running a Quick Test	111
10.3 Running a POST	111
10.4 Using Beacons	111
10.5 Running D_Port Tests	112
10.6 Using FC Trace Route	114
10.7 Creating Diagnostic Dumps	116
10.7.1 Creating Diagnostic Dumps for LPe12000-Series Adapters	116
10.7.2 Creating Diagnostic Dumps for All Other Adapters	117
10.8 Running Advanced Diagnostic Tests	119
10.8.1 Running Loopback Tests	121
10.8.2 Running End-to-End Tests	122
10.8.3 Saving the Log File	123

Chapter 11: Troubleshooting	125
11.1 General Situations	125
11.2 Emulex Driver for Linux and OneCommand Manager Application Situations	126
11.3 vPorts and OneCommand Manager Application Situations	128
Appendix A: License Notices	129

Chapter 1: Introduction

The Emulex® OneCommand® Manager application is a comprehensive management utility for Emulex adapters that provides a powerful, centralized adapter management suite. Adapter management includes discovery, reporting, and management of local and remote adapters from a single console anywhere in the network and across operating system platforms. Remote configuration capability is provided by Transmission Control Protocol/Internet Protocol (TCP/IP) access from the IP addresses of remote machines. The OneCommand Manager application contains a GUI and a CLI. Refer to the *Emulex OneCommand Manager Command Line Interface User Guide* for information about installing and using the CLI.

NOTE: Screen captures in this user guide are for illustrative purposes only. Your system information can vary.

The OneCommand Manager application can be installed on multiple operating systems, including Windows, Linux, and Solaris. For supported versions of operating systems, platforms, and adapters, go to www.broadcom.com.

For VMware hosts, use the OneCommand Manager application for VMware vCenter. For more details, refer to the *Emulex OneCommand Manager for VMware vCenter User Guide*. You can manage adapters using the OneCommand Manager application on Windows, but you must install and use the appropriate Emulex CIM Provider.

1.1 Abbreviations

AL_PA	Arbitrated Loop Physical Address
BOFM	BladeCenter Open Firmware Manager
CIM	common interface model
CLP	command line protocol
D_ID	destination identifier
DFClib	Documentum Foundation Classes library
DH	Diffie-Hellman
DHCHAP	Diffie-Hellman Challenge Handshake Authentication Protocol
EDD	enhanced disk drive
F_Port	fabric port
FABL	fabric assigned boot LUN
FA-PWWN	fabric assigned Port Word Wide Name
FC-SP	Fibre Channel Security Protocol
FEC	forward error correction
FI_Port	fabric loop port
FLOGI	Fabric Login
IANA	Internet Assigned Numbers Authority
I/O	input/output
IP	Internet Protocol
IPL	initial program load
JEDEC	Joint Electron Device Engineering Council
LIP	Loop Initialization Primitive
N_Port	node port
NL_Port	node loop port
NOS	network operating system
NPIV	N_Port ID Virtualization
OUI	Organizationally Unique Identifier
PAM	pluggable authentication module
PLOGI	Port Login
RMAPI	resource management application programming interface
RPI	remote port indicator
SFS	SAN foundation software
ULP	upper layer protocol
VF	virtual function
XRI	Extensible Resource Indicator

Chapter 2: Installing and Uninstalling OneCommand Manager Application Components

This chapter describes installing and uninstalling the OneCommand Manager application.

2.1 Installing the OneCommand Manager Application

2.1.1 In Windows

The OneCommand Manager application can be installed two ways:

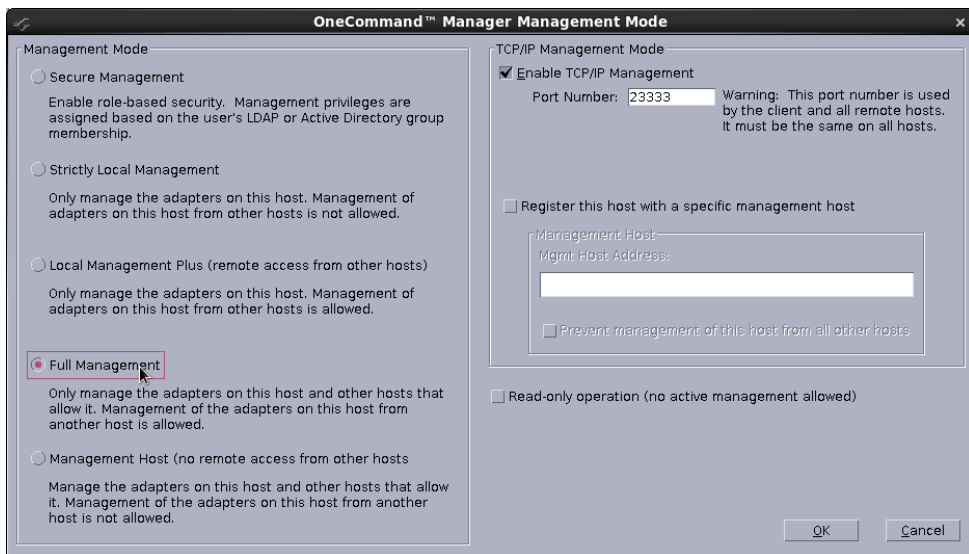
- An attended installation using the GUI.
- An unattended installation using the command line.

2.1.1.1 Attended Installation in Windows

To use the attended installation method in Windows, perform these steps:

1. Download the x64 or x86 OneCommand Manager enterprise kit installation file from the Support Documents and Downloads area of www.broadcom.com.
2. Navigate to the directory to which you downloaded the file.
3. Double-click the `elxocm-windows-<arch>-<version>.exe` file. The **Emulex OCManger Enterprise** window appears. Click **Next**. The **Installation Options** window appears.
4. Select the components that you want to install and click **Install**. After installing the OneCommand Manager application files, the **OneCommand Manager Management Mode** dialog appears ([Figure 1](#)).

Figure 1: Management Mode Dialog



The **OneCommand Manager Management Mode** dialog enables you to select **Secure Management** to assign the desired user privileges, or you can choose one of the other management modes. See [Section 4.2, Using OneCommand Manager Secure Management](#), or [Section 4.3, Changing Management and Read-Only Mode](#), for more information. Choose the management type you want

5. Select or clear the **Enable TCP/IP Management** check box to enable or disable remote management over TCP/IP. You can also change the TCP/IP port used (23333 is the IANA registered port for Broadcom®).
6. Click **OK**. The **Installation Completed** window appears when the installation is finished.
7. Click **Finish**. A shortcut is added to the **Start** menu. You do not need to reboot the system.

2.1.1.2 Unattended Installation in Windows

To use the unattended installation method in Windows, perform these steps:

1. Download the x64 or x86 OneCommand Manager enterprise kit installation file to your system from the Support Documents and Downloads area of www.broadcom.com.
2. Activate the kit with the `/q` or `/q2` switch.
 - The `/q` switch displays progress reports.
 - The `/q2` switch does not display progress reports.
3. You can enable Secure Management mode, or you can select a management mode:

Enable Secure Management mode by adding the `sec=1` argument, or disable it by adding the `sec=0` argument. If the `sec` argument is not entered, Secure Management mode is disabled by default. See [Section 4.2, Using OneCommand Manager Secure Management](#), for more information.

To enable Secure Management mode, type the following command at the command prompt:

```
elxocm-windows-x86-<version>.exe sec=1 /q2
```

To disable Secure Management mode, type the following command at the command prompt:

```
elxocm-windows-x86-<version>.exe sec=0 /q2
```

The management mode defaults for an unattended installation are:

- `mmode = 2` (Local Management Plus mode)
- `achange = 1` (Allow management mode to change)

NOTE: You cannot enable Secure Management mode and select another management mode, or a conflicting parameters error occurs.

Select a management mode by adding the `mmode` argument and the ability to change that management mode by adding the `achange` parameter with selected values as in the following example. See [Section 4.3, Changing Management and Read-Only Mode](#), for more information.

For example, type the following command at the command prompt:

```
elxocm-windows-x86-<version>.exe mmode=3 achange=1 /q2
```

The following are the possible `mmode` values:

- 1 – Local Only Management mode
- 2 – Local Plus Management mode
- 3 – Full Management mode
- 4 – Local Plus Management mode and read only
- 5 – Full Management mode and read only
- 6 – Management host

The following are the possible `achange` values:

- 0 – Do not allow the management mode to change.
- 1 – Allow the management mode to change.

You can also set the following optional parameters:

- `MHost` – This optional switch allows a non-management-host user to select a **Management Host** with which to register. If this switch is not specified, the default value of 0 is used, and the capability is disabled. If the switch is specified, the value can be a host name or an IP address that is validated by the installer. An error message appears if `/mmode` is set as Local Only or Management Host.
- `mExcl` – This optional switch allows the non-management-host user to select whether the OneCommand Manager application processes requests exclusively from the Management Host specified by the `MHost` parameter. This option is accepted only if accompanied by a valid `MHost` value; otherwise, an error message appears. If this switch is not specified, the default value of 0 is used. If the parameter is specified, the valid values are:
 - 0 – Remotely managed by other hosts.
 - 1 – Remotely managed by Management Host *only*.
- `Mtcp` – This optional parameter allows you to enable or disable remote management and to specify the TCP/IP port number over which management occurs. If this parameter is not specified, the default TCP/IP port number 23333 is used.

If the **Management Host** option is selected, you must either select the default port number or enter a valid TCP/IP port number on the command line. A value of 0 is not accepted.

If one of the non-management host options is selected, you can enter the TCP/IP port number on the command line.

2.1.2 In Linux

You must install the appropriate driver version for your operating system before you can install the OneCommand Manager application. Go to the Support Documents and Downloads area at www.broadcom.com for the latest drivers.

Previous versions of the Linux driver must be uninstalled. You must run the uninstall script that shipped with the version of the Linux driver that you want to remove.

2.1.2.1 Attended Installation in Linux

To install the OneCommand Manager application or to update an existing installation, perform these steps:

1. Log on as root.
2. Download the utilities from the Support Documents and Downloads area of www.broadcom.com.
3. Copy the OneCommand `elxocm-<Platform>-<version>.tgz` file to a directory on the installation machine.
4. Change to the directory to which you copied the tar file.

5. Untar the file.

- For RHEL 7 and RHEL 8, type the following command:
`tar zxvf elxocm-rhel7-rhel8-<version>-<rel>.tgz`
- For SLES 12 and SLES 15, type the following command:
`tar zxvf elxocm-sles12-sles15-<version>-<rel>.tgz`

6. Change to the `elxocm` directory created in [Step 3](#).

- For RHEL 7 and RHEL 8, type the following: command:
`cd elxocm-rhel7-rhel8-<version>-<rel>`
- For SLES 12 and SLES 15, type the following command:
`cd elxocm-sles12-sles15-<version>-<rel>`

NOTE: Prior to installation, you must configure OneCommand Manager application groups on the LDAP network or the local host machine for Secure Management operation. See [Section 4.2.1, OneCommand Manager Secure Management Configuration Requirements](#), for configuration instructions.

7. Run the install script. Type the following command:

```
./install.sh
```

8. When prompted, choose whether or not to enable Secure Management for OneCommand.

```
Do you want to enable Secure Management feature for OneCommand? (s/u)
Enter 's' to select secure management. (LDAP/NIS OCM group configuration required)
Enter 'u' to run without secure management (default).
Enter the letter 's' or 'u'.
```

If you enter `s`, proceed to [Step 11](#). You cannot choose a management mode as described in [Step 9](#).

NOTE: If you enable Secure Management in this step, you cannot configure a management mode in [Step 9](#).

9. When prompted, enter the type of management mode that you want to use.

```
Enter the type of management you want to use:
1 Local Mode : HBA's on this Platform can be managed by OneCommand clients on this Platform Only.
2 Managed Mode: HBA's on this Platform can be managed by local or
remote OneCommand clients.
3 Remote Mode : Same as '2' plus OneCommand clients on this Platform can manage local and remote
HBA's.
4 Management Host : Same as '1' plus OneCommand clients on this Platform can manage remote HBA's.
```

NOTE: If you enabled Secure Management in [Step 8](#), you cannot configure a management mode.

- If you select option 2, you are asked if you want to enable TCP/IP management from remote hosts.
- If you select option 3, you are asked if you want to enable TCP/IP management of remote hosts and enable TCP/IP management from remote hosts. You are prompted to enter the TCP/IP port number to use. (Leaving the field blank defaults to 23333.)
- If you select options 2 or 3, you are prompted for the management host address. (Leaving the field blanks means none.)
- You can enter an IP address or host name. If you enter a management host address, you are prompted to exclude management of this host from any other host.
- If you select option 4, management of remote hosts is automatically selected, and you are prompted to enter the TCP/IP port number to use. (Leaving the field blank defaults to 23333.)

NOTE: Management hosts cannot be managed by remote hosts.

10. If you answered 2, 3, or 4 in [Step 9](#), you must decide whether you want the OneCommand Manager application to operate in read-only mode. Read-only mode prevents users from performing some operations, such as resetting adapters, updating an adapter's firmware, or changing adapter driver properties and bindings. It affects only the local OneCommand Manager application interface. These operations can still be performed using remote management. Enter either `y` for yes to allow users to perform these operations or `n` for no if read-only mode is desired.
11. You are prompted whether to allow users to change the management mode after installation. Enter either `y` for yes or `n` for no.

2.1.2.2 Unattended Installation in Linux

For unattended or silent installation of the OneCommand Manager application for Linux, installation settings are defined using the installation script command line.

NOTE: Prior to installation, OneCommand groups must be configured on the LDAP network or the local host machine for Secure Management operation. See [Section 4.2.1, OneCommand Manager Secure Management Configuration Requirements](#), for configuration instructions.

To view the options for unattended installation, type the following command:

```
./install.sh --help
```

To perform an unattended, silent installation, type the following command:

```
./install.sh -q2
```

NOTE: The default management mode for unattended installation is Local Management Plus.

2.1.2.3 Updating an Installation in Linux

The OneCommand Manager application supports the following update paths:

- You can update from an earlier core kit to a later enterprise kit.
- You can update from an earlier enterprise kit to a later enterprise kit.

See [Section 2.1.2.1, Attended Installation in Linux](#), or [Section 2.1.2.2, Unattended Installation in Linux](#), for instructions.

2.1.3 In Solaris

CAUTION! OneCommand Manager installation overwrites some conflicting files that are required by the `fwupdate` tool. The `fwupdate` tool will work after OneCommand Manager is installed, but it might fail when uninstalling OneCommand Manager. To fix the `fwupdate` tool, execute the `pkg fix` command.

NOTE: Emulex HBAs support only the inbox Emulex driver for Solaris.

To install the OneCommand Manager application in Solaris, perform these steps:

1. Copy the Solaris utility kit to a temporary directory on your system.

2. Untar the utility kit by typing the following command:

```
tar xvf elxocm-solaris-<version>.tar
```

3. Change to the newly created `elxocm-solaris-<version>` directory:

```
cd ./elxocm-solaris-<version>/
```

NOTE: Prior to installation, OneCommand groups must be configured on the LDAP network or the local host machine for Secure Management operation. See [Section 4.2.1, OneCommand Manager Secure Management Configuration Requirements](#), for configuration instructions.

4. Run the installation script to begin installation by typing the following command.

```
./install
```

NOTE: If the HBAnyware utility, the OneCommand Manager core kit, the OneCommand Manager enterprise application, or the Solaris driver utilities are already present on the system, the installation script attempts to remove them before proceeding.

5. When prompted, choose whether to enable Secure Management for OneCommand:

```
Do you want to enable Secure Management feature for OneCommand? (s/u)
Enter 's' to select secure management. (LDAP/NIS OCM group configuration required)
Enter 'u' to run without secure management (default).
Enter the letter 's' or 'u'.
```

If you enter `s`, proceed to [Step 7](#). You cannot choose a management mode as described in [Step 6](#).

6. When prompted, enter the type of management that you want to use:

```
Enter the type of management you want to use:
1 Local Mode:HBA's on this Platform can be managed by OneCommand clients on this Platform Only.
2 Managed Mode:HBA's on this Platform can be managed by local or
remote OneCommand clients.
3 Remote Mode:Same as '2' plus OneCommand clients on this Platform can manage local and remote
HBA's.
4 Management Host:Same as '1' plus OneCommand clients on this Platform can manage remote HBA's.
```

NOTE: If you enabled Secure Management in [Step 5](#), you cannot configure a management mode.

- If you select option 2, you are asked if you want to enable TCP/IP management from remote hosts.
- If you select option 3, you are asked if you want to enable TCP/IP management of remote hosts and enable TCP/IP management from remote hosts. You are prompted to enter the TCP/IP port number to use. Leaving the field blank defaults to 23333.
- If you select options 2 or 3, you are prompted for the management host address. Leaving the field blank means none.
- You can enter an IP address or host name. If you enter a management host address, you are prompted to exclude management of this host from any other host.
- If you select option 4, management of remote hosts is automatically selected and you are prompted to enter the TCP/IP port number to use. Leaving the field blank defaults to 23333.

NOTE: Management hosts cannot be managed by remote hosts.

7. If you answered 2, 3, or 4 in [Step 6](#), you must decide whether you want the OneCommand Manager application to operate in read-only mode. Read-only mode prevents users from performing some operations such as resetting adapters, updating an adapter's firmware, or changing adapter driver properties and bindings. It affects only the local OneCommand Manager application interface. These operations can still be performed using remote management. Enter either `y` for yes to allow users to perform these operations, or enter `n` for no if read-only mode is desired.

8. You are prompted whether to allow users to change the management mode after installation. Enter either `y` for yes or `n` for no.

2.1.4 In VMware

For VMware hosts, you can manage adapters using the OneCommand Manager application on Windows, but you must install and use the appropriate Emulex CIM Provider.

The Emulex CIM Provider is available as an offline bundle in ESXi platforms. Use the offline bundle to update software on VMware platforms. For more information about the ESXi Patch Management activities, refer to the VMware website.

For the best real-time management of Emulex adapters in VMware ESXi environments, use the OneCommand Manager application for VMware vCenter. For more information, refer to the *Emulex OneCommand Manager for VMware vCenter User Guide*.

To install the Emulex CIM Provider in a VMware ESXi hypervisor environment, use the `esxcli` command line utility and perform these steps:

1. Copy the CIM Provider zip file to the `/var/log/vmware` directory.
2. Log in to the VMware hypervisor host, and execute the following command:

```
esxcli software vib install -d vmware-esx-provider-emulex-cim-provider-<version>.zip
```
3. Reboot the system.

2.2 Uninstalling the OneCommand Manager Application

To uninstall the OneCommand Manager application, perform these steps:

- In Windows:
 - a. Select **Start > Control Panel > Programs > Uninstall a Program**.
 - b. Select **Emulex OCManger Enterprise <version>**, and click either **Remove** or **Uninstall**.
- In Linux:
 - a. Log on as root.
 - b. Change to the `elxocm-<platform>-<version>` installation directory.
 - c. Type the following:

```
./uninstall
```
- In Solaris:
 - a. Log on as root.
 - b. Run the OneCommand Manager application uninstallation script:

```
/opt/ELXocm/scripts/uninstall
```
- In VMware
 - a. Type the following:

```
esxcli software vib remove -n emulex-cim-provider
```

Chapter 3: Starting and Stopping the OneCommand Manager Application

This chapter describes how to start and stop the OneCommand Manager application.

NOTE: For VMware systems, if you are running only a CIM client, you do not need to stop it.

3.1 In Windows

To start the OneCommand Manager application, from the Windows desktop, select **Start > All Programs > Emulex > OCManager**. If Secure Management is enabled, you are prompted for your user name and password. See [Section 4.2, Using OneCommand Manager Secure Management](#), for more information.

To stop the OneCommand Manager application, from the OneCommand Manager application menu bar, select **File > Exit**.

3.2 In Linux and Solaris

Linux and Solaris installations of the OneCommand Manager application include two basic daemon processes that are affected by the start and stop scripts:

- `elxhbamgrd` – The remote management daemon, which services requests from OneCommand Manager application clients running on remote host machines.
- `elxdiscoveryd` – The discovery daemon, which maintains all discovery data (remote and local) for OneCommand Manager application clients running on the local machine.

The `elxhbamgrd` daemon starts at system boot time. The `elxdiscoveryd` daemon starts whenever the OneCommand Manager application GUI process first runs on the host machine.

To start the OneCommand Manager application, use the `ocmanager` script. The script is located in the following OneCommand Manager installation directory:

- **Linux:** `/usr/sbin/ocmanager`
Example usage on Linux:

```
# /usr/sbin/ocmanager/ocmanager
```
- **Solaris:** `/opt/ELXocm`
Example usage on Solaris:

```
# /opt/ELXocm/start_ocmanager
```

If Secure Management is enabled, you are prompted for your user name and password. See [Section 4.2, Using OneCommand Manager Secure Management](#), for more information.

To stop the OneCommand Manager application, use one of the following methods:

- From the menu bar, select **File > Exit**.
- From the shell, use the `stop_ocmanager` script located in the OneCommand Manager installation directory.

Example usage on Linux:

```
# /usr/sbin/ocmanager/stop_ocmanager -n
```

Example usage on Solaris:

```
# /opt/ELXocm/stop_ocmanager -n
```


NOTE: The `-n` parameter stops only the OneCommand Manager application and the associated `elxdiscoveryd` daemon. When the `stop_ocmanager` script is run without the `-n` parameter, the script stops the OneCommand Manager application and all associated daemons.

Chapter 4: Using the OneCommand Manager Application

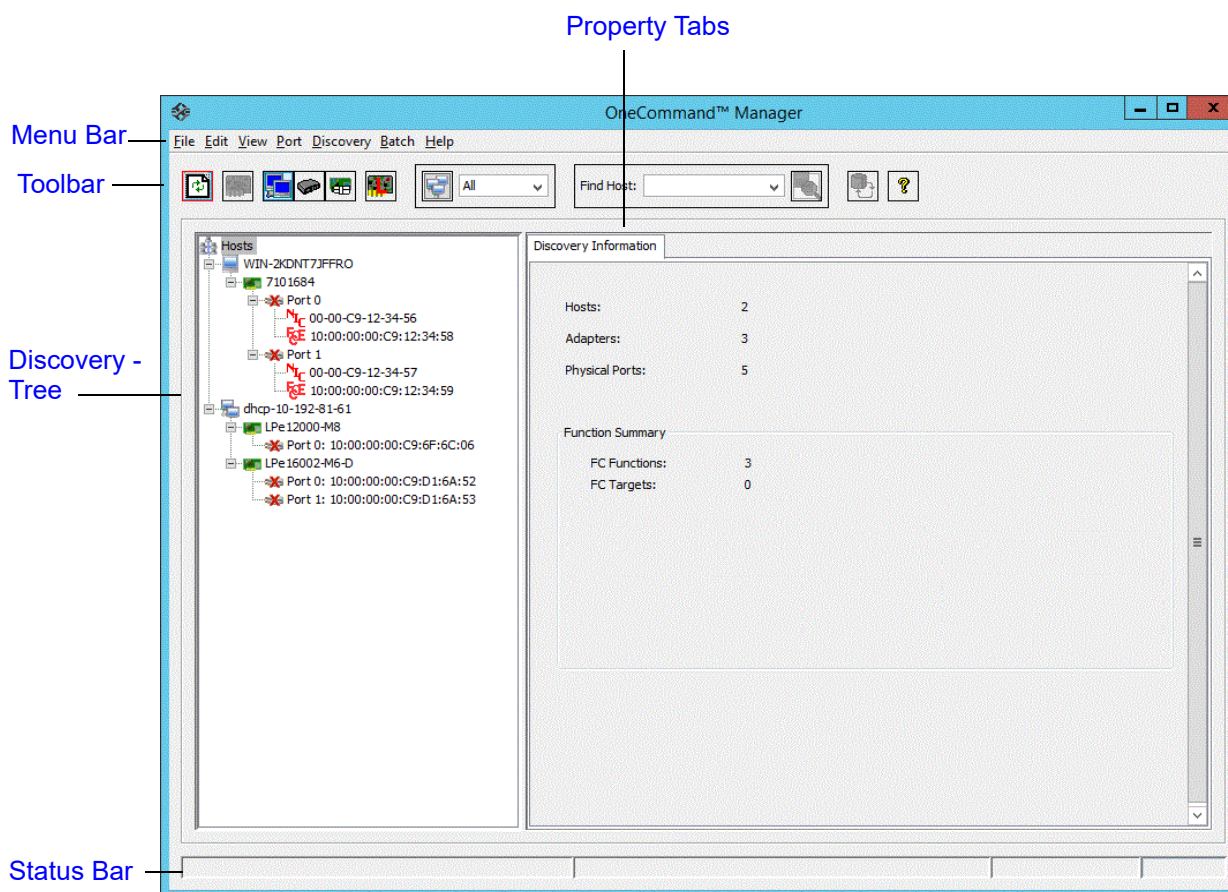
To properly view the OneCommand Manager application, make sure that your system meets the following display requirements:

- For Windows, Linux, and Solaris systems, the display resolution must be set to 1024 × 768 or higher. For Windows systems, use the default font size.
- The display must run in 256-color mode or higher. OneCommand Manager application icons use 256 colors. If the display is set for 16 color mode, OneCommand Manager application icons are not displayed.

4.1 The OneCommand Manager Application Window Element Definitions

The **OneCommand Manager** application window (Figure 2) contains five basic components: the menu bar, the toolbar, the discovery-tree, the property tabs, and the status bar.

Figure 2: OneCommand Manager Application Window



NOTE: The element that you select in the discovery-tree determines whether a menu item or toolbar button is active. For example, if you select the local host or other system host, the **Reset Port** item on the **Adapter** menu is unavailable. The **Reset Port** toolbar button is unavailable as well.

The capabilities displayed by your local interface match those of the remote server. When accessing a remote server running an older version of the OneCommand Manager application, capabilities that are not supported by the server's older version of the OneCommand Manager application are unavailable.

In some instances, the type of information displayed and available functionality is determined by the operating system in use.

4.1.1 Menu Bar

The menu bar contains commands that enable you to perform a variety of tasks, such as exiting the OneCommand Manager application, resetting adapters, and sorting items in the discovery-tree view. Many of the menu bar commands are also available from the toolbar.

4.1.2 Toolbar

The toolbar (Figure 3) contains buttons that enable you to refresh the discovery-tree, reset the selected adapter, and choose how you want to view discovered SAN elements in the discovery-tree. Many of the toolbar functions are also available from the menu bar.

Figure 3: Toolbar



The toolbar is visible by default. Use the **Toolbar** item in the **View** menu to hide the toolbar. If the item is checked, the toolbar is visible.

4.1.2.1 Toolbar Buttons

The toolbar buttons perform the following tasks.



Discovery Refresh button

- Initiates a discovery refresh cycle.



Reset Port button

- Resets the selected port.

NOTE: The three view buttons on the toolbar enable you to view SAN elements from the host, fabric, virtual ports, or by local or remote adapter perspective. By default, both local and remote adapters are displayed in the Host view. The OneCommand Manager application displays elements in ascending order.



Host View button (default)

- Displays the host system.

NOTE: You cannot change host names using the OneCommand Manager application; names must be changed locally on that system.

- Displays the installed adapters within each host system.
- Displays adapter ports and the port numbers; if available.
- Displays adapters by the WWNN if multiple adapters have the same model number.
- Displays the WWPN if targets are present. Multiple adapters can refer to the same target.
- Displays the LUN number if LUNs are present.



Fabric View button

- Displays the fabrics in the SAN with their fabric IDs.
- Displays the ports under each switch.
- If targets are present, displays each WWPN. Multiple adapters can refer to the same target.
- If LUNs are present, displays each LUN number.
- If the fabric ID is all zeros, no fabric is attached.



Virtual Ports View button

- Displays virtual ports in the SAN.



Local HBAs Only button

- Displays only local adapters.



Show Host Groups button and menu

- Displays hosts by their associated groups.
- Displays available host groups.



Find Host button and search field

- Enables you to search by host name for a particular host in the discovery-tree.



Refresh LUNS button

- Initiates a LUN discovery refresh cycle.



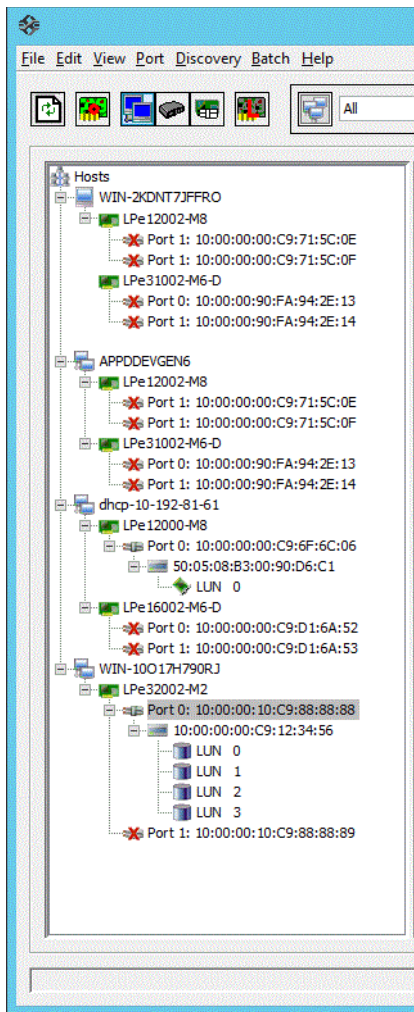
Help button

- Displays the OneCommand Manager application's online help.
- Displays the **About OneCommand Manager** dialog. The dialog displays version information including RMAPI, Discovery, DFClib, and Remote Management Agent Version (Windows). It also enables you to contact Broadcom Technical Support.

4.1.3 Discovery-Tree




The discovery-tree (Figure 4) has icons that represent discovered hosts, adapters, ports, virtual ports, fabrics, targets, and LUNs.

Figure 4: Discovery-Tree



4.1.3.1 Discovery-Tree Icons

Discovery-tree icons represent the following:

-  The local host.
-  Other hosts connected to the system.
-  A green **Adapter** icon with black descriptive text represents an online adapter. Blue text represents an adapter port that was previously discovered, but is not currently seen by the discovery engine (service). The adapter is removed from the discovery-tree if it still is not seen after the undiscovered adapter expiration time has elapsed (default is 1800 seconds, or 30 minutes). If the adapter is discovered again before the expiration time has elapsed, it reverts back to normal black text. See [Section 5.2, Configuring Discovery and Default CIM Credentials](#), for more information about discovery settings.



The **Port** icon represents an adapter port. A **Port** icon with a red X indicates the link is down.

NOTE: Multiport adapters are represented in the discovery-tree with separate port icons for each port, with the port number displayed next to the icon.



The **ASIC Node** icon, only displayed for dual ASIC adapters, represents each ASIC on the adapter. Each ASIC is managed independently. The ASIC node format **ASIC bus#-sub-adapter#** represents the PCI bus number and the sub-adapter number, which is a concatenation of the discovered port numbers for the ASIC. For example, in **ASIC 64-12**, **64** represents PCI bus number 64, and **12** represents ports 1 and 2. If there were no discovered functions for a port on that ASIC, the label would be **ASIC 64-2** (port 1 is missing).



The **Virtual Port** icon represents a virtual port.



The **Target** icon represents connections to individual storage devices.



The **LUN** icon represents connections to individual disk LUNs.



The **Masked LUN** icon represents a LUN not presented to the host.



The **Media Exchanger** icon represents connections to individual media exchangers. A media exchanger is a jukebox-type device that is capable of swapping various media device instances (such as records or CDs) in and out.



The **Tape LUN** icon represents LUNs that are tape devices.



The **Target Controller LUN** icon represents LUNs that are storage controllers.



The **Switch** icon represents connections to the switch.

4.1.3.2 Expanding or Collapsing the Discovery-Tree View

You can use the expand/collapse capability on the **View** menu to change the way discovered elements are displayed. By selecting one of the five levels, the discovery-tree ([Figure 4](#)) is expanded or collapsed to that level. You can choose hosts/fabrics (depending on the view), adapters, ports, PCI functions, and targets.

4.1.4 Property Tabs

The property tabs display configuration, statistical, and status information for network elements ([Figure 2](#)). The set of available tabs is context-sensitive, depending on the type of network element or adapter port currently selected in the discovery-tree ([Figure 4](#)).

4.1.5 Status Bar

The status bar is located near the bottom of the **OneCommand Manager application** window ([Figure 2](#)). The status bar displays messages about OneCommand Manager application functions, such as *Discovery in progress* or the progress when performing an Export SAN Info operation.

The status bar is visible by default. Use the **Status Bar** item in the **View** menu to hide the status bar. If checked, the status bar is visible.

4.2 Using OneCommand Manager Secure Management

OneCommand Manager Secure Management gives system administrators the ability to further enhance the active management security of their networks. Using Secure Management, administrators can define each user's privileges for managing both local and remote adapters. When running in Secure Management mode, users must log on with their user name and password to run the OneCommand Manager application. If users are authenticated, they can only perform the functions allowed by the OneCommand Manager user group to which they belong. If the systems are running in an LDAP or Active Directory domain, the OneCommand Manager application authenticates users with those defined for LDAP or Active Directory domains. For Linux and Solaris systems, this is accomplished using PAM.

NOTE: OneCommand Manager Secure Management is not supported on VMware hosts.

Administrators set up user accounts such that users belong to one of the OneCommand Manager application user groups. The user groups define the management capabilities for each user.

Table 1 defines the OneCommand Manager application user groups and each group's management capabilities.

Table 1: Secure Management User Privileges

Group Name	OneCommand Manager Capability
ocmadmin	Allows full active management of local and remote adapters.
ocmlocaladmin	Permits full active management of local adapters only.
ocmuser	Permits read-only access of local and remote adapters.
ocmlocaluser	Permits read-only access of local adapters.

On Linux or Solaris systems, the UNIX `getent group` utility can be run on the target host system's command shell to verify the correct configuration of the groups. The groups, and users within the groups, appear in the output of this command.

NOTE: Although users might belong to the administrator group or be root users, they do not have full privileges to run the OneCommand Manager application unless they are also members of the ocmadmin group. Otherwise, if Secure Management is enabled, root users or administrators can only manage local adapters (similar to the ocmlocaladmin users).

Remote management operations between two machines are allowed or denied depending on the OneCommand Manager Secure Management status of the machines, and on the domains to which the machines belong. The following tables list the behavior (assuming appropriate user credentials are used).

Table 2: Active Commands: Machines on Same Domain

Client	Remote Server (Secure)	Remote Server (Not Secure)
Secure	Allowed	Denied ^a
Not Secure	Denied	Allowed

a. To inform you of an unsecured server that you might want to secure.

Table 3: Active Commands: Machines on Different Domain

Client	Remote Server (Secure)	Remote Server (Not Secure)
Secure	Denied ^a	Denied ^b
Not Secure	Denied	Allowed

- a. Allowed if the username and password are the same on both domains.
- b. To inform you of an unsecured server that you might want to secure.

Table 4: Passive Commands: Machines on Any Domain

Client	Remote Server (Secure)	Remote Server (Not Secure)
Secure	Allowed	Allowed
Not Secure	Allowed	Allowed

4.2.1 OneCommand Manager Secure Management Configuration Requirements

For systems to run OneCommand Manager Secure Management, they must be configured to provide the following two capabilities:

- Authentication – On Linux and Solaris systems, this is accomplished using the PAM interface and must be configured as follows:
 - For Solaris systems, place the correct setting in the `auth` section of the `/etc/pam.d/other` file or its earlier equivalent, `/etc/pam.conf`.
 - For Linux systems, place the correct setting in the `auth` section of the `/etc/pam.d/password` file equivalent.
- User Group Membership – From the host machine, OneCommand Manager Secure Management must be able to access the OneCommand Manager group to which the user belongs. For Linux and Solaris systems, it uses the `getgrnam` and `getgrid` C-library API calls. The equivalent to the API calls can be obtained by typing `getent group` from the shell command line. If the four OneCommand Manager group names are listed with their member users, the machine is ready to use OneCommand Manager secure management.

For Solaris systems, you must use `useradd -G <groupname>` for authentication to work. You cannot use a lowercase `g`.

4.3 Changing Management and Read-Only Mode

NOTE: This functionality is available only to root users and administrators even when running in Secure Management mode.

During installation, a management and a read-only mode are selected. If you chose a **Secure Management** or **Full Management** option, you can change the management mode after installation.

The following options are available:

- Secure Management – The setting enables roles-based security. See [Section 4.2, Using OneCommand Manager Secure Management](#), for details.
- Strictly Local Management – This setting allows management of adapters on this host. Management of adapters on this host from other hosts is not allowed.
- Local Management Plus – This setting allows management of adapters on this host only, but management of adapters on this host from another host is possible.
- Full Management – This setting allows you to manage adapters on this host and other hosts that allow it.
- Management Host – This setting allows this host to manage other hosts, but prevents this host from being managed by other hosts.
- Enable TCP/IP Management (of or from the remote host) – This setting enables you to manage remote hosts or to manage this host remotely. If this setting is enabled, you must supply the port number (between 1024 and 65535). The default port number is 23333. If the port number or the **Enable TCP/IP Management** check box is changed, a set of warning messages might appear before the change is made. Click **Yes** to continue with the change.

If the IP port number is changed, the utility restarts the OneCommand Manager application discovery server and management agent to use the new settings. If the servers cannot be stopped and restarted, you are prompted to reboot the host for the new TCP/IP management settings to take effect.

CAUTION! The IP port number must be the same for all hosts that are to be managed. Setting an IP port number for one host to a different value than the other hosts makes the host unable to manage other hosts over TCP/IP using a different port. It also makes the host unmanageable over TCP/IP from other hosts using a different port.

- Register this host with specific management host – This setting enables you to register this host with a specific host for management. If this setting is enabled, you must supply the IP address or host name of the management host. You can also choose to prevent management of this host from any other host but the management host. See [Section 4.3.1, Management Host](#), for more information.

If Local Management Plus or Full Management mode is selected, you can also set read-only mode.

- Read-only operation – This setting prevents some operations from being performed, such as resetting adapters, updating the adapter firmware image, and changing adapter settings and driver properties. Dialog controls that pertain to these tasks are completely hidden or disabled.

4.3.1 Management Host

The OneCommand Manager management host provides enhanced discovery and security by enabling a managed host to register with a management host. The management host receives these registrations when the remote host is started and updates its hosts file so the discovery server discovers the remotely managed host. You do not need to manually add remote hosts to be managed.

If you choose to exclude management from all hosts except the management host, the managed host responds to requests from the management host only. All requests from other hosts are rejected. This TCP/IP management security solution allows only the management host to manage the remote host.

To change the management mode and read-only type, perform these steps:

NOTE: After making changes, you must restart the OneCommand Manager application to see the new management mode settings.

- In Windows:
 - a. From the **File** menu, select **Management Mode**. The **Management Mode** dialog appears ([Figure 1](#)).
 - b. Choose the management type and read-only mode that you want.
 - c. Click **OK**.

- In Solaris:
 - a. Run the following script:

```
/opt/ELXocm/set_operating_mode
```
 - b. Choose the management type and read-only mode that you want.

- In Linux:
 - a. Stop the OneCommand Manager application.
 - b. Run the following script:

```
/usr/sbin/ocmanager/set_operating_mode
```
 - c. Choose the management type and read-only mode that you want.

4.4 Using CIM (Windows Only)

VMware uses CIM as the only standard mechanism for device management. The OneCommand Manager application uses the standard CIM interfaces to manage the adapters in the Visor environment and supports CIM-based devices and HBA management.

To manage the adapters on a VMware host using the OneCommand Manager application, you must install the Emulex CIM Provider on the VMware host. Refer to the *Emulex CIM Provider Installation Guide* for additional information.

For more information about the VMware Patch Management activities, refer to the VMware website.

NOTE: For VMware hosts, if advanced adapter management capabilities are required (for example, port disable), use the OneCommand Manager application for VMware vCenter. For more details, refer to the *Emulex OneCommand Manager for VMware vCenter User Guide*.

