

# **OneCommand Manager Application Version 10.6**

## **User Manual**

Copyright © 2003-2015 Emulex. All rights reserved worldwide. No part of this document may be reproduced by any means or translated to any electronic medium without the prior written consent of Emulex.

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

Emulex, the Emulex logo, Emulex BladeEngine, Emulex InSpeed, Emulex LightPulse, Emulex OneCommand, Emulex OneConnect, and Emulex SLI are registered trademarks, and Emulex Advanced-8, Emulex Connect, Emulex CrossLink, Emulex Engine, Emulex Edge, Emulex ExpressLane, Emulex GreenState, Emulex OneCore, Emulex Pilot, Emulex SURF, Emulex Universal Multi-Channel, Emulex vEngine, Emulex Virtual Fabric, Emulex Virtual Network Exceleration, Emulex vPath, and Emulex vScale, are trademarks, of Emulex. All other brand or product names referenced herein are trademarks or registered trademarks of their respective companies or organizations.

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

Emulex, 3333 Susan Street  
Costa Mesa, CA 92626

# Table of Contents

List of Figures .....	11
List of Tables .....	15
1. Introduction .....	16
Abbreviations .....	16
2. Installing and Uninstalling OneCommand Manager Application Components .....	21
Installing the OneCommand Manager Application .....	21
In Windows .....	21
Attended Installation in Windows .....	21
Unattended Installation in Windows .....	22
In Linux .....	24
Attended Installation in Linux .....	24
Unattended Installation in Linux .....	26
Updating an Installation in Linux .....	27
In Solaris .....	27
In VMware .....	29
Installing the OneCommand Manager Application Web Launch Interface .....	29
Requirements .....	30
In Windows .....	30
In Linux and Solaris .....	31
Installation .....	31
Uninstalling the OneCommand Manager Application .....	32
Uninstalling the OneCommand Manager Application Web Launch Interface Only .....	33
3. Starting and Stopping the OneCommand Manager Application .....	34
In Windows .....	34
In Linux and Solaris .....	34
Starting the OneCommand Manager Application Web Launch Interface .....	35
Managing Files When Running the OneCommand Manager Application Web Launch Interface .....	36
4. Using the OneCommand Manager Application .....	37
The OneCommand Manager Application Window Element Definitions .....	37
Menu Bar .....	38
Toolbar .....	38
Toolbar Buttons .....	38

Discovery-Tree .....	40
Discovery-Tree Icons .....	40
Expanding or Collapsing the Discovery-Tree View .....	42
Property Tabs.....	42
Status Bar .....	42
<b>Using OneCommand Manager Secure Management.....</b>	<b>42</b>
OneCommand Manager Secure Management Configuration Requirements .....	44
<b>Changing Management and Read-Only Mode.....</b>	<b>44</b>
Management Host.....	45
<b>Using CIM (Windows Only) .....</b>	<b>46</b>
<b>5. Configuring Discovery .....</b>	<b>48</b>
Discovery Using the TCP/IP Access Protocol .....	48
Hosts File.....	48
Manually Editing the Hosts File .....	49
Copying the File.....	50
Adding a Single Host.....	50
Adding a Range of Hosts (IPv4 Only) .....	51
Removing Hosts .....	53
Configuring Discovery and Default CIM Credentials .....	53
Viewing Discovery Information .....	55
Discovery Information Field Definitions .....	56
<b>6. Managing Hosts .....</b>	<b>57</b>
Viewing Host Information.....	57
Host Information Field Definitions .....	58
Viewing Host Grouping Information .....	58
Host Group Information Field Definitions .....	59
Grouping Hosts.....	59
Managing Host Groups.....	60
Host Group Management Field Definitions.....	61
Creating a Host Group.....	62
Deleting a Host Group.....	62
Adding a Host to a Host Group.....	62
Removing a Host from a Host Group .....	63
Restoring a Host Group .....	63
Restoring all Host Groups.....	63
Exporting Host Grouping Configurations .....	63
Searching for Hosts in the Discovery-Tree .....	64

<b>7. Managing Adapters and Ports .....</b>	<b>65</b>
<b>FC Adapters .....</b>	<b>65</b>
Viewing FC Adapter Information.....	65
FC Adapter Information Field Definitions .....	65
Viewing FC Port Information .....	66
FC Port Information Field Definitions.....	67
<b>OneConnect Adapters .....</b>	<b>68</b>
Viewing OneConnect Adapter Information.....	68
OneConnect Adapter Information Field Definitions .....	69
Viewing Channel Management Information (OCe11102 Adapters Only) .....	71
Configuring UMC Channel Management.....	72
Viewing the Channel Management Tab for vNIC1 (Lenovo System X Only) .....	74
Viewing the Channel Management Tab for UFP (Lenovo System X Only).....	75
Viewing ASIC Information.....	76
ASIC Information Field Definitions .....	77
Viewing OneConnect Multi-ASIC Adapter Information.....	78
OneConnect Multi-ASIC Adapter Information Field Definitions .....	79
Showing and Installing Licenses for OneConnect OCe11102 Adapters .....	80
Showing Licenses.....	80
Installing Licenses .....	81
Changing Personalities on OneConnect OCe11102 Adapters .....	82
Viewing OneConnect Firmware Information .....	83
Firmware Tab Field Definitions.....	84
Viewing OneConnect Physical Port Information.....	85
Enabling and Disabling OneConnect Physical Ports .....	85
Setting Port Speed and DAC Cable Length (OneConnect OCe11102 and OCe14000-Series Adapters Only) .....	86
Viewing PHY Data (OneConnect 10GBASE-T series Adapters Only).....	87
PHY Data Field Definitions .....	88
Viewing OneConnect Transceiver Information .....	88
Transceiver Data Field Definitions .....	89
<b>OCe14000-Series Adapter Configuration .....</b>	<b>90</b>
OCe14000-series Adapter Configuration Tab Field Definitions .....	91
Configuring Single Personalities .....	92
Custom Configurations .....	94
Mixed Mode Configuration .....	95
Concurrent Mode.....	96
Configuring RoCE in a Custom View.....	97
UMC Configuration View .....	98
Default UMC Settings.....	99

Lenovo System X Multi-Channel Configuration View .....	100
vNIC Configuration.....	101
SIMode Configuration.....	102
UFP Configuration .....	104
Link Re-Configurability.....	106
Adapter Configuration Buttons with Link Re-Configurability.....	109
Dell NPar Configuration View .....	110
Configuring RoCE on NPar Adapters.....	111
Dell NPar Enabled.....	112
Dell NPar and NParEP Mode Enabled.....	114
NPar Bandwidth Settings.....	116
<b>8. Managing Protocols .....</b>	<b>117</b>
<b>NIC.....</b>	<b>117</b>
Viewing NIC Function Information.....	117
NIC Port Information Field Definitions.....	118
Viewing NIC VPD Information .....	119
VPD Table Definitions .....	119
Enabling and Disabling SR-IOV on NIC Functions (on OCe11000-series adapters) .....	120
Enabling and Disabling VEPA on NIC SR-IOV Functions (OCe14000-series adapters only) .....	121
Guest Operating System Discovery and Management from the Base Host Operating System .....	121
Port Information Field Definitions.....	123
Running the OneCommand Manager Application on a Guest Operating System.....	123
Viewing NIC PCI Registers .....	124
<b>FC/FCoE .....</b>	<b>124</b>
Viewing FCoE Function Information.....	124
FCoE Port Information Field Definitions.....	125
Viewing FC/FCoE Statistics.....	126
Port Statistics Field Definitions.....	127
Viewing FC/FCoE Virtual Port Information.....	129
Discovery Information for Virtual Ports Field Definitions.....	129
Creating and Deleting FC/FCoE Virtual Ports.....	129
Creating Virtual Ports .....	129
Deleting Virtual Ports .....	132
Viewing FC/FCoE Fabric Information .....	133
Discovery Information Field Definitions .....	134
Viewing FC Port Transceiver Information .....	134
Transceiver Data Field Definitions .....	135

Viewing FC/FCoE VPD Information .....	136
VPD Table Definitions .....	136
Viewing FC Maintenance Information .....	137
Maintenance Tab Field Definitions .....	137
Viewing FCoE Maintenance Information .....	138
Maintenance Tab Field Definitions .....	139
Viewing FC/FCoE Target Information .....	140
Target Information Field Definitions .....	140
Viewing FC/FCoE LUN Information .....	141
FC/FCoE LUN Information Field Definitions .....	142
Viewing FC/FCoE Target Mapping (Windows and Solaris) .....	142
Target Mapping Field Definitions .....	143
Viewing FC/FCoE Target Mapping (Linux and VMware ESXi) .....	144
Target Mapping Field Definitions .....	144
Using Automapping and Persistent Binding (Windows) .....	145
Changing Automapping Settings .....	145
Adding a Persistent Binding .....	146
Binding a Target that Does Not Appear in the Persistent Binding Table .....	147
Masking and Unmasking LUNs (Windows) .....	148
Managing FC/FCoE ExpressLane LUNS (LPe15000 and LPe16000 HBAs) .....	149
Changing FC/FCoE WWPN and WNNN .....	151
Changing FC/FCoE Names .....	154
Resetting FC/FCoE Functions .....	154
Configuring FC/FCoE Driver Parameters .....	155
Activation Requirements .....	156
The Host Driver Parameters Tab .....	156
Setting Driver Parameters .....	157
Setting Driver Parameters for a Single FC/FCoE Function .....	158
Restoring All Parameters to Their Earlier Values .....	159
Resetting All Default Values .....	159
Setting an Adapter Parameter Value to the Host Adapter Parameter Value .....	159
Saving Adapter Driver Parameters to a File .....	159
Setting Driver Parameters for All FC/FCoE Functions in a Host .....	160
Changing Non-dynamic Parameter Values (Linux 8.2) .....	161
Creating a Batch Mode Driver Parameters File .....	161
Assigning Batch Mode Parameters .....	162
Configuring FCoE Initialization Protocol (FIP) .....	163
FIP Dialog Box Field Descriptions .....	164
Configuring Boot from an FC/FCoE SAN .....	164
Boot Types .....	165
Boot Device Parameters .....	165

Configuring Advanced Settings (Boot from SAN) .....	169
x86 Boot Advanced Adapter Settings Dialog Box .....	170
EFIBoot Advanced Adapter Settings Dialog Box .....	172
OpenBoot Advanced Adapter Settings Dialog Box.....	173
Using FC-SP DHCHAP Authentication (Windows, Linux 8.2, and Solaris) .....	173
Linux Considerations .....	174
Enabling Authentication .....	174
fcauthd Daemon.....	174
DHCHAP Tab .....	176
Changing Authentication Configuration .....	177
Viewing the Error and Event Log .....	178
<b>iSCSI.....</b>	<b>179</b>
Viewing iSCSI Port Information .....	179
iSCSI Port Information Field Definitions .....	179
Viewing iSCSI Network Information .....	180
iSCSI Network Information Tab Field Definitions .....	180
Modifying Network Settings .....	181
Advanced TCP/IP Configuration .....	182
Viewing iSCSI PCI Registers.....	183
Viewing iSCSI VPD Information.....	184
VPD Table Definitions .....	184
Viewing iSCSI Statistics .....	185
iSCSI Statistics Field Definitions .....	185
Configuring iSCSI Target Discovery.....	187
Adding Target Portals .....	188
Removing a Target Portal.....	188
Configuring iSNS for iSCSI Target Discovery .....	188
Logging into Targets.....	190
Manually Adding an iSCSI Target.....	192
Removing Targets.....	192
Viewing Target Sessions.....	192
Logging out of Target Sessions.....	193
Viewing iSCSI Target Information .....	196
Target Information Field Definitions.....	196
Viewing iSCSI LUN Information.....	197
iSCSI LUN Information Field Definitions .....	197
Configuring iSCSI Initiator Login Options .....	198
Initiator Login Options Tab Field Definitions .....	198
<b>9. Updating Adapter Firmware .....</b>	<b>202</b>
<b>Updating Firmware for a Single Adapter .....</b>	<b>202</b>



Updating Firmware for Multiple Adapters.....	203
<b>10. DCB Configuration .....</b>	<b>206</b>
Configuring DCB Parameters for NIC Only Adapter Ports .....	206
DCB Tab Field Definitions.....	206
Configure DCB Dialog Box Field Definitions.....	208
Configuring DCB Parameters for NIC + iSCSI Adapter Ports .....	210
DCB Tab Field Definitions.....	211
Configure DCB Dialog Box Field Definitions.....	213
Configuring DCB Parameters for NIC + FCoE Adapter Ports .....	215
DCB Tab Field Definitions.....	216
Configure DCB Dialog Box Field Definitions.....	218
Configuring DCB Parameters for NIC + FCoE + iSCSI Adapter Ports.....	220
DCB Tab Field Definitions.....	221
Configure DCB Dialog Box Field Definitions.....	223
Configuring DCB Parameters for RoCE Adapter Ports .....	225
DCB Tab Field Definitions.....	226
Configure DCB Dialog Box Field Definitions.....	229
<b>11. Configuring HP Flex Adapters .....</b>	<b>232</b>
<b>12. Exporting SAN Information .....</b>	<b>234</b>
Creating a SAN Report.....	234
<b>13. Diagnostics .....</b>	<b>235</b>
LightPulse FC Diagnostics.....	235
Viewing Flash Contents, PCI Registers, and Wakeup Information.....	236
Viewing Flash Contents .....	236
Viewing Overlay Details.....	237
Viewing the PCI Registers.....	237
Running a Quick Test.....	237
Running a Power-On Self Test (POST) .....	238
Using Beaconing.....	238
Running D_Port Tests.....	238
D_Port Window Descriptions .....	240
Creating Diagnostic Dumps .....	241
Running Advanced Diagnostic Tests.....	242
Running Loopback Tests .....	243
Loopback Test Combinations.....	243
Error Action.....	244
Test Cycles.....	244
Test Pattern .....	244

Test Status .....	244
Running End-to-End (ECHO) Tests .....	245
Saving the Log File .....	246
<b>OneConnect Diagnostics .....</b>	<b>247</b>
OneConnect Loopback Test Combinations .....	248
FCoE End to End Echo Test .....	249
Error Action .....	249
Test Cycles .....	249
Test Pattern .....	249
Test Status .....	249
Using Beacons .....	250
Running TDR Tests (10GBASE-T Adapters Only) .....	251
Saving the Log File .....	251
Creating Diagnostic Dumps .....	252
<b>14. Troubleshooting .....</b>	<b>255</b>
General Situations .....	255
Emulex Driver for Linux and OneCommand Manager Application Situations .....	257
Emulex Driver for Solaris and OneCommand Manager Application Situations .....	260
VPorts and OneCommand Manager Application Situations .....	260
<b>Appendix A. License Notices .....</b>	<b>261</b>
Secure Hash Algorithm (SHA-1) Notice .....	261
OpenPegasus Licensing Notice .....	262
OpenSSL Notice .....	262

# List of Figures

<b>Figure 2-1</b>	Management Mode Dialog Box.....	22
<b>Figure 2-2</b>	OneCommand Manager Application Web Launch Uninstallation Screen .....	33
<b>Figure 4-1</b>	OneCommand Manager Application Window .....	37
<b>Figure 4-2</b>	Toolbar.....	38
<b>Figure 4-3</b>	Discovery-Tree.....	40
<b>Figure 4-4</b>	Management Mode Dialog Box.....	46
<b>Figure 5-1</b>	Discovery Information .....	48
<b>Figure 5-2</b>	Add Remote TCP/IP Host Dialog Box .....	50
<b>Figure 5-3</b>	Add Remote TCP/IP Host Dialog Box with CIM Credentials .....	51
<b>Figure 5-4</b>	Add Range of TCP/IP Hosts Dialog Box.....	52
<b>Figure 5-5</b>	Discovery Settings Dialog Box .....	54
<b>Figure 5-6</b>	Discovery Information (Host View Selected) .....	55
<b>Figure 6-1</b>	Host Information Tab .....	57
<b>Figure 6-2</b>	Host Group Information Tab.....	59
<b>Figure 6-3</b>	Host Group Management Dialog Box .....	61
<b>Figure 6-4</b>	Create New Host Group Dialog Box.....	62
<b>Figure 6-5</b>	Host Group Management Warning Dialog Box.....	62
<b>Figure 7-1</b>	FC Adapter Information Tab.....	65
<b>Figure 7-2</b>	FC Port Information Tab .....	66
<b>Figure 7-3</b>	Adapter Information Tab.....	69
<b>Figure 7-4</b>	Example of the Channel Management Tab (UMC management depicted) .....	72
<b>Figure 7-5</b>	UMC Channel Management Tab .....	73
<b>Figure 7-6</b>	Channel Management Tab for vNIC1 (Lenovo System X Only) .....	74
<b>Figure 7-7</b>	Channel Management Tab for UFP (Lenovo System X Only) .....	75
<b>Figure 7-8</b>	ASIC Information Tab .....	77
<b>Figure 7-9</b>	OneConnect Multi-ASIC Adapter Information .....	79
<b>Figure 7-10</b>	OneConnect OCe11102 Adapter Information Tab .....	80
<b>Figure 7-11</b>	Licensed Features Window .....	81
<b>Figure 7-12</b>	Install Feature Licenses Dialog Box.....	81
<b>Figure 7-13</b>	OneConnect OCe11102 Adapter Information Tab .....	83
<b>Figure 7-14</b>	OneConnect Firmware Tab .....	84
<b>Figure 7-15</b>	Physical Port Info Tab (OCe11102 Adapter Port Selected) .....	85
<b>Figure 7-16</b>	Change Port Speed Dialog box (Force mode/1Gb Ethernet speed selected).....	86
<b>Figure 7-17</b>	PHY Data Tab .....	87
<b>Figure 7-18</b>	OneConnect Transceiver Data Tab .....	89
<b>Figure 7-19</b>	OCe14000-series Adapter Configuration Tab (FCoE selected).....	91
<b>Figure 7-20</b>	SR-IOV Warning Message .....	92
<b>Figure 7-21</b>	Current Configuration Details example .....	92

<b>Figure 7-22</b>	Single Personality View (NIC selected) .....	93
<b>Figure 7-23</b>	Custom View .....	94
<b>Figure 7-24</b>	Mixed Mode Protocol pull-down menu .....	95
<b>Figure 7-25</b>	Concurrent Storage Configuration View .....	96
<b>Figure 7-26</b>	Concurrent Storage Configuration Choices for the Third Function .....	96
<b>Figure 7-27</b>	Custom NIC+RoCE Configuration .....	97
<b>Figure 7-28</b>	UMC View Adapter Configuration Tab (UMC, 2 ports, 8 functions/port, concurrent storage) .....	98
<b>Figure 7-29</b>	Multi-channel View (showing multi-channel type drop-down) .....	100
<b>Figure 7-30</b>	Multi-channel vNIC View (mix mode storage).....	101
<b>Figure 7-31</b>	Multi-channel SIMode View .....	102
<b>Figure 7-32</b>	2-Port Lenovo System X Adapter Configured for SIMODE (with 8 functions per port).....	103
<b>Figure 7-33</b>	Multi-channel UFP View (concurrent mode storage) .....	104
<b>Figure 7-34</b>	Adapter Configuration Tab Before Changing Link Configuration.....	106
<b>Figure 7-35</b>	Adapter Configuration Tab After Selecting a New Link Configuration .....	107
<b>Figure 7-36</b>	Adapter Configuration Tab After Changing Port Configuration for New Link Configuration .....	108
<b>Figure 7-37</b>	Adapter Configuration Tab After Reboot with Single Port Configuration .....	109
<b>Figure 7-38</b>	Active Adapter Configuration Window Showing Active Link Configuration .....	110
<b>Figure 7-39</b>	Adapter Configuration Tab for NPar Adapters (NPar disabled) .....	111
<b>Figure 7-40</b>	Adapter Configuration Tab with NPar Enabled (NParEP mode disabled) .....	112
<b>Figure 7-41</b>	Adapter Configuration Tab with NPar and NIC SR-IOV Enabled .....	113
<b>Figure 7-42</b>	Adapter Configuration Tab with NPar Disabled with a Single NIC+RoCE Port.....	114
<b>Figure 7-43</b>	Adapter Configuration Tab with NPar and NParEP Mode Enabled (2-port configuration) .....	115
<b>Figure 7-44</b>	Adapter Configuration Tab with NPar and NParEP Mode Enabled (4-port configuration) .....	115
<b>Figure 8-1</b>	NIC Port Information Tab .....	117
<b>Figure 8-2</b>	NIC VPD Tab.....	119
<b>Figure 8-3</b>	Port Information Tab, SR-IOV Enabled .....	121
<b>Figure 8-4</b>	OneCommand Manager Application Running on the Base Host Operating System after Discovering the Guest Host.....	122
<b>Figure 8-5</b>	VF Selected Showing the Port Information Tab for the Discovered NIC in the Guest Operating System.....	122
<b>Figure 8-6</b>	NIC PCI Registers Tab.....	124
<b>Figure 8-7</b>	FCoE Port Information Tab .....	125
<b>Figure 8-8</b>	Statistics Tab .....	127
<b>Figure 8-9</b>	Discovery Information Tab Displaying Virtual Ports Information.....	129
<b>Figure 8-10</b>	Virtual Ports Tab .....	130
<b>Figure 8-11</b>	Virtual Port Tab .....	132

<b>Figure 8-12</b>	Delete Virtual Port Warning .....	133
<b>Figure 8-13</b>	Fabric Discovery Information .....	134
<b>Figure 8-14</b>	FC Transceiver Data Tab .....	135
<b>Figure 8-15</b>	FC/FCoE VPD Tab .....	136
<b>Figure 8-16</b>	FC Maintenance Tab .....	137
<b>Figure 8-17</b>	FCoE Maintenance Tab .....	139
<b>Figure 8-18</b>	Target Information Tab .....	140
<b>Figure 8-19</b>	FC/FCoE LUN Information Tab .....	141
<b>Figure 8-20</b>	Target Mapping Tab.....	143
<b>Figure 8-21</b>	Target Mapping Tab.....	146
<b>Figure 8-22</b>	Add Persistent Binding Dialog Box .....	147
<b>Figure 8-23</b>	Bind New Target Dialog Box .....	147
<b>Figure 8-24</b>	LUN Masking Tab .....	148
<b>Figure 8-25</b>	Enabling an ExpressLane LUN .....	150
<b>Figure 8-26</b>	Maintenance Tab .....	152
<b>Figure 8-27</b>	Warning About Changing WWN .....	152
<b>Figure 8-28</b>	Change World Wide Name Configuration Dialog Box .....	153
<b>Figure 8-29</b>	Reset Warning .....	155
<b>Figure 8-30</b>	Host Driver Parameters Tab .....	156
<b>Figure 8-31</b>	Driver Parameters Tab - Adapter Selected.....	158
<b>Figure 8-32</b>	Host Driver Parameters Tab - Host Selected.....	160
<b>Figure 8-33</b>	Save Driver Parameters Dialog Box.....	162
<b>Figure 8-34</b>	Batch Driver Parameters Update Dialog Box .....	162
<b>Figure 8-35</b>	FIP Tab for FCoE Adapters .....	164
<b>Figure 8-36</b>	Boot from SAN Configuration Dialog Box .....	167
<b>Figure 8-37</b>	Select Boot Device Window (for x86 or EFIBoot) .....	169
<b>Figure 8-38</b>	x86 Boot Advanced Adapter Settings Dialog Box .....	170
<b>Figure 8-39</b>	EFIBoot Advanced Adapter Settings Dialog Box.....	172
<b>Figure 8-40</b>	OpenBoot Advanced Settings Dialog Box .....	173
<b>Figure 8-41</b>	DHCHAP Tab .....	176
<b>Figure 8-42</b>	iSCSI Port Information Tab.....	179
<b>Figure 8-43</b>	iSCSI Network Information Tab.....	180
<b>Figure 8-44</b>	Advanced TCP/IP Configuration Dialog Box .....	182
<b>Figure 8-45</b>	iSCSI PCI Registers Tab .....	183
<b>Figure 8-46</b>	iSCSI VPD Tab .....	184
<b>Figure 8-47</b>	iSCSI Statistics Tab.....	185
<b>Figure 8-48</b>	iSCSI Target Discovery Tab .....	187
<b>Figure 8-49</b>	iSCSI iSNS Tab.....	189
<b>Figure 8-50</b>	Add iSNS Server Dialog Box .....	189
<b>Figure 8-51</b>	Target Login Dialog Box .....	191
<b>Figure 8-52</b>	Target Sessions Dialog Box.....	193
<b>Figure 8-53</b>	iSCSI Target Information Tab.....	196

<b>Figure 8-54</b>	iSCSI LUN Information Tab .....	197
<b>Figure 8-55</b>	iSCSI Initiator Login Options Tab.....	198
<b>Figure 8-56</b>	Initiator Default Login Options Window .....	200
<b>Figure 9-1</b>	Firmware Download Dialog Box .....	203
<b>Figure 9-2</b>	Batch Firmware Download Dialog Box, Selecting Adapters to Update .....	204
<b>Figure 9-3</b>	Batch Firmware Download Dialog Box, Download Complete .....	205
<b>Figure 10-1</b>	DCB Tab for NIC Only Adapter Ports (NIC Only adapter port selected).....	206
<b>Figure 10-2</b>	Configure DCB Dialog Box for NIC Only Adapter Ports.....	208
<b>Figure 10-3</b>	DCB Tab for NIC + iSCSI Adapter Ports (NIC + iSCSI adapter port selected).....	211
<b>Figure 10-4</b>	Configure DCB Dialog Box for NIC + iSCSI Adapter Ports (DCBX enabled) .....	213
<b>Figure 10-5</b>	DCB Tab (NIC + FCoE adapter port selected) .....	216
<b>Figure 10-6</b>	Configure DCB Dialog Box for NIC + FCoE Adapter Ports (DCBX enabled) .....	218
<b>Figure 10-7</b>	DCB Tab (NIC + FCoE + iSCSI adapter port selected) .....	221
<b>Figure 10-8</b>	Configure DCB Dialog Box for NIC + FCoE + iSCSI Adapter Ports (DCBX enabled) .....	223
<b>Figure 10-9</b>	DCB Tab for RoCE Adapter Ports.....	226
<b>Figure 10-10</b>	Configure DCB Dialog Box for RoCE Adapter Ports (DCBX enabled) .....	228
<b>Figure 11-1</b>	HP Flex Adapter Configuration tab (Flex enabled) .....	232
<b>Figure 11-2</b>	Flex Adapter After Applying an iSCSI Single Personality Change .....	233
<b>Figure 13-1</b>	PCI Registers and Flash Contents of the Diagnostics Tab .....	236
<b>Figure 13-2</b>	Overlay Detail Window .....	237
<b>Figure 13-3</b>	Quick Test Warning .....	237
<b>Figure 13-4</b>	Diagnostics Tab for LPe16000-Series Adapters (D Port Tests... button depicted) .....	239
<b>Figure 13-5</b>	D_Port Tests Window .....	240
<b>Figure 13-6</b>	Diagnostic Dump Dialog Box .....	241
<b>Figure 13-7</b>	Diagnostic Dump File Transfer Dialog Box.....	242
<b>Figure 13-8</b>	Diagnostic Test Setup.....	243
<b>Figure 13-9</b>	Run Diagnostic Tests Warning .....	245
<b>Figure 13-10</b>	Advanced Diagnostic Tests Warning Window for External Loopback.....	245
<b>Figure 13-11</b>	Select Echo Test Target Window .....	246
<b>Figure 13-12</b>	Advanced Diagnostic Tests Warning Window .....	246
<b>Figure 13-13</b>	Example of a DiagTest.log Window.....	247
<b>Figure 13-14</b>	Diagnostics Tab (10GBASE-T adapter selected) .....	248
<b>Figure 13-15</b>	Run Diagnostic Tests Warning .....	250
<b>Figure 13-16</b>	Advanced Diagnostic Tests Warning Window for External Loopback.....	250
<b>Figure 13-17</b>	Example of a DiagTest.log Window.....	252
<b>Figure 13-18</b>	Diagnostic Dump Dialog Box .....	253
<b>Figure 13-19</b>	Diagnostic Dump File Transfer Dialog Box.....	254

## List of Tables

<b>Table 4-1</b>	Secure Management User Privileges.....	43
<b>Table 4-2</b>	Active Commands: machines on same domain .....	43
<b>Table 4-3</b>	Active Commands: machines on different domain.....	43
<b>Table 4-4</b>	Passive Commands: machines on any domain .....	43
<b>Table 7-1</b>	vNIC Maximum number of functions per port.....	102
<b>Table 7-2</b>	SIMode Maximum number of functions per port.....	103
<b>Table 7-3</b>	UFP maximum number of functions per port .....	105
<b>Table 14-1</b>	General Situations .....	255
<b>Table 14-2</b>	Emulex Driver for Linux and OneCommand Manager Application Situations .....	257
<b>Table 14-3</b>	Emulex Driver for Solaris and OneCommand Manager Application Situations .....	260
<b>Table 14-4</b>	VPorts and OneCommand Manager Application Situations.....	260

# 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 IP addresses of remote machines. The OneCommand Manager application contains a graphical user interface (GUI) and a command line interface (CLI). Refer to the *OneCommand Manager Command Line Interface User Manual* for information about installing and using the CLI.

**Note:** Screenshots in this manual 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, see the Emulex website.

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

## Abbreviations

AL_PA	Arbitrated Loop Physical Address
AP	access point
API	application programming interface
ARI	alternate routing-ID interpretation
ARP	Address Resolution Protocol
ASIC	application-specific integrated circuit
BIOS	basic input/output system
BOFM	Blade Open Firmware Management Protocol
CND	congestion notification domain
CHAP	Challenge Handshake Authentication Protocol
CIM	Common Interface Model
CIMOM	CIM object manager
CLI	command line interface
CLP	Command Line Protocol
CNA	Converged Network Adapter
CNPV	Congestion Notification Priority Value
CRC	cyclic redundancy check



CSV	comma-separated values
DAC	direct-attach copper
D_ID	destination identifier
DCB	Data Center Bridging
DCBX	Data Center Bridging Capabilities Exchange
DH	Diffie-Hellman
DHCHAP	Diffie-Hellman Challenge Handshake Authentication Protocol
DHCP	Dynamic Host Control Protocol
DID	device ID
DMA	direct memory access
EDD	Enhanced Disk Drive
EFI	Extensible Firmware Interface
EMI	electro-magnetic interference
ETO	extended timeout
F_BSY	FC port busy
FC	Fibre Channel
FC-SP	Fibre Channel Security Protocol
FCoE	Fibre Channel over Ethernet
FIP	FCoE Initialization Protocol
FLOGI	Fabric login
GFC	gigabit Fibre Channel
GUI	Graphical User Interface
HBA	host bus adapter
HTTP	Hypertext Transfer Protocol
I/O	input/output
iBFT	iSCSI boot firmware table
IEEE	Institute of Electrical and Electronics Engineers
IIS	Internet Information Services
IP	Internet Protocol
IPL	initial program load
iSCSI	internet Small Computer System Interface
ISID	initiator session identifier
iSNS	Internet Storage Name Service
IQN	iSCSI Qualified Name
JBOD	just a bunch of disks
JEDEC	Joint Electron Device Engineering Council
JNLP	Java Network Launching Protocol

JRE	Java Runtime Environment
KVM	kernel-based virtual machine
LAN	local area network
LDAP	Lightweight Directory Access Protocol
LDPC	low-density parity check
LDTO	link down timeout
LED	light-emitting diode
LIP	Loop Initialization Primitive
LLDP	Link Layer Discovery Protocol
LPVID	logical port VLAN ID
LUN	logical unit number
MAC	Media Access Control
MDI	Medium Dependent Interface
MIME	Multipurpose Internet Mail Extension
MSE	Mean Squared Error
MTU	maximum transmission unit
NCSI	Network Communication Services Interface
NIC	network interface card
NOS	network operating system
NPar	NIC partitioning
NPIV	N_Port ID Virtualization
NVP	normal velocity of propagation
NVRAM	non-volatile random access memory
OS	operating system
OUI	Organizationally Unique Identifier
PAM	pluggable authentication modules
PCI	Peripheral Component Interconnect (interface)
PDU	protocol data unit
PF	physical function
PFC	priority flow control
PG	priority group
PGID	priority group ID
PLL	phase-locked loop
POST	power-on self test
PXE	Preboot Execution Environment
QoS	quality of service
QCN	quantized congestion notification

RAID	redundant array of independent disks
RDMA	remote direct memory access
RHEL	Red Hat Enterprise Linux
RMAPI	Remote Management Application Programming Interface
RoCE	RDMA over Converged Ethernet
RPI	remote port indicator
SAN	storage area network
SCSI	Small Computer System Interface
SFP	small-form factor pluggable
SFS	Software Foundation Software
SIMode	Lenovo System X term for UMC. Previously known as vNIC2
SLES	SUSE Linux Enterprise Server
SMB	Server Message Block
SNMP	Simple Network Management Protocol
SNR	signal-to-noise ratio
SR-IOV	single root I/O virtualization
TCP	Transmission Control Protocol
TCP/IP	TCP over Internet Protocol
TDR	time-domain reflectometer
TSIH	iSCSI target session identifier handle
Tx	transmit
UEFI	Unified Extensible Firmware Interface
UFP	Universal Fabric Port
ULP	Upper Layer Protocol
UMC	Universal Multi-Channel
VEPA	virtual Ethernet port aggregator
VF	virtual function
VLAN	virtual local area network
VLAN ID	VLAN identifier
VM	virtual machine
vNIC	virtual network interface card
VPD	vital product data
vPort	virtual port
WLAN	wireless LAN
WWN	World Wide Name
WWNN	World Wide Node Name
WWPN	World Wide Port Name

XRI                      extensible resource indicator

## 2. Installing and Uninstalling OneCommand Manager Application Components

This section describes installing and uninstalling the OneCommand Manager application.

### Installing the OneCommand Manager Application

**Note:** If the OneCommand Vision application was previously installed on the system, you are prompted to remove it before installing the OneCommand Manager application.

#### In Windows

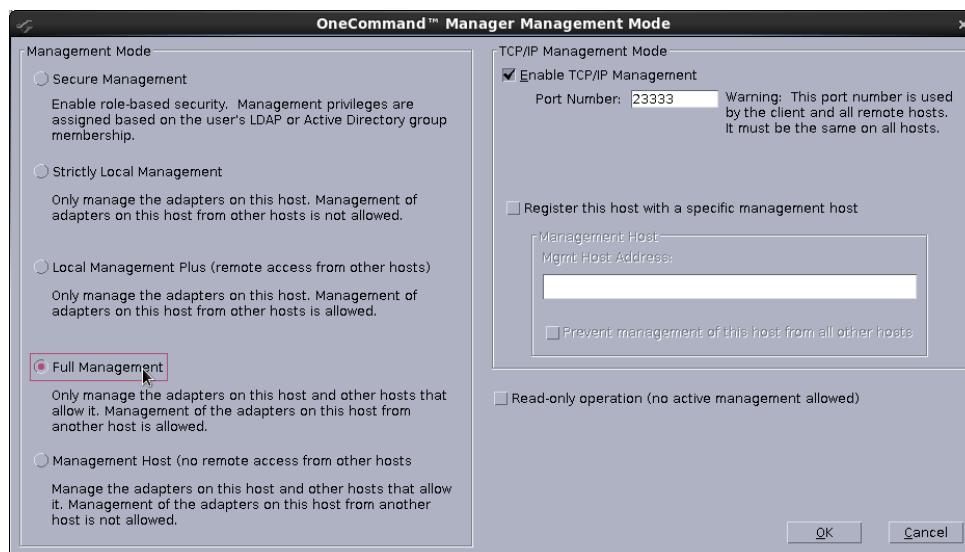
The OneCommand Manager application can be installed two ways:

- Attended installation using the GUI.
- Unattended installation using the command line.

#### Attended Installation in Windows

To install the OneCommand Manager application in Windows:

1. From the Emulex website, download the x64 or x86 OneCommand Manager Enterprise Kit installation file.
2. Navigate to the directory to which you downloaded the file.
3. Double-click the elxocm<version>.exe. The Emulex OCManager Enterprise window appears. Click **Next**. The Installation Options window appears.
4. Check the components that you want to install and click Install. After installing the OneCommand Manager application files, the Management Mode dialog box appears (Figure 2-1 on page 22).



**Figure 2-1** Management Mode Dialog Box

5. The Management Mode dialog box (Figure 2-1 on page 22) enables you to select Secure Management to assign the desired user privileges, or you can choose one of the other management modes. See “Using OneCommand Manager Secure Management” on page 42 or “Changing Management and Read-Only Mode” on page 44 for more information. Choose the management type you want and click **OK**.
6. Check or uncheck the Enable TCP/IP Management checkbox 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 Emulex).
7. The Installation Completed window appears when the installation is finished. Click **Finish**. A shortcut is added to the Start menu. You do not need to reboot the system.

## Unattended Installation in Windows

To install the OneCommand Manager application in Windows:

1. From the Emulex website, download the x64 or x86 OneCommand Manager Enterprise Kit installation file to your system.
2. Activate the kit with switch /q or /q2.
  - The /q switch displays progress reports.
  - The /q2 switch does not display progress reports.
3. You can enable Secure Management Mode by adding the sec=1 argument or disable it by sec=0. If the sec argument is not entered, Secure Management is disabled by default. See “Using OneCommand Manager Secure Management” on page 42 for more information.

To enable Secure Management, at the command prompt type

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

To disable Secure Management, at the command prompt type

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

**Note:** The Management Mode defaults for unattended installation are:

mmode = 2 (Local Plus Management Mode)

achange = 1

4. You can select a management mode by adding the mmode argument and the ability to change that management mode by adding the achange argument with selected values as in the example below. See “Changing Management and Read-Only Mode” on page 44 for more information.

**Note:** If you enabled Secure Management in step 3, you cannot enter an mmode value. Doing so results in a 'conflicting parameters' error.

For example, at the command prompt type

```
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 Management Mode to change
- 1 – Allow 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.
- excl – 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 switch. This option is only accepted 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 switch is specified, the valid values are:
  - 0 – Remotely managed by other hosts.
  - 1 – Remotely managed by Management Host ONLY.

- **Mtcp** – This optional switch allows you to enable or disable remote management and to specify the TCP/IP port number over which management occurs. If this switch is not specified, the default TCP/IP port number 2333 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.

## In Linux

**Note:** The OneCommand Manager application GUI is not supported on Citrix XenServer; however, the OneCommand Manager application CLI is supported. Refer to the *OneCommand Manager application CLI User Manual* for Citrix instructions.

The following must be installed before you can install the OneCommand Manager application:

- The appropriate driver version for your operating system:
  - Linux driver version 8.2.0.33.3p or later (for RHEL 5 operating systems).
  - Linux driver version 8.3.5.x or later (for RHEL 6 and SLES 11 SP1 operating systems).
  - Linux driver version 10.6.x.x for RHEL 7 operating systems.
  - Linux driver version 10.6.x.x for SLES 12 operating systems.

**Note:** The RHEL 6 Enterprise kit requires the installation of the `libstdc++-5.so` library. This library is available through the `compat-libstdc++-33-3.2.3-68.<arch>.rpm` or later. The PPC and x86\_64 builds require the 64-bit version, which is installed in `/usr/lib64`. The i386 build requires the 32-bit version, which is installed in `/usr/lib`.

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

## Attended Installation in Linux

To install the OneCommand Manager application, or update an existing installation, in Linux:

1. Log on as 'root'.
2. Download the utilities from the Emulex website.
3. Copy the OneCommand `elxocm-<Platform>-<AppsRev>.tgz` file to a directory on the install machine.
4. Change to the directory to which you copied the tar file.



## 5. Untar the file.

- For RHEL 5 and RHEL 6 type

```
tar zxvf elxocm-rhel5-rhel6-<apps_ver>-<rel>.tgz
```

- For RHEL 7 type

```
tar zxvf elxocm-rhel7-<apps_ver>-<rel>.tgz
```

- For SLES 11 type

```
tar zxvf elxocm-sles11-<apps_ver>-<rel>.tgz
```

- For SLES 12 type

```
tar zxvf elxocm-sles12-<apps_ver>-<rel>.tgz
```

## 6. Change to the elxocm directory created in step 3.

- For RHEL 5 and RHEL 6 type

```
cd elxocm-rhel5-rhel6-<apps_ver>-<rel>
```

- For RHEL 7 type

```
cd elxocm-rhel7-<apps_ver>-<rel>
```

- For SLES 11 type

```
cd elxocm-sles10-sles11-<apps_ver>-<rel>
```

- For SLES 12 type

```
cd elxocm-sles12-<apps_ver>-<rel>
```

**Note:** Prior to installation, OneCommand Manager application groups must be configured on the LDAP network or the local host machine for Secure Management operation. See “OneCommand Manager Secure Management Configuration Requirements” on page 44 for configuration instructions.

## 7. Run the install script. Type

```
./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.

## 9. When prompted, enter the type of management 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 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 certain operations such as resetting adapters, updating an adapter's firmware and changing adapter driver properties and bindings. It only affects the local OneCommand Manager application interface. These operations can still be performed using remote management. Enter **<y>** for yes to allow users to perform these operations, enter **<n>** for no if read-only mode is desired.
11. You are prompted about allowing users to change the management mode after installation. Enter **<y>** for yes, or **<n>** for no.

## Unattended Installation in Linux

For unattended or silent installation of the OneCommand Manager application for Linux, installation settings are defined using the install 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 "OneCommand Manager Secure Management Configuration Requirements" on page 44 for configuration instructions.

To view the options for unattended installation, type

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

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

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

**Note:** The Management Mode default for unattended installation is Local Management Plus.

## 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 “Attended Installation in Linux” on page 24 or “Unattended Installation in Linux” on page 26 for instructions.

## In Solaris

The following must be installed for the utilities to function properly:

- The Solaris FC/FCoE inbox driver version emlxs 2.80.8.0 or later or the out-of-box driver version elxfc 2.85.xx.xx must be installed for FC/FCoE management.
- The NIC inbox driver version oce 4.4.173.9.3S or later or the out-of-box driver version elxnic 4.1.xx.xx must be installed for CNA management.

**Note:** If Emulex CNAs are installed on the system, the NIC driver must be installed and reporting all NIC ports. Otherwise, the OneCommand Manager application cannot manage the CNAs.

To install the OneCommand Manager application in Solaris:

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

```
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 “OneCommand Manager Secure Management Configuration Requirements” on page 44 for configuration instructions.

4. Run the install script to begin installation. If the HBAnyware utility, OneCommand Manager Core or OneCommand Manager Enterprise applications or the Solaris driver utilities are already present on the system, the install script attempts to remove them first:

```
./install
```

5. 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 7. You cannot choose a management mode as described in step 6.

6. When prompted, enter the type of management 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 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.

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 certain operations such as resetting adapters, updating an adapter's firmware and changing adapter driver properties and bindings. It only affects the local OneCommand Manager application interface. These operations can still be performed using remote management. Enter **<y>** for yes to allow users to perform these operations, 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 **<y>** for yes, or **<n>** for no.

## 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.

VMware ESXi 5.1 comes with an inbox Emulex CIM Provider. The inbox Emulex CIM Provider enables you to manage Emulex LightPulse adapters, but not Emulex CNAs. For ESXi 5.5 and 6.0, to manage Emulex CNAs, you must install the out-of-box Emulex CIM Provider.

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

For the best real-time management of Emulex adapters in VMware ESXi environments, the OneCommand Manager application for VMware vCenter is highly recommended. For more information go to <http://www.emulex.com/products/software-solutions.html>.

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

1. Copy the CIM Provider zip file to `/var/log/vmware`.
2. Log into the VMware hypervisor host, and execute the following command all on one line:  

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

## Installing the OneCommand Manager Application Web Launch Interface

### Notes

- The OneCommand Manager application Web Launch Interface is not supported on VMware ESXi Servers.
- The Web Launch Interface is only supported on server systems; client systems are not supported.
- OneCommand Manager Secure Management mode is not supported for OneCommand Manager application Web Launch clients.

In addition to the driver and OneCommand Manager application, the following prerequisites must be met before you install the Web Launch Interface.

## Requirements

The following section details the requirements for installing the OneCommand Manager application Web Launch interface for the Windows, Linux, and Solaris operating systems.

### In Windows

- Microsoft IIS Server must be installed. See the Microsoft website for information on downloads and installation.
- The Windows Firewall may be enabled by default. If it is, you must add and enable three exceptions: HTTP port, java.exe, and rmiregistry.exe.

**Note:** Allowing programs or ports through the firewall may increase the security risks. Use at your own discretion.

To enable the HTTP port:

1. Click **Add Port...** The Add a Port dialog box is displayed.
2. On the Add a Port dialog box, type `HTTP` as the Name and `80` as the Port Number.
3. Leave **TCP** enabled and click **OK**.

To enable the java.exe program:

1. Click **Add Program...** The Add a Program dialog box is displayed.
2. Click **Browse...**
3. Specify java.exe located in the OneCommand Manager JRE installation path. For example:  
`C:\Program Files\Emulex\util\JRE\bin\java.exe.`
4. Click **OK**.

To enable the rmiregistry.exe program:

1. Click **Add Program...** The Add a Program dialog box is displayed.
2. Click **Browse...** and specify the rmiregistry.exe located in the OneCommand Manager JRE installation path. For example:  
`C:\Program Files\Emulex\util\JRE\bin\rmiregistry.exe.`
3. Click **OK**.
4. Click **OK** to apply the new firewall settings.

To add the MIME type:

1. Launch Server Manager.
2. Expand Roles.
3. Under Roles, expand Web Server (IIS).
4. Under Web Server (IIS), click **Internet Information Services (IIS) Manager**.
5. In the right pane, find your server name under "Start Page" and click on it.
6. Double-click **MIME Types** listed under IIS group.
7. A MIME Types page appears. Under "Actions", click **Add...** A popup dialog box appears.

8. Add "jnlp" (without quotes) to the File name extension field.
9. Add "application/x-java-jnlp-file" (without quotes) to the MIME type field.
10. Click **OK**.

## In Linux and Solaris

**Note:** Apache Web server must be installed and running on the server that is hosting the Web Launch Service software.

The server on which you are installing the Web Launch Service package requires:

- An HTTP server configured to handle the JNLP MIME file type. The following MIME file type/file extension must be added to your server configuration:  
MIME type: application/x-java-jnlp-file  
File Extension: jnlp
- The HTTP server must be running.

The client on which you are running the browser requires:

- Java must be installed. The specific requirements are:
  - Sun 32-bit Java 6.0 or later for Intel based systems (x86)
  - 32-bit Java 6.0 or later for x86-64 systems
  - 32-bit Java 6.0 or later for SLES 10 (ppc64)
  - 32-bit Java 6.0 or later for SLES 11 and 12

Refer to the appropriate vendor documentation for detailed instructions about configuring MIME types, configuring and starting the HTTP server and installing the JRE. See /opt/ELXocm/README\_WEBLAUNCH.txt (Solaris) or /usr/sbin/ocmanager/README\_WEBLAUNCH.txt (Linux) for more setup information.

## Installation

To install the OneCommand Manager application Web Launch Interface:

- In Windows  
Click **Programs>Emulex >OCManager WebLaunch Install**. Web Launch installation begins.
- In Solaris and Linux

**Note:** Citrix XenServer does not support the Web Launch Interface.

1. Log on as 'root'.
2. Navigate to the OneCommand Manager directory.
  - Solaris:  
cd /opt/ELXocm
  - Linux:  
cd /usr/sbin/ocmanager

3. Run the install script. Type:  

```
./wsinstall
```
4. When prompted, enter the web server's document root directory. For example:
  - Solaris:  

```
/var/apache/htdocs
```
  - Linux:  

```
/srv/www/htdocs
```
  - or-  

```
/var/www/html
```
5. Confirm that the IP address of the host is the IP address that the web server uses. Answer **<y>** or **<n>** as appropriate. If you answer **<n>**, you are prompted for the IP address you want to use.
6. When asked if your web server is listening on the normal default HTTP port (80), answer **<y>** or **<n>** as appropriate. If you answer **<n>**, you are prompted for the port you want to use.

After you have entered the necessary information, you are notified when the installation of the OneCommand Manager application Web Launch package is complete. The Web Launch configuration files are created and Web Launch Service automatically starts.
7. To verify the installation, locate another client, open a web browser window and enter the following URL:  

```
http://IP_ADDR:PORT_NUM/ocmanager.jnlp
```

where IP\_ADDR is the IP address of the host on which you installed the OneCommand Manager application Web Launch service, and PORT\_NUM is the TCP port number of the listening host's web server. The standard OneCommand Manager application user interface appears.

**Note:** It is not necessary to enter a port number if the standard HTTP port was chosen during configuration.

## Uninstalling the OneCommand Manager Application

To uninstall the OneCommand Manager application and OneCommand Manager application Web Launch Interface:

- In Windows
  1. Select **Start>Control Panel>Programs>Uninstall a Program**.
  2. Select **Emulex OCManager Enterprise [version]** and click **Remove** or **Uninstall**.
- In Linux
  1. Log on as 'root'.
  2. Change to the `elxocm-<platform>-<version>` installation directory.
  3. Type



```
./uninstall
```

- In Solaris
  1. Log on as 'root'.
  2. Run the OneCommand Manager application uninstall script:  

```
/opt/ELXocm/scripts/uninstall
```

## Uninstalling the OneCommand Manager Application Web Launch Interface Only

To uninstall the OneCommand Manager application Web Launch, but leave the OneCommand Manager application installed:

- In Windows
  1. Select **Start>Programs>Emulex>OCManager WebLaunch Uninstall**. The OneCommand Manager application Web Launch Uninstallation Screen appears (Figure 2-2 on page 33).




Figure 2-2 OneCommand Manager Application Web Launch Uninstallation Screen

2. The OneCommand Manager application Web Launch Interface is removed. Press any key to continue.
- In Linux and Solaris
    1. Log on as 'root'.
    2. Execute the uninstallation script.
      - Linux:  

```
/usr/sbin/ocmanager/wsuninstall
```
      - Solaris:  

```
/opt/ELXocm/wsuninstall
```

This script stops the OneCommand Manager application Web Launch Interface service daemons (if they are running) and removes all Web Launch related files from the host.

## 3. Starting and Stopping the OneCommand Manager Application

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

### 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 “Using OneCommand Manager Secure Management” on page 42 for more information.

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

### In Linux and Solaris

To start the OneCommand Manager application:

Use the ocmanager script to start the application. The script is located in the following OneCommand Manager installation directory:

- Linux - /usr/sbin/ocmanager
- Solaris - /opt/ELXocm

Example usage on Linux:

```
# /usr/sbin/ocmanager/ocmanager
```

Example usage on Solaris:

```
# /opt/ELXocm/ocmanager
```

To stop the OneCommand Manager application:

Use one of the following methods to stop the application:

- From the menu bar, select **File**, then **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” option stops only the OneCommand Manager application and associated elxdiscoveryd daemon. When the stop\_ocmanager script is run without options, the script stops the OneCommand Manager application and all associated daemons.

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

- elxhbamgrd – Remote management daemon that services requests from OneCommand Manager application clients running on remote host machines.
- elxdiscoveryd – Discovery daemon responsible for maintaining all discovery data (remote and local) for OneCommand Manager application clients running on the local machine.

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

Additionally, if the Web Launch component of OneCommand Manager application is installed, the daemon process rmiserver starts at system boot time. The start\_weblaunch script starts this daemon, while the stop\_weblaunch stops it.

## Starting the OneCommand Manager Application Web Launch Interface

After the OneCommand Manager application Web Launch Interface software is installed and the Web Launch server is initialized, you can launch the OneCommand Manager application directly with your web browser.

### Notes

- Only the OneCommand Manager application GUI is exported to the requesting client. All adapter discovery and remote management operations are performed by resources running on the remote host that served the GUI component. Therefore, the SAN view displayed by the GUI is not from the client running the GUI, but rather from the host from which this GUI was retrieved.
- Client operating systems do not support Web Launch, it is only supported on server systems.

To launch the OneCommand Manager application with your web browser:

1. Open your web browser. Linux and Solaris users must log on as 'root'.
2. If Secure Management is enabled, you are prompted for your user name and password. See "Using OneCommand Manager Secure Management" on page 42 for more information.
3. Enter the URL of the ocmanager.jnlp file. Make sure that the URL specifies a remote server which has the OneCommand Manager application Web Launch Interface software installed and running.

`http://IP_ADDR:PORT_NUM/ocmanager.jnlp`

where IP\_ADDR is the IP address of the host on which you installed the OneCommand Manager application Web Launch Service, and PORT\_NUM is the TCP port number of the listening hosts' Web server. If the port number is omitted,

the default port 80 is used. The standard OneCommand Manager application user interface is displayed.

## **Managing Files When Running the OneCommand Manager Application Web Launch Interface**

When running the OneCommand Manager application Web Launch Interface, all OneCommand Manager application files, such as log files, configuration files, and driver parameter files are located on the Web Launch server. User supplied files, such as firmware images and licenses, should be located on the Web Launch client.

## 4. Using the OneCommand Manager Application

**Note:** To properly view the OneCommand Manager application, ensure your system meets the following display requirements:

- For Windows, Linux, and Solaris systems, the display resolution must be set to 1024 x 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.

## The OneCommand Manager Application Window Element Definitions

The OneCommand Manager application window (Figure 4-1 on page 37) contains five basic components: the menu bar, the toolbar, the discovery-tree, the property tabs, and the status bar.

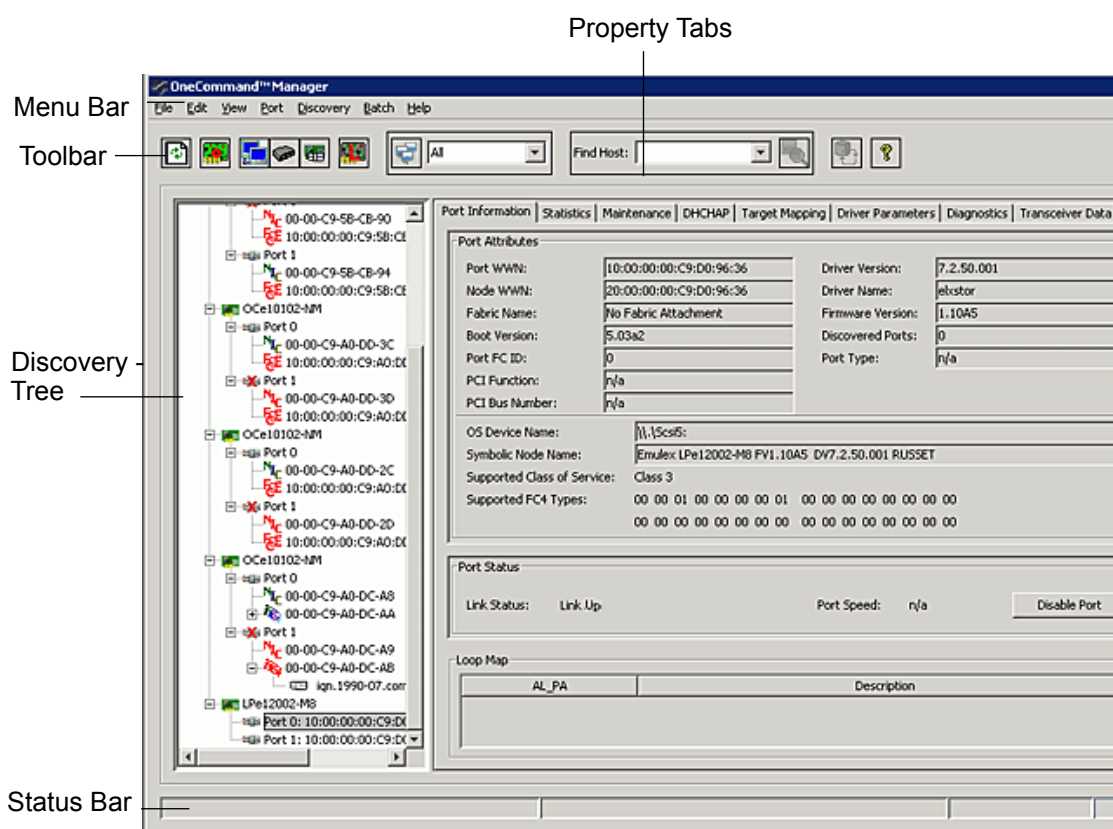


Figure 4-1 OneCommand Manager Application Window

## Notes

- The element you select in the discovery-tree determines whether a menu item or toolbar icon is active. For example, if you select the local host or other system host, the Reset Adapter item on the Adapter menu is unavailable. The Reset Adapter toolbar button is unavailable as well.
- The capabilities displayed by your local OneCommand Manager application interface matches 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.

## 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.

## Toolbar

The toolbar (Figure 4-2 on page 38) 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 4-2 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.

## Toolbar Buttons

The toolbar buttons perform the following tasks:



### Discovery Refresh button

- Initiates a discovery refresh cycle.



### Reset Adapter button

- Resets the selected adapter.

### View Buttons on the Toolbar

The 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 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 FC/FCoE 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.

**Note:** iSCSI and NIC ports are not displayed in Fabric View.



**Virtual Ports View button**

- Displays virtual ports in the SAN.

**Note:** SCSI and NIC ports are not displayed in Virtual Ports View.



**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 box. The dialog box displays version information including; RMAPI, Discovery, DFCLib, MILI Library Version (Windows) and Remote Management Agent Version (Windows). It also enables you to contact Emulex Technical Support.

## Discovery-Tree

The discovery-tree (Figure 4-3 on page 40) has icons that represent discovered hosts, adapters, ports, virtual ports, fabrics, targets, and LUNs.

Using the View menu, the OneCommand Manager application allows you to control the way iSCSI initiator and target ports are identified in the discovery-tree. The “iSCSI Names” option displays all iSCSI ports by their IQN. The “iSCSI Alias” option displays each port by its alias.

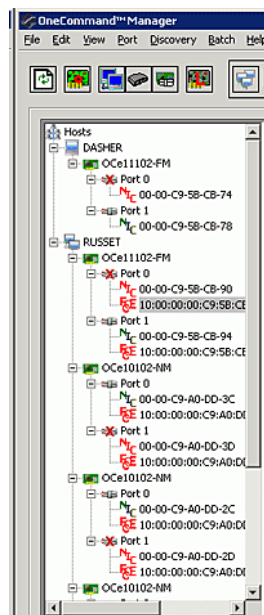


Figure 4-3 Discovery-Tree

## 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 had previously been discovered, but currently is not being 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 “Configuring Discovery and Default CIM Credentials” on page 53 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 iSCSI icon represents an iSCSI PCI function instance. iSCSI functions can support up to sixteen logical adapters, with each logical adapter appearing in the discovery-tree as a separate child node under the respective iSCSI function. A green iSCSI icon represents an iSCSI PCI function on-line instance. A black iSCSI icon represents an iSCSI PCI function port-disabled instance. A red iSCSI icon represents an iSCSI PCI function link down instance.



A green FCoE icon represents an FCoE PCI function on-line instance. A black FCoE icon represents an FCoE PCI function port-disabled instance. A red FCoE icon represents an FCoE PCI function link down instance.



The NIC icon represents a NIC-Only PCI function instance. A green icon indicates this function instance is on-line, black indicates it is disabled, and red indicates a link down instance.



The RoCE icon represents a NIC+RoCE function on the selected port.



The ASIC node icon, only displayed for OneConnect 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, "ASIC 64-12" 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 ExpressLane LUN icon represents a LUN with ExpressLane priority queueing enabled.



The Media Exchanger icon represents connections to individual media exchangers. A media exchanger is a jukebox-like 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.

## 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 four levels, the discovery-tree (Figure 4-3 on page 40) is expanded or collapsed to that level. You can choose Hosts/Fabrics (depending on the view), adapters, ports, PCI functions, and targets.

## Property Tabs

The property tabs display configuration, statistical, and status information for network elements (Figure 4-1 on page 37). 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 4-3 on page 40).

## Status Bar

The status bar is located near the bottom of the OneCommand Manager application window (Figure 4-1 on page 37). 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.

## 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 in that domain. For Linux and Solaris systems this is done using PAM.

**Note:** OneCommand Manager Secure Management is not supported on VMware hosts nor is it supported on OneCommand Manager application Web Launch clients.

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 4-1 on page 43 defines the OneCommand Manager application user groups and each group's management capabilities.

**Table 4-1** 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 may 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 is allowed or denied depending on the OneCommand Manager secure management status of the machines, and the domains to which the machines belong. The following tables list the behavior (assuming appropriate user credentials are used).

**Table 4-2** Active Commands: machines on same domain

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

**Table 4-3** Active Commands: machines on different domain

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

**Table 4-4** Passive Commands: machines on any domain

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

\* To inform you of an unsecured server that you may want to secure.

\*\* Allowed if the username and password are the same on both domains.

## OneCommand Manager Secure Management Configuration Requirements

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

1. Authentication – On Linux and Solaris systems this is done using the PAM interface and must be configured as follows:
  - For Solaris systems, place the correct setting in the “auth” section of /etc/pam.d/other file or its earlier equivalent /etc/pam.conf.
  - For Linux systems, this is the /etc/pam.d/passwd file “auth” section or equivalent.
2. 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.
3. For Solaris systems, you must use ‘useradd -G groupname’ for authentication to work. You cannot use a lowercase ‘g’.

## Changing Management and Read-Only Mode

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

During installation, a management and a read-only mode are selected. If modification of these settings after installation was selected, you can change the management mode:

- Secure Management - The setting enables roles-based security. See “Using OneCommand Manager Secure Management” on page 42 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 only allows management of adapters on this host, but management of adapters on this host from another host is possible.
- Full Management – This setting enables 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 it from being managed by other hosts.
- Enable TCP/IP Management (of/from remote host) – This setting enables you to manage remote hosts or to manage this host remotely. If 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 checkbox is

changed, a set of warning messages may appear before changes are 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, as well as make 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 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 “Management Host” on page 45 for more information.

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

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

## Management Host

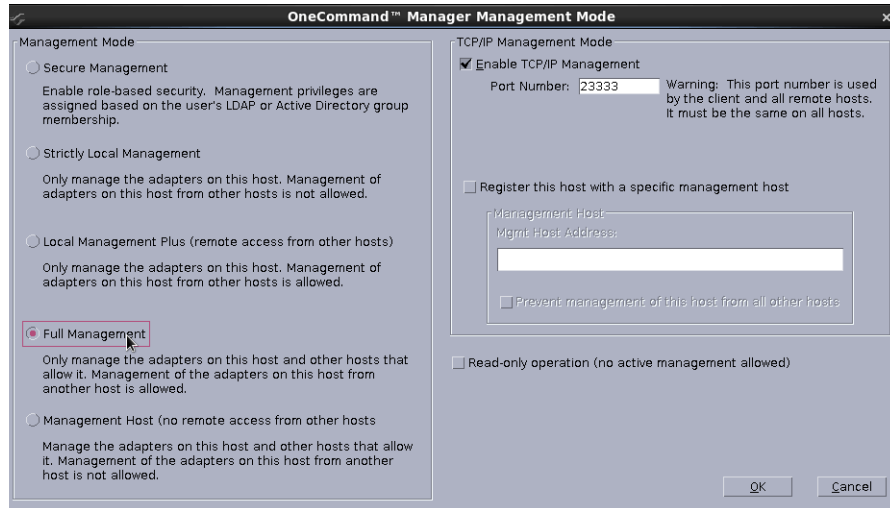
The OneCommand Manager application management host provides enhanced discovery and security by enabling a managed host to register with a management host. The management host receives these registrations when the remote host is started and updates its hosts file so the discovery server discovers the 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 only responds to requests from the management host. All requests from other hosts are rejected. This TCP/IP management security solution only allows the management host to manage the remote host.

To change management/read-only mode:

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

- In Windows
  1. From the **File** menu, select **Management Mode**. The Management Mode dialog box appears (Figure 4-4 on page 46).



**Figure 4-4** Management Mode Dialog Box

2. Choose the management type and read-only mode you want.
  3. Click **OK**.
- In Solaris
    1. Run the following script:
 

```
/opt/ELXocm/set_operating_mode
```
    2. Choose the management type and read-only mode you want.
  - In Linux
    1. Stop the OneCommand Manager application.
    2. Run the following script:
 

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

## Using CIM (Windows Only)

VMware on the Visor-based ESXi platforms 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 device and HBA management.

To manage the adapters on an ESXi host using the OneCommand Manager application, you must install the Emulex CIM Provider on the ESXi host.

ESXi comes with an inbox Emulex CIM Provider. The inbox Emulex CIM Provider enables you to manage Emulex LightPulse adapters, but not Emulex CNA adapters. To manage Emulex CNA adapters, you must install the out-of-box Emulex CIM Provider. The Emulex CIM Provider is available as an offline bundle in ESXi platforms. VMware recommends using the offline bundle to update software on VMware platforms.

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

**Note:** For VMware ESXi hosts, if advanced adapter management capabilities are required (for example, iSCSI Management and port disable), use the OneCommand Manager application for VMware vCenter. For more details, see the *OneCommand Manager for VMware vCenter User Manual*.

## 5. Configuring Discovery

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

### 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 5-1 on page 48). Remote SAN management over TCP/IP sends remote management requests using TCP/IP access protocol to remote hosts. TCP/IP access enables you to access adapters via their host IP-address or by the name of the host on which they reside.

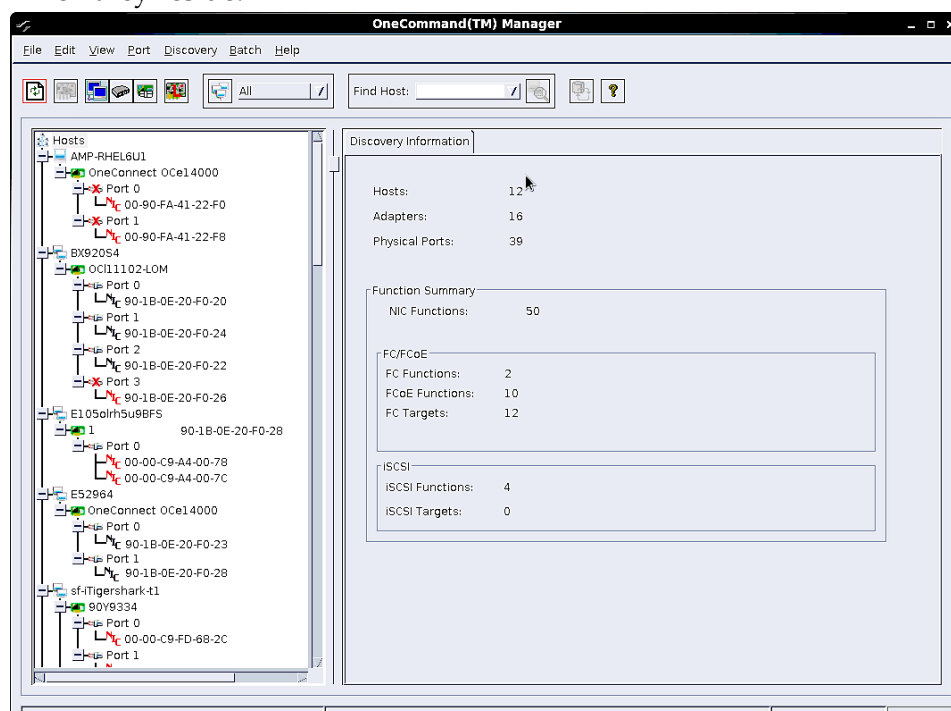


Figure 5-1 Discovery Information

**Note:** In Windows, if you are running a firewall you may 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\rmserver.exe
```

### 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 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 (Figure 5-2 on page 50). If the host is discovered, the OneCommand Manager application adds its IP address and name to the host file.
- Scanning a range of IP addresses for hosts that can be managed. This function is performed in the Add Remote Hosts window (Figure 5-2 on page 50). For each discovered host, the OneCommand Manager application adds its IP address and name to the host file.
- Removing a host from the host file using the Remove Remote Hosts window. For each removed host, the OneCommand Manager application removes its IP address and name from the host file.
- Adding or removing a host using the CLI.