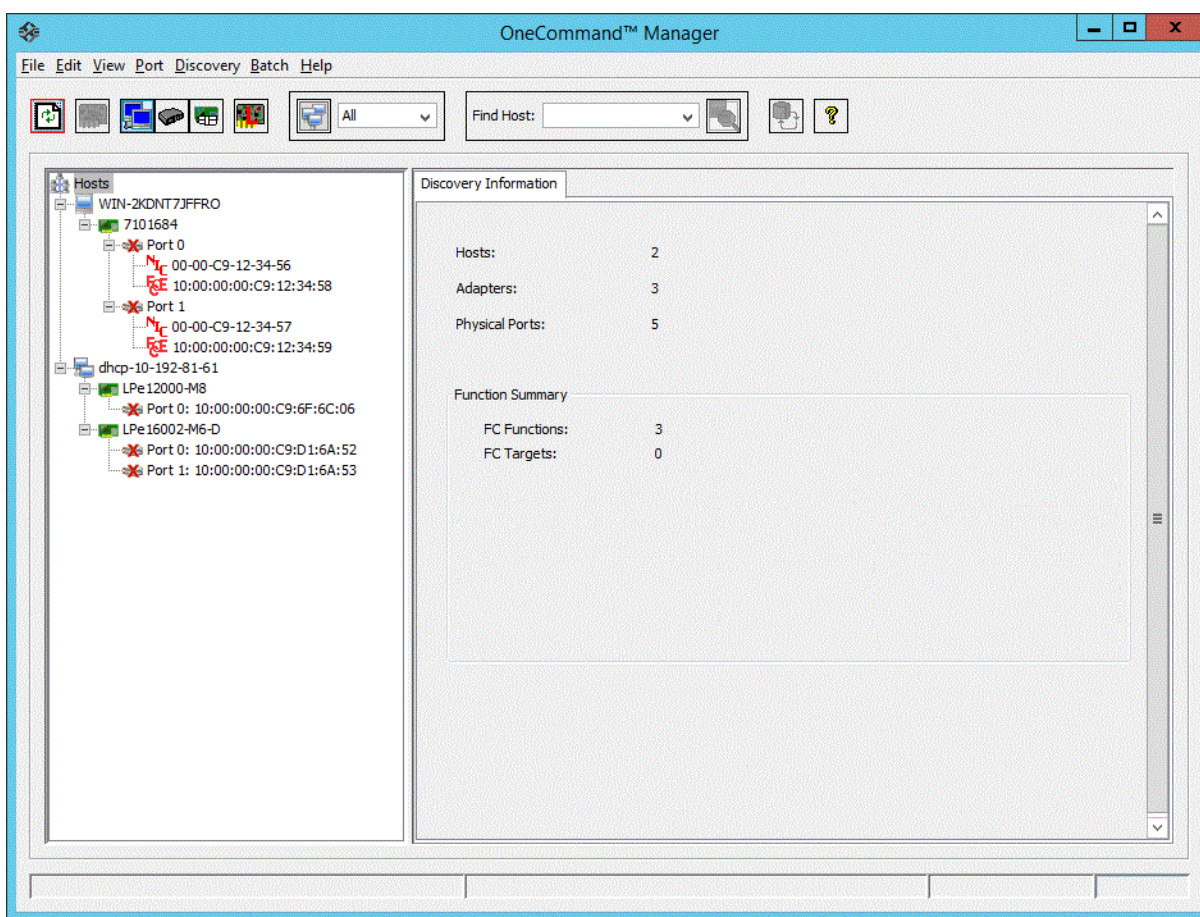
Chapter 5: Configuring Discovery

This chapter describes how to configure discovery to find Emulex adapters on remote hosts.

5.1 Discovery Using the TCP/IP Access Protocol

You can discover adapters on IPv4 and IPv6 TCP/IP hosts and on hosts configured to support the CIM interface that have the OneCommand Manager application installed (Figure 5). Remote SAN management over TCP/IP sends remote management requests using the TCP/IP access protocol to remote hosts. TCP/IP access enables you to access adapters by using their host IP address or the name of the host on which they reside.

Figure 5: Discovery Information



NOTE: In Windows, if you are running a firewall, you might need to add the OneCommand Manager application remote server to the firewall's exception list. This remote server's path is:

```
\Program Files\Emulex\Util\Common\rmsserver.exe
```

5.1.1 Hosts File

The TCP/IP discovery function of the OneCommand Manager application discovery server relies on a file called the `hosts` file. This plain-text file contains a list of hosts that the utility attempts to discover. The discovery server does not attempt to discover hosts over TCP/IP through any other mechanisms (such as ping sweeps and broadcasts).

The `hosts` file is automatically created or modified when you perform any of the following operations:

- Adding a single host from the **Add Remote Host** window (see [Section 5.1.2, Adding a Single Host](#)). If the host is discovered, the OneCommand Manager application adds its IP address and name to the `hosts` file.
- For IPv4, scanning a range of IP addresses for hosts that can be managed. This function is performed in the **Add Remote Hosts** window ([Section 5.1.3, Adding a Range of Hosts \(IPv4 Only\)](#)). For each discovered host, the OneCommand Manager application adds its IP address and name to the `hosts` file.
- Removing a host from the host file using the **Remove Remote Hosts** window (see [Section 5.1.4, Removing Hosts](#)). For each removed host, the OneCommand Manager application removes its IP address and name from the `hosts` file.
- Adding or removing a host using the CLI (refer to the *Emulex OneCommand Manager Command Line Interface User Guide*).

5.1.1.1 Manually Editing the `hosts` File

You can open the `hosts` file with any text editor, modify the contents, and save the file. The name of the `hosts` file is `hbahosts.lst`. After the file is modified and saved, the updated file is used after the next TCP/IP discovery cycle is complete. If the discovery server is running, it does not need to be restarted.

To manually edit the `hosts` file, perform these steps:

1. Locate and open the `hosts` file.
 - Windows – The file is located on the system drive in the directory `\Program Files\Emulex\Util`.
 - Solaris – The file is located in the directory `/opt/ELXocm`.
 - Linux – The file is located in the directory `/usr/sbin/ocmanager`.
2. Edit the file. Guidelines for editing the file are as follows:
 - Each line of the file starts with an IPv4 or IPv6 address. Following the IP address can be any number of tabs or spaces. These are followed by a `#` character, zero, or more tabs or spaces, and the name of the host for that IP address. The host name is not required for discovery. Its purpose is to make the file more readable, and it is used by the OneCommand Manager application to display the host name in the **Remove Remote Hosts** window when the host is not discovered. However, the discovery server only needs the IP address to discover the host.
 - IPv6 address tuples are delimited by colons and can be added in shortened notation as defined by the IPv6 address specification.
 - An IP port number can be specified after the IPv4 address by appending a colon and port number to the address (such as `10.192.80.24:23333`).
 - An IP port number can be specified after an IPv6 address by putting the IPv6 address in brackets and following it with a colon and the port number. For example, `[fe80::50f1:832:3ce4:8d30]:23333`
 - Each line in the file can be up to 1023 characters, although this is longer than is typically needed for a host IP address and host name. A line longer than 1023 characters is truncated, possibly causing discovery to not discover some of the hosts.
 - Blank lines are ignored.
3. Save the file.

5.1.1.2 Copying the File

A `hosts` file on one host can be copied and used on another host. This is useful when there are multiple hosts on the same network running the OneCommand Manager application. For example, after the remote hosts are added to the `hosts` file on one host, you can copy it to other hosts so you do not need to create another `hosts` file.

NOTE: Because of the line terminator differences between Windows, Solaris, and Linux hosts, `hosts` files cannot be shared between Windows, Solaris, or Linux hosts.

5.1.2 Adding a Single Host

NOTE: This option is not available in read-only mode.


The OneCommand Manager application enables you to specify a single TCP/IP host to manage. You can add an RMAPI host or a CIM host using the host name or IP address. If the host is successfully discovered, it is added to the `hosts` file. If it has not been discovered over FC already, the host and its adapter ports are added to the discovery-tree (Figure 4).

NOTE: The OneCommand Manager application must be installed on the remote host.

To add a single host, perform these steps:

1. From the **Discovery** menu, select **TCP/IP > Add Host**. The **Add Remote TCP/IP Host** dialog appears (Figure 6).

Figure 6: Add Remote TCP/IP Host Dialog



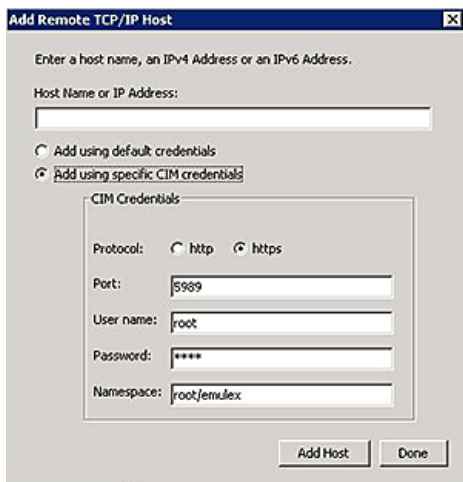
2. Enter the name or the IPv4 or IPv6 address of the host to be added.

NOTE: Entering the IP address to identify the host avoids possible name resolution issues. IPv6 address tuples are delimited by colons and can be entered in a shortened form suppressing 0s as defined by the IPv6 address specification.

3. Configure the discovery method:
 - If you want to add the host using the default discovery methods, check **Add using default credentials** and click **Add Host**. A message appears indicating whether the new host was successfully added.
 - If you want to add the new host using specific CIM credentials, check **Add using specific CIM credentials**, modify any additional CIM settings, and click **Add Host**. The **Add Remote TCP/IP Host** dialog appears with the default CIM settings (Figure 7).

NOTE: Remote CIM hosts can be managed only by Windows client systems.

Figure 7: Add Remote TCP/IP Host Dialog with CIM Credentials



4. Edit the default CIM settings if necessary and click **Add Host**. A message appears indicating that the new host was successfully added.

5.1.3 Adding a Range of Hosts (IPv4 Only)

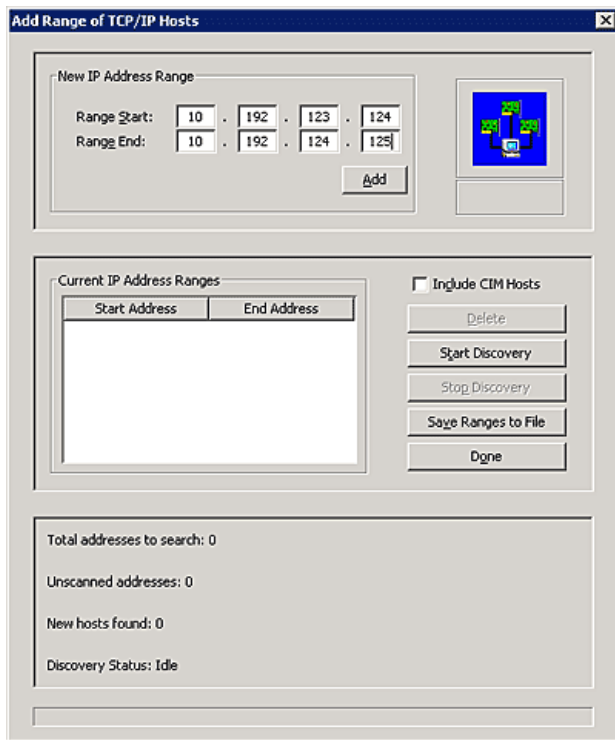
NOTE: This option is not available in Strictly Local or Local Plus Management modes.

You can find the TCP/IP-accessed manageable hosts by searching a range of IPv4 addresses. The **Add Range of TCP/IP Hosts** dialog (Figure 8) enables you to build the initial list of TCP/IP-accessed manageable hosts.

NOTE:

- The ranges of IP addresses are scanned only each time you open the **Add Remote TCP/IP Hosts** dialog and click **Start Discovery**. The ranges are not automatically scanned by the discovery server during its discovery cycles.
- Discovery of VMware (CIM) hosts is supported only on Windows systems. Adding a range of hosts is supported only for IPv4 addresses. It is not supported for IPv6 addresses.
- The OneCommand Manager application must be installed on all remote hosts.

Figure 8: Add Range of TCP/IP Hosts Dialog



To add a range of remote hosts, perform these steps:

1. From the **Discovery** menu, select **TCP/IP > Add Range of Hosts**. The **Add Range of TCP/IP Hosts** dialog appears (Figure 8).
2. Enter the complete start and end address range (IPv4 only) and click **Add**. The added address range appears in the dialog. Add any additional ranges that you want to search.
3. Click **Start Discovery**. If an address is remotely manageable, it is added to the list of addresses that the discovery server attempts to discover. The utility creates a `hosts` file if necessary, and checks each address in the range to determine if the host is available and remotely manageable. The number of addresses (of manageable hosts) discovered is periodically updated on the dialog.

NOTE: The number of hosts found does not correspond directly to the number of hosts added to the discovery-tree (Figure 4). A host can have more than one IP address assigned to it. If multiple IP addresses for a host are discovered during the search, the host is added to the discovery-tree only once.

4. You can save the IP address ranges. Click **Save Ranges to File** to save the specified ranges to a file so that these address ranges appear the next time you use the **Add Range of TCP/IP Hosts** dialog (Figure 8).

5.1.4 Removing Hosts

NOTE: This option is not available in read-only mode.

Removing hosts that are no longer discovered improves the operation of the discovery server. For example, you might want to remove a host when it is removed from the network.

To remove hosts, perform these steps:

1. From the **Discovery** menu, select **TCP/IP > Remove Host(s)**. The **Remove Hosts** dialog shows a list of discovered hosts. Any host that is not currently discovered appears in red. Click **Show Undiscovered Hosts Only** to display only currently undiscovered hosts.
2. From the **Remove Hosts** dialog, select the hosts that you want to remove. You can select all the displayed hosts by clicking **Select All**.
3. Click **Remove** to remove the selected hosts.

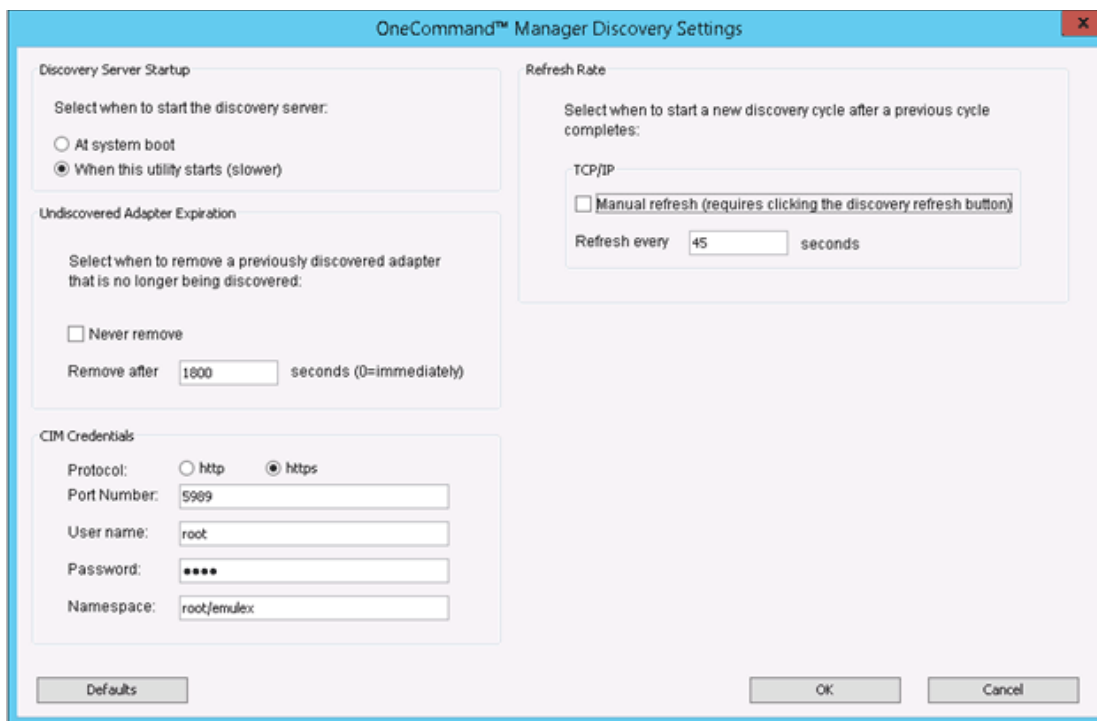
5.2 Configuring Discovery and Default CIM Credentials

Use the **Discovery Settings** dialog (Figure 9) in the OneCommand Manager application to configure several discovery server parameters. You can define when to remove previously discovered adapters that are no longer being discovered. You can also define default CIM credentials, such as the protocol, user name, port number, password, and name space.

NOTE: Management of CIM hosts is supported only on Windows systems.

A host can have more than one IP address assigned to it. If multiple IP addresses for a host are discovered during the search, the host is added to the discovery-tree (Figure 4) only once. If the same host name appears for more than one host, the adapters of all these hosts are displayed by the OneCommand Manager application as a single host entry.

Figure 9: Discovery Settings Dialog



To configure discovery settings, perform these steps:

1. From the **Discovery** menu, select **Modify Settings**. The **OneCommand Manager Discovery Settings** dialog appears (Figure 9).

2. Define the discovery properties that you want.
3. Set the default CIM credentials in the CIM credentials area that are used to connect to all the ESXi hosts that are managed through the CIM interface.
 - **Protocol** – The HTTP or HTTPS protocol can be used to connect to the VMware hosts.
 - **Port Number** – The default port numbers used for HTTP and HTTPS are 5988 and 5989, respectively. The port number changes automatically according to the protocol selected. You can also manually change the port number. By default, the HTTP protocol is disabled on `sfcb` in VMware host, so you must use HTTPS to communicate to the VMware host.
 - **User name** – The **User name** field contains the user name with which to connect to the VMware hosts. By default, this is `root`.
 - **Password** – The **Password** field contains the password of the user name that is used to connect to the VMware host.
 - **Namespace** – Namespace is the namespace of the Emulex provider.
The default namespace is `root/emulex`.
4. Choose the refresh rate settings that you want to apply.
5. Click **OK** to apply your changes. Click **Defaults** to return the discovery properties to their default settings.

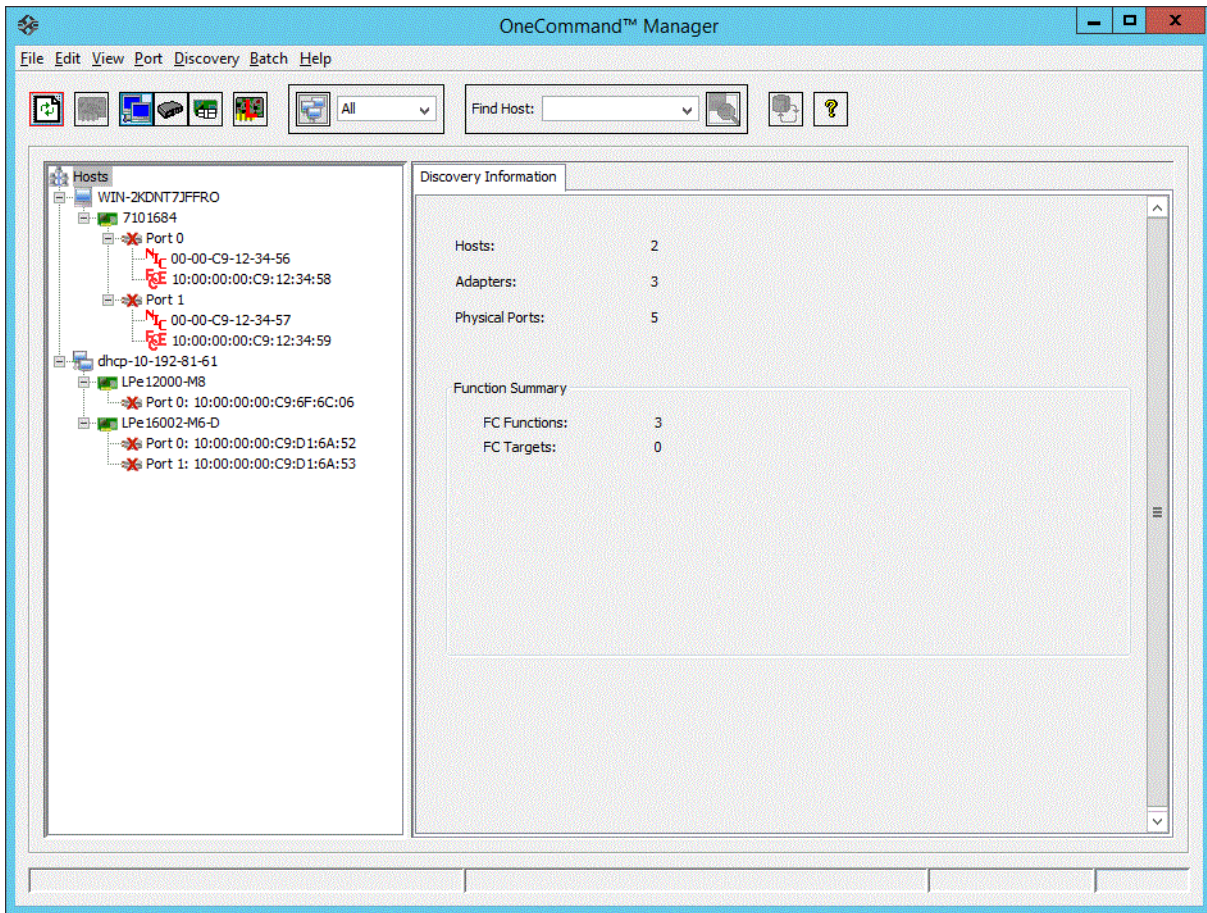
5.3 Viewing Discovery Information

The Discovery Information window (Figure 10) contains a general summary of the discovered elements. The **Host**, **Fabric**, or **Virtual Port** icon, depending upon which view you select, is the root of the discovery-tree, but it does not represent a specific network element. Expanding it reveals all hosts, LUNs, targets, adapter ports, and virtual ports that are visible on the SAN.

To view discovery information, perform these steps:

1. Click the **Hosts**, **Fabrics**, or **Virtual Port** icon at the root of the discovery-tree. Discovered SAN elements appear in the discovery-tree.
2. Select an element from the discovery-tree to learn more about it.

Figure 10: Discovery Information (Host View Selected)



The following **Discovery Information** fields are displayed:

- **Hosts** – The total number of discovered host computers containing manageable Emulex adapters. This number includes servers, workstations, personal computers, multiprocessor systems, and clustered computer complexes.
- **Adapters** – The total number of discovered adapters.
- **Physical Ports** – The number of discovered physical ports that can be managed by this host.
- **Function Summary** – Listed by protocol, the total number of discovered functions and targets.

Chapter 6: Managing Hosts

This chapter describes viewing host information, managing host groups, and searching for hosts.

6.1 Viewing Host Information

Two tabs show host information: the **Host Information** tab (Figure 11) and the **Host Driver Parameters** tab (Figure 46). The **Host Information** tab is read-only. The **Host Driver Parameters** tab enables you to view and define adapter driver settings for a specific host. See Section 7.17.2, [Host Driver Parameters Tab](#), for more information about the **Host Driver Parameters** tab.

To view the **Host Information** tab, perform these steps:


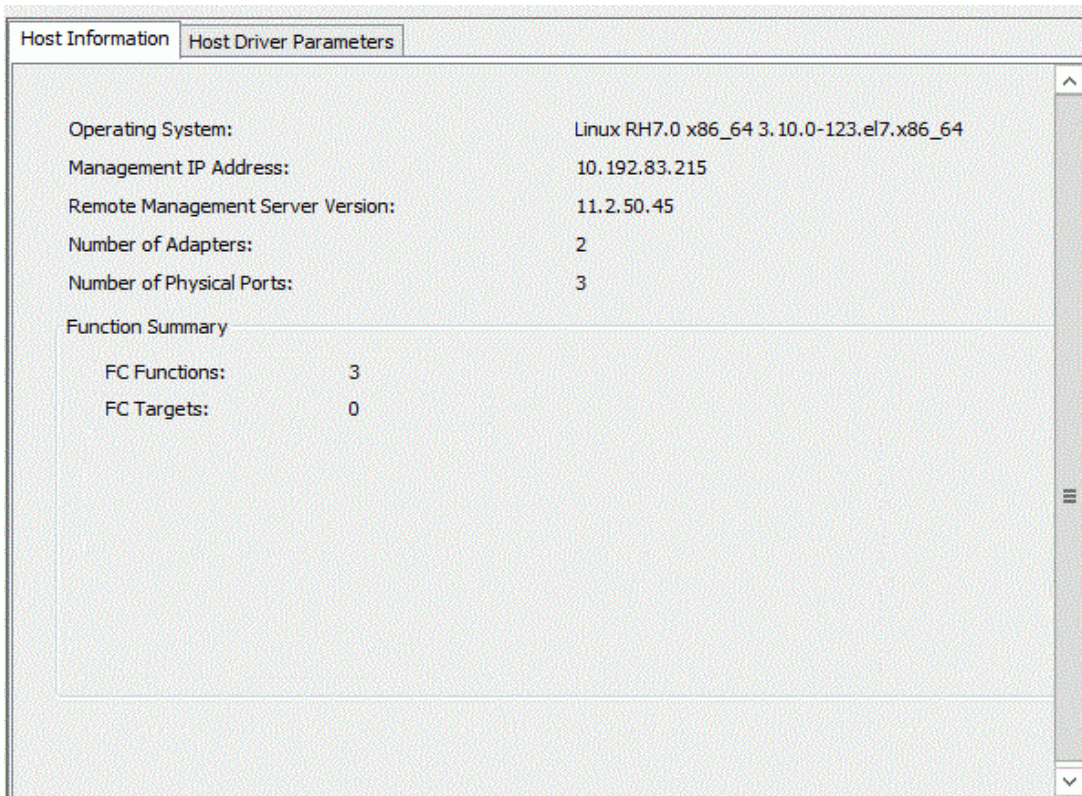
1. Perform one of the following tasks:
 - From the **View** menu, click **Group Adapters by Host Name**.
 - From the toolbar, click the  **Group Adapters by Host Name** button.
2. Select a host in the discovery-tree.
3. Select the **Host Information** tab (Figure 11).

Figure 11: Host Information Tab



The **Host Information** tab ([Figure 11](#)) displays the following fields:

- **Operating System** – Details about the installed operating system.
- **Management IP Address** – The host's IP address; for example, 138.239.82.131. **Local Host** is displayed if you selected the host from which you are running the OneCommand Manager application.
- **Remote Manager Server Version** – The version of the OneCommand Manager application server that is running on the host.
- **Number of Adapters** – The number of adapters installed in the host.
- **Number of Physical Ports** – The number of discovered physical ports that can be managed by this host.
- **CIM Provider Version** – If the host is being managed using the CIM interface, displays the version of the Emulex CIM Provider that is running on the remotely managed system.

NOTE: The **CIM Provider Version** field appears only if the host is managed through the CIM interface.

The Function Summary area has the following information:

- **FC Functions** – The number of FC functions running on the discovered adapters on this host.
- **FC Targets** – The number of FC targets discovered on the FC functions on this host.
- **VPorts** – The number of discovered virtual ports that can be managed by this host (not supported on VMware ESXi servers being managed through the CIM interface).

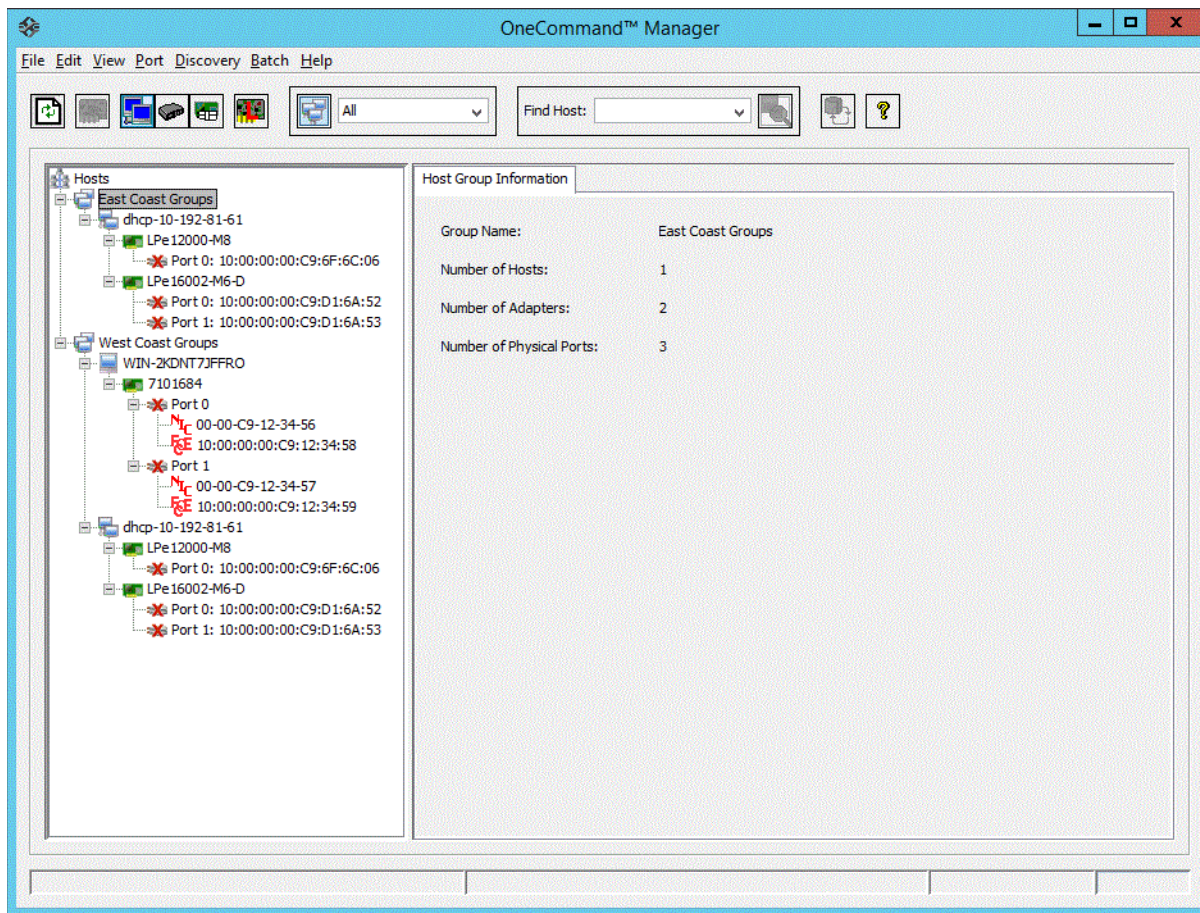
6.2 Viewing Host Grouping Information

The **Host Group Information** tab ([Figure 12](#)) displays information about the selected host group, such as the group name and the total number of hosts. See [Section 6.3, Grouping Hosts](#), to learn about creating host groups.

NOTE: Host grouping is not supported for VMware.

To view host grouping information, from the discovery-tree ([Figure 4](#)), select the host group whose information you want to view.

Figure 12: Host Group Information Tab



The following Host Group Information fields are displayed:

- **Group Name** – The name of the selected group.
- **Number Hosts** – The total number of hosts assigned to the group.
- **Number of Adapters** – The total number of discovered adapters in the group.
- **Number of Physical Ports** – The total number of ports in the group.

6.3 Grouping Hosts

The OneCommand Manager application enables you to assign related hosts to host groups. Typically, hosts within the same host group share some common functions, or they may simply reside within the same organizational unit within an enterprise, such as a Payroll group or a Shipping/Receiving group.


You can display the hosts in the discovery-tree (Figure 4) in either a group-centric format or in the host-based flat format. The Host grouping capability is available in the Host view, vPort view, or Fabric view mode.

NOTE: The same fabric can appear under more than one host group. For example, some ports on the fabric can be attached to ports and hosts in one host group, and other ports on the same fabric can be attached to ports and hosts in a different host group.


You can also perform batch operations, such as firmware download and driver parameter updates on a selected set of groups. See [Section 8.2, Updating Firmware for Multiple Adapters](#), for more information.

NOTE: Grouping hosts is not supported on VMware.


To display all hosts without grouping, perform one of the following tasks:

- From the **View** menu, clear **Show Groups**.
- From the toolbar, clear the  **Show Host Groups** button.

To display all hosts groups, perform these steps:

1. Perform one of the following tasks:
 - From the **View** menu, select **Show Groups**.
 - From the toolbar, click the  **Show Host Groups** button.
2. From the **Available Host Group** list, choose **All**.

To display all hosts assigned to a particular group, perform these steps:

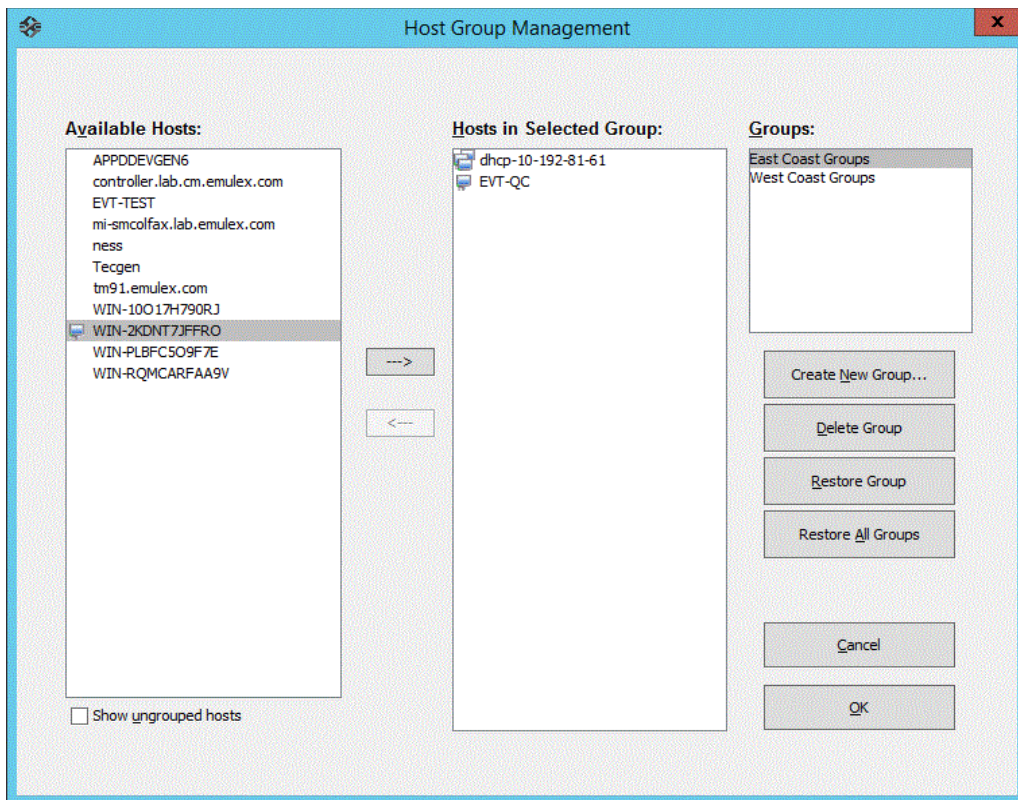
1. Perform one of the following tasks:
 - From the **View** menu, select **Show Groups**.
 - From the toolbar, click the  **Show Host Groups** button.
2. From the **Available Host Group** list, choose the group whose hosts you want to view.

6.3.1 Managing Host Groups

Use the **Host Group Management** dialog ([Figure 13](#)) to create and delete host groups, add and remove hosts, and restore host groups.

NOTE: Managing host groups is not supported on VMware.



Figure 13: Host Group Management Dialog



The following Host Group Management fields are displayed:

- **Available Hosts** – The list of hosts that can be added to a host group. You can select a host and right-click to see its group assignments.
- **Show ungrouped hosts** – If selected, displays only hosts that are currently assigned to a host group.
- **Hosts in Selected Group** – The list of hosts assigned to the currently selected host group.
- **Groups** – The list of currently defined host groups. If you select a group in this list, its host members appear in the **Hosts in Selected Group** list.

The following icons are used in the **Host Group Management** window:


-  Indicates that the host is currently assigned to a single host group.
-  Indicates that the host is currently assigned to multiple host groups.

6.3.2 Creating a Host Group

To create a new host group, perform these steps:

1. From the **View** menu, select **Manage Groups**. The **Host Group Management** dialog appears (Figure 13).
2. Click **Create New Group**. The **Create New Host Group** dialog appears (Figure 14).

Figure 14: Create New Host Group Dialog




3. Enter the name of the group you want to create and click **OK**. The new group appears in the Groups list on the **Host Group Management** dialog.

6.3.3 Deleting a Host Group

To delete a host group, perform these steps:

1. From the **View** menu, select **Manage Groups**. The **Host Group Management** dialog appears (Figure 13).
2. From the **Groups** list, select the group that you want to delete. The **Host Group Management** warning popup appears (Figure 15).

Figure 15: Host Group Management Warning Popup



3. Click **Yes** to delete the selected host group.

6.3.4 Adding a Host to a Host Group

To add a host to a group, perform these steps:

1. From the **View** menu, select **Manage Groups**. The **Host Group Management** dialog appears (Figure 13).
2. From the **Groups** list, select the group to which you want to add the host.
3. From the **Available Hosts** list, select the host you want to add (or select multiple hosts by using Ctrl-Click or Shift-Click), and click the Right Arrow. The selected host is removed from the Available Hosts list and is added to the Hosts in Selected Group list.
4. Click **OK** to commit your changes. The discovery-tree (Figure 4) displays the new configuration.

6.3.5 Removing a Host from a Host Group

To remove a host from a host group, perform these steps:

1. From the **View** menu, select **Manage Groups**. The **Host Group Management** dialog appears (Figure 13).
2. From the **Groups** list, select the group containing the host you want to remove.
3. From the **Hosts in Selected Group** list, select the host you want to remove and click the Left Arrow. The selected host is removed from the Hosts in Selected Group list and is added to the Available Hosts list.
4. Click **OK** to commit your changes. The discovery-tree (Figure 4) displays the new configuration.

6.3.6 Restoring a Host Group

To restore a host group, perform these steps:

1. From the **View** menu, select **Manage Groups**. The **Host Group Management** dialog appears (Figure 13).
2. Click **Restore Group** to return the configuration settings for the currently selected host group to those in use when the dialog was opened.

NOTE: If the currently selected group was created during the current configuration session, clicking **Restore Group** deletes the new group name.

6.3.7 Restoring All Host Groups

To restore all host groups, perform these steps:

1. From the **View** menu, select **Manage Groups**. The **Host Group Management** dialog appears (Figure 13).
2. Click **Restore All Groups** to return the entire host group configuration to the state that existed when the dialog was opened. All host group assignments are returned to their original configuration. Newly added host groups yet to be committed are removed, and host groups that were deleted are restored.

6.3.8 Exporting Host Grouping Configurations

To export the host grouping configuration to a remote host, you must copy the various host group configuration files from the host on which the configuration was created to the remote host. Copy the entire contents of the `config/hostgroups` subdirectory under the OneCommand installation directory to the equivalent location on the remote system.

The `hostgroups` configuration file locations for the supported platforms are:

- **Windows:** `<InstallationDriveLetter>:\Program Files\Emulex\Util\Config\hostgroups`
- **Linux:** `/usr/sbin/ocmanager/config/hostgroups`
- **Solaris:** `/opt/ELXocm/config/hostgroups`


The host group configuration files are completely interchangeable between different operating systems. For example, the host group configuration files created on a Solaris host can be copied directly to a Linux or Windows host, with no conversion required.

6.4 Searching for Hosts in the Discovery-Tree

The OneCommand Manager application enables you to search the discovery-tree (Figure 4) for a particular host by the host's name. If the specified host name is found, the discovery-tree scrolls up or down to bring the desired host name into view.

This capability is especially useful when you are searching for a host in a large installation with hundreds or thousands of hosts. It is also helpful in the Fabric view mode because the ports on a specific host can be dispersed among several fabrics, making the ports on that host difficult to find in the discovery-tree (Figure 4).

To search for a host, perform these steps:

1. Perform one of the following tasks:
 - From the **Edit** menu, select **Find** and enter the name of the host you are searching for into the **Find Host** field.
 - From the toolbar, enter the name of the host you are searching for into the **Find Host** field.
2. From the toolbar, click the  **Find Host** button or press **Enter** on the keyboard.

The host you are searching for is highlighted in the discovery-tree (Figure 4).

Selecting the **Find Next** option on the **Edit** menu, or pressing **F3**, enables you to continue searching for more instances of the name you specified.

Chapter 7: Managing Adapters and Ports

This chapter describes the various adapter and port management functions, you can perform using the OneCommand Manager application.

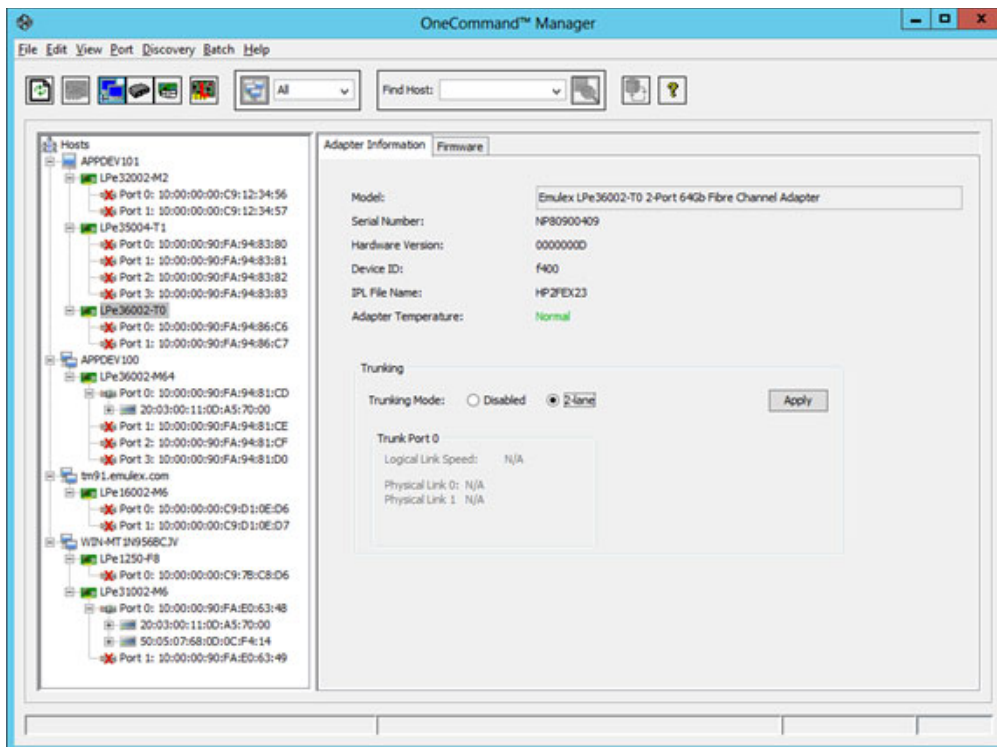
7.1 Viewing Adapter Information

When you select an adapter from the discovery-tree (Figure 4), the **Adapter Information** tab (Figure 16) contains general attributes associated with the selected adapter.

To view adapter information, perform these steps:

1. Select the **Host**, **Fabric**, or **Virtual Ports** view.
2. Select an adapter in the discovery-tree. The **Adapter Information** tab appears (Figure 16).

Figure 16: Adapter Information Tab (Two-Lane Trunking Supported)



The following **Adapter Information** tab fields are displayed:

- **Model** – The complete model name of the adapter.
- **Serial Number** – The manufacturer's serial number for the adapter.
- **Hardware Version** – The JEDEC ID.
- **Device ID** – The default device ID for the selected adapter.
- **IPL File Name** – The IPL file name for the selected adapter.

- **Adapter Temperature** – If the adapter's temperature is not available, **Not Supported** is displayed. If supported by the adapter, this field displays the following adapter temperature-related status messages:
 - **Normal:** The adapter's temperature is within normal operational range.
 - **Warning:** The adapter's temperature is beyond normal operational range. If the temperature continues to increase, the adapter shuts down. You must determine the cause of the temperature issue and fix it immediately. Check for system cooling issues. Common causes of system cooling issues include clogged air filters, inoperative fans, and air conditioning issues that cause high ambient air temperatures.
 - **Exceeds operational range – Adapter stopped:** The temperature has reached a critical limit, forcing the adapter to shut down. You must determine the cause of the temperature issue and fix it before resuming operation. Check for system cooling issues. Common causes of system cooling issues include clogged air filters, inoperative fans, and air conditioning issues that cause high ambient air temperatures.

After the system overheating issue is resolved and the adapter has cooled down, reboot the system, or if the system supports hot swapping, cycle the power of the adapter slot.

- **Trunking area:**
 - When supported by the adapter, you can combine multiple physical FC links to form a single logical link. Once created, you can view an aggregated port's logical link speed and physical link status. See [Section 7.1.1, Configuring Trunking](#), for additional information.

7.1.1 Configuring Trunking

NOTE:

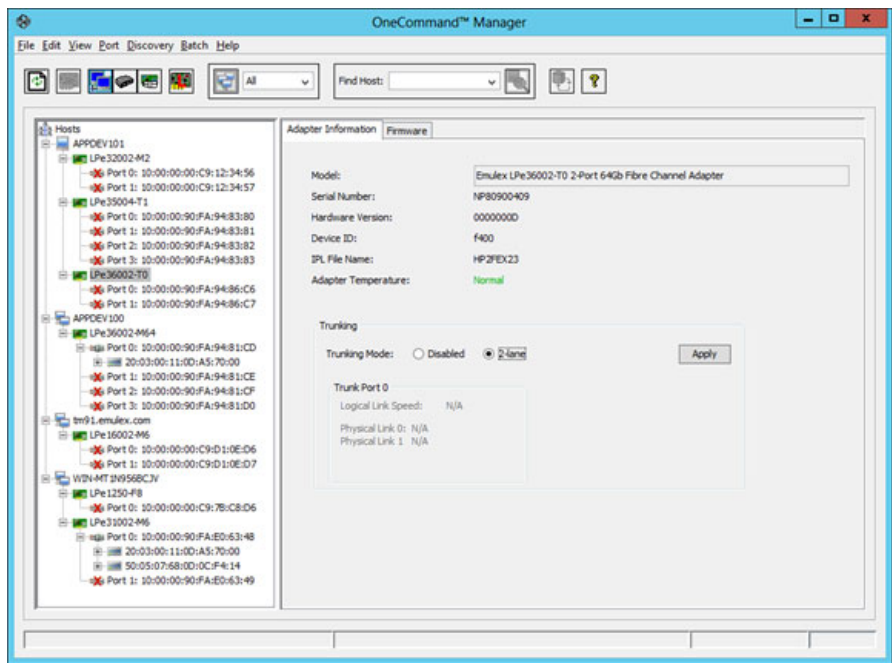
- Trunking is supported only on LPe35002 and LPe35004 adapters.
- Neither FA-PWWN nor Dynamic D_Port can coexist with the trunking feature on LPe35000-series adapters. If trunking is enabled, the firmware automatically disables FA-PWWN and Dynamic D_Port.
- Trunking is not supported at 8 Gb/s and the link will not come up at this speed.
- Before you configure trunking on the Emulex adapter, follow the instructions from Brocade® for configuring trunking on the switch.

The **Adapter Information** tab enables you to configure trunking (also called FC port aggregation), which combines multiple physical FC links to form a single logical link (aggregated port). The aggregated port's maximum link speed is the sum of the maximum link speeds of the individual physical links comprising the aggregated port. For example, an aggregated port comprised of two physical links running at 64 Gb/s each will have a potential logical (aggregate) link speed of 128 Gb/s. The actual link speed of the aggregated port depends on the states (active or non-active) of the individual physical links comprising the aggregated port.

The physical links comprising an aggregated port are referred to as lanes. Both 2-lane and 4-lane aggregated ports are supported. For dual-port adapters, only 2-lane port aggregation is possible. If 2-lane port aggregation is configured on a dual-port adapter, the two physical links are combined to form a single 2-lane aggregated port whose aggregate speed is potentially 128 Gb/s assuming both physical links are active.

LPe35004 adapters support both 2-lane port aggregation and 4-lane port aggregation. If 2-lane port aggregation is configured on an LPe35004 adapter, the four physical links on the adapter are divided among two separate aggregated ports. The two lowest numbered physical links form the first aggregated port, and the two highest number physical links will form the second aggregated port. If 4-lane port aggregation is configured on an LPe35004 adapter, all four physical links are combined to form a single 4-lane trunk whose aggregate speed is potentially 256 Gb/s, assuming all four links are active.

Figure 17: Adapter Information Tab (Two-Port Adapter with Two-Lane Trunking Enabled)




To set trunking, perform these steps:

1. In the discovery-tree (Figure 4), select the supported adapter port whose trunking you want to set.
2. Select the **Adapter Information** tab (Figure 17).
3. Select **Disabled**, **2-lane**, or **4-lane**.

NOTE: On an LPe35004 adapter, if 2-lane port aggregation is selected, the four physical links on the adapter are divided among two separate aggregated ports (Port 0 and Port 1). The two lowest numbered physical links form the first aggregated port (Port 0), and the two highest number physical links form the second aggregated port (Port 1).

4. Click **Apply**. The **Trunking Configuration** dialog appears notifying you that your changes require a system reboot.

Figure 18: Trunking Configuration Dialog



5. Click **OK** and reboot the system.

7.2 Viewing Port Information

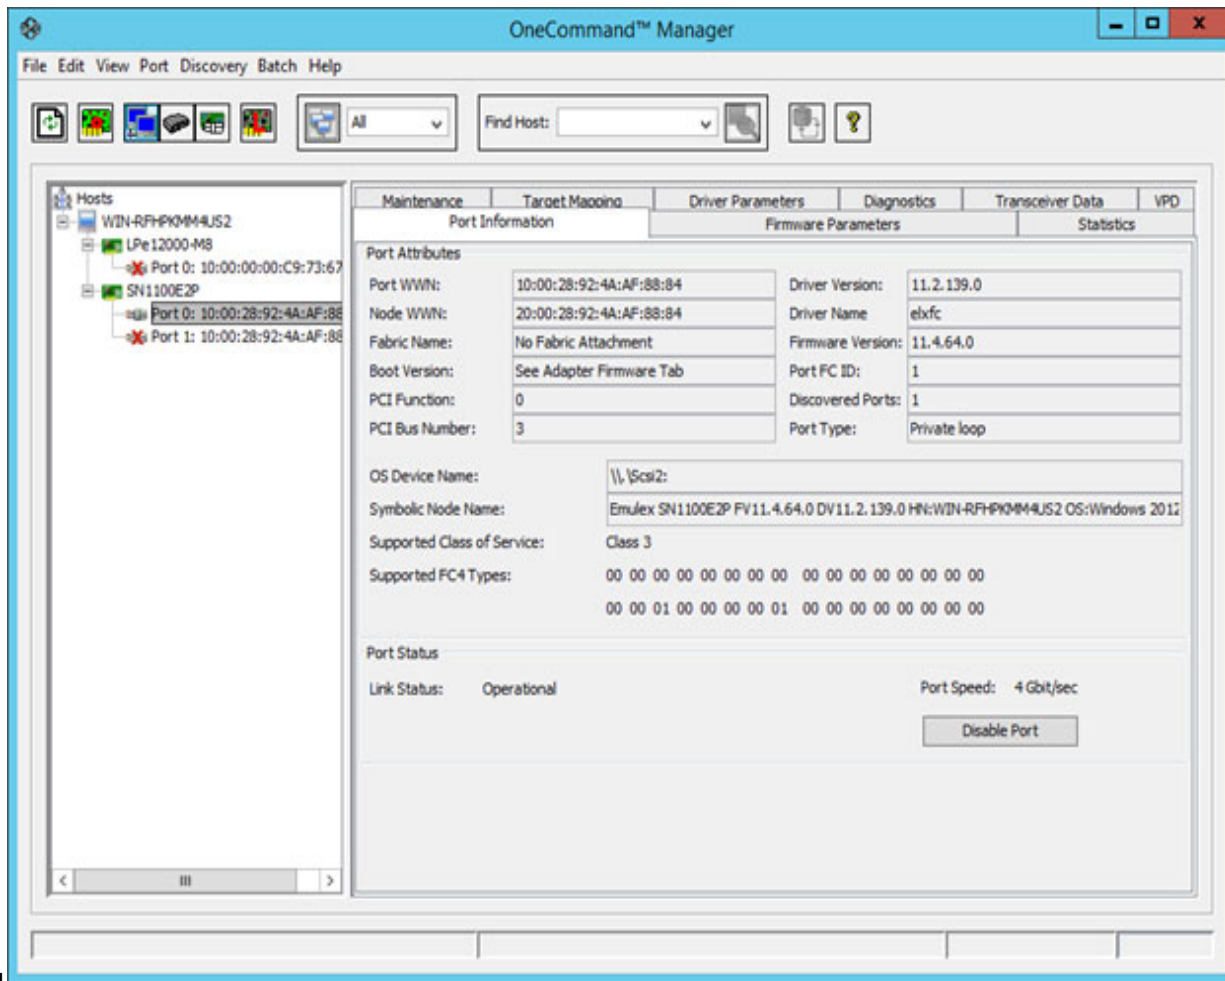
When you select a port from the discovery-tree, the **Port Information** tab (Figure 19) contains general attributes associated with the selected adapter.

NOTE: Some information cannot be retrieved when the port is offline.

To view port information, perform these steps:

1. Select the **Host** or **Fabric** view.
2. Select a port in the discovery-tree.
3. Select the **Port Information** tab (Figure 19).

Figure 19: Port Information Tab



The following **Port Information** tab fields are displayed:

- **Port Attributes** area:
 - **Port WWN** – The WWPN of the adapter.
 - **Node WWN** – The WWNN of the adapter.

- **Fabric Name** or **Host Name** – The **Fabric Name** field is displayed in the Host view. This is a 64-bit worldwide unique identifier assigned to the fabric. The **Host Name** field is displayed in the Fabric view and is the name of the host containing the adapter.
- **Boot Version** – The version of boot code installed on the selected adapter port. If the boot code is disabled, the field displays **Disabled**.
- **PCI Function** – The PCI function number assigned by the system.
- **PCI Bus Number** – The PCI bus number assigned to the FC function.
- **Driver Version** – The version of the driver installed for the adapter.
- **Driver Name** – The executable file image name for the driver as it appears in the Emulex driver download package.
- **Firmware Version** – The version of Emulex firmware currently active on the adapter port.
- **Port FC ID** – The FC ID for the selected adapter port.
- **Discovered Ports** – The number of ports found during discovery by the Emulex adapter driver.
- **Port Type** – The FC type of the selected adapter's port (not available if the port link is down).
- **OS Device Name** – The platform-specific name by which the selected adapter is known to the operating system.
- **Symbolic Node Name** – The FC name used to register the driver with the name server.
- **Supported Class of Service** – A frame delivery scheme exhibiting a set of delivery characteristics and attributes. Three classes of service include:
 - **Class 1** – Provides a dedicated connection between a pair of ports with confirmed delivery or notification of nondelivery.
 - **Class 2** – Provides a frame switched service with confirmed delivery or notification of nondelivery.
 - **Class 3** – Provides a frame switched service similar to Class 2 but without notification of frame delivery or nondelivery.
- **Supported FC4 Types** – A 256-bit (8-word) map of the FC-4 protocol types supported by the port containing the selected adapter.
- **Port Status** area:
 - **Link Status** – The status of the link on the selected adapter port.
 - **Port Speed** – The current port speed of the selected adapter port. For trunked ports, the maximum speed that the trunked port is capable of (if all ports in the trunk are up) is displayed.

7.2.1 Enabling and Disabling a Port

You can enable or disable a port from the **Port Information** tab. When you disable a port, you disable all functions for the port.

CAUTION! Do not disable a boot port; this could result in data loss or corruption.

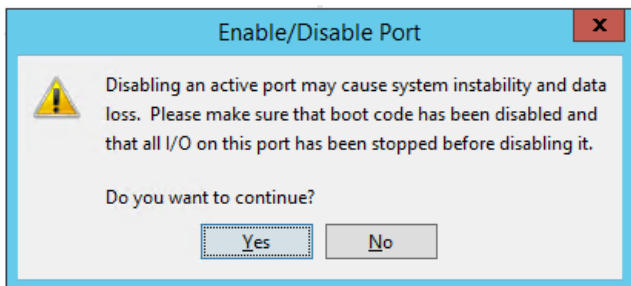
NOTE:

- Ensure that there is no I/O traffic on the port before disabling it.
- You must reset the adapter to activate the new value.

To enable or disable a port:

1. Select the **Host** or **Fabric** view.
2. Select a port in the discovery-tree.

3. From the **Port Information** tab, click **Enable Port** or **Disable Port**. The following popup appears.

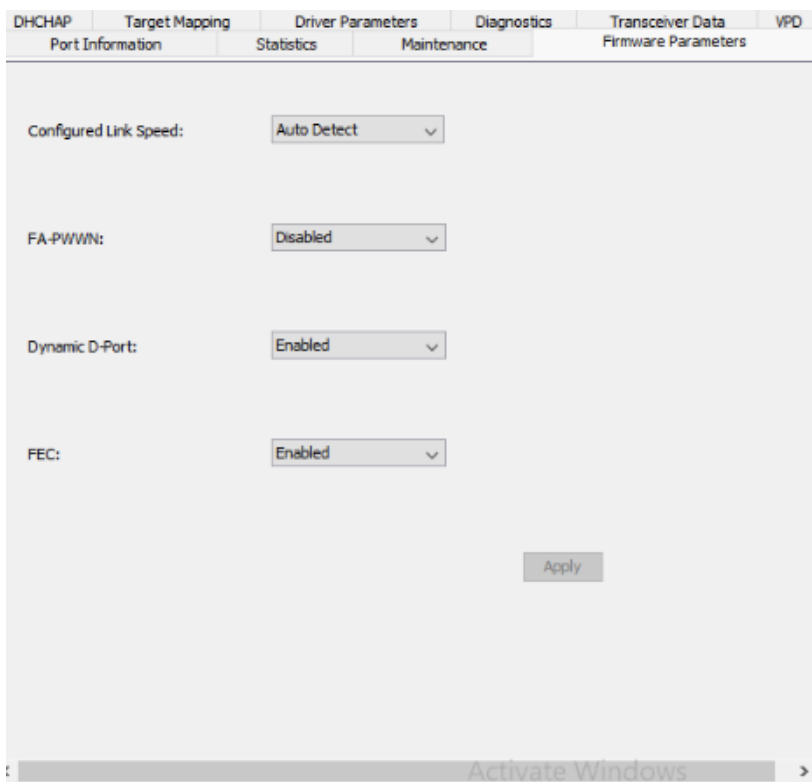


4. Click **Yes** to enable or disable the port.

7.3 Viewing Firmware Parameters

To view firmware parameters, select the **Firmware Parameters** tab (Figure 20).

Figure 20: Firmware Parameters Tab



The following **Firmware Parameters** tab fields are displayed:

- **Configured Link Speed** – This field displays the link speeds that are supported on the port. The list varies depending on the adapter type. The list also includes an **Auto Detect** option, which indicates that the link speed should be auto-negotiated.

NOTE: If an installed adapter does not support forced link speeds, the **Configured Link Speed** settings and the **Apply** button are not shown.

- **FA-PWWN** – This field displays the FA-PWWN status of the port. FA-PWWN allows a switch to assign a virtual WWPN to the initiator. **Disabled** is the default setting. (Not supported on LPe12000-series adapters.)
- **Dynamic D_Port** – This field displays the Dynamic D_Port status of the port. **Enabled** is the default setting. (Not supported on LPe12000-series adapters.)
- **FEC** – This field displays the FEC status of the port. **Enabled** is the default setting. (Not supported on LPe12000-series adapters.)

NOTE:

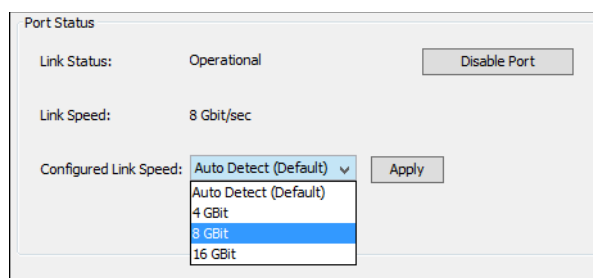
- FA-PWWN is not supported on VMware ESXi servers being managed through the CIM interface.
- The switch must support FA-PWWN. Refer to the documentation that accompanied the switch for instructions on configuring FA-PWWN on the switch.
- The link is toggled if FA-PWWN is enabled, but the switch does not support FA-PWWN.
- After enabling or disabling FA-PWWN, the port must be reset for changes to take effect.
- When a new WWPN is assigned using FA-PWWN, persistently stored configuration information associated with the original WWPN, such as driver parameters and LUN frame priority settings, is not applied to the newly assigned WWPN. The configuration information associated with the original WWPN must be reconfigured for the new WWPN.
- FA-PWWN is not available when trunking is enabled.
- Dynamic D_Port and FA-PWWN cannot be enabled simultaneously. If Dynamic D_Port is enabled and you want to enable FA-PWWN, you must first disable Dynamic D_Port. If FA-PWWN is enabled and you want to enable Dynamic D_Port, you must first disable FA-PWWN.
- Dynamic D_Port is not available when trunking is enabled.
- If Dynamic D_Port is enabled on an adapter, it is not supported in a direct-connect point-to-point environment. The adapter must be connected to a switch.

7.3.1 Configuring Link Speed

To configure a link speed, perform these steps:


1. Select the **Firmware Parameters** tab (Figure 20).
2. Select a link speed from the **Configured Link Speed** list (Figure 21).


Figure 21: Port Status Dialog



3. Click **Apply** to set the new link speed. The **Apply** button is enabled only if the currently selected link speed does not match the current speed.

If the speed has been set successfully, the following message states that the port must be reset to activate the new speed setting.



4. Click **OK**.
5. From the toolbar, click the  **Reset Port** button.

In some situations, the currently configured link speed is not in the supported speed list for the port. This situation can occur if a new SFP is installed that supports a different set of link speeds than the previously installed SFP. If the currently configured link speed is not in the supported speed list, the following message is displayed:

Warning: The currently configured port speed is not a valid supported speed.
Please select a link speed and click Apply.

The **Apply** button remains enabled until you select a valid port speed.

If the installed SFP is not supported by the adapter, you cannot configure a link speed. In this situation, the following message is displayed:

Unsupported optics installed.

If an adapter does not support forced link speed, the **Firmware Parameters** tab does not show a Link Speed list.

7.3.2 Enabling and Disabling FA-PWWN


NOTE:


- FA-PWWN is not supported on VMware ESXi servers being managed through the CIM interface.
- The switch must support FA-PWWN. Refer to the documentation that accompanied the switch for instructions on configuring FA-PWWN on the switch.
- The link is toggled if FA-PWWN is enabled, but the switch does not support FA-PWWN.
- After enabling or disabling FA-PWWN, the port must be reset for changes to take effect.
- When a new WWPN is assigned using FA-PWWN, persistently stored configuration information associated with the original WWPN, such as driver parameters and LUN frame priority settings, is not applied to the newly assigned WWPN. The configuration information associated with the original WWPN must be reconfigured for the new WWPN.
- FA-PWWN is not available when trunking is enabled.
- Dynamic D_Port and FA-PWWN cannot be enabled simultaneously. If Dynamic D_Port is enabled and you want to enable FA-PWWN, you must first disable Dynamic D_Port. If FA-PWWN is enabled and you want to enable Dynamic D_Port, you must first disable FA-PWWN.

To enable or disable FA-PWWN, perform these steps:

1. Select the **Firmware Parameters** tab ([Figure 20](#)).
2. Select **Enable** or **Disable** from the **FA-PWWN** list.
3. Click **Apply**. The **Change Firmware Parameter** popup appears ([Figure 22](#)).
4. Click **OK**.

Figure 22: Change Firmware Parameter Popup



5. From the toolbar, click the  **Reset Port** button.
6. Restart the OneCommand Manager application.

NOTE: The following popup appears if the FA-PWWN change was unsuccessful (Figure 23).

Figure 23: Change Firmware Parameter Popup



7.3.3 Disabling and Enabling Dynamic D_Port Tests

Dynamic D_Port is enabled on the HBA by default. See [Section 10.5, Running D_Port Tests](#) for more information about Dynamic D_Port.

NOTE: D_Port and FA-PWWN cannot be enabled simultaneously. If Dynamic D_Port is enabled and you want to enable FA-PWWN, you must first disable Dynamic D_Port. If FA-PWWN is enabled and you want to enable Dynamic D_Port, you must first disable FA-PWWN.


NOTE: If Dynamic D_Port is enabled on an adapter, it is not supported in a direct-connect point-to-point environment. The adapter must be connected to a switch.

To disable or enable Dynamic D_Port tests, perform these steps:

1. From the discovery-tree (Figure 4), select the port on which you want to enable or disable the Dynamic D_Port test.
2. Select the **Firmware Parameters** tab (Figure 20) and choose **Enabled** or **Disabled** from the **Dynamic D_Port** list.
3. Click **Apply**.

7.3.4 Enabling and Disabling FEC

To enable or disable FEC, perform these steps:

1. Select the **Firmware Parameters** tab (Figure 20).
2. Select **Enable** or **Disable** from the **FEC** list.
3. Click **Apply**. The **Change Firmware Parameter** popup (Figure 22) appears.
4. Click **OK**.
5. From the toolbar, click the  **Reset Port** button.

- Restart the OneCommand Manager application.

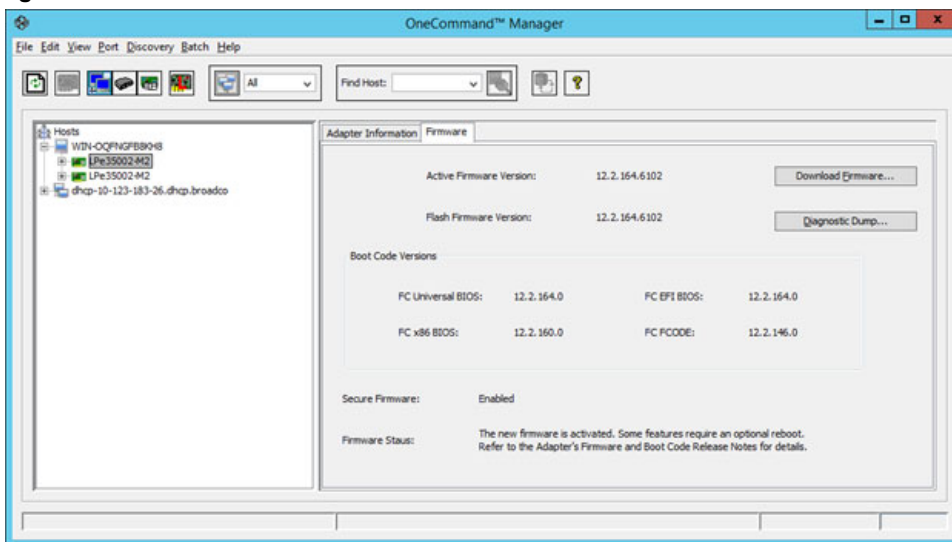
7.4 Viewing Firmware Information

Use the **Firmware** tab (Figure 24) to download firmware and create diagnostic dumps. (Not supported on LPe12000-series HBAs; see Section 7.11, Viewing Maintenance Information.)

To view firmware information, perform these steps:

- Select the **Host** view.
- Select an adapter in the discovery-tree (Figure 4).
- Select the **Firmware** tab (Figure 24).

Figure 24: Firmware Tab



The following **Firmware** tab fields are displayed:

- **Active Firmware Version** – Displays the firmware version currently being used by the adapter.
- **Flash Firmware Version** – Displays the flash firmware version currently being used by the adapter.
- **Boot Code Versions** – Displays the versions of the adapter boot code. It has no relation to the FC boot code versions.
- **Secure Firmware** – Displays the secure status of the firmware.
- **Firmware Status** – Displays any activation requirements necessary to use the new firmware.

NOTE: The buttons on the **Firmware** tab are not available in read-only mode.

See Chapter 8, [Updating Adapter Firmware](#), for information on updating firmware.

See Section 10.7, [Creating Diagnostic Dumps](#), for information about performing a diagnostic dump.

7.5 Viewing and Clearing Statistics

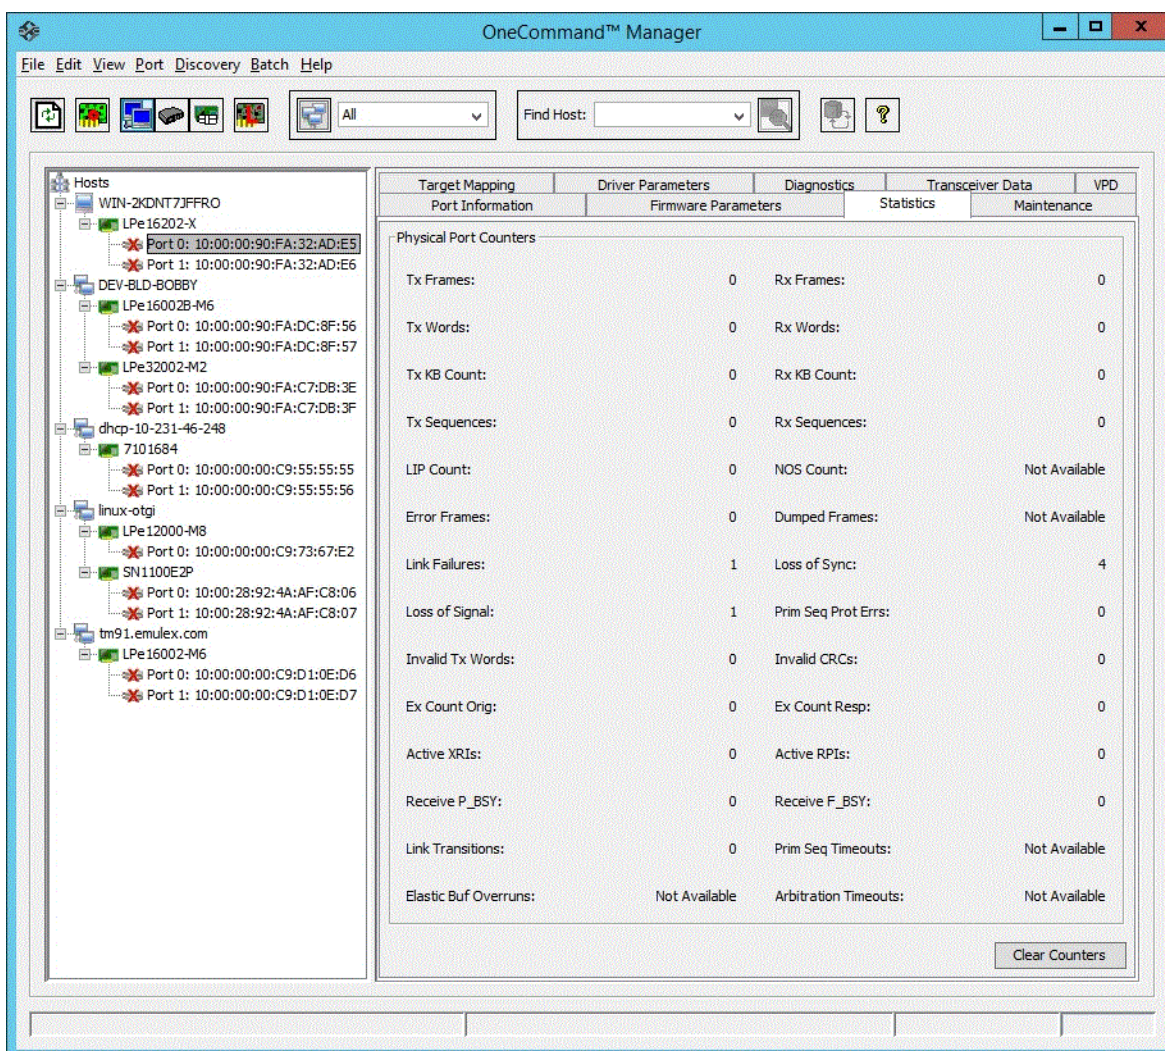
When you select a port from the discovery-tree, the **Statistics** tab (Figure 25) provides cumulative totals for various error events and statistics on the port. If supported by the adapter, you can also clear all the values displayed in the tab.

NOTE: Some statistics are cleared when the adapter is reset.

To view statistics, perform these steps:

1. Select the **Host** or **Fabric** view.
2. Select a port in the discovery-tree.
3. Select the **Statistics** tab.

Figure 25: Statistics Tab



The following **Statistics** tab fields are displayed:

- **Tx Frames** – FC frames transmitted by this FC function.
- **Tx Words** – FC words transmitted by this FC function.

- **Tx KB Count** – FC kilobytes transmitted by this FC function.
- **Tx Sequences** – FC sequences transmitted by this FC function.
- **LIP Count** – The number of LIP events that have occurred for the FC function. This field is not supported if the topology is not arbitrated loop. Loop initialization consists of the following:
 - Temporarily suspending loop operations.
 - Determining whether loop-capable ports are connected to the loop.
 - Assigning AL_PA IDs.
 - Providing notification of configuration changes and loop failures.
 - Placing loop ports in the monitoring state.
- **Error Frames** – The number of frames received with CRC errors.
- **Link Failures** – The number of times the link has failed. A link failure is a possible cause of a timeout.
- **Loss of Signal** – The number of times the signal was lost.
- **Invalid Tx Words** – The total number of invalid words transmitted by this FC function.
- **Ex Count Orig** – The number of FC exchanges originating on this FC function (not supported on VMware ESXi servers being managed through the CIM interface).
- **Active XRIs** – The number of active XRIs (not supported on VMware-based ESXi platforms using the CIM interface).
- **Received P_BSY** – The number of FC port-busy link response frames received.
- **Link Transitions** – The number of times the SLI port sent a link attention condition.
- **Elastic Buf Overruns** – The number of times the link interface has had its elastic buffer overrun.
- **Rx Frames** – The number of FC frames received by this FC function.
- **Rx Words** – The number of FC words received by this FC function.
- **Rx KB Count** – The received kilobyte count by this FC function.
- **Rx Sequences** – The number of FC sequences received by this FC function (not supported on VMware ESXi servers being managed through the CIM interface).
- **NOS Count** – The number of NOS events that have occurred on the switched fabric. The NOS count is not supported for Emulex Windows drivers or for arbitrated loop.
- **Dumped Frames** – The number of frames that were lost because of a lack of host buffers available. This option is not currently supported for the Storport Miniport driver or the driver for Solaris.
- **Loss of Sync** – The number of times loss of synchronization has occurred.
- **Prim Seq Prot Errs** – The Primitive Sequence Protocol error count. This counter is incremented whenever there is any type of protocol error.
- **Invalid CRCs** – The number of frames received that contain CRC failures.
- **Ex Count Resp** – The number of FC exchange responses made by this FC function (not supported on VMware ESXi servers being managed through the CIM interface).
- **Active RPIs** – The number of RPIs (not supported on VMware ESXi servers being managed through the CIM interface).
- **Receive F_BSY** – The number of FC fabric-busy link response frames received.
- **Primitive Seq Timeouts** – The number of times a primitive sequence event timed out (not supported on VMware ESXi servers being managed through the CIM interface).
- **Arbitration Timeouts** – The number of times the arbitration loop has timed out. Large counts could indicate a malfunction somewhere in the loop or heavy usage of the loop (not supported on VMware ESXi servers being managed through the CIM interface).

If supported by the adapter, click **Clear Counters** to clear all the values displayed on the tab.

7.6 Viewing Virtual Port Information

The **Discovery Information** tab (Figure 26) displays information about virtual ports and their associated targets and LUNs.

To view virtual port information, perform these steps:


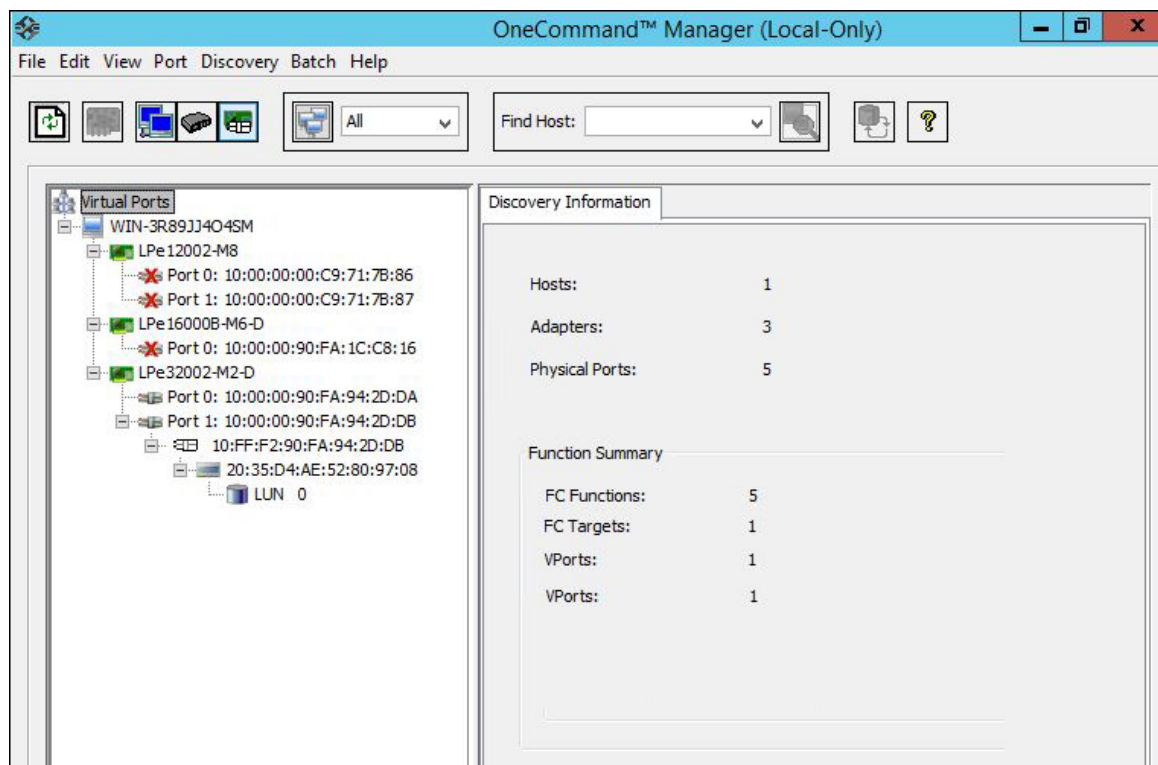
1. Perform one of the following tasks:
 - From the **View** menu, select **Group Adapters by Virtual Port**.
 - From the toolbar, click the  **Group Adapters by Virtual Port** button.The **Discovery Information** tab appears (Figure 26).

Figure 26: Discovery Information Tab



The following **Discovery Information** tab fields are displayed:

- **Hosts** – The total number of hosts discovered in the SAN.
- **Adapters** – The total number of adapters discovered in the SAN.
- **Physical Ports** – The total number of physical ports discovered in the SAN.
- **Function Summary** area:
 - **FC Functions** – The total number of FC functions discovered in the SAN.
 - **FC Targets** – The total number of FC targets discovered in the SAN.
 - **VPorts** – The total number of virtual ports discovered in the SAN.

7.7 Creating and Deleting Virtual Ports

This section describes how to create and delete virtual ports.

7.7.1 Creating Virtual Ports

Using the **Virtual Ports** tab (Figure 27), you can automatically generate the WWPN for the virtual port based on the WWPN for the physical port, or you can manually type the WWPN.

NOTE:

- The OneCommand Manager application cannot create or delete virtual ports on VMware ESXi server systems. Although VMware ESXi server supports NPIV, only VMware management tools can create or delete virtual ports.
- In Linux, virtual ports do not persist across system reboots.

The NPIV driver parameter must be enabled before you attempt to create a virtual port. The driver parameter name varies slightly depending upon your operating system:

- For Windows: From the **Driver Parameters** tab, highlight **enableNPIV**, and then select **Enable**. On the Storport Miniport system, the `SLIMode` driver parameter must also be set to 0 or 3.
- For Solaris: `enable-npiv`
- For Linux: `lpfc_enable_npiv`

See [Section 7.17, Configuring the Driver Parameters](#), for more information on enabling driver parameters.