## 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 host 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:

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. This is 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 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.

## 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:** Due to the line terminator differences between Windows, Solaris, and Linux hosts, host files cannot be shared between Windows, Solaris, or Linux hosts.

## Adding a Single Host

The OneCommand Manager application enables you to specify a single TCP/IP host to manage. You can add a RMAPI host or 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 4-3 on page 40). (Not available in read-only mode.)

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

To add a single host:

1. From the **Discovery** menu, select **TCP/IP>Add Host**. The Add Remote TCP/IP Host dialog box appears (Figure 5-2 on page 50).



**Figure 5-2** Add Remote TCP/IP Host Dialog Box

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 0's as defined by the IPv6 address specification.

3. Configure the discovery method:
  - If you want to add the host using 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 box appears with default CIM settings (Figure 5-3 on page 51).

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

**Figure 5-3** Add Remote TCP/IP Host Dialog Box with CIM Credentials

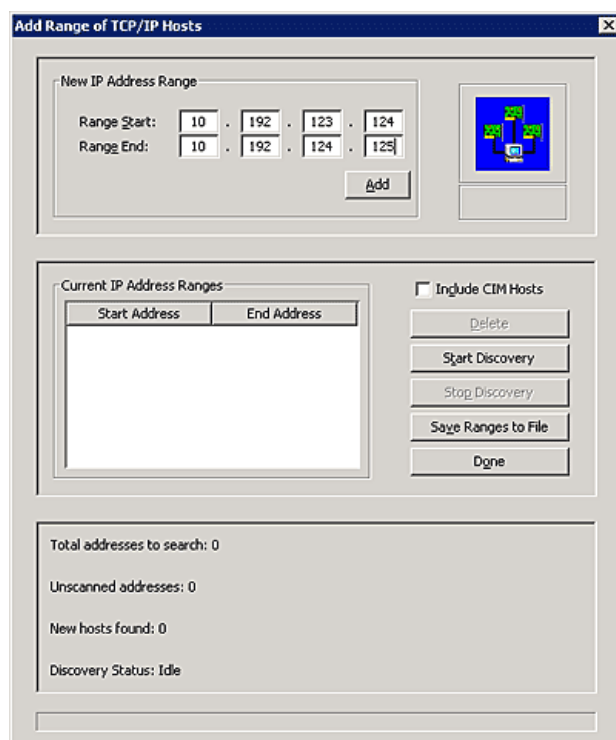
4. Edit the default CIM settings if necessary and click **Add Host**. A message appears indicating the new host was successfully added.

## Adding a Range of Hosts (IPv4 Only)

You can find the TCP/IP-accessed manageable hosts by searching a range of IPv4 addresses. The Add Range of TCP/IP Hosts dialog box (Figure 5-4 on page 52) enables you to build the initial list of TCP/IP accessed manageable hosts. (Not available in strictly local or local plus management modes.)

### Notes

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



**Figure 5-4** Add Range of TCP/IP Hosts Dialog Box

To add a range of remote hosts:

1. From the **Discovery** menu, select **TCP/IP>Add Range of Hosts**. The Add Range of TCP/IP Hosts dialog box appears (Figure 5-4 on page 52).
2. Enter the complete start and end address range (IPv4 only) and click **Add**. The added address range appears in the dialog box. Add any additional ranges 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 host 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 box.

**Note:** The number of hosts found does not correspond directly to the number of hosts added to the discovery-tree (Figure 4-3 on page 40). 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 box (Figure 5-4 on page 52).

## Removing Hosts

Removing hosts that are no longer discovered improves the operation of the discovery server. For example, you may want to remove a host when it is removed from the network. (Not available in read-only mode.)

To remove hosts:

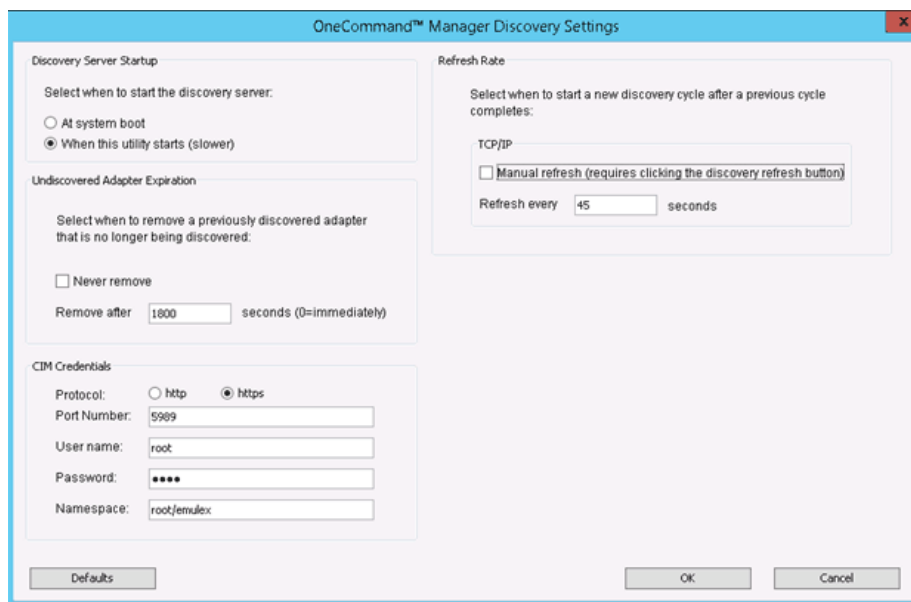
1. From the **Discovery** menu, select **TCP/IP>Remove Host(s)**. The Remove Hosts dialog box shows a list of discovered hosts. Any host that is not currently discovered appears in red. Click **Show Undiscovered Hosts Only** to display only currently undiscovered hosts.
2. From the Remove Hosts dialog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

## Configuring Discovery and Default CIM Credentials

Use the OneCommand Manager application Discovery Settings dialog box (Figure 5-5 on page 54) 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 only supported 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 4-3 on page 40) 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 5-5** Discovery Settings Dialog Box

To configure discovery settings:

1. From the Discovery menu, select **Modify Settings**. The Discovery Settings dialog box appears (Figure 5-5 on page 54).
2. Define the discovery properties you want.
3. The CIM credentials area can be used to set the default CIM credentials which 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 ESXi hosts. 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. Since, by default, the HTTP is disabled on sfcB in ESXi host, you should use HTTPS to communicate to the ESXi host.
  - User name: The user name field contains the username with which to connect to the ESXi hosts. By default this is 'root'.
  - Password: This password field contains the password of the user name that is used to connect to the ESXi host.
  - Namespace: Namespace is the namespace of the Emulex provider. The default namespace is 'root/emulex'.

**Note:** If the Emulex CIM Provider present in ESXi is an inbox provider, then the namespace should be "elxhBacmpi/cimv2". If the out-of-box CIM Provider is installed, then the namespace should be "root/emulex".

To check whether the CIM Provider is inbox or out-of-box, enter the following command on the ESXi host.

```
~ # esxcli software vib list | grep cim
```

Output similar to the following is returned:

```
emulex-cim-provider
10.0.514.6-01
Emulex VMwareAccepted
2013-06-10
```

4. Choose the refresh rate settings you want to apply.
5. Click **OK** to apply your changes. Click **Defaults** to return the discovery properties to their default settings.

## Viewing Discovery Information

The Discovery Information page 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:

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.

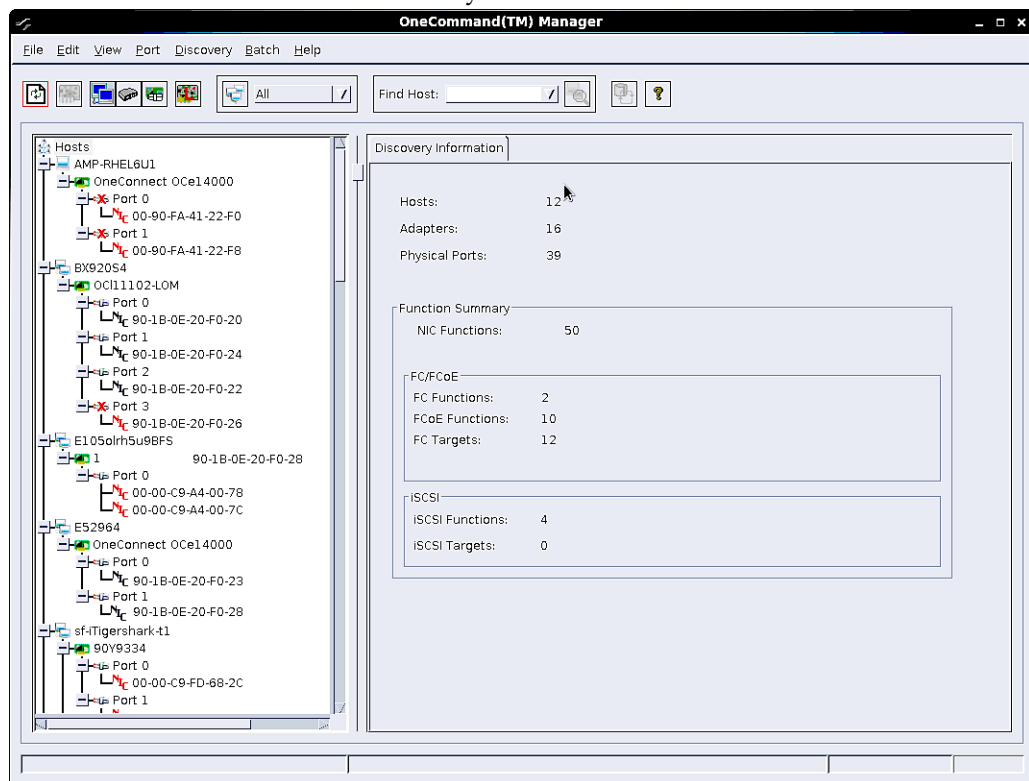


Figure 5-6 Discovery Information (Host View Selected)

## Discovery Information Field Definitions

- **Hosts** – The total number of discovered host computers containing manageable Emulex adapters. This 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.




## 6. Managing Hosts

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

### Viewing Host Information

Two tabs show host information: the Host Information tab (Figure 6-1 on page 57) and the Host Driver Parameters tab (Figure 8-30 on page 156). 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 “The Host Driver Parameters Tab” on page 156 for more information about the Host Driver Parameters tab.

To view the Host Information and Host Driver Parameters tabs:

1. Do one of the following:
  - From **View** menu, click **Group Adapters by Host Name**.
  - From the toolbar, click  **Group Adapters by Host Name**.
2. Select a host in the discovery-tree.
3. Select the **Host Information** tab (Figure 6-1 on page 57) or the **Host Driver Parameters** tab (Figure 8-30 on page 156).

The Host Information tab (Figure 6-1 on page 57) displays information for the selected host such as the number of adapters installed in the selected host and the number of fabrics to which it is connected.

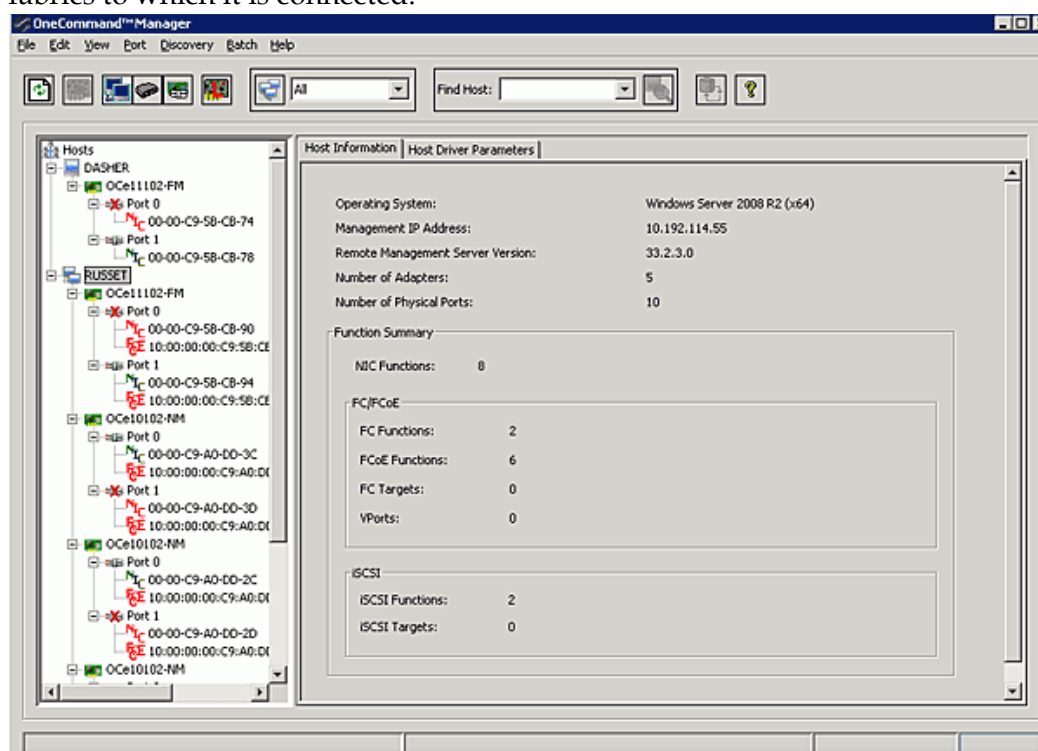


Figure 6-1 Host Information Tab

## Host Information Field Definitions

- Management IP Address – The Management IP Address field displays the host's IP address; for example, 138.239.82.131. “Local Host” is displayed if you selected the host from which you are actually 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, the “CIM Provider Version” field displays the version of the Emulex CIM Provider that is running on the remotely managed system.

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

## Function Summary Area

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

## Viewing Host Grouping Information

The Host Group Information (Figure 6-2 on page 59) tab displays information about the selected host group, such as the group name and the total number of hosts. See “Grouping Hosts” on page 59 to learn about creating host groups.

**Note:** Host grouping is not supported for VMware.

To view host grouping information, from the discovery-tree (Figure 4-3 on page 40), select the host group whose information you want to view.

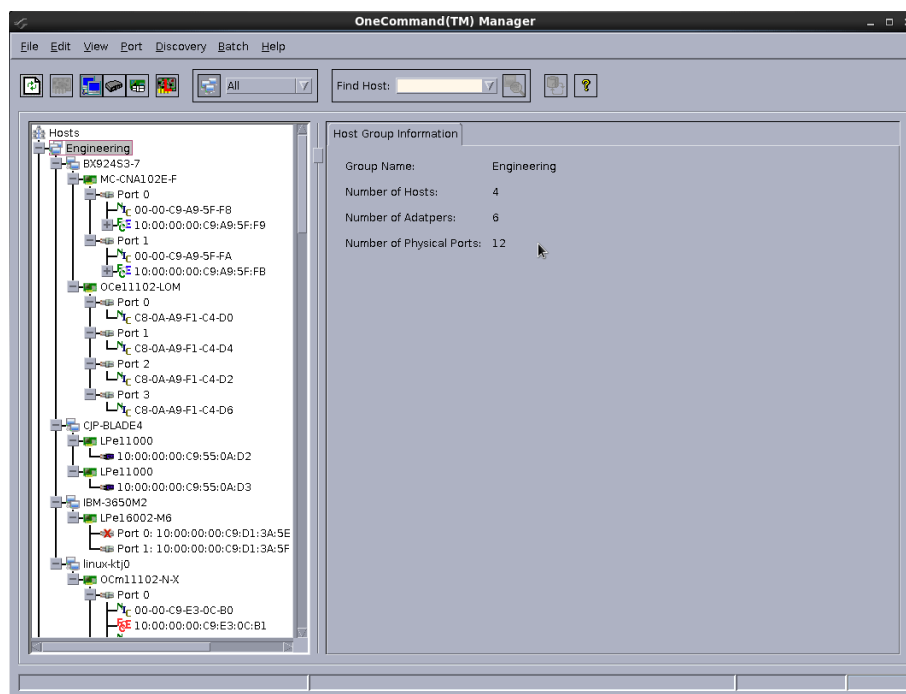


Figure 6-2 Host Group Information Tab

## Host Group Information Field Definitions

- 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.

## 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 function, 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 4-3 on page 40) in either a group centric format or in the host-based flat format. The Host grouping capability is available in Host view, Vport view, or Fabric view mode.


**Note:** The same fabric may appear under more than one host group. For example, some ports on the fabric may be attached to ports and hosts in one host group, and other ports on the same fabric may be attached 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 “Updating Firmware for Multiple


Adapters” on page 203 for more information.

**Note:** Grouping hosts is not supported on VMware.


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

- From the **View** menu, uncheck **Show Groups**.
- From the toolbar  unclick **Show Host Groups**.

To display all hosts groups:

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

To display all hosts assigned to a particular group:

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

## Managing Host Groups

Use the Host Group Management dialog box (Figure 6-3 on page 61) to create and delete host groups, add and remove hosts, and restore host groups.

**Note:** Managing host groups is not supported on VMware.

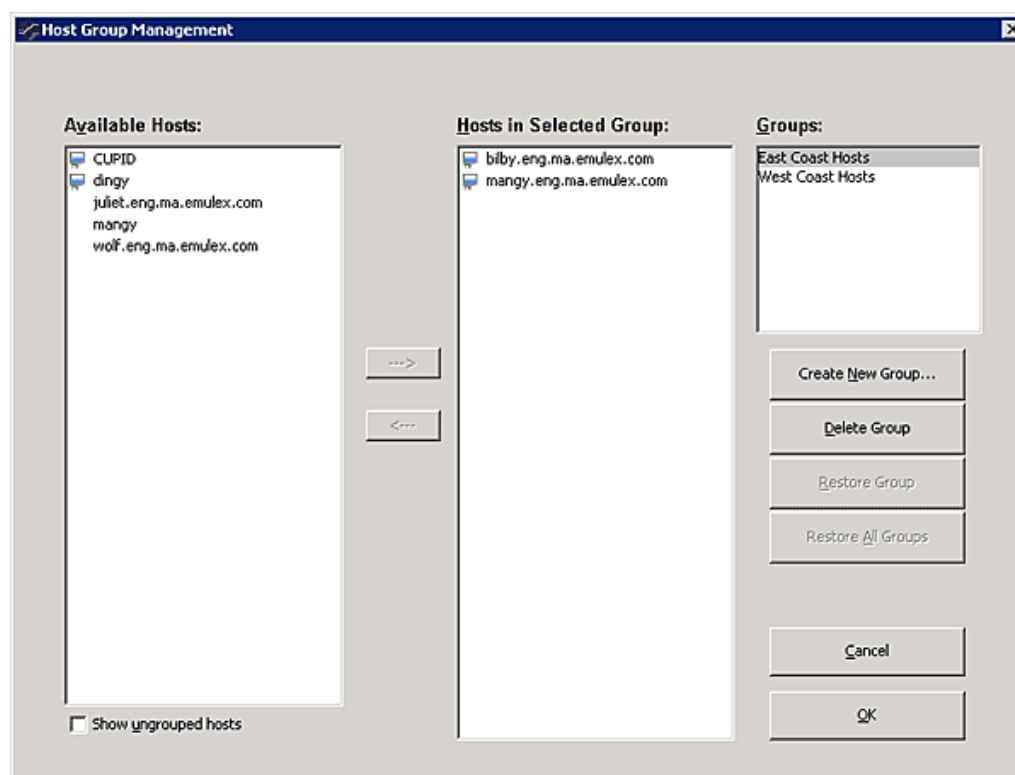


Figure 6-3 Host Group Management Dialog Box



## Host Group Management Field Definitions

- 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 checked, 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 the currently defined host groups. If you select a group in this list, its host members appear in the Hosts in Selected Group list.

## Host Group Management Buttons

- Right arrow – Adds selected available hosts to the currently selected group.
- Left arrow – Removes selected hosts from the currently selected group.
- Create New Group – Enables you to create a new host group.
- Delete Group – Removes the currently selected host group.
- Restore Group – Returns the selected group's configuration to its original state.
- Restore All Groups – Returns all groups to their original state.
- OK – Saves the current configuration changes and closes the dialog box.
- Cancel – Discards changes and closes the dialog box.


## Host Group Management Icons

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

## Creating a Host Group

To create a new host group:

1. From the **View** menu, select **Manage Groups**. The Host Group Management dialog box appears (Figure 6-3 on page 61).
2. Click **Create New Group**. The Create New Host Group dialog box is displayed (Figure 6-4 on page 62).




**Figure 6-4** Create New Host Group Dialog Box

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 box.

## Deleting a Host Group

To delete a host group:

1. From the **View** menu, select **Manage Groups**. The Host Group Management dialog box appears (Figure 6-3 on page 61).
2. From the **Groups** list, select the group you want to delete. The Host Group Management warning dialog box appears (Figure 6-5 on page 62).



**Figure 6-5** Host Group Management Warning Dialog Box

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

## Adding a Host to a Host Group

To add a host to a group:

1. From the **View** menu, select **Manage Groups**. The Host Group Management dialog box appears (Figure 6-3 on page 61).
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 4-3 on page 40) displays the new configuration.

## Removing a Host from a Host Group

To remove a host from a host group:

1. From the **View** menu, select **Manage Groups**. The Host Group Management dialog box appears (Figure 6-3 on page 61).
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 4-3 on page 40) displays the new configuration.

## Restoring a Host Group

To restore a host group:

1. From the **View** menu, select **Manage Groups**. The Host Group Management dialog box appears (Figure 6-3 on page 61).
2. Click **Restore Group** to return the configuration settings for the currently selected host group to those in use when the dialog box was opened.

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

## Restoring all Host Groups

To restore all host groups:

1. From the **View** menu, select **Manage Groups**. The Host Group Management dialog box appears (Figure 6-3 on page 61).
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.

## 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 host groups 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.

## Searching for Hosts in the Discovery-Tree

The OneCommand Manager application enables you to search the discovery-tree (Figure 4-3 on page 40) 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 large installation with hundreds or thousands of hosts. It is also helpful in Fabric view mode, since the ports on a specific host may be dispersed among several fabrics making the ports on that host difficult to find in the discovery-tree (Figure 4-3 on page 40).

To search for a host:

1. Do one of the following:
  - 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  **Find Host** or press <Enter> on the keyboard.

The host you are searching for is highlighted in the discovery-tree (Figure 4-3 on page 40).

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



## 7. Managing Adapters and Ports

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

### FC Adapters

This section pertains to viewing FC adapter and port information.

#### Viewing FC Adapter Information

When you select an FC adapter from the discovery-tree (Figure 4-3 on page 40), the Adapter Information tab (Figure 7-1 on page 65) contains general attributes associated with the selected FC adapter.

To view FC adapter information:

1. Select **Host**, **Fabric**, or **Virtual Ports** view.
2. Select an FC adapter in the discovery-tree. The Adapter Information tab appears (Figure 7-1 on page 65).

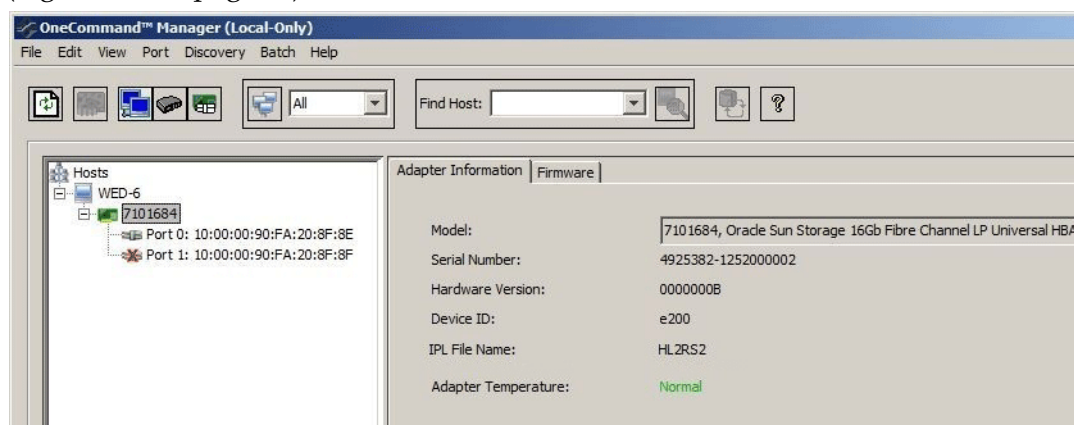


Figure 7-1 FC Adapter Information Tab

#### FC Adapter Information Field Definitions

- **Model** – The complete model name of the adapter.
- **Serial Number** – The manufacturer's serial number for the adapter.
- **Hardware Version** – Displays 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 adapter's temperature and one of the following 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, inoperable fans, and air conditioning issues that cause high ambient air temperatures.
- Exceeds operational range – Adapter stopped: The temperature has reached 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, inoperable 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.

## Viewing FC Port Information

When you select an FC port from the discovery-tree, the Port Information tab (Figure 7-2 on page 66) contains general attributes associated with the selected FC adapter.

To view FC Port information:

1. Select **Host** or **Fabric** view.
2. Select a FC port in the discovery-tree.
3. Select the **Port Information** tab (Figure 7-2 on page 66).

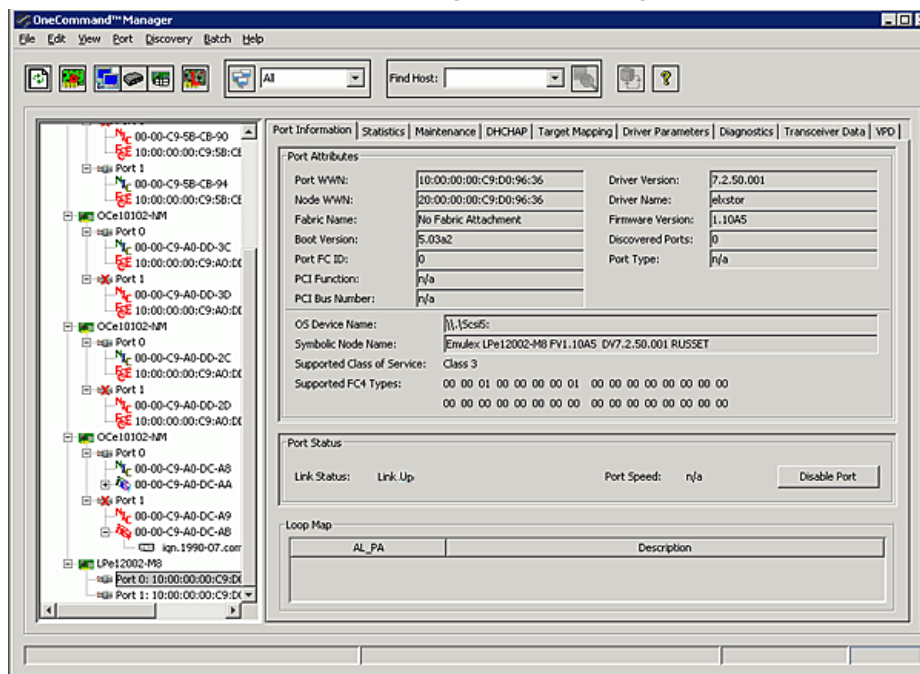


Figure 7-2 FC Port Information Tab

## FC Port Information Field Definitions

### Port Attributes Area Field Definitions

- Port WWN – The WWPNN of the adapter.
- Node WWN – The WWPNN of the adapter.
- Fabric Name or Host Name – The Fabric Name field is displayed in Host view. This is a 64-bit worldwide unique identifier assigned to the fabric. The Host Name is displayed in Fabric view. The host name 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”.
- Port FC ID – The FC ID for the selected adapter port.
- 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.
- Discovered Ports – The number of mapped and unmapped ports found during discovery by the Emulex adapter driver. The mapped ports are targets and the unmapped ports are non-targets such as switches or adapters.
- 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 non-delivery.
  - Class 2 – provides a frame switched service with confirmed delivery or notification of non-delivery.
  - Class 3 – provides a frame switched service similar to Class 2 but without notification of frame delivery or non-delivery.
- 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 Field Definitions

- Link Status – The status of the link on the selected adapter port.
- Port Speed – The current port speed of the selected adapter port.

## Loop Map Table Definitions

- The loop map shows the different ports present in the loop, and is present only if the port (adapter) is operating in loop mode. The simplest example would be to connect a JBOD directly to an adapter. When this is done, the port type is a private loop, and the loop map has an entry for the adapter, and one entry for each of the disks in the JBOD. (Not supported on VMware ESXi servers being managed through the CIM interface. )

## Port Information Buttons

- Enable\Disable Port - Click to enable or disable the selected FC port.

# OneConnect Adapters

This section describes viewing port and adapter information and managing OneConnect Adapters.

## Viewing OneConnect Adapter Information

When you select a OneConnect adapter from the discovery-tree, the Adapter Information tab contains general attributes associated with the selected OneConnect adapter. You can also use this tab to change an adapter's personality, enable or disable UMC, and to view and enable licenses. See "Configuring UMC Channel Management" on page 72, "Changing Personalities on OneConnect OCe11102 Adapters" on page 82, and "Showing and Installing Licenses for OneConnect OCe11102 Adapters" on page 80 for more information.

To view general OneConnect adapter information:

1. Select **Host** view.

2. Select a OneConnect adapter in the discovery-tree The Adapter Information tab appears (Figure 7-3 on page 69).

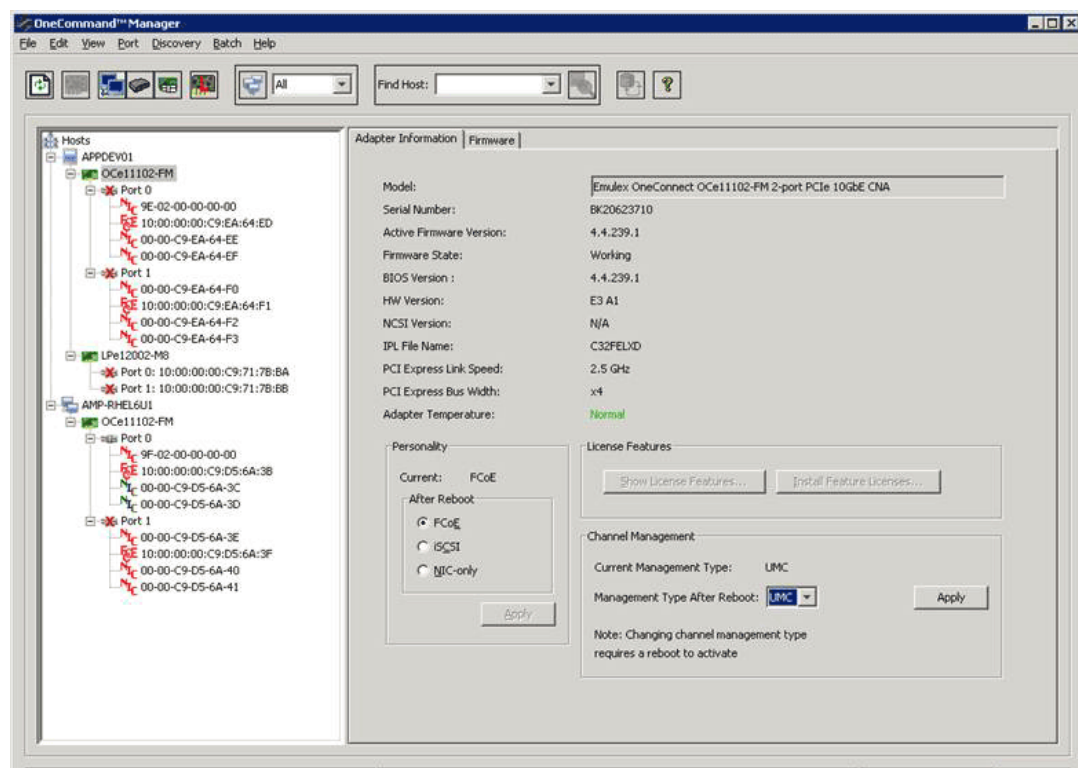


Figure 7-3 Adapter Information Tab

## OneConnect Adapter Information Field Definitions

- Model – The model of the selected adapter.
- Serial Number – The serial number of the selected adapter.
- Active Firmware Version – The version of the firmware running on the selected adapter.
- Firmware State – The condition of the firmware.
- BIOS Version – The version of the BIOS in use.
- HW Version – The hardware version of the selected adapter.
- NCSI Version – The NCSI version.
- IPL File Name – The name of the IPL file currently loaded.
- PCI Express Link Speed – The speed of the PCI bus in which the adapter running.
- PCI Express Bus Width – The number of lanes for the slot in which the adapter is running.
- Adapter Temperature – If the adapter's temperature is not available, "Not Supported" is displayed. If supported by the adapter, this field displays the adapter's temperature and one of the following 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, inoperable fans, and air conditioning issues that cause high ambient air temperatures.
- Exceeds operational range: The temperature has reached critical limit. 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, inoperable 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.

## Personality Area

**Note:** Use the Adapter Configuration tab to manage personalities on OCe14000-series adapters.

- Current – The current personality in use by the adapter.
- After Reboot – The type of management that is used after a reboot.

**Note:** Not available on OneConnect OCe11101-EM/EX or OCe11102-EM/EX adapters.

- FCoE – Check to choose the FCoE personality.
- iSCSI – Check to choose the iSCSI personality.
- NIC Only – Check to choose the NIC only personality.

**Note:** Some of the personalities may be disabled if the personality is not available on the adapter.

- Apply button – Click to apply the personality you choose. The system must be rebooted for your selection to take effect.

## License Features Area

**Note:** Not available on OneConnect OCe11101-EM/EX, OCe11102-EM/EX or OCe14000-series adapters.

- Show License Features button – Click to show available licenses. See “Showing and Installing Licenses for OneConnect OCe11102 Adapters” on page 80 for more information.
- Install License Features button – Click to install licenses. See “Showing and Installing Licenses for OneConnect OCe11102 Adapters” on page 80 for more information.

## Channel Management Area

**Note:** Not available on OneConnect OCe11101-EM/EX or OCe11102-EM/EX adapters.

**Note:** Channel management for OCe14000-series adapters is done using the Adapter Configuration tab.

- Type – The type of channel management in use.
- Management Type After Reboot – A dropdown box to select a channel management type or to disable channel management.

**Note:** A reboot is required for the change to take effect.

## Viewing Channel Management Information (OCe11102 Adapters Only)

The Channel Management tab is displayed when an adapter port is selected in the discovery-tree (Figure 4-3 on page 40) and Channel Management is enabled on the Adapter Information tab.

The Channel Management tab shows the channel management type and the properties for the particular channel management type.

**Note:** If Channel Management was enabled on the Adapter Information tab without rebooting, the Channel Management tab is not displayed.

To view general channel management information:

1. From the discovery-tree (Figure 4-3 on page 40), select the adapter port whose channel information you want to view.
2. Select the **Channel Management** tab (Figure 7-4 on page 72). The appearance of the tab varies depending on the channel management type in use.



**Note:** Channel Management must be enabled on the Adapter Information tab for the Channel Management tab to appear.

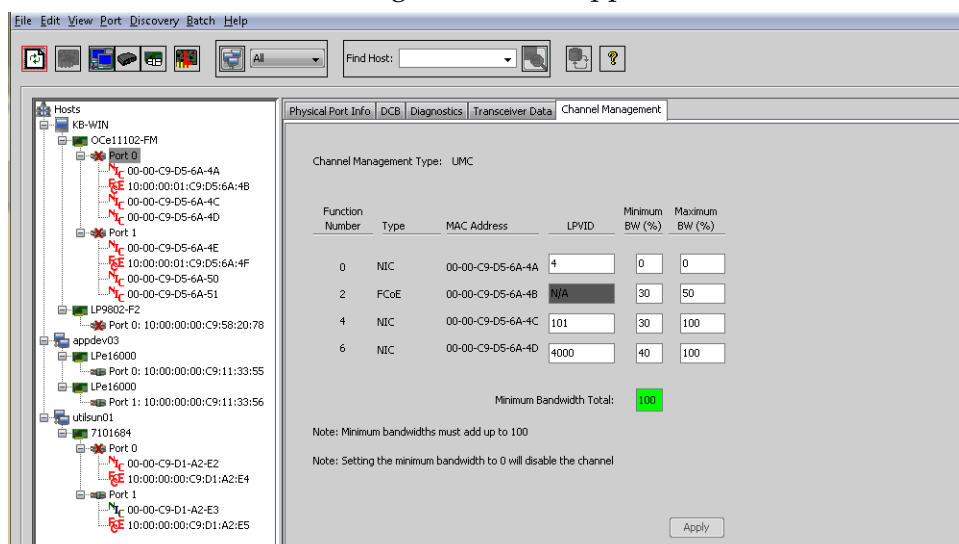


Figure 7-4 Example of the Channel Management Tab (UMC management depicted)

## Configuring UMC Channel Management

**Note:** Lenovo System X refers to UMC as SIMode.

Using the Channel Management tab (Figure 7-5 on page 73), each physical port can be partitioned into a maximum of four isolated channels providing a converged conduit for network and storage traffic. Each channel has its own unique MAC address and provides traffic management and provisioning capabilities such as minimum and maximum transmit rates and LPVIDs (for untagged packets).

Refer to the *Emulex Universal Multi-Channel Reference Guide* for additional information on UMC.

### Notes

- Properties for all channels on a port can be viewed and modified if UMC is enabled on the Adapter Information tab, even before rebooting to activate UMC on the adapter. This allows you to enable and configure UMC (on all channels), reboot and run UMC without further configuration.
- For Lenovo System X adapters, UMC mode may be referred to as “vNIC2”. The OneCommand Manager application displays the Channel Management Type “SIMode”.
- SR-IOV is not supported with UMC.



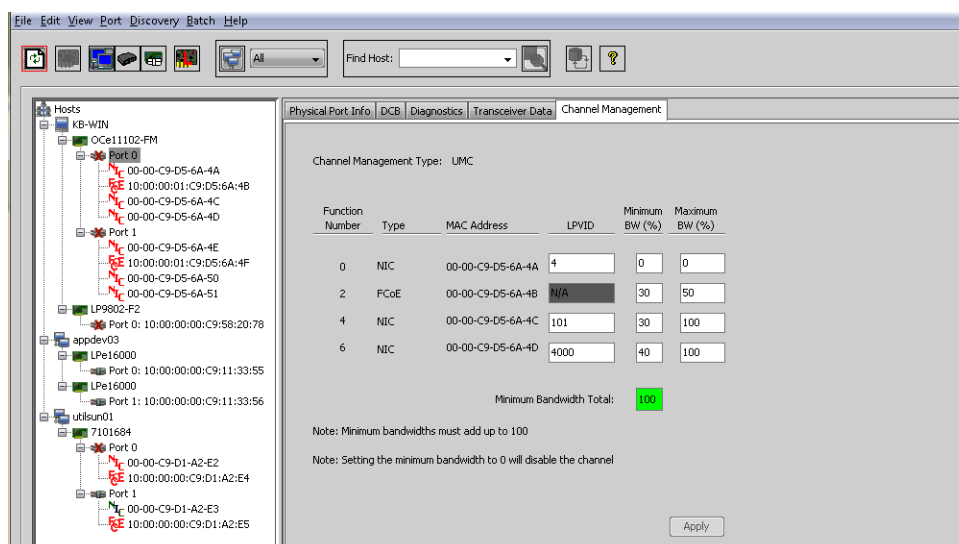


Figure 7-5 UMC Channel Management Tab

To configure UMC channel management:

1. From the discovery-tree, select the adapter port whose channel information you want to configure.
2. Select the **Channel Management** tab (Figure 7-5 on page 73).

**Note:** Channel Management must be enabled on the Adapter Information tab for the Channel Management tab to appear.

3. Assign an LPVID to all NIC channels. LPVID values can be specified between 2 and 4094 and must be unique across all channels on a port. However, LPVIDs do not need to be unique across the adapter's ports. You can assign the same LPVIDs on each port on the adapter.
4. Set the **Minimum Bandwidth** and **Maximum Bandwidth** values for each channel. Setting a channel's minimum and maximum bandwidth values to 0 disables the channel. The minimum bandwidth values of a channel can be set between 0 and 100, but the total of all of the minimum bandwidths must add up to 100 percent. The one exception is if all channel minimum and maximum bandwidths are set to 0. The maximum bandwidth values of a channel must be greater than or equal to the minimum bandwidth values for that channel.
5. Click **Apply**. A reboot is not required.

### UMC Channel Management Field Definitions

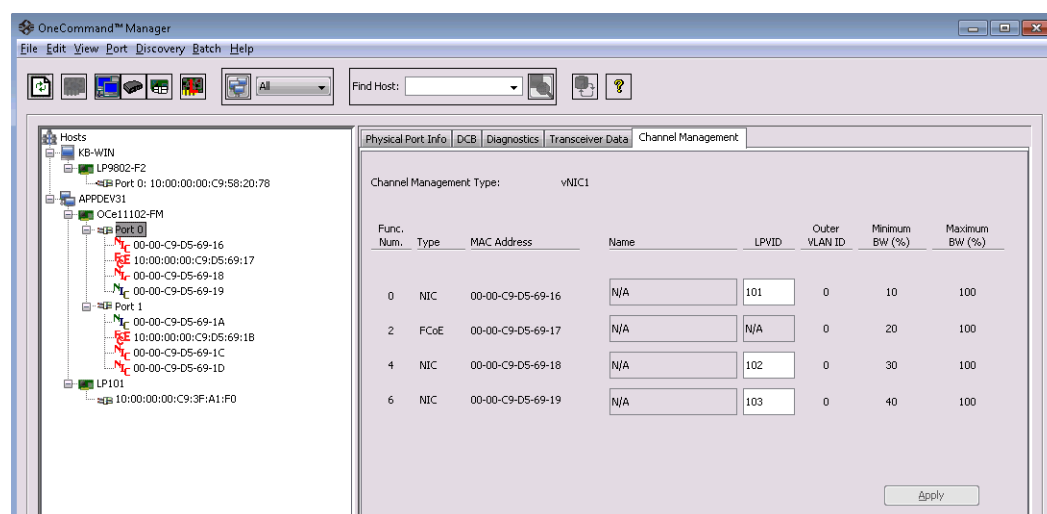
- **Function Number (Read-only)** – The function number assigned by the system to the channel. Channels 2 and 4 are read-only on OCe11102 adapters.
- **Type (Read-only)** – The protocol type in use.
- **MAC address (Read-only)** – The MAC address assigned to the channel.
- **LPVID** – Each channel can be assigned a default VLAN ID. Each egress untagged packet is tagged with this default value.

- Minimum BW (%) – The minimum percentage of the port's bandwidth at which the channel is guaranteed to run.
- Maximum BW (%) – The maximum percentage of the port's bandwidth that can be used by the channel.

## Viewing the Channel Management Tab for vNIC1 (Lenovo System X Only)

For Lenovo System X adapters that support vNIC1, the vNIC properties are displayed in the Channel Management tab if vNIC1 is enabled. The Channel Management Tab for vNIC1 (Figure 7-6 on page 74) is read-only except for the LPVID, which can be modified.

**Note:** vNIC is supported only on Lenovo System X virtual fabric adapters. For specific information as to whether it is supported on a specific adapter, see the release notes that are available on the Lenovo System X adapter pages on the Emulex website.



**Figure 7-6** Channel Management Tab for vNIC1 (Lenovo System X Only)

To view vNIC channel management:

1. From the discovery-tree, select the vNIC adapter port whose channel information you want to view.
2. Select the **Channel Management** tab (Figure 7-6 on page 74).

## Channel Management Field Definitions for vNIC1 (Lenovo System X Only)

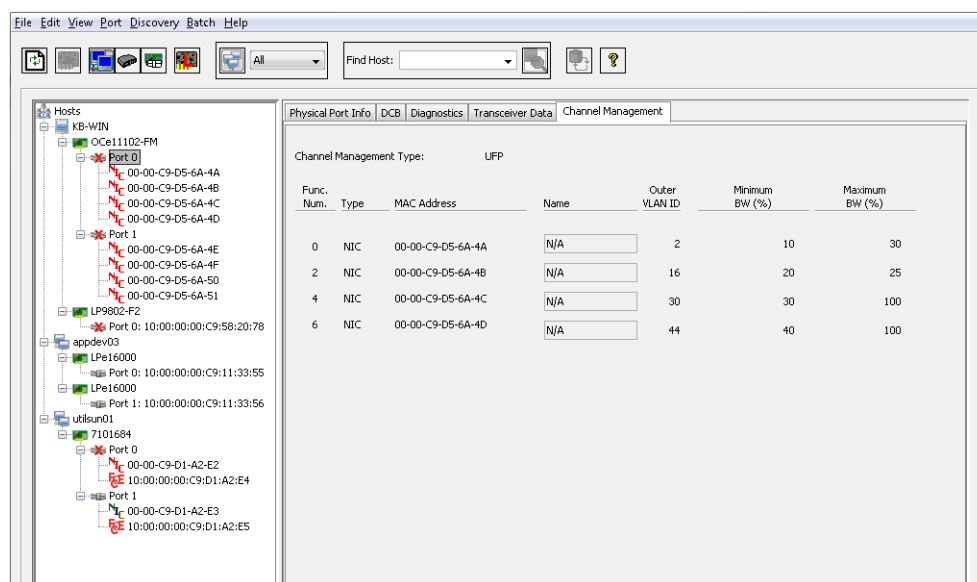
- Type (Read-only) – The protocol type in use.
- Func. Num. (Read-only) – The function number assigned by the system to the channel.
- Type (Read-only) – The protocol type in use by the channel.
- MAC address (Read-only) – The MAC address assigned to the channel.

- Name – The name assigned to the vNIC by an administrator during switch configuration.
- LPVID – Each channel can be assigned a default VLAN ID. Each egress untagged packet is tagged with this default value.
- Outer VLAN ID – The VLAN identifier used between the NIC port and the switch. The switch maps this value into the VLAN ID used on the network.
- Minimum BW (%) – The minimum percentage of the port's bandwidth at which the channel is guaranteed to run.
- Maximum BW (%) – The maximum percentage of the port's bandwidth that can be used by the channel.

## Viewing the Channel Management Tab for UFP (Lenovo System X Only)

For Lenovo System X adapters that support UFP, the UFP properties are displayed in the Channel Management tab (when UFP is enabled). The Port Channel Management Tab for UFP (Figure 7-7 on page 75) is read-only.

**Note:** UFP is supported only on Lenovo System X virtual fabric adapters. For specific information as to whether it is supported on a specific adapter, see the release notes that are available on the Lenovo System X adapter pages on the Emulex website.



**Figure 7-7** Channel Management Tab for UFP (Lenovo System X Only)

To view UFP channel management:

1. From the discovery-tree, select the UFP adapter port whose channel information you want to view.
2. Select the **Channel Management** tab (Figure 7-7 on page 75).

### Channel Management Field Definitions for UFP (Lenovo System X Only)

- Channel Management Type (Read-only) – The type of channel management type in use.
- Func. Num. (Read-only) – The function number assigned by the system to the channel.
- Type (Read-only) – The protocol type in use by the channel.
- MAC address (Read-only) – The MAC address assigned to the channel.
- Name – The name assigned to the vNIC by an administrator during switch configuration.
- Outer VLAN ID – The VLAN identifier used between the NIC port and the switch. The switch maps this value into the VLAN ID used on the network.
- Minimum BW (%) – The minimum percentage of the port's bandwidth at which the channel is guaranteed to run.
- Maximum BW (%) – The maximum percentage of the port's bandwidth that can be used by the channel.

## Viewing ASIC Information

When you select a OneConnect multi-ASIC adapter from the discovery-tree, the ASIC Information tab (Figure 7-8 on page 77) contains general attributes associated with the selected ASIC. You can also use this tab to view and enable licenses. See “Showing and Installing Licenses for OneConnect OCe11102 Adapters” on page 80 for more information.

To view general ASIC information:

1. Select **Host** view.
2. Select a OneConnect four-port adapter ASIC in the discovery-tree.

The ASIC Information tab is displayed (Figure 7-8 on page 77).

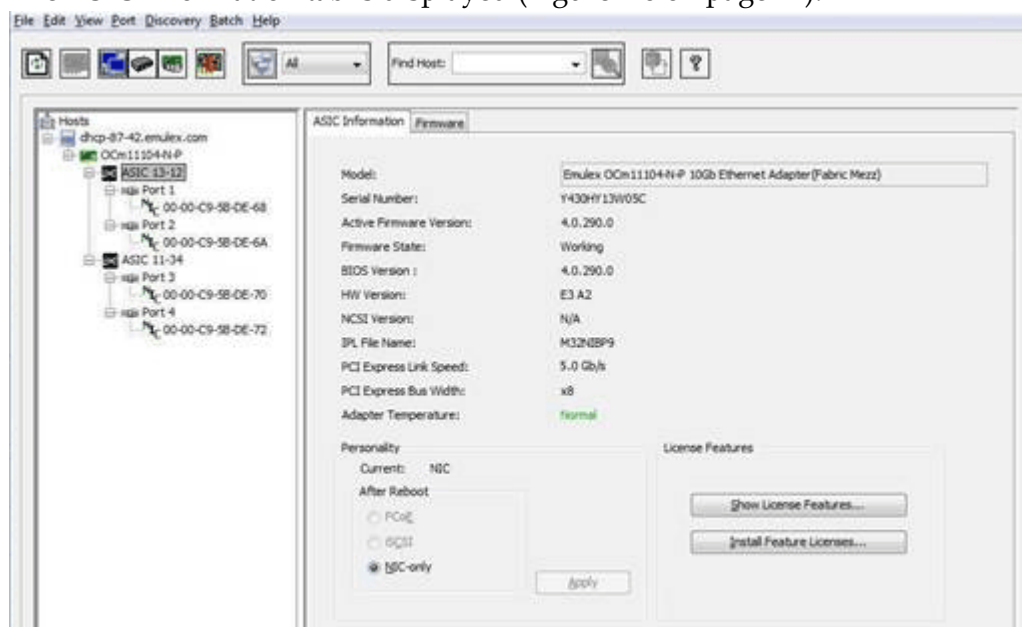


Figure 7-8 ASIC Information Tab

## ASIC Information Field Definitions

- Model – The model of the selected adapter.
- Serial Number – The serial number of the selected adapter.
- Active Firmware Version – The version of the firmware running on the selected adapter.
- Firmware State – The condition of the firmware.
- BIOS Version – The version of the BIOS in use.
- HW Version – The hardware version of the selected adapter.
- NCSI Version – The NCSI version.
- IPL File Name – The name of the IPL file currently loaded.
- PCI Express Link Speed – The speed of the PCI bus in which the adapter is running.
- PCI Express Bus Width – The number of lanes for the slot in which the adapter is running.
- Adapter Temperature – If the adapter's temperature is not available, "Not Supported" is displayed. If supported by the adapter, this field displays the adapter's temperature and one of the following temperature-related status messages:
  - Normal: The adapter's temperature is within normal operational range.
  - Exceeded operational range – Critical: 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, inoperable fans, and air conditioning issues that cause high ambient air temperatures.

- Exceeds operational range: The temperature has reached critical limit. 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, inoperable 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.

### Personality Area (OCe11102 Adapters Only)

- Current – The current personality in use by the adapter.
- After Reboot
  - FCoE – Check to choose the FCoE personality.
  - iSCSI – Check to choose the iSCSI personality.
  - NIC Only – Check to choose the NIC only personality

**Note:** Some of the personalities may be disabled if the personality is not available on the adapter.

- Apply button – Click to apply the personality you choose. The system must be rebooted for your selection to take effect.

### License Features Area (OCe11102 Adapters Only)

- Show License Features button – Click to show available licenses. See “Showing and Installing Licenses for OneConnect OCe11102 Adapters” on page 80 for more information.
- Install License Features button – Click to install licenses. See “Showing and Installing Licenses for OneConnect OCe11102 Adapters” on page 80 for more information.

## Viewing OneConnect Multi-ASIC Adapter Information

When you select a OneConnect multi-ASIC adapter from the discovery-tree, the Adapter Information tab (Figure 7-9 on page 79) contains general attributes associated with the selected dual ASIC four-port OneConnect adapter.

To view general OneConnect multi-ASIC adapter information:

1. Select **Host** view.
2. Select a OneConnect multi-ASIC adapter in the discovery-tree.

The Adapter Information tab is displayed (Figure 7-9 on page 79).



Figure 7-9 OneConnect Multi-ASIC Adapter Information

## OneConnect Multi-ASIC Adapter Information Field Definitions

- Model – The model of the selected adapter.
- Serial Number – The serial number of the selected adapter.
- HW Version – The hardware version of the selected adapter.
- ASICs – The number of ASICs on the selected adapter.

## Showing and Installing Licenses for OneConnect OCe11102 Adapters

The OneCommand Manager application allows you to view available licenses and install licenses to enable capabilities such as FCoE or iSCSI personalities on OCe11102 adapters without having to “re-wire” the adapter.

Using the Adapter Information tab (Figure 7-10 on page 80), you can view what licenses are available and install licenses for the selected adapter.

**Note:** If licenses or licensable features are not available for the adapter, license information is not displayed.

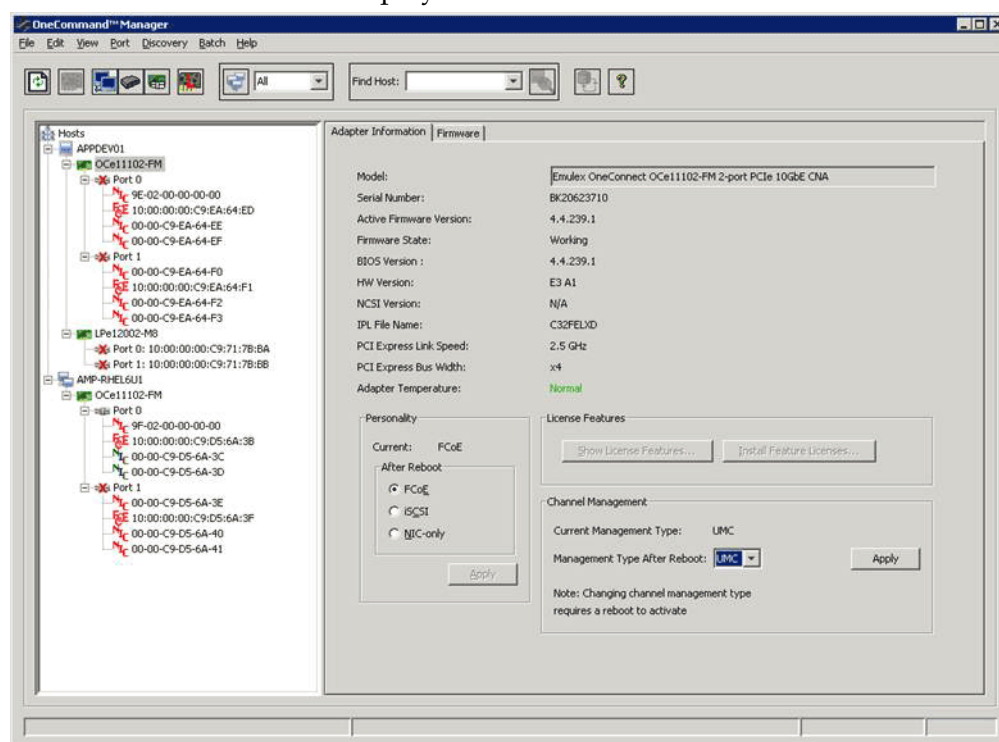


Figure 7-10 OneConnect OCe11102 Adapter Information Tab

## Showing Licenses

To view the available licenses for OCe11102 adapters:

1. From the discovery-tree, select the OCe11102 adapter whose licenses you want to view. The Adapter Information tab is displayed (Figure 7-10 on page 80).
2. On the Adapter Information tab, click **Show License Features**. The License Features window appears (Figure 7-11 on page 81). An X in the Enabled column indicates that the capability is licensed and enabled for that adapter.

**Note:** An empty feature list means the adapter has no licensable capabilities.



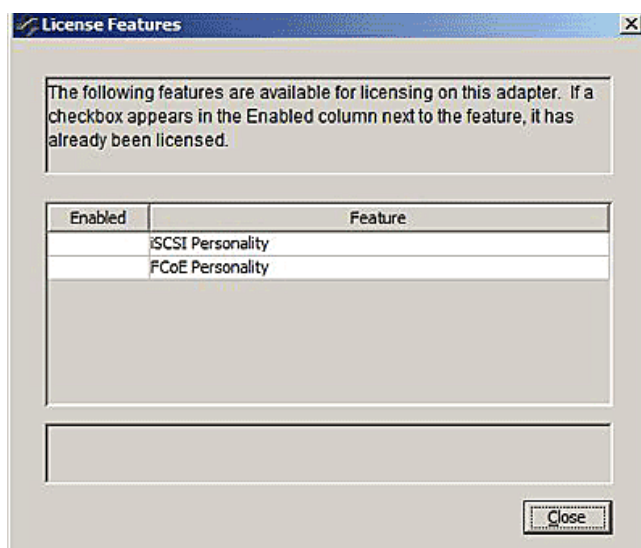


Figure 7-11 Licensed Features Window

## Installing Licenses

To install licenses for OCe11102 adapters:

1. From the discovery-tree (Figure 4-3 on page 40), select the OCe11102 adapter whose licenses you want to install. The Adapter Information tab is displayed (Figure 7-10 on page 80).
2. From the Adapter Information tab, click **Install Feature Licenses**. The Install Feature Licenses dialog box appears displaying the AdapterID (Figure 7-12 on page 81).

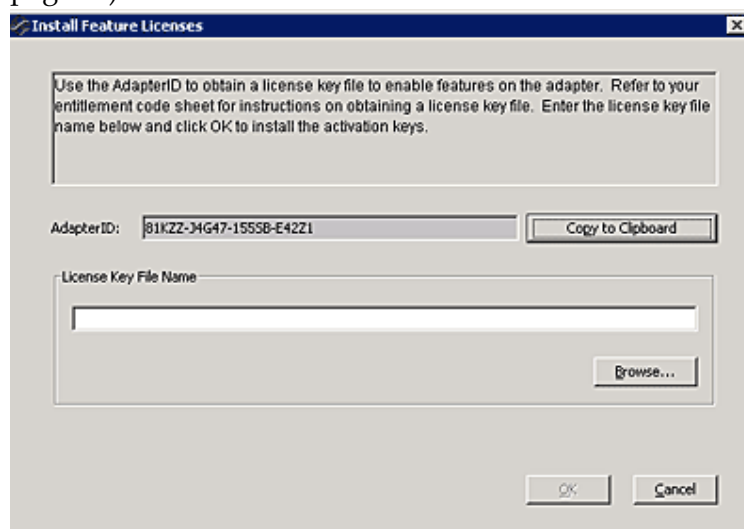


Figure 7-12 Install Feature Licenses Dialog Box

3. Following the instructions you received with the Entitlement Code, go to the License website and enter the AdapterID and Entitlement Code.

**Note:** The Copy to Clipboard button enables you to copy the AdapterID to the clipboard so you can paste it into a file or in the AdapterID field at the License website.

4. When the AdapterID and Entitlement Code are successfully validated, download a License Key File containing one or more activation keys.
5. Using the Install Feature Licenses dialog box (Figure 7-12 on page 81), enter the name of the License Key File (or click **Browse** to use a file browser to find the file) and click **OK**.
6. A dialog box appears confirming that you want to install the licenses. Click **OK**.
7. A dialog box appears notifying you that the installation was successful or why it failed. Click **OK**.

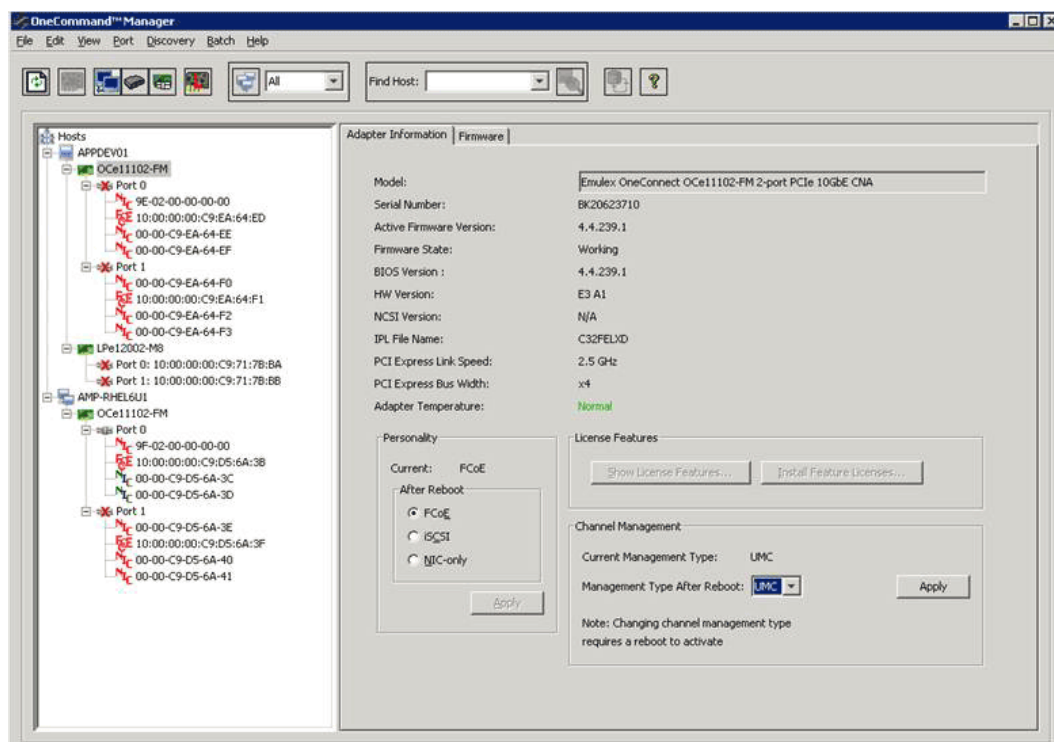
## Changing Personalities on OneConnect OCe11102 Adapters

The OneCommand Manager application enables you to change the personality or protocol running on OneConnect OCe11102 adapters.

When you change the personality of the adapter and reboot the host, the adapter starts running the new personality or protocol. The personalities that OCe11102 adapters currently run are NIC-Only, NIC + FCoE, and NIC + iSCSI. In some cases the adapters are pre-configured to support multiple personalities. In other cases you must install a license key before the adapter can support multiple personalities. See “Showing and Installing Licenses for OneConnect OCe11102 Adapters” on page 80 for more information.

**Note:** The three different personalities may not always be available on an adapter. For example, a NIC + FCoE adapter can change to a NIC-Only or NIC + iSCSI adapter, but an iSCSI adapter may not be able to change to a NIC + FCoE adapter.

Use the Adapter Information tab (Figure 7-13 on page 83) to make personality changes.



**Figure 7-13** OneConnect OCe11102 Adapter Information Tab

To change the personality of OCe11102 adapters:

1. From the discovery-tree, select the OCe11102 adapter whose personality you want to change. The Adapter Information tab is displayed (Figure 7-13 on page 83).
2. From the Personality area of the Adapter Information tab, select the personality type you want and click **Apply**.

**Note:** If the adapter does not support personalities, personality controls are not displayed. Also, if the adapter does not support a particular personality type that control is disabled.

3. Reboot the host for the personality change to take effect.

## Viewing OneConnect Firmware Information

Unlike LightPulse adapters, OneConnect adapter firmware is maintained on an adapter-specific, instead of port-specific, basis. Use the Firmware tab (Figure 7-14 on page 84) to download firmware and create diagnostic dumps for the selected adapter.

To view OneConnect firmware information:

1. Select **Host** view.

**Note:** iSCSI and NIC-Only adapters do not appear in Fabric view.

2. Select a OneConnect adapter in the discovery-tree (Figure 4-3 on page 40).

3. Select the **Firmware** tab (Figure 7-14 on page 84).

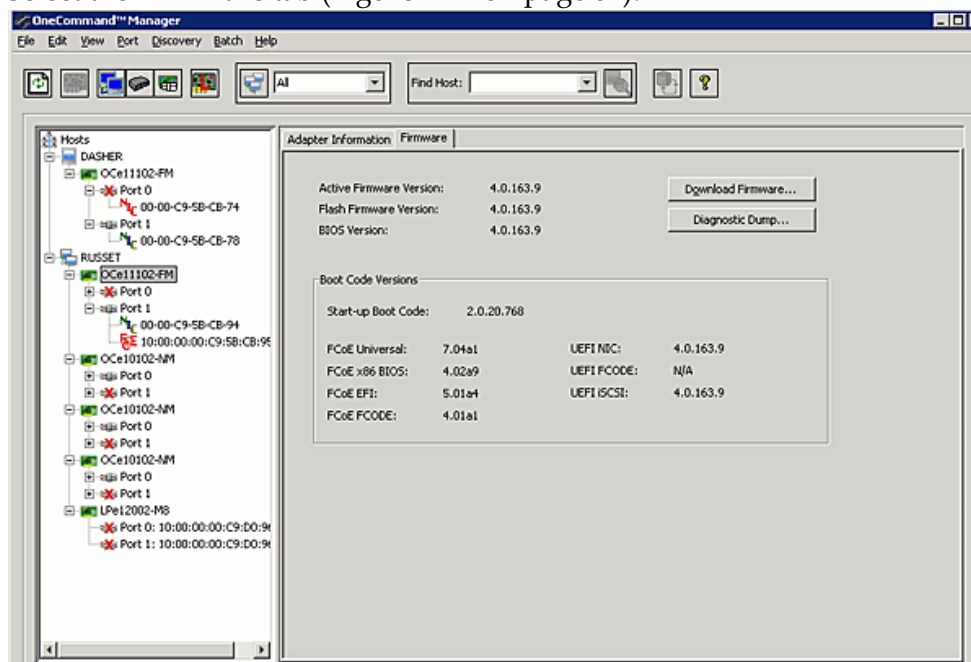


Figure 7-14 OneConnect Firmware Tab

## Firmware Tab Field Definitions

- **Active Firmware Version** – The firmware version currently being used by the adapter.
- **Flash Firmware Version** – The flash firmware version currently being used by the adapter.
- **BIOS Version** – The version of the BIOS currently being used by the adapter.

## Boot Code Versions Area

- **Startup-up Boot Code** – The boot code version currently being used by the adapter.

**Note:** This is the version of the code that boots the adapter. It has no relation to the FC, iSCSI, or PXE boot code versions.

- **FCoE Universal** – The combined flash image that includes three system specific FCoE Boot images (Open Boot, x86, and EFI 2.0).
- **FCoE x86 BIOS** – The single flash image containing x86 Boot for FCoE only.
- **FCoE EFI** – The single flash image containing EFI for FCoE only.
- **FCoE FCODE** – The single flash image containing Open Boot FCode for FCoE only.
- **UEFI NIC** – The single flash image containing UEFI for NIC and PXE Boot.
- **UEFI FCODE** – The single flash image containing Fcode for NIC only.
- **UEFI iSCSI** – The single flash image containing UEFI for iSCSI only.

### Firmware Tab Buttons (Not available in read-only mode.)

- Download Firmware – Click to update firmware on the selected adapter. See “Updating Adapter Firmware” on page 202 for more information.
- Diagnostic Dump – Click to create a diagnostic dump for the selected adapter. See “Creating Diagnostic Dumps” on page 252 for more information.

## Viewing OneConnect Physical Port Information

The Physical Port Info tab (Figure 7-15 on page 85) contains a general summarization of the PCI functions under that physical port and the current physical port status.

OneConnect OCe11100-series adapters also display additional Physical Port Status information including interface type, configured speed, and DAC cable length. You can set the port speed and DAC cable length. See “Setting Port Speed and DAC Cable Length (OneConnect OCe11102 and OCe14000-Series Adapters Only)” on page 86 for more information.

The Physical Port Info tab (Figure 7-15 on page 85) also allows you to enable or disable the physical port. See “Configuring iSCSI Initiator Login Options” on page 198 for more information.

To view physical port information:

1. Select **Host** view.
2. Select a OneConnect adapter port in the discovery-tree.
3. Select the **Physical Port Info** tab (Figure 7-15 on page 85).

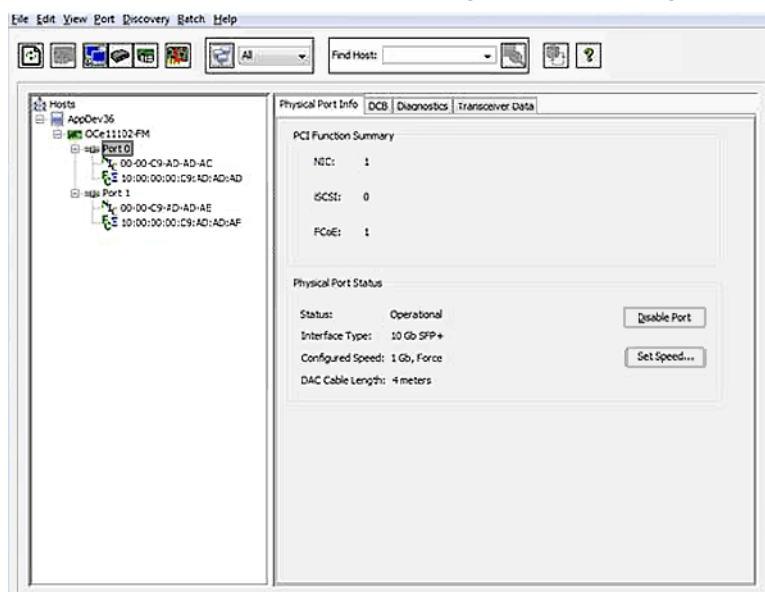


Figure 7-15 Physical Port Info Tab (OCe11102 Adapter Port Selected)

## Enabling and Disabling OneConnect Physical Ports

Using the Physical Port Info tab (Figure 7-15 on page 85) you can enable or disable the physical port. If you disable a physical port, you disable all functions, such as iSCSI and NIC, for the port. Disabled ports appear in the discovery-tree as a black port icon.

**Note:** You cannot disable a port if PXE Boot is enabled or if any of the iSCSI target sessions are boot sessions.

To enable or disable a physical port:


1. In the discovery-tree (Figure 4-3 on page 40), select the physical port you want to enable or disable.
2. Select the **Physical Port Info** tab (Figure 7-15 on page 85).
3. Click **Enable Port** or **Disable Port**.

## Setting Port Speed and DAC Cable Length (OneConnect OCe11102 and OCe14000-Series Adapters Only)

The Physical Port Info tab enables you to set port speed and DAC cable lengths for OCe11102 and OCe14000-series adapters.

To set the port speed for OCe11102-series adapters:

1. From the discovery-tree (Figure 4-3 on page 40), select the OCe11102-series adapter port whose speed you want to change.
2. Click **Set Speed** on the Physical Port Info tab (Figure 7-15 on page 85). The Change Port Speed dialog box appears (Figure 7-16 on page 86).



**Figure 7-16** Change Port Speed Dialog box (Force mode/1Gb Ethernet speed selected)

3. Set the desired mode and port speed. The (port speed) mode setting can be either "Default", "Force", or Auto-negotiate".
  - Default – Sets the port speed to the factory default configured speed of the adapter (from IPL).
  - Force – Sets the speed to a single speed value.
  - Auto-negotiate – Requires a speed setting to single speed or multiple speed choices that the port will use to auto-negotiate the port speed with the switch port.

**Note:** If the adapter's port speed setting and the switch's port speed setting conflict, the link will not come up.

4. If you set the Mode to "Force" and the Speed to "10 GB SFP+" or "QSFP+" you must set the DAC cable length in the range of 0-10 meters.

For 10GB SFP+, the length is the actual DAC cable length and 0 indicates an optical cable.

For the QSFP+ module type, a length of 0 indicates an optical cable and any non-zero length indicates a DAC cable (that is, the length does need to be the actual DAC cable length).

5. Click **OK**.

**Note:** For an embedded mezzanine adapter linked to an embedded switch on the internal port, the DAC Cable Length value is ignored.

## Viewing PHY Data (OneConnect 10GBASE-T series Adapters Only)

The PHY Data Tab (Figure 7-17 on page 87) displays port level operational parameters, error rates, and counters that are protocol and personality independent for OneConnect 10GBASE-T adapter ports.

**Note:** PHY data is not supported through the CIM interface.

To view OneConnect 10GBASE-T adapter port PHY information:

1. Select **Host** or **Fabric** view.

**Note:** iSCSI and NIC-Only adapters do not appear in Fabric view.

2. In the discovery-tree, select the OneConnect 10GBASE-T adapter port whose PHY information you want to view.
3. Select the **PHY Data** tab (Figure 7-17 on page 87).

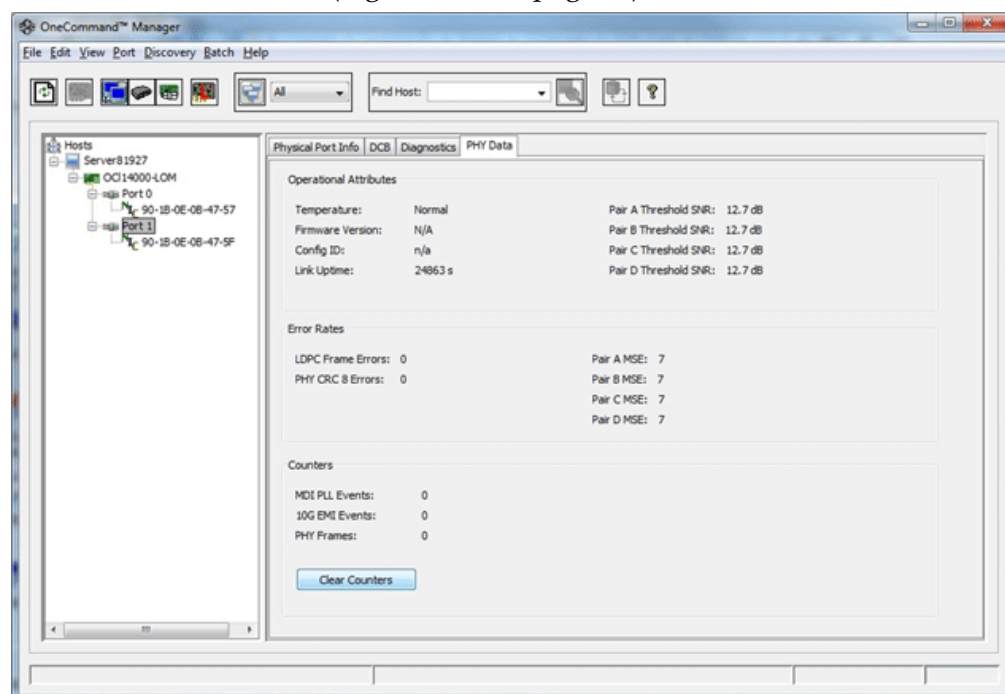


Figure 7-17 PHY Data Tab



## PHY Data Field Definitions

### Operational Attributes Area

- Temperature – The temperature of the selected port.
- Pair A/B/C/D SNR Margin – Displays the CNA's MDI interface average SNR margin for twisted pairs A, B, C, and D.

### Error Rates Area

- LDPC Frame Errors – The LDPC counter tracks the number of LDPC frames received by CNA's MDI interface that can not be corrected. This counter self-clears at MDI link down.
- Pair A/B/C/D MSE – Displays the CNA's MDI interface average MSE relative to the transmitted codewords for twisted pairs A, B, C, and D.

### Counters Area

- MDI PLL Events – The MDI PLL Event counter tracks events that affect CNA's normal operation. This counter self-clears at MDI link down.
- 10G EMI Events – The 10G EMI Event counter tracks the number of single-tone interference detected by CNA's MDI signals. This counter holds its value at MDI link down and self-clears at the next link up.
- PHY Frames – Counts the number of PHY frames transmitted and received since the MDI link has been established. This counter holds its value at MDI link down and self-clears at the next link up.

**Note:** For the OCe14000-series adapters, the PHY Frames counter value has increased to 64-bits to allow for larger frame counter values before rolling over to zero. The frame counter is unchanged for non-OCe14000-series adapters; specifically 32-bit.

- Clear Counters – Click the Clear Counters button to reset the counter fields in the Counters Area to zero for the selected port (not available for OCe11000-series adapters).

## Viewing OneConnect Transceiver Information

When you select a OneConnect adapter port from the discovery-tree (Figure 4-3 on page 40), the Transceiver Data tab (Figure 7-18 on page 89) enables you to view transceiver information such as vendor name, serial number, part number and so on. If the adapter/transceiver does not support some or all of the transceiver data, the fields display N/A.

To view OneConnect transceiver information:

1. Select **Host** or **Fabric** view.

**Note:** iSCSI and NIC-Only adapters do not appear in Fabric view.



2. In the discovery-tree (Figure 4-3 on page 40), select the OneConnect adapter port whose transceiver information you want to view.
3. Select the **Transceiver Data** tab (Figure 7-18 on page 89).

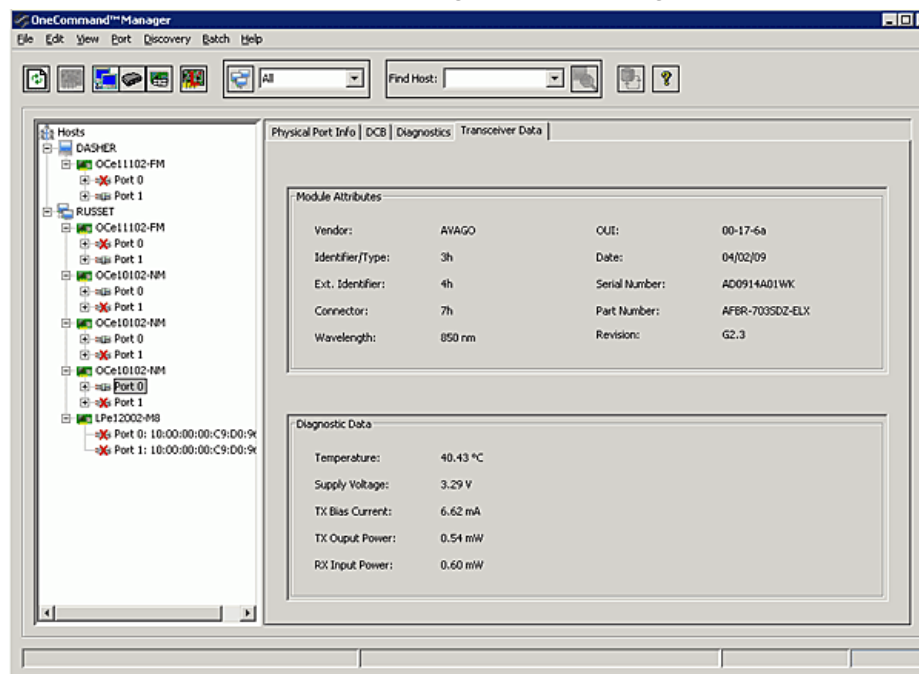


Figure 7-18 OneConnect Transceiver Data Tab

## Transceiver Data Field Definitions

### 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.
- OUI – The vendor's OUI. It is also known as the IEEE Company Identifier for the vendor.
- Date – The vendor's date code in the 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 TX bias current.
- TX Output Power – The measured TX output power.
- RX Input Power – The measured RX input power.

## OCe14000-Series Adapter Configuration

When you select an OCe14000-series adapter from the discovery-tree (Figure 4-3 on page 40), the Adapter Configuration tab (Figure 7-19 on page 91) is displayed. The type of information that appears depends on the protocols and capabilities that are available on the adapter. (The IPL defines the protocols and capabilities.) Additionally, you can enable SR-IOV on an adapter-wide basis.

For all adapters except NPar, there are three configuration modes available:

- Single personality – Select a single protocol (along with NIC, or NIC-only) to run on all ports on the adapter.
- Custom – Choose the protocols to run on each port independently.
- multi-channel – View the channel properties and change the configurable properties (based on the multi-channel type).

### Notes

- For NPar, the adapter is either running with NPar disabled (meaning the adapter is running as a single NIC or NIC+RoCE function on each port), or with NPar enabled, which is similar to a UMC configuration. See “Dell NPar Configuration View” on page 110 for more information.
- RoCE configurations are not supported with SR-IOV.

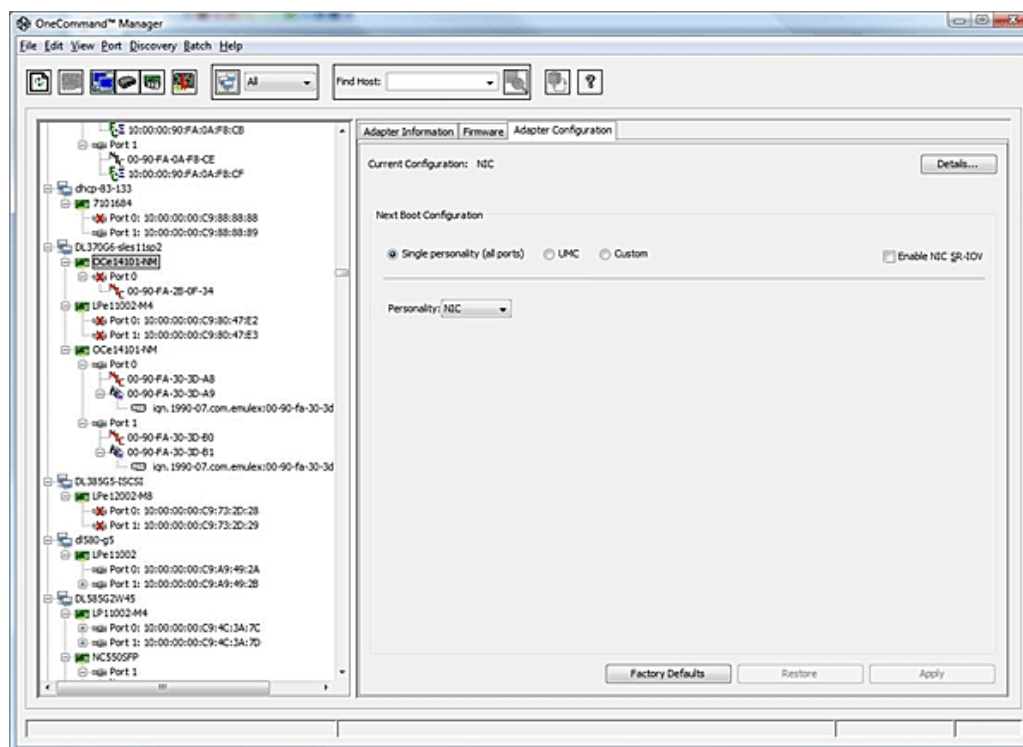



Figure 7-19 OCe14000-series Adapter Configuration Tab (FCoE selected)

## OCe14000-series Adapter Configuration Tab Field Definitions

- **Current Configuration** – Displays the protocol running on the adapter.
- **Next Boot Configuration** – Displays different checkboxes or radios buttons subject to the protocol running on the adapter. Select a checkbox or radio button to change the tab's view and next boot configuration method.
- **Enable NIC SR-IOV** – Select this checkbox to enable NIC SR-IOV on the adapter.

### Notes

- SR-IOV can only be enabled for non-multi-channel and non-NIC+RoCE configurations.
- If the NIC+RoCE personality is selected, the Enable NIC SR-IOV checkbox is automatically cleared and disabled.
- If you have an active configuration where only some of the NIC functions have SR-IOV enabled and you click the Adapter Configuration tab, a warning message appears, as shown in Figure 7-20 on page 92, prompting you to click the Apply button to save the configuration with SR-IOV enabled on all NIC functions.

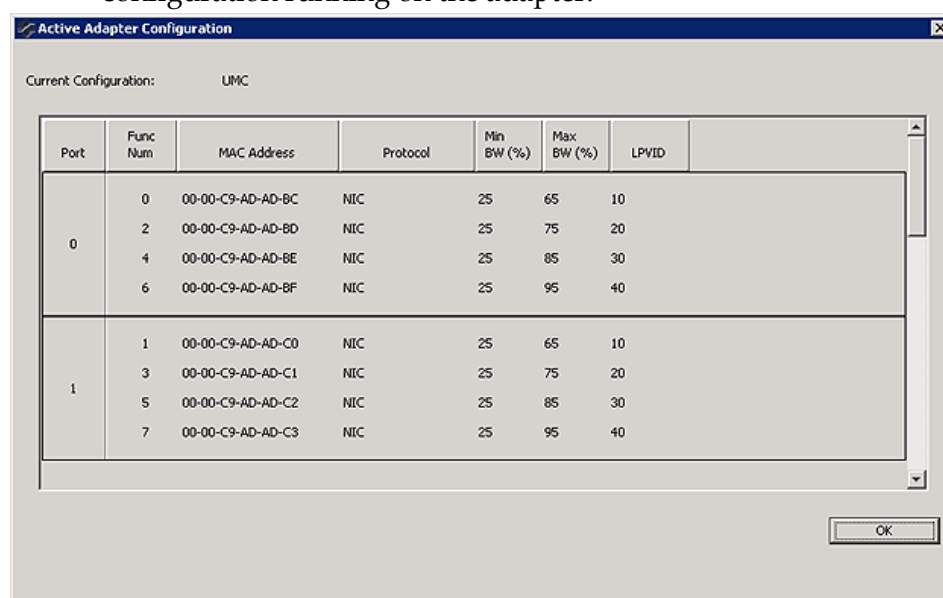


**Figure 7-20** SR-IOV Warning Message

- The display in the lower area of the tab depends upon the selected view of the next boot configuration. You can configure the adapter for a single personality (for example, FCoE), multi-channel, or a custom configuration (for mix and concurrent modes).

### OCe14000-series Adapter Configuration Tab Buttons

- Apply - Saves the changes to the next boot configuration to the adapter.
- Restore - Resets any changes to the next boot configuration back to the currently saved settings on the adapter.
- Factory Defaults - Returns the adapter's ports to their factory default profile and settings for that adapter. A reboot is required.
- Details - Displays a pop-up window with more information about the current configuration running on the adapter.

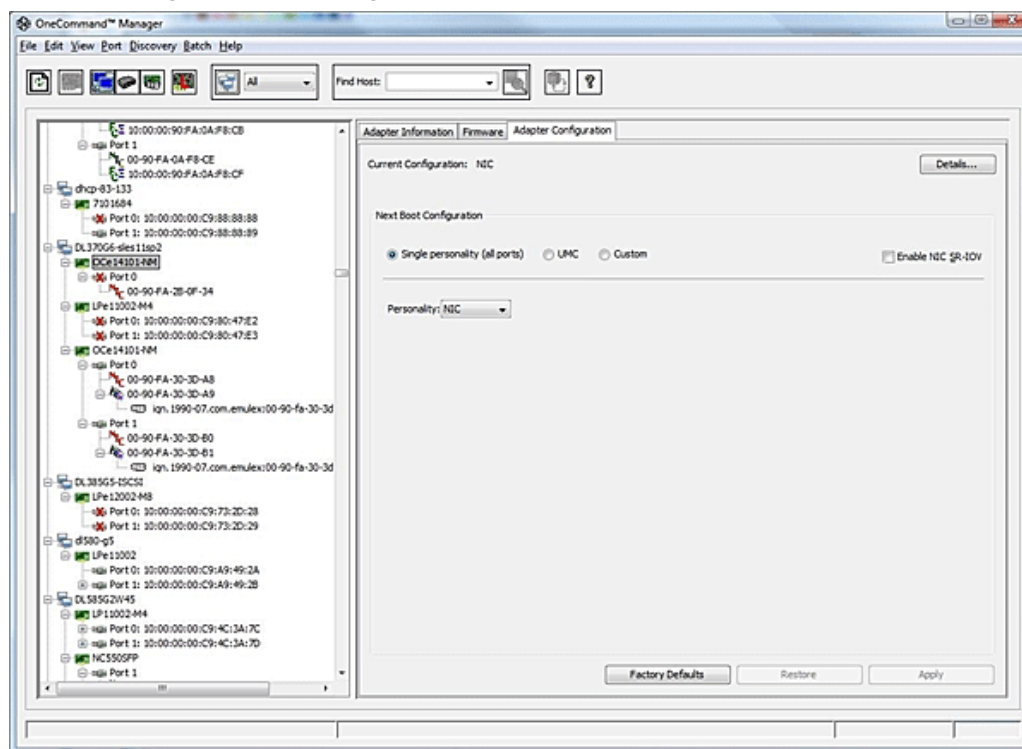


**Figure 7-21** Current Configuration Details example

## Configuring Single Personalities

The Single personality view allows you to select the same protocol to run on all ports of the adapter. You can also enable SR-IOV on all ports of the adapter, except when NIC+RoCE is selected.

When the Single personality radio button is selected, the Adapter Configuration tab looks like Figure 7-22 on page 93.



**Figure 7-22** Single Personality View (NIC selected)

When you configure a single personality for an adapter, all ports run the same number of functions and protocols on those functions.

To configure a single personality:

1. From the discovery-tree (Figure 4-3 on page 40), select the OCe14000-series adapter whose personality you want to configure.
2. From the Adapter Configuration tab (Figure 7-19 on page 91), check the **Single personality (all ports)** radio button.
3. Select the personality you want to apply from the Personality pull-down menu.

If you choose the NIC+RoCE configuration, the NIC+RoCE configuration list displays the available RoCE profiles.

- Choose the RoCE-2 profile for Windows SMB Direct or Linux iSER, NFS, or VM Migration support.
- For RoCE-1, check the Implementer's Lab on the Emulex website for any updated information on additional use cases.

After rebooting a NIC+RoCE configuration, the Current Configuration field displays NIC+RoCE along with the description of the selected RoCE profile.

## Notes

- RoCE configurations are not supported with SR-IOV.

- If you choose a NIC-Only configuration, there may be multiple NIC-Only configurations available. If the adapter supports multiple NIC configurations (profiles), you must select a NIC-Only configuration from the NIC-Only configuration list.
4. Click **Apply**. A message appears notifying you of the profiles activation requirements. Not all selections require a reboot.

## Custom Configurations

The Custom configuration view allows you to customize the protocols running on each port of an adapter. The Port Configuration table only displays the available ports. If the Custom radio button is selected, the Adapter Configuration tab looks like Figure 7-23 on page 94.

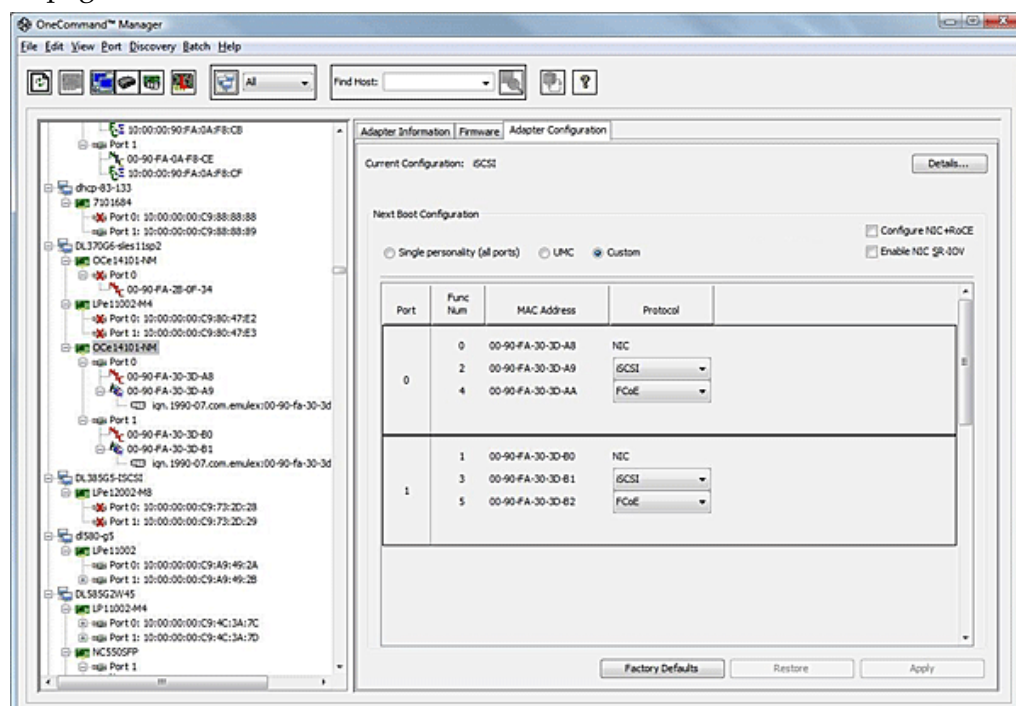


Figure 7-23 Custom View

### Notes

- Custom configurations are only available if the adapter supports mixed or concurrent mode storage.
- If you configure all ports the same, you essentially configured a single personality.
- NIC+RoCE and SR-IOV are mutually exclusive. If you select the Configure NIC+RoCE check box, you cannot select the Enable NIC SR-IOV check box because it is disabled. Similarly, if you select the Enable NIC SR-IOV check box, you cannot select the Configure NIC+RoCE check box because it is disabled.

To set up a custom configuration:

1. From the discovery-tree, select the OCe14000-series adapter whose personality you want to configure.
2. From the Adapter Configuration tab (Figure 7-19 on page 91), check the **Custom** radio button.
3. Select the protocol you want to apply from the Protocol pull-down menu for each function on each port.

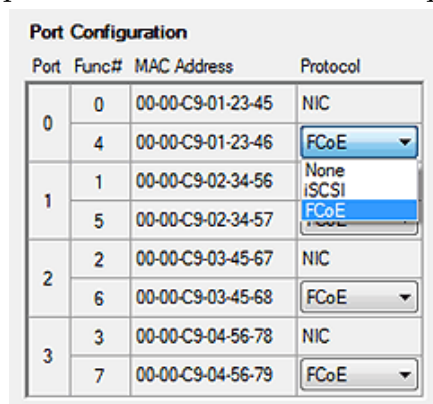
Or, check the NIC+RoCE checkbox to configure NIC+RoCE or only NIC on each port.

When the NIC+RoCE checkbox is not checked, you can check the SR-IOV checkbox to enable SR-IOV on each NIC function on the adapter.

4. Click **Apply**. A message appears notifying you of the profiles activation requirements. Not all selections require a reboot.

## Mixed Mode Configuration

If mixed mode is available, up to two functions per port can be configured. The first function is always NIC. The second function can be a storage protocol or “None”. When you click on the second function's pull-down menu, the available storage protocols for the function are displayed as depicted in Figure 7-24 on page 95.



Port	Func#	MAC Address	Protocol
0	0	00-00-C9-01-23-45	NIC
	4	00-00-C9-01-23-46	FCoE
1	1	00-00-C9-02-34-56	None
	5	00-00-C9-02-34-57	FCoE
2	2	00-00-C9-03-45-67	NIC
	6	00-00-C9-03-45-68	FCoE
3	3	00-00-C9-04-56-78	NIC
	7	00-00-C9-04-56-79	FCoE

**Figure 7-24** Mixed Mode Protocol pull-down menu

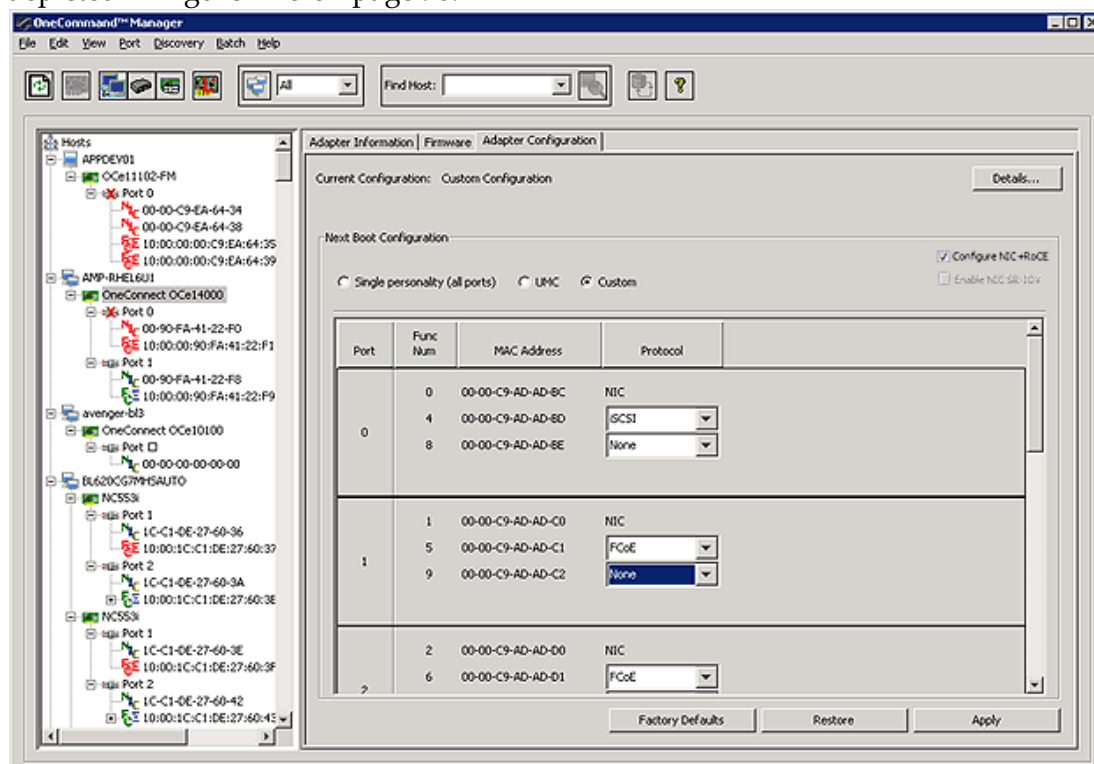
You can select any of the available storage protocols or “None” to indicate that the second function will not be available on that port. After selecting the protocols to run on each port, click **Apply** to save the configuration.

**Note:** If you select the same protocol or “None” on all ports, you essentially configured a single personality. In this case, after clicking **Apply**, the tab switches to the single personality view showing the selected personality in the personality pull-down menu.



## Concurrent Mode

If concurrent mode is available, up to three functions per port are configurable as depicted in Figure 7-25 on page 96.



**Figure 7-25** Concurrent Storage Configuration View

When configuring storage, the first function must be set to NIC. The pull-down menu for selecting a storage protocol to run on the second and third functions is the same as the pull-down menu for mixed mode storage. If two storage functions are configured on a port, they must be different storage protocols (that is, concurrent mode). When you select the storage protocol for the second function, the choices in the pull-down menu for the third function are limited to the other storage protocol or “None”. (Figure 7-26 on page 96).

Port Configuration			
Port	Func#	MAC Address	Protocol
0	0	00-00-C9-01-23-45	NIC
	4	00-00-C9-01-23-46	FCoE
	8	00-00-C9-01-23-47	iSCSI
1	1	00-00-C9-02-34-56	None
	5	00-00-C9-02-34-57	FCoE
	9	00-00-C9-02-34-58	None
2	2	00-00-C9-03-45-67	NIC
	6	00-00-C9-03-45-68	None
	10	00-00-C9-03-45-69	None
3	3	00-00-C9-04-56-78	NIC
	7	00-00-C9-04-56-79	iSCSI
	11	00-00-C9-04-56-7A	FCoE

**Figure 7-26** Concurrent Storage Configuration Choices for the Third Function



**Note:** If you select the same protocol or “None” for all second functions on all ports and “None” for all third functions, you essentially configured a single personality. In this case, after clicking the **Apply** button, the tab switches to the single personality view and the selected personality is shown in the personality pull-down menu.

## Configuring RoCE in a Custom View

In Custom view, NIC+RoCE can be configured on the first function of a port (see Figure 7-27 on page 97). After checking Configure NIC+RoCE, only the first function of each port is available and you can select NIC+RoCE or just NIC to run on each port. Storage functions cannot be configured if any of the ports on the adapter are configured to run NIC+RoCE. You must also choose a RoCE configuration (that is, a profile) from the drop-down menu below the port table.

- Choose the RoCE-2 profile for Windows SMB Direct or Linux iSER, NFS, or VM Migration support.
- For RoCE-1, check the Implementer's Lab on the Emulex website for any updated information on additional use.

**Note:** RoCE configurations are not supported with SR-IOV.

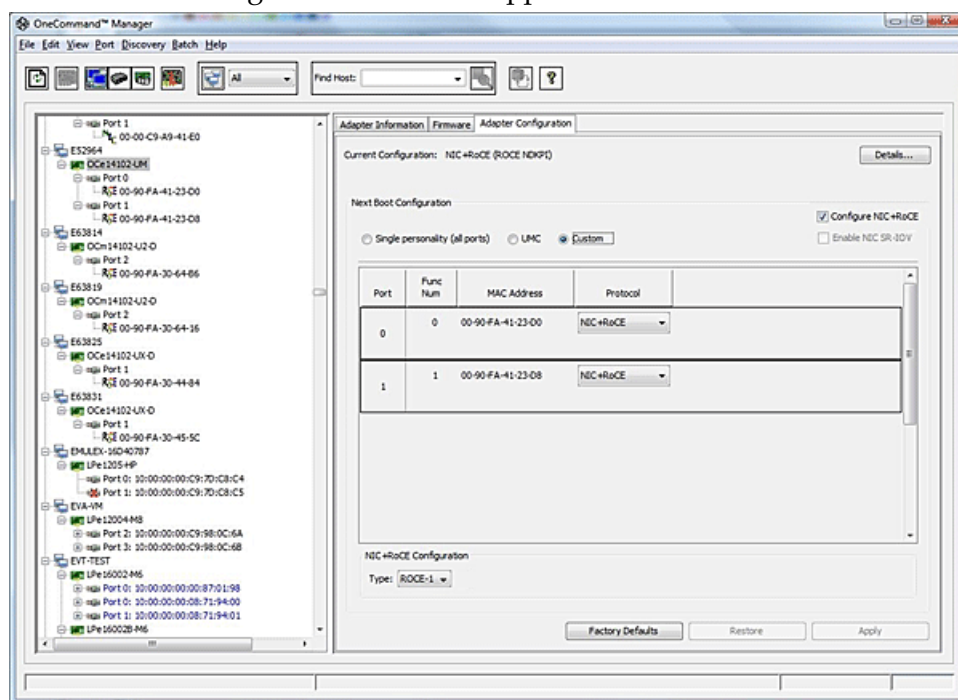
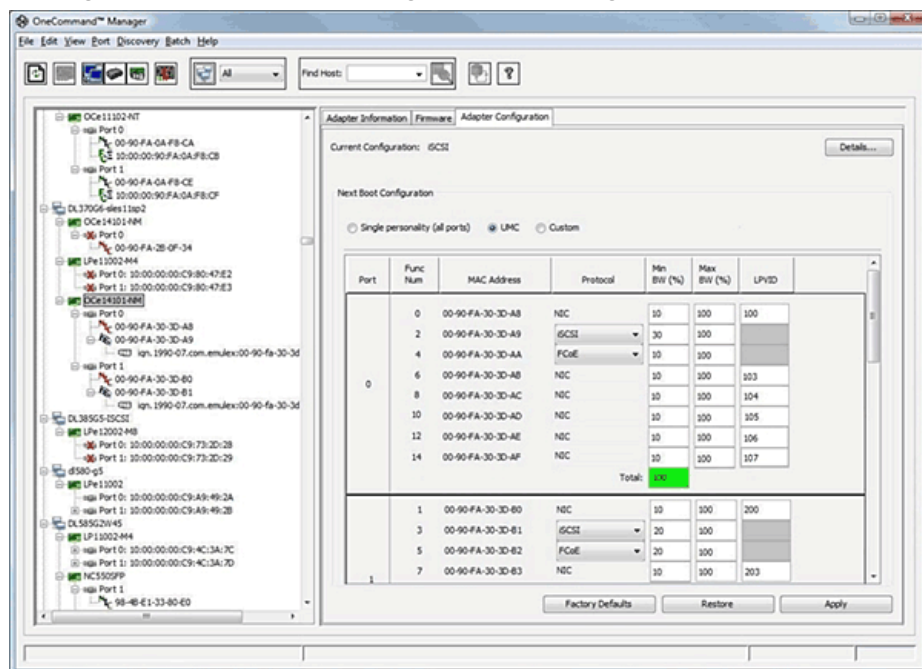


Figure 7-27 Custom NIC+RoCE Configuration

## UMC Configuration View

When UMC is available and the UMC radio button is selected, the Adapter Configuration tab looks like Figure 7-28 on page 98.



**Figure 7-28** UMC View Adapter Configuration Tab (UMC, 2 ports, 8 functions/port, concurrent storage)

For UMC, the OneCommand Manager application allows you to configure up to sixteen functions on a single port adapter, up to eight functions per port on a two-port adapter, and up to four functions per port on a four-port adapter. The maximum number of functions allowed on an adapter is controlled by the adapter's IPL file.

### Notes

- ARI must be available to support up to sixteen functions on an adapter. OCe14000-series adapters automatically support ARI. However, the system's motherboard must support ARI, it must be enabled in the system BIOS, and the operating system must support ARI. If these conditions are not met, although you can configure all sixteen functions, only eight functions will be present and discovered by the OneCommand Manager application after a reboot.
- SR-IOV is not supported with UMC.

UMC channel protocol assignments are subject to the following rules:

- The first channel is always the NIC protocol.
- NIC+RoCE cannot be configured with UMC.
- The second and third channels can be assigned the following:
  - NIC
  - None (channel disabled)
  - Storage protocol

- ◆ (Note: On the third channel only if concurrent mode storage is available.)
- If storage protocols are assigned to the second and third channels, they must be different storage protocols.
- The fourth and higher channels on a port (if available) can only be assigned NIC or None.

Bandwidth and LPVID assignments are subject to the following rules:

- A channel's logical link can be set to down by setting both the minimum and maximum bandwidths to 0.
- The total of the minimum bandwidths for all enabled non-NIC+RoCE channels on a port must add up to 100 percent, except in the case where all the maximum bandwidths are also 0 (that is, setting logical link down on all channels).
- The maximum bandwidth for each channel must be greater than or equal to the minimum bandwidth for that channel.
- LPVIDs (range 2-4094) must be assigned to all NIC and NIC+RoCE channels and must be unique across all NIC/RoCE channels on a port. However, any or all of the LPVIDs on one port can be repeated on another port.
- After changing the LPVIDs or bandwidths on any of channels and clicking **Apply**, the changes take effect immediately. No reboot is required.

## Default UMC Settings

Since most hosts today do not support ARI, the default UMC settings are designed, where possible, to prevent enabling functions that require ARI. This prevents you from enabling channels that will not be seen on the PCI bus when the system boots.

### Notes

- On a four-port adapter, if UMC is disabled, if the third function on each port was already enabled when UMC is enabled, that function will still be enabled. This case assumes that you enabled the third function previously because ARI was available.
- SR-IOV is not supported with UMC.

If UMC is enabled, the channels are assigned default values based upon the following rules:

For each port:

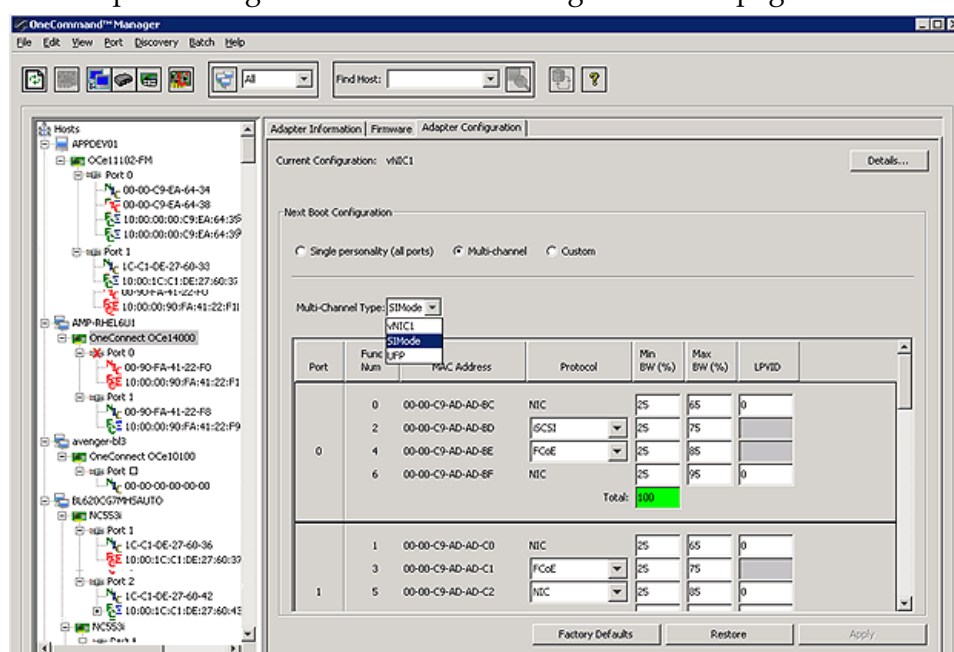
1. All enabled functions (that is, functions not set to "None") will continue to be enabled running the same protocol before UMC was enabled even if its function number is greater than 7 (only possible on four-port adapters).
2. Any channel with a function number less than 8 will be enabled with NIC.
3. Any channel with function number greater than 7 will be set to "None" (that is, channel disabled).
4. The minimum bandwidths will be assigned by dividing up 100% as evenly as possible between the enabled channels.
5. The maximum bandwidths for the enabled channels will be set to 100.

6. Both the minimum and maximum bandwidth for the channels assigned to “None” will be set to 0.
7. The LPVIDs will be set to 0. You must assign a valid LPVID to the NIC channel.

**Note:** These rules create the default settings for the channels. You can change these settings before saving them.

## Lenovo System X Multi-Channel Configuration View

For Lenovo multi-channel configuration, there are three different multi-channel types available: vNIC, SIMode, and UFP. When you check the Multi-channel radio button, the Adapter Configuration tab looks like Figure 7-29 on page 100.



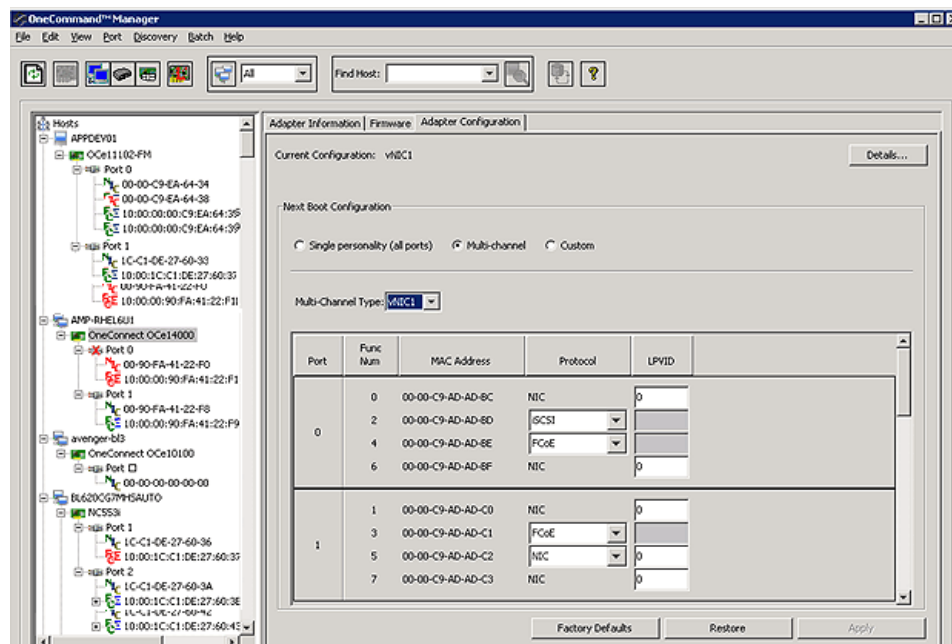
**Figure 7-29** Multi-channel View (showing multi-channel type drop-down)

A Multi-channel Type drop-down menu is available to select the type of multi-channel configuration on the adapter. The contents of the Channel Configuration table depend upon the selected multi-channel type.

**Note:** If Lenovo System X multi-channel is enabled, RoCE cannot be configured on any function.

## vNIC Configuration

When vNIC is selected on the adapter, the Adapter Configuration tab looks like Figure 7-30 on page 101.



**Figure 7-30** Multi-channel vNIC View (mix mode storage)

The Channel Configuration table shows the channel properties for vNIC. Additionally, the protocol can be configured on the second and third channels including NIC+RoCE.

The NIC channels can be assigned an LPVID (or inner VLAN ID). It can be assigned any value from 2-4094. The outer VLAN ID is assigned by the switch and can be viewed by clicking **Details**.

**Note:** After changing the LPVIDs on any of the channels and clicking **Apply**, the changes take effect immediately. No reboot is required.

## vNIC Configuration Rules

Drop-down menus for each port's second and third channels are available. vNIC channel protocol assignments are subject to the following rules:

- The first channel is always NIC.
- NIC+RoCE cannot be configured with vNIC.
- The second and third channels can be assigned the following:
  - NIC
  - None (channel disabled)
  - Storage protocol

**Note:** On the third channel only when concurrent mode storage is available.

- A maximum of four channels per port can be configured.
- The fourth channel of a port (if available) can only be assigned NIC or None.
- LPVIDs (range 2-4094) must be assigned to all NIC channels and must be unique across all NIC channels on a port. However, any or all of the LPVIDs on one port can be repeated on another port.

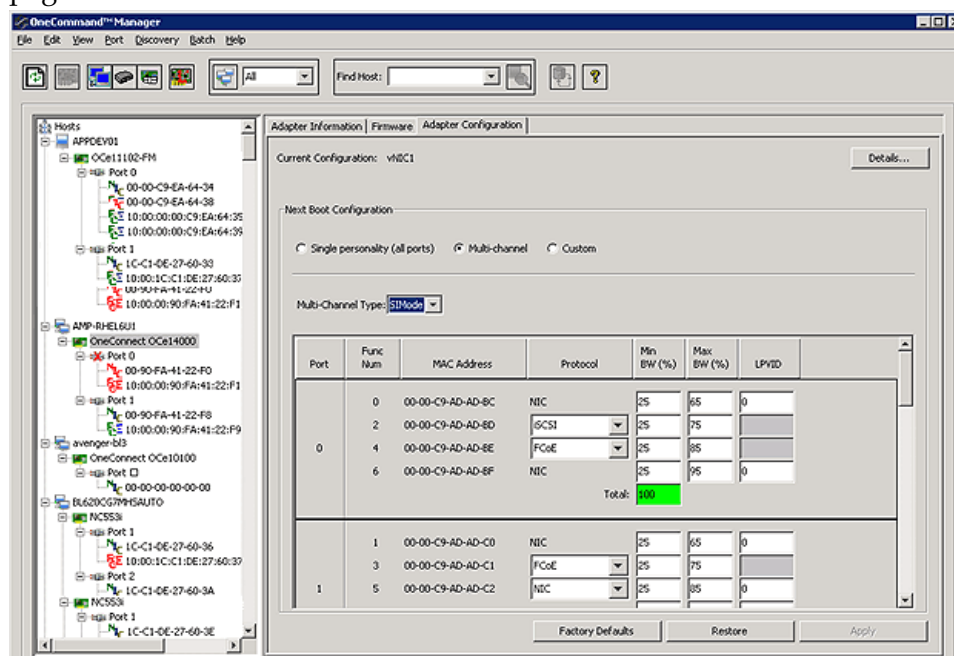
The maximum number of functions per port is shown in Table 7-1 on page 102.

**Table 7-1** vNIC Maximum number of functions per port

Number of Ports	Port Speed	Default functions per port	Maximum functions per port
1	10Gbps	4	4
2	10Gbps	4	4
4	10Gbps	4	4

## SIMode Configuration

When SIMode is selected, the Adapter Configuration tab looks like Figure 7-31 on page 102.



**Figure 7-31** Multi-channel SIMode View

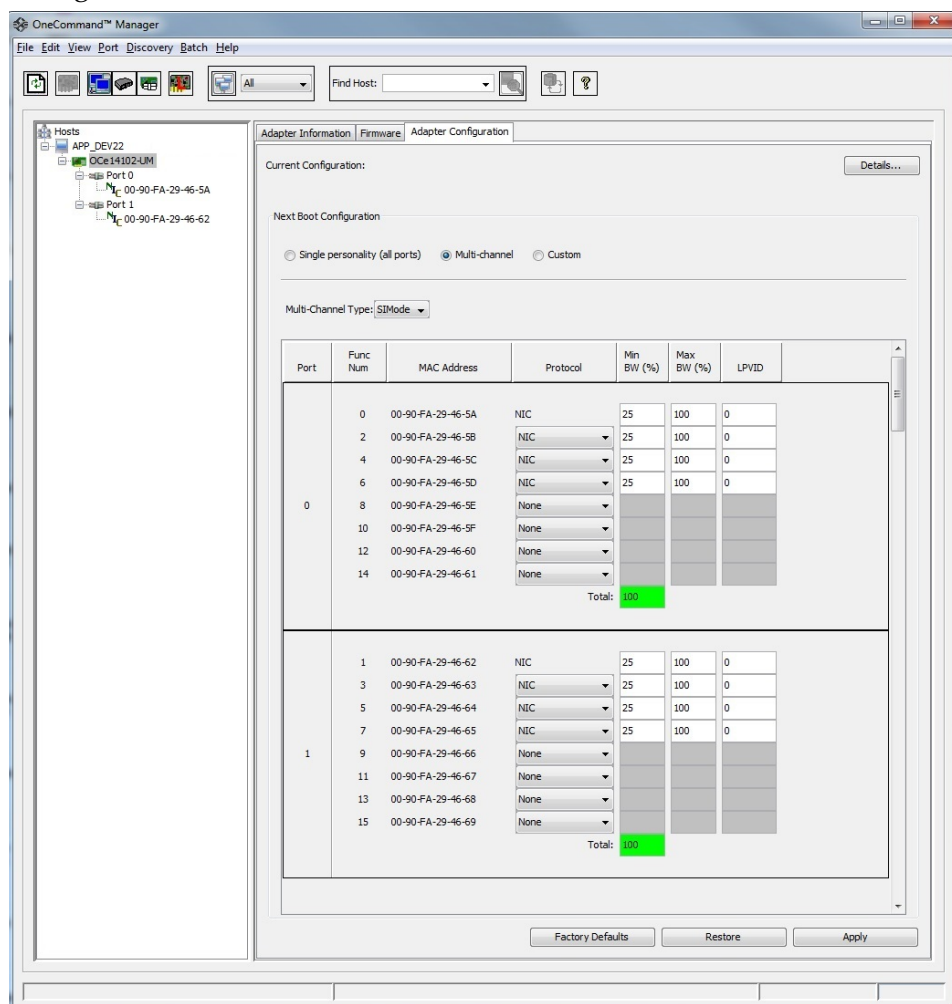
SIMode and UMC configuration are exactly the same. See “UMC Configuration View” on page 98 for more information.

For two-port adapters configured to run SIMode, up to eight functions can be configured on each port.

When switching from a non-multi-channel to a SIMode configuration, the default configuration presented in the tab of the channels is four channels set to NIC with minimum and maximum bandwidths set to 25 and 100 respectively. The upper four



channels on each port are set to “None”. Figure 7-32 on page 103 shows this configuration.