To create a virtual port, perform these steps:


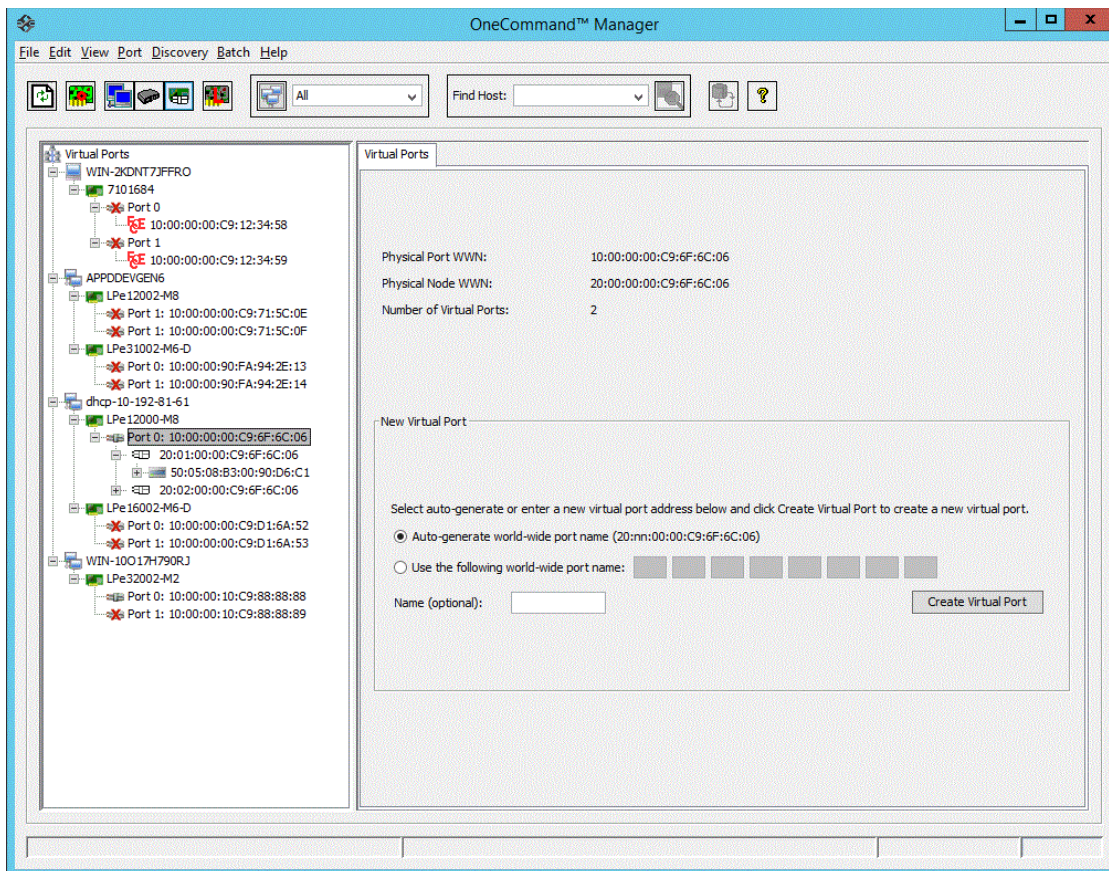
1. Perform one of the following tasks:
 - From the **View** menu, select **Group Adapters by Virtual Ports**.
 - From the toolbar, click  **Group Adapters by Virtual Ports**.
2. From the discovery-tree, select the FC function on which you want to create a virtual port. The **Virtual Ports** tab appears (Figure 27).

Figure 27: Virtual Ports Tab



3. Perform one of the following tasks:

- Select **Auto-generate world-wide port name**. The OneCommand Manager application creates the unique WWPN for the new virtual port based on the WWPN of the FC function. This option allows you to automatically create up to 255 unique virtual ports for each physical port. It also has the advantage that the new WWPN is unique.

NOTE: After auto-generating 255 unique virtual ports, you cannot auto-generate any more virtual ports even if you delete existing auto-generated virtual ports. However, you can still enter your own WWPN to create a virtual port.

- Select **Use the following world-wide port name** and enter a unique WWPN. You can create as many virtual ports as you want. A valid port name must have one of the following formats:

10:00:xx:xx:xx:xx:xx:xx
2x:xx:xx:xx:xx:xx:xx:xx
3x:xx:xx:xx:xx:xx:xx:xx
5x:xx:xx:xx:xx:xx:xx:xx

where *x* is a hexadecimal value.

CAUTION! Make sure that a manually entered WWPN is unique to your particular SAN. Otherwise, a non-functioning SAN and data loss could occur.

- 4. Enter an optional name for the virtual port if you want. You can give the new virtual port any name you want up to 99 characters in length. This name is used as part of the Symbolic Node Name for the vPort.

5. Click **Create Virtual Port**. A popup notifies you that the virtual port was created. The popup also displays the new virtual port's WWPN. Each virtual port has its own WWPN, but its WWNN is the same as the physical port's WWNN.

NOTE: If you entered a WWPN that is already in use, you are prompted to enter another WWPN.

6. Click **OK**. The new virtual port is added to the discovery-tree ([Figure 4](#)) under the physical port where it was created, and the **Number of Virtual Ports** field is updated.

NOTE: The OneCommand Manager application automatically refreshes its discovery after a virtual port is created. However, targets for a new virtual port might not be discovered during the refresh. Therefore, you must refresh the discovery until the targets appear under the virtual port in the discovery-tree ([Figure 4](#)).

7.7.2 Deleting Virtual Ports

NOTE: The OneCommand Manager application cannot create or delete virtual ports on VMware ESXi server systems. Although VMware ESXi server supports NPIV, only VMware management tools can create or delete virtual ports.

To delete a virtual port, perform these steps:


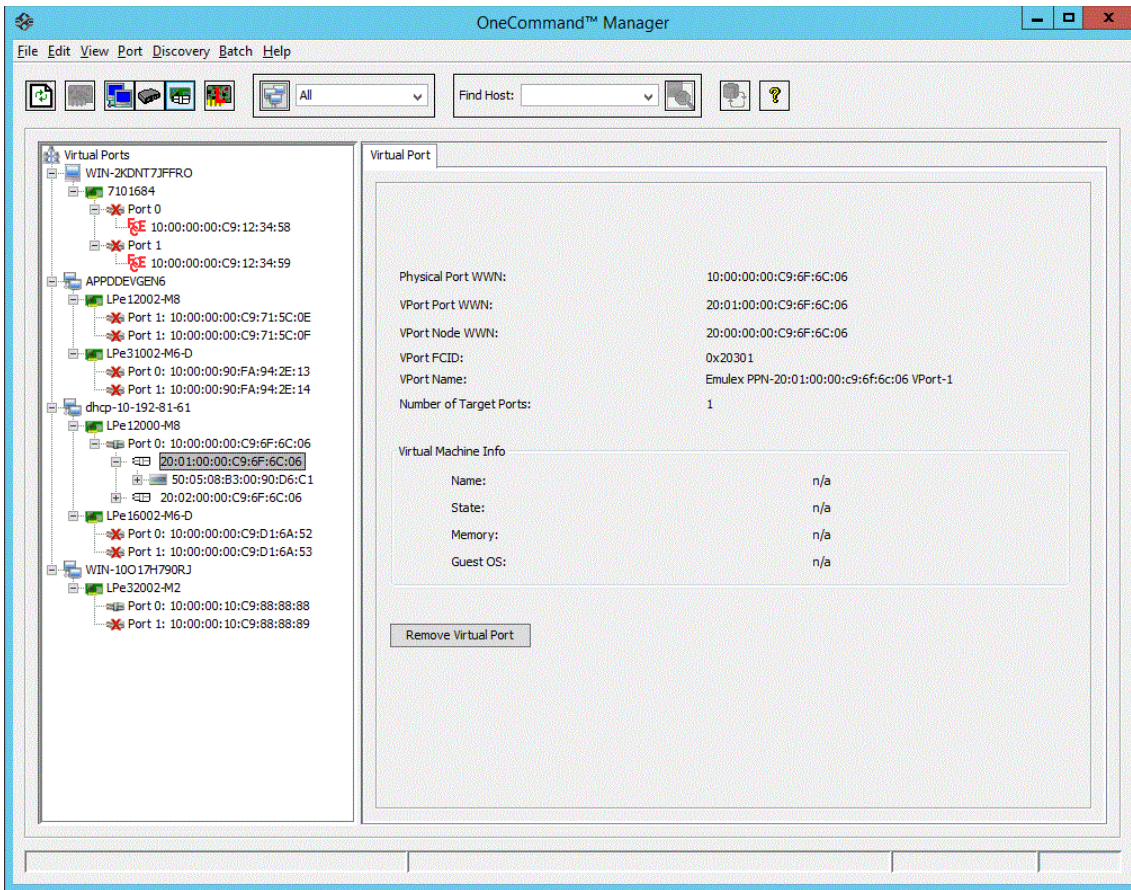
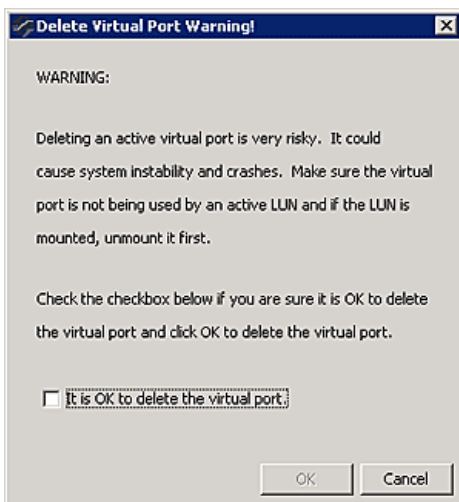
1. Perform one of the following tasks:
 - From the **View** menu, select **Group Adapters by Virtual Ports**.
 - From the toolbar, click the  **Group Adapters by Virtual Ports** button.
2. From the discovery-tree, select the virtual port that you want to delete. The **Virtual Port** tab appears ([Figure 28](#)).

Figure 28: Virtual Port Tab



3. Click **Remove Virtual Port**. The **Delete Virtual Port Warning** popup appears (Figure 29).

Figure 29: Delete Virtual Port Warning Popup



NOTE: The link on the physical port must be up to delete a virtual port. The **Remove Virtual Port** button on the **Virtual Port** tab is disabled if the link is down.

4. Select **It is OK to delete the virtual port** and click **OK**. You are notified that the virtual port is no longer available and that it was removed from the discovery-tree (Figure 4).
5. Click **OK**.

7.8 Viewing Fabric Information

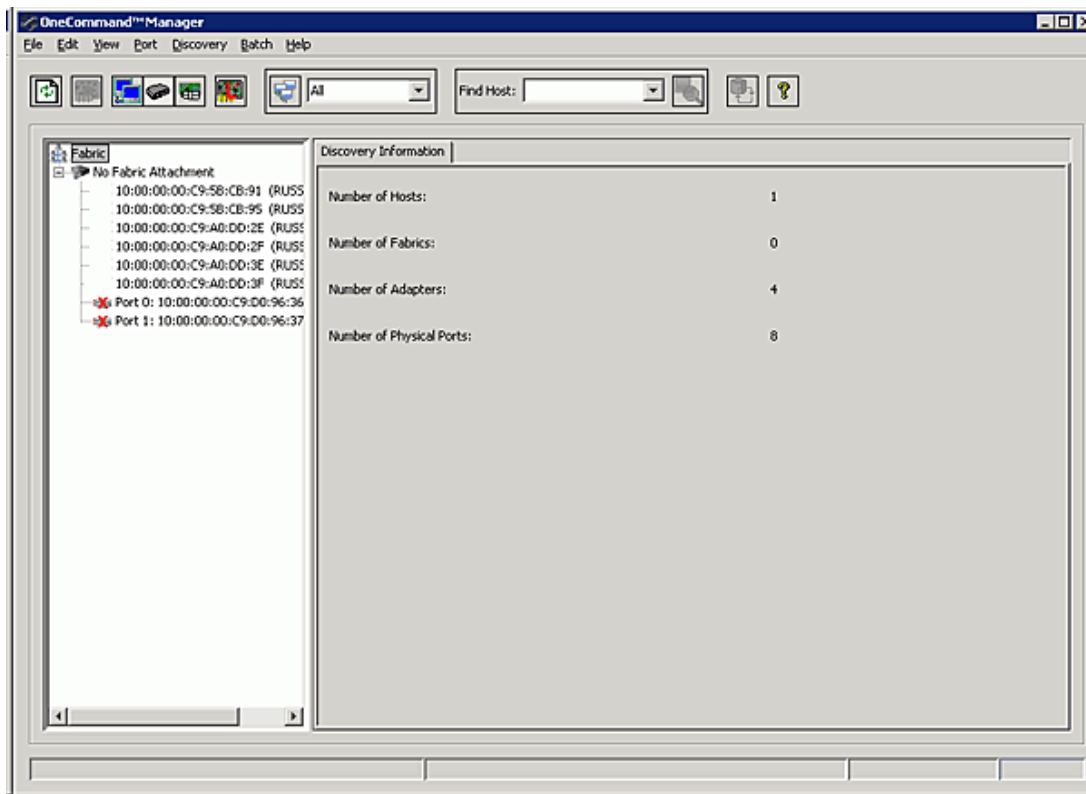
The **Discovery Information** tab (Figure 30) contains information about the selected fabric.

To view fabric discovery information, perform one of the following tasks:

- From the **View** menu, select **Group Adapters by Fabric Address**.
- From the toolbar, click the  **Group Adapters by Fabric Address** button.

The **Discovery Information** tab is displayed (Figure 30).

Figure 30: Fabric Discovery Information



The following **Discovery Information** tab fields are displayed:

- **Number of Hosts** – The number of hosts discovered or seen by this host on the selected fabric.
- **Number of Fabrics** – The number fabrics identified during discovery.
- **Number of Adapters** – The number of adapters discovered by this host on the selected fabric.
- **Number of Physical Ports** – The number of discovered physical ports on this host that can be managed by this host.

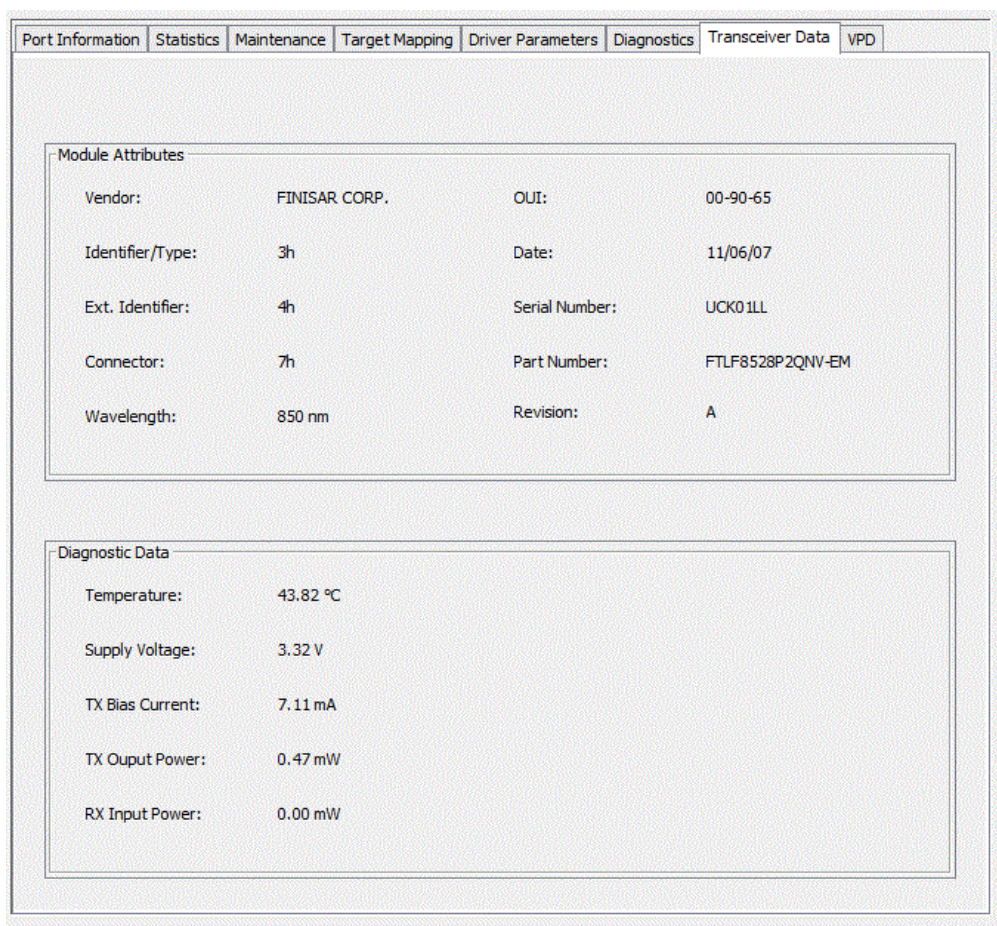
7.9 Viewing Port Transceiver Information

When you select a port from the discovery-tree (Figure 4), the **Transceiver Data** tab (Figure 31) enables you to view transceiver information, such as vendor name, serial number, and part number. If the adapter does not support some or all of the transceiver data, the fields display **N/A**.

To view transceiver information, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4), select the port whose transceiver information you want to view.
3. Select the **Transceiver Data** tab (Figure 31).

Figure 31: Transceiver Data Tab



The following **Transceiver Data** tab fields are displayed:

- **Module Attributes** area:
 - **Vendor** – The name of the vendor.
 - **Identifier/Type** – The identifier value that specifies the physical device described by the serial information.
 - **Ext. Identifier** – Additional information about the transceiver.
 - **Connector** – The external optical or electrical cable connector provided as the media interface.
 - **Wavelength** – The nominal transmitter output wavelength at room temperature.

- **Speed** – The speed, or speeds, at which the selected port can run.
- **OUI** – The vendor’s OUI, also known as the IEEE Company Identifier for the vendor.
- **Date** – The vendor’s date code in MM/DD/YY format.
- **Serial Number** – The serial number provided by the vendor.
- **Part Number** – The part number provided by the SFP vendor.
- **Revision** – The vendor revision level.
- **Diagnostic Data** area:
 - **Temperature** – The internally measured module temperature.
 - **Supply Voltage** – The internally measured supply voltage in the transceiver.
 - **TX Bias Current** – The internally measured transmitted bias current.
 - **TX Output Power** – The measured transmitted output power.
 - **RX Input Power** – The measured received input power.

7.10 Viewing VPD Information

The **VPD** tab (Figure 32) displays vital product data (if available) for the selected adapter port, such as the product name, part number, and serial number.

To view VPD information, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree, select the FC function whose VPD information you want to view.
3. Select the **VPD** tab.

Figure 32: VPD Tab

Item Name	Item Value
Product Name	Emulex LightPulse LPe16002-M6-D 2-Port 16Gb Fibre Channel Adapter, F...
PN (Part Number)	LPe16002-M6-D
SN (Serial Number)	99999999
MN (ManufactureID)	1028
VS	DSV1028VPDR.VER1.0
VP	NPY2
VM	PMT7
VV	NMVEMULEX CORPORATION
VT	DTINIC
VO	99999999
V1	Emulex LightPulse LPe16002-M6-D 2-Port 16Gb Fibre Channel Adapter
V2	LPe16002-M6-D

The following VPD fields are displayed:

- **Product Name** – Product information about the selected FC function.
- **PN (Part Number)** – The adapter's part number.
- **SN (Serial Number)** – The adapter's serial number.
- **MN (Manufacture ID)** – The manufacturer's identification number.
- **Vx** – Vendor unique data. **V** indicates a vendor-specific field. An adapter can have none, one, or more of these fields defined. Valid values for this field are VS, VP, VM, VV, VT, VO (the letter O, not the number zero) and Vx (where x is a number).

NOTE: Some adapters might show additional VPD information, such as EC (EC level), MN (Manufacturer ID), and XY data. Data in the **XY** field is a vendor-specific hexadecimal dump.

7.11 Viewing Maintenance Information

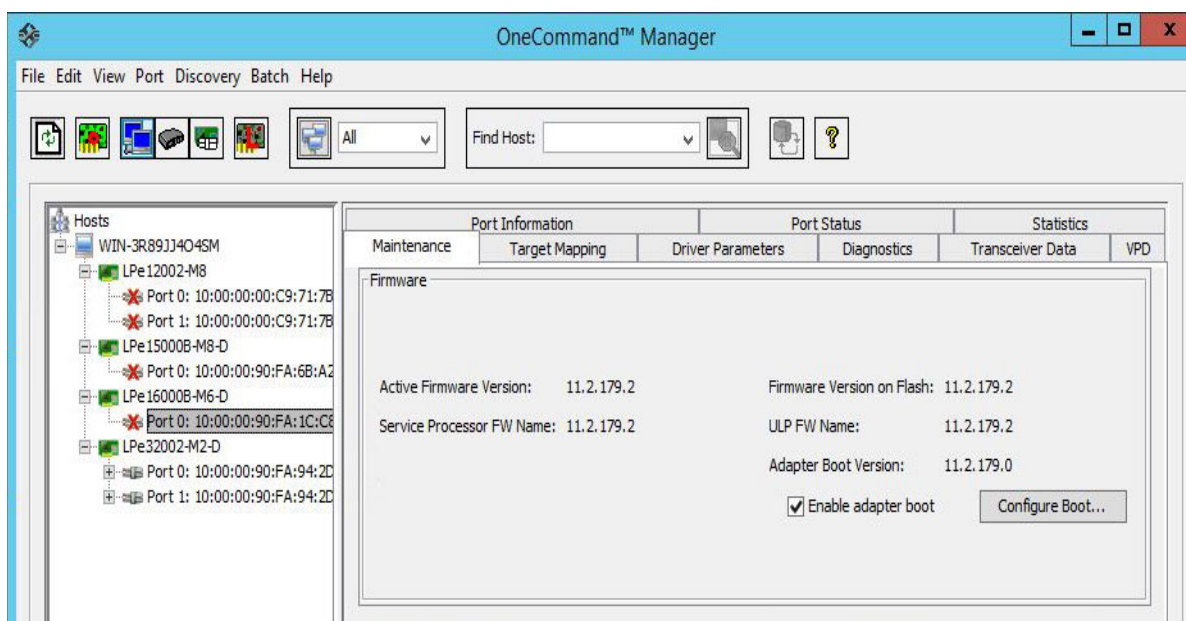
NOTE: This option is not available in read-only mode.

Use the **Maintenance** tab (Figure 33) to view firmware information and update adapter firmware. You can also configure boot from SAN and change WWPN and WWNN information for the selected adapter port.

To view firmware information, perform these steps:

1. Select the **Host** or **Fabric** view.
2. Select a port in the discovery-tree.
3. Select the **Maintenance** tab (Figure 33).

Figure 33: Maintenance Tab



The following **Maintenance** tab fields are displayed:

- **Firmware** area:
 - **Active Firmware Version** – The Emulex firmware version number for this port.
 - **Service Processor FW Name** – The Emulex firmware name for this port.
 - **Firmware Version on Flash** – The flash firmware version currently being used by the adapter.
 - **ULP FW Name** – The firmware version running on the ULP processors within the ASIC.
 - **Adapter Boot Version** – Displays one of the following:
 - The selected adapter port's boot code version if the boot code is present.
 - **Disabled** if the boot code is disabled.
 - **Not Present** if the boot code is not loaded. If the boot code is not loaded, **Enable adapter boot** is not visible, and you cannot configure the selected port to boot from SAN.
 - **Enable adapter boot** – Select this option if you want the port to load and execute boot code during system startup.
 - Click **Configure Boot** to configure boot from SAN (not available in read-only mode).

NOTE: Enabling adapter boot only causes the port to load the boot code and run it during system startup. It does not mean that the port boots from SAN. To boot from SAN, the boot type must be enabled. Enable this in the **Boot from SAN configuration** window for each boot type. (See [Section 7.17.5, Configuring Boot from SAN.](#))

■ **WWN Management** area:

NOTE: The WWN Management area is disabled when FA-PWWN is enabled on an adapter port.

– **Current:**

- **WWPN** – The World Wide Port Name for the selected port.
- **WWNN** – The World Wide Node Name for the selected port.

– **Pending Changes:**

- **WWPN** – Works with the **Change WWN** button. It displays the WWPN that you assigned for the selected port, but the system must be rebooted for these changes to take effect and appear under the Current listing. See [Section 7.16, Changing the WWPN and WWNN](#), for more information.
- **WWNN** – Works with the **Change WWN** button. It displays the WWNN that you assigned for the selected port, but the system must be rebooted for these changes to take effect and appear under the Current listing. See [Section 7.16, Changing the WWPN and WWNN](#), for more information.

For LPe12000-series adapters, the tab includes a **Download Firmware** button. For instructions on updating firmware on a port of an LPe12000-series adapter, see [Chapter 8, Updating Adapter Firmware](#).

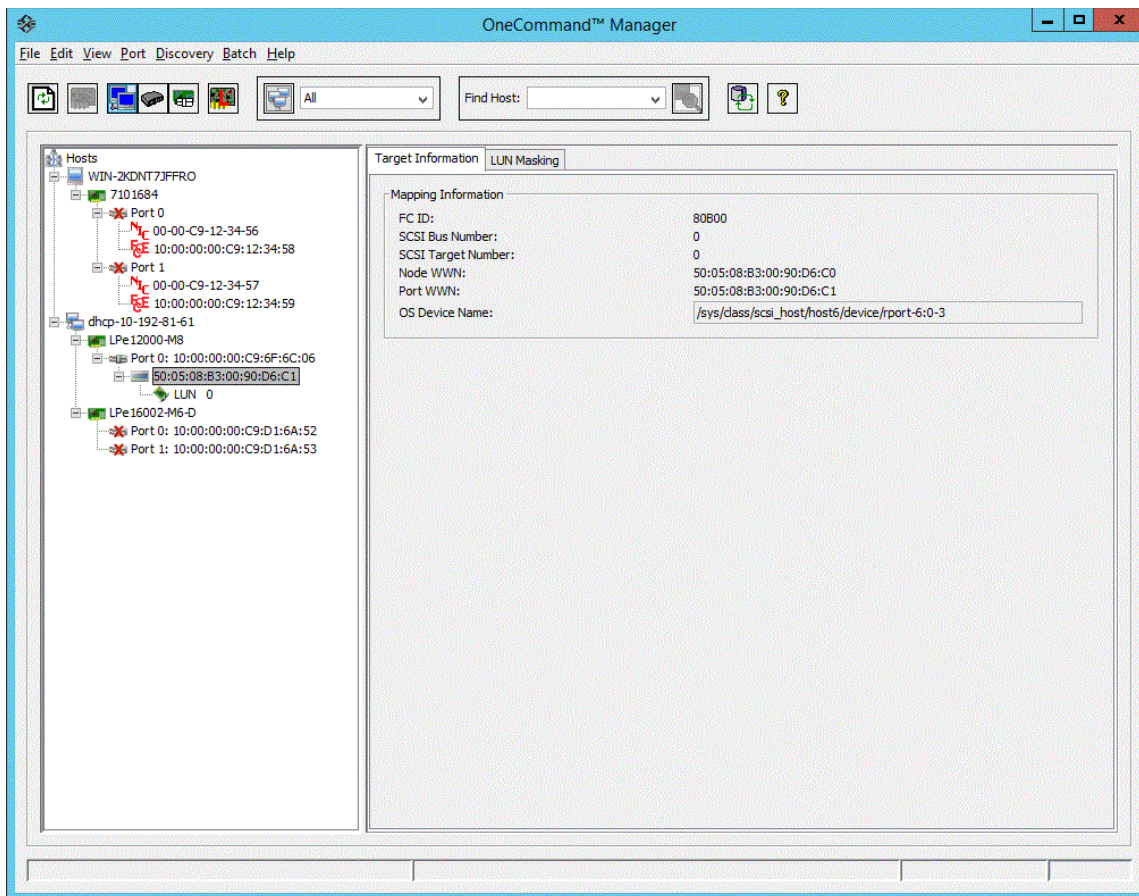
7.12 Viewing Target Information

When you select a target associated with an adapter from the discovery-tree ([Figure 4](#)), the **Target Information** tab ([Figure 34](#)) displays information associated with that target.

To view target information, perform these steps:

1. Select the **Host**, **Fabric**, or **Virtual Port** view.
2. In the discovery-tree ([Figure 4](#)), select the target whose information you want to view. The **Target Information** tab appears ([Figure 34](#)).

Figure 34: Target Information Tab



The following **Target Information** tab fields are displayed:

- **Mapping Information** area:
 - **FC ID** – The FC ID for the target; assigned automatically in the firmware.
 - **SCSI Bus Number** – The SCSI bus number to which the target is mapped.
 - **SCSI Target Number** – The target's identifier on the SCSI bus.
 - **Node WWN** – A unique 64-bit number, in hexadecimal, for the target (N_PORT or NL_PORT).
 - **Port WWN** – A unique 64-bit number, in hexadecimal, for the fabric (F_PORT or FL_PORT).
 - **OS Device Name** – The operating system device name.

7.13 Viewing LUN Information

When you select a LUN associated with a target from the discovery-tree (Figure 35), the **LUN** tab displays information associated with that LUN.

NOTE: The **Refresh LUNs** button on the toolbar refreshes only the LUN list for the currently selected target.

NOTE: On Linux systems, to make LUNs that are newly added to a storage array appear on the host, the following script must run from the command shell:

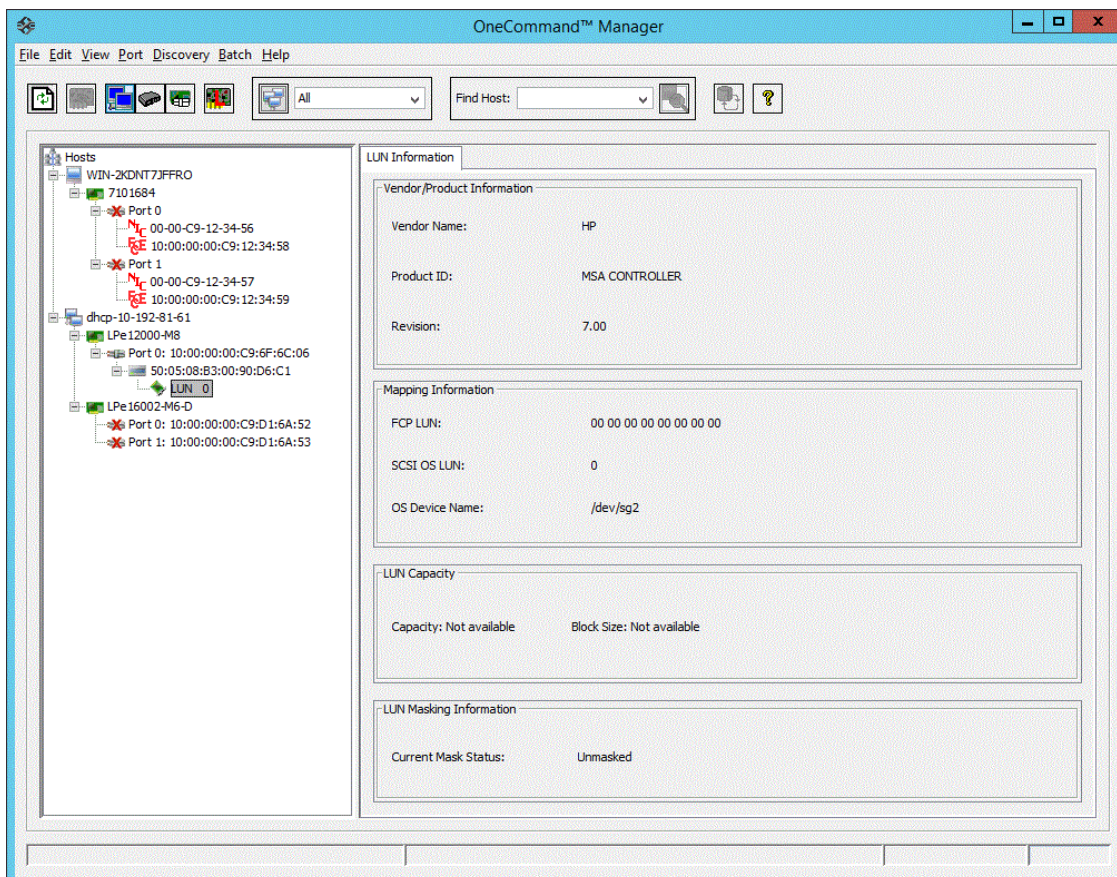
```
/usr/sbin/lpfc/lun_scan all
```

This script prevents you from having to reboot. If the host machine is rebooted after the LUN is added to the target array, you do not need to run the script.

To view the LUN information, perform these steps:

1. Select the **Host**, **Fabric**, or **Virtual Port** view.
2. From the discovery-tree, select a LUN. The **LUN Information** tab appears (Figure 35).

Figure 35: LUN Information Tab



The following **LUN Information** tab fields are displayed:

- **Vendor/Product Information** area:
 - **Vendor Name** – The name of the vendor of the LUN.
 - **Product ID** – The vendor-specific ID for the LUN.
 - **Revision** – The vendor-specific revision number for the LUN.
- **Mapping Information** area:
 - **FCP LUN** – The FC identifier used by the adapter to map to the SCSI OS LUN.
 - **SCSI OS LUN** – The SCSI identifier used by the operating system to map to the specific LUN.
 - **OS Device Name** – The name assigned by the operating system to the LUN.
- **LUN Capacity** area:

NOTE: LUN capacity information is provided only if the LUN is a mass-storage (disk) device. Other devices, such as tapes and scanners, do not display capacity.

- **Capacity** – The capacity of the LUN, in MB.
- **Block Size** – The length of a logical unit block, in bytes.
- **LUN Masking Information** area:
 - **Current Mask Status** – Possible states are **Masked** or **Unmasked**. See [Section 7.15, Masking and Unmasking LUNs \(Windows\)](#), for more information on LUN masking.

7.14 Viewing Target Mapping

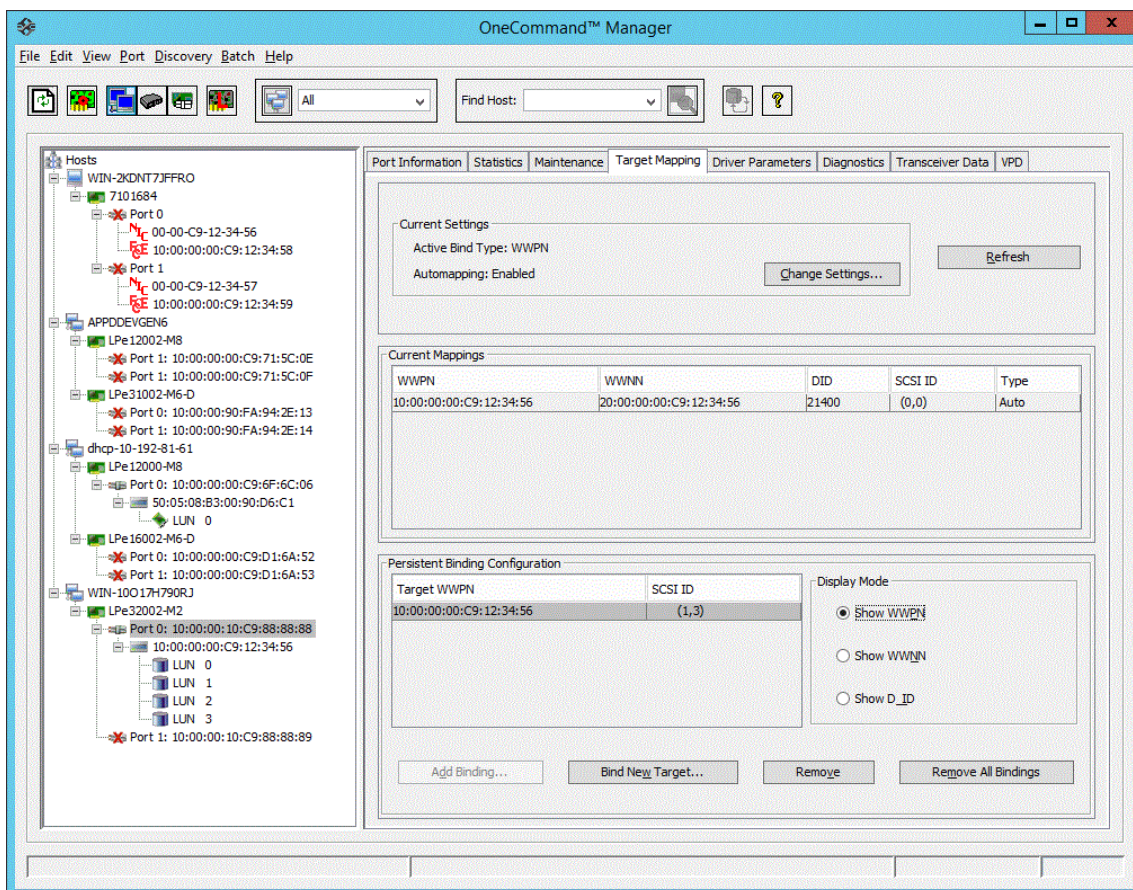
The **Target Mapping** tab ([Figure 36](#)) enables you to view current target mapping and to set up persistent binding.

NOTE: Persistent binding is not supported on Solaris systems.

To view target mapping, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree, select the FC function whose target mapping information you want to view.
3. Select the **Target Mapping** tab ([Figure 36](#)).

Figure 36: Target Mapping Tab



The following **Target Mapping** tab fields are displayed:

- **Current Settings** area:

NOTE: For Linux and VMware ESXi, this area is N/A.

- **Active Bind Type** – WWPN, WWNN, or a D_ID.
- **Automapping** – The current state of SCSI device automapping: **Enabled** (default) or **Disabled**.

- **Current Mappings** area:

- Lists current mapping information for the selected FC function. Shows the WWPN, WWNN, D_ID, and SCSI ID of the currently mapped device. The type can be **Auto** (automapped target) or **PB** (mapped from persistent binding),

- **Persistent Binding Configuration** area:

This table lists persistent binding information for the selected FC function (not available on VMware ESXi servers being managed through the CIM interface).

NOTE: For Linux, this area is N/A.

For information on changing settings, see [Section 7.14.1.1, Changing Automapping Settings](#).

- **Display Mode** area:

- Select whether you want to display information in the Persistent Binding Configuration table.

For information on adding a binding, see [Section 7.14.1.2, Adding a Persistent Binding](#).

For information on binding a new target, see [Section 7.14.1.3, Binding a Target That Does Not Appear in the Persistent Binding Table](#).

To remove a single binding, select the binding and click **Remove**.

To remove all bindings, click **Remove All Bindings**.

7.14.1 Using Automapping and Persistent Binding (Windows Only)

NOTE: This option is not available in read-only mode.

You can set up persistent binding on remote and local adapters. Global automapping assigns a binding type, target ID, SCSI Bus, and SCSI ID to the device. The binding type, SCSI bus, and SCSI ID can change when the system is rebooted. With persistent binding applied to one of these targets, the WWPN, SCSI bus, and SCSI ID remain the same when the system is rebooted.

The driver refers to the binding information during system boot. When you create a persistent binding, the OneCommand Manager application tries to make that binding dynamic. However, the binding must meet all of the following criteria to be dynamic:

- The SCSI ID (target/bus combination) specified in the binding request must not be mapped to another target. For example, the SCSI ID must not already appear in the Current Mappings table under SCSI ID. If the SCSI ID is already in use, the binding cannot be made dynamic, and a reboot is required.
- The target (WWPN, WWNN, or D_ID) specified in the binding request must not be mapped to a SCSI ID. If the desired target is already mapped, a reboot is required.
- The bind type (WWPN, WWNN, or D_ID) specified in the binding request must match the currently active bind type shown in the Current Settings area of the **Target Mapping** tab. If they do not match, the binding cannot be made active.

7.14.1.1 Changing Automapping Settings

To change automapping settings, perform these steps:

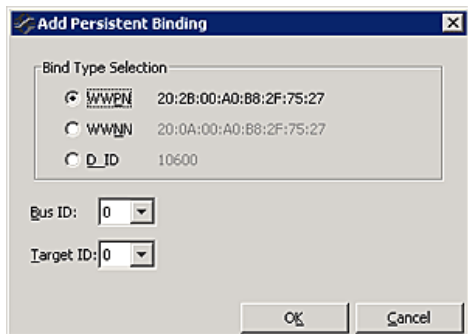
1. Select the **Host** or **Fabric** view.
2. In the discovery-tree, select the FC function you want to set up with persistent binding.
3. Select the **Target Mapping** tab (Figure 36). All targets are displayed.
4. If you want to make changes, click **Change Settings**. The **Mapped Target Settings** dialog appears. You can enable or disable auto-mapping and change the active bind type. Click **OK**.
5. Click **Refresh** to see the changes.
6. Reboot the system for changes to take effect.

7.14.1.2 Adding a Persistent Binding

To add a persistent binding, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4), select the FC function you want to set up with persistent binding.
3. Select the **Target Mapping** tab (Figure 36). All targets are displayed. In the Persistent Binding Configuration table, click the target that you want to bind.
4. Click **Add Binding**. The **Add Persistent Binding** dialog (Figure 37) is displayed.

Figure 37: Add Persistent Binding Dialog



5. Select the bind type that you want to use (**WWPN**, **WWNN**, or **D_ID**).
6. Select the **Bus ID** and **Target ID** that you want to bind, and click **OK**.

NOTE: Automapped targets have entries only in the second column of the Current Mappings table. Persistently bound targets have entries in the second and third columns. In this case, the third column contains the SCSI bus and target numbers that you specified in the **Add Persistent Binding** dialog. This binding takes effect only after the local machine is rebooted.

7.14.1.3 Binding a Target That Does Not Appear in the Persistent Binding Table


NOTE: It is possible to specify a SCSI bus and target that have already been used on behalf of a different FC target. Attempting to bind a target already in the Persistent Binding table on the **Target Mapping** tab results in the following error message:

```
Target already in target list.
```

To bind a target that does not appear in the Persistent Binding table on the **Target Mapping** tab, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4), select the FC function you want to set up with persistent binding.
3. Select the **Target Mapping** tab (Figure 36). All targets are displayed.
4. Click **Bind New Target**. The **Bind New Target** dialog is displayed (Figure 38).