**Figure 7-32** 2-Port Lenovo System X Adapter Configured for SIMODE (with 8 functions per port)

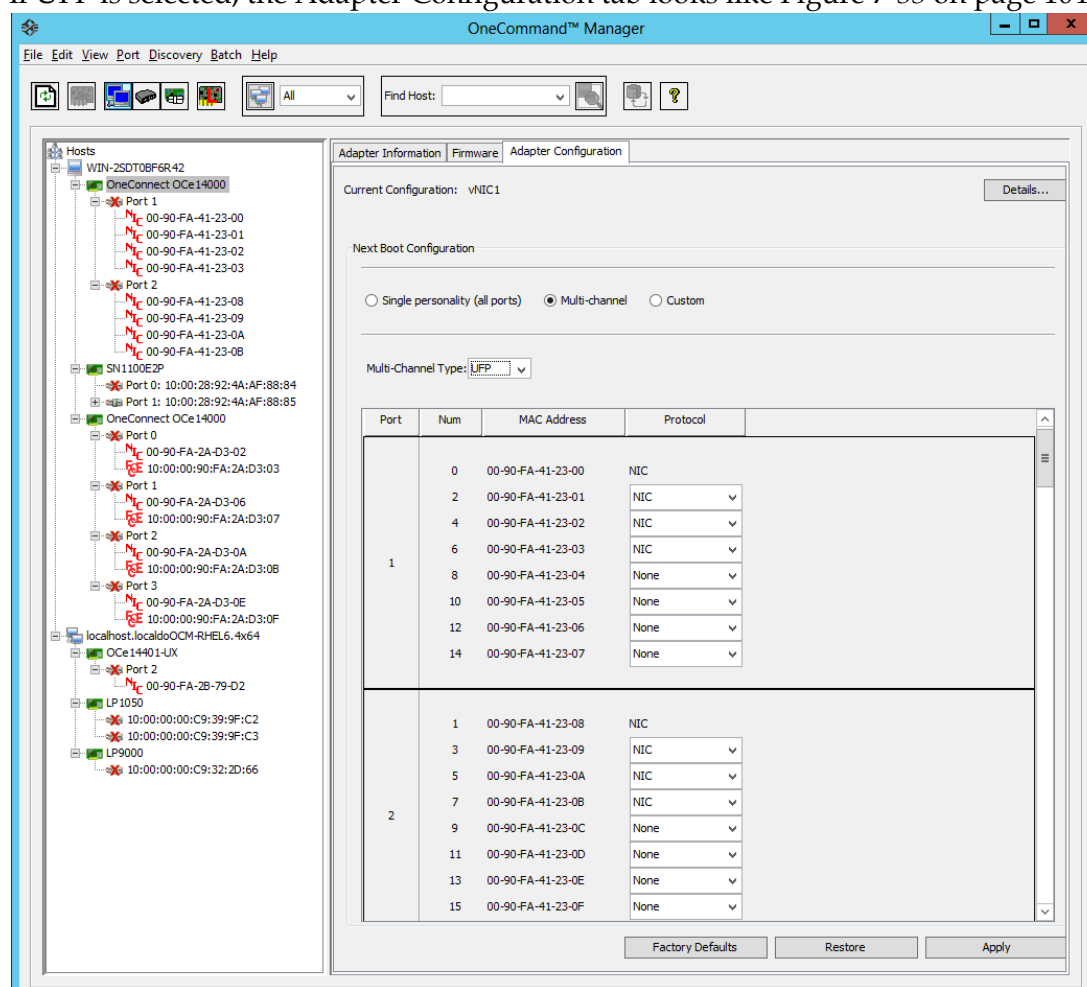
The maximum number of functions per port is shown in Table 7-2 on page 103.

**Table 7-2** SIMode Maximum number of functions per port

Number of Ports	Port Speed	Default functions per port	Maximum functions per port
1	10Gbps	4	8
2	10Gbps	4	8
4	10Gbps	4	4
1	40Gbps	4	16

## UFP Configuration

If UFP is selected, the Adapter Configuration tab looks like Figure 7-33 on page 104.



**Figure 7-33** Multi-channel UFP View (concurrent mode storage)

The Channel Configuration table shows the channel properties for UFP. Additionally, the protocol can be configured on the second and third channels

UFP switch support provides two modes of QoS determined by the switch on a per-port basis:

- ETS mode – uses switch-specified bandwidths based on priority groups.
- Minimum bandwidth mode – uses switch-specified minimum guaranteed bandwidths for each channel on a port.

You can configure a variable number of functions per port, from one up to the maximum number supported by the port. A different number of functions per port can be configured on the same adapter.

For example, you can configure 4 functions per port on one port and 6 functions per port on another. If you configure 6 functions per port on both ports of a two-port adapter, then 12 functions are enumerated on the PCI bus (not 16).



When UFP is enabled on an adapter, the default is 4 functions per port. The maximum number of functions per port is shown in Table 7-3 on page 105.

**Table 7-3** UFP maximum number of functions per port

Number of Ports	Port Speed	Default functions per port	Maximum functions per port
1	10Gbps	4	8
2	10Gbps	4	8
4	10Gbps	4	4
1	40Gbps	4	16

### UFP Configuration Rules

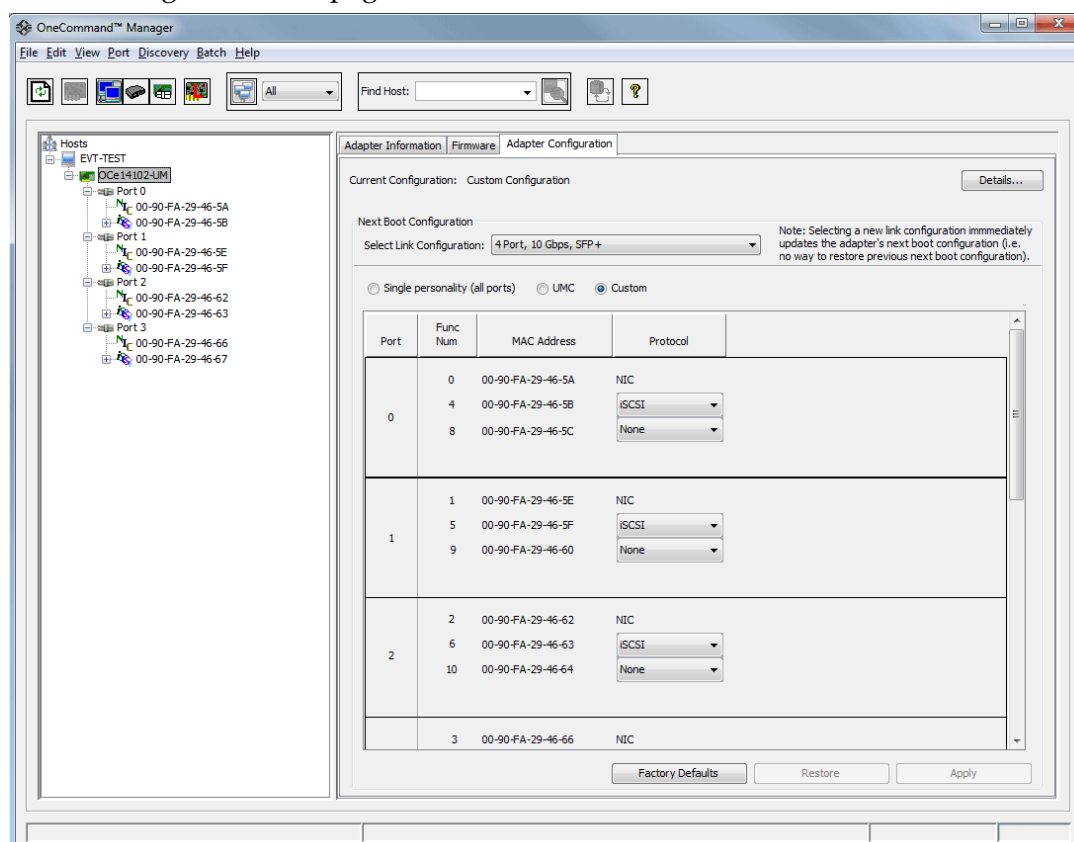
Menus for each port's second and third channels are available. UFP channel protocol assignments are subject to the following rules:

- The first channel is always NIC.
- NIC+RoCE cannot be configured with UFP.
- The second and third channels can be assigned the following:
  - NIC
  - None (channel disabled)
  - Storage protocol
    - ◆ (Note: On the third channel only when concurrent mode storage is available.)
- The fourth and higher channels of a port (if available) can be assigned NIC or None.

## Link Re-Configurability

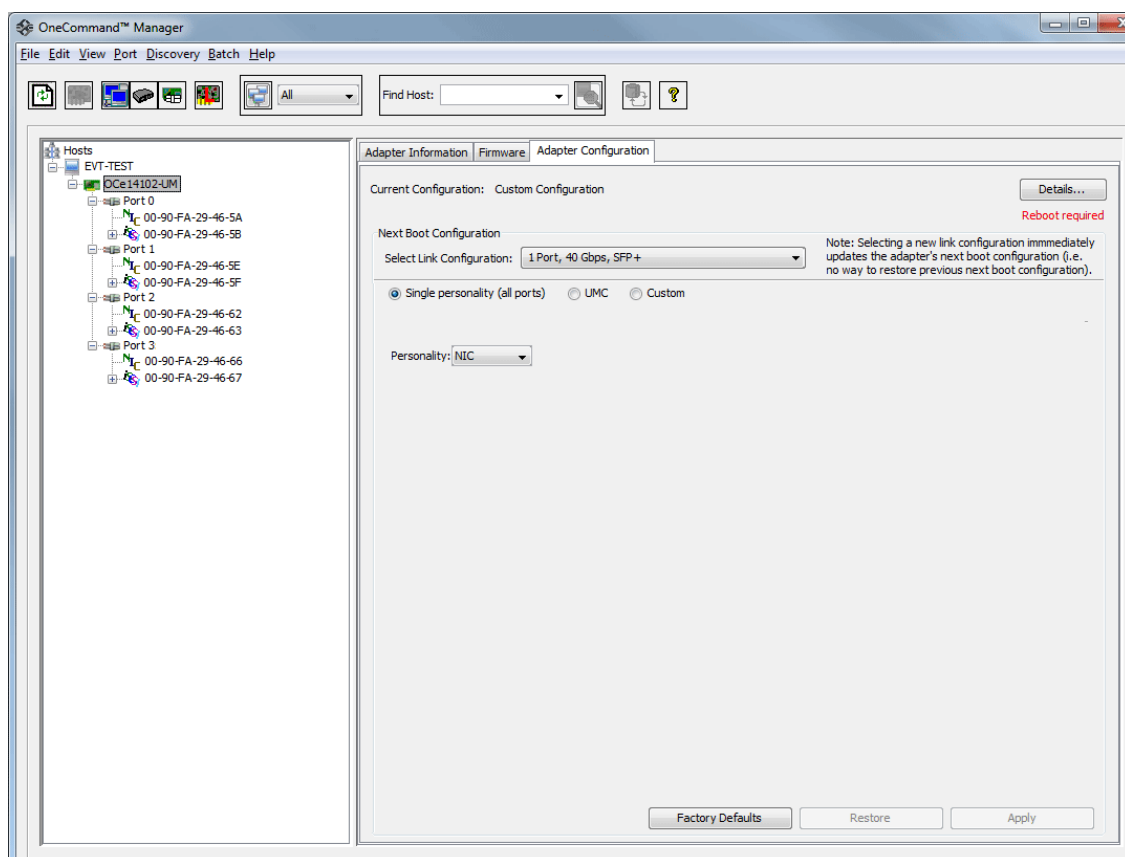
For adapters that support it, Link Re-configurability allows you to adjust the number of ports and port speeds on the adapter by selecting a specific Link Re-configurable setting. The adapter provides a list of Link Re-configurable descriptions (for example, one port, 40Gbps, SFP+) and Link IDs.

For supported adapters, the Select Link Configuration menu on the Adapter Information tab enables you to select the link configuration. An example of this is shown in Figure 7-34 on page 106.



**Figure 7-34** Adapter Configuration Tab Before Changing Link Configuration

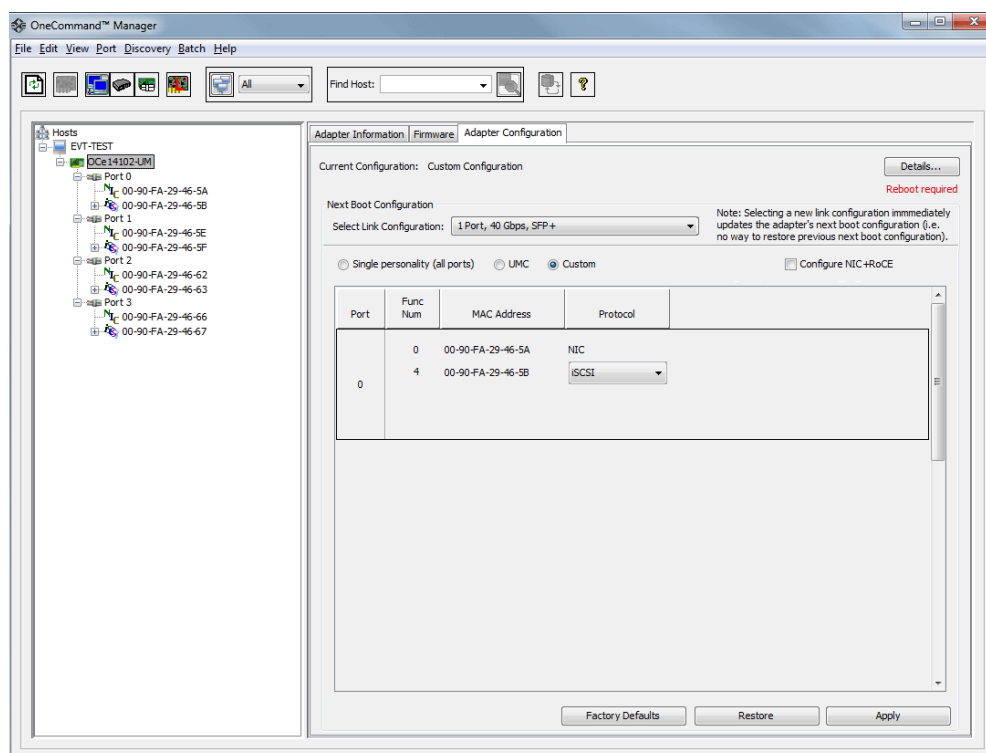
Any changes to the Select Link Configuration setting take effect immediately; you do not need to click Apply. After changing the link configuration, you can see the new port configuration and change the functions running on that port before performing a reboot to activate the new link and port configuration. For example, after changing from a four-port to a one-port configuration, the Adapter Configuration tab looks like Figure 7-35 on page 107.



**Figure 7-35** Adapter Configuration Tab After Selecting a New Link Configuration

The adapter's configured port settings and profile are changed to the default values defined in the adapter's IPL for the selected link configuration. The next boot configuration in the tab reflects these default settings (in Figure 7-35 on page 107 the default is NIC-only). The "Reboot required" message indicates that you must reboot the system to activate the new configuration. You cannot restore the previously saved next boot configuration (the Restore button is disabled).

Before rebooting, you can change the default port configuration that was set on the adapter by changing the link configuration. For example, you can change the single port configuration from a NIC-only configuration to NIC+iSCSI as shown in Figure 7-36 on page 108.



**Figure 7-36** Adapter Configuration Tab After Changing Port Configuration for New Link Configuration

The Apply button at the bottom of the tab must still be clicked to save the updated port configuration (in this case, adding the iSCSI function). If a reboot is performed before clicking **Apply**, the default port configuration for the new link configuration is activated instead of the updated port configuration. After applying the port configuration changes and rebooting the system, the discovery-tree and Adapter Configuration tab should look like Figure 7-37 on page 109.

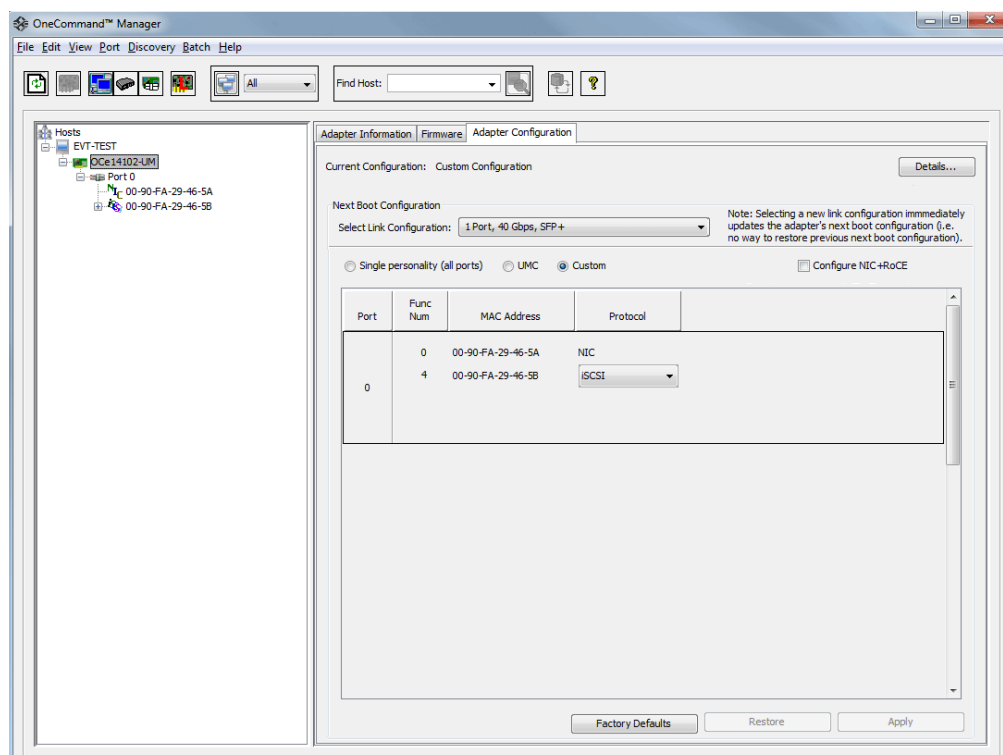
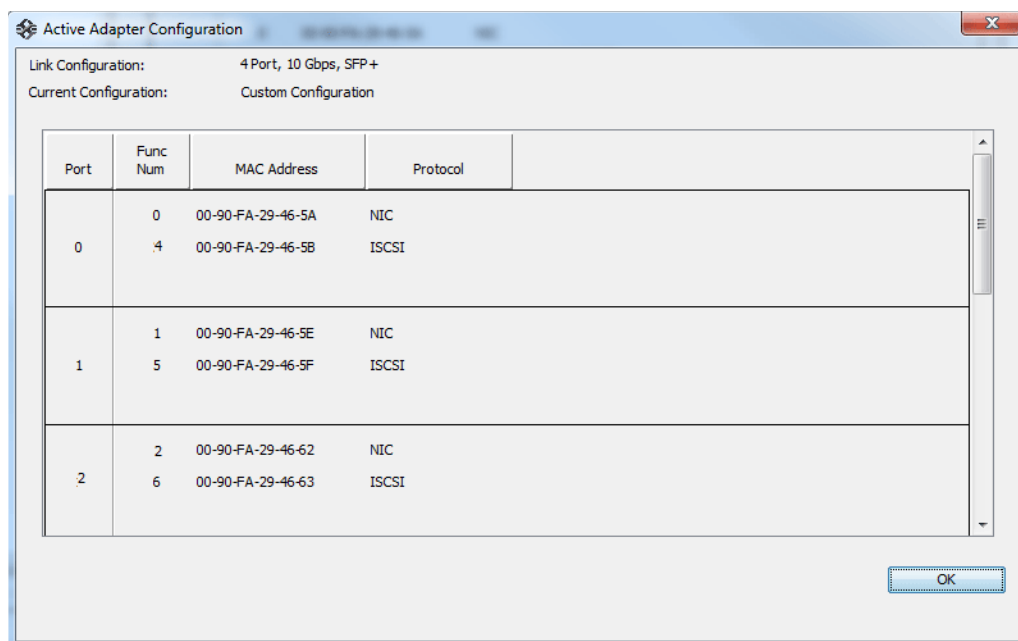


Figure 7-37 Adapter Configuration Tab After Reboot with Single Port Configuration

## Adapter Configuration Buttons with Link Re-Configurability

- **Factory Defaults** – Clicking the Factory Defaults button not only resets the adapter to the factory default port settings, but also returns the adapter to its default Link Configuration. For example, if the current adapter's link configuration was for four ports, but the factory default link configuration for the adapter is one port, clicking **Factory Defaults** changes the adapter's link configuration to one port.
- **Restore** – The Restore button changes the port configuration back to the currently saved configuration. Since changing the link configuration immediately updates the adapter's next boot configuration, clicking the **Restore** button does not change the adapter back to the previous link configuration. Therefore, after selecting a new link configuration, the Restore button is disabled.
- **Apply** – There is no change to how the Apply button operates when Link Re-configurability is available on the adapter. Since selecting a new link configuration results in an immediate update to the adapter, the Apply button is not needed to save the new link configuration setting. It is disabled after selecting a new link configuration.
- **Details** – When the **Details** button is clicked, details about the active configuration are shown in a pop-up window. If Link Re-configurability is available, the current link configuration description is displayed at the top of the window as shown in Figure 7-38 on page 110.



**Figure 7-38** Active Adapter Configuration Window Showing Active Link Configuration

## Dell NPar Configuration View

**Note:** For NPar, functions are referred to as partitions.

For Dell adapters, the Adapter Configuration tab (Figure 7-39 on page 111) is different than other adapters. Single personality, multi-channel, or custom configurations are not available in the Next Boot Configuration. Instead, there is a single checkbox to enable or disable NPar and an NPar Configuration table to configure NPar partition protocol and bandwidth assignments.

If NPar is enabled, the OneCommand Manager application allows you (subject to restrictions in the adapter's IPL) to configure up to sixteen functions on a single port adapter and up to eight partitions per port on a two-port adapter.

If NPar is enabled, the minimum bandwidth for each partition is the same or nearly the same for all partitions. For four-partition ports, the minimum bandwidth for each partition is 25%. For eight-partition ports, the first four partitions are allocated 13% and the last four partitions are allocated 12%. The maximum bandwidth for all partitions will default to 100%. You can change the bandwidths before applying the changes. However, the minimum bandwidths must add up to 100%.

**Note:** NParEP must be available to support up to sixteen functions on an adapter. OCe14000-series adapters automatically support NParEP. However, the system's motherboard must support NParEP, it must be enabled in the system BIOS, and the operating system must support NParEP. On Linux and VMware systems, BIOS SR-IOV must be enabled. See the Dell instructions for enabling BIOS SR-IOV.

If these conditions are not met, although you may configure all sixteen functions, only eight functions will be present and discovered by the OneCommand Manager application after a reboot.

For storage protocols, any of the second through fourth partitions can be configured to run storage. However, if more than one storage protocol is configured, they must be different storage protocols. The same storage protocol cannot be run on two partitions on the same port.

If NPar is disabled (NPAR Enabled checkbox is unchecked), only a single partition per port is available. This partition can be configured to run NIC or NIC+RoCE.

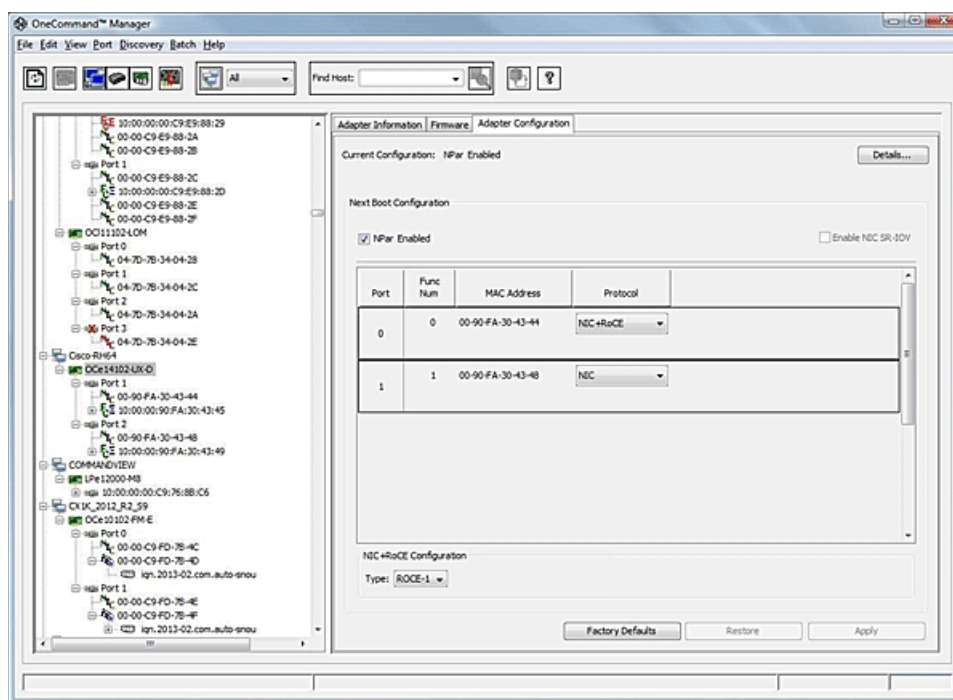


Figure 7-39 Adapter Configuration Tab for NPar Adapters (NPar disabled)

**Note:** SR-IOV can be enabled on NPar adapters only when NPar is disabled and all ports are configured to run NIC functions.

## Configuring RoCE on NPar Adapters

You can configure NIC+RoCE on NPar adapters if NPar is disabled (Figure 7-39 on page 111).

Use the Protocol pull-down menu for each port to select NIC or NIC+RoCE. If NIC+RoCE is chosen, you must also select a NIC+RoCE configuration.

- Choose the RoCE-2 profile for Windows SMB Direct or Linux iSER, NFS, or VM Migration support.
- For RoCE-1, check the Implementer's Lab on the Emulex website for any updated information on additional use.

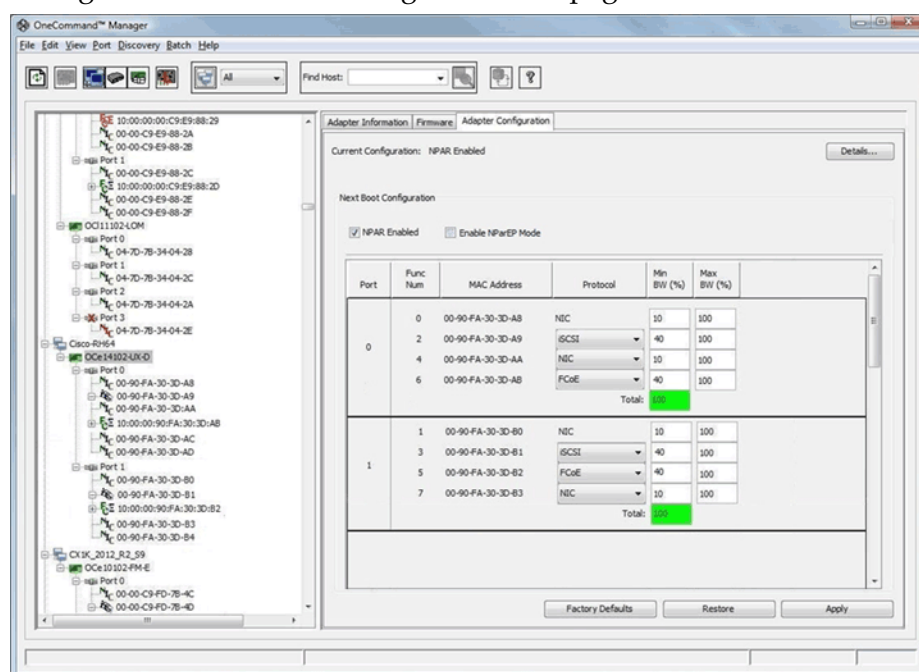
**Note:** RoCE configurations are not supported with SR-IOV.

## Dell NPar Enabled

### Notes

- On Linux and VMware systems, SR-IOV must be enabled on the system BIOS if NParEP is used. See the documentation that accompanied your Dell server for more information.
- NParEP support is available only on Dell 13G or newer systems.

SR-IOV is available when NPar is enabled. If NPar is enabled, the Adapter Configuration tab looks like Figure 7-40 on page 112.



**Figure 7-40** Adapter Configuration Tab with NPar Enabled (NParEP mode disabled)

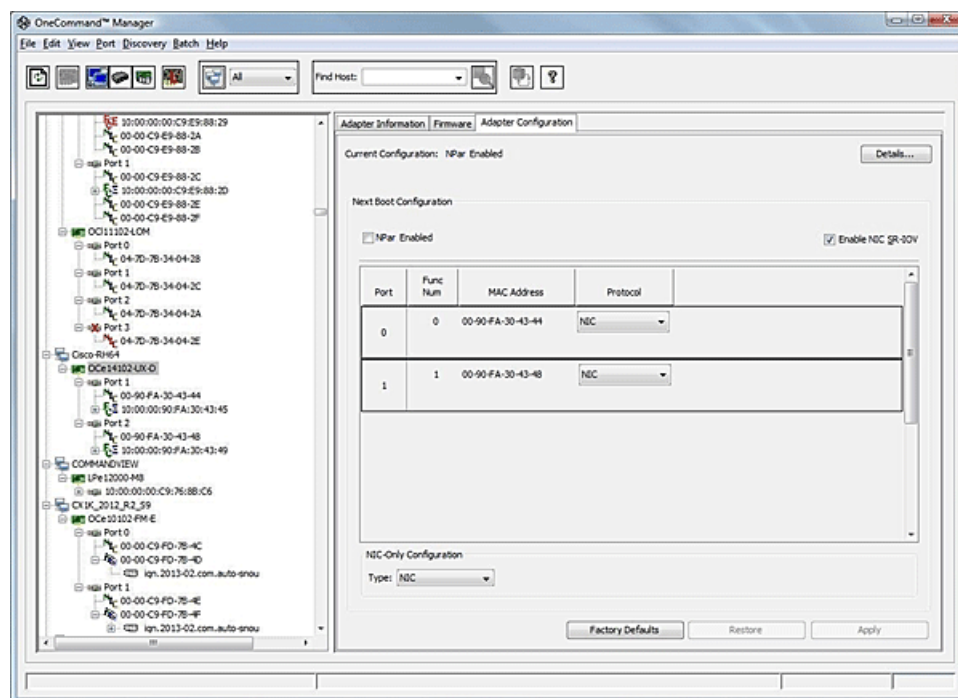
For adapters that support sixteen functions, on two port adapters you can configure up to eight partitions per port. Four-port adapters allow up to four partitions per port. The “None” setting can be configured on any but the first partition.

**Note:** If NPar is enabled, RoCE cannot be configured on any partition.

SR-IOV can be enabled on NPar adapters only when NPar is disabled and all ports are configured to run NIC functions. The Adapter Configuration tab looks like Figure 7-41 on page 113.

Select the Enable NIC SR-IOV checkbox to enable SR-IOV on the NIC function on each port. When Enable NIC SR-IOV is checked, the NIC-Only Configuration Type field may be displayed where you can select the NIC profile.

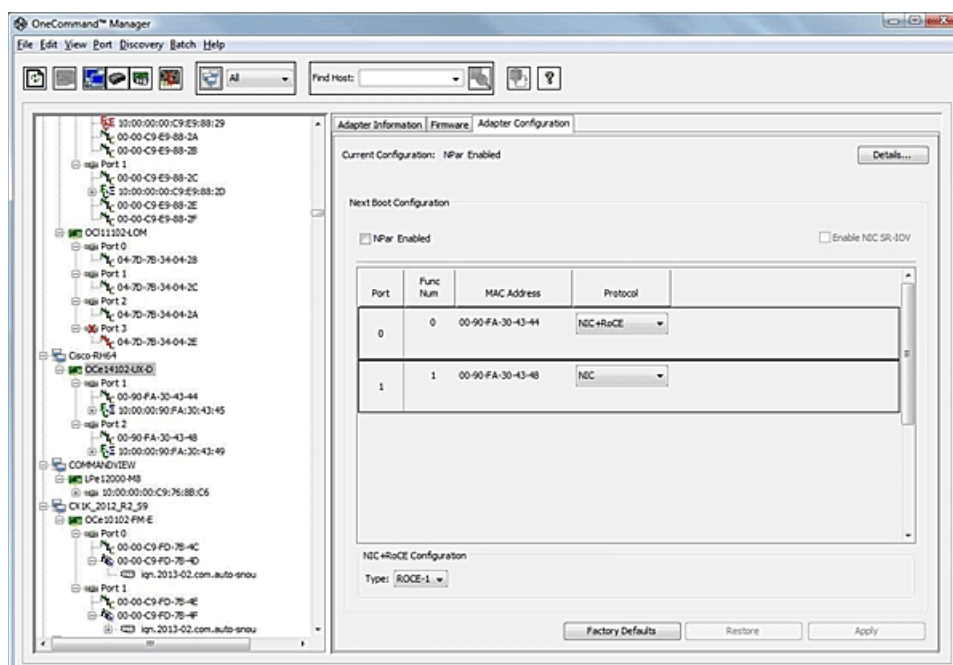




**Figure 7-41** Adapter Configuration Tab with NPar and NIC SR-IOV Enabled

When NPar is disabled and one or more ports are configured for NIC+RoCE functions, the Adapter Configuration tab looks like Figure 7-42 on page 114.

**Note:** The Enable NIC SR-IOV checkbox is unavailable. If you selected the checkbox previously and you change one of the ports to NIC+RoCE, the checkbox will be automatically unchecked and disabled.



**Figure 7-42** Adapter Configuration Tab with NPar Disabled with a Single NIC+RoCE Port

If you have an active configuration where only some of the NIC functions have SR-IOV enabled and you click the Adapter Configuration tab, a warning message appears, as shown in Figure 7-20 on page 92, prompting you to click the **Apply** button to save the configuration with SR-IOV enabled on all NIC functions.

If you have at least one of the NIC functions configured for SR-IOV, the Enable SR-IOV checkbox is checked. You must click **Apply** in the Adapter Configuration tab to save the SR-IOV changes to all NIC functions. You can also clear the SR-IOV checkbox and click **Apply** to disable SR-IOV. However, if you do nothing, each time the tab is selected, the warning message is displayed.

**Note:** If an adapter on a remote host is selected which is running OneCommand Manager, the SR-IOV state in the Active Configuration window is displayed as N/A when you display the current configuration using the **Details** button. See Figure 7-21 on page 92. However, you can enable or disable SR-IOV on the entire adapter on a OneCommand Manager host by checking or unchecking the Enable NIC SR-IOV checkbox and clicking the **Apply** button.

## Dell NPar and NParEP Mode Enabled

### Notes

- On Linux and VMware systems, SR-IOV must be enabled on the system BIOS if NParEP is used. See the documentation that accompanied your Dell server for more information.
- NParEP support is available only on Dell 13G or newer systems.

NPar adapters have an NParEP mode setting that changes the total number of partitions displayed and configured (between eight and sixteen) by the OneCommand Manager application. If NParEP Mode is enabled, you can view and configure sixteen partitions per adapter.

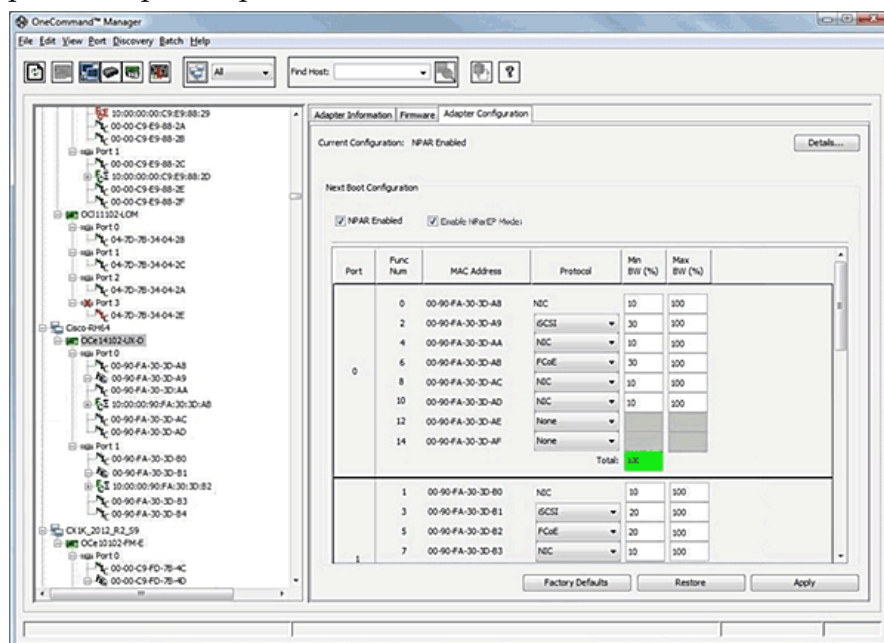


Figure 7-43 Adapter Configuration Tab with NPar and NParEP Mode Enabled (2-port configuration)

If NParEP Mode is enabled on two port adapters, up to eight partitions per port can be configured (Figure 7-43 on page 115).

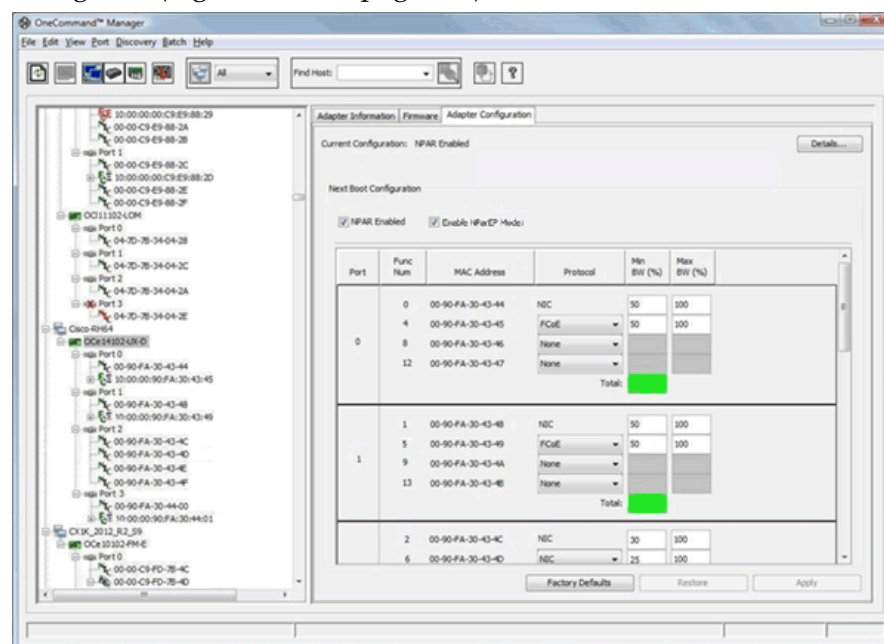


Figure 7-44 Adapter Configuration Tab with NPar and NParEP Mode Enabled (4-port configuration)

Four-port adapters allow up to four partitions per port to be configured if NParEP Mode is enabled (Figure 7-44 on page 115).

NPar protocol assignments are subject to the following rules:

**Note:** These rules apply if NParEP mode is enabled or disabled.

- The first partition on a port is always NIC.
- Up to sixteen partitions can be configured on NPar adapters if NParEP Mode is enabled.
- If NPar is enabled, RoCE cannot be configured on any partition.
- A storage protocol can be assigned to any of the second through fourth partitions. Any partitions above the fourth partition can only be assigned NIC or None.
- Any partition other than the first partitions on a port can be set to NIC or None.
- The total of the minimum bandwidths of the enabled functions for each port must add up to 100 percent. The maximum bandwidth for each partition must be greater than, or equal to, the minimum bandwidth for that partition.
- After changing the NPar configuration and applying the changes, except for bandwidth changes, a reboot is required to activate the new configuration.

## NPar Bandwidth Settings

Similar to UMC, a minimum and maximum bandwidth is assigned to each enabled (that is, not set to None) partition. The sum of the minimum bandwidths must add up to 100 percent. The maximum bandwidth for each partition must be equal to or greater than the minimum bandwidth on that partition and no larger than 100 percent.

Unlike UMC, if the minimum and maximum bandwidths are set to 0, the logical link on the adapter is not brought down. A small amount of data will pass through the port for that partition if bandwidth is available. Bandwidth changes take effect immediately. A reboot is not required.

**Note:** When ETS Priority Groups are enabled and available from the switch, the bandwidths in those priority groups are used instead of the bandwidths specified for each NPar partition.

## 8. Managing Protocols

This sections describes how to manage the available protocols.

### NIC

This section describes managing NIC functions.

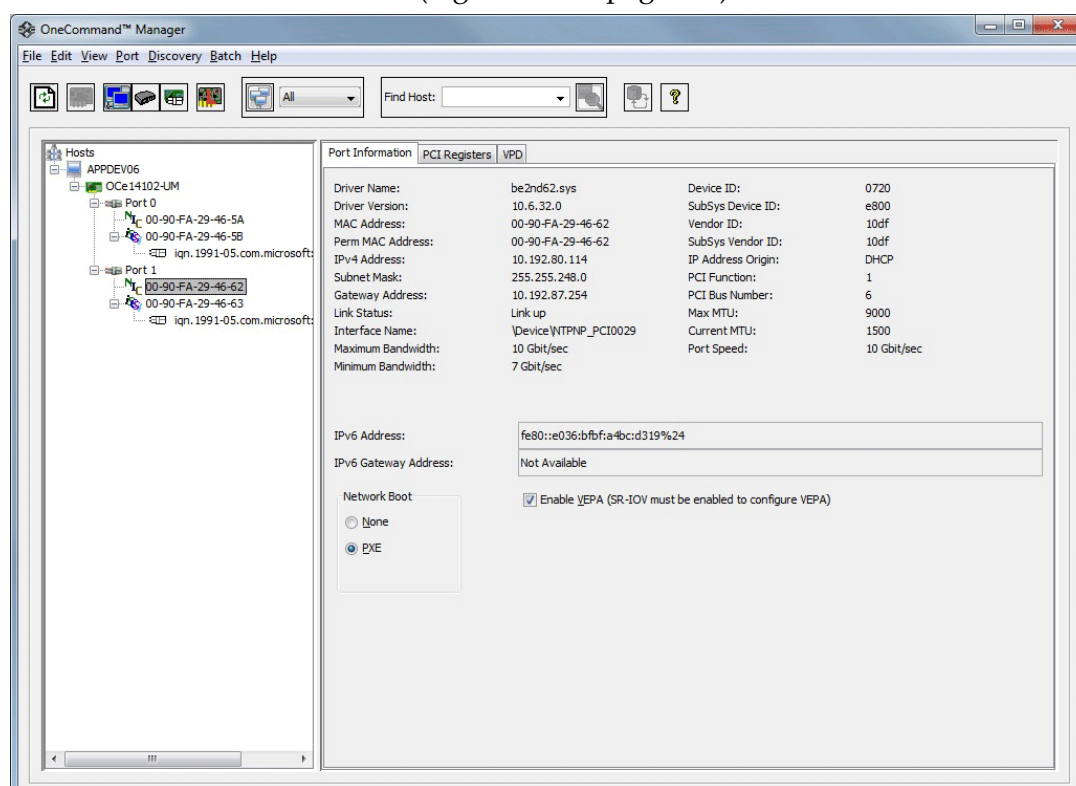
### Viewing NIC Function Information

When you select a NIC function from the discovery-tree (Figure 4-3 on page 40), the NIC Port Information tab (Figure 8-1 on page 117) contains general attributes associated with the selected NIC function.

The NIC Port Information tab (Figure 8-1 on page 117) also allows you to enable or disable network boot on the selected port.

To view general NIC function information:

1. Select **Host** view.
2. Select NIC function in the discovery-tree (Figure 4-3 on page 40).
3. Select the **Port Information** tab (Figure 8-1 on page 117).



**Figure 8-1** NIC Port Information Tab

## NIC Port Information Field Definitions

- Driver Name – The NIC driver file name.
- Driver Version – The NIC driver version.
- MAC Address – The NIC MAC address currently assigned to the NIC function.
- Perm MAC Address – The original factory assigned NIC MAC address.
- IPv4 Address – The IPv4 address for the NIC function.
- Subnet Mask – The subnet mask for the NIC function.
- Gateway Address – The NIC initiator gateway address.
- Link Status – The status of the link on the selected NIC function.
- Interface Name – The interface assigned to this NIC function by the host operating system.
- Maximum Bandwidth – The maximum bandwidth (that is, speed) at which the NIC function is guaranteed to run.
- Minimum Bandwidth – The minimum bandwidth (that is, speed) at which the NIC function is guaranteed to run.
- Device ID – The PCI device ID assigned to the NIC function.
- Subsys Device ID – The PCI subsystem ID assigned to the NIC function.
- Vendor ID – The PCI vendor ID assigned to the NIC function.
- Subsys Vendor ID – The PCI subsystem vendor ID assigned to the NIC function.
- IP Address Origin – The origin of the IP address (DHCP or Static).
- PCI Function – The PCI function number assigned to the NIC function.
- PCI Bus Number – The PCI BUS number assigned to the NIC function.
- Max MTU – The maximum transmission unit for iSCSI traffic.
- Current MTU – The current transmission unit for iSCSI traffic.
- Port Speed – The current port speed of the selected port.
- Bandwidth Limit – The QoS bandwidth restriction on the NIC function. (Non vNIC adapters only)
- IPv6 Address – The IPv6 address for the NIC function.
- IPv6 Gateway Address – The IPv6 gateway address for the NIC function.

## Network Boot

- None – Select this option to disable network boot on the selected port.
- PXE – Select this option to enable PXE boot on the selected NIC function.

**Note:** PXE boot is only supported on the first function of each physical port on the adapter.

- iBFT – If supported by the adapter, select this option to enable iBFT on the selected NIC function.

## Checkboxes

- Enable VEPA – Check to enable VEPA. Clear to disable VEPA.

**Note:** The feature is only available on OCe14000-series adapter NIC functions if SR-IOV is enabled on the NIC function. If SR-IOV is not enabled, this checkbox is not visible.

## Viewing NIC VPD Information

The VPD tab (Figure 8-2 on page 119) displays vital product data (if available) for the selected NIC function such as the product name, part number, serial number, and so on.

To view VPD information:

1. Select **Host** view.
2. In the discovery-tree, select the NIC function whose VPD information you want to view.
3. Select the **VPD** tab (Figure 8-2 on page 119).




Figure 8-2 NIC VPD Tab

## VPD Table Definitions

- Product Name – Product information about the selected adapter port.
- PN (Part Number) – The adapter's part number.
- SN (Serial Number) – The adapter's serial number.
- VO – Vendor unique data. “V” indicates a vendor-specific field. An adapter may have none, one, or more of these fields defined. Valid values for this field are “VO” (the letter “O”, not the number zero) and “Vx” (where “x” is a number).

**Note:** Some adapters may show additional VPD information such as EC (EC level) and MN (manufacturer ID).



## Enabling and Disabling SR-IOV on NIC Functions (on OCe11000-series adapters)

### Notes

- For OCe14000-series adapters, refer to “OCe14000-Series Adapter Configuration” on page 90.
- This section only applies when running the OneCommand Manager application on supported SR-IOV adapters and operating system platforms.
- SR-IOV is not supported with UMC.

If SR-IOV is available, the NIC Port Information tab (Figure 8-1) allows you to enable SR-IOV on NIC adapter functions of an adapter. If SR-IOV is enabled, multiple VFs can be created on a NIC PF for an adapter port. These VFs are assigned by virtual operating systems such as Microsoft Hyper-V and Linux KVM to virtual machines (VMs). Each VM can be assigned one or more VFs by the guest operating system running on the VM. To the guest operating system, the VF is an independent NIC function with its own MAC address and is available for network I/O.

The discovery-tree (Figure 4-3 on page 40) displays the VFs running on the discovered PFs. It also shows the selected VF's MAC Address, VLAN ID, Link Status, and Transmit Rate.

To enable or disable SR-IOV:

1. Select **Host** view.
2. Select a NIC function in the discovery-tree (Figure 4-3 on page 40).
3. Select the **Port Information** tab (Figure 8-3 on page 121).
4. Check or uncheck **Enable SR-IOV**.
5. Reboot to change the SR-IOV state.



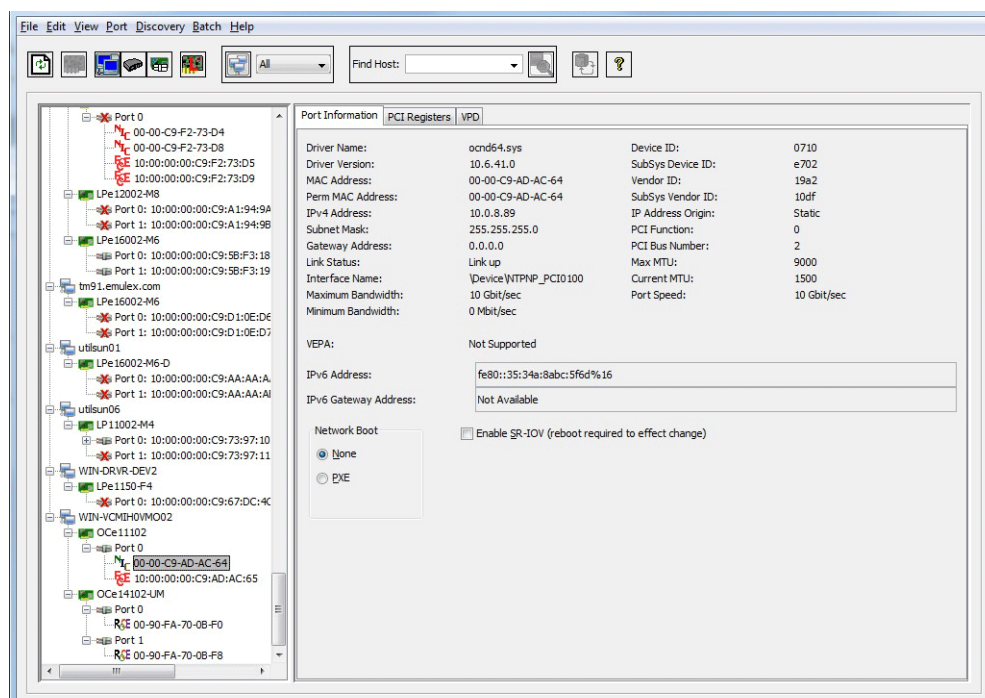


Figure 8-3 Port Information Tab, SR-IOV Enabled

## Enabling and Disabling VEPA on NIC SR-IOV Functions (OCe14000-series adapters only)

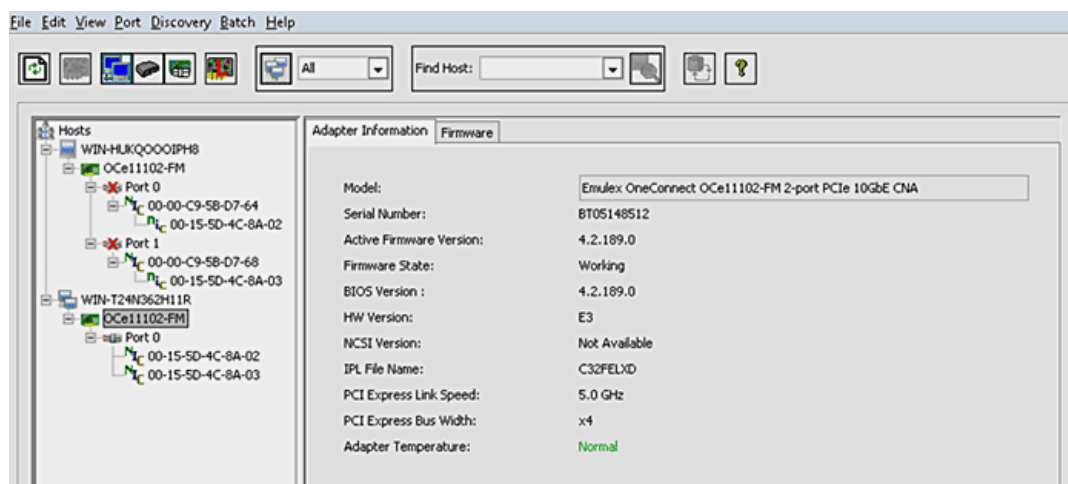
The VEPA checkbox is displayed on the Port Information tab when SR-IOV is currently enabled on the NIC function. Check or uncheck the Enable VEPA checkbox to enable or disable VEPA for the virtual functions on the NIC port. See Figure 8-1 on page 117. If SR-IOV is not enabled on the adapter, the VEPA checkbox is not visible.

**Note:** VEPA is only available on OCe14000-series adapters.

## 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 and VF are discovered by the OneCommand Manager application running on the host operating system. The guest operating system host

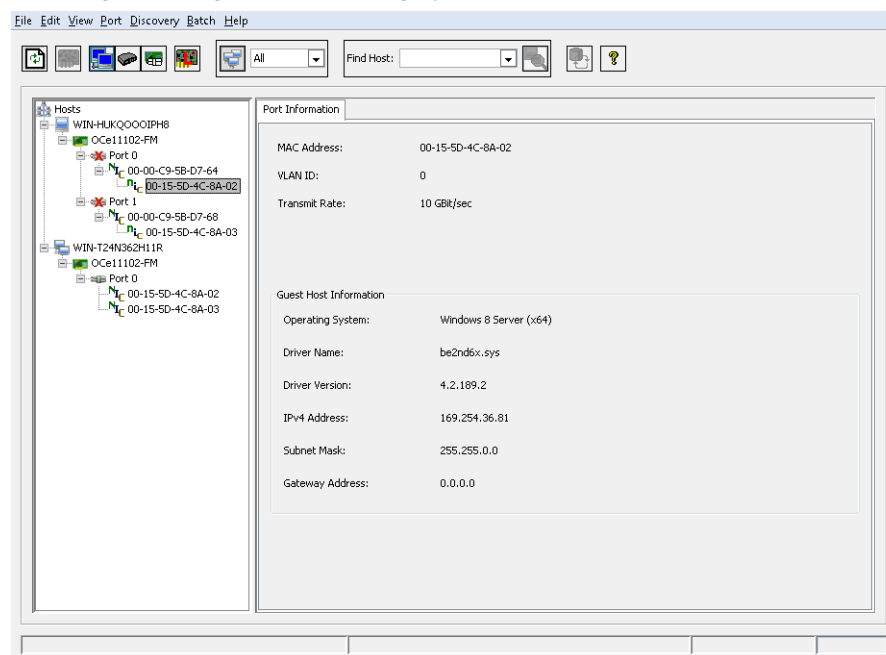
appears as a remote host in the discovery-tree (Figure 8-4 on page 122).



**Figure 8-4** OneCommand Manager Application Running on the Base Host Operating System after Discovering the Guest Host

The NIC function for the discovered guest operating system matches one of the VFs in the base host operating system as seen in Figure 8-5 on page 122.

Select the VF in under the base operating system host in the discovery-tree to display the Port Information tab (Figure 8-5 on page 122). Some of the information displayed in the tab is obtained from the OneCommand Manager remote management agent running on the guest operating system.



**Figure 8-5** VF Selected Showing the Port Information Tab for the Discovered NIC in the Guest Operating System

## Port Information Field Definitions

- MAC Address – The NIC MAC address currently assigned to the port.
- VLAN-ID – The VLAN identifier used by the NIC port.
- Transmit Rate – The rate at which data is transmitted over the port in MBps.

## Guest Host Information Area

**Note:** The following fields are only available when the OneCommand Manager application is installed on the guest operating system running on this VF and the guest host has been discovered (over TCP/IP) by the OneCommand Manager application running on the base host operating system. See “Discovery Using the TCP/IP Access Protocol” on page 48.

- Operating System – The operating system and version installed on the selected host.
- Driver Name – The NIC driver file name.
- Driver Version – The NIC driver version.
- IPv4 Address – The IPv4 address for the NIC port.
- Subnet Mask – The subnet mask for the NIC port.
- Gateway Address – The NIC initiator gateway address.

## Running the OneCommand Manager Application on a Guest Operating System

When the OneCommand Manager application is installed and runs on a VM's guest operating system, it runs in a local-only/read-only mode. Therefore, only the guest host containing the adapter and NIC port is displayed in the discovery-tree. Active management of the NIC properties is not available, thereby preventing OneCommand Manager running on the guest operating system from performing operations that could adversely affect the host operating system or other guest operating systems; such as firmware download, diagnostics, or DCB changes from the guest operating system. By preventing remote access to the host operating systems, the guest operating system cannot affect the operation of the adapter.

## Viewing NIC PCI Registers

The NIC PCI Registers tab (Figure 8-6 on page 124) displays base PCI registers. See “Viewing the PCI Registers” on page 237 for FC PCI register information.

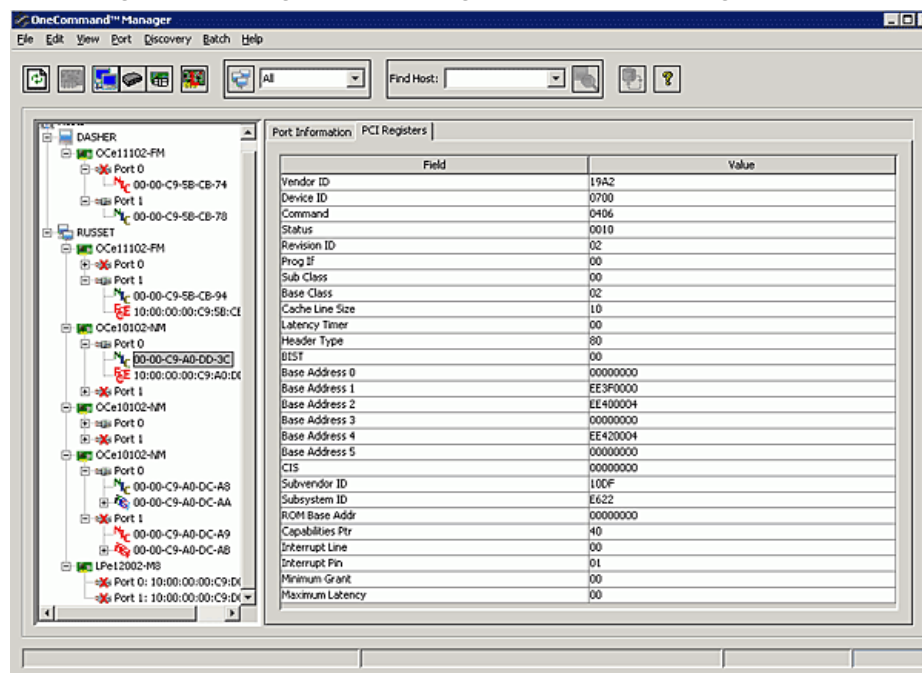


Figure 8-6 NIC PCI Registers Tab

To view NIC PCI registers:

1. From the discovery-tree, select the NIC function whose PCI information you want to view.
2. Select the **NIC PCI Registers** tab (Figure 8-6 on page 124).

## FC/FCoE

This section describes managing FC/FCoE ports.

### Viewing FCoE Function Information

When you select an FCoE function from the discovery-tree, the Port Information tab (Figure 8-7 on page 125) contains general attributes associated with the selected FCoE port.

To view FCoE Port information:

1. Select **Host** or **Fabric** view.
2. Select an FCoE function in the discovery-tree.

3. Select the **Port Information** tab (Figure 8-7 on page 125).

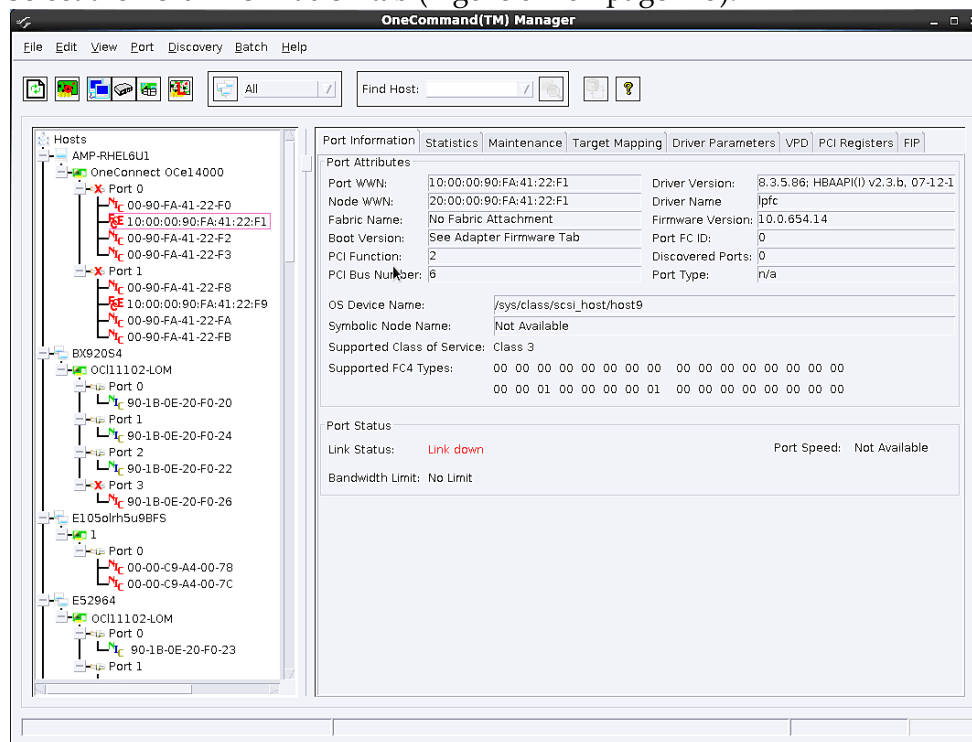


Figure 8-7 FCoE Port Information Tab

## FCoE Port Information Field Definitions

### Port Attributes Area Field Definitions

- Port WWN – The WWPNN of the FCoE function.
- Node WWN – The WWNN of the FCoE function.
- Fabric Name or Host Name – The Fabric Name field is displayed in Host view. This is a 64-bit worldwide unique identifier assigned to the fabric. The Host Name is displayed in Fabric view. The host name is the name of the host containing the FCoE function.
- Boot Version – The version of boot code installed on the selected FCoE function. If the boot code is disabled, the field displays “Disabled”.
- Port FC ID – The FCoE ID for the selected FCoE function.
- PCI Function – The PCI function number assigned by the system to the FCoE function.
- PCI Bus number – The PCI BUS number assigned by the system to the FCoE function.
- Driver Version – The version of the driver installed for the FCoE function.
- 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 FCoE function.