Figure 38: Bind New Target Dialog



5. Select the type of binding that you want to use, and type the WWPN, WWNN, or D_ID that you want to bind to the target.
6. Select the **Bus ID** and **Target ID** that you want to bind, and click **OK**.

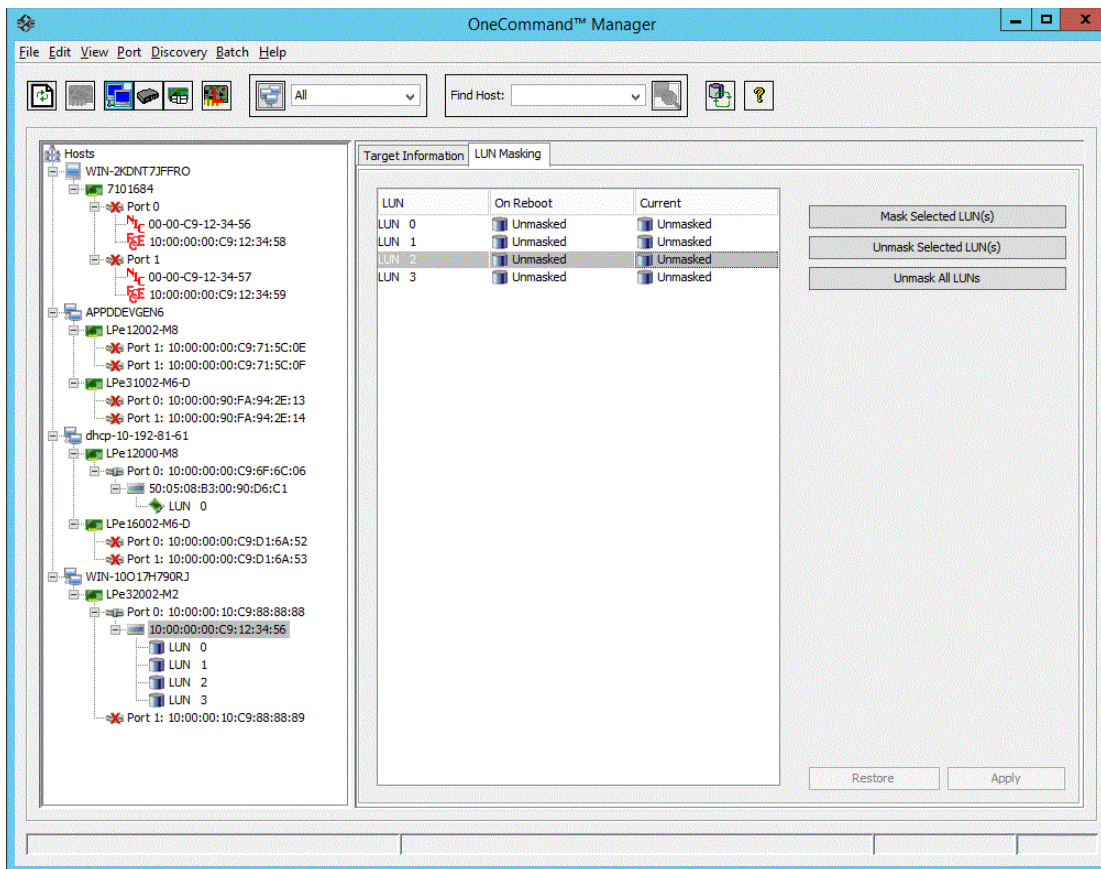
NOTE: A target does not appear on the target list if automapping is disabled and the target is not already persistently bound.

7.15 Masking and Unmasking LUNs (Windows)

LUN masking refers to whether a LUN is visible to the operating system. A masked LUN is not available and is not visible to the operating system. You can use the **LUN Masking** tab (Figure 39) to mask or unmask LUNs at the host level.

NOTE: The **LUN Masking** tab is not shown in the Virtual Port view because LUN masking is not available for virtual ports.

Figure 39: LUN Masking Tab



7.15.1 LUN Masking Conventions and Guidelines

LUN icons in the discovery-tree (Figure 4) reflect the live mask state used by the driver. Green LUN icons indicate unmasked LUNs. Gray LUN icons indicate masked LUNs. Red text indicates that a LUN mask has been changed, but not applied (saved).

The following **LUN Masking** tab information is displayed:

- **LUN** – The FC LUN number.
- **On Reboot** – The **On Reboot** column shows the mask configuration currently saved to the configuration file on disk (Solaris) or to the registry (Windows). Usually, for a specific LUN, the states reported in the **On Reboot** and **Current** columns are identical. However, there are times when these do not match. For example, the `hbacmd` utility can be used to change only the current mask state for a LUN and not touch the **On Reboot** mask state contained in the configuration file.
- **Current** – The **Current** column displays the live mask state currently in use by the driver. When you first see the **LUN Masking** tab, the mask states displayed in the **Current** column are identical to the mask states for the corresponding LUNs in the discovery-tree (Figure 4).

To change the mask status of a LUN, perform these steps:

1. Select the **Host** view.
2. From the discovery-tree (Figure 4), select the target whose LUN masking state that you want to change. A set of LUNs appears below the selected target.

3. Select the **LUN Masking** tab (Figure 39). This tab contains a list of the same set of LUNs that appear below the FC target in the discovery-tree (Figure 4).
4. In the LUN list of the **LUN Masking** tab, select one or more LUNs. The **Mask Selected LUN(s)**, **Unmask Selected LUN(s)**, **Unmask All LUNs**, **Restore**, and **Apply** buttons become active as appropriate. For example, if the LUN is currently unmasked, only the **Mask Selected LUN(s)** button is active.
5. Change the mask status: click **Mask Selected LUN(s)**, **Unmask Selected LUN(s)** or **Unmask All LUNs** as appropriate. Mask status changes appear in red text.

NOTE: To return all mask settings to their status before you started this procedure, click **Restore** before you click **Apply**. If you click **Apply**, changes cannot be cancelled by clicking **Restore**. To unmask all LUNs, click **Unmask All LUNs**. This button is always active.

6. Click **Apply** to commit the changes. An informational message confirms that the mask status has changed and the red text changes to black.

7.16 Changing the WWPN and WWNN

The **Maintenance** tab (Figure 40) enables you to change the WWPN and the WWNN of a selected FC function. For example, you can use an installed adapter as a standby in case another installed adapter fails. By changing the standby adapter's WWPN or WWNN, it can assume the identity and configuration (for example, driver parameters, persistent binding settings, and so on) of the failed adapter.

NOTE: You cannot change WWPN and WWNN when FA-PWWN is enabled on the adapter port.

Three options exist for referencing WWNs:

- **Factory Default WWN** – The value that shipped from the factory. This value cannot be changed.
- **Non-volatile WWN** – A value that is saved in the adapter's non-volatile flash memory, and that survives a reboot or a power outage.
- **Volatile WWN** – A temporary value that is saved in volatile memory on the flash. If volatile WWNs are set, they are used instead of the non-volatile WWNs.

NOTE: Volatile WWN changes require a warm system reboot to take effect. Volatile WWN changes are lost on systems that power-cycle the adapters during the reboot.

To avoid address conflicts, do not assign a WWPN with the same WWPN as another FC function on your SAN. The OneCommand Manager application checks the WWPN you specify against all the other detected WWPNs and, if a duplicate is found, an error is displayed and the WWPN is not changed.

CAUTION! Changing volatile WWNs takes the selected adapter offline. Make sure that this adapter is not controlling a boot device and that all I/O activity on this adapter is stopped before proceeding, or unexpected behavior or data loss can result.

In an environment where preboot management exists, a WWPN or WWNN modified by the OneCommand Manager application can be overridden by preboot management, such as BOFM and industry-standard CLP.

For example, in an environment with CLP or BOFM, the OneCommand Manager application modifies the WWNN or WWPN. The OneCommand Manager application requires a reboot to complete the change. After reboot, the CLP string is sent during system boot and rewrites the WWNN or WWPN, or EFIBoot finds the BOFM and uses the default WWNN or WWPN by the command from the BOFM.

To change an FC function's WWPN or WWNN, perform these steps:


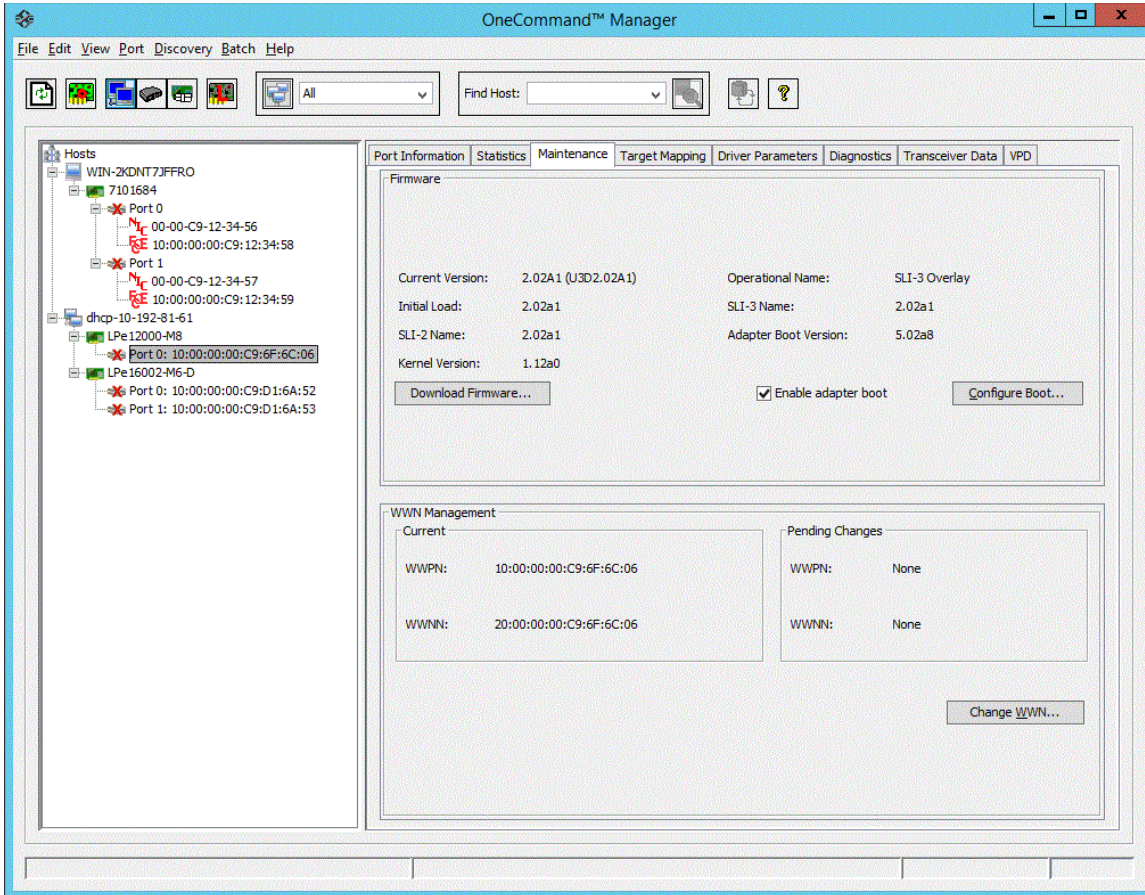

1. Perform one of the following tasks:
 - From the **View** menu, select **Group Adapters by Host Name**.
 - From the toolbar, click the  **Group Adapters by Host Name** button.
 - From the **Host Grouping** menu, select **Group Adapter by Fabric Names**.
2. In the discovery-tree, select the FC function that you want to change.
3. Select the **Maintenance** tab ([Figure 40](#)).

Figure 40: Maintenance Tab



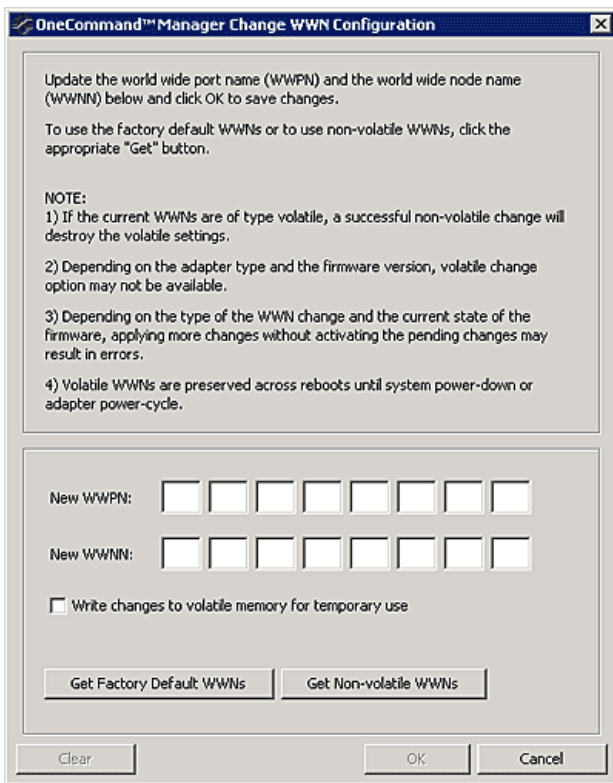
4. Click **Change WWN**. The popup in [Figure 41](#) appears.

Figure 41: Warning About Changing WWN



5. Click **Yes**. The **OneCommand Manager Change WWN Configuration** dialog appears (Figure 42).

Figure 42: OneCommand Manager Change WWN Configuration Dialog



6. Perform one of the following tasks:
 - Enter a new WWPN and WWNN.
 - Click **Get Factory Default WWNs** to load the settings that were assigned to the FC function when the adapter was manufactured to the **New WWPN** and **New WWNN** fields. These values can then be modified if desired and saved as volatile or non-volatile WWNs.
 - Click **Get Non-volatile WWNs** to load the current non-volatile WWN settings to the **New WWPN** and **New WWNN** fields. These values can then be modified if desired and saved to volatile or non-volatile memory. These values can then be modified if desired and saved as volatile or non-volatile WWNs.
7. Select **Write changes to volatile memory for temporary use** to save the new WWPN and new WWNN settings as volatile WWNs. If **Write changes to volatile memory for temporary use** is cleared, the new WWPN and new WWNN settings are saved as non-volatile WWNs.

NOTE: If the adapter or firmware does not support volatile WWNs, the **Write changes to volatile memory for temporary use** check box is dimmed.

8. Click **OK**. After checking for a duplicate WWPN, the new WWPN and new WWNN values are saved for volatile or non-volatile use. The new WWPN and WWNN appear in the Pending Changes section in the WWN Management area of the **Maintenance** tab until the system is rebooted.
9. Reboot the system for the changes to take effect. After rebooting, the changes are applied and appear in the Current area of the **Maintenance** dialog.

7.16.1 Changing Port Names

NOTE: This option is not available in read-only mode.

The OneCommand Manager application allows you to change the adapter port names in the discovery-tree.

For example, you might want to identify a particular FC function with the role it supports, such as a tape drive, scanner, or some other device. Use any characters you want for names, and names can be up to 255 characters in length. You can also revert to the adapter's default name.

NOTE: Although you can change the FC function's displayed name from the default WWPN, the change occurs in the discovery-tree (Figure 4) only. The function's WWPN is still active; it is replaced for display purposes with the name you enter. For example, the **Port WWN** field of the **Port Information** tab is not changed. Also, any changes you make to the names in your discovery-tree are seen only by you; users running the OneCommand Manager application on another host do not see your name changes.

To change the name of an FC function, perform these steps:

1. From the discovery-tree (Figure 4), select the port that you want to change by performing one of the following tasks:
 - Select **Edit Name** from the **Port** menu.
 - From the discovery-tree, right-click the port that you want to change and select **Change Name**.
2. Edit the name in the discovery-tree.

To use the FC function's default name, perform these steps:

1. From the discovery-tree (Figure 4), select the FC function that you want to change.
2. Perform one of the following tasks:
 - Select **Use Default Name** from the **Port** menu.
 - From the discovery-tree, right-click the port that you want to change and select **Restore Default Name**.

7.16.2 Resetting the FC Functions

You can reset remote and local functions.

CAUTION! Do not reset functions while copying or writing files. This action could result in data loss.

To reset the FC function, perform these steps:


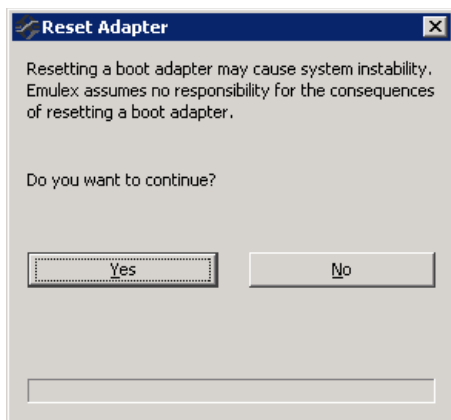
1. In the discovery-tree (Figure 4), select the FC function that you want to reset.
2. Perform one of the following tasks:
 - From the **Port** menu, select **Reset Port**.
 - From the toolbar, click the  **Reset** button.The **Reset Adapter** popup appears (Figure 43).

Figure 43: Reset Adapter Popup



3. Click **Yes** to perform the reset.

The reset can require several seconds to complete. During the reset, the status bar shows **Reset in progress**. When the reset is finished, the status bar shows **Reset Completed**.

7.17 Configuring the Driver Parameters

NOTE: This option is not available in read-only mode.

The OneCommand Manager application displays available driver parameters along with their defaults and maximum and minimum settings. A description of the selected parameter is also provided. This section contains information you must be aware of when working with driver parameters. For a more detailed description of specific driver parameters, refer to the appropriate Emulex driver user guide.

NOTE: In Solaris and Linux, you can also specify parameters when loading the driver manually. Refer to the appropriate driver user guide for instructions.

7.17.1 Activation Requirements

A parameter has one of the following activation requirements:

- **Dynamic** – Takes effect while the system is running.
- **Reset** – Requires a reset from the utility before the change takes effect.
- **Reboot** – Requires 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.

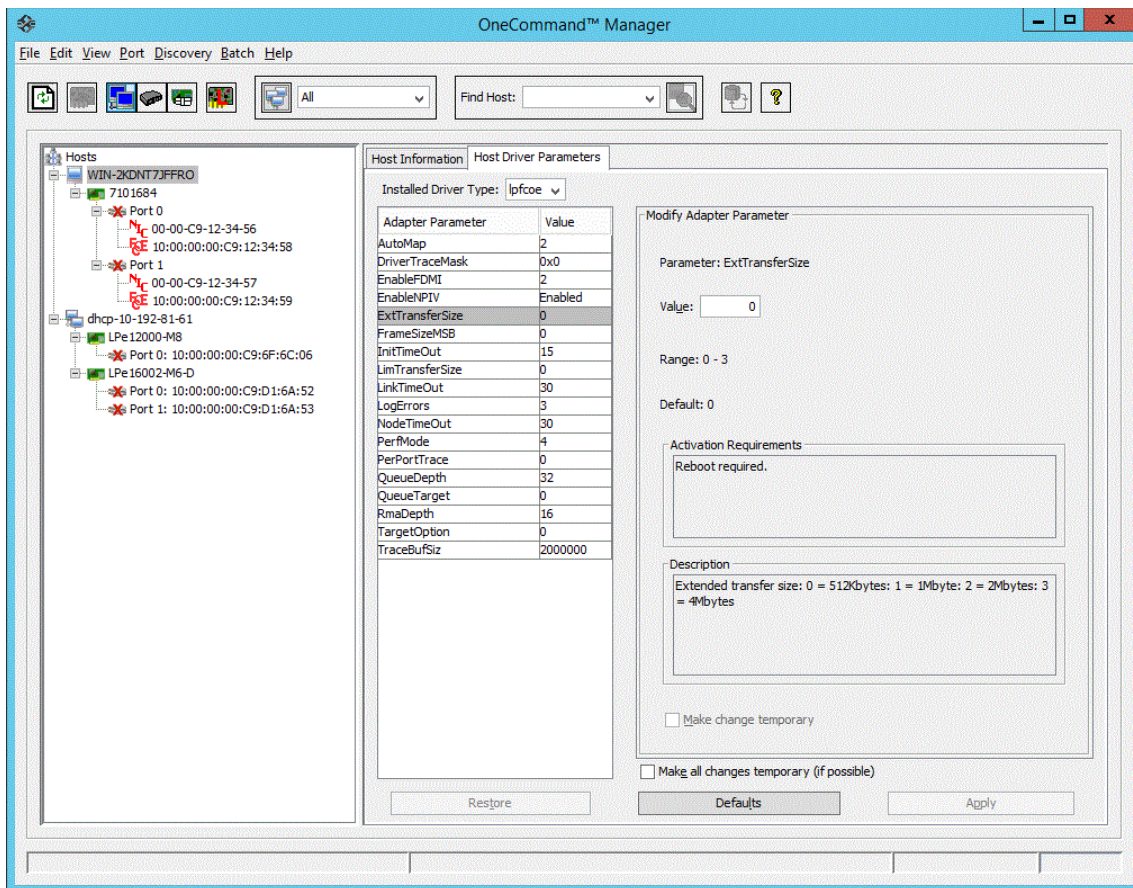
7.17.2 Host Driver Parameters Tab

The **Host Driver Parameters** tab (Figure 44) enables you to view and edit the adapter driver parameter settings contained in a specific host. The host driver parameters are global values, and apply to all adapters in that host unless they are overridden by parameters assigned to a specific adapter using the adapter **Driver Parameters** tab. For each parameter, the tab shows the current value, the range of acceptable values, the default value, and whether the parameter is dynamic. A dynamic parameter allows the change to take effect without resetting the adapter or rebooting the system.

For information on changing parameters for a single adapter, see [Section 7.17.3, Setting the Driver Parameters](#). For information on changing parameters for the host, see [Section 7.17.3.5, Setting the Driver Parameters for All FC Functions in a Host](#).

NOTE: Setting a parameter for a specific port overrides all host parameters for the port.

Figure 44: Host Driver Parameters Tab



The following **Host Driver Parameters** tab fields are displayed:

- **Installed Driver Type** – The current drivers installed on this host. If more than one driver type is installed, the **Installed Driver Types** drop-down list shows a list of all driver types that are installed on the adapters in the host and enables you to select the particular driver type to configure.
- **Adapter Parameter table** – A list of adapter driver parameters for the selected driver type and their current values.
- **Modify Adapter Parameter** area:
 - Adapter-specific information is displayed in this area, which includes value, range, default, activation requirements, and description.

7.17.3 Setting the Driver Parameters

The **Driver Parameters** tab for FC functions and hosts enables you to modify driver parameters for a specific FC function or all FC functions in a host.

For example, if you select a host in the discovery-tree ([Figure 4](#)), you can globally change the parameters for all FC functions in that host. If you select an FC function in the discovery-tree, you can change parameters for only that FC function.

For each parameter, the **Driver Parameters** tab shows the current value, the range of acceptable values, the default value, and the activation requirement. You can also restore parameters to their default settings.

You can apply driver parameters for one FC function to other FC functions in the system using the **Driver Parameters** tab, thereby simplifying multiple adapter configuration. See [Section 7.17.4, Creating a Batch Mode Driver Parameters File](#), for more information.

NOTE: The Linux 2.6 kernel only supports setting some of the driver parameters for individual FC functions. Some driver parameters must be applied to all FC functions contained in the host. Refer to the *Emulex Driver for Linux User Guide* for more information.

7.17.3.1 Setting the Driver Parameters for a Single FC Function

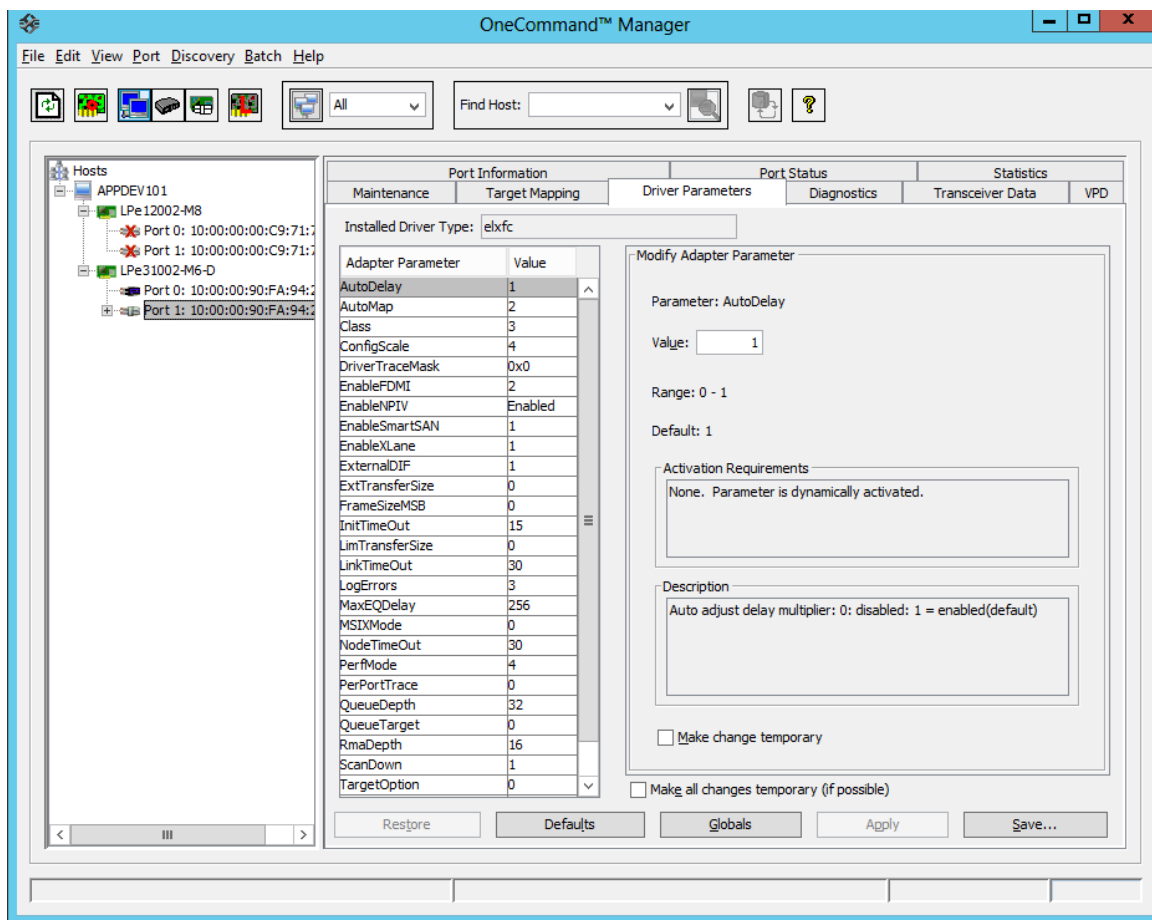
NOTE: Setting a parameter for a specific port overrides all host parameters for the port.

To change the driver parameters for a single FC function, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree ([Figure 4](#)), select the FC function that you want to change.
3. Select the **Driver Parameters** tab ([Figure 45](#)). The parameter values for the selected FC function are displayed.

NOTE: The LinkSpeed (Windows) or link-speed (Linux/Solaris) driver parameters are not shown if the adapter supports forced link speed. The link speed is configured using the **Firmware Parameters** tab. See [Section 7.3, Viewing Firmware Parameters](#), for more information.

Figure 45: Driver Parameters Tab (Adapter Selected)



4. Click the parameter that you want to change. A description of the parameter appears on the right side of the tab.
5. Enter a new value in the **Value** field in the same hexadecimal or decimal format as the current value or select a value from the list. If the current value is in hexadecimal format, it is prefaced by 0x (for example, 0x2d). You can enter a new hexadecimal value without the 0x. For example, if you enter ff10, this value is interpreted and displayed as 0xff10.
6. If you want the change to be temporary (causing the parameter to revert to its last permanent setting when the system is rebooted), select the **Make change temporary** check box. This option is available only for dynamic parameters.
7. If you are making changes to multiple parameters, and you want all the changes to be temporary, select **Make all changes temporary**. This setting overrides the setting of **Make change temporary**. Only dynamic parameters can be made temporary.
8. Click **Apply**.

7.17.3.2 Restoring All Parameters to Their Earlier Values

If you changed parameters, but did not click **Apply** in the **Driver Parameters** tab (Figure 45) and you want to restore the parameters to their last saved values, click **Restore**.

7.17.3.3 Resetting All Default Values

To reset all parameter values to their default (factory) values, click **Defaults** in the **Driver Parameters** tab (Figure 45).

7.17.3.4 Setting an Adapter Parameter Value to the Host Adapter Parameter Value

To set an adapter parameter value to the corresponding host parameter value, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4), select the adapter port.
3. Select the **Driver Parameters** tab (Figure 45).
4. Click **Globals**. All parameter values are now the same as the global, or host, values.
5. To apply the global values, click **Apply**.

7.17.3.5 Setting the Driver Parameters for All FC Functions in a Host

To change the driver parameters for all FC functions installed in a host, perform these steps:


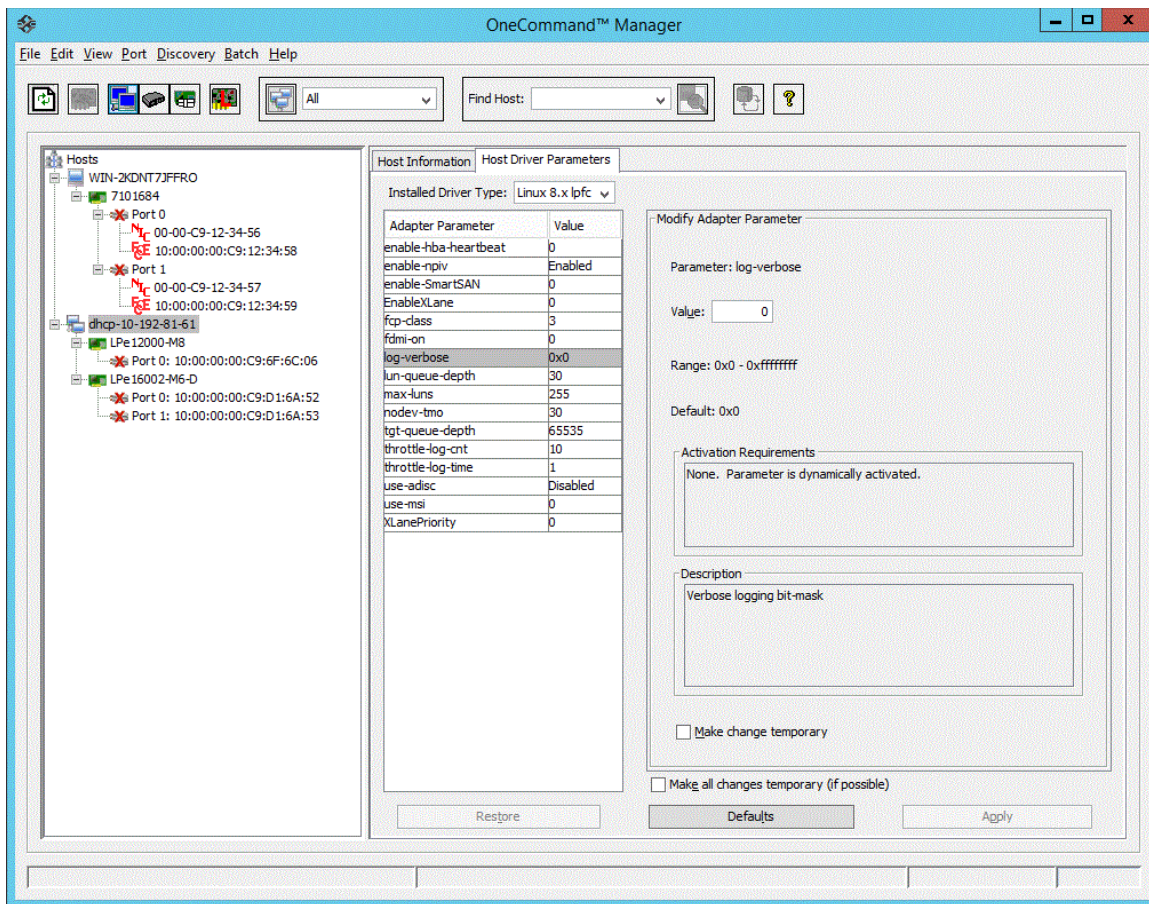
1. Perform one of the following tasks:
 - From the View menu, select **Group Adapters by Host Name**.
 - From the toolbar, click the  **Group Adapters by Host Name** button.
2. In the discovery-tree, click the host whose adapter driver parameters you want to change.
3. Select the **Host Driver Parameters** tab (Figure 46). If adapters with different driver types are installed, the **Installed Driver Types** menu shows a list of all driver types and driver versions that are installed. Select the driver whose parameters you want to change. This menu does not appear if all the adapters are using the same driver.
4. Click the parameter that you want to change. A description of the parameter appears on the right side of the tab.

Figure 46: Host Driver Parameters Tab – Host Selected



5. Enter a new value in the **Value** field in decimal or hexadecimal format, depending on how the current value is presented. If the value is in hexadecimal format, it is prefaced by 0x (for example, 0x2d).
6. To make a change temporary (the parameter reverts to its last permanent setting when the system is rebooted), select **Make changes temporary**. This option is available only for dynamic parameters.
7. To make changes to multiple parameters, select **Make all changes temporary (if possible)**. Only dynamic parameters can be made temporary.
8. Click **Apply**.

7.17.3.6 Changing Non-Dynamic Parameter Values (Linux)

To change non-dynamic parameter values for Linux, perform these steps:

1. Navigate to the `/usr/sbin/ocmanager` directory, and run the scripts to stop the OneCommand Manager application processes. Type the following command:

```
./stop_ocmanager
```
2. Stop all I/O to FC attached devices.
3. Unload the FC driver. Type the following command:

```
modprobe -r lpfc
```

4. Reload the driver. Type the following command:

```
modprobe lpfc
```

5. Start the `elxhbamgr` service (remote service). Type the following command:

```
./start_ocmanager
```

The OneCommand Manager application discovery service starts automatically when you start the application.

NOTE: For changes to persist after a reboot, you must create a new ramdisk image. Refer to the *Emulex Drivers for Linux User Guide* for more information.

7.17.4 Creating a Batch Mode Driver Parameters File

NOTE: This option is not available in read-only mode.

You can apply driver parameters for one FC function to other FC functions in the system using the **Driver Parameters** tab. When you save the driver parameters for an adapter, you create a `.dpv` file that contains parameters for that adapter. After you create the `.dpv` file, the OneCommand Manager application enables you to assign the `.dpv` file parameters to multiple adapters in the system.

To create the `.dpv` file, perform these steps:

1. Select the **Host** or **Fabric** view.
2. Select the FC function whose parameters you want to apply to other FC functions from the discovery-tree ([Figure 4](#)).
3. Select the **Driver Parameters** tab ([Figure 45](#)).
4. Set the driver parameters.
5. After you define the parameters for the selected adapter, click **Apply**.
6. Click **Save**. The **Save Driver Parameters** dialog appears ([Figure 47](#)). You can save the file to a different directory or change its name.

Each definition is saved in a comma-delimited file with the following format:

```
<parameter-name>=<parameter-value>
```

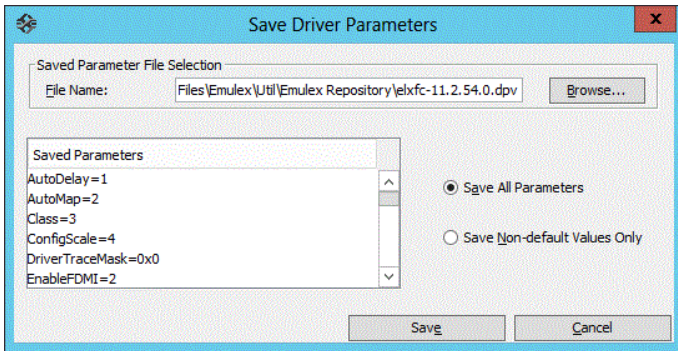
The file is saved in the Emulex repository directory.

- In Windows: `\Program Files\Emulex\Util\Emulex Repository`
- In Linux: `/usr/sbin/ocmanager/RMRepository`
- In VMware ESXi: `/tmp/RMRepository`
- In Solaris: `/opt/ELXocm/RMRepository`

NOTE: Host driver parameters and persistent binding settings cannot be saved.

7. The OneCommand Manager application then uses the Batch Driver Parameter Update function to apply these saved settings to all compatible adapters on the SAN.

Figure 47: Save Driver Parameters Dialog



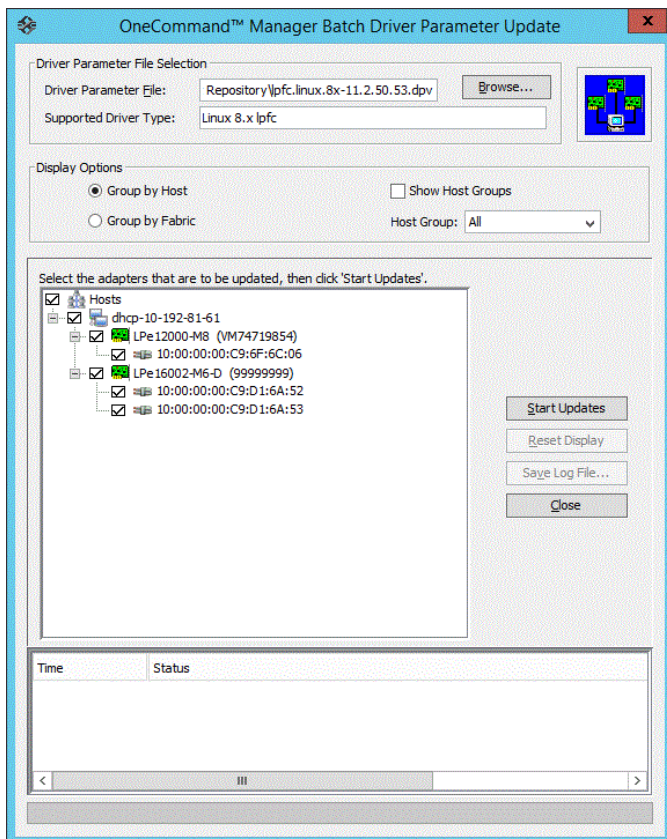
8. Choose whether to save all parameters or only those parameters whose current values differ from their corresponding default values.
A list of the saved parameters and their current values appear in the **Saved Parameters** list.
9. Click **Save**.

7.17.4.1 Assigning Batch Mode Parameters

To assign batch mode parameters to adapters, perform these steps:

1. From the **Batch** menu, select **Update Driver Parameters**. (You do not need to select any discovery-tree [\[Figure 4\]](#) elements at this time.)
2. When the **Batch Driver Parameter Update** dialog appears, click **Browse**.

Figure 48: Batch Driver Parameter Update Dialog



3. The **Batch Driver Parameter Update** dialog appears (Figure 48). Select the file that you want to use and click **OK**. A dialog notifies you that the OneCommand Manager application is searching for compatible adapters.
After compatible FC functions are found, the **Driver Parameter File** field of the **Batch Driver Parameter Update** dialog displays the selected file's path. The **Supported Driver Type** field displays the selected driver. The set of compatible adapters is displayed in the dialog's discovery-tree.
Using the **Display Options** settings, you can choose how adapters are displayed in the discovery-tree. Select **Group by Host** to display adapters in a host-centric view. Select **Group by Fabric** to display hosts in a fabric-centric view with their fabric addresses. The WWPN and host name for each available FC function is displayed under its respective fabric.
You can also display host groups by selecting **Show Host Groups**. To display a particular host group, select that group from the **Host Group** list.
You can select or clear the host, adapter, and FC functions entries in the discovery-tree. Selecting an adapter selects or removes all FC functions on that adapter. Checking a host removes or selects all eligible adapters for that host.
4. Make your selections, and click **Start Updates**. The **Batch Driver Parameter Update** dialog (Figure 48) shows the current status of the update. When the update is finished, a final summary shows the number of FC functions that were successfully processed, and the number of FC functions for which one or more parameter updates failed.
You can click **Save Log File** to save a report of the update. If you change any parameter settings, click **Reset Display** to refresh the dialog.

7.17.5 Configuring Boot from SAN

You can use the OneCommand Manager application to configure a system to boot from an attached LUN. Boot from SAN allows servers on a storage network to boot their operating systems directly from a SAN storage device, typically identified by its WWPN and a LUN located on the device. By extending the server system BIOS, boot from SAN functionality is provided by the BootBIOS contained on an Emulex adapter in the server. When properly configured, the adapter then permanently directs the server to boot from a LUN on the SAN as if it were a local disk.

NOTE: Boot from SAN is not supported through the CIM interface.

7.17.5.1 Boot Types

Using the **Maintenance** tab, you can enable, disable, or configure boot from SAN for x86 BootBIOS, EFIBoot, and OpenBoot (also known as FCode).

NOTE:

- x86 BootBIOS works with the existing BIOS on x64 and x86 systems.
- OpenBoot (FCode) works with the existing system BIOS on Solaris SPARC systems using the SFS driver.
- EFIBoot works with x64-based systems and provides 64-bit system boot capability through the use of the EFI Shell.
- When x86 settings are changed, the same changes are also made to EFI. Conversely, when EFI settings are changed, the changes are also made to x86.

For LPe12000-series adapters, Emulex provides Universal Boot and Pair Boot code images that contain multiple types of boot code. These images provide multiplatform support for boot from SAN. Universal Boot and Pair Boot transparently determine your system platform type and automatically run the proper boot code image in the adapter. These code images reside in adapter flash memory, allowing easier adapter portability and configuration between servers.

For all other adapters, the firmware image includes all supported types of boot code.

The adapters store the boot configuration data for each of these boot types.

NOTE: x86 and OpenBoot share the same configuration memory space. You cannot configure an adapter for both x86 and OpenBoot at the same time. If you try, a message is displayed, stating that the existing boot type configuration will be overwritten by the new configuration.

NOTE: Boot from SAN configuration does not affect current system operation. The changes take effect only upon reboot if you have configured it correctly.

7.17.5.2 Boot Device Parameters

The boot LUN for all three boot types is in the range of 0 to 255. EFIBoot and OpenBoot (FCode) also support an 8-byte LUN, which you can use instead of the single-byte LUN. You must select which LUN type to configure.

- For OpenBoot, you must also provide a Target ID parameter.
- You must boot the host to configure boot from SAN with the OneCommand Manager application.
- You must work from a running host that supports the OneCommand Manager application. Often, this host has booted from a direct-attached drive. With the OneCommand Manager application, you can configure a direct boot host to boot from a SAN. You can modify an existing boot from SAN configuration or configure boot from SAN on an adapter for installation in another host so it can boot from SAN.
- You must know the boot code type that the adapter has; the OneCommand Manager application cannot detect this information. Without knowing this information, you could configure a boot type but not be able to boot from it because the adapter lacks the correct boot code.

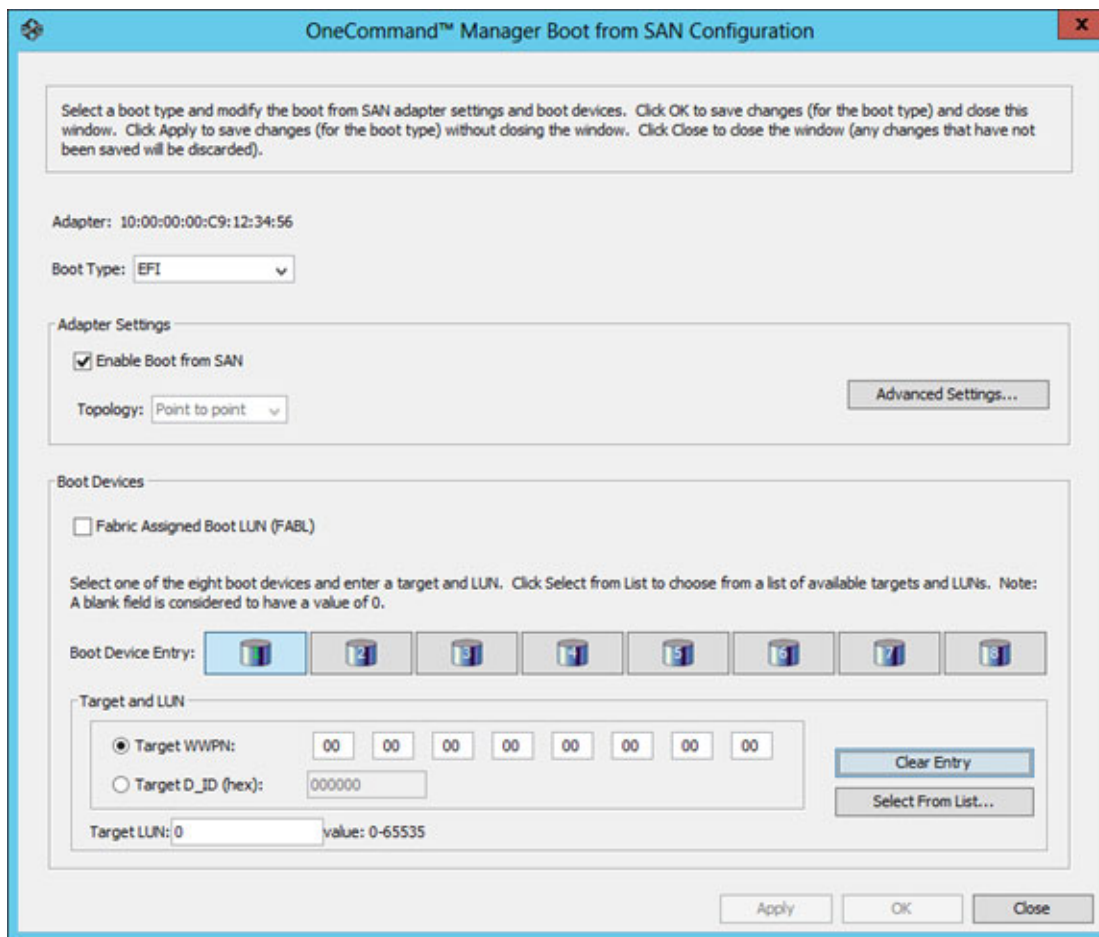
- You must know the boot code type that the system supports; the OneCommand Manager application cannot detect this information. You can configure any boot type, but if the system does not support that type, it cannot boot from SAN.
- One of the following adapter drivers must be installed:
 - Windows: Storport Miniport driver
 - Linux: Emulex driver
 - Solaris: inbox driver
 - VMware: Emulex driver

To configure boot from SAN, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4), click the FC adapter function on which you want to enable boot from SAN.
3. Select the **Maintenance** tab (Figure 33), select **Enable adapter boot** (if available), and click **Configure Boot**. The **Boot from SAN Configuration** dialog appears (Figure 49).

NOTE: The **Configure Boot** button is dimmed if **Enable Adapter Boot** is not selected.

Figure 49: EFIBoot Boot from SAN Configuration Dialog



The **Boot from SAN Configuration** dialog varies for each boot type. Figure 49 depicts the boot from SAN configuration for the EFIBoot type boot.

4. Verify that the **Adapter** field contains the WWPN of the FC function, and ensure that you configure the correct adapter FC function.
5. From the **Boot Type** menu, select **X86**, **EFI**, or **OpenBoot**.

NOTE: x86 and OpenBoot share the same configuration memory space. You cannot configure an adapter for both x86 and OpenBoot at the same time. If you select one of these boot types and the configuration region is configured for the other boot type, a message warns you that making changes overwrites the other boot-type configuration.

If you modified the settings for the current boot type and then change to a new boot type, a message reminds you to save the current settings before changing to the new boot type.

6. Select **Enable Boot from SAN**, and select the topology.

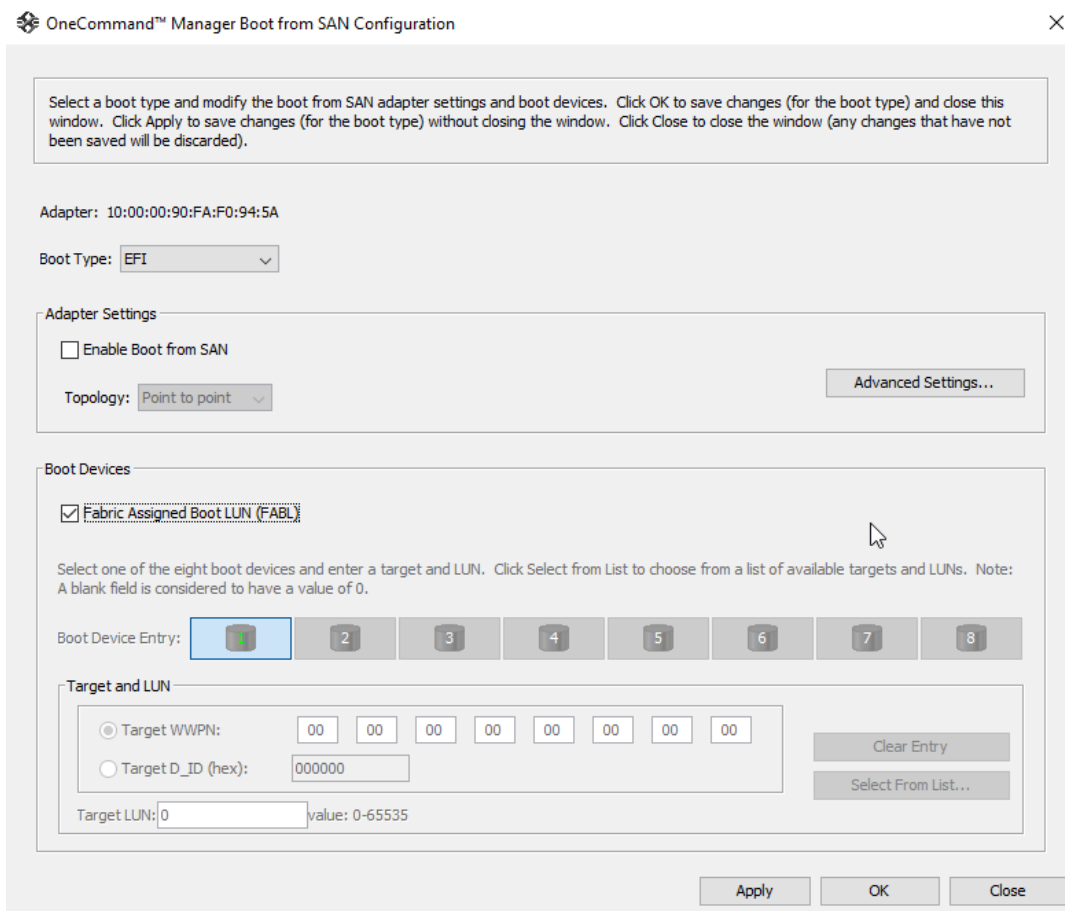
Topology options are:

- **Auto, Loop First** (default)
 - **Auto, Point to Point First**
 - **Loop**
 - **Point to Point**
7. To configure autoscan, spinup delay, and other advanced settings, see [Section 7.18, Configuring Advanced Settings \(Boot from SAN\)](#).
 8. For EFIBoot, you can select Brocade **Fabric Assigned Boot LUN (FABL)**. You are prompted to reboot after enabling FABL. (FABL is not supported on LPe12000-series adapters.)

NOTE: Switch configuration is required when using FABL.

NOTE: Boot LUNs allocated using FABL will override the boot device configuration currently defined for the port ([Figure 50](#)).

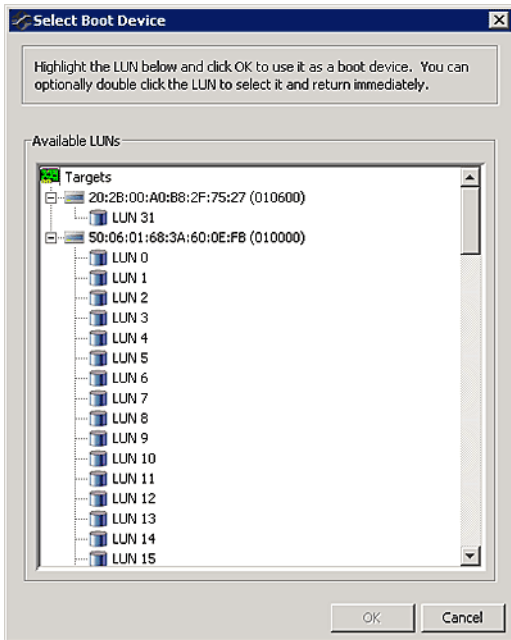
Figure 50: EFIBoot Boot from SAN Configuration Dialog: FABL Enabled



9. For x86 and EFIBoot, select one or more boot devices. For OpenBoot, select only one boot device.
10. Perform one of the following tasks on the **Boot from SAN Configuration** dialog (Figure 49):
 - Select **Target WWPN**, type the desired WWPN, and click **OK**.
 - Select **Target D_ID**, type desired D_ID, and click **OK**.
 - Select **Target LUN**, type the desired LUN, and click **OK**.

For EFIBoot and OpenBoot, enter the LUN value in big endian order (most-significant byte, or big endian first) and enter all 16 characters, including leading zeros, for the Target WWPN. Type in an 8-byte LUN (hexadecimal) Target ID for the LUN.
 - Click **Select from List**, select the target from a list of discovered LUNs (if available), and click **OK** on the **Select Boot Device** dialog (Figure 51). You can manually enter the target and LUN from the **Boot from SAN Configuration** dialog; however, it is easier to select an existing LUN from the **Select Boot Device** dialog (Figure 51.) The OneCommand Manager application attempts to update the boot parameters. If the update is successful, a popup displays a confirmation message. Click **OK**.

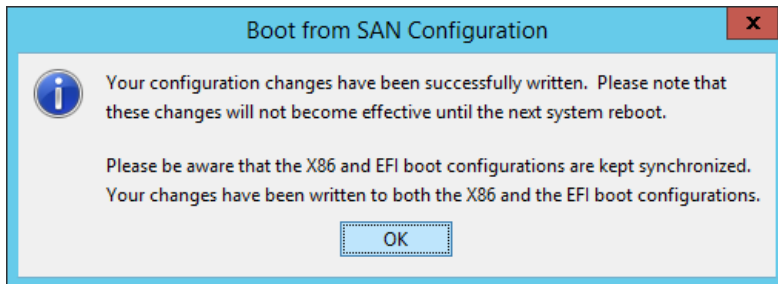
Figure 51: Select Boot Device Window (for x86 or EFIBoot)



11. On the **Boot from SAN Configuration** dialog (Figure 49), click **Apply** to save your changes but leave the dialog open, or click **OK** to apply the changes and close the dialog.

NOTE: Click **Close** to close the **Boot from SAN Configuration** dialog without saving your changes. A message prompts you to discard your changes.

12. If you changed x86 or EFI boot settings, the following popup appears. Click **OK**.



13. Reboot the system for your changes to take effect.

7.18 Configuring Advanced Settings (Boot from SAN)

The OneCommand Manager application provides advanced settings for each boot type. From the **Boot from SAN Configuration** dialog (Figure 49), click **Advanced Settings**. A boot type-specific dialog allows you to enable options, such as spinup delay and autoscan.

If you make changes, you must click **OK** to save the changes and close the dialog. You can click **Cancel** and close the dialog without saving the changes.

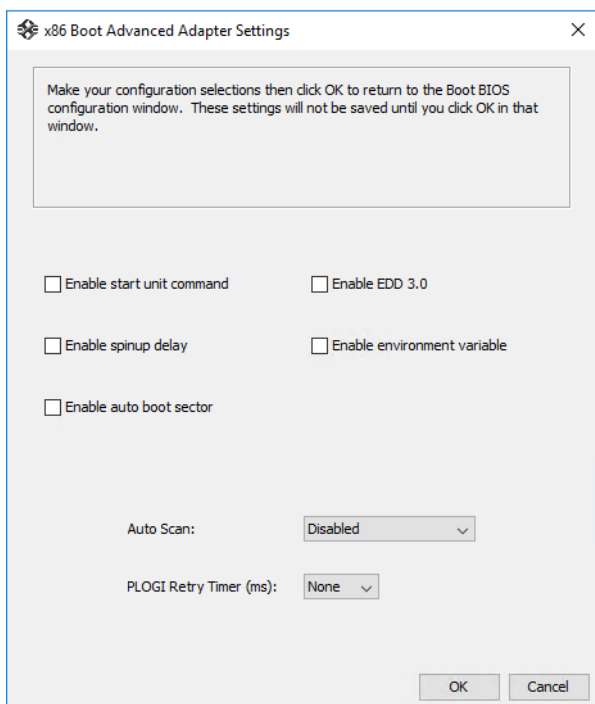
NOTE: If you do not enter the advanced settings and the configuration for the boot type is new, default values are used. The default settings are given with descriptions of the Advanced Adapter Settings dialogs in the following sections.

7.18.1 x86 Boot Advanced Adapter Settings Dialog

Use the **x86 Boot Advanced Adapter Settings** dialog (Figure 52) to configure advanced settings for the selected x86 adapter. All selections are cleared by default. All changes require a reboot to activate.

NOTE: When x86 settings are changed, the same changes are also made to EFI. Conversely, when EFI settings are changed, the changes are also made to x86.

Figure 52: x86 Boot Advanced Adapter Settings Dialog



The following x86 Advanced Adapter Settings fields are displayed:

- **Enable start unit command** – Issues the SCSI start unit command.
- **Enable EDD 3.0** – Enables the EDD option showing the path to the boot device. (Available on Intel Itanium servers only.)
- **Enable spinup delay** – If at least one boot device has been defined, and the spinup delay is enabled, the BIOS searches for the first available boot device.
 - If a boot device is present, the BIOS boots from it immediately.
 - If a boot device is not ready, the BIOS waits for the spinup delay and, for up to three additional minutes, continues the boot scanning algorithm to find another multi-boot device.
 - If no boot devices have been defined and auto scan is enabled, the BIOS waits for five minutes before scanning for devices.
 - In a private loop, the BIOS attempts to boot from the lowest target AL_PA it finds.
 - In an attached fabric, the BIOS attempts to boot from the first target found in the NameServer data.

NOTE: The default topology is auto topology with loop first. Change this topology setting, if necessary, before configuring boot devices.

- **Enable environment variable** – Sets the boot controller order if the system supports the environment variable.

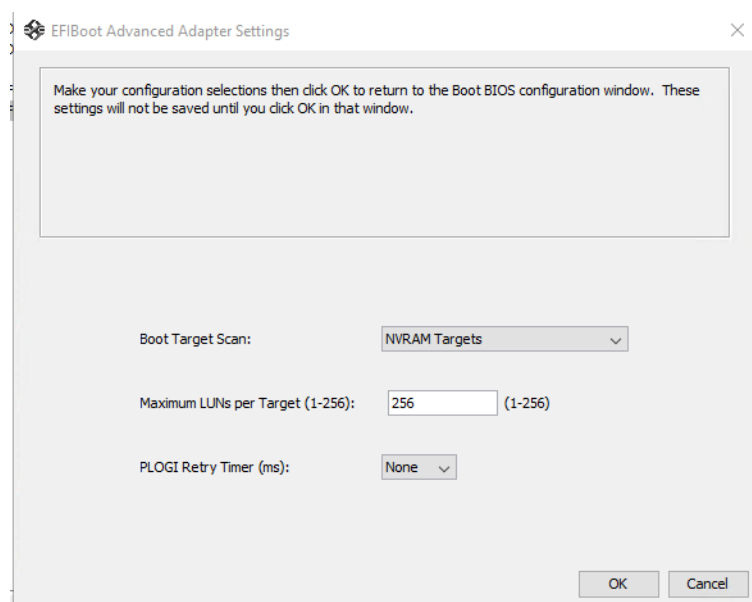
- **Enable auto boot sector** – Automatically defines the boot sector of the target disk for the migration boot process, which applies only to HPE MSA1000 arrays. If no partition exists on the target, the default boot sector format is 63 sectors.
- **Auto Scan** – With auto scan enabled, the first device issues a name server inquiry. The boot device is the first D_ID, LUN 0, or non-LUN 0 device returned, depending on the option you select. Only this device is the boot device, and it is the only device exported to the multiboot menu; which appears at boot time. Auto scan is available only if none of the eight boot entries is configured to boot through D_ID or WWPN. Select one of the following options:
 - **Disabled** (default)
 - **Any First Device**
 - **First LUN 0 Device**
 - **First non-LUN 0 Device**
- **PLOGI Retry Timer (ms)** – Sets the interval for the PLOGI retry timer. This option is especially useful for Tachyon-based RAID arrays. Under very rare occasions, a Tachyon-based RAID array resets itself, and the port goes offline temporarily in the loop. When the port comes back online, the PLOGI retry interval scans the loop to discover this device. The default setting is None (0 ms). Select one of the following values:
 - **None** (default)
 - **50 ms**
 - **100 ms**
 - **200 ms**

7.18.2 EFIBoot Advanced Adapter Settings Dialog

Use the **EFIBoot Advanced Adapter Settings** dialog (Figure 53) to configure the advanced settings for the selected EFIBoot adapter.

NOTE: When x86 settings are changed, the same changes are also made to EFI. Conversely, when EFI settings are changed, the changes are also made to x86.

Figure 53: EFIBoot Advanced Adapter Settings Dialog



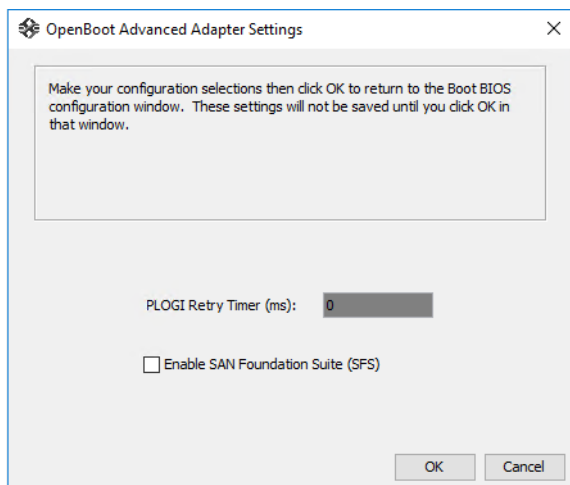
EFIBoot Advanced Adapter Settings field definitions follow:

- **Boot Target Scan** – This option is available only if none of the eight boot entries are configured to boot by using D_ID or WWPN.
 - **NVRAM Targets** (default) – Discovers only LUNs that are saved to the adapter's NVRAM.
 - **Discovered Targets** – Discovers all devices that are attached to the port. Discovery can take a long time on large SANs.
 - **None**.
 - **EFIBootFCScanLevel: NVRAM Targets and EFIBootFCScanLevel: Discovered Targets** – Allows third-party software to toggle between boot path from NVRAM and boot path from Discovered Targets by manipulating an EFI system NVRAM variable.
- **Maximum LUNs per Target** – Sets the maximum number of LUNs that are polled during device discovery. The range is 1 to 4096. The default is 256.
- **PLOGI Retry Timer** – Sets the interval for the PLOGI retry timer. This option is especially useful for Tachyon-based RAID arrays. Under rare occasions, a Tachyon-based RAID array resets itself and the port goes offline temporarily in the loop. When the port comes online again, the PLOGI retry interval scans the loop to discover this device.
 - **None** (default)
 - **50 ms**
 - **100 ms**
 - **200 ms**

7.18.3 OpenBoot Advanced Adapter Settings Dialog

Use the **OpenBoot Advanced Adapter Settings** dialog (Figure 54) to configure the advanced adapter settings for the selected OpenBoot adapter.

Figure 54: OpenBoot Advanced Settings Dialog



OpenBoot Advanced Adapter field definitions follow:

- **PLOGI Retry Timer (ms)** – Sets the PLOGI Retry timer value. The range is 0 to 0xFF.
- **Enable SAN Foundation Suite (SFS)** – Select this check box to enable the SFS driver (the inbox Emulex driver for Solaris).

7.19 Using FC-SP DHCHAP Authentication

Use the **DHCHAP** tab to view and configure FC-SP DHCHAP authentication between an adapter and a switch. FC-SP-2 DHCHAP authentication is disabled by default.

NOTE:

- Boot from SAN is not supported when DHCHAP authentication is enabled.
- DHCHAP is supported only on Windows, Linux, and VMware operating systems.
- DHCHAP is available only for physical ports, not for virtual ports.
- The authentication driver parameters are available only on local hosts. This parameter is not displayed for any remote hosts.
- DHCHAP is not supported on FA-PWWN ports.
- Authentication must be enabled at the driver level. Enable the `lpfc_enable_auth` parameter for Linux, the `EnableAuth` parameter for Windows, or the `lpfc_enable_auth` parameter for VMware before attempting to configure DHCHAP. See [Section 7.17, Configuring the Driver Parameters](#), for instructions on changing driver parameters. Authentication is disabled by default.

To enable DHCHAP on Linux systems, the `lpfc_enable_auth` parameter must be passed to the driver by typing the following command:

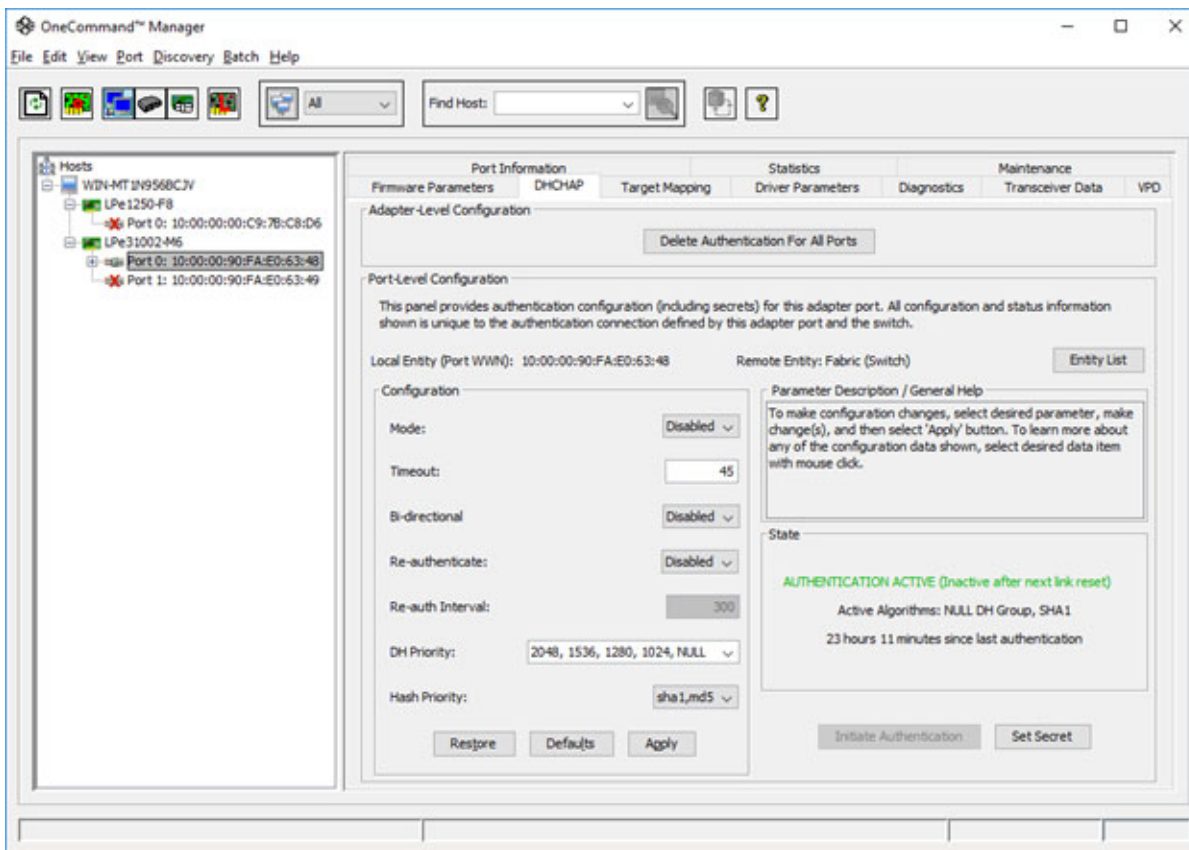
```
insmod lpfc.ko lpfc_enable_auth=1
```

For a permanent configuration on Linux systems that will persist across system reboots, create the `/etc/modprobe.d/lpfc.conf` file, and place the following line into it:

```
options lpfc lpfc_enable_auth=1
```

The **DHCHAP** tab ([Figure 55](#)) enables you to configure authentication.

Figure 55: DHCHAP Tab (LPe31000-series Adapter Depicted)



The following **DHCHAP** tab fields and buttons are displayed:

- **Adapter-Level Configuration** area (not supported on LPe12000-series adapters):
 - Click **Delete Authentication For All Ports** to permanently delete the entire authentication configuration for all the ports on the adapter.
- **Port-Level Configuration** area (not supported on LPe12000-series adapters):
 - Click **Entity List** to see the list of entity pairs with a saved authentication configuration.
- **Configuration** area:
 - **Mode** – The mode of operation. Three modes are available:
 - **Enabled** – The FC function initiates authentication after issuing a FLOGI to the switch. If the connecting device does not support DHCHAP authentication, the software still continues with the rest of the initialization sequence.
 - **Passive** – The FC function does not initiate authentication, but participates in the authentication process if the connecting device initiates an authentication request.
 - **Disabled** – The FC function does not initiate authentication or participate in the authentication process if it is initiated by a connecting device. This is the default mode.
 - **Timeout** – During the DHCHAP protocol exchange, if the switch does not receive the expected DHCHAP message within a specified time interval, authentication failure is assumed (no authentication is performed). The time value ranges from 20 to 999 seconds, and the default is 45 seconds.
 - **Bi-directional** – If this setting is enabled, the FC driver supports authentication initiated by either the switch or the FC function. If this setting is disabled, the driver supports only FC function-initiated authentication. The remote password must be configured to enable this setting. See [Section 7.19.3, Setting or Changing Secrets](#), for instructions.

- **Re-authenticate** – If this setting is enabled, the FC driver can periodically initiate authentication.
- **Re-auth Interval** – The value, in minutes, that the FC driver uses to periodically re-initiate authentication. Valid interval ranges are 10 to 3600 minutes. The default is 300 minutes.
- **DH Priority** – The priority of the five supported DH Groups (Null group, and groups 1, 2, 3, and 4) that the FC driver presents during the DHCHAP authentication negotiation with the switch.
- **Hash Priority** – The priority of the supported hash algorithms (MD5 and SHAx) that the FC driver presents during the DHCHAP authentication negotiation with the switch (default is MD5 first, then SHA1, SHA2, SHA3...).
- **Restore** – Click this button to return parameters to their previous settings.
- **Defaults** – Click this button to return parameters to their default settings.
- **Apply** – Click this button to apply new the parameter settings before initiating authentication.

NOTE: Clicking **Defaults** removes all current configuration settings, including port secrets and this switch/target connection.

- **Parameter Description / General Help** area:
 - This area contains a brief description of the selected parameter and the options available for the parameter.
- **State** area:
 - This area displays the current authentication state. Possible states are Not Authenticated, Authentication In Progress, Authentication Success, Authentication Active, and Authentication Failed.
- **Initiate Authentication** – Click this button after configuring parameters, and clicking the **Apply** button, to perform immediate authentication.
- **Set Secret** – Click this button to set a new local or remote secret in ASCII or hexadecimal (binary). See [Section 7.19.3, Setting or Changing Secrets](#), for instructions.

After DHCHAP has been activated and configured, manually initiate authentication per adapter by clicking **Initiate Authentication** or by inducing a FLOGI time in accordance with the FC-SP standard to the switch. A FLOGI can also be caused by bringing the link between the switch and adapter down and then up (not available in read-only mode).

7.19.1 Deleting Authentication for All Ports

NOTE: The driver authentication parameter `lpfc_enable_auth` (Linux), `lpfc_enable_auth` (VMware), or `EnableAuth` (Windows) must be disabled before deleting authentication for all ports. See [Section 7.17, Configuring the Driver Parameters](#), for instructions on changing driver parameters.

NOTE: This command deletes the authentication configuration, including secrets, from the adapter flash memory. You must reload the driver or reboot the system to activate the new driver settings.

To delete authentication for all ports, perform these steps:

1. In the discovery-tree ([Figure 4](#)), select the adapter port whose authentication you want to delete.
2. Select the **DHCHAP** tab ([Figure 55](#)).
3. Click **Delete Authentication For All Ports**.

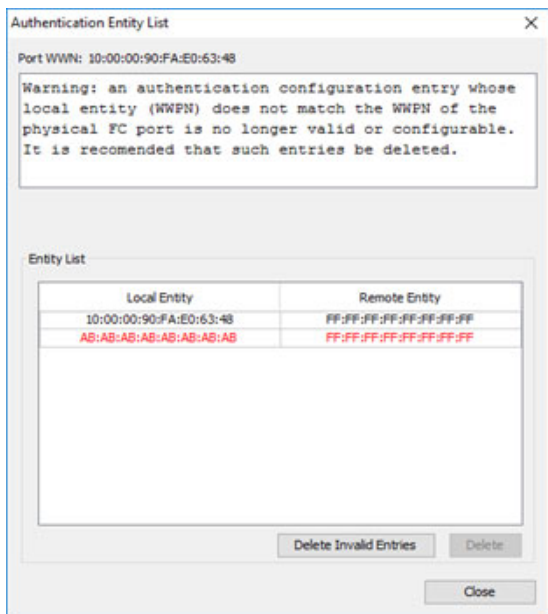
7.19.2 Viewing Saved Authentication Configuration Entities

The Entity List displays a list of entity pairs that have a saved authentication configuration. The list might include entity pairs for authentication configurations that are no longer valid or configurable. For example, the list would contain an entity pair whose configuration become obsolete and invalid after a port WWN change.

To view saved authentication configuration entities, perform these steps:

1. In the discovery-tree (Figure 4), select the adapter port whose authentication configuration entities you want to view.
2. Select the **DHCHAP** tab (Figure 55).
3. Click **Entity List**. The **Authentication Entity List** dialog appears (Figure 56).

Figure 56: Authentication Entity List Dialog



7.19.2.1 Deleting Authentication Entities

You can delete all invalid entities or particular entities.

To delete saved authentication configuration entities, perform these steps:

1. In the discovery-tree (Figure 4), select the adapter port whose authentication configuration entities you want to delete.
2. Select the **DHCHAP** tab (Figure 55).
3. Click **Entity List**. The **Authentication Entity List** dialog appears (Figure 56).
4. Click **Delete Invalid Entries** to remove all invalid entities (red), or select single or multiple entities and click **Delete**.

7.19.3 Setting or Changing Secrets

You can change or set the local or remote secret (password). The local secret is typically used by the driver when the adapter initiates authentication to the switch. The remote secret is used by the driver if the switch attempts to authenticate with the adapter. Bi-directional authentication requires the remote secret.

To set or change secrets, perform these steps:

1. In the discovery-tree (Figure 4), select the adapter port whose secrets you want to set or change.
2. Select the **DHCHAP** tab (Figure 55).
3. Click **Set Secret**. The **Set Secret** dialog appears (Figure 57).

Figure 57: Set Secret Dialog



4. Choose **Set Local Secret** or **Set Remote Secret**.
 - The FC driver uses the local password if the adapter initiates authentication to the switch (typical use).
 - The FC driver uses the remote password if the switch authenticates with the adapter. This situation is possible only when **Bi-Directional** is enabled on the **DHCHAP** tab (Figure 55).
5. To see the password characters entered in the dialog, select the **Show Characters** check box.
6. Enter the new secret. Secrets must contain at least 12 bytes, and local and remote secrets must be different.
7. Re-enter the new values.
8. Select **Alpha-Numeric** or **Hex (binary)** format.
9. Click **OK**.

CAUTION! Do not forget the secret after one has been assigned. After a secret is assigned to an adapter, subsequent DHCHAP configuration settings for that adapter, including the default configuration or new secrets, require you to enter the existing secret to validate your request. No further changes can be made without the secret.

NOTE: Click **Help** on the **Set Secret** dialog for assistance with secrets.

7.19.4 Changing the Authentication Configuration

NOTE: You can configure DHCHAP only on the local host.

To view or change the authentication configuration, perform these steps:

1. In the discovery-tree ([Figure 4](#)), select the adapter port whose configuration you want to view or change.
2. Select the **DHCHAP** tab ([Figure 55](#)).

NOTE: If the fields on this tab are dimmed, either authentication has not been enabled at the driver level, or the local secret has not been set.

- For instructions on enabling the driver authentication parameter `lpfc_enable_auth` (Linux), `lpfc_enable_auth` (VMware), or `EnableAuth` (Windows), see [Section 7.17, Configuring the Driver Parameters](#).
- For instructions on setting the local secret, see [Section 7.19.3, Setting or Changing Secrets](#).

3. Change the configuration values you want.
4. Click **Apply**.

NOTE: If you click **Apply**, changes cannot be canceled.

To return settings to the status before you started this procedure, click **Restore** before you click **Apply**.

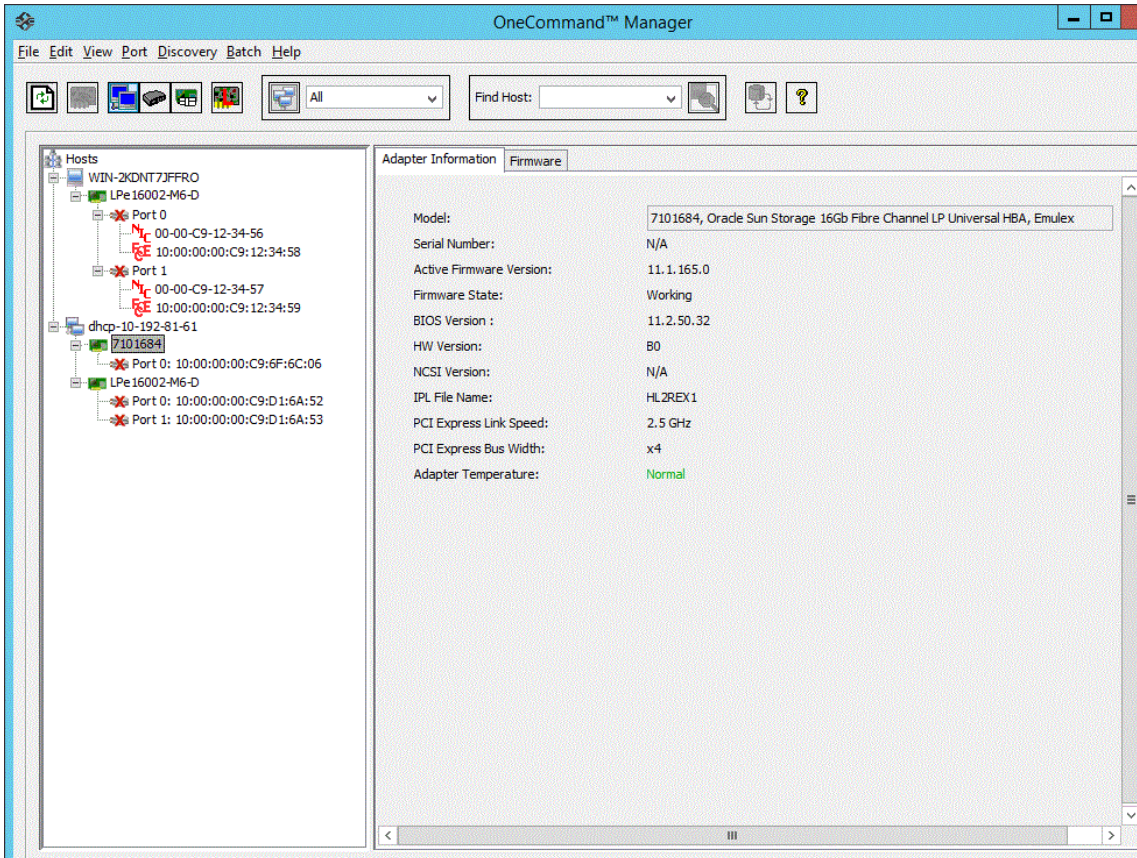
To return all settings to the default configuration, click **Defaults**.

CAUTION! This action also resets any passwords to NULL for this configuration.

7.20 Guest Operating System Discovery and Management from the Base Host Operating System

When the OneCommand Manager application is installed on a guest operating system, the guest operating system is discovered by the OneCommand Manager application running on the host operating system. The guest operating system host appears as a remote host in the discovery-tree ([Figure 58](#)).

Figure 58: OneCommand Manager Application Running on the Base Host after Discovering the Guest Host



Chapter 8: Updating Adapter Firmware

The OneCommand Manager application enables you to update firmware for a single adapter or simultaneously across multiple adapters.

CAUTION! Updating firmware or boot code on an LPe12000-series adapter that is being used to boot from SAN can cause unpredictable behavior. After the update is completed, an adapter reset is issued, which can cause a loss of connectivity to the SAN and possible loss of data.

To update firmware on an LPe12000-series adapter, make sure that the adapter is not currently being used to boot from SAN. Do one of the following:

- Move the adapter to be updated to a non-boot from SAN host, and perform the update from that location.
- If the host with the target adapter is also hosting other boot from SAN adapters, perform a boot from SAN using one of the other boot from SAN adapters. The target adapter can now be updated.