- Discovered Ports – The number of mapped and unmapped FC/FCoE ports found during discovery by the Emulex adapter driver. The mapped ports are targets and the unmapped ports are non-targets such as switches or adapters.
- Port Type – The current operational mode of the selected adapter's port.
- OS Device Name – The platform-specific name by which the selected FCoE function 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 non-delivery.
  - Class 2 – provides a frame switched service with confirmed delivery or notification of non-delivery.
  - Class 3 – provides a frame switched service similar to Class 2 but without notification of frame delivery or non-delivery.
- Supported FC4 Types – A 256-bit (8-word) map of the FC-4 protocol types supported by the FCoE function containing the selected adapter.

### Port Status Area Field Definitions

- Link Status – The status of the link on the selected FCoE function.
- Port Speed – The current speed of the selected FCoE function.
- Bandwidth Limit – The QoS bandwidth restriction on the FCoE function.

## Viewing FC/FCoE Statistics

When you select an FC port or FCoE function from the discovery-tree, the Statistics tab (Figure 8-8 on page 127) provides cumulative totals for various error events and statistics on the port. Some statistics are cleared when the adapter is reset.

To view FC/FCoE statistics:

1. Select **Host** or **Fabric** view.
2. Select an FC/FCoE function in the discovery-tree.

3. Select the **Statistics** tab (Figure 8-8 on page 127).

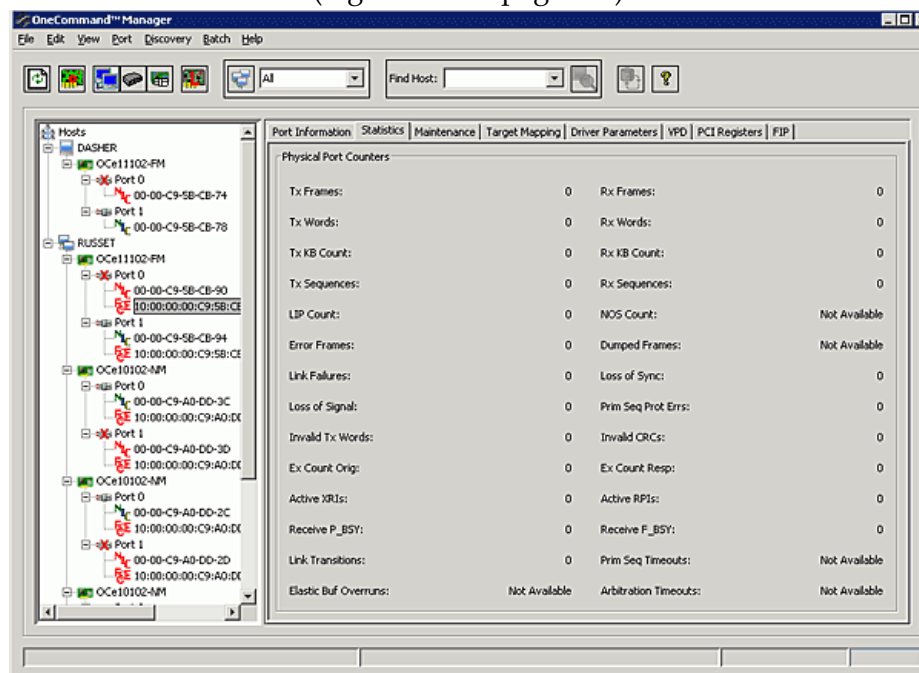


Figure 8-8 Statistics Tab

## Port Statistics Field Definitions

- Tx Frames – FC frames transmitted by this FC/FCoE function.
- Tx Words – FC words transmitted by this FC/FCoE function.
- Tx KB Count – FC kilobytes transmitted by this FC/FCoE function.
- Tx Sequences – FC sequences transmitted by this FC/FCoE function.
- LIP count – The number of loop initialization primitive (LIP) events that have occurred for the FC/FCoE 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/FCoE function.




- Ex Count Orig – The number of FC exchanges originating on this FC/FCoE function. (Not supported on VMware ESXi servers being managed through the CIM interface.)
- Active XRIs – The number of active exchange resource indicators. (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/FCoE function.
- Rx Words – The number of FC words received by this FC/FCoE function.
- Rx KB Count – The received kilobyte count by this FC/FCoE function.
- Rx Sequences – The number of FC sequences received by this FC/FCoE 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. (Not currently supported for Emulex Windows drivers or arbitrated loop.)
- Dumped Frames – The number of frames that were lost due to a lack of host buffers available. (Not currently supported for the SCSIport Miniport driver, 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/FCoE 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 port-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.)



## Viewing FC/FCoE Virtual Port Information

The Discovery Information tab (Figure 8-9 on page 129) can display information about FC/FCoE virtual ports and their associated targets and LUNs.

To view virtual port information:

1. Do one of the following:
  - From the **View** menu, select **Group Adapters by Virtual Port**.
  - From the toolbar, click  **Group Adapters by Virtual Port**.

The Discovery Information tab appears displaying virtual port information (Figure 8-9 on page 129).

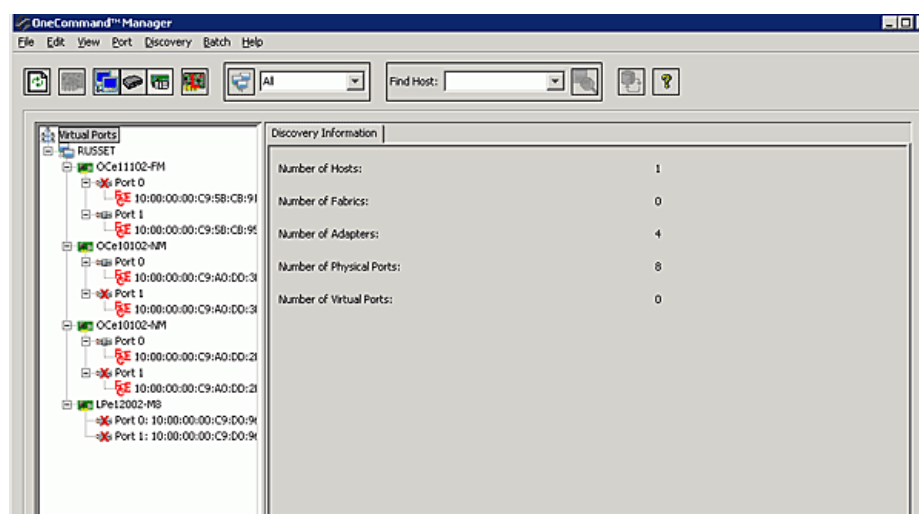


Figure 8-9 Discovery Information Tab Displaying Virtual Ports Information

### Discovery Information for Virtual Ports Field Definitions

- Number of Hosts – The total number of hosts discovered in the SAN.
- Number of Fabrics – The total number of fabrics discovered in the SAN.
- Number of Adapters – The total number of adapters discovered in the SAN.
- Number of Physical Ports – The total number of physical ports discovered in the SAN.
- Number of Virtual Ports – The total number of virtual FC/FCoE ports discovered in the SAN.

## Creating and Deleting FC/FCoE Virtual Ports

This section describes how to create and delete virtual ports.

### Creating Virtual Ports

Using the Virtual Ports tab (Figure 8-10 on page 130), 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. You cannot generate virtual ports on 1Gb Fibre Channel and 2GFC adapters.

## Notes


- 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 attempting 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, 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 8.2: lpfc\_enable\_npiv

See “Configuring FC/FCoE Driver Parameters” on page 155 for more information on enabling driver parameters.

To create a virtual port:

1. Do one of the following:
  - 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/FCoE function on which you want to create a virtual port. The Virtual Ports tab appears (Figure 8-10 on page 130).

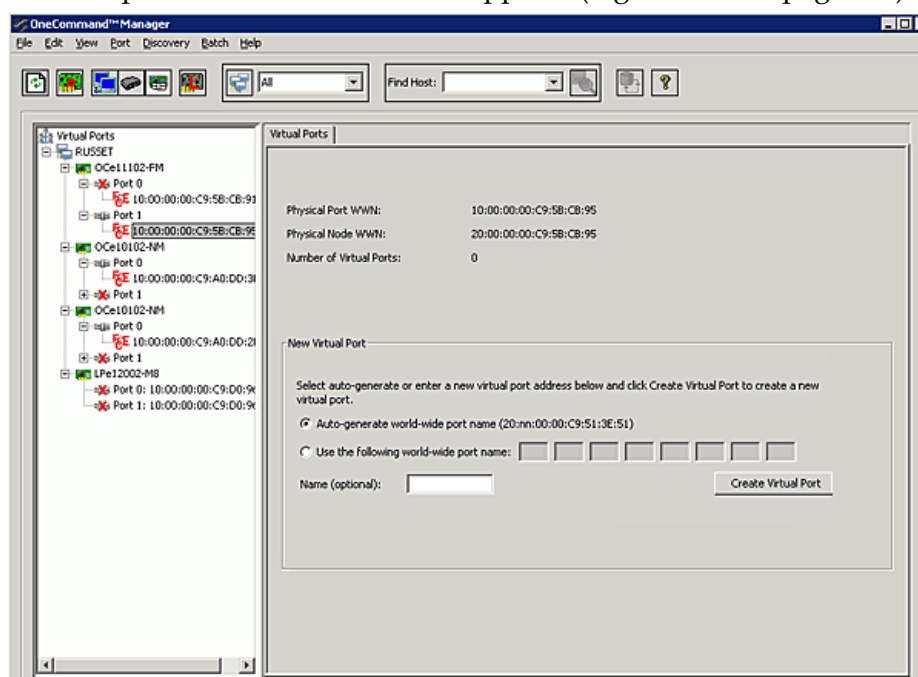


Figure 8-10 Virtual Ports Tab

3. Do one of the following:

- Check **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/FCoE 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.

- Check **Use the following world-wide port name** and enter a unique WWPN you want to use. 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

**Note:** Ensure that a manually entered WWPN is unique to your particular SAN. Failure to do so could result in a non-functioning SAN and data loss.

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 dialog box appears notifying you that the virtual port was created. The dialog box 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 4-3 on page 40) 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 may 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 4-3 on page 40).

## 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:

1. Do one of the following:
  - 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 virtual port you want to delete. The Virtual Ports tab appears (Figure 8-11 on page 132).

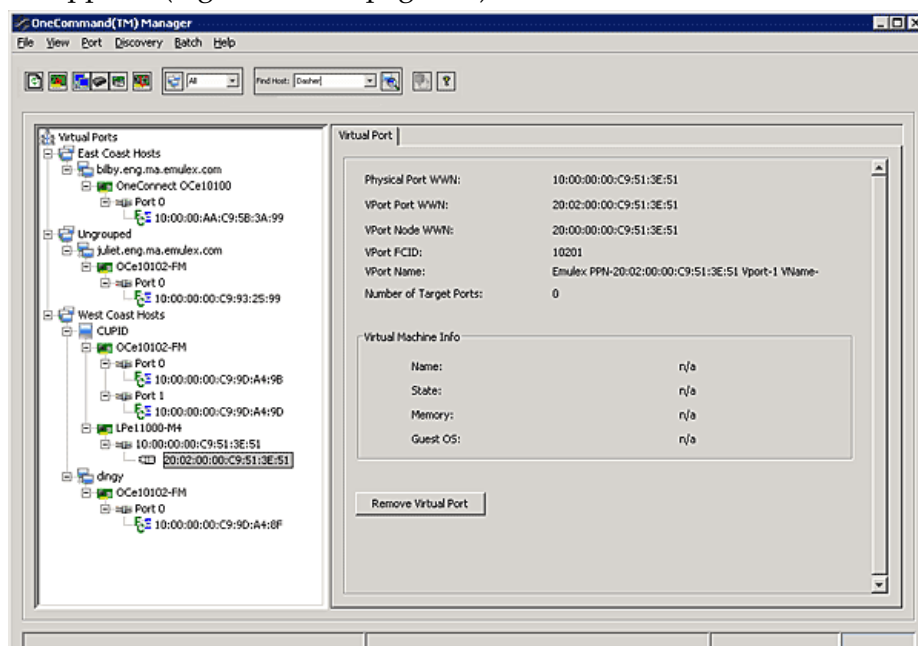
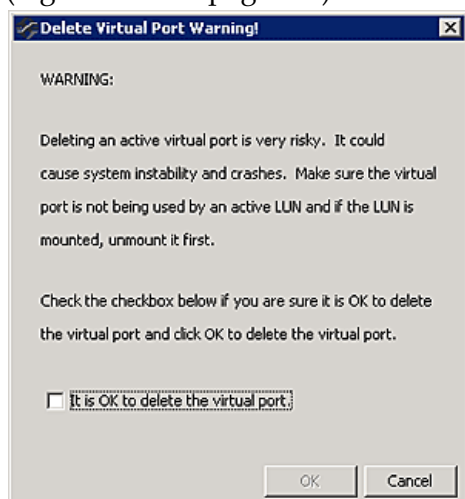


Figure 8-11 Virtual Port Tab

3. Click **Remove Virtual Port**. The Delete Virtual Port Warning dialog box appears (Figure 8-12 on page 133).



**Figure 8-12** Delete Virtual Port Warning

**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. Check **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 4-3 on page 40).
5. Click **OK**.

## Viewing FC/FCoE Fabric Information

The Discovery Information tab (Figure 8-13 on page 134) contains information about the selected fabric.

To view fabric discovery information, do one of the following:

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

The Discovery Information tab shows information about the fabric (Figure 8-13 on page 134).

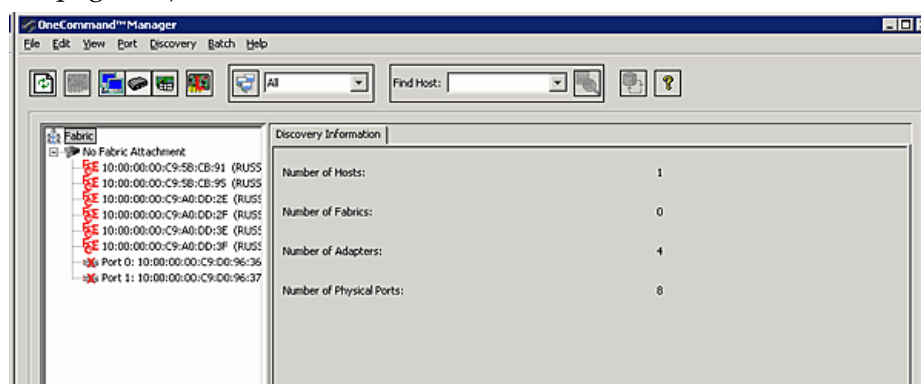


Figure 8-13 Fabric Discovery Information

## Discovery Information Field Definitions

- **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.

## Viewing FC Port Transceiver Information

**Note:** To view transceiver data for a OneConnect adapter, see “Viewing OneConnect Transceiver Information” on page 88.

When you select an FC port from the discovery-tree (Figure 4-3 on page 40), the Transceiver Data tab (Figure 8-14 on page 135) enables you to view transceiver information such as vendor name, serial number, part number and so on. If the adapter/transceiver does not support some or all of the transceiver data, the fields display N/A.

To view FC transceiver information:

1. Select **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4-3 on page 40), select the FC port whose transceiver information you want to view.

3. Select the **Transceiver Data** tab (Figure 8-14 on page 135).

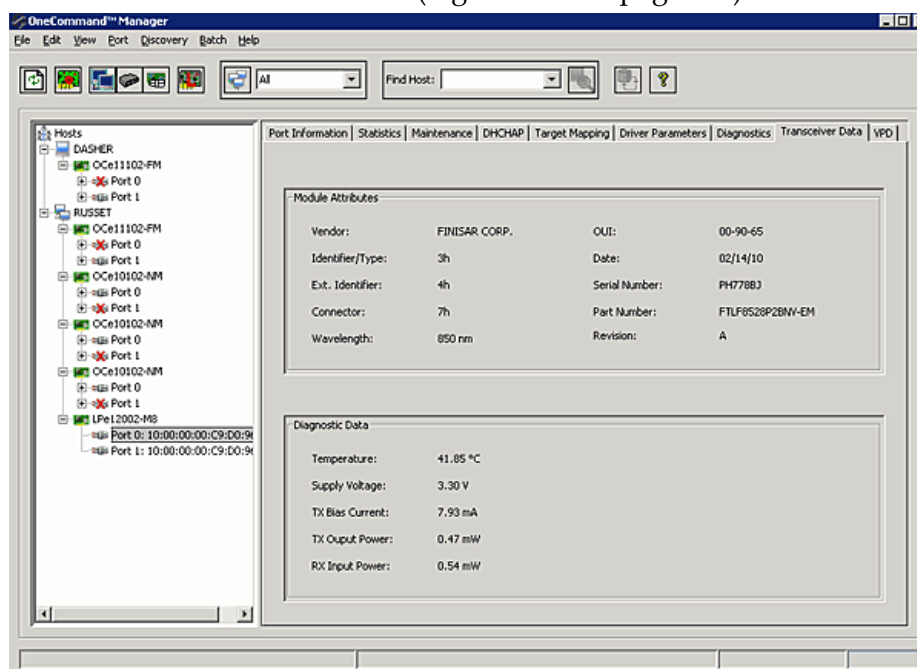


Figure 8-14 FC Transceiver Data Tab

## Transceiver Data Field Definitions

### 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.
- OUI – The vendor's OUI. It is also known as the IEEE Company Identifier for the vendor.
- Date – The vendor's date code in the 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.

## Viewing FC/FCoE VPD Information

The VPD tab (Figure 8-15 on page 136) displays vital product data (if available) for the selected FC adapter port such as the product name, part number, serial number, and so on.

To view VPD information:

1. Select **Host** or **Fabric** view.
2. In the discovery-tree, select the FC/FCoE function whose VPD information you want to view.
3. Select the **VPD** tab (Figure 8-15 on page 136).

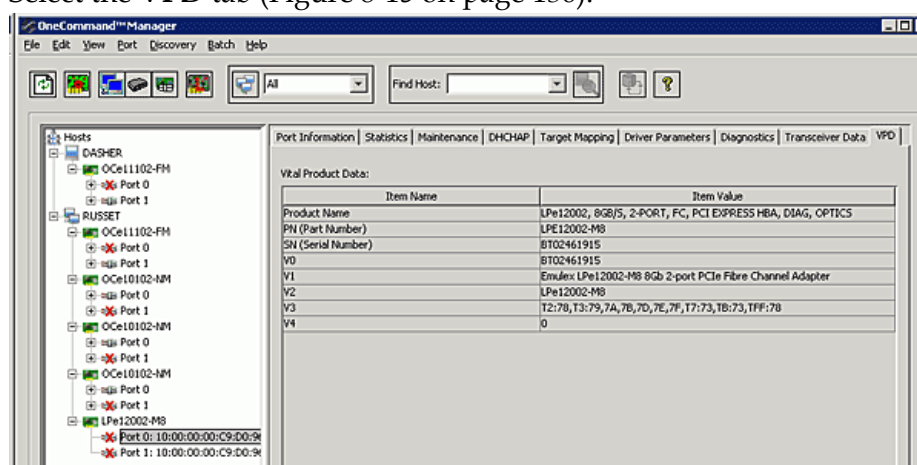


Figure 8-15 FC/FCoE VPD Tab

## VPD Table Definitions

- Product Name – Product information about the selected FC/FCoE function.
- PN (Part Number) – The adapter's part number.
- SN (Serial Number) – The adapter's serial number.
- VO – Vendor unique data. “V” indicates a vendor-specific field. An adapter may have none, one, or more of these fields defined. Valid values for this field are “VO” (the letter “O”, not the number zero) and “Vx” (where “x” is a number).

**Note:** Some adapters may 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.



## Viewing FC Maintenance Information

Use the Maintenance tab (Figure 8-16 on page 137) 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. (Not available in read-only mode.)

To view FC firmware information:

1. Select **Host** or **Fabric** view.
2. Select an FC port in the discovery-tree.
3. Select the **Maintenance** tab (Figure 8-16 on page 137).

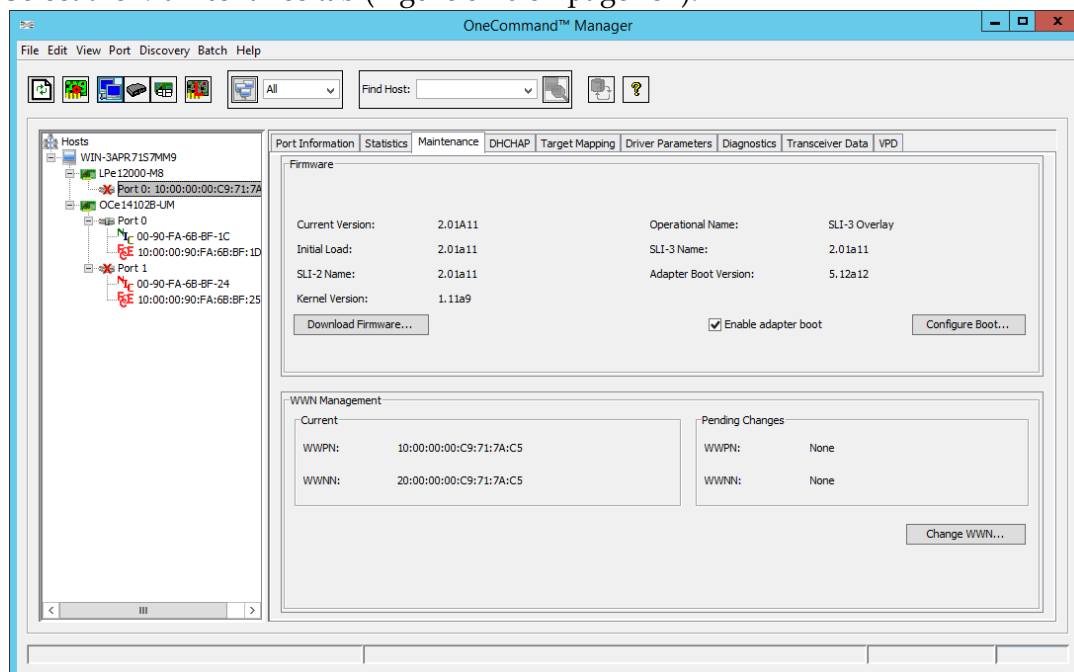


Figure 8-16 FC Maintenance Tab

## Maintenance Tab Field Definitions

### Firmware Area

- Current Version – The Emulex firmware version number for this FC port.
- Initial Load – The firmware version stub responsible for installing SLI code into its proper slot.
- SLI-2 Name – The name of the SLI-2 firmware overlay.
- Kernel Version – The version of the firmware responsible for starting the driver.
- Operational Name – The name of the operational firmware for the selected FC port.
- SLI-3 Name – The name of the SLI-3 firmware overlay.
- Adapter Boot Version – Displays one of the following:
  - The selected adapter port's boot code version if boot code is present.

- “Disabled” if the boot code is disabled.
- “Not Present” if boot code is not loaded. If boot code is not loaded, the Enable Adapter boot checkbox is not visible and you cannot configure the selected port to boot from SAN.
- Enable adapter boot checkbox – Check this box if you want the FC 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 FC port to load the boot code and execute it during system startup. It does not mean that the FC port will boot from SAN. To boot from SAN, the boot type must be enabled. Do this in the Boot from SAN configuration window for each boot type.

## WWN Management Area

### Current

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

### Pending Changes

- WWPN – Works in conjunction with the Change WWN button. Displays the WWPN you assigned for the selected FC port, but the system must be rebooted for these changes to take effect and appear under the “Current” listing. See “Changing FC/FCoE WWPN and WWNN” on page 151 for more information.
- WWNN – Works in conjunction with the Change WWN button. Displays the WWNN you assigned for the selected FC port, but the system must be rebooted for these changes to take effect and appear under the “Current” listing. See “Changing FC/FCoE WWPN and WWNN” on page 151 for more information.

## Maintenance Tab Buttons (Not available in read-only mode.)

- Download Firmware – Click to update firmware on the selected port. See “Updating Adapter Firmware” on page 202 for more information.
- Configure Boot – Check **Enable adapter boot** and click **Configure Boot** to configure boot from SAN. See “Exporting SAN Information” on page 234 for more information. (Not available on VMware ESXi servers being managed through the CIM interface.)
- Change WWN – Click to change the selected adapter port's WWNN or WWPN.

## Viewing FCoE Maintenance Information

Use the Maintenance tab (Figure 8-17 on page 139) to view firmware information. You can also configure boot from SAN and change WWPN and WWNN information for the selected FCoE function. (Not available in read-only mode.)

To view FCoE firmware information:

1. Select **Host** or **Fabric** view.
2. Select an FCoE function in the discovery-tree.

3. Select the **Maintenance** tab (Figure 8-17 on page 139).

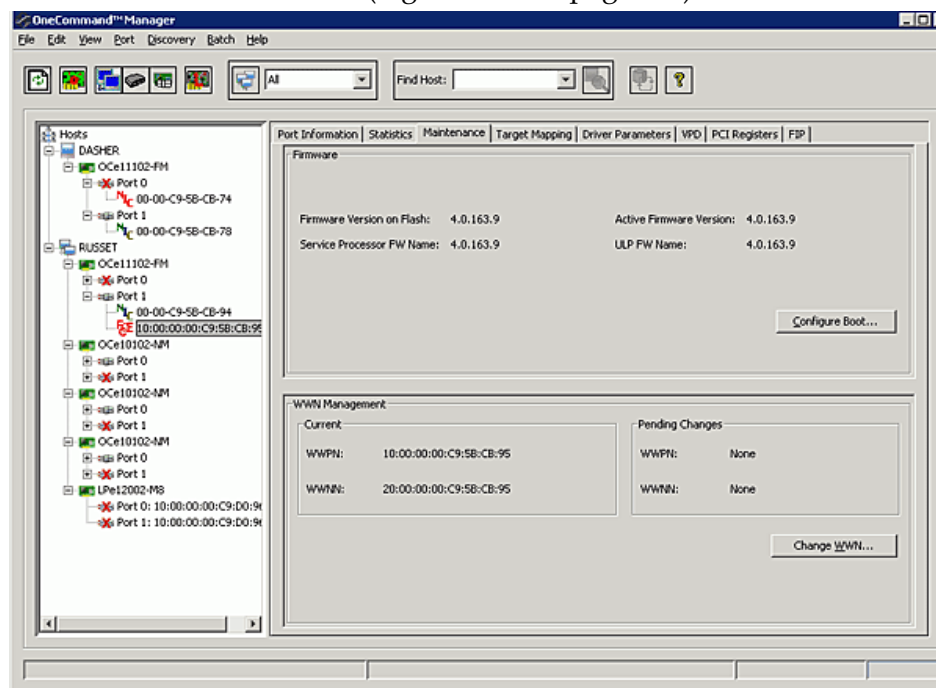


Figure 8-17 FCoE Maintenance Tab

## Maintenance Tab Field Definitions

### Firmware Area

- **Firmware Version on Flash** – The firmware version stored on the adapter’s non-volatile storage. When the system restarts, this version becomes the active firmware version.
- **Service Processor FW Version** – The firmware version that is currently operational on the adapter.
- **Active Firmware Version** – The version of firmware running on the adapter.
- **ULP FW Name** – The firmware version running on the ULP processors within the ASIC.

### WWN Management Area

#### Current

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

#### Pending Changes

- **WWPN** – Works in conjunction with the Change WWN button. Displays the WWPN you assigned for the selected FCoE function, but the system must be rebooted for these changes to take effect and appear under the “Current” listing. See “Changing FC/FCoE WWPN and WWNN” on page 151 for more information.

- **WWNN** – Works in conjunction with the Change WWN button. Displays the WWNN you assigned for the selected FCoE function, but the system must be rebooted for these changes to take effect and appear under the “Current” listing. See “Changing FC/FCoE WWPN and WWNN” on page 151 for more information.

## Maintenance Tab Buttons

**Note:** Not available in read-only mode.

- **Configure Boot** – Click **Configure Boot** to configure boot from SAN. (Not available on VMware ESXi servers being managed through the CIM interface.)
- **Change WWN** – Click to change the selected FCoE function’s WWNN or WWPN.

## Viewing FC/FCoE Target Information

When you select a target associated with an FC/FCoE adapter from the discovery-tree (Figure 4-3 on page 40), the Target Information tab (Figure 8-18 on page 140) displays information associated with that target.

To view FC/FCoE target information:

1. Select **Host**, **Fabric** or **Virtual Port** view.
2. In the discovery-tree (Figure 4-3 on page 40), select the FC/FCoE target whose information you want to view. The Target Information tab appears (Figure 8-18 on page 140).

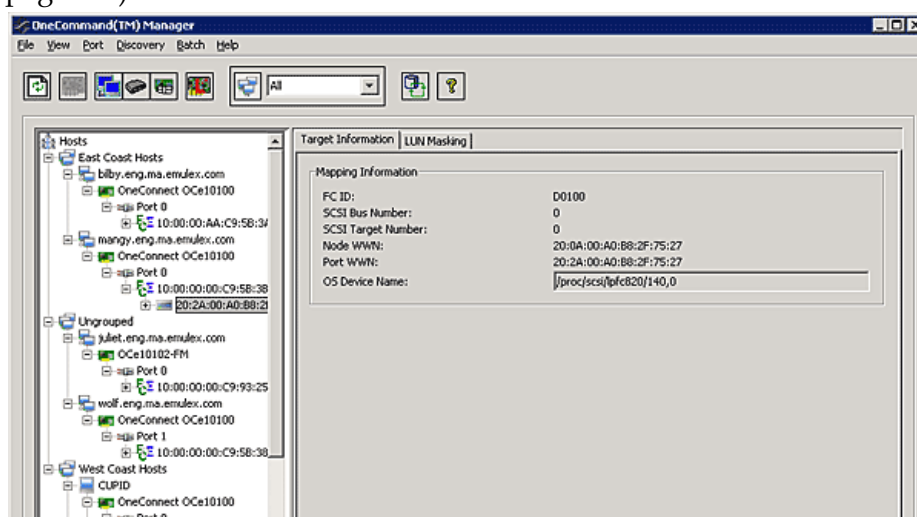


Figure 8-18 Target Information Tab

## Target Information Field Definitions

### 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 Switched Fabric Loop Port [FL\_PORT]).
- OS Device Name – The operating system device name.

## Viewing FC/FCoE LUN Information

When you select a LUN associated with an FC/FCoE target from the discovery-tree (Figure 8-19 on page 141), the LUN tab displays information associated with that LUN.

### Notes

- The Refresh LUNs button only refreshes the LUN list for the currently selected target.
- 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 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:

1. Select **Host**, **Fabric** or **Virtual Port** view.
2. From the discovery-tree, select a LUN.
3. The LUN Information tab appears (Figure 8-19 on page 141).

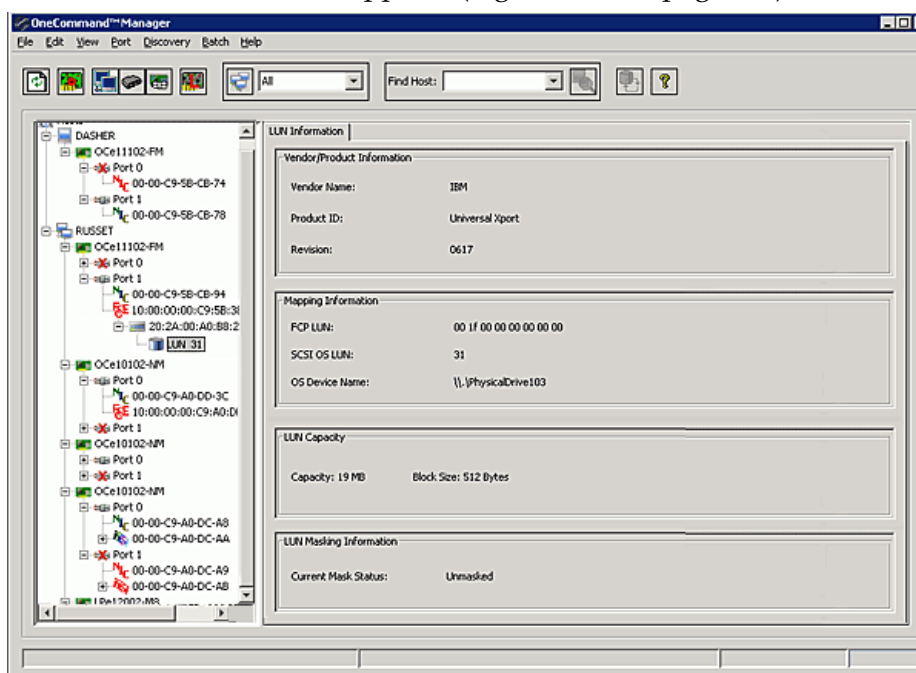


Figure 8-19 FC/FCoE LUN Information Tab

## FC/FCoE LUN Information Field Definitions

### 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 only provided 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 megabytes.
- Block Size – The length of a logical unit block in bytes.

### LUN Masking Area

- Current Mask Status – Possible states are masked or unmasked.

See “Masking and Unmasking LUNs (Windows)” on page 148 for more information on LUN Masking.

## Viewing FC/FCoE Target Mapping (Windows and Solaris)

The Target Mapping tab (Figure 8-20 on page 143) 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:

1. Select **Host** or **Fabric** view.
2. In the discovery-tree, select the FC/FCoE function whose target mapping information you want to view.

3. Select the **Target Mapping** tab (Figure 8-20 on page 143).

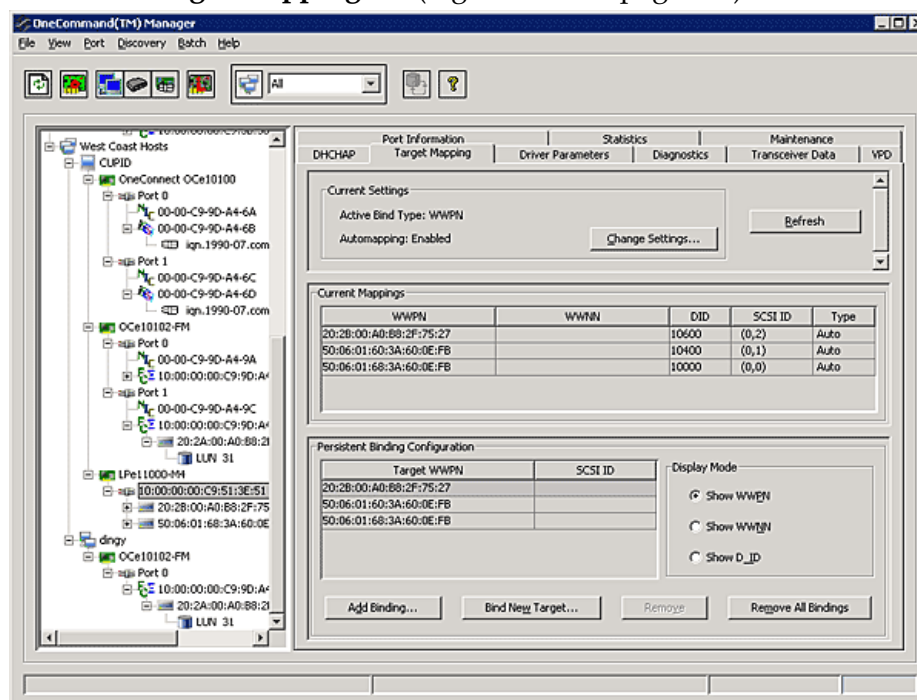


Figure 8-20 Target Mapping Tab

## Target Mapping Field Definitions

### Current Settings Area

- Active Bind Type – WWPN, WWNN, or a destination identifier (D\_ID).
- Automapping – The current state of SCSI device automapping: enabled (default) or disabled.

### Current Mappings Area

- This table lists current mapping information for the selected FC/FCoE function.

### Persistent Binding Configuration Area

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

### Display Mode Radio Buttons

- Show WWPN, Show WWNN or Show D\_ID options enable you to choose how to display information in the Persistent Binding Configuration table.

### Target Mapping Buttons

- Refresh – Click to refresh the Target Mapping tab.

- **Change Settings** – Click to enable or disable automapping, choose a bind type, and enable or disable LUN mapping and unmasking. (Not available on VMware ESXi servers being managed through the CIM interface.)
- **Add Binding** – Click to add a persistent binding.
- **Bind New Target** – Click to add a target that does not appear in the Persistent Binding table.
- **Remove** – Click to remove the selected binding.
- **Remove All Bindings** – Click to remove all persistent bindings that are displayed.

## Viewing FC/FCoE Target Mapping (Linux and VMware ESXi)

Use the Target Mapping tab to view target mapping. The Target Mapping tab is read-only.

**Note:** Persistent binding is not supported by the Linux 2.6 kernel, the Emulex 8.2 version of the driver for Linux, or by VMware ESXi Server.

To view target mapping:

1. Select **Host** or **Fabric** view.
2. Select the FC/FCoE function in the discovery-tree (Figure 4-3 on page 40) whose target mapping information you want to view.
3. Select the **Target Mapping** tab (Figure 8-20 on page 143).

## Target Mapping Field Definitions

### Current Settings Area

- Active Bind Type – N/A
- Automapping – N/A

### Current Mappings Area

- This table lists current mapping information for the selected FC/FCoE function.

### Persistent Binding Configuration Area

- N/A

### Display Mode Radio Buttons

- N/A

### Target Mapping Buttons

- N/A



## Using Automapping and Persistent Binding (Windows)

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. (Not available in read-only mode.)

The driver refers to the binding information at 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:

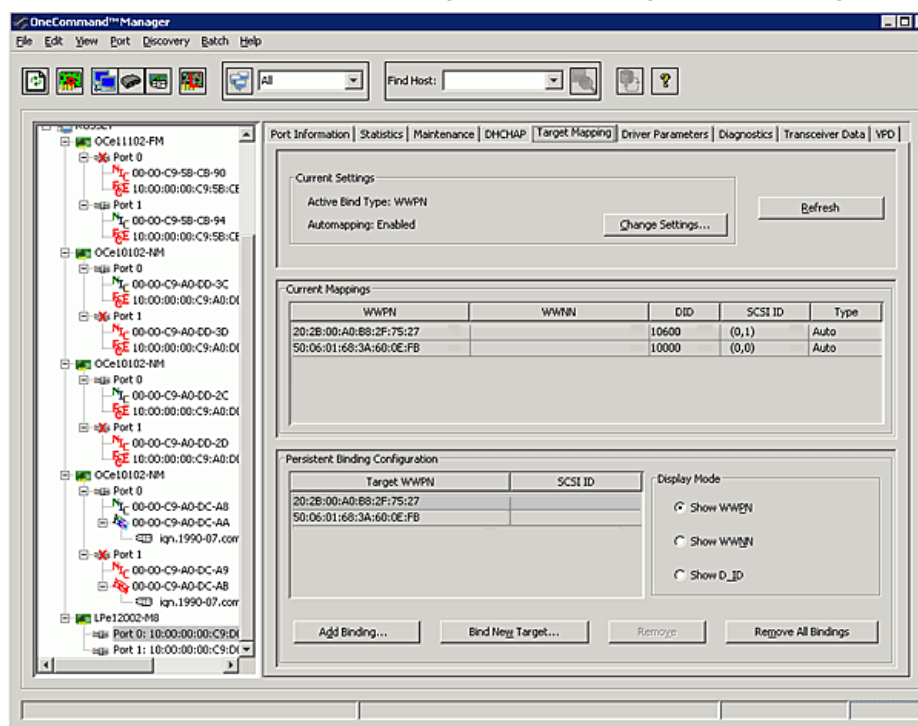
- The SCSI ID (target/bus combination) specified in the binding request must not be mapped to another target. For example, the SCSI ID must not already appear in the 'Current Mappings' table under 'SCSI ID'. If the SCSI ID is already in use, then the binding cannot be made dynamic, and a reboot is required.
- The target (WWPN, WWNN, or DID) specified in the binding request must not be mapped to a SCSI ID. If the desired target is already mapped, then a reboot is required.
- The bind type (WWPN, WWNN, or DID) 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, then the binding cannot be made active.

## Changing Automapping Settings

To change automapping settings:

1. Select **Host** or **Fabric** view.
2. In the discovery-tree, select the FC/FCoE function you want to set up with persistent binding.

3. Select the **Target Mapping** tab (Figure 8-21 on page 146). All targets are displayed.



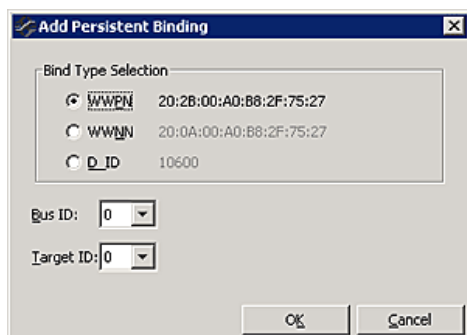
**Figure 8-21** Target Mapping Tab

4. Target mappings are displayed by WWP, WWNN, or D\_ID. “PB”, indicates mapping from persistent binding, while “Auto”, indicates an automapped target. In the Display Mode section, choose the display mode you want to use.
5. If you want click **Change Settings**. The Mapped Target Settings dialog box appears. You can enable or disable auto-mapping and change the active bind type. Click **OK**.
6. Reboot the system for changes to take effect.

## Adding a Persistent Binding

To add a persistent binding:

1. Select **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4-3 on page 40), select the FC/FCoE function you want to set up with persistent binding.
3. Select the **Target Mapping** tab (Figure 8-21 on page 146). All targets are displayed. In the Targets Table, click the target that you want to bind.
4. Click **Add Binding**. The Add Persistent Binding dialog box (Figure 8-22 on page 147) is displayed.



**Figure 8-22** Add Persistent Binding Dialog Box

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 Targets 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 you specified in the Add Persistent Binding dialog box. This binding takes effect only after the local machine is rebooted.

## 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 an error message, “Target already in target list. Use the Add Binding button.”

To bind a target that does not appear in the Persistent Binding table on the Target Mapping tab:

1. Select **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4-3 on page 40), select the FC/FCoE function you want to set up with persistent binding.
3. Select the **Target Mapping** tab (Figure 8-20 on page 143). All targets are displayed.
4. Click **Bind New Target**. The Bind New Target dialog box is displayed (Figure 8-23 on page 147).



**Figure 8-23** Bind New Target Dialog Box

5. Click the type of binding you want to use, and type the WWPN, WWNN, or D\_ID 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.

## Masking and Unmasking LUNs (Windows)

LUN masking refers to whether or not 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 8-24 on page 148) to mask or unmask LUNs at the host level.

**Note:** The LUN Masking tab (Figure 8-24 on page 148) is not shown in Virtual Port view because LUN masking is not available for virtual ports.

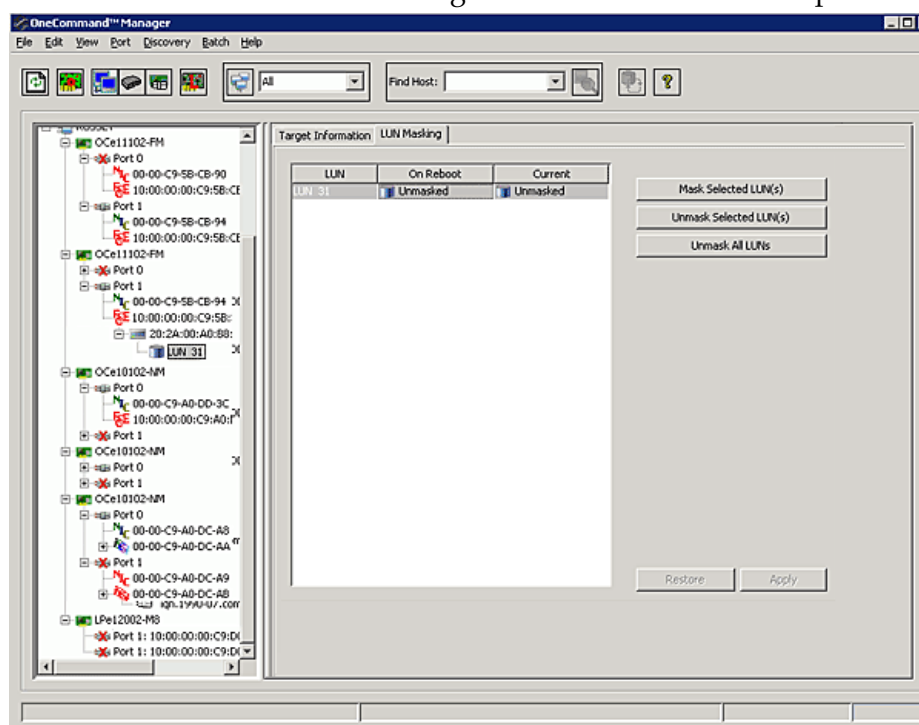


Figure 8-24 LUN Masking Tab

## LUN Masking Conventions and Guidelines

LUN icons in the discovery-tree (Figure 4-3 on page 40) reflect the live mask state currently in use 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).

## LUN Masking Column Definitions

- 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). Normally, for a specific LUN, the states reported in the 'On Reboot' and 'Current' column are identical. However, there can be times where 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 4-3 on page 40).

To change the mask status of a LUN:

1. Select **Host** view.
2. From the discovery-tree (Figure 4-3 on page 40), select the FC/FCoE target whose LUN masking state you want to change. A set of LUNs appears below the selected FC/FCoE target.
3. Select the **LUN Masking** tab (Figure 8-24 on page 148). This tab contains a list of the same set of LUNs that appear below the FC/FCoE target in the discovery-tree (Figure 4-3 on page 40).
4. In the LUN list of the LUN Masking tab, select one or more LUNs. The Mask Selected LUNs, Unmask Selected LUNs, Unmask All LUNs, Restore, and Apply buttons become active as appropriate. For example, if the LUN is currently unmasked, only the Mask Selected LUNs 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. Be sure to also click **Apply** to commit the changes.

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

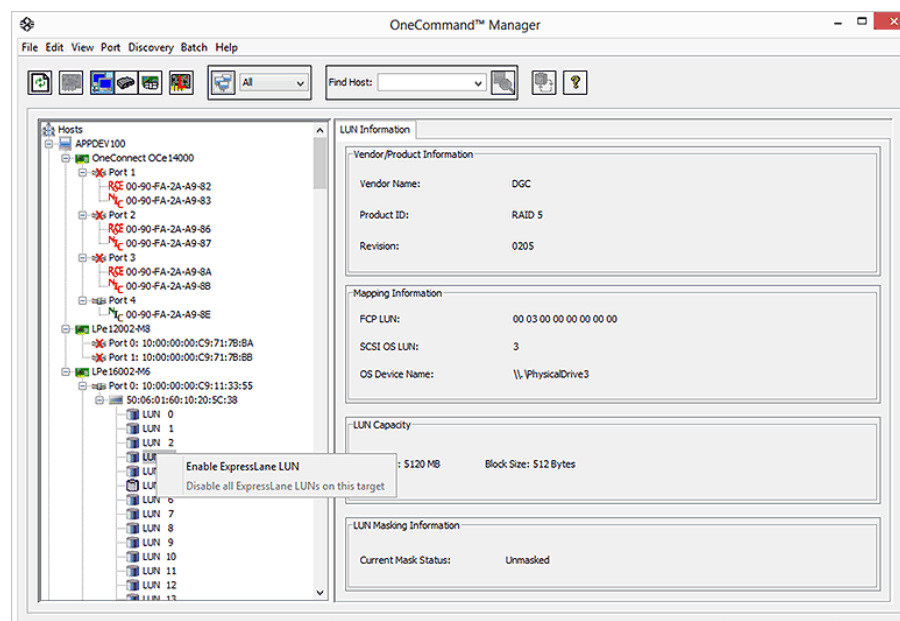
## Managing FC/FCoE ExpressLane LUNS (LPe15000 and LPe16000 HBAs)

The OneCommand Manager application allows you set special priority queuing for selected LUNs by making them ExpressLane LUNs (Figure 8-25 on page 150). ExpressLane LUN performance is superior to that of regular LUNs. You can enable ExpressLane LUNs attached to both physical and virtual ports.

ExpressLane LUN assignments persist across reboots.

## Notes

- Masked LUNs cannot be ExpressLane enabled because they are not presented to the host. Conversely, ExpressLane LUNs cannot be masked
- For Linux operating systems, if ExpressLane LUNs are created, the VPort must be recreated after a system boot because VPorts do not persist across system reboots. If the VPort is re-created with the same WWPN that the ExpressLane LUN was previously assigned to, and the same LUN is then discovered, it becomes an ExpressLane LUN again.



**Figure 8-25** Enabling an ExpressLane LUN

To enable an ExpressLane LUN:

1. Select **Host**, **Fabric** or **Virtual Port** view.
2. From the discovery-tree, select a LUN under an LPe15000 or LPe16000-series adapter on which you want to enable ExpressLane.
3. The LUN Information tab appears (Figure 8-25 on page 150).
4. Right-click the selected LUN in the discovery-tree and choose **Enable ExpressLane LUN**. The LUN's icon in the discovery-tree changes to the ExpressLane LUN icon.

To disable an ExpressLane LUN or LUNs:

1. Select **Host**, **Fabric**, or **Virtual Port** view.
2. From the discovery-tree, select a LUN under an LPe15000 or LPe16000-series adapter on which you want to enable ExpressLane.
3. The LUN Information tab appears (Figure 8-24 on page 148).
4. Right-click on the selected LUN in the discovery-tree and choose **Disable ExpressLane LUN** to disable the selected LUN or **Disable all ExpressLane LUNs on this target**. The ExpressLane LUN's icon in the discovery-tree changes to the regular LUN icon.

## Changing FC/FCoE WWPN and WWNN

The Maintenance tab (Figure 8-26 on page 152) enables you to change the WWPN and the WWNN of a selected FC/FCoE function. For example, you might want to 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.

Three options exist for referencing WWNs:

- **Factory Default WWN** – As shipped from the factory. This value cannot be changed.
- **Non-Volatile WWN** – Values that are saved in non-volatile adapter's flash memory that survives a reboot and/or 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 in order to take effect. Volatile WWN changes are lost on systems that power cycle the adapters during the reboot.

**Caution:** Changing volatile WWNs takes the selected adapter offline. Ensure that this adapter is not controlling a boot device and all I/O activity on this adapter is stopped before proceeding. Emulex assumes no responsibility for the consequences of making volatile WWN changes on a boot adapter.

### Notes

- To avoid address conflicts, do not assign a WWPN with the same WWPN as another FC/FCoE function on your SAN. The OneCommand Manager application checks the WWPN you specify against all the other discovered WWPNs and if a duplicate is found, an error is displayed and the WWPN is not be changed.
- 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 Lenovo System X 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 protocol and uses the default WWNN or WWPN per BOFM's command.
- ◆ The OneCommand Manager application modifies the WWNN or WWPN. The OneCommand Manager application requires a reboot to



complete the change. The system comes up and the OneCommand Manager application-modified WWNN or WWPN is used.

To change an FC/FCoE function's WWPN or WWNN:

1. Do one of the following:
  - From the **View** menu, click **Group Adapters by Host Name**.
  - From the toolbar, click  **Group Adapters by Host Name**.
  - From the **Host Grouping** menu, select **Group Adapter by Fabric Names**.
2. In the discovery-tree, select the FC/FCoE function whose information you want to change.
3. Select the **Maintenance** tab (Figure 8-26 on page 152).

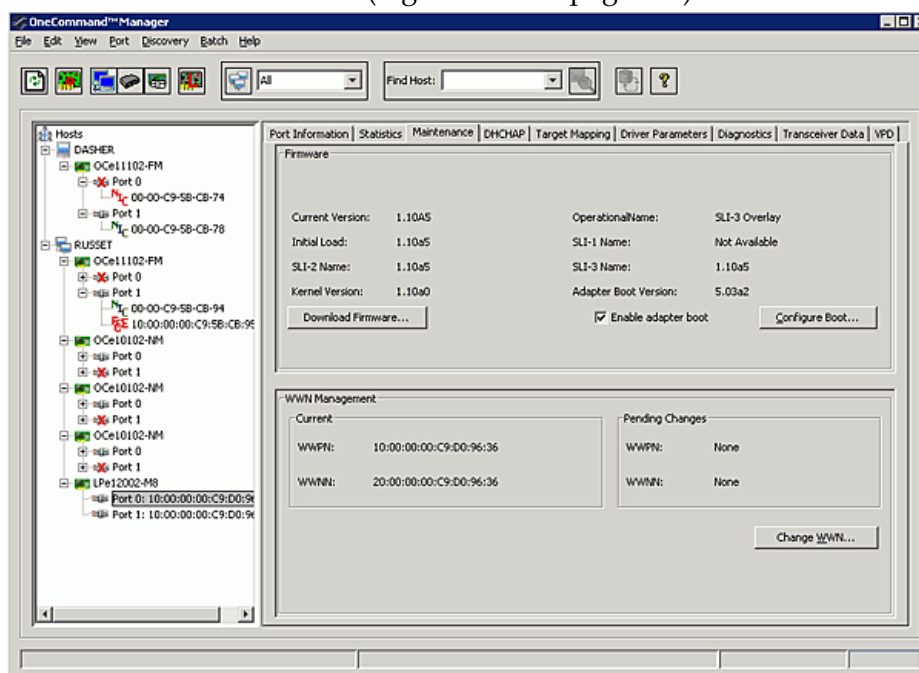


Figure 8-26 Maintenance Tab

4. Click **Change WWN**. The warning in Figure 8-27 on page 152 appears:


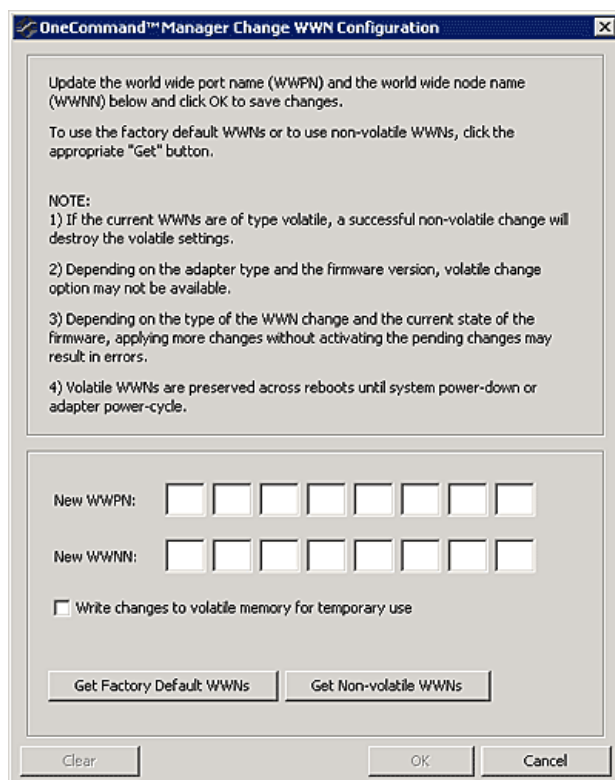


Figure 8-27 Warning About Changing WWN



5. Click **Yes**. The Change World Wide Name Configuration dialog box appears (Figure 8-28 on page 153).



**Figure 8-28** Change World Wide Name Configuration Dialog Box

6. Do one of the following:
    - Enter a new WWPN and WWNN.
    - Click **Get Factory Default WWNs** to load the settings that were assigned to the FC/FCoE function when the adapter was manufactured to the New WWPN and WWNN settings. 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 WWNN settings. These values can then be modified if desired and saved to volatile or non-volatile memory. You can edit the data returned from the button.
  7. Check **Write changes to volatile memory for temporary use** to save the New WWPN and New WWNN settings as Volatile WWNs. If unchecked, 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" checkbox is disabled.
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 section of the Maintenance dialog box.

## Changing FC/FCoE Names

The OneCommand Manager application allows you to change FC/FCoE adapter port names in the discovery-tree. (Not available in read-only mode.)

For example, you may want to identify a particular FC/FCoE 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/FCoE function's displayed name from the default WWPN, the change occurs in the discovery-tree (Figure 4-3 on page 40) only. The functions's WWPN is still active, it is simply 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 change 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/FCoE function:

1. From the discovery-tree (Figure 4-3 on page 40), select the FC port whose name you want to change.
2. Do one of the following:
  - Select **Edit Name** from the **Port** menu.
  - From the discovery-tree, right-click the port whose name you want to change and select **Change Name**.
3. Edit the in the discovery-tree.

To use the FC/FCoE function's default name:

1. From the discovery-tree (Figure 4-3 on page 40), select the FC FC/FCoE function whose name you want to change.
2. Do one of the following:
  - Select **Use Default Name** from the **Port** menu.
  - From the discovery-tree, right-click the port whose name you want to change and select **Restore Default Name**.


## Resetting FC/FCoE Functions

You can reset remote and local FC/FCoE functions. (Not available in read-only mode or on NIC or iSCSI functions.)

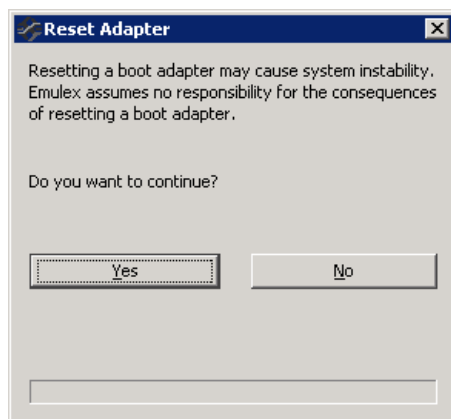
**Caution:** Do not reset your functions while copying or writing files. This could result in data loss or corruption.

**Note:** For FCoE functions, a reset is only necessary to activate updated driver parameters (that require a reset) or FIP settings. It does not actually perform an adapter function level reset.

To reset the FC/FCoE function:

1. In the discovery-tree (Figure 4-3 on page 40), select the FC/FCoE you want to reset.
2. Do one of the following:
  - From the **Port** menu, click **Reset Port**.
  - From the toolbar, click  **Reset**.

The Reset Adapter warning appears (Figure 8-29 on page 155).



**Figure 8-29** Reset Warning

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

The reset can require several seconds to complete. While resetting, the status bar shows “Reset in progress.” When the reset is finished, the status bar shows “Reset Completed”.

## Configuring FC/FCoE Driver Parameters

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 should be aware of when working with driver parameters. (Not available in read-only mode.) For a more detailed description of specific driver parameters, refer to the appropriate Emulex driver user manual.

### Notes

- This section only applies to FC and FCoE drivers. It does not apply to NIC, RoCE, or iSCSI drivers.
- In Solaris and Linux, you can also specify parameters when loading the driver manually. (Not available in read-only mode.) Refer to the appropriate driver manual for instructions.

## Activation Requirements

A parameter has one of the following activation requirements:

- Dynamic – The change takes effect while the system is running.
- Reset – Requires an FC/FCoE 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.

## The Host Driver Parameters Tab

The Host Driver Parameters tab (Figure 8-30 on page 156) 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 “Setting Driver Parameters” on page 157. For information on changing parameters for the host, see “Setting Driver Parameters for All FC/FCoE Functions in a Host” on page 160.

**Note:** If there are no discovered FC or FCoE functions, the entire Host Driver Parameters tab (Figure 8-30 on page 156) is grayed-out. This occurs because there are no drivers to which the host driver parameters apply.

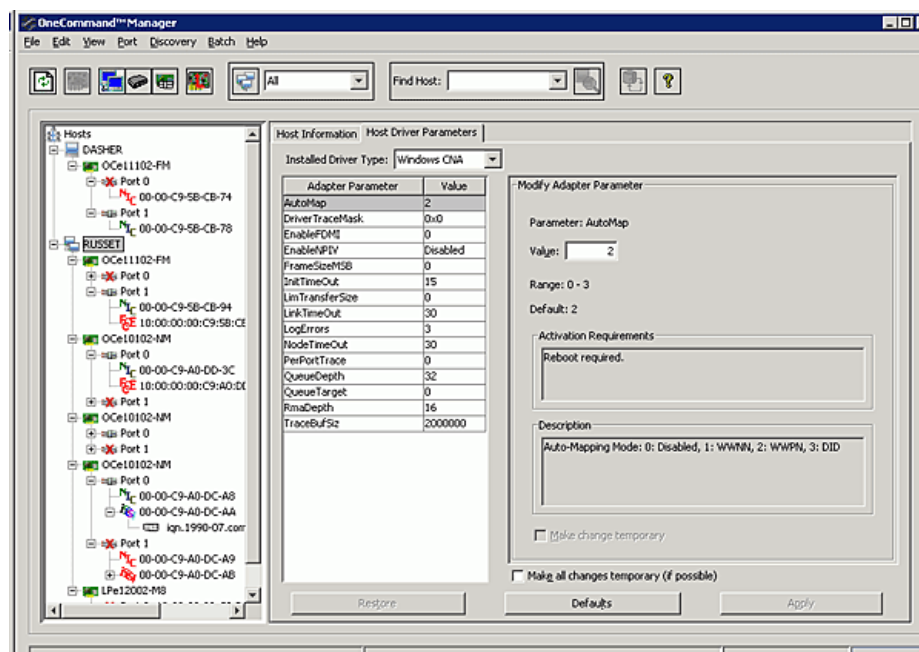


Figure 8-30 Host Driver Parameters Tab

## Host Driver Parameters Tab Field Definitions

- **Installed Driver Type** – The current drivers installed on this host. If there is more than one driver type installed, the Installed Driver Types drop-down menu 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. This can include value, range, default, activation requirements, and description.

### Driver Parameters Tab Buttons (Not available in read-only mode.)

- **Restore** – If you changed driver parameters, but did not click **Apply** and you want to restore the parameters to their last saved values, click **Restore**.
- **Defaults** – Click to reset all driver parameter values to their default (out-of-box) values.
- **Apply** – Click to apply any driver parameter changes. If you changed a driver parameter that is not dynamic, you may need to reset the FC/FCoE function or reboot the system.

## Setting Driver Parameters

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

For example, if you select a host in the discovery-tree (Figure 4-3 on page 40), you can globally change the parameters for all FC/FCoE functions in that host. If you select an FC/FCoE function in the discovery-tree, you can change parameters for only that FC/FCoE function.

For each parameter, the Driver Parameters tabs show 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/FCoE function to other FC/FCoE functions in the system using the Driver Parameters tab, thereby simplifying multiple adapter configuration. See “Creating a Batch Mode Driver Parameters File” on page 161 for more information.

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

## Setting Driver Parameters for a Single FC/FCoE Function

To change the driver parameters for a single FC/FCoE function:

1. Select **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4-3 on page 40), select the FC or FCoE function whose parameters you want to change.
3. Select the **Driver Parameters** tab (Figure 8-31 on page 158). The parameter values for the selected FC/FCoE functions are displayed.