NOTE: For LPe35000-series adapters only:

In some cases, a firmware update requires either a firmware reset or a basic PCIe reset, depending on the features available in the new firmware. A firmware reset is performed automatically if it is needed, regardless of the operating system. A basic PCIe reset is also performed automatically, but only on the following Linux operating systems:

- SLES 12 SP2 and later
- SLES 15 and later
- RHEL 7.6 and later

If a firmware reset or basic PCIe reset occurs after the firmware is downloaded, a message similar to the following appears:

```
Download and pci reset successfully completed.
```

In some cases, a full reboot is required to activate new firmware or to enable a new feature. In that case, a message similar to one of the following messages appears after the firmware download is complete:

```
Download successfully completed. Please reboot the system to activate new firmware.
```

```
Download completed. Some features require an optional reboot. Refer to the Adapter's  
Firmware and Boot Code Release Notes for details.
```

For a list of features that require a reboot in order to be enabled, refer to the *Emulex LPe35000-Series HBA Firmware and Boot Code Release Notes*.

8.1 Updating Firmware for a Single Adapter

NOTE: This option is not available in read-only mode.

Using the **Maintenance** or **Firmware** tab, you can update firmware on local and remote adapters. Before you can perform this procedure, do the following:

- Download the firmware file from www.broadcom.com to a local drive.
- Make sure that the Emulex driver is installed.
- Make sure that the OneCommand Manager application is installed.

- If the adapter is already connected to a boot device, check that the system is in a state in which this type of maintenance can be performed:
 - I/O activity on the bus has been stopped.
 - Cluster software, or any other software that relies on the adapter to be available, is stopped or paused.

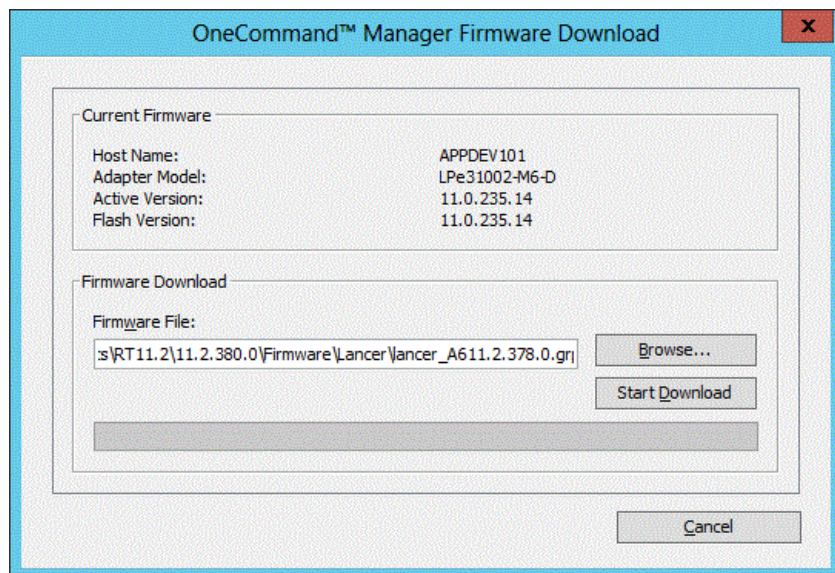
NOTE:

- For vendor-branded adapters, refer to the vendor's website or contact the vendor's technical support for the firmware files.
- You cannot update firmware with the OneCommand Manager application on an Oracle-branded adapter.
- For LPe12000-series adapters, you update the firmware and boot code on each FC port or function. The firmware and boot code are two separate binaries. You must flash both the firmware and boot binaries to update LPe12000-series adapters.
- For all other adapters, you update the firmware and boot code for the entire adapter with a single firmware file.

To update firmware for a single adapter or an adapter port, perform these steps:

1. Select the **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4), select the adapter or port whose firmware you want to update.
3. Select the **Maintenance** tab for LPe12000-series adapters (Figure 33) or the **Firmware** tab for all other adapters (Figure 24), and click **Download Firmware**. The **Firmware Download** dialog appears (Figure 59).

Figure 59: Firmware Download Dialog




4. Using the **Firmware Download** dialog (Figure 59), navigate to the image file you want to download. The firmware image can be specified either by entering the image file's full path name in the **Firmware File** field or by clicking **Browse**. If you click **Browse**, the **Firmware File Selection** dialog appears. Select the file you want to use and click **OK**. The **Firmware Download** dialog appears.
5. Click **Start Download**. A warning popup appears.
6. Click **Yes** to continue.

A status bar shows the progress of the download. The adapter in the discovery-tree (Figure 4) is displayed in black text when the update is complete.

NOTE: The adapter in the discovery-tree might change to blue during the download, but this is normal.

7. Click **Close**. The **Firmware** tab displays the updated firmware information for the selected adapter. You are notified by a popup if a reset is necessary to activate the firmware.



8. Click **OK** to close the popup.

If you are updating the firmware on a multiport LPe12000-series adapter, repeat steps 1 through 7 to update the firmware on the second port, or use the [Section 8.2, Updating Firmware for Multiple Adapters](#), procedure.

8.2 Updating Firmware for Multiple Adapters

Use batch mode to install firmware on multiple adapters in a single step. Batch firmware loading is restricted to a single firmware file and to all accessible, compatible adapters. Batch mode is not available in read-only mode.

NOTE: Stop other OneCommand Manager application functions while batch loading is in progress.

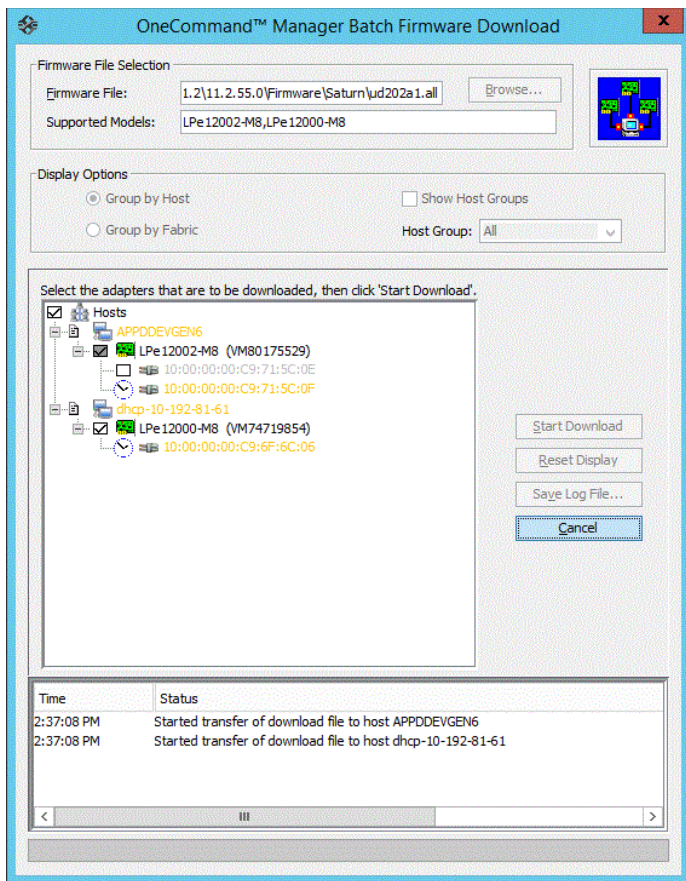
Before you can perform a batch update, the firmware file must be downloaded from www.broadcom.com to a directory on your local drive.

NOTE: VMware ESXi hosts managed through the CIM interface list all adapters regardless of whether the selected firmware can update the adapter. You must manually clear the nonmatching adapters.

To update firmware for multiple adapters, perform these steps:

1. From the **Batch** menu, select **Download Firmware**.
You do not need to select a particular discovery-tree element for this operation.
2. When the **Batch Firmware Download** dialog appears (Figure 60), click **Browse**.

Figure 60: Batch Firmware Download Dialog: Selecting Adapters to Update



3. The **Firmware File Selection** dialog appears. Select the file that you want to use and click **OK**. A dialog notifies you that the OneCommand Manager application is searching for compatible adapters.

After compatible adapters are found, the Firmware File field of the main **Batch Firmware Download** dialog displays the selected image file's path. The **Supported Models** text field displays a list of all adapter models that are compatible with the selected image file. The compatible adapters appears in the tree-view.

Using the **Display Options** settings, you can choose how adapters are displayed in the tree-view. Select **Group by Host** to display adapters in a host-centric view. Select **Group by Fabric** to display hosts in a fabric-centric view with their fabric addresses. The WWPN and host name for each downloadable port is displayed under its respective fabric.

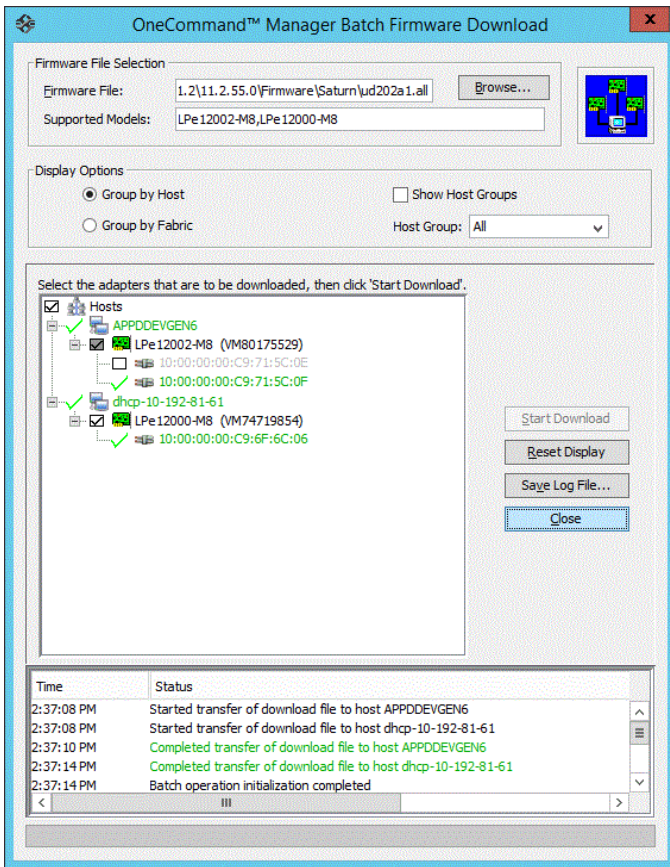
You can also display host groups by selecting the **Show Host Groups** check box. To display a particular host group, choose that group from the **Host Group** list.

You can select and clear host and adapter entries in the **Batch Firmware Download** discovery-tree. Selecting an adapter selects or clears that adapter; selecting a host clears or selects all eligible adapters for that host.


For adapters where each individual port or ASIC can have new firmware downloaded, you can select the ports or ASICs on the adapter to which you want to download firmware.

4. Make your selections, and click **Start Download**. When downloading begins, the tree-view displays the progress. As firmware for a selected adapter is being downloaded, it appears orange in the tree-view. After successful downloading is complete, the entry changes to green. If the download fails, the entry changes to red (Figure 61).

Figure 61: Batch Firmware Download Dialog: Download Complete



A popup notifies you if a reset is necessary to activate the firmware.



5. Click **OK** to close the popup.

To save a copy of the activity log when downloading is finished, click **Save Log File**.

To refresh the **Batch Firmware Download** dialog information when downloading is finished, click **Reset Display**.

Chapter 9: Exporting SAN Information

9.1 Creating a SAN Report

The OneCommand Manager application enables you to create reports about discovered SAN elements. Reports are generated in `.xml` and `.csv` format and include all the SAN information that is displayed through the various OneCommand Manager application tabs.

NOTE: Creating a SAN report can take several minutes for a large SAN.

To create a SAN report, perform these steps:

1. From the **File** menu, select **Export SAN Info**.
2. Browse to a folder and enter a file name with the `.xml` or `.csv` extension.
3. Click **Save** to start the export process.

During the export process, progress is displayed in the lower-right side of the progress bar. On Windows, you cannot change views, reset, or download firmware during the export process.

Chapter 10: Diagnostics

This chapter describes the diagnostic tests that can be run on Emulex adapters.

Use the **Diagnostics** tab to perform the following tasks:

- View PCI registers and Flash Contents information.
- Run the following tests on Emulex adapters installed in the system:
 - PCI loopback
 - Internal loopback
 - External loopback
 - POST
 - Echo (end-to-end)
 - Quick Test

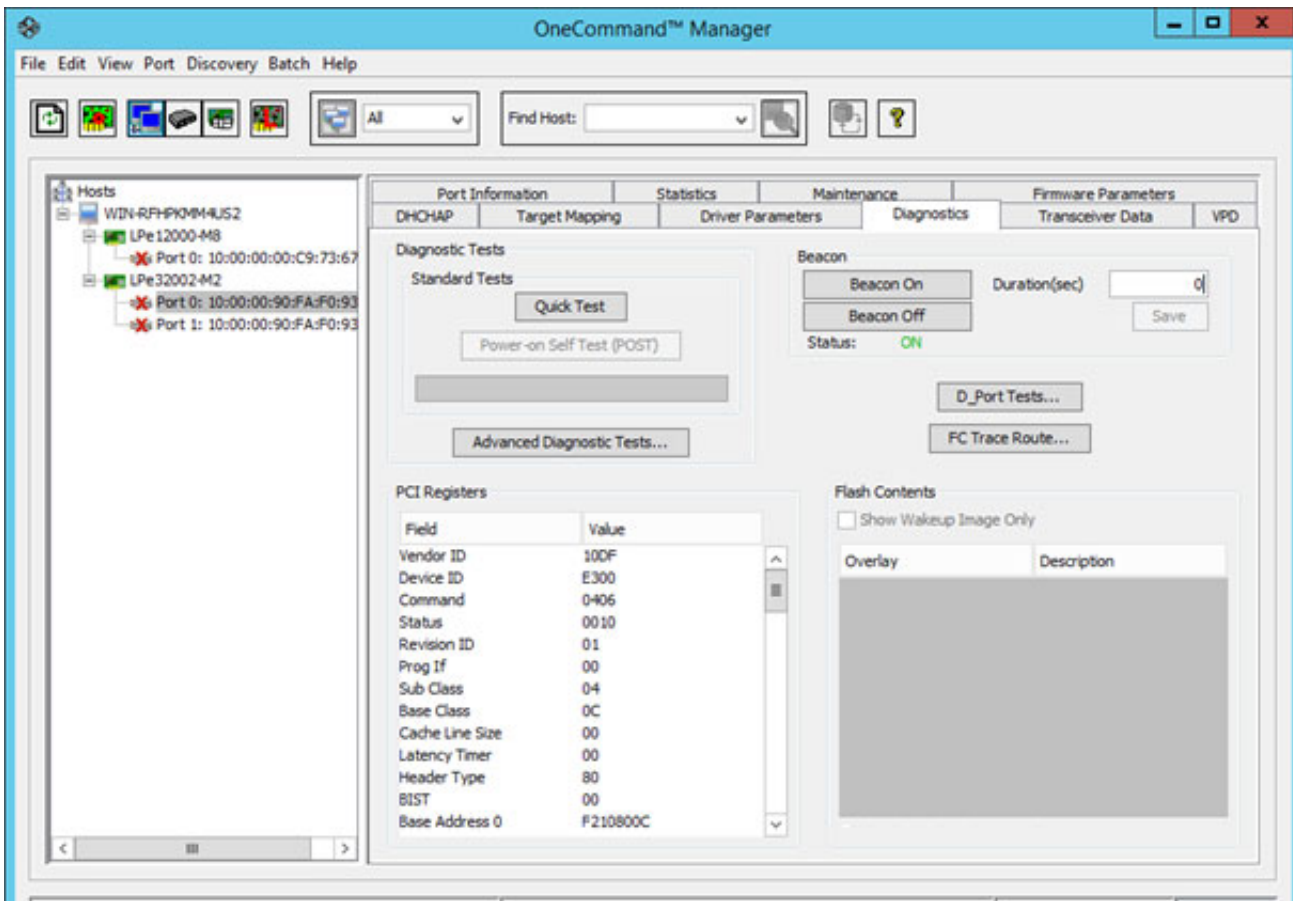
These tests are not available in read-only mode.

- Run a diagnostic dump and retrieve dump files from remote hosts (this option is not available in read-only mode).
- Control adapter beaconing (this option is not available in read-only mode).

CAUTION! Running a POST or loopback test on an LPe12000-series adapter (for example, LPe12000) that is being used to boot from SAN is not advisable. After the tests are complete, the system performs an adapter reset, which might cause a loss of connectivity to the SAN and possible loss of data. To perform these tests on an LPe12000-series adapter, you must make sure that the adapter is not currently being used to boot from SAN. Perform one of the following actions:

- Move the adapter to be tested to a non-boot from SAN host, and perform the tests from that location.
- If the host with the adapter that needs to be tested also hosts other boot from SAN adapters, perform a boot from SAN using one of the other boot from SAN adapters. The target adapter can now be tested, because it is no longer being used for boot from SAN.

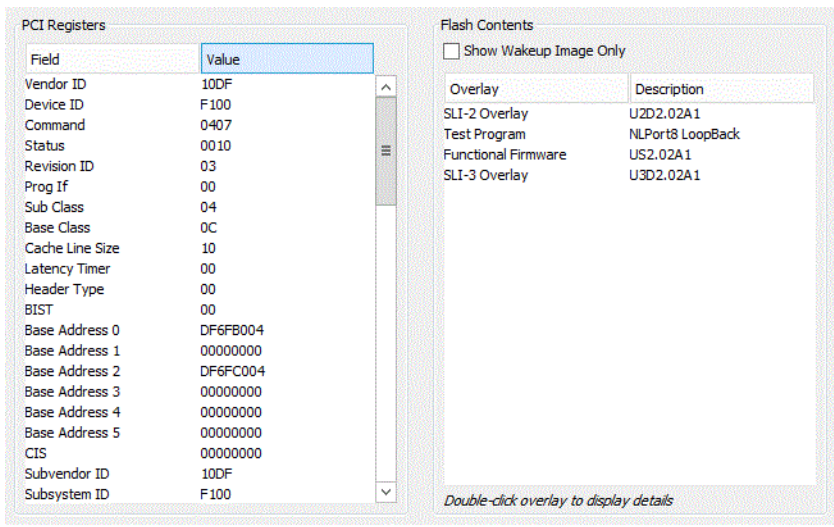
Figure 62: Diagnostics Tab



10.1 Viewing Flash Contents, Overlay Details, PCI Registers, and Wakeup Information

The **Diagnostics** tab shows PCI register dump information and flash memory contents. The information is read-only and is outlined in the following section (Figure 63).

Figure 63: PCI and Flash Information on the Diagnostics Tab




10.1.1 Viewing Flash Contents

By selecting **Show Wakeup Image Only** in the **Flash Contents** area, the flash overlays that are not loaded when the system is booted are not displayed. The default setting is not selected.

10.1.2 Viewing Overlay Details

If you double-click a flash overlay, a popup displays details about that overlay (Figure 64).

Figure 64: Overlay Detail Popup



To see the details of a different flash overlay image, you can either close the details window and double-click another overlay name or choose a different overlay name from the **Overlay Detail** popup.

10.1.3 Viewing the PCI Registers

The PCI registers appear in the PCI Registers area on the **Diagnostics** tab (Figure 63).


10.2 Running a Quick Test

The **Diagnostics** tab enables you to run a quick diagnostics test on a selected port. The quick test consists of 50 PCI loopback test cycles and 50 internal loopback test cycles. This test is not available in read-only mode or on Emulex adapters in ESXi hosts.

To use quick test, perform these steps:

1. From the discovery-tree (Figure 4), select the port on which you want to run the quick test.
2. Select the **Diagnostics** tab (Figure 63) and click **Quick Test**. A warning message appears (Figure 65).

Figure 65: Run Quick Test Warning



3. Click **OK** to run the test. The **Quick Diagnostic Test** window displays the test results.

10.3 Running a POST

NOTE: The POST is supported only on LPe12000-series adapters.

The POST is a firmware test normally performed on an adapter after a reset or restart. The POST does not require any configuration to run. This test is not available in read-only mode.

To run the POST, perform these steps:

1. From the discovery-tree (Figure 4), select the port on which you want to run the POST.
2. Select the **Diagnostics** tab (Figure 63) and click **Power-on Self Test (POST)**. A warning popup appears.
3. Click **OK**. A **POST** popup displays POST information.

NOTE: After the test starts, it cannot be cancelled. It must run to completion.

10.4 Using Beaconing

The beaconing capability enables you to force a specific adapter's LEDs to blink in a particular sequence. The blinking pattern acts as a beacon, making it easier to locate a specific adapter among racks of other adapters. This option is not available in read-only mode.

If you enable beaconing, the two LEDs blink rapidly in unison for 24 seconds, after which the LEDs report the adapter health status for 8 seconds. After the 8 seconds are up, the adapter returns to Beaconing mode. This cycle repeats indefinitely until you turn beaconing off or you reset the adapter.

On supported adapters, you can also specify a specific beaconing duration, in seconds.

NOTE: The beaconing buttons are disabled if the selected adapter does not support beaconing.

To enable beaconing, perform these steps:

1. From the discovery-tree (Figure 4), select the port whose LEDs you want to set.
2. Select the **Diagnostics** tab (Figure 63) and click **Beacon On**. The beacon **Status** changes to **On**.
On supported adapters, you can also enter an optional **Duration** time, in seconds, for the LEDs to blink. Enter the duration time and click **Save**.

To disable beaconing, perform these steps:

1. From the discovery-tree (Figure 4), select the port whose LEDs you want to set.
2. Select the **Diagnostics** tab (Figure 63) and click **Beacon Off**. The beacon **Status** changes to **Off**.

10.5 Running D_Port Tests

D_Port is a diagnostic mode supported by 16GFC and faster Brocade switches. D_Port tests enable you to detect physical cabling issues that result in increased error rates and intermittent behavior. If activated, D_Port runs a series of tests including local electrical loopback, loopback to the remote optics, loopback from the remote port to the local optics, and a full device loopback test with data integrity checks. It also provides an estimate of cable length to validate that a proper buffering scheme is in place. The various loopback tests allow some level of fault isolation so you can distinguish faults from marginal cable, optics modules, and connector or optics seating.

Bidirectional D_Port testing is supported. The switch or initiator can initiate D_Port testing.

NOTE:

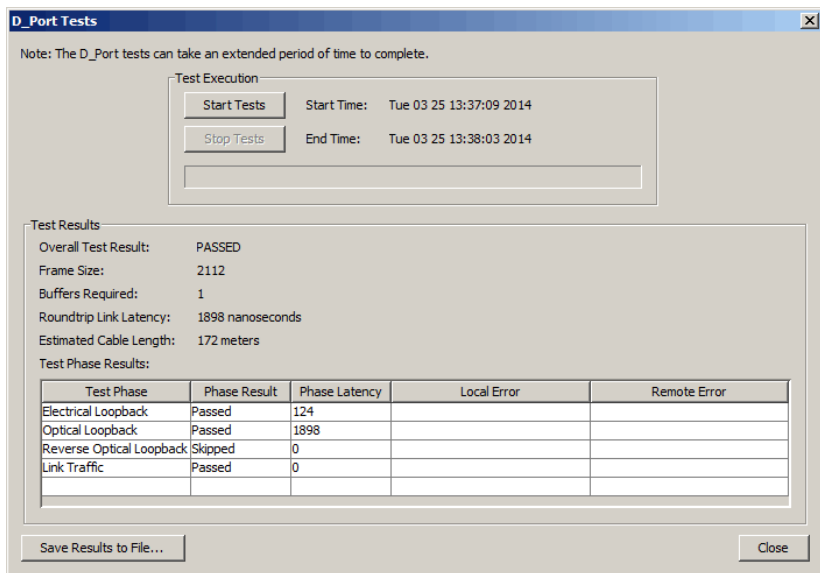
- D_Port must be enabled on the HBA port. (D_Port is enabled by default.) See [Section 7.3.3, Disabling and Enabling Dynamic D_Port Tests](#), for more information.
- D_Port testing is not available when FC port trunking (aggregation) is enabled.
- Basic connectivity diagnostics are already supported by Emulex HBAs. The OneCommand Manager application has diagnostic modes that support validation of the connection to the switch. The functionality that Brocade offers provides the ability to diagnose marginal cable conditions (for example, dust in the optics) that result in higher error rates.
- Do not enable D_Port on the switch port.
- D_Port tests run with the physical connection in an offline diagnostic state, so normal I/O cannot be sent through the physical port while the test is in progress. While the port is in D_Port mode, the link appears down on that port, similar to an unplugged cable.
- When using D_Port in a boot from SAN configuration, the configuration must have redundant paths to the boot LUN, and only one of the redundant adapter ports should be set to D_Port.
- For more information about D_Port, refer to the Brocade section of the Broadcom website at www.broadcom.com.
- D_Port is also referred to as ClearLink®.

The **D_Port Tests** button on the **Diagnostics** tab (Figure 63) enables you to run D_Ports tests on supported adapters.

To run a D_Port test, perform these steps:

1. From the discovery-tree (Figure 4), select the port on which you want to run the D_Port test.
2. Select the **Diagnostics** tab (Figure 63) and click **D_Port Tests**. The **D_Port Tests** dialog appears (Figure 66).
3. Click **Start Tests**. The start time is displayed.

Figure 66: D_Port Tests Dialog



The following **D_Port Tests** dialog fields are displayed:

- **Test Results** area:
 - **Overall Test Result** – **PASSED** or **FAILED**, depending upon the outcome of all the test phases.
 - **Frame Size** – The size of the frames used in each test phase.
 - **Buffers Required** – The number of buffers required for each test phase.
 - **Roundtrip Link Latency** – Estimated roundtrip link latency calculated by the switch during the execution of all tests.
 - **Estimated Cable Length** – Estimated cable length calculated by the switch during the execution of all tests.
- **Test Phase Results** area:
 - **Test Phase** – The name of the test run.
 - **Phase Result** – The result of the test run. Possible results are Passed, Failed, or Skipped.
 - **Phase Latency** – The round trip legacy (in ns.) calculated during the execution of the test.
 - **Local Error** – The errors, if any, detected on the local side of the test.
 - **Remote Error** – The errors, if any, detected on the remote side of the test.

To stop running D_Port tests, click **Stop Tests**. The stop time is displayed.

To save the test results to a file, click **Save Results to File**. You are prompted to enter the file name.

NOTE: If the SFP or adapter firmware does not support running D_Port diagnostics, clicking **Start Tests** causes an error message to be displayed, and the tests are not run.

10.6 Using FC Trace Route

The OneCommand Manager application allows you to trace the communication route for FC packets transmitted between an FC initiator port and an FC target port.

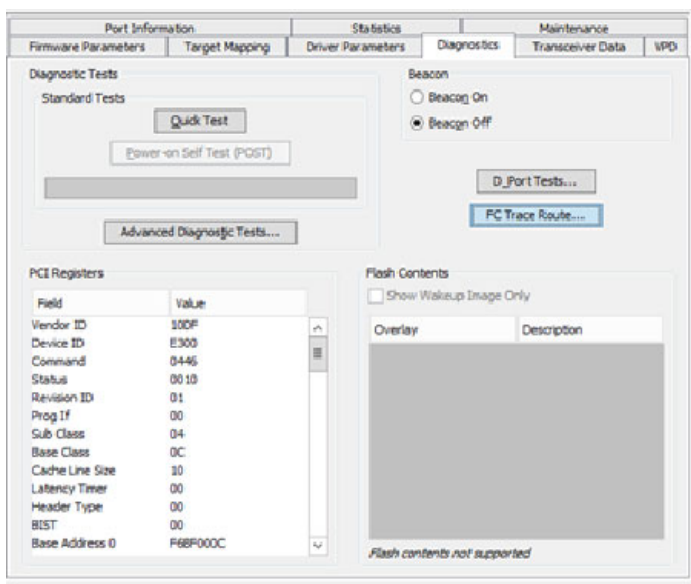
Communication route information, such as the switch name, domain ID, ingress and egress port name, and ingress and egress physical port number, is accumulated for all switch ports through which packets are routed. Data for both the outward bound route from the initiator to the target, and the inbound route from the target to the initiator, is collected.

The **FC Trace Route** button on the **Diagnostics** tab enables you collect an adapter's FC trace route information (Figure 67).

NOTE: The following considerations apply to FC trace route:

- FC trace route is not supported on LPe12000-series adapters.
- Both local and remote support for the FC trace route must be provided.
- FC Trace Route support is provided on Windows, Linux and ESXi operating system platforms.

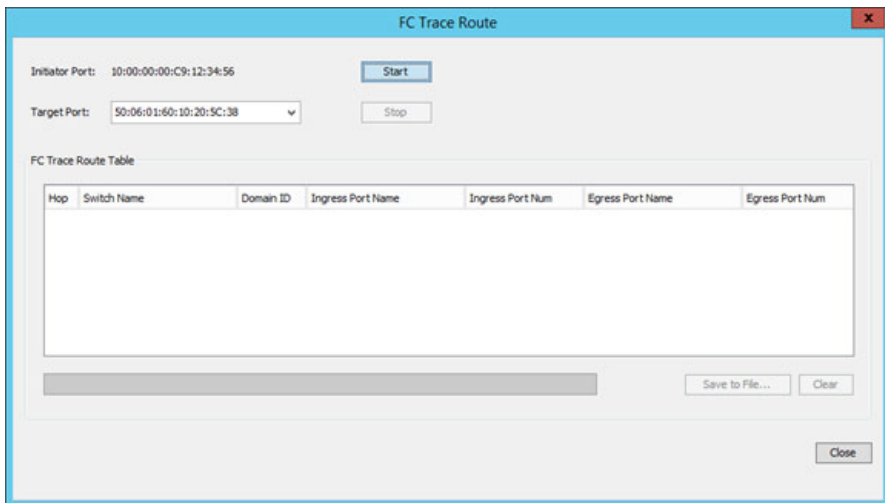
Figure 67: Diagnostics Tab (FC Trace Route Button Depicted)



To enable FC trace route, perform these steps:

1. From the discovery-tree (Figure 4), select the port on which you want to enable FC trace route.
2. Select the **Diagnostics** tab (Figure 67) and click **FC Trace Route**. The **FC Trace Route** dialog appears (Figure 68).

Figure 68: FC Trace Route Dialog



3. The **Target Port** list displays the WWPNs of all targets that are seen by the initiator port. Select a target port and Click **Start**.

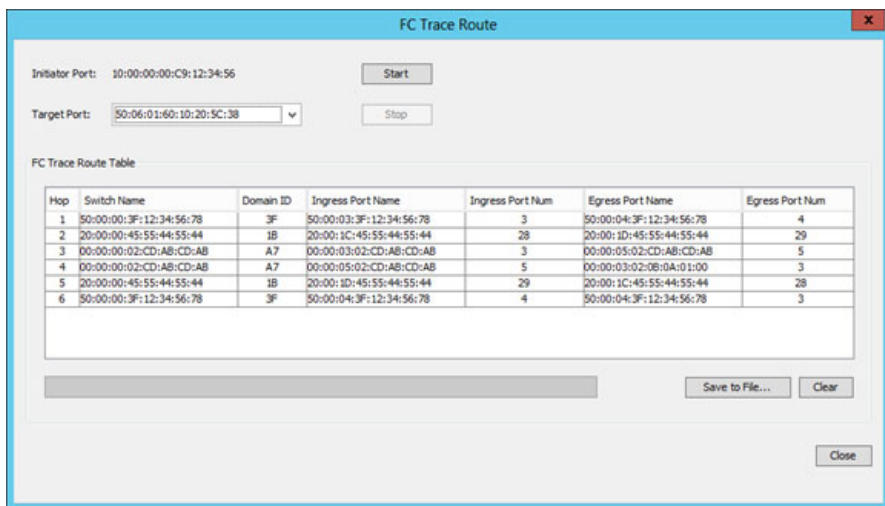
The **FC Trace Route** dialog displays trace route information for selected initiator and target ports (Figure 69).

NOTE: If an error occurs when processing the FC trace route request, a message is displayed at the bottom of the dialog. Click **Stop** to stop the trace route request.

Click **Save to File** to save the results of the most recent FC trace route operation to a log text file. The default file name for the log text file is `FCTraceRte_IWWPN_TWWPN` (where *IWWPN* is the initiator WWPN and *TWWPN* is the target WWPN). You can change the file name.

Click **Clear** to erase the FC Trace Route Table information from the dialog.

Figure 69: FC Trace Route Dialog with Route Information Displayed



The following information is collected for each trace route:

- **Switch Name** – The switch chassis WWN.
- **Domain ID** – A number used to uniquely identify a switch in a fabric. This number is assigned by a fabric administrator as part of fabric configuration. The Domain ID is an 8-bit field whose value ranges from 0 to 255.

- **Ingress Port Name** – The port WWN of the physical port through which an FC packet enters a specific switch.
- **Ingress Port Num** – The physical port number of the port through which an FC packet enters a specific switch.
- **Egress Port Name** – The port WWN of the physical port through which an FC packet exits a specific switch.
- **Egress Port Num** – The physical port number of the port through which an FC packet exits a specific switch.

10.7 Creating Diagnostic Dumps

NOTE: This option is not available in read-only mode.

The diagnostic dump capability enables you to create a dump file for a selected port. Dump files contain various information, such as firmware version and driver version, that is useful when troubleshooting an adapter. You can also retrieve dump files from remote hosts.

For LPe12000-series HBAs, see [Section 10.7.1, Creating Diagnostic Dumps for LPe12000-Series Adapters](#). For other adapters, see [Section 10.7.2, Creating Diagnostic Dumps for All Other Adapters](#).

CAUTION! Disruption of service can occur if a diagnostic dump is run during I/O activity.

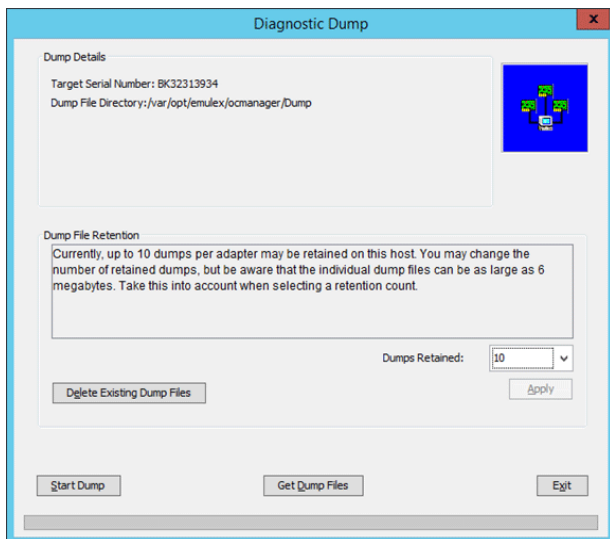
NOTE: For VMware systems, you must set a dump directory before initiating a dump. The dump directory must be a storage partition (a datastore) under the directory `/vmfs/volumes`.

10.7.1 Creating Diagnostic Dumps for LPe12000-Series Adapters

To start a diagnostic dump, perform these steps:

1. From the discovery-tree ([Figure 4](#)), select a port whose diagnostic information you want to dump.
2. Select the **Diagnostics** tab ([Figure 63](#)) and click **Diagnostic Dump**. The **Diagnostic Dump** dialog appears ([Figure 70](#)).
3. Specify how many files you want to retain by selecting the number from the **Dumps Retained** list and clicking **Apply**.
4. Click **Delete Existing Dump** files to remove existing dump files from the selected port on your system.

Figure 70: Diagnostic Dump Dialog



5. Click **Start Dump**. A warning to take the adapter offline appears.

6. Click **Start Dump**. Dump files are created. The file location depends upon your operating system:
 - Windows – In the `Dump` directory under the OneCommand Manager Installation Directory `Util\Dump\`.
 - Solaris – In the `/opt/ELXocm/Dump` directory.
 - Linux – In the `/var/log/emulex/ocmanager/Dump` directory.
 - VMware – In the dump directory that you created under `/vmfs/volumes`.

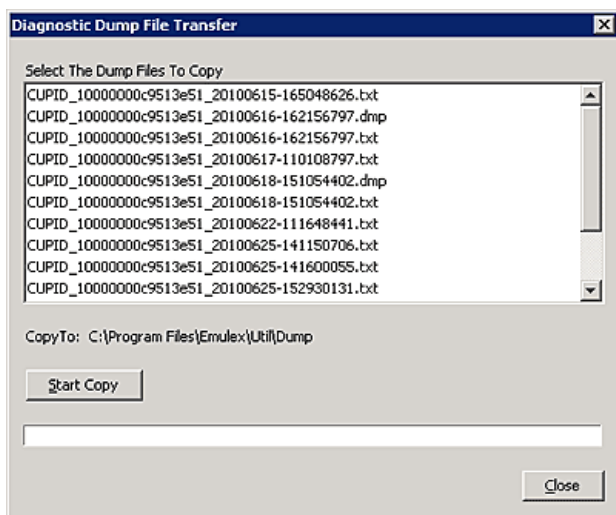
Two files are created:

- `<Hostname_WWPN_Date-Time>.dmp`
- `<Hostname_WWPN_Date-Time>.txt`

7. To list the dump files in the local system or to obtain remote host dump files and copy them to your local system, click **Get Dump Files**. The **Diagnostic Dump File Transfer** dialog appears (Figure 71).

NOTE: The **Start Copy** button is dimmed when a local adapter port is selected.

Figure 71: Diagnostic Dump File Transfer Dialog



8. Select the files that you want to copy (multiple selections are available), and click **Start Copy**. The remote dump files are copied to your local dump folder. The local dump folder locations are described in [Step 6](#).

10.7.2 Creating Diagnostic Dumps for All Other Adapters

NOTE: This option is not available in read-only mode.

The diagnostic dump capability enables you to create a dump file for a selected adapter. Dump files contain various information, such as firmware version and driver version, that is useful when troubleshooting an adapter.

You can retrieve user initiated and driver initiated driver dump files, delete the dump files, or repeat the process on all resident dump files. You can also retrieve or delete dump files from remote hosts.


To start a diagnostic dump, perform these steps:

1. From the discovery-tree (Figure 4), select an adapter whose diagnostic information you want to dump.
2. Select the **Firmware** tab (Figure 24), and click **Diagnostic Dump**. The **Diagnostic Dump** dialog appears (Figure 72).

For hosts being managed through the CIM interface, the **Set Dump Directory** button enables you to set the dump directory for ESXi host dumps (VMware only).

3. Specify how many files that you want to retain by selecting a number from the **Dumps Retained** list and clicking **Apply**.
4. Click **Delete Existing Dump Files** to remove existing dump files for the selected adapter from your system.

Figure 72: Diagnostic Dump Dialog



5. Click **Start Dump**. Dump files are created. The file location depends upon your operating system:
 - Windows – %ProgramFiles%Util\Dump\.
 - Solaris – /opt/ELXocm/Dump.
 - Linux – /var/log/emulex/ocmanager/Dump.
 - VMware – A dump directory you create under /vmfs/volumes.

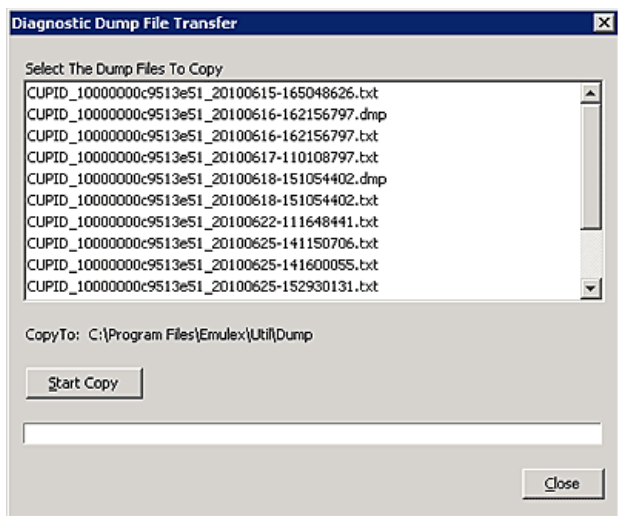
Two files are created:

- <Hostname_WWPN_Date-Time>.efd
- <Hostname_WWPN_Date-Time>.txt

6. To obtain remote host dump files and copy them to your local system, click **Get Dump Files**. The **Diagnostic Dump File Transfer** dialog appears (Figure 73).

NOTE: The **Get Dump Files** button is dimmed if a local adapter port is selected.

Figure 73: Diagnostic Dump File Transfer Dialog



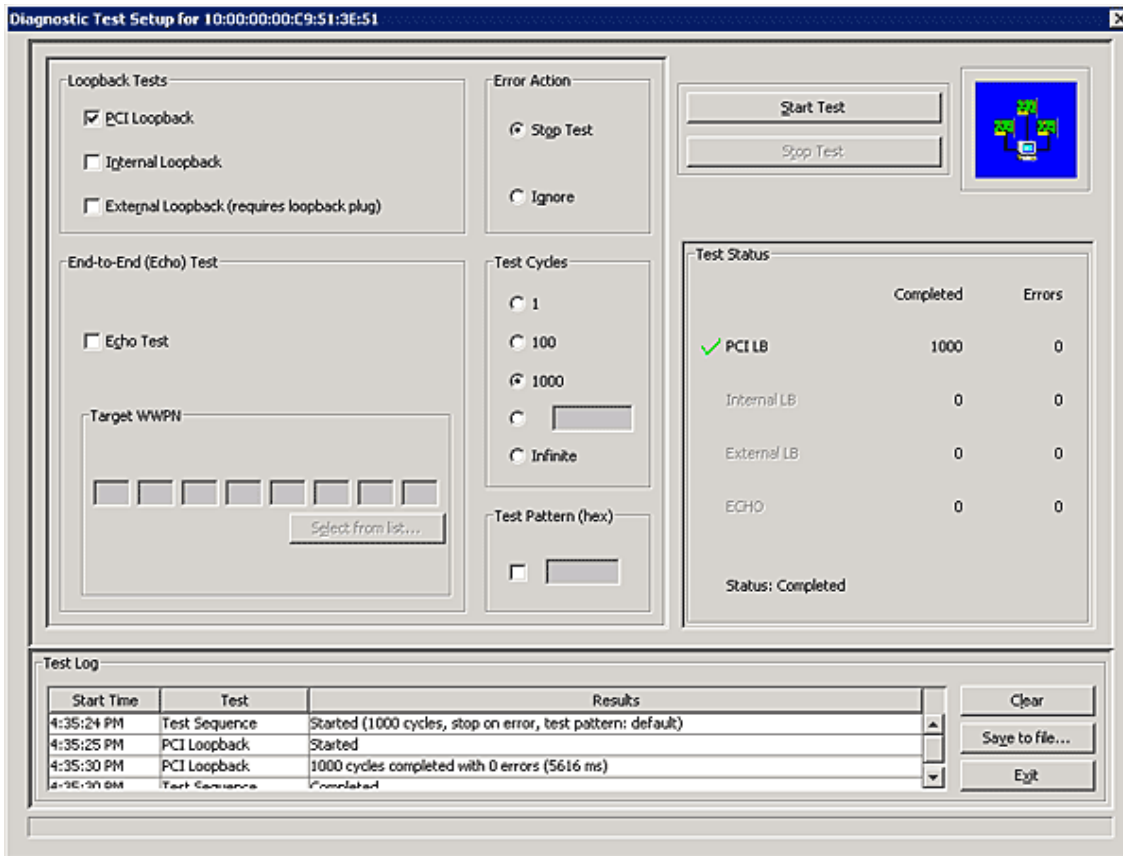
7. Select the files you want to copy (multiple selections are available), and click **Start Copy**. The remote dump files are copied to your local dump folder. The local dump folder locations are described in [Step 6](#).

10.8 Running Advanced Diagnostic Tests

The advanced diagnostics capability gives you greater control than the quick test over the type of diagnostics tests that run. Through advanced diagnostics, you can specify which tests to run, the number of cycles to run, and actions to take in the event of a test failure. Advanced diagnostics capability is not available in read-only mode.

To run advanced diagnostics tests, click **Advanced Diagnostic Tests** on the **Diagnostics** tab ([Figure 63](#)). The **Diagnostic Test Setup** dialog appears ([Figure 74](#)).

Figure 74: Diagnostic Test Setup Dialog



The following **Diagnostic Test Setup** dialog fields are displayed:

- **Loopback Tests** area:
 - **PCI Loopback**
 - **Internal Loopback**
 - **External Loopback**

NOTE: For details about these tests, see [Section 10.8.1, Running Loopback Tests](#).

- **End-to-End (Echo) Test** area:
 - **Echo Test**
 - **Target WWPN**

NOTE:

- For details about this test, see [Section 10.8.2, Running End-to-End Tests](#).
- You cannot run the External Loopback test and the Echo test concurrently. If you select **External Loopback**, the Echo test section is dimmed, and if you select the **Echo Test** check box, the External Loopback section is dimmed.

- **Error Action** area:

Error action enables you to define the actions to be performed in the event of a test failure.

- **Stop Test** – Do not log the error, and halt the test. No further tests are run.
- **Ignore** – Log the error, and proceed with the next test cycle.

■ **Test Cycles** area:

Test cycles enables you to specify test cycles in three ways:

- Select an established cycle count by choosing the desired number.
- Enter a custom cycle count in the blank field in the Test Cycles area.
- Set the test to run until you click **Stop Test**, by selecting **Infinite**.

■ **Test Pattern** area:

Enter a custom test pattern to be used in tests that transfer data. The test pattern can be up to 8 hexadecimal bytes.

■ **Test Status** area:

The Test Status area displays the number of completed cycles of each test run, in addition to the number of errors for each test.

■ **Test Log** area:

For details about test logs, see [Section 10.8.3, Saving the Log File](#).

10.8.1 Running Loopback Tests

You can run the following loopback tests:

- **PCI loopback test** – A firmware-controlled diagnostic test in which a random data pattern is routed through the PCI bus without being sent to an adapter link port. The returned data is subsequently validated for integrity.
- **Internal loopback test** – A diagnostic test in which a random data pattern is sent to an adapter link port, then is immediately returned without actually going out on the port. The returned data is subsequently validated for integrity.
- **External loopback test** – A diagnostic test in which a random data pattern is sent to an adapter link port. The data goes out the port and immediately returns by way of a loopback connector. The returned data is subsequently validated for integrity.

NOTE:

- You cannot run the external loopback test and echo test concurrently. If you select **External Loopback**, the echo Test section is dimmed, and if you select **Echo Test**, **External Loopback** is dimmed.
- Adapters and port information are not available during diagnostic loopback tests.
- Internal and external loopback tests on trunking enabled ports do not support Infinite test cycles.
- Internal and external loopback test results are displayed for each physical port.
- Each physical port must have a loopback connector when performing external loopback tests on trunking enabled ports.


To run loopback tests, perform these steps:

1. From the **Diagnostics** tab ([Figure 63](#)), click **Advanced Diagnostics Tests** ([Figure 74](#)). From the **Loopback Test** section of the dialog, choose the type of loopback test you want to run, and specify the error action, the test cycles, and test pattern.

NOTE: You must insert a loopback plug in the selected port before running an external loopback test.

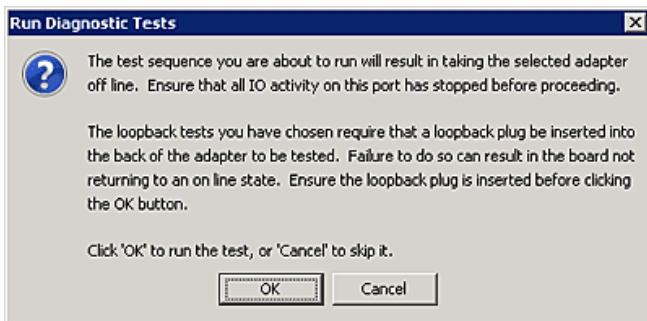
2. Click **Start**. The following popup appears ([Figure 71](#)).

Figure 75: Run Diagnostic Tests Popup



3. Click **OK**. If you chose to run an external loopback test, the following popup appears (Figure 76).

Figure 76: Run Diagnostic Tests Popup for External Loopback



4. Click **OK**. A progress bar indicates that the test is running.

Periodic test feedback, consisting of the current loopback test/cycle plus the completion status of each type of test, is displayed in the **Test Log** section of the dialog. Either click **Clear** to erase the contents of the log display, or click **Save to File** to save the log file.

After starting the tests, you can click **Stop Tests** to stop the tests before they complete. Depending upon the tests being run, it might take some time before they stop.

10.8.2 Running End-to-End Tests

The End-to-End (Echo) test enables you send an `echo` command and response sequence between an adapter port and a target port.

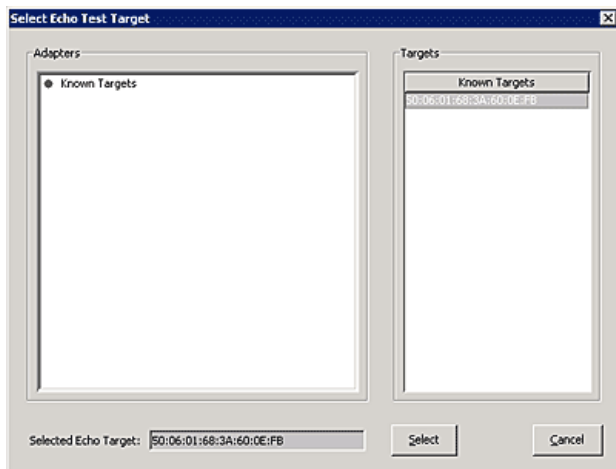
NOTE:

- Not all remote devices respond to an `echo` command.
- You cannot run the echo test and the external loopback test concurrently. If you select **Echo Test**, **External Loopback** is dimmed.

To run echo tests, perform these steps:


1. From the **Diagnostics** tab (Figure 63), click **Advanced Diagnostic Tests** (Figure 74).
2. Select **Echo Test**, and enter the WWPN for the target.
 - a. Click **Select From List** if you do not know the WWPN of the test target. The **Select Echo Test Target** dialog appears (Figure 77).
 - b. Select the port to test from the tree-view and click **Select**. All relevant information for the selected port is automatically added to the Target Identifier section of the **Diagnostics** tab.

Figure 77: Select Echo Test Target Window



3. Define the other parameters you want to use and click **Start Test**. The following warning popup appears ([Figure 78](#)).

Figure 78: Run Diagnostic Tests Popup



4. Click **OK**. A result screen appears, and the test results appear in the test log. Either click **Clear** to erase the contents of the log display, or click **Save to File** to save the log file.

10.8.3 Saving the Log File


You can save the test log to a log file for later viewing or printing. When new data is written to a saved file, the data is appended to the end of the file. Each entry has a two-line header that contains the identifier of the adapter being tested and the date and time of the test. Over time, the data accumulates to form a chronological history of the diagnostics performed on the adapter.

The default location is as follows:

- In Windows: The OneCommand Manager application install directory on your local drive.
- In Solaris: `/opt/ELXocm/Dump`
- In Linux: `/var/opt/emulex/ocmanager/Dump`
- In VMware Server: A default directory does not exist for VMware.

After writing an entry into the log, you are prompted to clear the display. The default name of the saved file is `DiagTest.log`. An example of a saved log file is shown in [Figure 79](#).

Figure 79: Example of a DiagTest.log Window



To save the log file, perform these steps:

1. After running a test from the **Diagnostic Test Setup** dialog (Figure 74), click **Save to File**. The **Select Diagnostic Log File Name** dialog appears. The default name of a saved file is `DiagTest.log`.
2. Browse to the desired directory, change the log file name if you want, and click **Save** (Figure 74).

Chapter 11: Troubleshooting

Your system might operate in an unexpected manner in some circumstances. The Troubleshooting chapter explains many of these circumstances and offers one or more workarounds for each situation.

11.1 General Situations

The following table lists general situations and their resolution.

Table 5: General Situations

Situation	Resolution
After installing and starting the OneCommand Manager application, the status bar shows <code>Initializing discovery engine...</code> , but after waiting for a while, nothing is displayed in the discovery-tree.	It is possible that the discovery server was not installed properly and therefore is not running. Try uninstalling and re-installing the OneCommand Manager application package.
The FC link fails to come up.	Verify that an 8GFC adapter is not attempting to connect to a 1GFC device or that a 16GFC adapter is not attempting to connect to a 1GFC or 2GFC device. Only 8GFC devices are supported on 8GFC adapters, and only 16GFC devices are supported on 16GFC adapters.
The other utilities can be installed, but the OneCommand Manager application cannot.	You have attempted to install the utilities before installing the Emulex driver. Perform the installation tasks in the following order: 1. Install the Emulex driver (refer to the Installation section of the driver manual). 2. Install the utilities (refer to the Installation section of this manual).
An operating error, such as a system hang, occurs when attempting to run the OneCommand Manager application.	Reboot the system.
Unwanted remote servers appear in the OneCommand Manager application.	To prevent remote servers from appearing in the OneCommand Manager application, perform one of the following tasks on the remote systems: <ul style="list-style-type: none"> ■ In Windows, disable the OneCommand Manager application service. ■ In Linux, stop the <code>elxhbamgr</code> daemon by running the following script: <pre>/usr/sbin/ocmanager/stop_ocmanager</pre> ■ In Solaris, stop the <code>elxhbamgr</code> service by issuing the following command: <pre>svcadm disable elxhbamgr</pre> NOTE Disabling this service or process prevents the local servers from being seen remotely.
If Help > Contents is selected in the OneCommand Manager application, the online help is not opened in a web browser. The <code>OCManager_Help.htm</code> file can be opened in a text editor (displaying HTML code) or by some other application. This occurs when the operating system has associated <code>.html</code> files with an application other than a web browser.	On Windows systems, this can be fixed using the following steps: <ol style="list-style-type: none"> 1. In Windows Explorer, navigate to the <code>C:\Program Files\Emulex\Util\OCManager\OCManager_help\</code> directory. 2. Right-click on OCManager_Help.htm. 3. Select Open With > Choose default program. 4. Select a web browser, such as Internet Explorer. 5. Check Always use the selected program to open this kind of file. 6. Click OK. On Linux and Solaris, the preceding steps are very similar, with the <code>OCManager_Help.htm</code> file located in the following directories: <pre>/usr/sbin/ocmanager/ocmanager_help/OCManager_Help.htm</pre> and <pre>/opt/ELXocm/ocmanager_help/OCManager_Help.htm</pre>

11.2 Emulex Driver for Linux and OneCommand Manager Application Situations

The following lists possible situations and resolutions involving the OneCommand Manager application for Linux.

Table 6: OneCommand Manager for Linux Situations

Situation	Resolution
The OneCommand Manager application software package is installed, but an error message says: inserv Service Elxlpfc has to be enabled for service ElxDiscSrvinserv: exiting now/sbin/inserv failed exit code 1.	Reinstall the driver for Linux with the <code>elx_lpfc_install</code> script.
<code>rmmod</code> fails to unload the driver because the device or resource is busy. This message occurs when you attempt to remove the driver without first stopping the OneCommand Manager application, when the OneCommand Manager application is installed and running, or when FC disks connected to an Emulex adapter are mounted.	Stop the OneCommand Manager application before attempting to unload the driver. The script is located in the <code>/usr/sbin/ocmanager</code> directory. <ol style="list-style-type: none"> 1. Type the following command: <code>./stop_ocmanager</code> 2. Unmount any disks connected to the adapter. 3. Unload the driver. 4. Type the following command: <code>rmmod lpfc</code>
The driver uninstallation fails. The <code>elx_lpfc_install --uninstall</code> script fails with an error.	Try the following solutions: <ul style="list-style-type: none"> ■ Uninstall the OneCommand Manager application by running the <code>/uninstall</code> script from the OneCommand Manager application installation directory. ■ Unmount all FC disk drives. ■ Unload the FC driver.
The OneCommand Manager application cannot be installed. The following error message is displayed: inserv Service Elxlpfc has to be enabled for service ElxDiscSrvinserv: exiting now/sbin/inserv failed exit code 1.	Reinstall the driver with the <code>elx_lpfc_install</code> script.
The OneCommand Manager application cannot see any adapters.	Try the following solutions: <ul style="list-style-type: none"> ■ Perform an <code>lsmod</code> to see if the Emulex drivers are loaded. Look for an error message on the command line stating that the LPFC driver is not loaded. If this is the case, do an <code>lsmod</code> of the FC driver and restart the OneCommand Manager application. ■ Exit the OneCommand Manager application and run the following scripts in this order: <ol style="list-style-type: none"> 1. <code>/usr/sbin/ocmanager/stop_ocmanager</code> – Stops the OneCommand Manager application daemons. 2. <code>/usr/sbin/ocmanager/start_ocmanager</code> – Starts the OneCommand Manager application daemons. 3. <code>/usr/sbin/ocmanager/ocmanager</code> – Starts the OneCommand Manager application GUI. <p>The adapters should be visible. If they are not visible, reboot your system.</p>

Table 6: OneCommand Manager for Linux Situations (Continued)

Situation	Resolution
The OneCommand Manager application cannot see new LUNs.	<p>Try the following:</p> <ol style="list-style-type: none">1. Click Refresh LUNs in the toolbar.2. Exit the OneCommand Manager application and restart it. If new LUNs are visible, you are finished. <p>If that does not work, try the following:</p> <ol style="list-style-type: none">1. Exit the OneCommand Manager application.2. Navigate to <code>/usr/sbin/ocmanager</code>.3. Run <code>./stop_ocmanager</code> to stop both the <code>elxhbamgr</code> and <code>elxdiscovery</code> processes.4. Run <code>./start_ocmanager</code> and <code>./start_elxdiscovery</code> to restart both processes.5. Start the OneCommand Manager application.

11.3 vPorts and OneCommand Manager Application Situations

The following table lists possible vPort situations and their resolution.

Table 7: vPort and OneCommand Manager Application Situations

Situation	Resolution
The OneCommand Manager application fails to create vPorts.	If an error occurs during vPort creation, an error message indicates the failure. Several conditions must be met before a virtual port can be created. To view a detailed list of unsatisfied conditions, perform the following steps: <ol style="list-style-type: none">1. Start the OneCommand Manager application.2. Select View > Group Adapters by Virtual Port from the Main menu.3. In the discovery-tree, select the FC function on which you want to create a virtual port. The Virtual Ports tab should contain a list of unsatisfied conditions (if any) that are preventing a virtual port from being created.4. If there are no unsatisfied conditions, yet vPort creation still fails, contact Broadcom technical support.
The port is not ready.	The controls in the New Virtual Port area of the Virtual Port tab are replaced by a list of reasons why vPorts cannot be created. The reasons can be one or more of the following: <ul style="list-style-type: none">■ The driver NPIV parameter is disabled.■ SLI-3 is not being used by a port.■ The adapter port is out of resources for additional virtual ports.■ The port is not connected to a fabric.■ The fabric switch does not support virtual ports.■ The fabric switch is out of resources for additional virtual ports.■ The port link state is down.

Appendix A: License Notices

A.1 Oracle Binary Code License Agreement for the Java SE Platform Products and JavaFX

ORACLE AMERICA, INC. ("ORACLE"), FOR AND ON BEHALF OF ITSELF AND ITS SUBSIDIARIES AND AFFILIATES UNDER COMMON CONTROL, IS WILLING TO LICENSE THE SOFTWARE TO YOU ONLY UPON THE CONDITION THAT YOU ACCEPT ALL OF THE TERMS CONTAINED IN THIS BINARY CODE LICENSE AGREEMENT AND SUPPLEMENTAL LICENSE TERMS (COLLECTIVELY "AGREEMENT"). PLEASE READ THE AGREEMENT CAREFULLY. BY SELECTING THE "ACCEPT LICENSE AGREEMENT" (OR THE EQUIVALENT) BUTTON AND/OR BY USING THE SOFTWARE YOU ACKNOWLEDGE THAT YOU HAVE READ THE TERMS AND AGREE TO THEM. IF YOU ARE AGREEING TO THESE TERMS ON BEHALF OF A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT THAT YOU HAVE THE LEGAL AUTHORITY TO BIND THE LEGAL ENTITY TO THESE TERMS. IF YOU DO NOT HAVE SUCH AUTHORITY, OR IF YOU DO NOT WISH TO BE BOUND BY THE TERMS, THEN SELECT THE "DECLINE LICENSE AGREEMENT" (OR THE EQUIVALENT) BUTTON AND YOU MUST NOT USE THE SOFTWARE ON THIS SITE OR ANY OTHER MEDIA ON WHICH THE SOFTWARE IS CONTAINED.

1. DEFINITIONS. "Software" means the software identified above in binary form that you selected for download, install or use (in the version You selected for download, install or use) from Oracle or its authorized licensees and/or those portions of such software produced by jlink as output using a Program's code, when such output is in unmodified form in combination, and for sole use with, that Program, as well as any other machine readable materials (including, but not limited to, libraries, source files, header files, and data files), any updates or error corrections provided by Oracle, and any user manuals, programming guides and other documentation provided to you by Oracle under this Agreement. The Java Linker (jlink) is available with Java 9 and later versions. "General Purpose Desktop Computers and Servers" means computers, including desktop and laptop computers, or servers, used for general computing functions under end user control (such as but not specifically limited to email, general purpose Internet browsing, and office suite productivity tools). The use of Software in systems and solutions that provide dedicated functionality (other than as mentioned above) or designed for use in embedded or function-specific software applications, for example but not limited to: Software embedded in or bundled with industrial control systems, wireless mobile telephones, wireless handheld devices, kiosks, TV/STB, Blu-ray Disc devices, telematics and network control switching equipment, printers and storage management systems, and other related systems are excluded from this definition and not licensed under this Agreement. "Programs" means (a) Java technology applets and applications intended to run on the Java Platform, Standard Edition platform on Java-enabled General Purpose Desktop Computers and Servers; and (b) JavaFX technology applications intended to run on the JavaFX Runtime on JavaFX-enabled General Purpose Desktop Computers and Servers. "Java SE LIUM" means the Licensing Information User Manual – Oracle Java SE and Oracle Java Embedded Products Document accessible at <http://www.oracle.com/technetwork/java/javase/documentation/index.html>. "Commercial Features" means those features that are identified as such in the Java SE LIUM under the "Description of Product Editions and Permitted Features" section.

2. LICENSE TO USE. Subject to the terms and conditions of this Agreement including, but not limited to, the Java Technology Restrictions of the Supplemental License Terms, Oracle grants you a non-exclusive, non-transferable, limited license without license fees to reproduce and use internally the Software complete and unmodified for the sole purpose of running Programs. THE LICENSE SET FORTH IN THIS SECTION 2 DOES NOT EXTEND TO THE COMMERCIAL FEATURES. YOUR RIGHTS AND OBLIGATIONS RELATED TO THE COMMERCIAL FEATURES ARE AS SET FORTH IN THE SUPPLEMENTAL TERMS ALONG WITH ADDITIONAL LICENSES FOR DEVELOPERS AND PUBLISHERS.

3. RESTRICTIONS. Software is copyrighted. Title to Software and all associated intellectual property rights is retained by Oracle and/or its licensors. Unless enforcement is prohibited by applicable law, you may not modify, decompile, or reverse engineer Software. You acknowledge that the Software is developed for general use in a variety of information management applications; it is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use the Software in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle disclaims any express or implied warranty of fitness for such uses. No right, title or interest in or to any trademark, service mark, logo or trade name of Oracle or its licensors is granted under this Agreement. Additional restrictions for developers and/or publishers licenses are set forth in the Supplemental License Terms.

4. DISCLAIMER OF WARRANTY. THE SOFTWARE IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND. ORACLE FURTHER DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT.

5. LIMITATION OF LIABILITY. IN NO EVENT SHALL ORACLE BE LIABLE FOR ANY INDIRECT, INCIDENTAL, SPECIAL, PUNITIVE OR CONSEQUENTIAL DAMAGES, OR DAMAGES FOR LOSS OF PROFITS, REVENUE, DATA OR DATA USE, INCURRED BY YOU OR ANY THIRD PARTY, WHETHER IN AN ACTION IN CONTRACT OR TORT, EVEN IF ORACLE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. ORACLE'S ENTIRE LIABILITY FOR DAMAGES HEREUNDER SHALL IN NO EVENT EXCEED ONE THOUSAND DOLLARS (U.S. \$1,000).

6. TERMINATION. This Agreement is effective until terminated. You may terminate this Agreement at any time by destroying all copies of Software. This Agreement will terminate immediately without notice from Oracle if you fail to comply with any provision of this Agreement. Either party may terminate this Agreement immediately should any Software become, or in either party's opinion be likely to become, the subject of a claim of infringement of any intellectual property right. Upon termination, you must destroy all copies of Software.

7. EXPORT REGULATIONS. You agree that U.S. export control laws and other applicable export and import laws govern your use of the Software, including technical data; additional information can be found on Oracle's Global Trade Compliance web site (<http://www.oracle.com/us/products/export>). You agree that neither the Software nor any direct product thereof will be exported, directly, or indirectly, in violation of these laws, or will be used for any purpose prohibited by these laws including, without limitation, nuclear, chemical, or biological weapons proliferation.

8. TRADEMARKS AND LOGOS. You acknowledge and agree as between you and Oracle that Oracle owns the ORACLE and JAVA trademarks and all ORACLE- and JAVA-related trademarks, service marks, logos and other brand designations ("Oracle Marks"), and you agree to comply with the Third Party Usage Guidelines for Oracle Trademarks currently located at <http://www.oracle.com/us/legal/third-party-trademarks/index.html>. Any use you make of the Oracle Marks inures to Oracle's benefit.

9. U.S. GOVERNMENT LICENSE RIGHTS. If Software is being acquired by or on behalf of the U.S. Government or by a U.S. Government prime contractor or subcontractor (at any tier), then the Government's rights in Software and accompanying documentation shall be only those set forth in this Agreement.

10. **GOVERNING LAW.** This agreement is governed by the substantive and procedural laws of California. You and Oracle agree to submit to the exclusive jurisdiction of, and venue in, the courts of San Francisco, or Santa Clara counties in California in any dispute arising out of or relating to this agreement.

11. **SEVERABILITY.** If any provision of this Agreement is held to be unenforceable, this Agreement will remain in effect with the provision omitted, unless omission would frustrate the intent of the parties, in which case this Agreement will immediately terminate.

12. **INTEGRATION.** This Agreement is the entire agreement between you and Oracle relating to its subject matter. It supersedes all prior or contemporaneous oral or written communications, proposals, representations and warranties and prevails over any conflicting or additional terms of any quote, order, acknowledgment, or other communication between the parties relating to its subject matter during the term of this Agreement. No modification of this Agreement will be binding, unless in writing and signed by an authorized representative of each party.

SUPPLEMENTAL LICENSE TERMS

These Supplemental License Terms add to or modify the terms of the Binary Code License Agreement. Capitalized terms not defined in these Supplemental Terms shall have the same meanings ascribed to them in the Binary Code License Agreement. These Supplemental Terms shall supersede any inconsistent or conflicting terms in the Binary Code License Agreement, or in any license contained within the Software.

A. COMMERCIAL FEATURES. You may not use the Commercial Features for running Programs, Java applets or applications in your internal business operations or for any commercial or production purpose, or for any purpose other than as set forth in Sections B, C, D and E of these Supplemental Terms. If You want to use the Commercial Features for any purpose other than as permitted in this Agreement, You must obtain a separate license from Oracle.

B. SOFTWARE INTERNAL USE FOR DEVELOPMENT LICENSE GRANT. Subject to the terms and conditions of this Agreement and restrictions and exceptions set forth in the Java SE LIUM incorporated herein by reference, including, but not limited to the Java Technology Restrictions of these Supplemental Terms, Oracle grants you a non-exclusive, non-transferable, limited license without fees to reproduce internally and use internally the Software complete and unmodified for the purpose of designing, developing, and testing your Programs.

C. LICENSE TO DISTRIBUTE SOFTWARE. Subject to the terms and conditions of this Agreement and restrictions and exceptions set forth in the Java SE LIUM, including, but not limited to the Java Technology Restrictions and Limitations on Redistribution of these Supplemental Terms, Oracle grants you a non-exclusive, non-transferable, limited license without fees to reproduce and distribute the Software, provided that (i) you distribute the Software complete and unmodified and only bundled as part of, and for the sole purpose of running, your Programs, (ii) the Programs add significant and primary functionality to the Software, (iii) you do not distribute additional software intended to replace any component(s) of the Software, (iv) you do not remove or alter any proprietary legends or notices contained in the Software, (v) you only distribute the Software subject to a license agreement that: (a) is a complete, unmodified reproduction of this Agreement; or (b) protects Oracle's interests consistent with the terms contained in this Agreement and that includes the notice set forth in Section H, and (vi) you agree to defend and indemnify Oracle and its licensors from and against any damages, costs, liabilities, settlement amounts and/or expenses (including attorneys' fees) incurred in connection with any claim, lawsuit or action by any third party that arises or results from the use or distribution of any and all Programs and/or Software. The license set forth in this Section C does not extend to the Software identified in Section G.

D. LICENSE TO DISTRIBUTE REDISTRIBUTABLES. Subject to the terms and conditions of this Agreement and restrictions and exceptions set forth in the Java SE LIUM, including but not limited to the Java Technology Restrictions and Limitations on Redistribution of these Supplemental Terms, Oracle grants you a non-exclusive, non-transferable, limited license without fees to reproduce and distribute those files specifically identified as redistributable in the Java SE LIUM ("Redistributables") provided that: (i) you distribute the Redistributables complete and unmodified, and only bundled as part of Programs, (ii) the Programs add significant and primary functionality to the Redistributables, (iii) you do not distribute additional software intended to supersede any component(s) of the Redistributables (unless otherwise specified in the applicable Java SE LIUM), (iv) you do not remove or alter any proprietary legends or notices contained in or on the Redistributables, (v) you only distribute the Redistributables pursuant to a license agreement that: (a) is a complete, unmodified reproduction of this Agreement; or (b) protects Oracle's interests consistent with the terms contained in the Agreement and includes the notice set forth in Section H, (vi) you agree to defend and indemnify Oracle and its licensors from and against any damages, costs, liabilities, settlement amounts and/or expenses (including attorneys' fees) incurred in connection with any claim, lawsuit or action by any third party that arises or results from the use or distribution of any and all Programs and/or Software. The license set forth in this Section D does not extend to the Software identified in Section G.

E. DISTRIBUTION BY PUBLISHERS. This section pertains to your distribution of the Java™ SE Development Kit Software ("JDK") with your printed book or magazine (as those terms are commonly used in the industry) relating to Java technology ("Publication"). Subject to and conditioned upon your compliance with the restrictions and obligations contained in the Agreement, Oracle hereby grants to you a non-exclusive, nontransferable limited right to reproduce complete and unmodified copies of the JDK on electronic media (the "Media") for the sole purpose of inclusion and distribution with your Publication(s), subject to the following terms: (i) You may not distribute the JDK on a stand-alone basis; it must be distributed with your Publication(s); (ii) You are responsible for downloading the JDK from the applicable Oracle web site; (iii) You must refer to the JDK as Java™ SE Development Kit; (iv) The JDK must be reproduced in its entirety and without any modification whatsoever (including with respect to all proprietary notices) and distributed with your Publication subject to a license agreement that is a complete, unmodified reproduction of this Agreement; (v) The Media label shall include the following information: "Copyright [YEAR], Oracle America, Inc. All rights reserved. Use is subject to license terms. ORACLE and JAVA trademarks and all ORACLE- and JAVA-related trademarks, service marks, logos and other brand designations are trademarks or registered trademarks of Oracle in the U.S. and other countries." [YEAR] is the year of Oracle's release of the Software; the year information can typically be found in the Software's "About" box or screen. This information must be placed on the Media label in such a manner as to only apply to the JDK; (vi) You must clearly identify the JDK as Oracle's product on the Media holder or Media label, and you may not state or imply that Oracle is responsible for any third-party software contained on the Media; (vii) You may not include any third party software on the Media which is intended to be a replacement or substitute for the JDK; (viii) You agree to defend and indemnify Oracle and its licensors from and against any damages, costs, liabilities, settlement amounts and/or expenses (including attorneys' fees) incurred in connection with any claim, lawsuit or action by any third party that arises or results from the use or distribution of the JDK and/or the Publication; ; and (ix) You shall provide Oracle with a written notice for each Publication; such notice shall include the following information: (1) title of Publication, (2) author(s), (3) date of Publication, and (4) ISBN or ISSN numbers. Such notice shall be sent to Oracle America, Inc., 500 Oracle Parkway, Redwood Shores, California 94065 U.S.A., Attention: General Counsel.

F. JAVA TECHNOLOGY RESTRICTIONS. You may not create, modify, or change the behavior of, or authorize your licensees to create, modify, or change the behavior of, classes, interfaces, or subpackages that are in any way identified as "java", "javax", "sun", "oracle" or similar convention as specified by Oracle in any naming convention designation.

G. LIMITATIONS ON REDISTRIBUTION. You may not redistribute or otherwise transfer patches, bug fixes or updates made available by Oracle through Oracle Premier Support, including those made available under Oracle's Java SE Support program.

H. COMMERCIAL FEATURES NOTICE. For purpose of complying with Supplemental Term Section C.(v)(b) and D.(v)(b), your license agreement shall include the following notice, where the notice is displayed in a manner that anyone using the Software will see the notice:

Use of the Commercial Features for any commercial or production purpose requires a separate license from Oracle. "Commercial Features" means those features that are identified as such in the Licensing Information User Manual – Oracle Java SE and Oracle Java Embedded Products Document, accessible at <http://www.oracle.com/technetwork/java/javase/documentation/index.html>, under the "Description of Product Editions and Permitted Features" section.

I. SOURCE CODE. Software may contain source code that, unless expressly licensed for other purposes, is provided solely for reference purposes pursuant to the terms of this Agreement. Source code may not be redistributed unless expressly provided for in this Agreement.

J. THIRD PARTY CODE. Additional copyright notices and license terms applicable to portions of the Software are set forth in the Java SE LIUM accessible at <http://www.oracle.com/technetwork/java/javase/documentation/index.html>. In addition to any terms and conditions of any third party opensource/freeware license identified in the Java SE LIUM, the disclaimer of warranty and limitation of liability provisions in paragraphs 4 and 5 of the Binary Code License Agreement shall apply to all Software in this distribution.

K. TERMINATION FOR INFRINGEMENT. Either party may terminate this Agreement immediately should any Software become, or in either party's opinion be likely to become, the subject of a claim of infringement of any intellectual property right.

L. INSTALLATION AND AUTO-UPDATE. The Software's installation and auto-update processes transmit a limited amount of data to Oracle (or its service provider) about those specific processes to help Oracle understand and optimize them. Oracle does not associate the data with personally identifiable information. You can find more information about the data Oracle collects as a result of your Software download at <http://www.oracle.com/technetwork/java/javase/documentation/index.html>.

For inquiries please contact: Oracle America, Inc., 500 Oracle Parkway,
Redwood Shores, California 94065, USA

Last updated 21 September 2017

A.2 OpenPegasus Licensing Notice

Licensed to The Open Group (TOG) under one or more contributor license agreements. Refer to the OpenPegasusNOTICE.txt file distributed with this work for additional information regarding copyright ownership. Each contributor licenses this file to you under the OpenPegasus Open Source License; you may not use this file except in compliance with the License.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

A.3 OpenSSL Notice

This is a copy of the current LICENSE file inside the CVS repository.

LICENSE ISSUES

=====

The OpenSSL toolkit stays under a dual license, i.e. both the conditions of the OpenSSL License and the original SSLeay license apply to the toolkit. See below for the actual license texts. Actually both licenses are BSD-style Open Source licenses. In case of any license issues related to OpenSSL please contact openssl-core@openssl.org.

OpenSSL License

```
/* =====
```

```
* Copyright (c) 1998-2011 The OpenSSL Project. All rights reserved.
```

```
*
```

```
* Redistribution and use in source and binary forms, with or without  
* modification, are permitted provided that the following conditions  
* are met:
```

```
*
```

```
* 1. Redistributions of source code must retain the above copyright  
* notice, this list of conditions and the following disclaimer.
```

```
*
```

```
* 2. Redistributions in binary form must reproduce the above copyright  
* notice, this list of conditions and the following disclaimer in  
* the documentation and/or other materials provided with the  
* distribution.
```

```
*
```

```
* 3. All advertising materials mentioning features or use of this
```

```
* software must display the following acknowledgment:
* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit. (http://www.openssl.org/)"
*
* 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to
* endorse or promote products derived from this software without
* prior written permission. For written permission, please contact
* openssl-core@openssl.org.
*
* 5. Products derived from this software may not be called "OpenSSL"
* nor may "OpenSSL" appear in their names without prior written
* permission of the OpenSSL Project.
*
* 6. Redistributions of any form whatsoever must retain the following
* acknowledgment:
* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit (http://www.openssl.org/)"
*
* THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY
* EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
* PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR
* ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
* SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT
* NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
* LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
* STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
* ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
* OF THE POSSIBILITY OF SUCH DAMAGE.
* =====
*
* This product includes cryptographic software written by Eric Young
* (eay@cryptsoft.com). This product includes software written by Tim
* Hudson (tjh@cryptsoft.com).
*
*/

Original SSLeay License
-----

/* Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)
* All rights reserved.
*
* This package is an SSL implementation written
* by Eric Young (eay@cryptsoft.com).
* The implementation was written so as to conform with Netscapes SSL.
*
* This library is free for commercial and non-commercial use as long as
* the following conditions are aheared to. The following conditions
* apply to all code found in this distribution, be it the RC4, RSA,
* lhash, DES, etc., code; not just the SSL code. The SSL documentation
* included with this distribution is covered by the same copyright terms
* except that the holder is Tim Hudson (tjh@cryptsoft.com).
*
* Copyright remains Eric Young's, and as such any Copyright notices in
* the code are not to be removed.
* If this package is used in a product, Eric Young should be given attribution
* as the author of the parts of the library used.
* This can be in the form of a textual message at program startup or
* in documentation (online or textual) provided with the package.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
* 1. Redistributions of source code must retain the copyright
* notice, this list of conditions and the following disclaimer.
* 2. Redistributions in binary form must reproduce the above copyright
* notice, this list of conditions and the following disclaimer in the
* documentation and/or other materials provided with the distribution.
* 3. All advertising materials mentioning features or use of this software
```

```
*      must display the following acknowledgement:
*      "This product includes cryptographic software written by
*      Eric Young (eay@cryptsoft.com)"
*      The word 'cryptographic' can be left out if the routines from the library
*      being used are not cryptographic related :-).
* 4. If you include any Windows specific code (or a derivative thereof) from
*      the apps directory (application code) you must include an acknowledgement:
*      "This product includes software written by Tim Hudson (tjh@cryptsoft.com)"
*
* THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND
* ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
* ARE DISCLAIMED.  IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE
* FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
* DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
* OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
* LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
* OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
* SUCH DAMAGE.
*
* The licence and distribution terms for any publically available version or
* derivative of this code cannot be changed.  i.e. this code cannot simply be
* copied and put under another distribution licence
* [including the GNU Public Licence.]
*/
```

A.4 Secure Hash Algorithm (SHA-1) Notice

```
/*
* Written by Aaron D. Gifford <me@aarongifford.com>
*
* Copyright 1998, 2000 Aaron D. Gifford.  All rights reserved.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
* 1. Redistributions of source code must retain the above copyright
*    notice, this list of conditions and the following disclaimer.
* 2. Redistributions in binary form must reproduce the above copyright
*    notice, this list of conditions and the following disclaimer in the
*    documentation and/or other materials provided with the distribution.
* 3. Neither the name of the copyright holder nor the names of contributors
*    may be used to endorse or promote products derived from this software
*    without specific prior written permission.
*
* THIS SOFTWARE IS PROVIDED BY THE AUTHOR(S) AND CONTRIBUTORS ``AS IS'' AND
* ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
* ARE DISCLAIMED.  IN NO EVENT SHALL THE AUTHOR(S) OR CONTRIBUTORS BE LIABLE
* FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
* DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
* OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
* LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
* OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
* SUCH DAMAGE.
*/
```

A.5 Secure Hash Algorithm (SHA-2) Notice

```
/*
* AUTHOR:Aaron D. Gifford - http://www.aarongifford.com/
*
* Copyright (c) 2000-2001, Aaron D. Gifford
* All rights reserved.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
```

```
* are met:
* 1. Redistributions of source code must retain the above copyright
*    notice, this list of conditions and the following disclaimer.
* 2. Redistributions in binary form must reproduce the above copyright
*    notice, this list of conditions and the following disclaimer in the
*    documentation and/or other materials provided with the distribution.
* 3. Neither the name of the copyright holder nor the names of contributors
*    may be used to endorse or promote products derived from this software
*    without specific prior written permission.
*
* THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTOR(S) ``AS IS'' AND
* ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
* ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTOR(S) BE LIABLE
* FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
* DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
* OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
* LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
* OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
* SUCH DAMAGE.
*
*/
```

A.6 VI Java SDK

Copyright (c) 2012 Steve Jin. All Rights Reserved.

Copyright (c) 2008 VMware, Inc. All Rights Reserved.

Copyright (c) 2009 Altor Networks. All Rights Reserved.

Copyright (c) 2009 NetApp. All Rights Reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

– Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

– Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

– Neither the name of VMware, Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED.

IN NO EVENT SHALL VMWARE, INC. OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