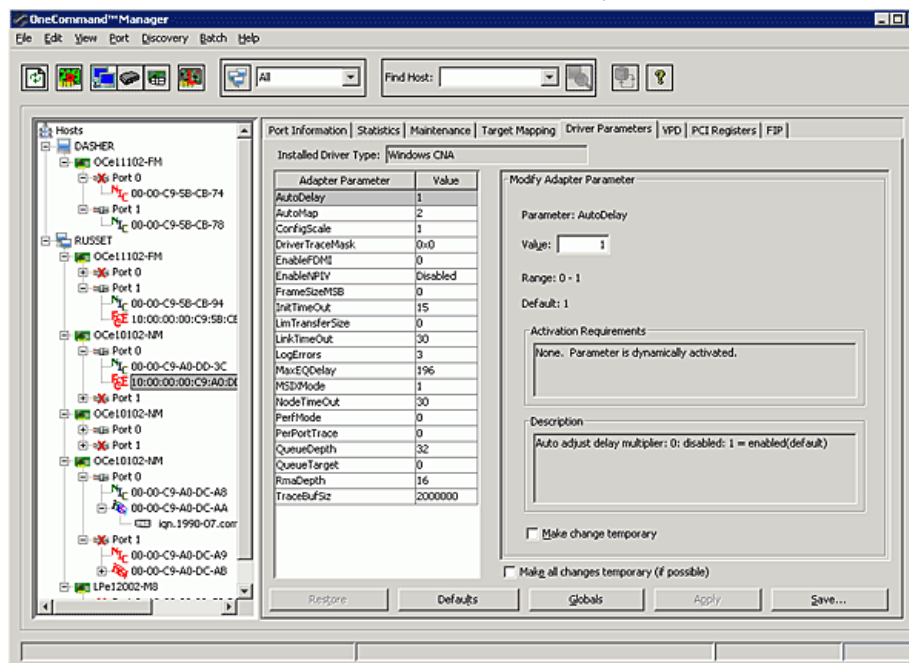


Figure 8-31 Driver Parameters Tab - Adapter Selected

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 drop-down menu. If you enter a value and 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), check the **Make change temporary** 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, check the **Make all changes temporary** box. This setting overrides the setting of the **Make change temporary** box. Only dynamic parameters can be made temporary.
8. Click **Apply**.

## Restoring All Parameters to Their Earlier Values

If you changed parameters, but did not click **Apply** in the **Driver Parameters** tab (Figure 8-31 on page 158) and you want to restore the parameters to their last saved values, click **Restore**.

## Resetting All Default Values

To reset all parameter values to their default (factory) values, click **Defaults** in the **Driver Parameters** tab (Figure 8-31 on page 158).

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

To set an adapter parameter value to the corresponding host parameter value:

1. Select **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4-3 on page 40), select the adapter port.
3. Select the **Driver Parameters** tab (Figure 8-31 on page 158).
4. Click **Globals**. All parameter values are now the same as the global, or host, values.
5. To apply the global values, click **Apply**.

## Saving Adapter Driver Parameters to a File

To save a desired adapter parameter configuration for using with the Batch Driver Parameter Update feature, click **Save** in the **Driver Parameters** tab (Figure 8-31 on page 158). To apply your configuration changes, click **Apply**.

**Note:** OneCommand Manager application Web Launch Interface driver parameters files are saved on the host that the browser was launched from, not the host IP specified in browser.

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


The OneCommand Manager application can then use the Batch Driver Parameter Update function to apply these saved settings to any or all compatible adapters on the SAN.

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



## Setting Driver Parameters for All FC/FCoE Functions in a Host

To change the driver parameters for all FC/FCoE functions installed in a host:

1. Do one of the following:
  - From the **View** menu, click **Group Adapters by Host Name**.
  - From the toolbar, click  **Group Adapters by Host Name**.
2. In the discovery-tree, click the host whose adapter driver parameters you want to change.
3. Select the **Host Driver Parameters** tab (Figure 8-32 on page 160). If there are adapters with different driver types 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.

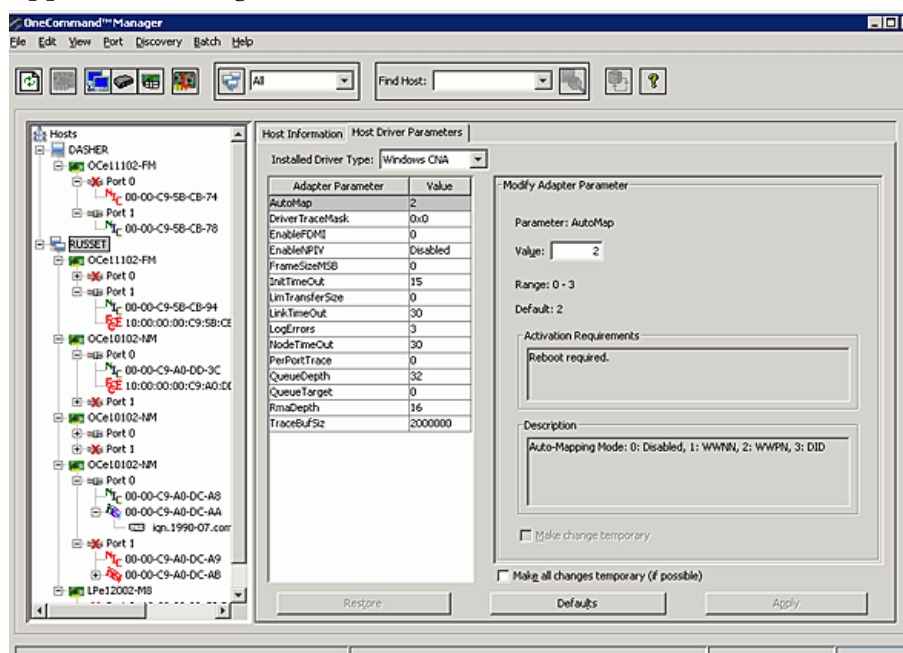


Figure 8-32 Host Driver Parameters Tab - Host Selected

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 to revert to its last permanent setting when the system is rebooted), check **Make changes temporary**. This option is available only for dynamic parameters.
7. To make changes to multiple parameters, check **Make all changes temporary**. Only dynamic parameters can be made temporary.
8. Click **Apply**.



## Changing Non-dynamic Parameter Values (Linux 8.2)

To change non-dynamic parameter values for Linux version 8.2:

1. Navigate to the /usr/sbin/ocmanager directory and run the scripts to stop the OneCommand Manager application processes. Type

```
./stop_ocmanager
```

2. Stop all I/O to FC, FCoE, or FC/FCoE attached devices.

3. Unload the FC, FCoE, or FC/FCoE driver. Type

```
modprobe -r lpfc
```

4. Reload the driver. Type:

```
modprobe lpfc
```

5. If DHCHAP authentication is currently employed on this machine, start up the Emulex FC authentication service. Type

```
/etc/init.d/fcauthd start
```

6. Start the elxhbmgr service (remote service). Type

```
./start_ocmanager
```

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

**Note:** If DHCHAP authentication is currently employed on Emulex adapters on this machine, you must type “/etc/init.d/fcauthd start” to restart the authentication daemon.

If the machine has the OneCommand Manager application Web Launch Interface installed, the RMI services must be restarted. Type

```
./start_weblaunch
```

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

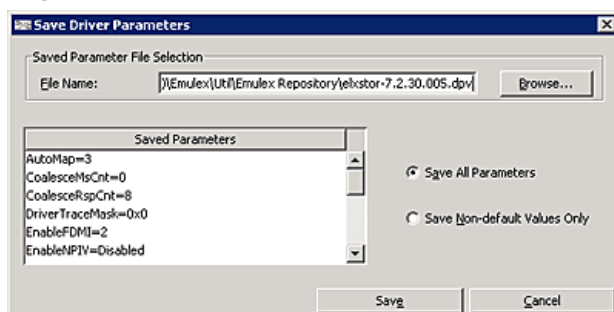
## Creating a Batch Mode Driver Parameters File

You can apply driver parameters for one FC/FCoE function to other FC/FCoE functions in the system using the Driver Parameters tab. When you save the driver parameters for an adapter, you create a .dpv file. The .dpv file 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. (Not available in read-only mode.)

To create the .dpv file:

1. Select **Host** or **Fabric** view.
2. Select the FC/FCoE function whose parameters you want to apply to other FC/FCoE functions from the discovery-tree (Figure 4-3 on page 40).
3. Select the **Driver Parameters** tab (Figure 8-31 on page 158).
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 box appears (Figure 8-33 on page 162). You can save the file to a different directory or change its name.



**Figure 8-33** Save Driver Parameters Dialog Box

7. Use the two radio buttons to choose the type of parameters to save. You can save all parameters or only those parameters whose current values differ from their corresponding default values.

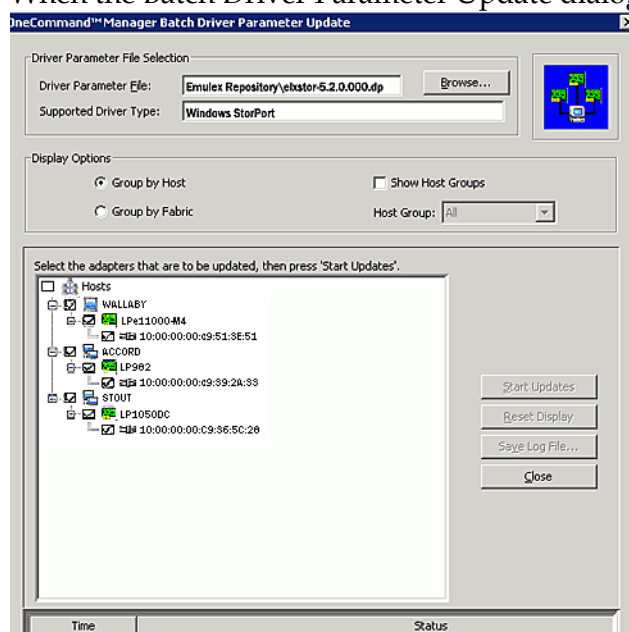
A list of the saved parameters and their current values show in the Saved Parameters box.

8. Click **Save**.

## Assigning Batch Mode Parameters

To assign batch mode parameters to adapters:

1. From the **Batch** menu, select **Update Driver Parameters**. (You do not need to select any discovery-tree (Figure 4-3 on page 40) elements at this time.)
2. When the Batch Driver Parameter Update dialog box appears, click **Browse**.



**Figure 8-34** Batch Driver Parameters Update Dialog Box

3. The Driver Parameter File Selection dialog box appears (Figure 8-34 on page 162). Select the file you want to use and click **OK**. A dialog box appears notifying you that the OneCommand Manager application is searching for compatible adapters.

After compatible FC/FCoE functions are found, the Driver Parameter File field of the Batch Driver Parameter Update dialog box displays the selected file's path. The "Supported Models" text field displays a list of all adapter models that are compatible with the selected file. The set of compatible adapters appears in the dialog box's discovery-tree.

Using the Display Options settings you can choose how adapters are displayed in the discovery-tree. Clicking **Group by Host** displays adapters in a host-centric view. Clicking **Group by Fabric** shows hosts in a fabric-centric view with their fabric addresses. The WWPN and host name for each available FC/FCoE function is displayed under its respective fabric.

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

Checkboxes next to the host, adapter, and FC/FCoE functions entries are used to select or clear an entry. Checking an adapter selects or removes all FC/FCoE 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 OneCommand Manager application Batch Driver Parameter Update dialog box (Figure 8-34 on page 162) shows the current status of the update. When the update completes, a final summary shows the number of FC/FCoE functions that were successfully processed, and the number of FC/FCoE functions for which one or more parameter updates failed.
5. You can click **Save Log File** to save a report of the update.

## Configuring FCoE Initialization Protocol (FIP)

The FIP tab enables you to configure FIP for FCoE functions.

To configure FIP:

1. From the discovery-tree, select the FCoE function whose FIP properties you want to configure.
2. Select the **FIP** tab (Figure 8-35 on page 164).

- Set the parameters you want and click **Apply Changes**.

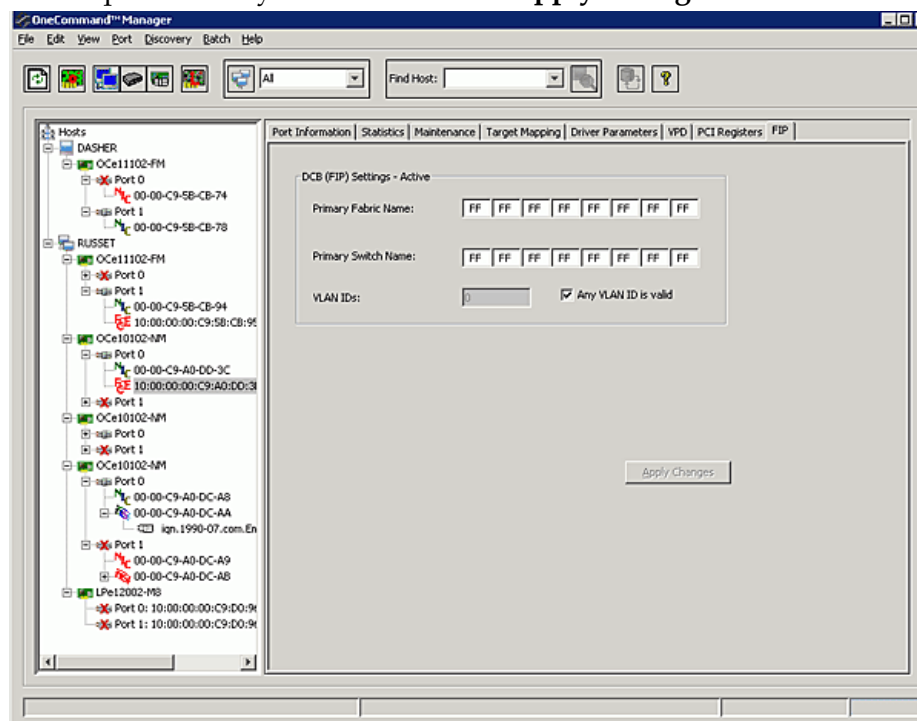


Figure 8-35 FIP Tab for FCoE Adapters

## FIP Dialog Box Field Descriptions

- Primary Fabric Name – Indicates the FC Fabric's WWN to which to connect. If the Primary Fabric Name is wild, that is, all 0xFFs, then connection to any fabric name is allowed.
- Primary Switch Name – Indicates the FC Switch's WWN to which to connect. If the Primary Switch Name is wild, that is, all 0xFFs, then connection to any switch name is allowed.
- VLAN ID – Determines the VLAN where the adapter FCoE services are available. It can have a value from 0-4095 and supports wild card values if "Any" is checked.
- Any VLAN ID is valid check box – If checked, the VLAN ID field of the FCoE forwarder can be any valid value.

## Configuring Boot from an FC/FCoE SAN

You can use the OneCommand Manager application to configure a system to boot from an attached FC/FCoE 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.

## 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).

- 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 and on Linux PowerPC systems. OpenBoot is also called FCode.
- EFIBoot works with x64-based systems and provides 64-bit system boot capability through the use of the EFI Shell.

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

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

## Notes

- 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 appears that the existing boot type configuration will be overwritten by the new configuration.
- Boot from SAN configuration does not affect current system operation. The changes only take effect upon reboot if you have configured it correctly.

## Boot Device Parameters

The boot LUN for all three boot types is in the range of 0-255. 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.
- The OneCommand Manager application runs on a running operating system, so 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. Without knowing this, you could

configure a boot type but not be able to boot from it since the adapter lacks the correct boot code.

- You must know what boot code type the system supports; the OneCommand Manager application cannot detect this. You can configure any boot type, but if the system does not support that type, it cannot boot from SAN.
- If you manage adapters on a remote host that is running a version of the OneCommand Manager application that does not support boot from SAN, the Configure Boot button does not appear.

**Note:** You can configure boot from SAN before boot by using the Emulex Boot BIOS setup command line interface that runs during system startup. See the Emulex Boot BIOS setup program documentation for details.

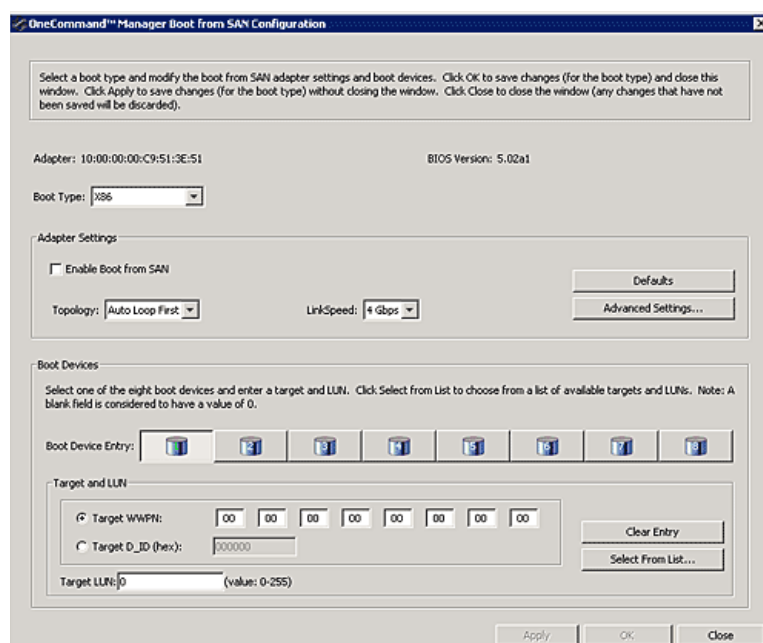
- One of the following FC or FCoE adapter drivers must be installed:
  - Storport Miniport or CNA driver for Windows
  - Emulex driver for Linux
  - Solaris emlxs FCA Driver
  - Emulex driver for VMware

To configure boot from SAN:

1. Select **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4-3 on page 40), click the FC or FCoE adapter function on which you want to enable boot from SAN.
3. Select the **Maintenance** tab (Figure 8-16 on page 137 or Figure 8-17 on page 139), check **enable adapter boot** (if available) and click **Configure Boot**. The Boot from SAN Configuration dialog box appears (Figure 8-36 on page 167).

### Notes

- The Configure Boot button is disabled if the Enable Adapter Boot checkbox is not checked. If boot code is not present on the adapter, the Enable Adapter Boot checkbox and Configure Boot button are not displayed on the Maintenance tab.
- For OneConnect adapters, boot is always enabled and cannot be disabled.



**Figure 8-36** Boot from SAN Configuration Dialog Box

The Boot from SAN Configuration dialog box varies for each boot type. Figure 8-36 on page 167 depicts the boot from SAN configuration for the x86 type boot.

4. Verify that the adapter field contains the WWPN of the FC/FCoE function and boot BIOS version to make sure you configure the correct adapter FC/FCoE function and that it has the BIOS boot code version you want.
5. From the **Boot Type** menu, select x86, EFIBoot, or OpenBoot.

## Notes

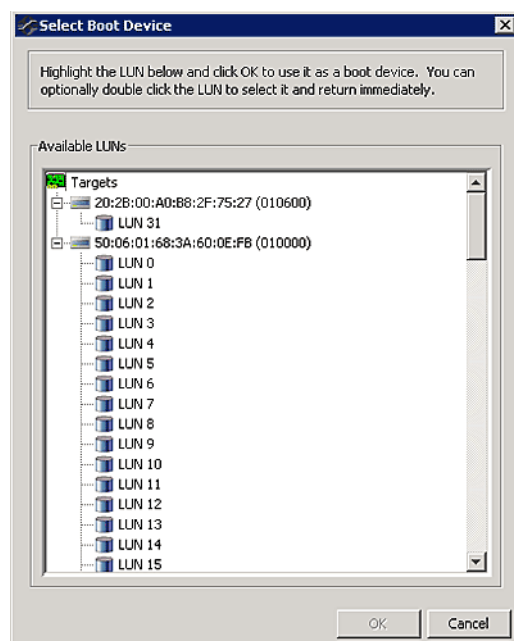
- x86 and OpenBoot share the same configuration memory space. You cannot configure an adapter for both x86 and OpenBoot at the same time. When you select one of these boot types and the configuration region is configured for the other boot type, a message appears warning that making changes overwrites 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 appears telling you to save the current settings before changing to the new boot type.
6. Check **Enable Boot from SAN** and for FC functions, set the Topology and Link Speed.

**Note:** Topology and link speed are not available for OneConnect adapters.

- Topology options are:
  - Auto, Loop First (default)
  - Auto, Point to Point First
  - Loop

- Point to Point
- Link speed options are:
  - Auto (default)
  - 1 Gbps (if available)
  - 2 Gbps (if available)
  - 4 Gbps (if available)
  - 8 Gbps (if available)
  - 16 Gbps (if available)
- 7. If you want, click **Advanced Settings** to configure autoscan, spinup delay and so on. See “Configuring Advanced Settings (Boot from SAN)” on page 169 for more information.
- 8. For x86 and EFIBoot, select one or more boot devices. For OpenBoot, select only one boot device.
- 9. Do one of the following on the Boot from SAN Configuration dialog box (Figure 8-36 on page 167):
  - Select **Target WorldWide Port Names**, type the numbers, and click **OK**.
  - Select **Target D\_ID**, type the numbers, and click **OK**.
  - Select **Target LUN**, type the number, and click **OK**.
    - For EFIBoot and OpenBoot, type in an 8-byte LUN (hex) and a target ID for the LUN. Also, you must enter the LUN value in “big endian” (most-significant byte, or “big end” first) order and enter all 16 characters including leading zeroes.
  - Click **Select from List**, select the target from a list of discovered LUNs (if available), and click **OK** on the Select Boot Device window (Figure 8-37 on page 169). While you can manually enter the target and LUN from the Boot from SAN Configuration dialog box, it is easier to select an existing LUN from this window. (See Figure 8-37.) The OneCommand Manager application attempts to update the boot parameters. If successful, a window appears with a confirmation message. Click **OK** on this confirmation window.





**Figure 8-37** Select Boot Device Window (for x86 or EFIBoot)

10. On the Boot from SAN Configuration dialog box (Figure 8-36 on page 167), click **Apply** to save your changes but leave the dialog box open, or click **OK** to apply the changes and close the dialog box.

**Note:** Click **Close** to close the Boot from SAN Configuration dialog box without saving your changes. A message appears to discard your changes.

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

## Configuring Advanced Settings (Boot from SAN)

The OneCommand Manager application provides advanced settings for each boot type. From the Boot from SAN Configuration dialog box (Figure 8-36 on page 167), click **Advanced Settings**. A boot type-specific dialog box allows you to enable options such as spinup delay and autoscan. If you do not use advanced settings, the default values are used.

If you make changes you must click **OK** to save the changes and close the dialog box. You can click **Cancel** and close the dialog box 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 dialog boxes in the following sections.

## x86 Boot Advanced Adapter Settings Dialog Box

Use the x86 Boot Advanced Adapter Settings dialog box (Figure 8-38 on page 170) to configure **advanced settings** for the selected x86 adapter. All checkboxes are cleared (off) by default. All changes require a reboot to activate.

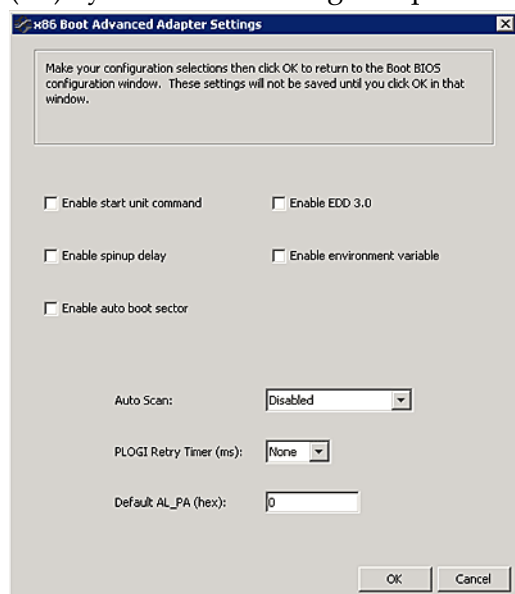


Figure 8-38 x86 Boot Advanced Adapter Settings Dialog Box

### x86 Boot Advanced Adapter Settings Definitions

- **Enable start unit command** – Issues the SCSI start unit command. You must know the specific LUN to issue.
- **Enable EDD 3.0** – Enables the EDD option showing the path to the boot device. (Available on Intel Itanium servers only.)

**Note:** An x86 series system could hang during Windows 2000 Server installation if EDD 3.0 is enabled.

- **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.

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

- If no boot devices have been defined and auto scan is enabled, then 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.
- 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 HP MSA1000 arrays. If there is no partition on the target, the default boot sector format is 63 sectors.
- Set Auto Scan – With auto scan enabled, the first device issues a Name Server Inquiry. The boot device is the first DID, LUN 0, or not 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 Multi-boot menu. Auto Scan is available only if none of the eight boot entries is configured to boot through DID or WWPN. Emulex strongly recommends that you use the Configure Boot Devices menu to configure eight boot entries for fabric point-to-point, public loop, or private loop configurations. Set to one of the following:
  - Disabled (default)
  - Any First Device
  - First LUN 0 Device
  - First non-LUN 0 Device
- Set the PLOGI Retry Timer – Sets the interval for the PLOGI (port log in) 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 to life, the PLOGI retry interval scans the loop to discover this device. This default setting is None (0 msec). Set to one of the following:
  - None (default)
  - 50 ms
  - 100 ms
  - 200 ms
- Type the Default AL\_PA number – It has a range of 00-EF (default=0). Changes the AL\_PA of the selected adapter. (Not available for OneConnect adapters.)

## EFIBoot Advanced Adapter Settings Dialog Box

Use the EFIBoot Advanced Adapter Settings dialog box (Figure 8-39 on page 172) to configure the advanced settings for the selected EFIBoot adapter.

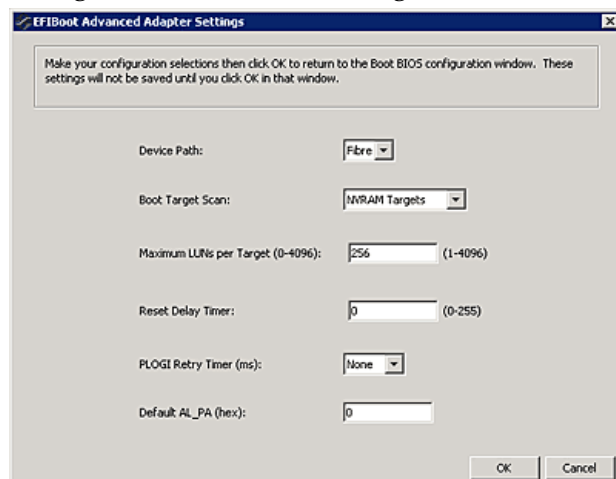


Figure 8-39 EFIBoot Advanced Adapter Settings Dialog Box

## EFIBoot Advanced Adapter Settings Field Definitions

- Device Path – Makes the Fibre driver appear as a SCSI driver.
  - Fibre (default)
  - SCSI
- Boot Target Scan – This option is available only if none of the eight boot entries are configured to boot via DID or WWPN.
  - NVRAM Targets (default) – Discovers only LUNs that are saved to the adapter Non-Volatile Random Access Memory (NVRAM).
  - Discovered Targets – Discovers all devices that are attached to the FC 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.
- Reset Delay Timer in seconds – Sets a value for delay device discovery. The range is 0 to 255. The default is 0.
- PLOGI Retry Timer – Sets the interval for the PLOGI (port log in) 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 online again the PLOGI retry interval scans the loop to discover this device.
  - 50 ms

- 100 ms
- 200 ms
- Default AL\_PA number – The range is 0x 00-EF. The default is 0x00. This option changes the AL\_PA (Arbitrated Loop Physical Address) of the selected adapter. (Not available for OneConnect adapters.)

## OpenBoot Advanced Adapter Settings Dialog Box

Use the OpenBoot Advanced Adapter Settings dialog box (Figure 8-40 on page 173) to configure the advanced adapter settings for the selected OpenBoot adapter.

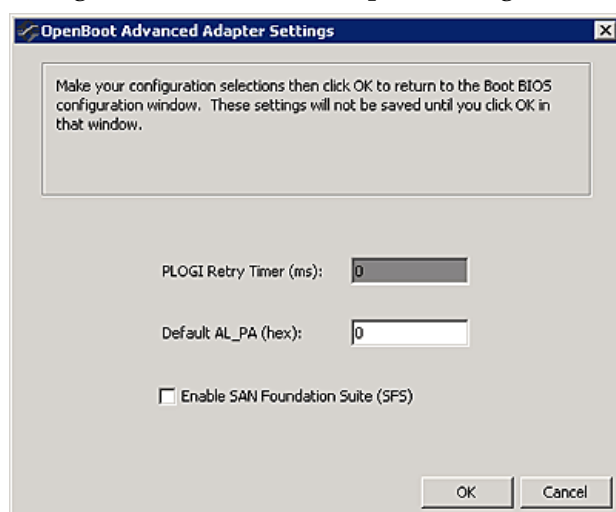


Figure 8-40 OpenBoot Advanced Settings Dialog Box

## OpenBoot Advanced Adapter Field Definitions

- PLOGI Retry Timer – Sets the PLOGI Retry timer value. Range is 0 to 0xFF.
- Default AL\_PA (hex) – Sets the default AL\_PA. The range is 0 to 0xEF. The default is 0. (Not available for OneConnect adapters.)
- Enable the Software Foundation Suite (SFS) – Check to enable the Software Foundation Suite (SFS) driver (the emlxs driver). The default is the FC, FCoE, or FC/FCoE driver.

## Using FC-SP DHCHAP Authentication (Windows, Linux 8.2, and Solaris)

Use the DHCHAP tab to view and configure FC-SP DHCHAP. You can authenticate an adapter to a switch.

### Notes

- DHCHAP is available only for FC functions, not FCoE functions.
- DHCHAP is not available on LPe15000 and LPe16000-series adapters.
- DHCHAP is available only for physical FC functions, not for virtual ports.
- DHCHAP is not supported on RHEL6+ and SLES11-SP1+.

- DHCHAP is not supported on OneConnect adapters.
- The authentication driver parameters are only available on local hosts. The OneCommand Manager application GUI does not display this driver parameter for any remote hosts.

After DHCHAP has been activated and configured, manually initiate authentication per adapter by clicking **Initiate Authentication** or by inducing a fabric login (FLOGI) time per 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.)

Authentication must be enabled at the driver level. Authentication is disabled by default. To enable DHCHAP using the Driver Parameters tab, enable one of the following parameters: enable-auth (in Windows), enable-auth (Solaris), or enable-auth (in Linux 8.2).

## Linux Considerations

To activate FC-SP/ Authentication between the adapter host port and fabric F\_Port using DHCHAP, you must modify the DHCHAP-associated driver properties in the driver configuration file.

The Emulex driver for Linux version 8.2.0.x supports MD5 and SHA-1 hash functions and supports the following DH groups: Null, 1024, 1280, 1536, and 2048.

**Note:** This version of the driver supports N-Port to F-Port authentication only and does not support N-Port to N-Port authentication.

## Enabling Authentication

Enabling authentication is a two step process. To enable authentication:

- The fcauthd daemon must be running.
- The lpfc\_enable\_auth module parameter must be set to enabled.

### lpfc\_enable\_auth Module Parameter

Use the lpfc\_enable\_auth module parameter to enable or disable authentication support. This module parameter can be set when loading the driver to enable or disable authentication on all Emulex adapters in the system, or it can be set dynamically after the driver is loaded to enable or disable authentication for each port (physical and virtual). The default setting for the lpfc-enable-auth module parameter is disabled.

### fcauthd Daemon

The Emulex LPFC driver requires the fcauthd daemon to perform authentication tasks for it. To enable authentication, you must have this daemon running. If you want to load the driver with authentication enabled, the fcauthd daemon should be running prior to driver load. The driver can start with authentication enabled if the daemon is not running, but all ports are placed into an error state. When the daemon is started the driver should discover the daemon and reset the adapter to enable the driver to perform authentication. To test if this daemon is running, start the daemon, or stop the

daemon, you must use the `/etc/init.d/fcauthd` script. This script accepts the standard daemon parameters: `start`, `stop`, `reload`, `status`, `restart`, and `condrestart`.

The script syntax is `/etc/init.d/fcauthd <parameter>`.

**Note:** The 8.2.0.X driver connects directly to the `fcauthd` daemon. To unload the driver, you must first stop the `fcauthd` daemon. This closes the netlink connection and allows the LPFC driver to unload.

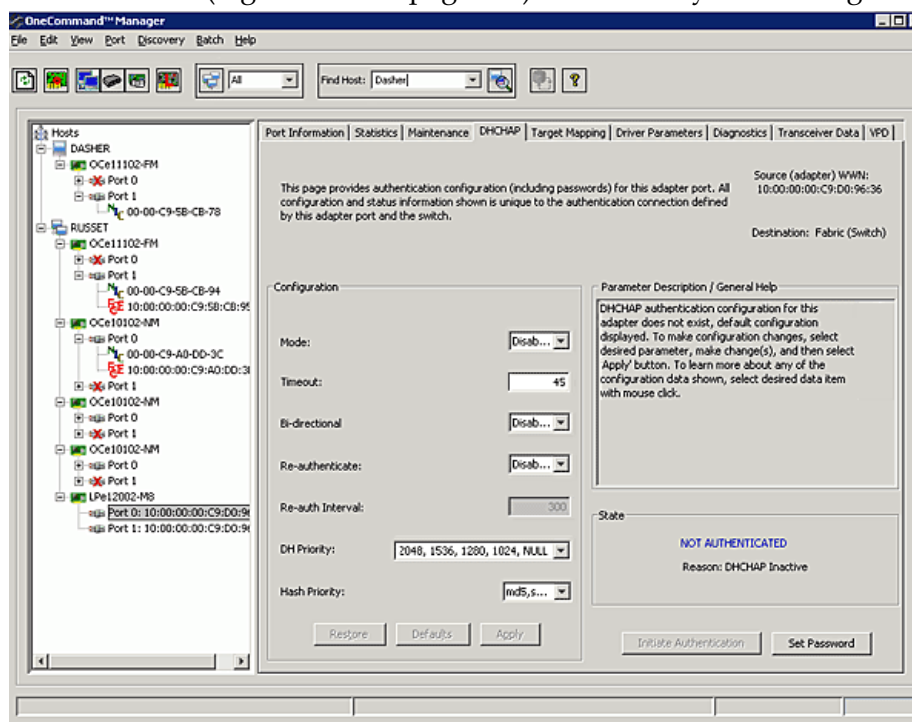
### **fcauthd Daemon Parameters**

The `fcauthd` daemon supports the following parameters:

- `start` – To start the `fcauthd` daemon, pass the `start` command to the `fcauthd` script. This command loads the daemon into memory, opens a netlink connection to the driver, and reads the authentication configuration database into memory for use by the LPFC driver.
- `stop` – To stop the `fcauthd` daemon, pass the `stop` command to the `fcauthd` script. This command takes down the netlink connection between the `fcauthd` daemon and the LPFC driver and stops the `fcauthd` daemon.
- `reload` – The `reload` command reloads the authentication configuration database into memory. This is done whenever the database is changed by another application (the OneCommand Manager application) or by you. If the database is changed, the new configuration information is not used until the `fcauthd` daemon reloads the database.
- `status` – This command is used to show the current status of the `fcauthd` daemon. The status should be either `running` or `stopped`.
- `restart` – The `restart` command performs a `stop` and then a `start`.
- `condrestart` – The conditional restart command checks the status of the `fcauthd` daemon. If it is `running`, it issues a `stop` and then a `start` command. If the `fcauthd` daemon is not `running`, nothing happens.

## DHCHAP Tab

The DHCHAP (Figure 8-41 on page 176) tab enables you to configure authentication.



**Figure 8-41** DHCHAP Tab

### DHCHAP Tab Field Definitions

- Source – The WWPN of the FC function.
- Destination – The fabric (switch).

#### Configuration Area

- Mode – The mode of operation. Three modes are available: Enabled, Passive, and Disabled.
  - Enabled – The FC function initiates authentication after issuing an 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 when 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.



- **Bi-Directional** – If enabled, the FC driver supports authentication initiated by either the switch or the FC function. If disabled, the driver supports FC function initiated authentication only.
- **Re-authenticate** – If enabled, the FC driver can periodically initiate authentication.
- **Re-auth Interval** – The value in minutes that the FC driver uses to periodically 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 two supported hash algorithms (MD5 and SHA1) that the FC driver presents during the DHCHAP authentication negotiation with the switch (default is MD5 first, then SHA1,2,3...).

#### State Area

- **State** – Possible states are Not Authenticated, Authentication In Progress, Authentication Success, and Authentication Failed.

## Changing Authentication Configuration

To view or change authentication configuration:

1. In the discovery-tree (Figure 4-3 on page 40), select the FC function whose configuration you want to view or change.
2. Select the **DHCHAP** tab (Figure 8-41 on page 176).

**Note:** If the fields on this tab are grayed-out, authentication has not been enabled at the driver level.

3. Change configuration values as you want.

**Note:** You can only configure DHCHAP on the local host.

4. Click **Apply**. You are prompted for the current password (local password) to validate the configuration change request. The verification request only appears if a local password has been defined for this adapter.
5. Enter the password and click **OK**.

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

**Note:** If you click **Apply**, changes cannot be cancelled.

To return all settings to the default configuration, click **Defaults**. Be careful as this also resets the password(s) to NULL for this configuration.

To initiate an immediate authentication, click **Initiate Authentication**. This request is sent to the driver, even if you have not made any changes to the setup.

**Note:** To successfully authenticate with the switch using DHCHAP, you only need to set the configuration mode to enabled and set the local password. The local password must be set to the identical value as the switch for the DHCHAP authentication to succeed.

## Changing Your Password

To change your password:

1. From the discovery-tree (Figure 4-3 on page 40), select the FC function whose password you want to change.
2. Select the **DHCHAP** tab (Figure 8-41 on page 176) and click **Set Password**. The Password dialog box is displayed.
3. Choose **Set Local Password** or **Set Remote Password**.
  - Local password is used by the FC driver when the adapter initiates authentication to the switch (typical use).
  - Remote password is used by the FC driver if the switch authenticates with the adapter. This is only possible when bi-directional is checked on the DHCHAP tab (Figure 8-41 on page 176).
4. If you want to see the password characters entered in the dialog box, check **Show Characters**.
5. Provide the current value for the password to validate the 'set new password' request (unnecessary if this is the first time the password is set for a given adapter).
6. Enter the new password.
7. Select alpha-numeric or hexadecimal format.
8. Click **OK**.

**Caution:** Do not forget the password after one has been assigned. After a password is assigned to an adapter, subsequent DHCHAP configuration settings for that adapter including 'default configuration' or new passwords require you to enter the existing password to validate your request (that is, no further changes can be made without the password).

**Note:** Additional help is available by clicking **Help** on the Set Password dialog box.

## Viewing the Error and Event Log

For Solaris and Linux systems, a simple shell script checks the /var/adm/messages and /var/log/messages files respectively for recent Emulex driver DHCHAP events and outputs them to a default location.

To view the error and event log:

1. Click **Event Log History** on the DHCHAP tab (Figure 8-41 on page 176).

## iSCSI

This section describes configuring iSCSI functions.

### Viewing iSCSI Port Information

When you select an iSCSI function from the discovery-tree (Figure 4-3 on page 40), the iSCSI Port Information tab (Figure 8-42 on page 179) contains general attributes associated with the selected iSCSI function..

**Note:** iSCSI functions are not discovered in the OneCommand Manager application if the adapter is managed by the CIM interface.

To view iSCSI function information:

1. Select **Host** view.
2. Select an iSCSI function in the discovery-tree (Figure 4-3 on page 40).
3. Select the **iSCSI Port Information** tab (Figure 8-42 on page 179).

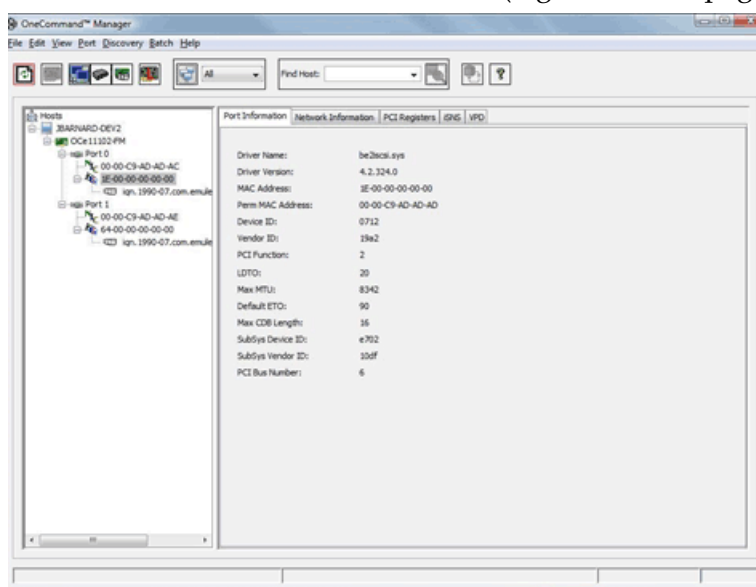


Figure 8-42 iSCSI Port Information Tab

### iSCSI Port Information Field Definitions

- Driver Name – The iSCSI driver file name.
- Driver Version – The iSCSI driver version.
- MAC Address – The iSCSI MAC address currently assigned to the function.
- Perm MAC Address – The original factory-assigned iSCSI MAC address.
- Device ID – The PCI device ID assigned to the iSCSI function.
- Vendor ID – The PCI vendor ID assigned to the iSCSI function.
- PCI Function – The PCI function number assigned to the iSCSI function.
- LDTO – The amount of time in seconds that the iSCSI driver delays reporting a link down error to the operating system.

- Max MTU – Maximum transmission unit for iSCSI traffic.
- Default ETO – The default extended timeout.
- Max CDB Length – Maximum SCSI command descriptor block size.
- SubSys Device ID – The PCI subsystem ID assigned to the iSCSI function.
- SubSys Vendor ID – The PCI subsystem vendor ID assigned to the iSCSI function.
- PCI Bus number – The PCI bus number assigned to the iSCSI function.

## Viewing iSCSI Network Information

The Network Information tab (Figure 8-43 on page 180) displays connection information such as link status and port speed. The tab also allows you to enable or disable DHCP authentication and VLANs, assign IP addresses, subnet masks, VLAN IDs and priorities, and more for the selected iSCSI port. You can also enable or disable iSCSI boot.

To view iSCSI network information:

1. Select **Host** view.
2. Select an iSCSI function in the discovery-tree.
3. Select the **iSCSI Network Information** tab (Figure 8-43 on page 180).

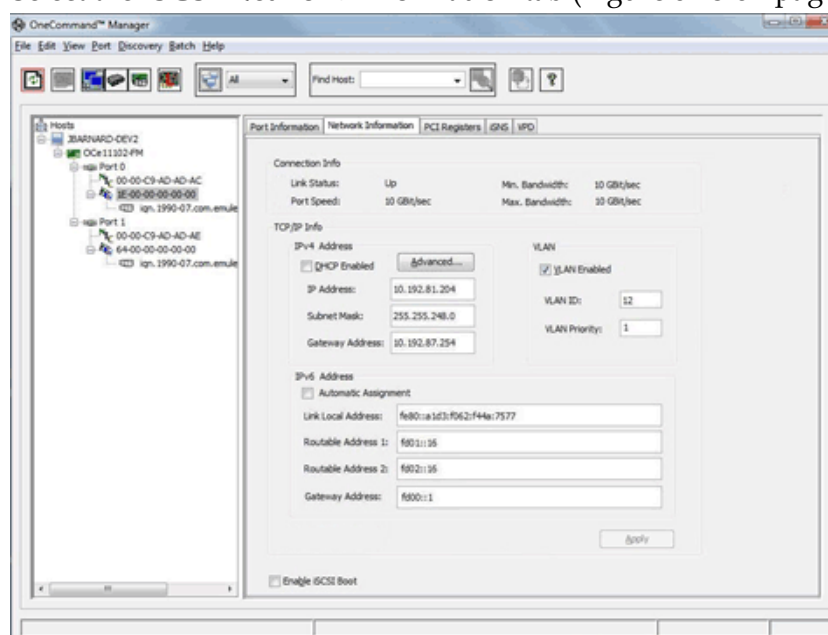


Figure 8-43 iSCSI Network Information Tab

## iSCSI Network Information Tab Field Definitions

### Connection Info Area

- Link Status – The status of the link on the selected iSCSI function.
- Port Speed – The speed at which the selected iSCSI function is running.

- **Min. Bandwidth** – The minimum bandwidth (that is, speed) at which the iSCSI function is guaranteed to run.
- **Max. Bandwidth** – The maximum bandwidth (that is, speed) at which the iSCSI function is guaranteed to run.

## TCP/IP Configuration Area

### IPv4 Address

- **DHCP Enabled** – Check the box to use DHCP to assign the IPv4 network settings on the selected iSCSI function.
- **IP Address** – The iSCSI initiator IP address.
- **Subnet Mask** – The iSCSI initiator subnet mask.
- **Gateway Address** – The iSCSI initiator gateway address.
- **VLAN Enabled** – Check the box to allow VLAN for the iSCSI interface.
- **VLAN ID** – The VLAN identifier to use 0-4094 (only valid if VLAN is enabled). 0 indicates the VLAN is disabled.
- **VLAN Priority** – The VLAN priority for the iSCSI interface.

### IPv6 Address (OCe14000-series adapters only)

- **Automatic assignment** – If enabled, unique routable addresses are assigned in conjunction with the router (similar for DHCP assignment).
- **Link Local Address** – The unique address assigned to the iSCSI function which is available for use inside the local network but not outside of the local network.
- **Routable Address 1** – A routable address assigned to the iSCSI function.
- **Routable Address 2** – A second routable address assigned to the iSCSI function.
- **Gateway Address** – The iSCSI initiator gateway address.

### General area

- **Enable iSCSI Boot checkbox** - Check or uncheck this box to enable or disable iSCSI boot on the iSCSI function.

## Network Information Buttons (Not available in read-only mode)

- **Advanced** – Click to display the TCP/IP Configuration dialog box. See “Advanced TCP/IP Configuration” on page 182 for more information.
- **Apply** – Click to make and save your changes.

## Modifying Network Settings

### Notes


- Checking DHCP Enabled disables the IP Address, Subnet Mask, and Gateway Address fields.
- If you enable both VLAN and DHCP, you must have a VLAN-enabled DHCP server to receive a valid IP address. If not, the request to the DHCP server for an IP address will fail.

To modify TCP/IP configurations for iSCSI ports:

1. From the discovery-tree (Figure 4-3 on page 40), select the iSCSI function whose configuration you want to modify.
2. Select the **Network Information** tab (Figure 8-43 on page 180).
3. Make your selections.
4. Click **Apply**.

## Advanced TCP/IP Configuration

The Advanced TCP/IP Configuration dialog box (Figure 8-44 on page 182) enables you to add and remove Route and ARP table entries (applicable to IPv4 only) for the selected iSCSI function.



**Figure 8-44** Advanced TCP/IP Configuration Dialog Box

To add table entries:

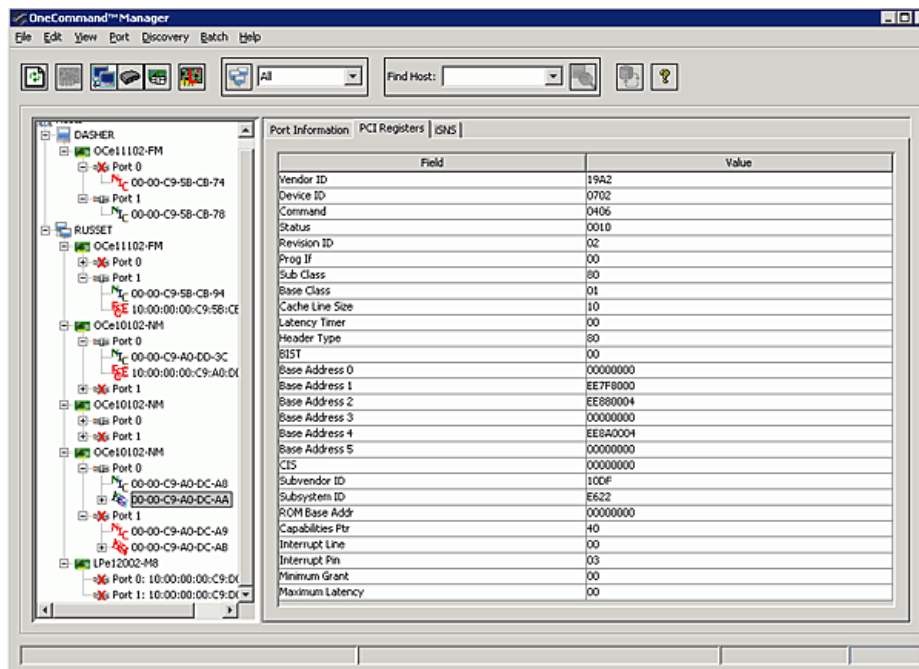
1. From the discovery-tree (Figure 4-3 on page 40), select the iSCSI function whose configuration you want to modify.
2. Select the **Network Information** tab (Figure 8-43 on page 180).
3. Click **Advanced**. The Advanced TCP/IP Configuration dialog box appears (Figure 8-44 on page 182).
4. From the Route Table or ARP Table sections, click **Add Entry**.
5. Enter the Route Table or ARP Table information and click **OK**. The entry you added appears in the table.

To delete table entries:

1. From the discovery-tree (Figure 4-3 on page 40), select the iSCSI function whose configuration you want to modify.
2. Select the **Network Information** tab (Figure 8-43 on page 180).
3. Click **Advanced**. The Advanced TCP/IP Configuration dialog box appears (Figure 8-44 on page 182).
4. From the Route Table or ARP Table sections, select the entry you want to delete and click **Remove Entry**. The entry you removed is deleted from the table.

## Viewing iSCSI PCI Registers

The iSCSI PCI Registers tab (Figure 8-45 on page 183) displays base PCI registers. See “Viewing the PCI Registers” on page 237 for FC PCI register information.



**Figure 8-45** iSCSI PCI Registers Tab

To view iSCSI PCI registers:

1. From the discovery-tree, select the iSCSI function whose PCI information you want to view.
2. Select the iSCSI **PCI Registers** tab (Figure 8-45 on page 183).

## Viewing iSCSI VPD Information

The VPD tab (Figure 8-46 on page 184) displays vital product data (if available) for the selected iSCSI function such as the product name, part number, serial number, and so on.

To view VPD information:

1. Select **Host** view.
2. In the discovery-tree, select the iSCSI function whose VPD information you want to view.
3. Select the **VPD** tab (Figure 8-46 on page 184).

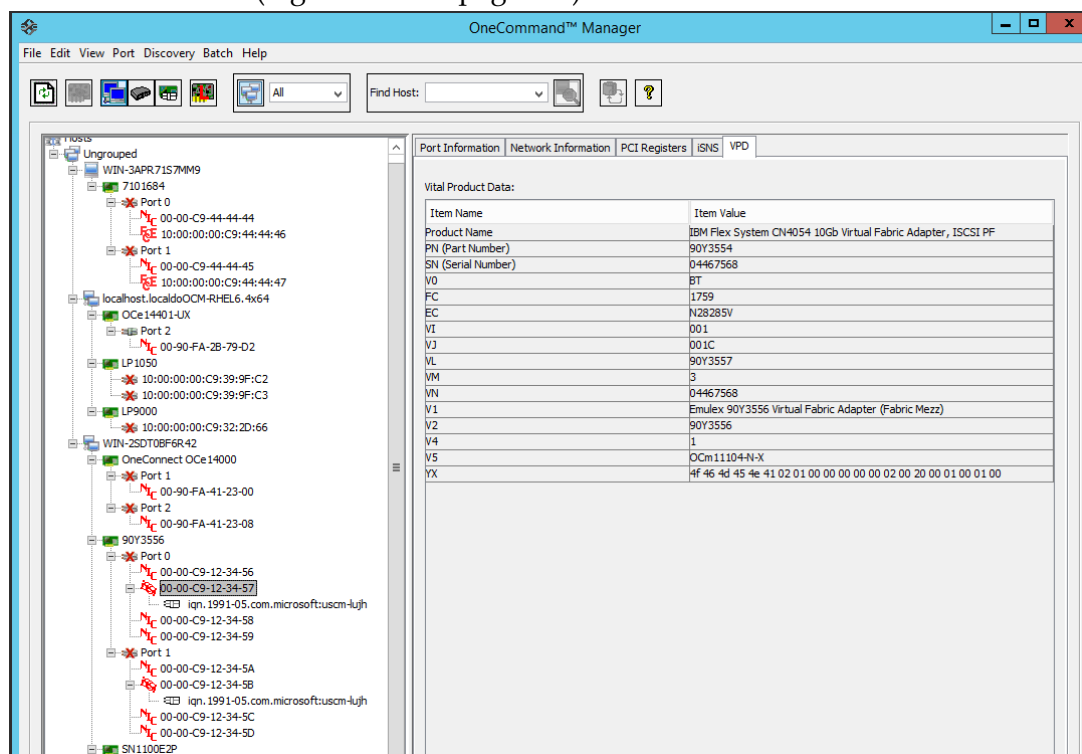


Figure 8-46 iSCSI VPD Tab

## VPD Table Definitions

- Product Name – Product information about the selected adapter port.
- PN (Part Number) – The adapter's part number.
- SN (Serial Number) – The adapter's serial number.
- VO – Vendor unique data. “V” indicates a vendor-specific field. An adapter may have none, one, or more of these fields defined. Valid values for this field are “VO” (the letter “O”, not the number zero) and “Vx” (where “x” is a number).

**Note:** Some adapters may show additional VPD information such as EC (EC level) and MN (manufacturer ID).



## Viewing iSCSI Statistics

When you select an iSCSI initiator node from the discovery-tree (Figure 4-3 on page 40), the iSCSI Statistics tab (Figure 8-47 on page 185) provides cumulative totals for various error events and statistics on the port.

**Note:** iSCSI statistical information is not available via the CIM interface.

To view iSCSI port statistics:

1. Select **Host** view.
2. Select an iSCSI initiator node in the discovery-tree.
3. Click the **iSCSI Statistics** tab (Figure 8-47 on page 185).

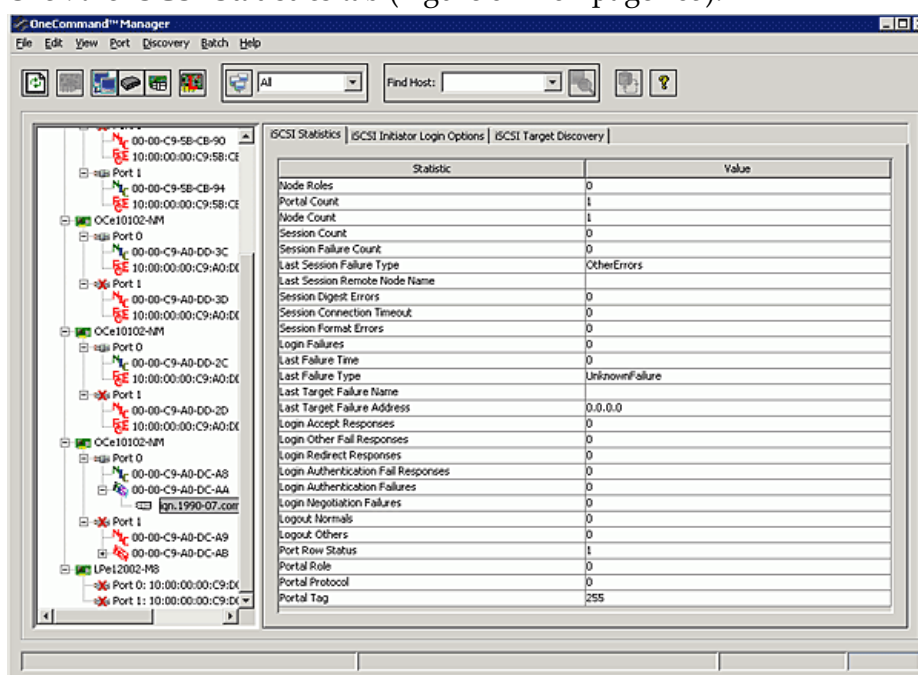


Figure 8-47 iSCSI Statistics Tab

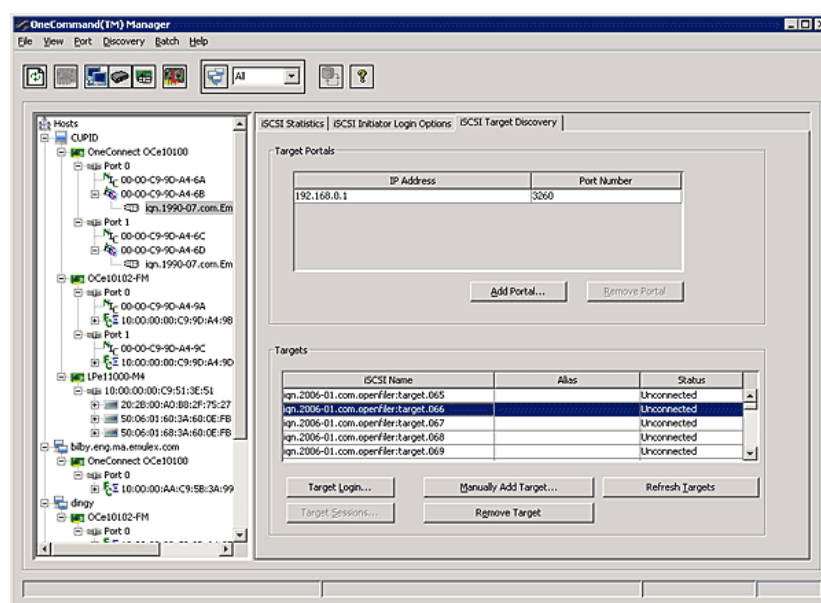
## iSCSI Statistics Field Definitions

- **Node Roles** – The node role for this iSCSI initiator.
- **Portal Count** – The number of rows in the `iscsiPortaltypesTable` that are currently associated with this iSCSI instance.
- **Node Count** – The number of rows in the `iscsiNodetypesTable` that are currently associated with this iSCSI instance.
- **Session Count** – The number of rows in the `iscsiSessiontypesTable` that are currently associated with this iSCSI instance.
- **Session Failure Count** – The number of times a session belonging to this initiator has failed.
- **Last Session Failure Type** – The type of failure encountered in the last session failure.

- Last Session Remote Node Name – The iSCSI name of the remote node from the failed session.
- Session Digest Errors – The count of sessions that failed due to receipt of a PDU containing header or data digest errors.
- Session Connection Timeout – The count of sessions that failed due to a sequence exceeding a time limit.
- Session Format Errors – The count of sessions that failed due to receipt of an iSCSI PDU that contained a format error.
- Login Failures – The number of times a login from this initiator failed.
- Last Failure Time – The timestamp of the most recent failure of a login attempt from this initiator. A value of 0 indicates that no failures have occurred.
- Last Failure Type – A description of the last failure.
- Last Target Failure Name – The UTF-8 string name of the target that most recently failed a login request from this initiator.
- Last Target Failure Address – The Internet Network Address of the target that most recently failed.
- Login Accept Responses – The count of Login Response PDUs received by this initiator that were accepted.
- Login Other Fail Responses – The count of Login Response PDUs received by this initiator with any status code not counted by the other objects.
- Login Redirect Responses – The count of Login Response PDUs received by the initiator with status class Redirection.
- Login Authentication Fail Responses – The count of Login Response PDUs with status class 0x201 Authentication Failed received by this initiator.
- Login Authentication Failures – The number of times the initiator has aborted a login because the target could not be authenticated.
- Login Negotiation Failures – The number of times the initiator has aborted a login because parameter negotiation with the target failed.
- Logout Normals – The count of Logout Command PDUs generated by this initiator with reason code normal.
- Logout Others – The count of Logout Command PDUs generated by this initiator with any status code other than normal.
- Port Row Status – This field allows entries to be dynamically added and removed from this table using SNMP.
- Portal Role – The role of a portal. A portal can operate in either one of two roles as a target portal or an initiator portal.
- Portal Protocol – The portal's transport protocol.
- Portal Tag – The portal's aggregation tag if the portal is used as an initiator.

## Configuring iSCSI Target Discovery

The iSCSI Target Discovery tab allows you to configure iSCSI target discovery related parameters (Figure 8-48 on page 187) .



**Figure 8-48** iSCSI Target Discovery Tab

To display the iSCSI Target Discovery tab:

1. From the discovery-tree, select the iSCSI initiator node whose discovery settings you want to configure.
2. Select the **iSCSI Target Discovery** tab (Figure 5-6 on page 53).

### Target Discovery Field Definitions

- **Target Portals** – The Target Portals table contains all target portals that are queried for targets. Depending on the SAN setup, the contents of this table may be a subset of the available target portals, or it could contain the full set of target portals for all iSCSI targets.
- **Targets** – The Targets table contains all currently discovered targets. Targets in this table come from one of three possible sources:
  - The target was manually added.
  - The target was discovered via a target portal.
  - The target was found through an iSNS server query.

### Target Discovery Buttons

- **Add Portal** – Click to add a target portal. See “Adding Target Portals” on page 188 for more information.
- **Remove Portal** – Click to remove a portal. See “Removing a Target Portal” on page 188 for more information.

- **Target Login** – Click to log into a selected target. See “Logging into Targets” on page 190 for more information.
- **Target Sessions** – Click to view active sessions for the selected target. See “Viewing Target Sessions” on page 192 for more information.
- **Manually Add Target** – Click to manually add an iSCSI target. See “Manually Adding an iSCSI Target” on page 192 for more information.
- **Remove Target** – Click to manually remove an iSCSI target. See “Removing Targets” on page 192 for more information.
- **Refresh Targets** – Click to manually force a complete rediscovery of the targets, querying all configured iSNS servers and target portals.

## Adding Target Portals

To add a target portal:

1. From the iSCSI Target Discovery tab (Figure 8-48 on page 187), click **Add Portal**. The Add Target Portal dialog box appears.
2. Enter the server IP address and TCP port number and click **OK**. After successfully adding a target portal, that target portal's targets are discovered and appear in the target list.
3. Specify the Portal Login Options and Authentication type you want to use. For more information, see “Initiator Login Options Tab Field Definitions” on page 198.
4. Click **OK**.

**Note:** When running the open-iSCSI driver, the default open-iSCSI configuration is for “automatic target login” at boot time for all targets discovered through the target portal.

## Removing a Target Portal

To remove a target portal:

1. From the iSCSI Target Discovery tab (Figure 8-48 on page 187), select the target portal you want to remove in the Target Portals table.
2. Click **Remove Portal**.

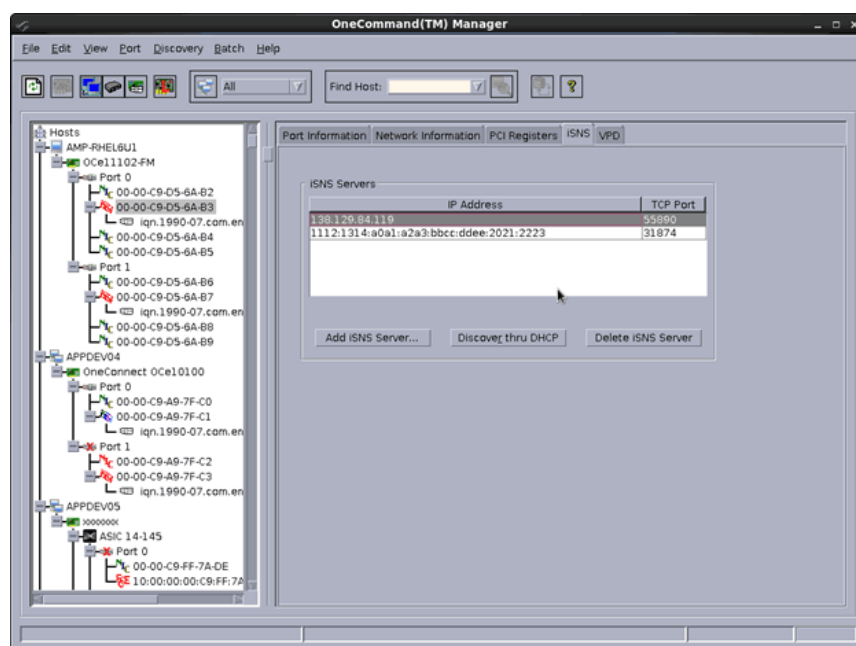
**Note:** The discovered targets are not removed by clearing the target portal. They must be specifically removed using the iSCSI Target Discovery tab. However, Windows targets that are not logged in to when the system is rebooted are removed.

## Configuring iSNS for iSCSI Target Discovery

An iSNS maintains a database of storage network elements that can be queried for iSCSI targets by other hosts within the SAN. iSCSI storage devices in particular can register targets with the iSNS for efficient discovery by iSCSI clients, such as the OneCommand Manager application.

**Note:** On Linux hosts running the Open-iSCSI driver, iSNS is not supported.


Use the iSCSI iSNS tab (Figure 8-49 on page 189) to configure the iSNS server or to discover the server using DHCP.



**Figure 8-49** iSCSI iSNS Tab

To add a server:

1. Select an iSCSI function in the discovery-tree and the iSNS tab.
2. Click **Add iSNS Server**. The Add iSNS Server dialog box appears (Figure 8-50 on page 189).



**Figure 8-50** Add iSNS Server Dialog Box

3. Enter the IPv4 or IPv6 address of the iSNS server and the TCP port number.

### Notes

- On OCE14000-series adapters, an IPv4 and a IPv6 iSNS address can be added. For other adapters, only an IPv4 address can be added.
  - On OCE14000-series adapters, up to four iSNS servers can be added. For other adapters, only one iSNS server can be added.
4. Click **OK**. The server is pinged and the new server appears in the iSNS Server list.

**Note:** If the server cannot be pinged, a popup message appears indicating this and you must choose whether or not to add the address.

The new iSNS server is also queried for iSCSI targets and any discovered targets are added to the Target's table on the main Target Discovery tab.

To remove a server:

Select the server from the iSNS Server list and click **Delete iSNS Server** (Figure 8-49 on page 189). The iSNS server is removed and no longer queried during a target refresh.

**Note:** The targets discovered using iSNS are not removed by clearing the iSNS server. They must be specifically removed in the iSCSI Target Discovery tab. However (except on ESXi hosts), targets that are not logged in to when the system is rebooted are removed.

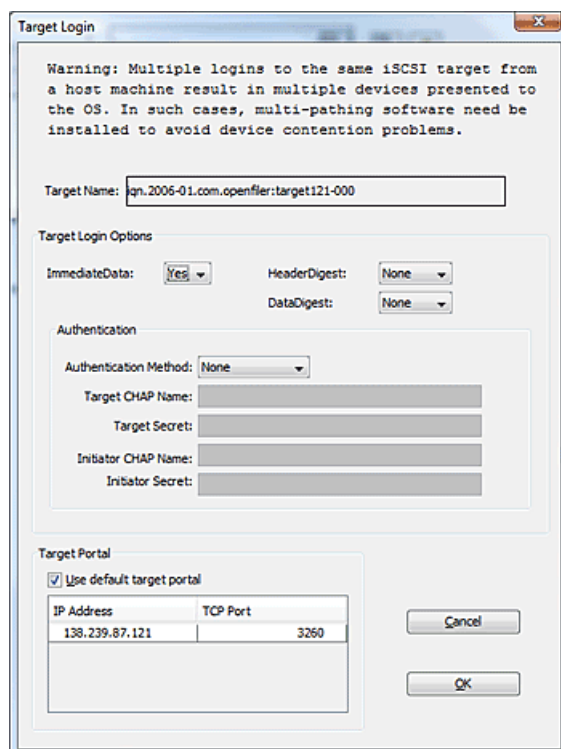
To discover iSNS servers using DHCP:

Click Discover thru DHCP (Figure 8-49 on page 189). If an iSNS server can be discovered through a DHCP server, it is added to the iSNS Server list. After an iSNS is added, its discovered targets should appear in the Targets list of the iSCSI Target Discovery tab.

## Logging into Targets

Only connected targets (targets that are successfully logged in to) are displayed in the discovery-tree (Figure 4-3 on page 40). However, the Targets table in the iSCSI Target Discovery tab (Figure 8-48 on page 187) is composed of all discovered targets regardless of their connection status. The connection status of each target is displayed in the 'Status' column of the Targets table. Disconnected targets are targets that have not yet been logged in to by the initiator.

**Note:** The target's login options are set at the time they are discovered from the target portal and match the target portal's login options. Changing the login options in the Initiator Login Options tab does not change the discovered targets login options.



**Figure 8-51** Target Login Dialog Box

To log into a target:

1. From the iSCSI Target Discovery tab (Figure 8-48 on page 187), select the target from the Targets table.
2. Click **Target Login**. The Target Login dialog box appears (Figure 8-51 on page 191). The dialog box displays the Target Name and Target Alias of the target. When you log into a target and reboot the system, the adapter automatically logs in to that target after the reboot is complete.
3. Specify the Target Login Options and Authentication method you want to use. For more information, see “Initiator Login Options Tab Field Definitions” on page 198.
4. If more than one Target Portal is available to log into the target, you can select the target portal you want to use from the Target Portal list. To use the default Target Portal, check “Use default target portal.”
5. Click **OK**. If the target was successfully logged into, the target's status in the Targets table changes to 'Connected'.

## Notes

- If you are logging into a target more than once, or you are logging into the same target from multiple iSCSI ports, you must have multi-pathing software installed to properly present the target's LUNs to the operating system.
- When running the open-iSCSI driver, only one login session is allowed per target. Subsequent login attempts to a target with an existing session results in an error.



## Manually Adding an iSCSI Target

The iSCSI Target Discovery tab (Figure 8-48 on page 187) enables you to manually add and log into iSCSI targets.

To manually add an iSCSI target:

1. From the iSCSI Target Discovery tab (Figure 8-48 on page 187), click **Manually Add Target**. The Add iSCSI Target dialog box appears.
2. Enter the target iSCSI name, target IP address, and TCP port number.
3. Specify the Target Login Options and Authentication type you want to use. For more information, see “Initiator Login Options Tab Field Definitions” on page 198.
4. Check the boot checkbox to add the target as a boot target.
5. Check the auto-login checkbox to log into the target after adding it.
6. Click **OK**. If the target was successfully added and logged into, the target appears as ‘Connected’ in the Targets table.

**Note:** If you add a target with the auto-login checkbox unchecked, a “closed” target session is created for the target. This allows you to change a manually added target to a boot target. Unlike “discovered” targets added through a target portal or iSNS, upon reboot the system attempts to log in to the target even if the target was not logged in to before the reboot. It is not automatically removed as is the case with “discovered” targets.

## Removing Targets

To remove a target:

1. Log out of all sessions for the target you want to remove.
2. From the iSCSI Target Discovery tab (Figure 8-48 on page 187), select the target you want to remove and click **Remove Target**.

## Viewing Target Sessions

The Target Sessions dialog box (Figure 8-52 on page 193) enables you to view sessions for a target.

**Note:** Making a target session a boot session for open-iSCSI implementation in the OneCommand Manager application is not supported. Therefore, when running the Emulex open-iSCSI driver, the OneCommand Manager application does not display the open-iSCSI boot checkbox on the target session data screen.

To view sessions for a target:

1. From the iSCSI Target Discovery tab (Figure 8-48 on page 187), select the target whose sessions you want to view and click **Target Sessions**. The Target Sessions dialog box appears (Figure 8-52 on page 193).
2. Use the Session pull-down menu to select the session whose information you want to view.
3. Click **Close** to close the dialog box.



## Logging out of Target Sessions

The Target Sessions dialog box enables you to log out of active sessions for a currently connected target.

To log out of active sessions for a connected target:

1. From the iSCSI Target Discovery tab (Figure 8-48 on page 187), select the target whose sessions you want to log out of and click **Target Sessions**. The Target Sessions dialog box appears (Figure 8-52 on page 193).
2. From the **Session** pull-down menu, select the session from which you want to log out.
3. Click **Logout** to log out of the session.

**Note:** If all sessions are logged out, the target is disconnected and removed from the discovery-tree (Figure 4-3 on page 40). However, the target is still available for login later.

The Target Sessions dialog box displays the following information:

- Session:** 0 (dropdown menu)
- Logout:** (button)
- Initiator Name:** iqn.1990-07.com.Emulex.00-00-c9-5b-3b-b9
- iSCSI Boot:** ☒
- Status:** Logged in
- ISID:** 0x4000006c0000
- ISIDQualifier:** 0x0
- TSIH:** 1792
- Session Negotiated Login Options:**

Initial R2T:	Yes	DefaultTimeToWait:	2
Immediate Data:	Yes	DefaultTimeToRetain:	20
Max Connections:	1	ErrorRecoveryLevel:	0
MaxOutstandingR2T:	1	DataPDUInOrder:	0
FirstBurstLength:	8192	DataSequenceInOrder:	0
MaxBurstLength:	262144		
- Session Statistics:**

Session Direction:	0
Cmd PDUs:	28
Response PDUs:	28
Xmt Data Octets:	0
Recv Data Octets:	4911
Digest Errors:	0
Connection Timeouts:	0
Session Target Alias:	
- Connection Information:**

iSCSI Connection ID:	1	Status:	Connected
Source IP Address:	192.168.0.248	Source Port:	39890
Destination IP Address:	192.168.0.1	Destination Port:	3260
Redirected Destination:	0.0.0.0	Redirected Destination Port:	0
- Connection Negotiated Login Properties:**

Authentication Method:	None	MaxRecvDataSegmentLength:	16384
Header Digest:	None	TCPMSS:	1436
Data Digest:	None		

Figure 8-52 Target Sessions Dialog Box

## Target Sessions Field Definitions

- **Initiator Name** – The initiator named used to log into the session.
- **Status** – The session status (logged in, login in progress, login failed, recovery, unknown).
- **ISID** – The ISID (unique for each session).
- **ISID Qualifier** – The first two bytes of the ISID (unique for each session).

- **TSIH** – The target session identifier handle. A tag generated by an iSCSI target to identify an iSCSI session with a specific iSCSI initiator. (Not available when running the open-iSCSI driver.)
- **iSCSI Boot** – If checked, the target is enabled for boot. However, it does not necessarily mean that the target is the current boot device. Check this box to enable boot from this target. Uncheck it to disable boot from this target. You must uncheck this box to log out from a target. However, even if you uncheck the box, logout is not possible from the currently booted target.

**Note:** Not supported when running the open-iSCSI driver.

#### Session Negotiated Login Options Area

- **InitialR2T** – The initial request to transmit. If set to Yes, the initiator has to wait for the target to solicit SCSI data before sending it. If set to No, it allows the initiator to send a burst of unsolicited FirstBurstLength bytes.
- **Immediate Data** – If set to Yes, it allows the initiator to append unsolicited data to a command.
- **Max Connections** – The maximum number of connections to targets that are allowed within a single session.
- **MaxOutstandingR2T** – The maximum number of outstanding request to transmits (R2Ts) per task within a session, each up to MaxBurstLength bytes.
- **FirstBurstLength** – The maximum amount of unsolicited data (in bytes) the initiator can send to the target during the execution of a single iSCSI command.
- **MaxBurstLength** – The maximum amount of either unsolicited or solicited data the initiator may send in a single burst. Any amount of data exceeding this value must be explicitly solicited by the target.
- **DefaultTimeToWait** – The minimum time to wait, in seconds, before the initiator attempts to reconnect or reassign a connection (or task) that has been dropped after an unexpected connection termination or reset. The initiator and target negotiate to determine this value.
- **DefaultTimeToRetain** – The maximum time, in seconds, to reassign a connection after the initial wait that is indicated in DefaultTimeToWait has elapsed. The initiator and target negotiate to determine this value.
- **ErrorRecoveryLevel** – The operational ErrorRecoveryLevel for the session. 0 indicates recovery only by session restart. 1 indicates recovery by reissuing commands, data, or status. 2 indicates connection failure recovery.
- **DataPDUInOrder** – The order of PDUs within a sequence.
- **DataSequenceInOrder** – The order between sequences.

#### Session Statistics Area

- **Session Direction** – The direction of iSCSI session. Valid values are InboundSession and OutboundSession.
- **Cmd PDUs** – The count of Command PDUs transferred on this session.
- **Response PDUs** – The count of Response PDUs transferred on this session.

- Xmt Data Octets – The count of data octets that were transmitted by the local iSCSI node on this session.
- Recv Data Octets – The count of data octets that were received by the local iSCSI node on this session.
- Digest Errors – The count of PDUs that were received on the session and contained header or data digest errors.
- Connection Timeouts – The count of connections within this session that have been terminated due to a timeout.
- Session Target Alias – The target alias for the session.

#### Connection Information Area

- iSCSI Connection ID – The iSCSI Connection ID assigned to the connection.
- Status – The status of the connection. Valid values are connected and unknown.
- Source IP Address – The source IP address for the connection.
- Source Port – The source TCP port number for the connection.
- Destination IP Address – The destination IP address for the connection.
- Destination Port – The destination TCP port number for the connection.
- Redirected Destination – The redirected IP address for the target.
- Redirected Destination Port – The redirected port number for the target.

#### Connection Negotiated Login Options

- Authentication Method – The authentication method used for connection. Valid values are None, Mutual CHAP, and One-Way CHAP.
- MaxRecdDataSegmentLength – The maximum data segment length in bytes an initiator or target can receive in an iSCSI PDU.
- Header Digest – If set to CRC32C, the integrity of an iSCSI PDU's header segments is protected by a CRC32C checksum.
- Data Digest – If set to CRC32C, the integrity of an iSCSI PDU's data segments is protected by a CRC32C checksum.
- TCPMSS – The maximum segment size for this connection. The driver uses this to determine the size of the data PDU whenever it is required to transmit the entire PDU with a single iSCSI header.

## Viewing iSCSI Target Information

When you select a target associated with a iSCSI adapter from the discovery-tree (Figure 4-3 on page 40), the Target Information tab (Figure 8-53 on page 196) displays information associated with that target.

To view iSCSI target information:

1. Select **Host** view.
2. In the discovery-tree, select the iSCSI target whose information you want to view. The Target Information tab appears (Figure 8-53 on page 196).

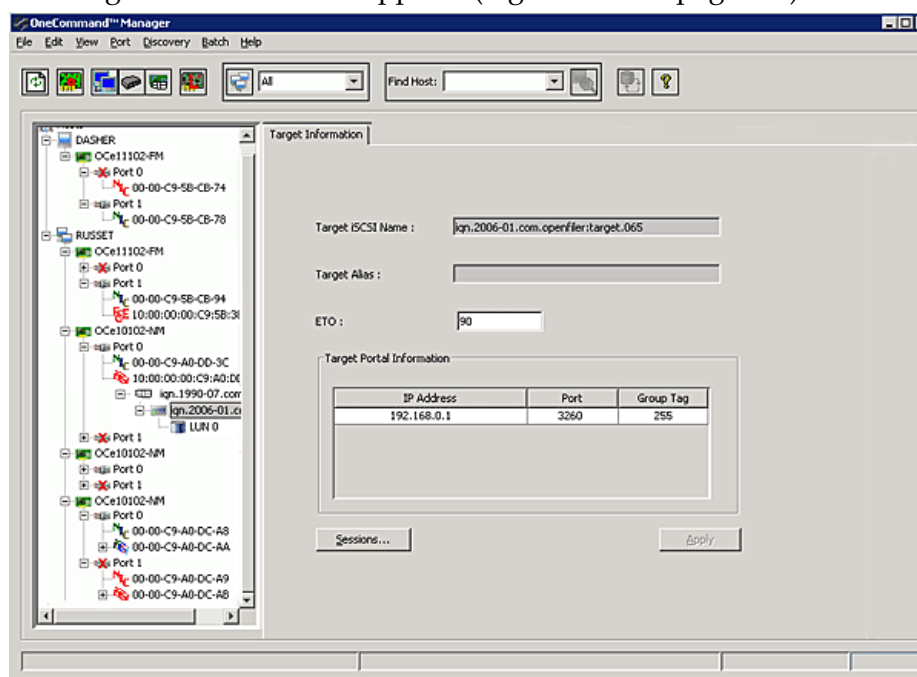


Figure 8-53 iSCSI Target Information Tab

## Target Information Field Definitions

- Target iSCSI Name – The iSCSI name assigned to the target.
- Target Alias – The iSCSI alias assigned to the target. This is assigned at the target portal, not by the OneCommand Manager application.
- ETO – The ETO for the target. The driver ensures that I/Os are not “timed out” until this time has expired (from the time the target stopped responding). You can change this value if you want.

## Target Portal Information Area

- IP Address – The IP address through which the initiator communicates with the target.
- Port – The TCP port through which the initiator communicates with the target.
- Group Tag – The tag of the group for which sub-groups must be refreshed.

## Target Information Buttons

- Sessions... – Click to view the currently active sessions for the target. See “Viewing Target Sessions” on page 192 for more information.
- Apply – Click to save and apply your ETO changes.

## Viewing iSCSI LUN Information

When you select a LUN associated with an iSCSI adapter from the discovery-tree, the LUN Information tab (Figure 8-54 on page 197) displays information associated with that LUN.

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

To view the LUN information:

1. Select **Host** view.
2. From the discovery-tree, select the iSCSI LUN whose information you want to view. The LUN Information tab appears (Figure 8-54 on page 197).

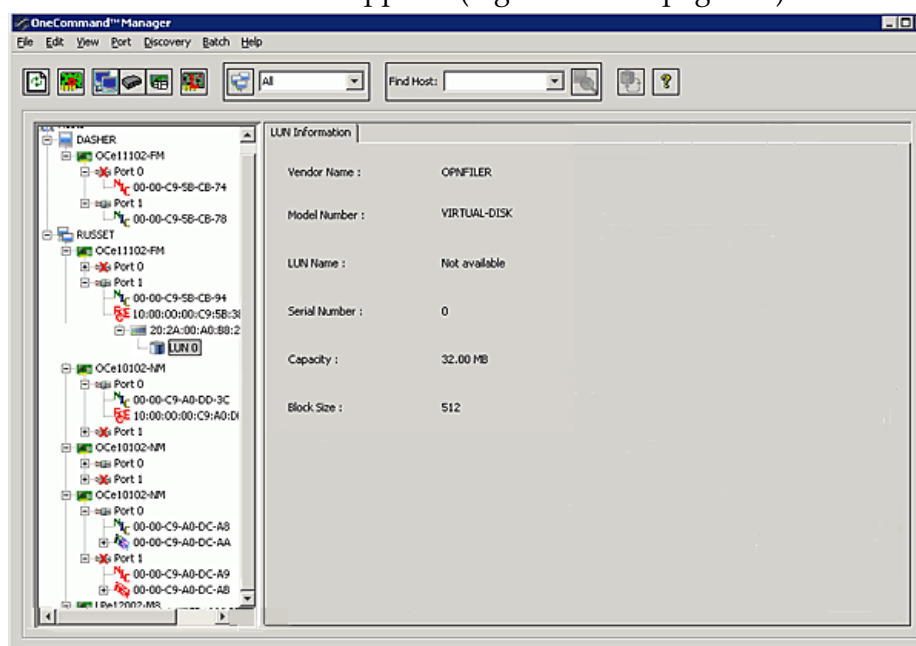


Figure 8-54 iSCSI LUN Information Tab

## iSCSI LUN Information Field Definitions

- Vendor Name – The name of the vendor of the LUN.
- Model Number – The vendor’s model number for the LUN.
- LUN Name – The name of the LUN. (Available only on ESXi platforms.)
- Serial Number – The vendor’s serial number for the LUN.
- Capacity – The unformatted size of the LUN.
- Block Size – The size of a logical unit block in bytes.

## Configuring iSCSI Initiator Login Options

The iSCSI Initiator Login Options dialog box (Figure 8-55 on page 198) enables you to configure the set of login options used by the iSCSI initiator when logging into a target portal or by the target portal when it is discovering targets. The discovered targets inherit the login options used during this discovery. Target portals discovered using iSNS also use these login options. The dialog box contains the initiator IQN and fields for manually entering the IQN and an optional initiator alias. Initiator login options are controlled using several drop down boxes. You can also configure the initiator authentication method and view the factory default login options.

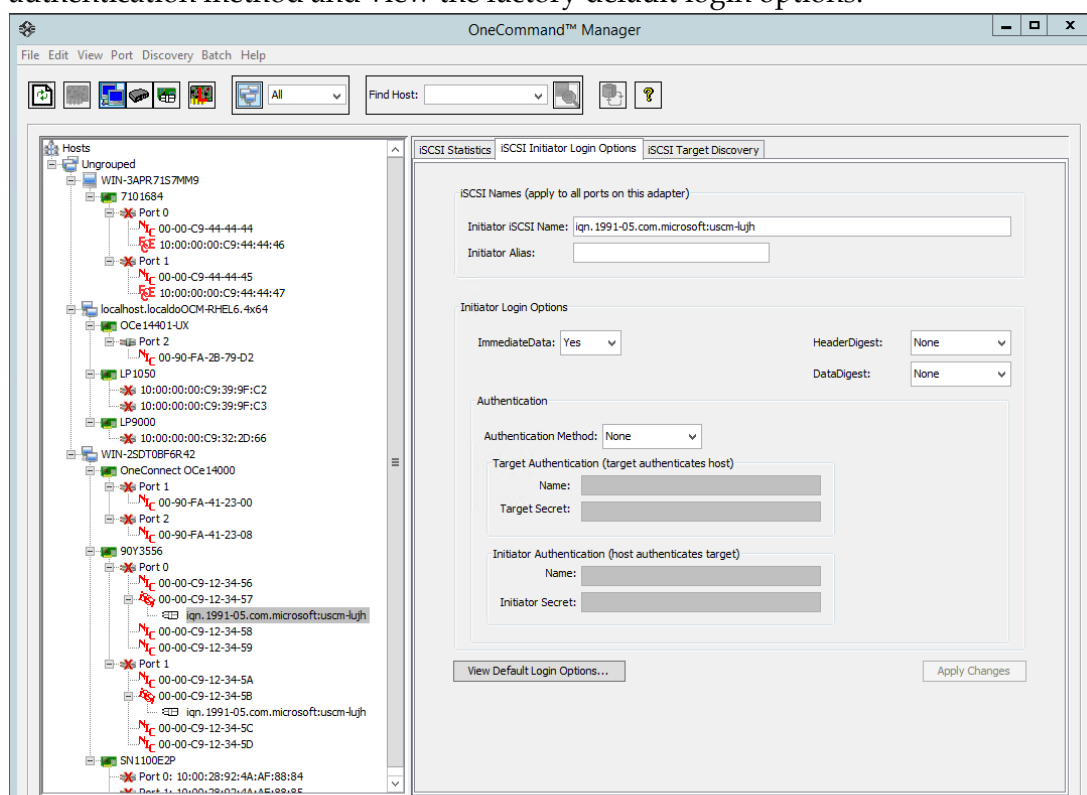


Figure 8-55 iSCSI Initiator Login Options Tab

## Initiator Login Options Tab Field Definitions

### iSCSI Names Area

- Initiator iSCSI Name – The iSCSI qualifier name of the initiator.
- Initiator Alias – An optional non-unique string used to identify the initiator.

### Initiator Login Options Area

- ImmediateData – Defines whether the initiator may append unsolicited data to a SCSI command. Possible values are “Yes” and “No”.
- HeaderDigest – If set to “CRC32C”, and the initiator is configured accordingly, the integrity of an iSCSI PDU's header segments are protected by a CRC32C checksum. Possible values are “CRC32C” and “None”.

- **DataDigest** – If set to “CRC32C” and the initiator is configured accordingly, the integrity of an iSCSI PDU's data segment is protected by a CRC32C checksum. Possible values are “CRC32C” and “None”.

### Authentication Area

- **Authentication Method** – Three options are available for the Authentication method: “None”, “One-Way CHAP” and “Mutual CHAP”. One-Way CHAP requires only that the authenticator (iSCSI target) authenticate the iSCSI initiator host. Mutual CHAP requires that both the iSCSI target and iSCSI initiator host authenticate each other. When “None” is selected, no authentication is performed.
- **Target Authentication Name** – The name sent to the target for authentication. This parameter is required for both One-Way CHAP and Mutual CHAP authentication. It can be any sequence of characters and numbers. The minimum length of the name is 1 character and the maximum length is 256 characters.
- **Target (Authentication) Secret** – The secret sent to the target for the target to use to authenticate the host. This parameter is required for both One-Way CHAP and Mutual CHAP authentication. It can be any sequence of characters and numbers. The minimum length of the secret is 12 characters and maximum length is 16 characters.
- **Initiator (Authentication) Name** – The name that is received from the target so the initiator can authenticate the target. This parameter is only required for Mutual CHAP authentication. It can be any sequence of characters and numbers. The minimum length of the name is 1 character and the maximum length is 256 characters.
- **Initiator (Authentication) Secret** – The secret that is received from the target so the initiator can authenticate the target. This parameter is only required for Mutual CHAP authentication. It can be any sequence of characters and numbers. The minimum length of the secret is 12 characters and the maximum length is 16 characters.

### Target Information Tab Buttons

- **View Default Login Options** – Click this button to see the default login settings.
- **Apply Changes** – Click this button to save and apply your changes.

To configure iSCSI port initiator login:

1. In the discovery-tree (Figure 4-3 on page 40), select the iSCSI initiator node you want to configure.
2. Select the **iSCSI Initiator Login Options** tab (Figure 8-55 on page 198) and make your changes.
3. Click **Apply Changes**.

### Notes


- Any changes to the iSCSI initiator name and alias apply to all ports on the adapter (that is, all iSCSI ports share the iSCSI initiator name and alias).



- On Windows platforms running the Microsoft iSCSI initiator, the initiator iSCSI name is the Microsoft iSCSI IQN. If you change it, the change remains in effect until the system is rebooted. After reboot, the Microsoft IQN is used again as the iSCSI initiator name.

To view default login options:

- In the discovery-tree (Figure 4-3 on page 40), select the iSCSI initiator node whose default login settings you want to view.
- Select the **iSCSI Initiator Login Options** tab (Figure 8-55 on page 198) and click **View Default Login Options**. The Initiator Default Login Options window appears (Figure 8-56 on page 200).



**Figure 8-56** Initiator Default Login Options Window

### Initiator Default Login Options Field Definitions

- InitialR2T** – The initial request to transmit. If set to “Yes”, the initiator has to wait for the target to solicit SCSI data before sending it. If set to “No”, it allows the initiator to send a burst of unsolicited FirstBurstLength bytes.
- Immediate Data** – If set to “Yes”, allows the initiator to append unsolicited data to a command.
- MaxOutstandingR2T** – The maximum number of outstanding request to transmit's (R2T's) per task within a session, each up to MaxBurstLength bytes.
- FirstBurstLength** – The maximum amount of unsolicited data (in bytes) the initiator can send to the target during the execution of a single iSCSI command.
- MaxBurstLength** – The maximum amount of either unsolicited or solicited data the initiator may send in a single burst. Any amount of data exceeding this value must be explicitly solicited by the target.
- DefaultTime2Wait** – The minimum time to wait, in seconds, before the initiator attempts to reconnect or reassign a connection (or task) that has been dropped after an unexpected connection termination or reset. The initiator and target negotiate to determine this value.
- DefaultTime2Retain** – The maximum time, in seconds, to reassign a connection after the initial wait that is indicated in DefaultTime2Wait has elapsed. The initiator and target negotiate to determine this value.
- DataPDUInOrder** – The order of data PDUs within a sequence.
- DataSequenceInOrder** – The order between sequences.



- **HeaderDigest** – The valid values for this property are “CRC32C” or “None”. If set to “CRC32C” and the initiator is configured accordingly, the integrity of an iSCSI PDU's header segments is protected by a CRC32C checksum.
- **DataDigest** – The valid values for this property are “CRC32C” or “None”. If set to CRC32C and the initiator is configured accordingly, the integrity of an iSCSI PDU's data segment is protected by a CRC32C checksum.
- **MaxConnections** – The maximum number of connections to targets that are allowed within a single session.
- **MaxRecvDataSegmentLength** – The maximum data segment length in bytes an initiator or target can receive in an iSCSI PDU.
- **ErrorRecoveryLevel** – The operational ErrorRecoveryLevel for the session. 0 indicates recovery only by session restart. 1 indicates recovery by reissuing commands, data, or status. 2 indicates connection failure recovery.

## 9. Updating Adapter Firmware

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

### Updating Firmware for a Single Adapter

Using the Maintenance or Firmware tab, you can update firmware on local and remote adapters. The firmware file must be downloaded from the Emulex website and extracted to a local drive before you can perform this procedure. (Not available in read-only mode.)

- The Emulex driver must be installed.
- The OneCommand Manager application must be installed.
- The firmware zip file must be downloaded from the Emulex website, unzipped and extracted to a folder on a local drive.
- If the adapter is already connected to a boot device, the system must be 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.

#### Notes


- For vendor-branded adapters, see 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 a Sun-branded adapter.

To update firmware for a single adapter, adapter port, or ASIC:

**Note:** For some older FC adapter models, you update the firmware on each FC port/function. For OneConnect CNAs and 16Gb HBAs, you update the firmware for the entire adapter. For OneConnect dual ASIC adapters, you update the firmware on each ASIC. (For example, dual ASIC 4-port adapters require a firmware download on each ASIC.)

1. Select **Host** or **Fabric** view.
2. In the discovery-tree (Figure 4-3 on page 40), select the adapter, FC port, or ASIC whose firmware you want to update.
3. Select the **Maintenance tab for older FC adapters** (Figure 8-26 on page 152) or the **Firmware tab** (Figure 7-14 on page 84) and click **Download Firmware**. The

Firmware Download dialog box appears (Figure 9-1 on page 203).



**Figure 9-1** Firmware Download Dialog Box

4. Using the Firmware Download dialog box (Figure 9-1 on page 203), navigate to the unzipped, extracted image file you want to download. The firmware image may be specified either by entering the image file's full pathname in the "Firmware File" field or by clicking the **Browse** button.

If you click **Browse**, the Firmware File Selection dialog box appears. Select the file you want to use and click **OK**. The Firmware Download dialog box appears.

5. Click **Start Download**. A warning dialog box appears.
6. Click **Yes to continue**. A status bar shows the progress of the download. The adapter in the discovery-tree (Figure 4-3 on page 40) is displayed in black text when the update is complete.

**Note:** The adapter in the discovery-tree may 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.

If you are updating the firmware on a multi-port FC adapter, repeat steps 1 through 7 to update the firmware on the second port or use the "Updating Firmware for Multiple Adapters" procedure.

## 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 adapters for which that file is compatible. (Not available in read-only mode).

### Notes

- Stop other OneCommand Manager application functions while batch loading is in progress.
- When using the OneCommand Manager application Web Launch Interface the firmware file must reside on the host where the browser window was launched from, not the host that was specified in the web address.

- VMware ESXi hosts managed through the CIM interface lists all the adapters regardless of whether the selected firmware can update the adapter. You must manually deselect the non-matching adapters.

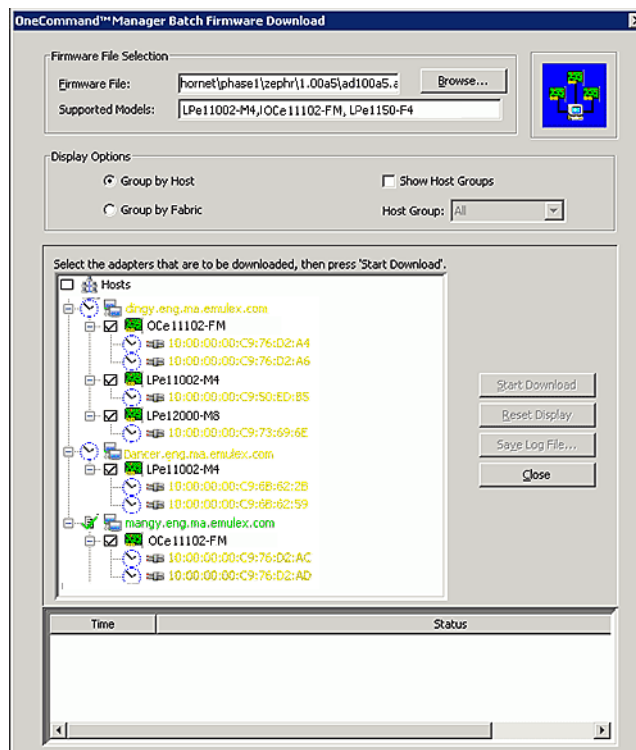
Before you can perform a batch update, the firmware file must be downloaded from the Emulex website and extracted to a directory on your local drive.

To update firmware for multiple adapters:

1. From the **Batch** menu, select **Download Firmware**.

**Note:** You do not need to select a particular discovery-tree element for this operation.

2. When the Batch Firmware Download dialog box appears (Figure 9-2 on page 204), click **Browse**.



**Figure 9-2** Batch Firmware Download Dialog Box, Selecting Adapters to Update

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

After compatible adapters are found, the “Firmware File” text area of the main Batch 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 set of compatible adapters appears in the dialog box's discovery-tree.

Using the Display Options settings you can choose how adapters are displayed in the discovery-tree. Clicking **Group by Host** displays adapters in a host-centric view. Clicking **Group by Fabric** shows hosts in a fabric-centric view with their

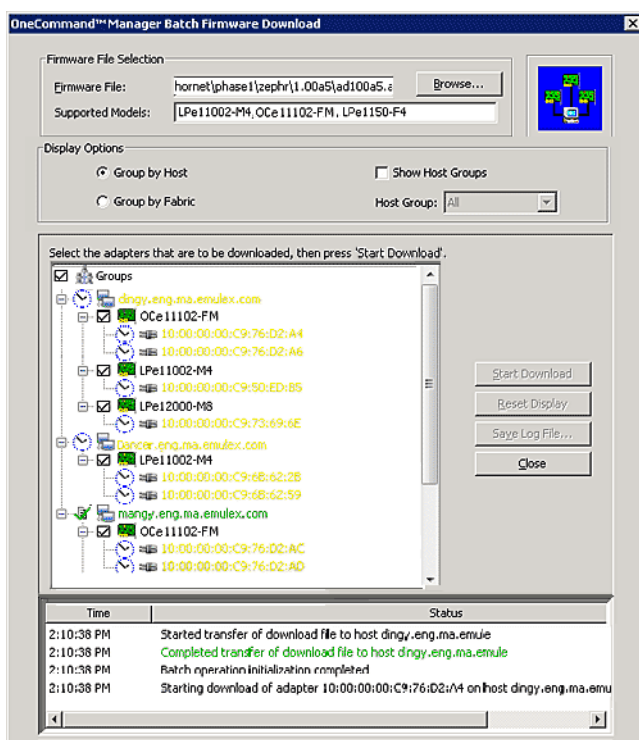
fabric addresses. The WWPN and host name for each downloadable port is displayed under its respective fabric.

You can also display host groups by checking **Show Host Groups**. To display a particular host group, choose that group from the **Host Group** menu.

Checkboxes next to the host, adapter, and ASIC entries are used to select or deselect an entry. Checking an adapter selects or removes that adapter; checking a host removes 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 9-3 on page 205).



**Figure 9-3** Batch Firmware Download Dialog Box, Download Complete

When downloading is finished, you can click **Save Log File** to save copy of the activity log.

## 10. DCB Configuration

The following section describes how to configure DCB parameters on OneConnect adapter ports.

### Configuring DCB Parameters for NIC Only Adapter Ports

The DCB tab (Figure 10-1 on page 206) displays parameters for NIC-only adapter ports. To view the DCB parameters for NIC-only adapter ports:

1. From the discovery-tree, select the NIC only adapter port whose DCB properties you want to view.
2. Select the **DCB** tab (Figure 10-1 on page 206).

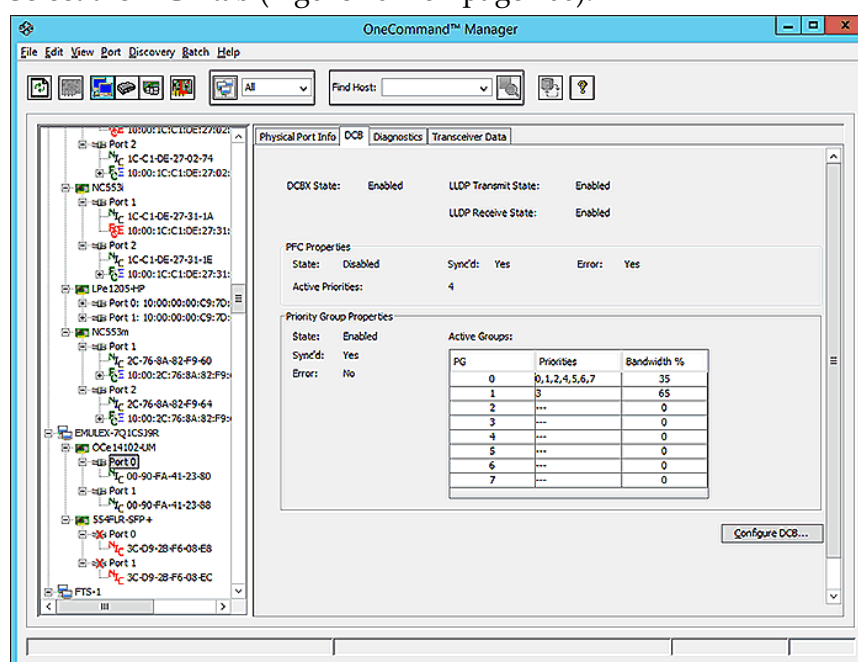


Figure 10-1 DCB Tab for NIC Only Adapter Ports (NIC Only adapter port selected)

### DCB Tab Field Definitions

- **DCBX State** – The current DCBX state (enabled or disabled).
- **LLDP Transmit State** – DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP Rx and Tx must be enabled. If either Rx or Tx is disabled, DCBX is disabled.
- **LLDP Receive States** – DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP Rx and Tx must be enabled. If either Rx or Tx is disabled, DCBX is disabled.

### PFC Properties Area

- State – Enabled means that flow control in both directions (Tx and Rx) is enabled. Disabled means that priority-flow control is currently disabled. The priority value, if shown, is not applicable. This may be caused by:
  - The switch port priority-flow control being set to “On” instead of “Auto”
  - Switch port using port flow control instead of priority flow control
  - PFC disabled at adapter or switch
- Active Priorities – Lists the priorities with PFC set to enabled.
- Sync'd – If yes, the PFC priorities have been set by the peer. This parameter cannot be set.
- Error – The error state. This capability indicates whether an error has occurred during the configuration exchange with the peer or when the compatible method for the capability fails.

### NIC Properties Area

- State – The NIC state. It can be enabled or disabled.
- Active Priority – The current active priority assigned for NIC.
- Sync'd – If yes, the NIC priority has been set by the peer. This parameter cannot be set.
- Error – The NIC error state. This capability indicates whether an error has occurred during the configuration exchange with the peer.

### ETS Priority Group Properties Area

**Note:** Not displayed if multi-channel is enabled on the adapter with the exception of NPar and UFP (if enabled by the switch).

- State – The current priority group state. It can be enabled or disabled.
- Sync'd – If yes, the priority groups have been set by the peer. This parameter cannot be set.
- Error – The error state. This capability indicates whether an error has occurred during the configuration exchange with the peer.

### Active Groups

- PG – The priority group number. It can be 0 to 7.
- Priorities – The priorities that are assigned to each priority group. It is represented in comma separated format.
- Bandwidth % – The percentage of available link bandwidth allocated to a particular priority group.
- Max Configurable PGs – This field indicates maximum number of priority groups that can be configured on the selected adapter port.

## DCB Tab Buttons

- **Configure DCB** – Click to configure DCB parameters. See the instructions below.

To configure DCB for NIC only adapter ports:

1. From the discovery-tree (Figure 4-3 on page 40), select the NIC only adapter port whose DCB properties you want to configure.
2. Select the **DCB** tab (Figure 10-1 on page 206).
3. Click **Configure DCB**. The Configure DCB dialog box appears (Figure 10-2 on page 208).
4. Configure the settings you want and click **OK**.

**Note:** An error message is displayed if you try to configure more priority groups than the adapter supports. The “Max Configurable PGs” field shows the number of priority groups supported by the adapter.

**Configure DCB**

**DCBX Settings**

☒ **Enabled**

DCBX Mode: **DCB**

Operating Version: 0

Maximum Version: 0

**LLDP Settings**

☒ **Transmit Enabled**

☒ **Transmit Port Description Enabled** ☒ **Transmit System Description Enabled**

☒ **Transmit System Name Enabled** ☒ **Transmit System Capabilities Enabled**

☒ **Receive Enabled**

Note: Configured values are only used when not provided by the switch.

**PFC Priorities**

Active Priorities: 3 ☒ **Enable** Configured Priorities: ☐ 0 ☐ 1 ☐ 2 ☒ 3 ☐ 4 ☐ 5 ☐ 6 ☐ 7

**ETS Priority Groups**

Active Groups:

Group ID	Priority Membership	Bandwidth[%]
0	0,1,2,4,5,6,7	50
1	3	50
2	---	0
3	---	0
4	---	0
5	---	0
6	---	0
7	---	0

Configured Groups: Max Configurable PGs: 8

Group ID	0	1	2	3	4	5	6	7	Bandwidth[%]
0	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	50
1	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	50
2	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	0
3	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	0
4	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	0
5	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	0
6	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	0
7	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	0
Total:									100

Buttons: Defaults, Configuration Rules, OK, Cancel

Figure 10-2 Configure DCB Dialog Box for NIC Only Adapter Ports

## Configure DCB Dialog Box Field Definitions

### DCBX Settings Area

- **Enabled** – DCBX can be enabled or disabled. With DCBX enabled, the configured values are used only if the switch does not provide them. With DCBX disabled, the configured values are used. Changes to the DCBX state require a reboot of the host.
- **Operating Version** – The operating version of the DCBX protocol. The system adjusts as needed to operate at the highest version supported by both link partners. This setting cannot be changed.



- **Maximum Version** – The highest DCBX protocol version supported by the system. Version numbers start at zero. The DCBX protocol must be backward compatible with all previous versions. This setting cannot be changed.

### LLDP Settings Area

- **Transmit Enabled** – LLDP Transmit can be enabled or disabled.
- **Transmit Port Description Enabled** – Provides a description of the port in an alpha-numeric format. The value equals the ifDescr object, if the LAN device supports RFC 2863.
- **Transmit System Name Enabled** – Provides the system's assigned name in an alpha-numeric format. The value equals the sysName object, if the LAN device supports RFC 3418.
- **Receive Enabled** – LLDP Receive can be enabled or disabled.
- **Transmit System Description Enabled** – Provides a description of the network entity in an alpha-numeric format. This includes the system's name and versions of hardware, operating system, and networking software supported by the device. The value equals the sysDescr object, if the LAN device supports RFC 3418.
- **Transmit System Capabilities Enabled** – Indicates the primary function(s) of the device and whether or not these functions are enabled on the device. The capabilities are indicated by two octets. Bits 0 through 7 indicate Other, Repeater, Bridge, WLAN AP, Router, Telephone, DOCSIS cable device, and Station respectively. Bits 8 through 15 are reserved.

### PFC Priorities Area

- **Active Priorities** – The priorities that are marked active for PFC.
- **Enable** – If checked, PFC is enabled.
- **Configured Priorities** – The priorities that are configured, but might not yet be active. A maximum of two PFC priority check boxes can be selected, out of which one of them must match the iSCSI priority. The additional PFC priority would be for the Ethernet traffic. This additional PFC priority must be assigned to a priority group which has no other priorities.

### ETS Priority Groups Area

**Note:** Not shown if any multi-channel type is enabled on the adapter.

- **Active Groups**
  - **Group ID** – The priority group ID.
  - **Priority Membership** – The different priorities that are assigned to the various priority groups. This is the currently active configuration.
  - **Bandwidth %** – The bandwidths that are assigned to different priority groups. This is the currently active configuration.
- **Configured Groups**
  - **Group ID** – The priority group ID.

- Priority Membership – The configured priority membership grouping.
- Bandwidth % – The configured value of bandwidth for the different priority groups.
- Max Configurable PGs – The maximum number of priority groups that can be configured.

**Note:** Configured Groups are only made active when the switch settings are not being used.

### Configure DCB Dialog Box Buttons

- Defaults – Click to return to the factory settings.
- Configuration Rules – Click to display the NIC Priority window that lists the rules for configuring NIC priorities.

You must observe the following rules when configuring priority groups for NIC-Only adapter ports:

1. Only one PFC priority can be configured.
  2. The PFC Priority must be assigned to a priority group that has no other priorities.
  3. Bandwidths of all the priority groups must add up to 100%.
- OK – Click to apply and save your changes.
  - Cancel – Click to discard any changes you made.

## Configuring DCB Parameters for NIC + iSCSI Adapter Ports

The DCB tab (Figure 10-3 on page 211) displays parameters for NIC + iSCSI adapter ports.

**Note:** For ports running both iSCSI and FCoE, see “Configuring DCB Parameters for NIC + FCoE + iSCSI Adapter Ports” on page 220.

To view the DCB parameters for NIC + iSCSI adapter ports:

1. From the discovery-tree, select the NIC + iSCSI adapter port whose DCB properties you want to view.
2. Select the **DCB** tab (Figure 10-3 on page 211).

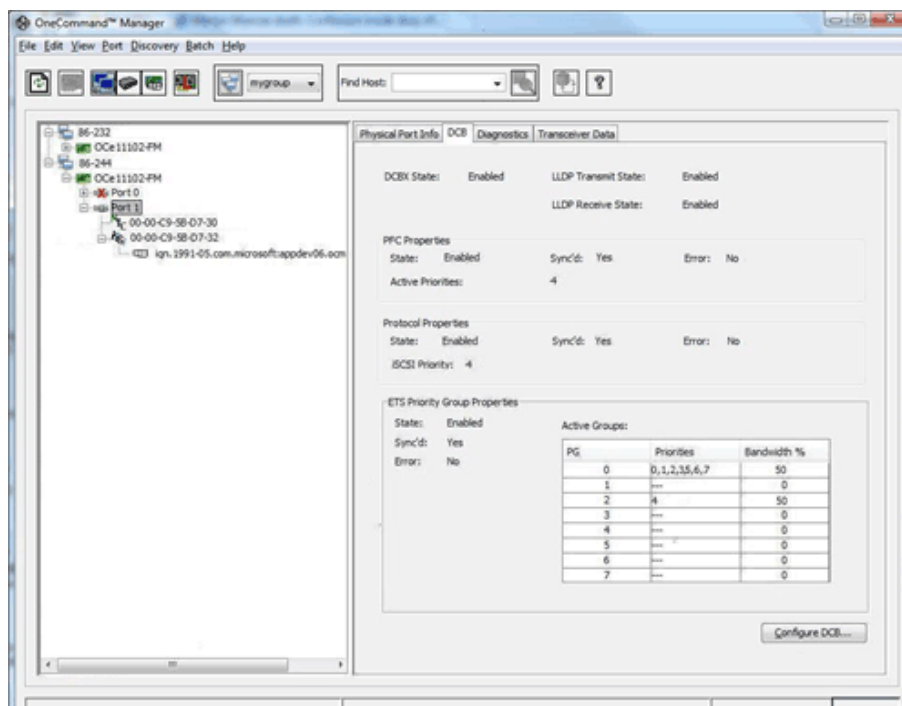


Figure 10-3 DCB Tab for NIC + iSCSI Adapter Ports (NIC + iSCSI adapter port selected)

## DCB Tab Field Definitions

- DCBX State – The current DCBX state (enabled or disabled).
- LLDP Transmit State – DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP Rx and Tx must be enabled. If either Rx or Tx is disabled, DCBX is disabled.
- LLDP Receive States – DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP Rx and Tx must be enabled. If either Rx or Tx is disabled, DCBX is disabled.

## PFC Properties Area

**Note:** PFC is not supported on all the iSCSI adapter ports.

- State – Enabled means that flow control in both directions (Tx and Rx) is enabled. Disabled means that priority-flow control is currently disabled. The priority value, if shown, is not applicable. This may be caused by:
  - The switch port priority-flow control being set to “On” instead of “Auto”
  - Switch port using port flow control instead of priority flow control
  - PFC disabled at adapter or switch
- Active Priority – Lists the priorities with PFC set to enabled.
- Sync'd – If yes, the PFC priorities have been set by the peer. This parameter cannot be set.

- **Error** – The error state. This capability indicates whether an error has occurred during the configuration exchange with the peer or when the compatible method for the capability fails.

### iSCSI Properties Area

- **State** – The iSCSI state. It can be enabled or disabled.
- **Active Priority** – The current active priority assigned for iSCSI.
- **Sync'd** – If yes, the iSCSI priority has been set by the peer. This parameter cannot be set.
- **Error** – The iSCSI error state. This capability indicates whether an error has occurred during the configuration exchange with the peer.

### ETS Priority Group Properties Area

**Note:** Not displayed if multi-channel is enabled on the adapter with the exception of NPar and UFP (if enabled by the switch).

- **State** – The current priority group state. It can be enabled or disabled.
- **Sync'd** – If yes, the priority groups have been set by the peer. This parameter cannot be set.
- **Error** – The error state. This capability indicates whether an error has occurred during the configuration exchange with the peer.

### Active Groups

- **PG** – The priority group number. It can be 0 to 7.
- **Priorities** – The priorities that are assigned to each priority group. It is represented in comma separated format.
- **Bandwidth %** – The percentage of available link bandwidth allocated to a particular priority group.

### DCB Tab Buttons

- **Configure DCB** – Click to configure DCB parameters. See the following instructions.

To configure DCB for NIC + iSCSI adapter ports:

1. From the discovery-tree (Figure 4-3 on page 40), select the NIC + iSCSI adapter port whose DCB properties you want to configure.
2. Select the **DCB** tab (Figure 10-3 on page 211).
3. Click **Configure DCB**. The Configure DCB dialog box appears (Figure 10-4 on page 213).
4. Configure the settings you want and click **OK**.

**Note:** An error message is displayed if you try to configure more priority groups than the adapter supports. The “Max Configurable PGs” field shows the number of priority groups supported by the adapter.

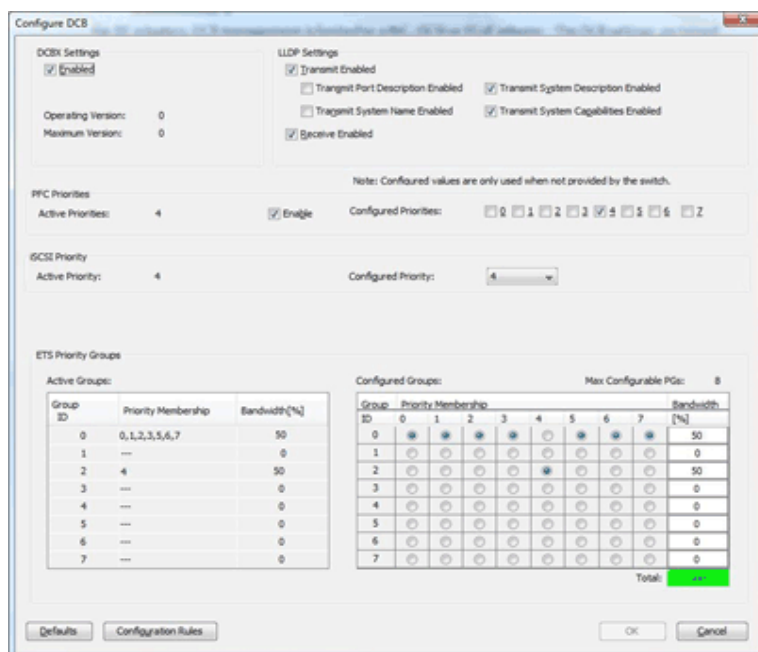


Figure 10-4 Configure DCB Dialog Box for NIC + iSCSI Adapter Ports (DCBX enabled)

## Configure DCB Dialog Box Field Definitions

### DCBX Settings Area

- **Enabled** – DCBX can be enabled or disabled. With DCBX enabled, the configured values are used only if the switch does not provide them. With DCBX disabled, the configured values are used. Changes to the DCBX state require a reboot of the host.
- **Operating Version** – The operating version of the DCBX protocol. The system adjusts as needed to operate at the highest version supported by both link partners. This setting cannot be changed.
- **Maximum Version** – The highest DCBX protocol version supported by the system. Version numbers start at zero. The DCBX protocol must be backward compatible with all previous versions. This setting cannot be changed.

### LLDP Settings Area

- **Transmit Enabled** – LLDP Transmit can be enabled or disabled.
- **Transmit Port Description Enabled** – Provides a description of the port in an alpha-numeric format. The value equals the ifDescr object, if the LAN device supports RFC 2863.
- **Transmit System Name Enabled** – Provides the system's assigned name in an alpha-numeric format. The value equals the sysName object, if the LAN device supports RFC 3418.
- **Receive Enabled** – LLDP Receive can be enabled or disabled.

- **Transmit System Description Enabled** – Provides a description of the network entity in an alpha-numeric format. This includes the system's name and versions of hardware, operating system, and networking software supported by the device. The value equals the sysDescr object, if the LAN device supports RFC 3418.
- **Transmit System Capabilities Enabled** – Indicates the primary function(s) of the device and whether or not these functions are enabled on the device. The capabilities are indicated by two octets. Bits 0 through 7 indicate Other, Repeater, Bridge, WLAN AP, Router, Telephone, DOCSIS cable device, and Station respectively. Bits 8 through 15 are reserved.

### **PFC Priorities Area**

**Note:** PFC is not supported on all the iSCSI adapter ports.

- **Active Priorities** – The priorities that are marked active for PFC.
- **Enable** – If checked, PFC is enabled.
- **Configured Priorities** – The priorities that are configured, but might not yet be active. A maximum of two PFC priority check boxes can be selected, out of which one of them must match the iSCSI priority. The additional PFC priority would be for the Ethernet traffic. This additional PFC priority must be assigned to a priority group which has no other priorities.

### **iSCSI Priority Area**

- **Active Priority** – The active iSCSI priority.
- **Configured Priority** – The configured iSCSI priority.

### **ETS Priority Groups Area**

**Note:** Not displayed if any type of multi-channel is enabled on the adapter.

- **Active Groups**
  - **Group ID** – The priority group ID.
  - **Priority Membership** – The different priorities that are assigned to the various priority groups. This is the currently active configuration.
  - **Bandwidth %** – The bandwidths that are assigned to different priority groups. This is the currently active configuration.
- **Configured Groups**
  - **Group ID** – The priority group ID.
  - **Priority Membership** – The configured priority membership grouping.
  - **Bandwidth %** – The configured value of bandwidth for the different priority groups.
  - **Max Configurable PGs** – The maximum number of priority groups that can be configured.

## Configure DCB Dialog Box Buttons

- Defaults – Click to return parameters to default iSCSI DCB settings.
- Configuration Rules – Click to display the iSCSI Priority window that lists the rules for configuring iSCSI priorities.

You must observe the following rules when configuring priority groups for iSCSI adapter ports:

1. Only one priority can be configured as the iSCSI priority.
2. A maximum of two PFC priorities can be selected and one of them must match the iSCSI priority.

**Note:** Not all adapters support two PFC priorities. Adapters that do not support two PFC priorities display an error message if you try to configure more than one PFC priority.

3. The priority group to which the iSCSI priority is assigned must contain no other priorities.
  4. The additional PFC priority must be assigned to a priority group that has no other priorities.
  5. Bandwidths of all the priority groups must add up to 100%.
- OK – Click to apply and save your changes.
  - Cancel – Click to discard any changes you made.

## Configuring DCB Parameters for NIC + FCoE Adapter Ports

**Note:** For ports running both iSCSI and FCoE, see “Configuring DCB Parameters for NIC + FCoE + iSCSI Adapter Ports” on page 220.

The DCB tab displays parameters for NIC + FCoE adapter ports.

To view the DCB parameters for NIC + FCoE adapter ports:

1. From the discovery-tree, select the NIC + FCoE adapter port whose DCB properties you want to view.
2. Select the **DCB** tab (Figure 10-5 on page 216).



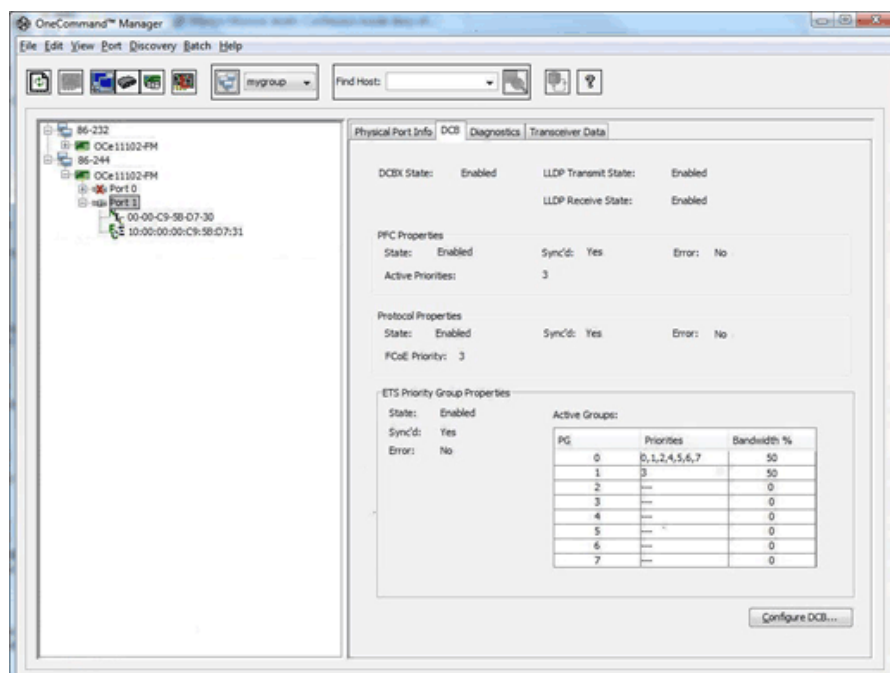


Figure 10-5 DCB Tab (NIC + FCoE adapter port selected)

## DCB Tab Field Definitions

- **DCBX State** – The current DCBX state (enabled or disabled).
- **LLDP Transmit State** – DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP Rx and Tx must be enabled. If either Rx or Tx is disabled, DCBX is disabled.
- **LLDP Receive State** – DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP Rx and Tx must be enabled. If either Rx or Tx is disabled, DCBX is disabled.

## PFC Properties Area

- **State** – Enabled means that flow control in both directions (Tx and Rx) is enabled. Disabled means that priority-flow control is currently disabled. The priority value, if shown, is not applicable. This may be caused by:
  - The switch port priority-flow control being set to “On” instead of “Auto”
  - Switch port using port flow control instead of priority flow control
  - PFC disabled at adapter or switch
- **Active Priorities** – Lists the priorities with PFC set to enabled.
- **Sync'd** – If yes, the PFC priorities have been set by the peer. This parameter cannot be set.
- **Error** – The error state. This capability indicates whether an error has occurred during the configuration exchange with the peer. Error is also set to YES when the Compatible method for the capability fails.



### FCoE Properties Area (FCoE Ports Only)

- State – The FCoE state. It can be enabled or disabled.
- Active Priority – The current active priority assigned for FCoE.
- Sync'd – If yes, the FCoE priority has been set by the peer. This parameter cannot be set.
- Error – The FCoE error state. This capability indicates whether an error has occurred during the configuration exchange with the peer. Error is also set to YES when the Compatible method for the capability fails.

### ETS Priority Group Properties Area

**Note:** Not displayed if multi-channel is enabled on the adapter with the exception of NPar and UFP (if enabled by the switch).

- State – The priority group state. It can be enabled or disabled.
- Sync'd – If yes, the priority groups have been set by the peer. This parameter cannot be set.
- Error – The error state. This capability indicates whether an error has occurred during the configuration exchange with the peer. Error is also set to YES when the Compatible method for the capability fails.

### Active Groups

- PG – The priority group number. It can be 0 to 7.
- Priorities – The priorities that are assigned to each priority group. It is represented in comma separated format.
- Bandwidth % – The percentage of available link bandwidth allocated to a particular priority group.

### DCB Tab Buttons

- Configure DCB – Click to configure DCB parameters. See the instructions below.

To configure DCB for NIC + FCoE adapter ports:

1. From the discovery-tree (Figure 4-3 on page 40), select the NIC + FCoE adapter port whose DCB properties you want to configure.
2. Select the **DCB** tab (Figure 10-5 on page 216).
3. Click **Configure DCB**. The Configure DCB dialog box appears (Figure 10-6 on page 218).
4. Configure the settings you want and click **OK**.

**Note:** An error message is displayed if you try to configure more priority groups than the adapter supports. The “Max Configurable PGs” field in the ETS priority groups area shows the number of priority groups supported by the adapter.

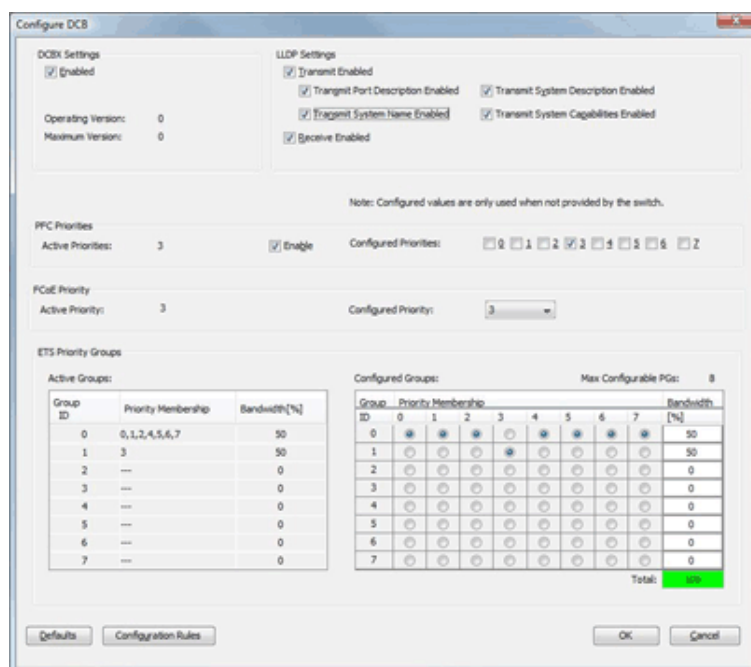


Figure 10-6 Configure DCB Dialog Box for NIC + FCoE Adapter Ports (DCBX enabled)

## Configure DCB Dialog Box Field Definitions

### DCBX Settings Area

- **Enabled** – DCBX can be enabled or disabled. With DCBX enabled, the configured values are used only if the switch does not provide them. With DCBX disabled, the configured values are used.
- **Operating Version** – Operating version of the DCBX protocol. The system adjusts as needed to operate at the highest version supported by both link partners. This setting cannot be changed.
- **Maximum Version** – The highest DCBX protocol version supported by the system. Version numbers start at zero. The DCBX protocol must be backward compatible with all previous versions. This setting cannot be changed.

### LLDP Settings Area

- **Transmit Enabled** – LLDP Transmit can be enabled or disabled.
- **Transmit Port Description Enabled** – Provides a description of the port in an alpha-numeric format. The value equals the ifDescr object, if the LAN device supports RFC 2863.
- **Transmit System Name Enabled** – Provides the system's assigned name in an alpha-numeric format. The value equals the sysName object, if the LAN device supports RFC 3418.
- **Transmit System Description Enabled** – Provides a description of the network entity in an alpha-numeric format. This includes system's name and versions of

hardware, operating system, and networking software supported by the device. The value equals the sysDescr object, if the LAN device supports RFC 3418.

- Transmit System Capabilities Enabled – Indicates the primary function(s) of the device and whether or not these functions are enabled on the device. The capabilities are indicated by two octets. Bits 0 through 7 indicate Other, Repeater, Bridge, WLAN AP, Router, Telephone, DOCSIS cable device, and Station respectively. Bits 8 through 15 are reserved.
- Receive Enabled – LLDP Receive can be enabled or disabled.

### **PFC Priorities Area**

- Active Priorities – The priorities that are marked active for PFC.
- Enable – If checked, PFC is enabled.
- Configured Priorities – The priorities that are configured, but might not yet be active.

### **FCoE Priority Area (FCoE Ports Only)**

- Active Priority – The active FCoE priority.
- Configured Priority – The configured FCoE priority.

### **ETS Priority Groups Area**

**Note:** Not displayed if any type of multi-channel is enabled on the adapter.

- Active Groups
  - Group ID – The priority group ID.
  - Priority Membership – The different priorities that are assigned to the various priority groups. This is the currently active configuration.
  - Bandwidth – The bandwidths that are assigned to different priority groups. This is the currently active configuration.
- Configured Groups
  - Group ID – The priority group ID.
  - Priority Membership – The configured priority membership grouping.
  - Bandwidth % – The configured value of bandwidth for the different priority groups.
  - Max Configurable PGs – The maximum number of priority groups that can be configured.

### **Configure DCB Dialog Box Buttons**

- Defaults – Click to return parameters to default FCoE DCB settings.
- Configuration Rules – Click to display the window that lists the rules for configuring FCoE priority group information.

You must observe the following rules when configuring priority groups for FCoE adapter ports:

1. One and only one priority is configured for the FCoE priority.

2. A maximum of two PFC priorities can be selected and one of them must match the FCoE priority.

**Note:** Not all adapters support two PFC priorities. Adapters that do not support two PFC priorities display an error message if you try to configure more than one PFC priority.

3. The priority group to which the FCoE priority is assigned must contain no other priorities.
  4. The additional PFC priority must be assigned to a priority group which has no other priorities.
  5. Bandwidths of all the priority groups must add up to 100%.
- OK – Click to apply and save your changes.
  - Cancel – Click to discard any changes you made.

## Configuring DCB Parameters for NIC + FCoE + iSCSI Adapter Ports

The DCB tab displays parameters for NIC + FCoE + iSCSI adapter ports.

**Note:** For adapter ports running only FCoE or only iSCSI, refer to “Configuring DCB Parameters for NIC + FCoE Adapter Ports” on page 215 and to “Configuring DCB Parameters for NIC + iSCSI Adapter Ports” on page 210.

To view the DCB parameters for NIC + FCoE + iSCSI adapter ports:

1. From the discovery-tree (Figure 4-3 on page 40), select the NIC + FCoE + iSCSI adapter port whose DCB properties you want to view.
2. Select the **DCB** tab (Figure 10-7 on page 221).

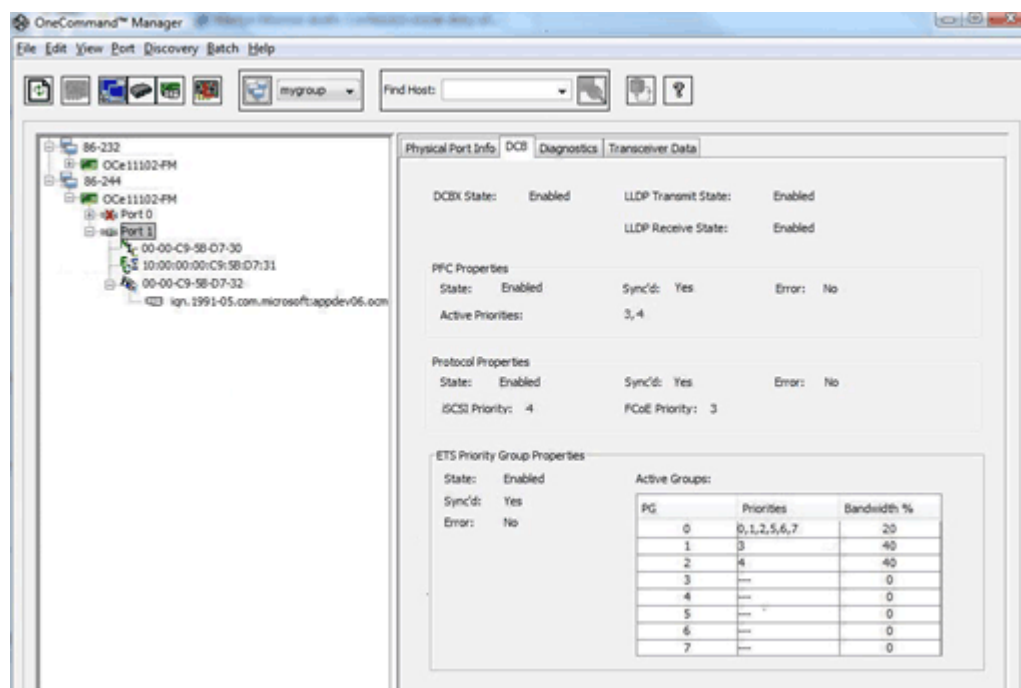


Figure 10-7 DCB Tab (NIC + FCoE + iSCSI adapter port selected)

## DCB Tab Field Definitions

- **DCBX State** – The current DCBX state (enabled or disabled).
- **LLDP Transmit State** – DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP Rx and Tx must be enabled. If either Rx or Tx is disabled, DCBX is disabled.
- **LLDP Receive State** – DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP Rx and Tx must be enabled. If either Rx or Tx is disabled, DCBX is disabled.

## PFC Properties Area

- **State** – Enabled means that flow control in both directions (Tx and Rx) is enabled. Disabled means that priority-flow control is currently disabled. The priority value, if Shown, is not applicable. This may be caused by:
  - The switch port priority-flow control being set to “On” instead of “Auto”
  - Switch port using port flow control instead of priority flow control
  - PFC disabled at adapter or switch
- **Active Priorities** – Lists the priorities with PFC set to enabled.
- **Sync’d** – If yes, the PFC priorities have been set by the peer. This parameter cannot be set.
- **Error** – The error state. This capability indicates whether an error has occurred during the configuration exchange with the peer. Error is also set to YES when the Compatible method for the capability fails.

## Properties Area

- State – The iSCSI or FCoE state. It can be enabled or disabled.
- Active Priority – The current active priority assigned for iSCSI or FCoE.
- Sync'd – If yes, the iSCSI or FCoE priority has been set by the peer. This parameter cannot be set.
- Error – The iSCSI or FCoE error state. This capability indicates whether an error has occurred during the configuration exchange with the peer. Error is also set to YES when the Compatible method for the capability fails.

## ETS Priority Group Properties Area

**Note:** Not displayed if multi-channel is enabled on the adapter with the exception of NPar and UFP (if enabled by the switch).

- State – The current priority group state. It can be enabled or disabled.
- Sync'd – If yes, the priority groups have been set by the peer. This parameter cannot be set.

**Note:** If NPar is enabled and Sync'd is set “Yes”, the protocols are limited to the bandwidths shown in the priority groups instead of the bandwidths configured in the Adapter Configuration tab.

- Error – The error state. This capability indicates whether an error has occurred during the configuration exchange with the peer. Error is also set to “Yes” when the Compatible method for the capability fails.

## Active Groups

- PG – The priority group number. It can be 0 to 7.
- Priorities – The priorities that are assigned to each priority group. It is represented in comma separated format.
- Bandwidth % – The percentage of available link bandwidth allocated to a particular priority group.

## DCB Tab Buttons

- Configure DCB – Click to configure DCB parameters. See the instructions below.

To configure DCB for NIC + FCoE + iSCSI adapter ports:

1. From the discovery-tree (Figure 4-3 on page 40), select the NIC + FCoE + iSCSI adapter port whose DCB properties you want to configure.
2. Select the **DCB** tab (Figure 10-7 on page 221).
3. Click **Configure DCB**. The Configure DCB dialog box appears (Figure 10-8 on page 223).
4. Configure the settings you want and click **OK**.

**Note:** An error message is displayed if you try to configure more priority groups than the adapter supports. The “Max Configurable PGs” field in the ETS

priority groups area shows the number of priority groups supported by the adapter.

**Configure DCB**

**DCBX Settings**  
☒ Enabled  
 Operating Version: 0  
 Maximum Version: 0

**LLDP Settings**  
☒ Transmit Enabled  
☒ Transmit Port Description Enabled  
☒ Transmit System Name Enabled  
☒ Transmit System Capabilities Enabled  
☒ Receive Enabled

Note: Configured values are only used when not provided by the switch.

**PFC Priorities**  
 Active Priorities: 3 ☒ Enable  
 Configured Priorities: ☐ 0 ☐ 1 ☐ 2 ☒ 3 ☐ 4 ☐ 5 ☐ 6 ☐ 7

**FCoE Priority**  
 Active Priority: 3  
 Configured Priority: 3

**iSCSI Priority**  
 Active Priority: 4  
 Configured Priority: 4

**ETS Priority Groups**

Active Groups:

Group ID	Priority Membership	Bandwidth(%)
0	0,1,2,3,6,7	20
1	3	40
2	4	40
3	---	0
4	---	0
5	---	0
6	---	0
7	---	0

Configured Groups: Max Configurable PGs: 8

Group ID	Priority Membership	Bandwidth(%)
0	<input checked="" type="checkbox"/> 0 <input checked="" type="checkbox"/> 1 <input checked="" type="checkbox"/> 2 <input checked="" type="checkbox"/> 3 <input checked="" type="checkbox"/> 4 <input checked="" type="checkbox"/> 5 <input checked="" type="checkbox"/> 6 <input checked="" type="checkbox"/> 7	20
1	<input type="checkbox"/> 0 <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input checked="" type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7	40
2	<input type="checkbox"/> 0 <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input checked="" type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7	40
3	<input type="checkbox"/> 0 <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7	0
4	<input type="checkbox"/> 0 <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7	0
5	<input type="checkbox"/> 0 <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7	0
6	<input type="checkbox"/> 0 <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7	0
7	<input type="checkbox"/> 0 <input type="checkbox"/> 1 <input type="checkbox"/> 2 <input type="checkbox"/> 3 <input type="checkbox"/> 4 <input type="checkbox"/> 5 <input type="checkbox"/> 6 <input type="checkbox"/> 7	0
Total:		100%

Defaults Configuration Rules OK Cancel

Figure 10-8 Configure DCB Dialog Box for NIC + FCoE + iSCSI Adapter Ports (DCBX enabled)

## Configure DCB Dialog Box Field Definitions

### DCBX Settings Area

- Enabled – DCBX can be enabled or disabled. With DCBX enabled, the configured values are used only if the switch does not provide them. With DCBX disabled, the configured values are used.
- Operating Version – Operating version of the DCBX protocol. The system adjusts as needed to operate at the highest version supported by both link partners. This setting cannot be changed.
- Maximum Version – The highest DCBX protocol version supported by the system. Version numbers start at zero. The DCBX protocol must be backward compatible with all previous versions. This setting cannot be changed.

### LLDP Settings Area

- Transmit Enabled – LLDP Transmit can be enabled or disabled.
- Transmit Port Description Enabled – Provides a description of the port in an alpha-numeric format. The value equals the ifDescr object, if the LAN device supports RFC 2863.
- Transmit System Name Enabled – Provides the system's assigned name in an alpha-numeric format. The value equals the sysName object, if the LAN device supports RFC 3418.



- **Transmit System Description Enabled** – Provides a description of the network entity in an alpha-numeric format. This includes system's name and versions of hardware, operating system, and networking software supported by the device. The value equals the sysDescr object, if the LAN device supports RFC 3418.
- **Transmit System Capabilities Enabled** – Indicates the primary function(s) of the device and whether or not these functions are enabled on the device. The capabilities are indicated by two octets. Bits 0 through 7 indicate Other, Repeater, Bridge, WLAN AP, Router, Telephone, DOCSIS cable device, and Station respectively. Bits 8 through 15 are reserved.
- **Receive Enabled** – LLDP Receive can be enabled or disabled.

### **PFC Priorities Area**

- **Active Priorities** – The priorities that are marked active for PFC.
- **Enable** – If checked, PFC is enabled.
- **Configured Priorities** – The priorities that are configured, but might not yet be active. Since only one PFC priority is supported, only the FCoE or iSCSI priority should be checked, not both.

### **FCoE Priority Area**

- **Active Priority** – The active FCoE priority.
- **Configured Priority** – The configured FCoE priority.

### **iSCSI Priority Area**

- **Active Priority** – The active iSCSI priority.
- **Configured Priority** – The configured iSCSI priority.

### **ETS Priority Groups Area**

**Note:** Not shown if any type of multi-channel is enabled on the adapter.

- **Active Groups**
  - **Group ID** – The priority group ID.
  - **Priority Membership** – The different priorities that are assigned to the various priority groups. This is the currently active configuration.
  - **Bandwidth** – The bandwidths that are assigned to different priority groups. This is the currently active configuration.
- **Configured Groups**
  - **Group ID** – The priority group ID.
  - **Priority Membership** – The configured priority membership grouping.
  - **Bandwidth %** – The configured value of bandwidth for the different priority groups.
  - **Max Configurable PGs** – The maximum number of priority groups that can be configured.



## Configure DCB Dialog Box Buttons

- Defaults – Click to return parameters to default FCoE / iSCSI DCB settings.
- Configuration Rules – Click to display the window that lists the rules for configuring FCoE priority group information.

You must observe the following rules when configuring priority groups for FCoE adapter ports:

1. One and only one priority is configured for the FCoE priority and one and only one priority can be the iSCSI priority and it cannot match the FCoE priority.
  2. A maximum of two PFC priorities can be selected and one of them must match the FCoE priority and the other must match the iSCSI priority.
  3. The priority group to which the FCoE priority is assigned must contain no other priorities.
  4. The priority group to which the iSCSI priority is assigned must contain no other priorities.
  5. The additional PFC priority must be assigned to a priority group which has no other priorities.
  6. Do not exceed the maximum number of configurable priority groups displayed above the priority groups box when assigning priorities to priority groups.
  7. Bandwidths of all the priority groups must add up to 100%.
- OK – Click to apply and save your changes.
  - Cancel – Click to discard any changes you made.

## Configuring DCB Parameters for RoCE Adapter Ports

The DCB tab (Figure 10-9 on page 226) displays parameters for RoCE adapter ports.

**Note:** RoCE configurations are not supported with SR-IOV.

To view the DCB parameters for RoCE adapter ports:

1. From the discovery-tree (Figure 4-3 on page 40), select the RoCE adapter port whose DCB properties you want to view.

2. Select the **DCB** tab (Figure 10-9 on page 226).

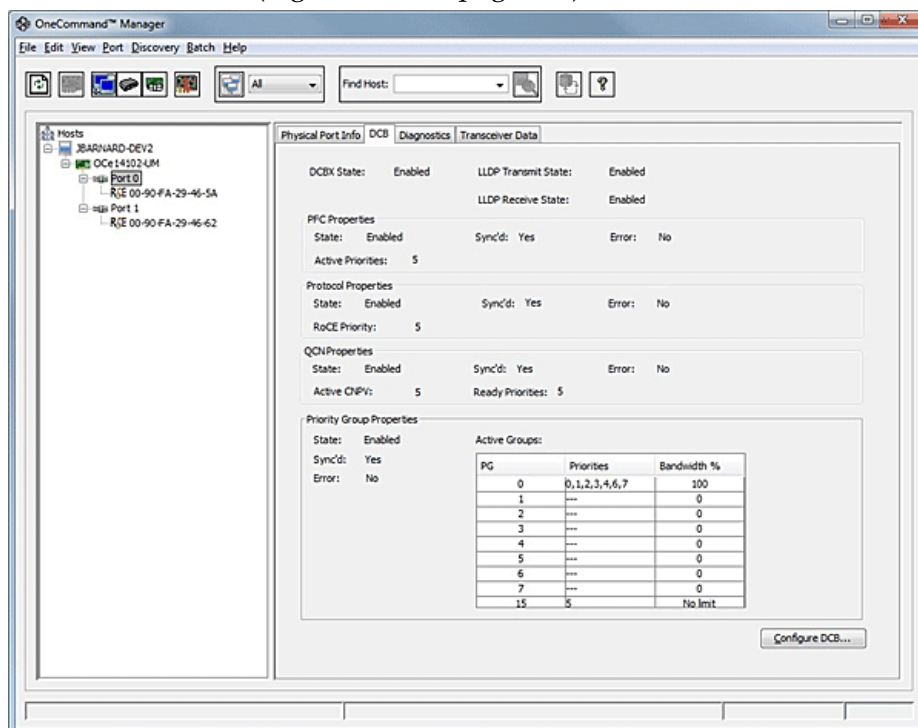


Figure 10-9 DCB Tab for RoCE Adapter Ports

## DCB Tab Field Definitions

- **DCBX State** – The current DCBX state (enabled or disabled).
- **LLDP Transmit State** – DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP Rx and Tx must be enabled. If either Rx or Tx is disabled, DCBX is disabled.
- **LLDP Receive States** – DCBX uses LLDP to exchange parameters between two link peers. For the DCBX protocol to operate correctly, both LLDP Rx and Tx must be enabled. If either Rx or Tx is disabled, DCBX is disabled.

## PFC Properties Area

- **State** – Enabled means that flow control in both directions (Tx and Rx) is enabled. Disabled means that priority-flow control is currently disabled. The priority value, if Shown, is not applicable. This may be caused by:
  - The switch port priority-flow control being set to “On” instead of “Auto”
  - Switch port using port flow control instead of priority flow control
  - PFC disabled at adapter or switch
- **Active Priorities** – Lists the priorities with PFC set to enabled.
- **Sync’d** – If “yes”, the PFC priorities have been set by the peer. This parameter cannot be set.

- **Error** – The error state. This capability indicates whether an error has occurred during the configuration exchange with the peer or when the compatible method for the capability fails.

### Protocol Properties Area

- **State** – The NIC state. It can be enabled or disabled.
- **RoCE Priority** – The current active priority assigned for RoCE.
- **Sync'd** – If yes, the RoCE priority has been set by the peer. This parameter cannot be set.
- **Error** – The RoCE error state. This capability indicates whether an error has occurred during the configuration exchange with the peer.

**Note:** In some cases, the switch may return a protocol priority other than RoCE. The non-RoCE protocol's priority value is also displayed to indicate that the switch may be mis-configured.

### QCN Properties Area

QCN is a congestion notification protocol for virtual bridged LAN networks as defined in the IEEE 802.1Qau standard. QCN spans the layer 2 subnet, and is an end-to-end congestion notification standard. In order to leverage the benefits of QCN, every network entity in the layer 2 subnet must be QCN aware.

**Note:** QCN is available only for the OCe14000-series adapters running NIC+RoCE protocol on at least one port and is used for RoCE traffic only.

- **State** – The current QCN state (enabled or disabled)
- **Active CNPV** – The current Active CNPV being used for QCN.
- **Sync'd** – If yes, the CNPV priority has been set by the peer. This parameter cannot be set. It is set by the firmware based upon whether or not is synchronizes with the peer or switch for the setting.
- **Ready Priorities** – The Ready Priorities field shows a comma separated list of the CND priorities advertised for QCN. It is a read-only field and cannot be modified in Configure DCB dialog box.
- **Error** – The QCN error state. This capability indicates whether an error has occurred during the configuration exchange with the peer.

### ETS Priority Group Properties Area

- **State** – The current priority group state. It can be enabled or disabled.
- **Sync'd** – If yes, the priority groups have been set by the peer. This parameter cannot be set.
- **Error** – The error state. This capability indicates whether an error has occurred during the configuration exchange with the peer.

### Active Groups

- **PG** – The priority group number. It can be 0-7, and 15.

- **Priorities** – The priorities that are assigned to each priority group. It is represented in comma separated format.
- **Bandwidth %** – The percentage of available link bandwidth allocated to a particular priority group.

**Note:** PG15 is only available for NIC+RoCE and has no bandwidth limitation.

## DCB Tab Buttons

- **Configure DCB** – Click to configure DCB parameters. See the instructions below.

To configure DCB for RoCE adapter ports:

1. From the discovery-tree (Figure 4-3 on page 40), select the RoCE adapter port whose DCB properties you want to configure.
2. Select the **DCB** tab (Figure 10-9 on page 226).
3. Click **Configure DCB**. The Configure DCB dialog box appears (Figure 10-10 on page 228).
4. Configure the settings you want and click **OK**.

**Note:** An error message is displayed if you try to configure more priority groups than the adapter supports. The “Max Configurable PGs” field shows the number of priority groups supported by the adapter.

**Configure DCB**

**DCBX Settings**

☒ Enabled

Operating Version: 0  
Maximum Version: 0

**LLDP Settings**

☒ Transmit Enabled  
☐ Transmit Port Description Enabled  
☒ Transmit System Description Enabled  
☐ Transmit System Name Enabled  
☒ Transmit System Capabilities Enabled  
☒ Receive Enabled

**PFC Priorities**

Note: Configured values are only used when not provided by the switch.

Active Priorities: 5 ☒ Enable  
Configured Priorities: ☐ 0 ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☒ 5 ☐ 6 ☐ 7

**RoCE Priority**

Active Priority: 5  
Configured Priority: 5

**QCN**

Active CNPV: 5 ☒ Enable  
Configured CNPV: 5 (same as RoCE Priority)

**ETS Priority Groups**

Active Groups:

Group ID	Priority Membership	Bandwidth[%]
0	0, 1, 2, 3, 4, 6, 7	100
1	---	0
2	---	0
3	---	0
4	---	0
5	---	0
6	---	0
7	---	0
15	5	No limit

Configured Groups: Maximum Configurable PGs: 3

Group ID	0	1	2	3	4	5	6	7	Bandwidth [%]
0	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="radio"/>	100
1	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	0
2	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	0
3	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	0
4	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	0
5	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	0
6	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	0
7	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	0
15	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="radio"/>	<input type="radio"/>	No limit

Total: 100

Defaults Configuration Rules OK Cancel

**Figure 10-10** Configure DCB Dialog Box for RoCE Adapter Ports (DCBX enabled)

## Configure DCB Dialog Box Field Definitions

### DCBX Settings Area

- **Enabled** – DCBX can be enabled or disabled. With DCBX enabled, the configured values are not used. The switch settings are used. With DCBX disabled, the configured values are used. Changes to the DCBX state require a reboot of the host.
- **Operating Version** – The operating version of the DCBX protocol. The system adjusts as needed to operate at the highest version supported by both link partners. This setting cannot be changed.
- **Maximum Version** – The highest DCBX protocol version supported by the system. Version numbers start at zero. The DCBX protocol must be backward compatible with all previous versions. This setting cannot be changed.

### LLDP Settings Area

- **Transmit Enabled** – LLDP Transmit can be enabled or disabled.
- **Transmit Port Description Enabled** – Provides a description of the port in an alpha-numeric format. The value equals the ifDescr object, if the LAN device supports RFC 2863.
- **Transmit System Name Enabled** – Provides the system's assigned name in an alpha-numeric format. The value equals the sysName object, if the LAN device supports RFC 3418.
- **Receive Enabled** – LLDP Receive can be enabled or disabled.
- **Transmit System Description Enabled** – Provides a description of the network entity in an alpha-numeric format. This includes the system's name and versions of hardware, operating system, and networking software supported by the device. The value equals the sysDescr object, if the LAN device supports RFC 3418.
- **Transmit System Capabilities Enabled** – Indicates the primary function(s) of the device and whether or not these functions are enabled on the device. The capabilities are indicated by two octets. Bits 0 through 7 indicate Other, Repeater, Bridge, WLAN AP, Router, Telephone, DOCSIS cable device, and Station respectively. Bits 8 through 15 are reserved.

### PFC Priorities Area

- **Active Priorities** – The priorities that are marked active for PFC.
- **Enable** – If checked, PFC is enabled.
- **Configured Priorities** – The priorities that are configured, but might not yet be active. Only the RoCE priority should be selected.

### RoCE Priorities Area

- **Active Priorities** – The priority that is marked active for RoCE.
- **Configured Priorities** – The priorities that are configured, but might not yet be active.

## QCN Area

When QCN is available, the Configure DCB Dialog contains the QCN area where you can enable the configured CNPV.

- Active CNPV – 5 is the RoCE default matching the default RoCE priority, the current Active CNPV being used for QCN.
- Enable – If checked, the configured QCN is enabled. If the Enable checkbox is cleared, the configured CNPV value is grayed-out, indicating that it is not applicable.
- Configured CNPV – The Configured CNPV value always matches the configured RoCE priority, and cannot be changed independently.

**Note:** The switch can override the configured QCN's CNPV value.

## ETS Priority Groups Area

- Active Groups
  - Group ID – The priority group ID.
  - Priority Membership – The different priorities that are assigned to the various priority groups. This is the currently active configuration.
  - Bandwidth % – The bandwidths that are assigned to different priority groups. This is the currently active configuration.
- Configured Groups
  - Group ID – The priority group ID.
  - Priority Membership – The configured priority membership grouping.
  - Bandwidth % – The configured value of bandwidth for the different priority groups.

**Note:** PG15 is only available for NIC+RoCE ports and has no bandwidth setting.

- Max Configurable PGs – The maximum number of priority groups that can be configured. Bandwidths of all the priority groups must add up to 100%.

## Configure DCB Dialog Box Buttons

- Defaults – Click to return to the factory settings.
- Configuration Rules – Click to display the NIC Priority window that lists the rules for configuring NIC priorities.

You must observe the following rules when configuring priority groups for RoCE adapter ports:

1. The RoCE priority must be assigned as a PFC priority too.
2. Only the RoCE priority can be assigned to PG15.
3. The RoCE priority can be assigned to the PG15 as a single priority or any other PG as the single priority to the PG (that is, no other priorities can be assigned to the PG where the RoCE priority is assigned). All other priorities should be assigned to a single PG.

4. No bandwidth is specified for PG15. The bandwidths for PG0-PG7 must add up to 100.
  - OK – Click to apply and save your changes.
  - Cancel – Click to discard any changes you made.

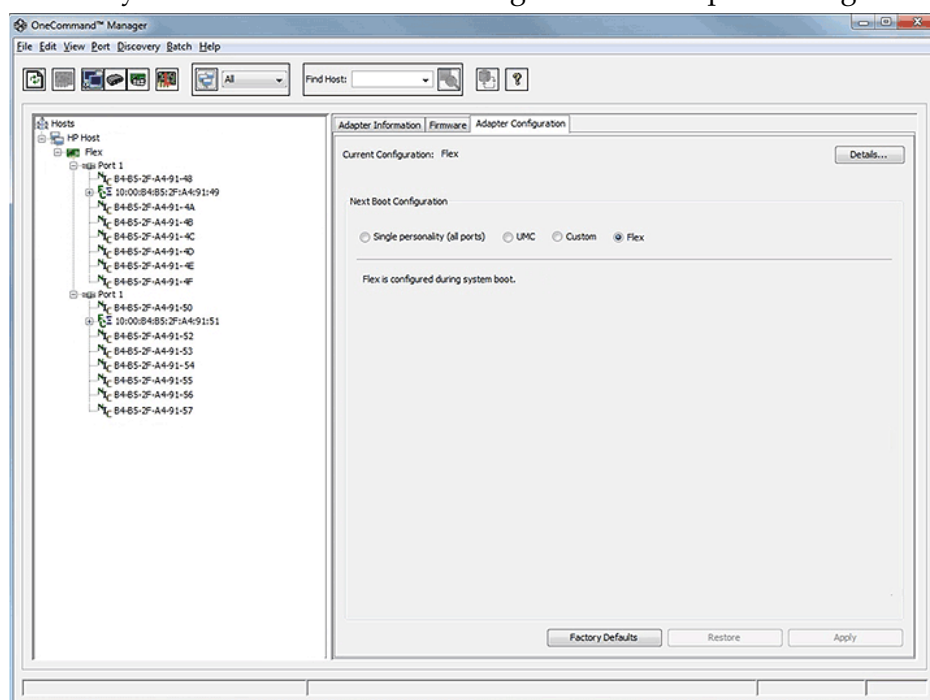


# 11. Configuring HP Flex Adapters

The OneCommand Manager application can actively manage HP Flex adapter configurations, including UMC channel management.

Configuring HP Flex adapters is very similar to configuring Emulex adapters. However, there are some minor differences due the Flex mode capability of the adapter.

**Note:** Even though the OneCommand Manager application may change the port configuration on the adapter, these changes can be overridden on the next boot by the commands used to configure a Flex adapter during the boot process.

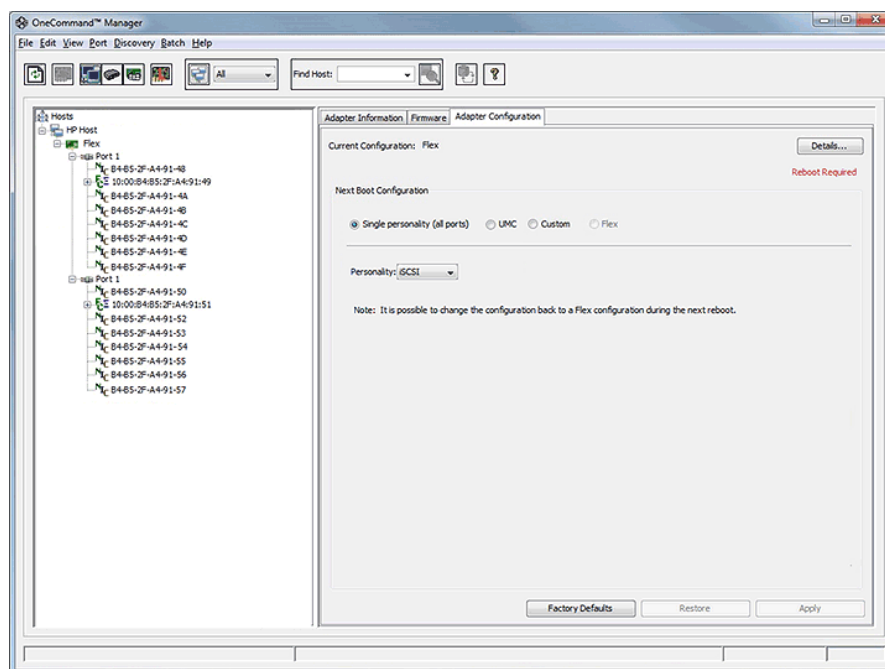


**Figure 11-1** HP Flex Adapter Configuration tab (Flex enabled)

When a Flex adapter is selected in the discovery-tree, the Adapter Configuration tab shows the Current Configuration as “Flex”.

Clicking **Details...** displays the Next Boot Configuration options as depicted in Figure 11-1 on page 232.





**Figure 11-2** Flex Adapter After Applying an iSCSI Single Personality Change

You can change the adapter's configuration by selecting Single Personality, UMC (if available), or Custom and then clicking **Apply**.

**Note:** If you click the **Apply** button to save the new configuration, it cannot be returned to a Flex configuration.

Checking Single Personality enables the Personality menu so you can choose the personality type you want to assign.

See "UMC Configuration View" on page 98 for UMC configuration instructions.

See "Custom Configurations" on page 94 for instructions on custom configurations.

## 12. Exporting SAN Information

### 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:

1. From the **File** menu, select **Export SAN Info**.
2. Browse to a folder and enter a filename 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 hand side of the progress bar. On Windows, you cannot change views, reset, or download firmware during the export process.

## 13. Diagnostics

This section describes the diagnostic tests that can be run on LightPulse and OneConnect adapters.

### LightPulse FC Diagnostics

This section describes the diagnostics (Figure 13-1 on page 236) available for LightPulse FC adapters. For OneConnect adapter diagnostics, see “OneConnect Diagnostics” on page 247.

Use the Diagnostics tab to do the following:

- View flash load list, PCI registers, and wakeup parameter information.
- Run the following tests on Emulex adapters installed in the system: (Not available in read-only mode.)
  - PCI Loopback
  - Internal Loopback
  - External Loopback
  - Power-on self test (POST)
  - Echo (End-to-End)
  - Quick Test
- Perform a diagnostic dump and retrieve dump files from remote hosts. (Not available in read-only mode. For 16GFC HBAs, refer to “Creating Diagnostic Dumps” on page 252 in the “OneConnect Diagnostics” section.)
- Control adapter beaconing. (Not available in read-only mode.)

## Viewing Flash Contents, PCI Registers, and Wakeup Information

The Diagnostics tab (Figure 13-1 on page 236) shows PCI register dump information and flash memory contents. The information is read-only and is depicted below.

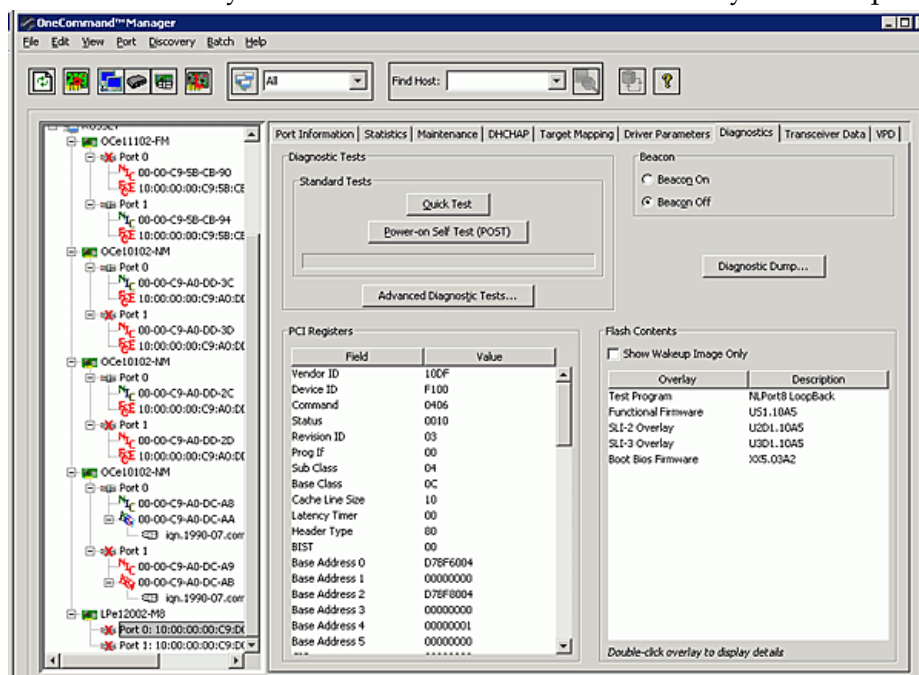


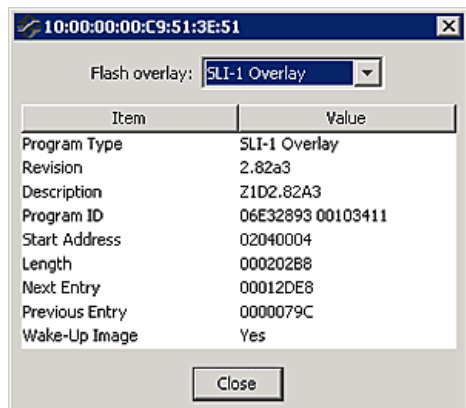
Figure 13-1 PCI Registers and Flash Contents of the Diagnostics Tab

## Viewing Flash Contents

If you check the **Show Wakeup Image Only** checkbox, the flash overlays that are not loaded when the system is booted no longer display. This checkbox defaults to unchecked.

## Viewing Overlay Details

If you double-click on a flash overlay, another window appears with details about that overlay (Figure 13-2 on page 237).



**Figure 13-2** Overlay Detail Window

To see the details of a different flash overlay image, you can either close the details window and double-click on another overlay name, or choose a different overlay name from the Flash overlay menu.

## Viewing the PCI Registers


The PCI Registers appear directly on the Diagnostics tab (Figure 13-1 on page 236).

## Running a Quick Test

The Diagnostics tab enables you to run a “quick” diagnostics test on a selected FC port. The Quick Test consists of 50 PCI Loopback test cycles and 50 Internal Loopback test cycles. (Not available in read-only mode or on LightPulse adapters in ESXi hosts.)

To use quick test:

1. From the discovery-tree (Figure 4-3 on page 40), select the FC port on which you want to run the Quick Test.
2. Select the **Diagnostics** tab (Figure 13-1 on page 236) and click **Quick Test**. A warning message appears (Figure 13-3 on page 237).



**Figure 13-3** Quick Test Warning

3. Click **OK** to run the test. The Quick Diagnostic Test window appears displaying the PCI Loopback and Internal Loopback test results.

## Running a Power-On Self Test (POST)

**Note:** Not supported on LPe15000 or LPe16000-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. (Not available in read-only mode.)

To run the POST:

1. From the discovery-tree (Figure 4-3 on page 40), select the FC port on which you want to run the POST.
2. Select the **Diagnostics** tab (Figure 13-1 on page 236) and click **Power-on Self Test (POST)**. A warning dialog box appears.
3. Click **OK**. A POST window appears displaying POST information.

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

## Using Beaconing

The beaoning 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. (Not available in read-only mode.)

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

**Note:** The beaoning buttons are disabled if the selected adapter does not support beaoning.

To enable or disable beaoning:

1. From the discovery-tree (Figure 4-3 on page 40), select the FC port whose LEDs you want to set.
2. Select the **Diagnostics** tab (Figure 13-1 on page 236) and click **Beacon On** or **Beacon Off**.

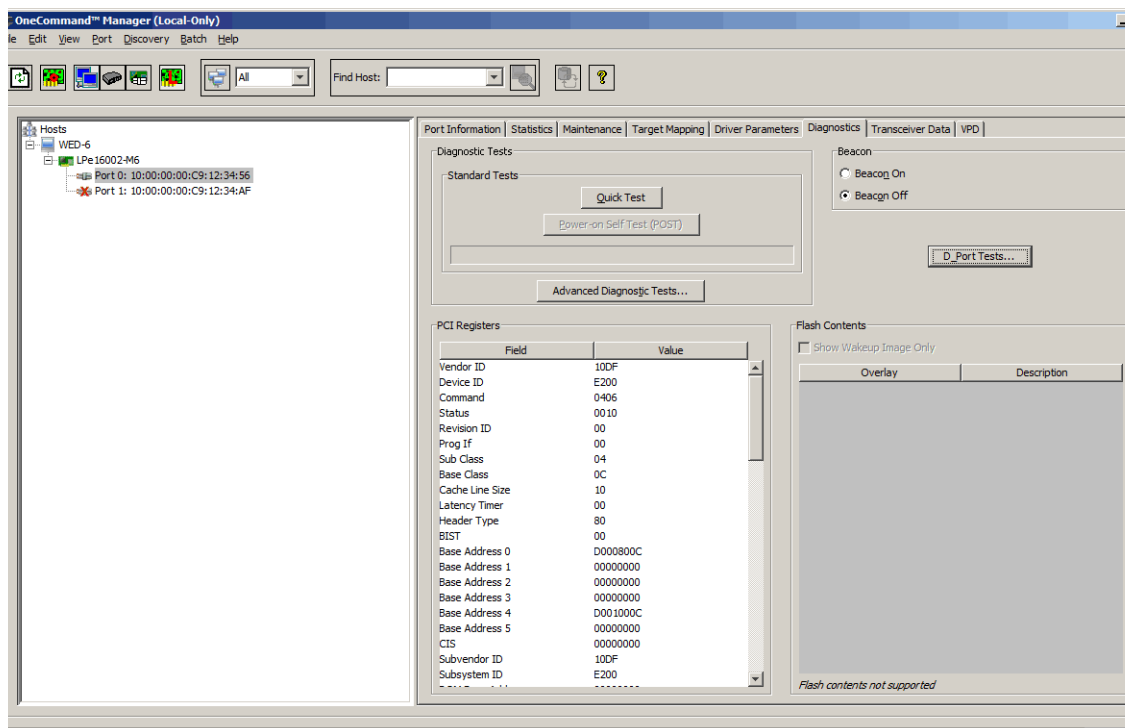
## Running D\_Port Tests

D\_Port is a diagnostic mode supported by Brocade switches for 16GFC. 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 due to marginal cable, optics modules, and connector or optics seating.

## Notes

- Basic connectivity diagnostics are already supported by Emulex HBAs. The OneCommand Manager application has diagnostic modes that support validation of 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.
- D\_Port should not be enabled 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 will appear 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 website at [www.brocade.com](http://www.brocade.com).
- D\_Port is also referred to as ClearLink.

The D\_Port Tests button on the Diagnostics tab enables you to run D\_Ports tests on LPe16000-series adapters (Figure 13-4 on page 239).



**Figure 13-4** Diagnostics Tab for LPe16000-Series Adapters (D\_Port Tests... button depicted)

To run a D\_Port test:

1. From the discovery-tree, select the LPe16000-series FC port on which you want to run the D\_Port test.

2. Select the **Diagnostics** tab (Figure 13-1 on page 236) and click **D\_Port Tests...** The D\_Port Tests window appears (Figure 13-5 on page 240).
3. Click **Start Tests**.

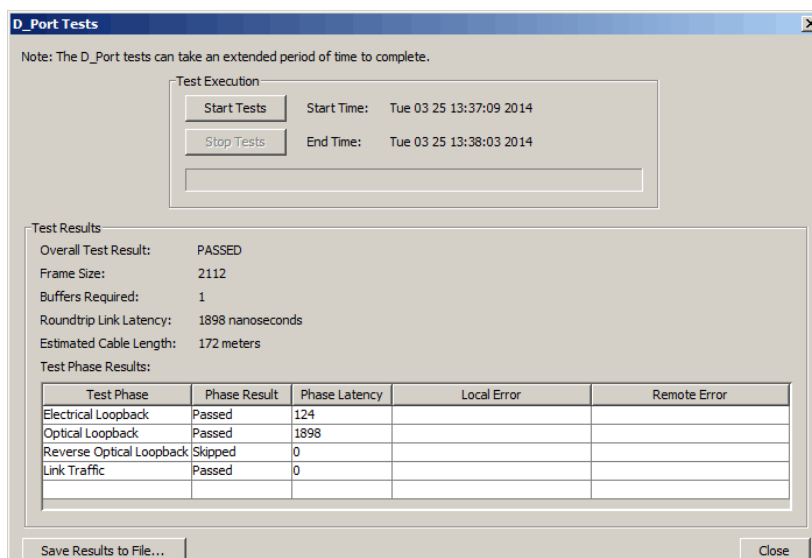


Figure 13-5 D\_Port Tests Window

## D\_Port Window Descriptions

This section provides details on D-Port window field definitions, test phase results, and D\_Port window buttons.

### D\_Port Window Field Definitions

- Overall Test Result – Displays PASSED or FAILED depending upon the outcome of all the test phases.
- Frame Size – The size of the frames used in each test phase.
- Frame Count – The number of frames generated during each test phase.
- Roundtrip Link Latency – Estimated roundtrip link latency calculated by switch during the execution of all tests.
- Estimated Cable Length – Estimated cable length calculated by switch during the execution of all tests.

### D\_Port Window Test Phase Results List

- Test Phase – The name of the test run.
- Phase Result – The result of the test run. Possible results are Pass, Fail, or Skipped.
- Phase Latency – The round trip legacy (in ns.) calculated during the execution of the test.
- Local Error – The error(s), if any, detected on the local side of the test.
- Remote Error – The error(s), if any, detected on the remote side of the test.



## D\_Port Window Buttons

- Start Tests – Click to start D\_Port tests. The start time is displayed.
- Stop Tests – Click to stop running D\_Port tests. The stop time is displayed.
- Save Results to File - Click to save test results to a file you specify.
- Close - Click to close the window (disabled while tests are running).

**Note:** If the SFP or adapter firmware do not support running D\_Port diagnostics, clicking **Start Tests** causes an error message to be displayed indicating this and the tests will not be executed.

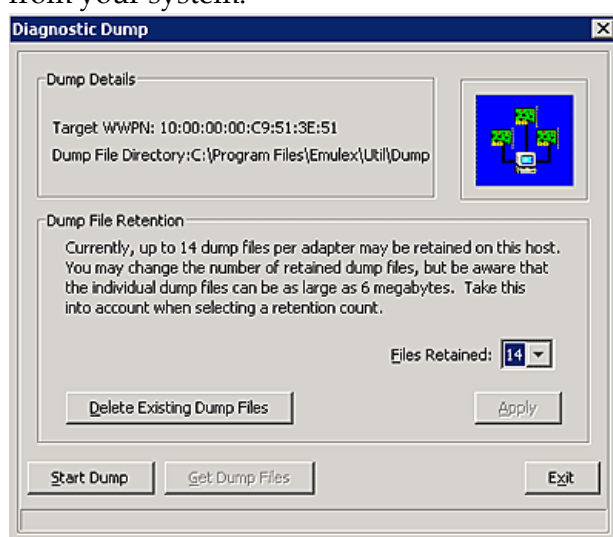
## Creating Diagnostic Dumps

The diagnostic dump capability enables you to create a “dump” file for a selected FC port. Dump files contain various information such as firmware version, driver version, and so on, that is particularly useful when troubleshooting an adapter. You can also retrieve dump files from remote hosts. (Not available in read-only mode.) For 16GFC HBAs, see the OneConnect section “Creating Diagnostic Dumps” on page 252.

**Caution:** Disruption of service can occur if a diagnostic dump is run during I/O activity.

To start a diagnostic dump:

1. From the discovery-tree (Figure 4-3 on page 40), select an FC port whose diagnostic information you want to dump.
2. Select the **Diagnostics** tab (Figure 13-1 on page 236) and click **Diagnostic Dump**. The Diagnostic Dump dialog box appears (Figure 13-6 on page 241). You can specify how many files you want to retain using the Files Retained counter. Click **Delete Existing Dump Files** to remove existing dump files for the selected FC port from your system.



**Figure 13-6** Diagnostic Dump Dialog Box

3. Click **Start Dump**. A warning message appears about taking the adapter offline.

**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.

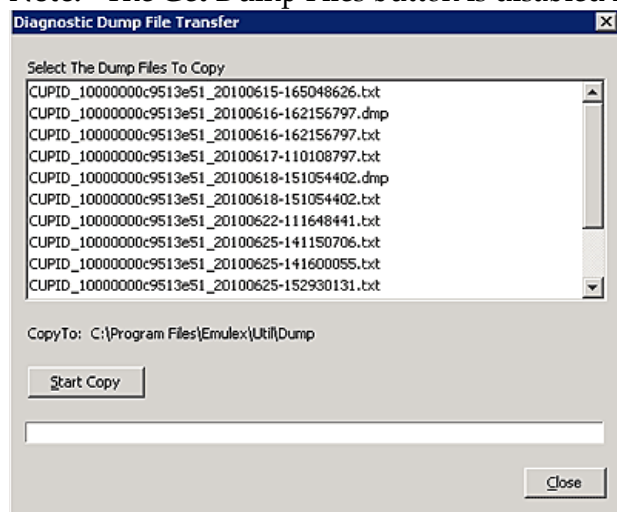
4. Click **OK**. Dump files are created. Where these files are created depends upon your operating system:
  - Windows – In the Dump directory under the OneCommand Manager Installation Directory Util\Dump\
  - Solaris – /opt/ELXocm/Dump
  - Linux – /var/log/emulex/ocmanager/Dump
  - VMware – a dump directory you created under /vmfs/volumes.

Two files are created:

- <Hostname\_WWPN\_Date-Time>.dmp
- <Hostname\_WWPN\_Date-Time>.txt

5. To obtain remote host dump files and copy them to your local system, click **Get Dump Files**. The Diagnostic Dump File Transfer dialog box appears (Figure 13-7 on page 242).

**Note:** The Get Dump Files button is disabled if a local adapter port is selected.



**Figure 13-7** Diagnostic Dump File Transfer Dialog Box

6. 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 4.

## 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 what to do in the event of a test failure. (Not available in read-only mode.)

To run advanced diagnostics tests:

Click **Advanced Diagnostic Tests** on the Diagnostics tab (Figure 13-1 on page 236) to view the Diagnostic Test Setup dialog box (Figure 13-8 on page 243).

You can run four types of tests:

- PCI Loopback
- Internal Loopback
- External Loopback
- End-to-End (ECHO)

**Note:** You cannot run the External Loopback test and ECHO test concurrently. If you select External Loopback, the ECHO test section is disabled and vice versa.

Test results and the status of running tests are time stamped and appear in the Test Log area.

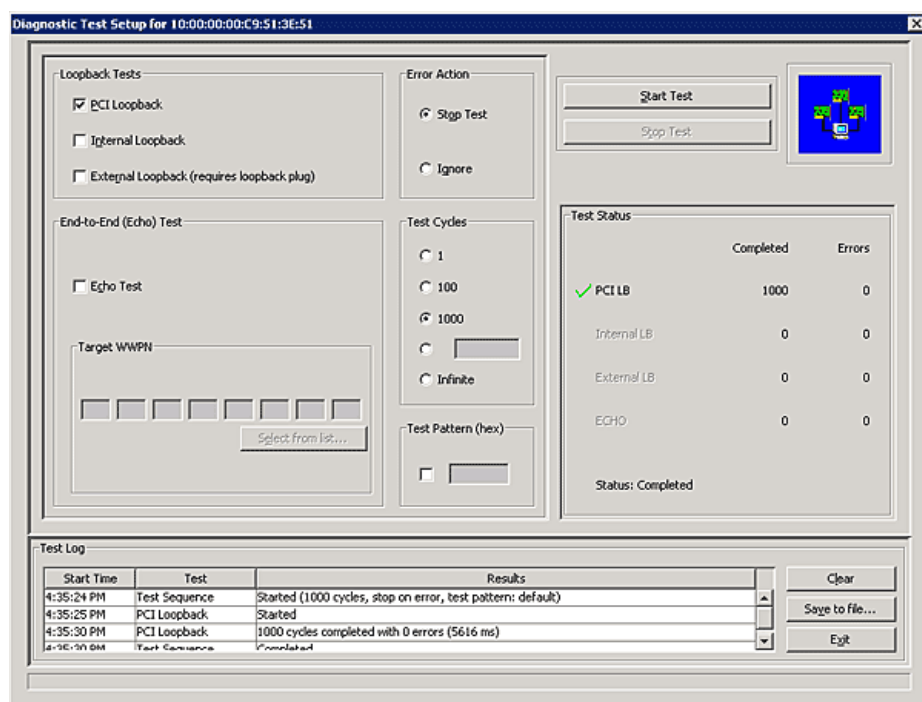


Figure 13-8 Diagnostic Test Setup

## Running Loopback Tests

To run a loopback test, use the Loopback Test section of the Diagnostic Test Setup dialog box (Figure 13-8 on page 243).

### Loopback Test Combinations

Run the following loopback test combinations using the appropriate checkboxes:

- **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 down 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 down to an adapter link port. The data goes out the port and immediately returns via 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 disabled and vice versa.

## Error Action

Error Action enables you to define what is to be done in the event of a test failure. Two error action options exist:

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

## Test Cycles

Test Cycles enables you to specify test cycles three ways:

- Select an established cycle count by clicking on the corresponding radio button.
- Enter a custom cycle count in the blank field in the Test Cycles area.
- Set the test to run until you manually click Stop Test, by selecting the Infinite radio button.

## Test Pattern

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

## Test Status


The Test Status area displays how many completed cycles of each test ran, as well as the number of errors.

To run loopback tests:

1. From the discovery-tree (Figure 4-3 on page 40), select the adapter port on which you want to run the Loopback Test.
2. Select the **Diagnostics** tab (Figure 13-1 on page 236) and click **Advanced Diagnostics Tests** (Figure 13-8 on page 243). From the Loopback Test section of the dialog box, choose the type of Loopback test you want to run and define the loopback test parameters.


**Note:** You must insert a loopback plug in the selected adapter before running an External Loopback test.

3. Click **Start**. The following warning appears (Figure 13-19 on page 254):



**Figure 13-9** Run Diagnostic Tests Warning

4. Click **OK**. If you choose to run an External Loopback test the following window appears (Figure 13-10 on page 245):



**Figure 13-10** Advanced Diagnostic Tests Warning Window for External Loopback

5. Click **OK**. The 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 box. 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 may take some time before they stop.

## Running End-to-End (ECHO) Tests

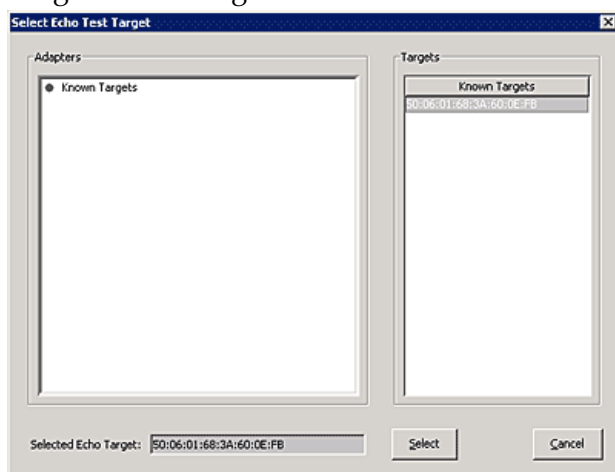
Run echo tests using the End-to-End (ECHO) Test section of the Diagnostics tab. The end-to-end test enables you send an ECHO command/response sequence between an adapter port and a target port. (Not available in read-only mode.)

**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 the ECHO Test, the External Loopback test is disabled.

To run end-to-end echo tests:


1. From the discovery-tree (Figure 4-3 on page 40), select the adapter port from which to initiate the End-to-End (ECHO) Test.
2. Select the **Diagnostics** tab (Figure 13-1 on page 236) and click **Advanced Diagnostic Tests** (Figure 13-8 on page 243).
3. Check **Echo Test**. Enter the WWPN for the target.

4. Click **Select From List** if you do not know the actual WWPN of the test target. The Select Echo Test Target dialog box appears (Figure 13-11 on page 246). 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 dialog box.



**Figure 13-11** Select Echo Test Target Window

5. Define the other parameters you want to use and click **Start Test**. The following warning window appears (Figure 13-12 on page 246):



**Figure 13-12** Advanced Diagnostic Tests Warning Window

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


## 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. (Not available in read-only mode.)

The default location is:

- 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 appears below:



**Figure 13-13** Example of a DiagTest.log Window

To save the log file:

1. After running a test from the Diagnostic Test Setup dialog box (Figure 13-8 on page 243), click **Save to File**. The Select Diagnostic Log file Name dialog box 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**. See Figure 13-8.

## OneConnect Diagnostics

This section describes the diagnostics for OneConnect adapters. For FC adapter diagnostics, see “LightPulse FC Diagnostics” on page 235.

**Note:** Diagnostics are not available in read-only mode. See “Changing Management and Read-Only Mode” on page 44 for more information.

Use the Diagnostics tab (Figure 13-14 on page 248) to do the following:

- Run the following tests on OneConnect adapters installed in the system:
  - Host to DDR DMA Loopback
  - PHY Loopback
  - MAC Loopback
  - End-to-End (ECHO) (FCoE only)
  - External Loopback
- Perform a diagnostic dump and retrieve dump files from remote hosts.
- Control adapter beaconing.



All functions are supported locally and remotely on hosts managed with TCP/IP access. Test results and the status of running tests are time stamped and appear in the Test Status area.

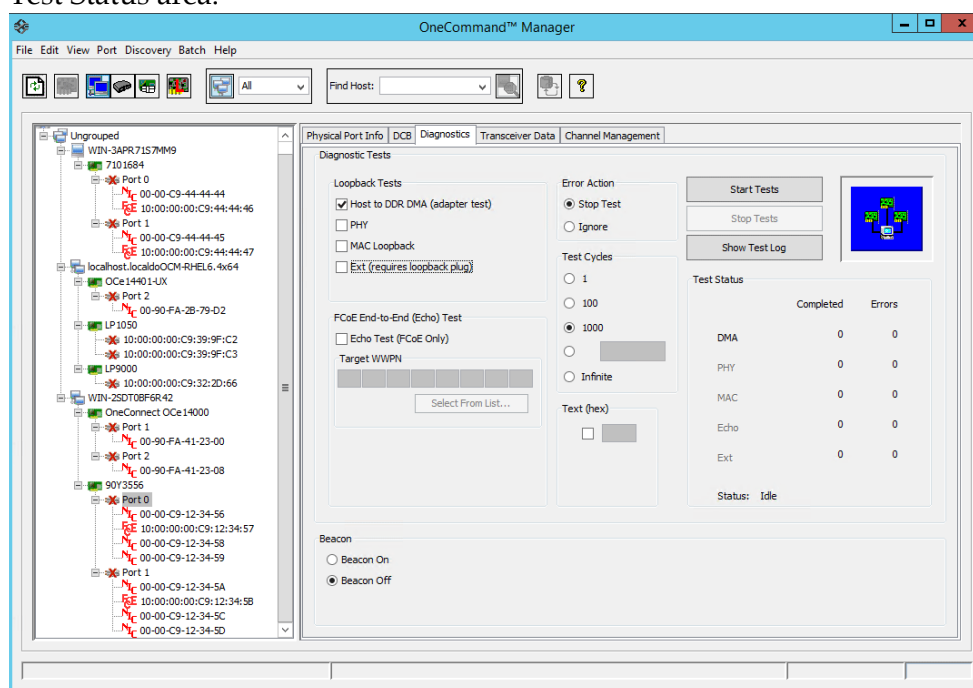


Figure 13-14 Diagnostics Tab (10GBASE-T adapter selected)

## OneConnect Loopback Test Combinations

Run the following loopback test combinations using the appropriate checkboxes:

- **Host to DDR DMA Loopback Test** – The Host to DDR DMA loopback test sends data from the host to the adapter, then back to the host, where it is checked for data miscompute errors. All tests except one are run on the currently selected port. This test is run across the entire adapter. The same diagnostic is therefore executed regardless of the currently selected physical port. Also, unlike other diagnostics, this test affects the operation of all ports on the adapter. (Not available on ESXi systems.)
- **PHY Loopback Test** – The PHY loopback test connects the transmit output of the physical layer to the receive input of the physical layer. The data is transmitted, received and checked for data miscompute errors.

**Note:** PHY diagnostics are not supported on mezzanine cards and blade network daughter cards because they do not contain PHYs.

- **External Loopback Test** – A diagnostic test in which a random data pattern is sent down to an adapter link port. The data goes out the port and immediately returns via a loopback connector. The returned data is subsequently validated for integrity. (Not available on 10GBASE-T adapters.)
- **MAC Loopback** – MAC loopback connects the transmit output of the MAC controller to the receive input of the MAC controller (bypassing the PHY).



## FCoE End to End Echo Test

The end-to-end test enables you send an ECHO command/response sequence between an adapter port and a target port. (Not available on ESXi systems.)

**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 the ECHO Test, the External Loopback test is disabled.

## Error Action

Error Action enables you to define what is to be done in the event of a test failure. Two error action option exists:

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

## Test Cycles

Test Cycles enables you to specify test cycles three ways:

- Select an established cycle count by clicking on the corresponding radio button.
- Enter a custom cycle count in the blank field in the Test Cycles area.
- Set the test to run until you manually click Stop Test, by selecting the Infinite radio button.

## Test Pattern

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

## Test Status


The Test Status area displays how many completed cycles of each test ran, as well as the number of errors.

To run loopback tests:

1. From the discovery-tree (Figure 4-3 on page 40), select the adapter port on which you want to run the Loopback Test.
2. Select the **Diagnostics** tab (Figure 13-14 on page 248). From the Loopback Test section of the dialog box, choose the type of Loopback test you want to run and define the loopback test parameters.

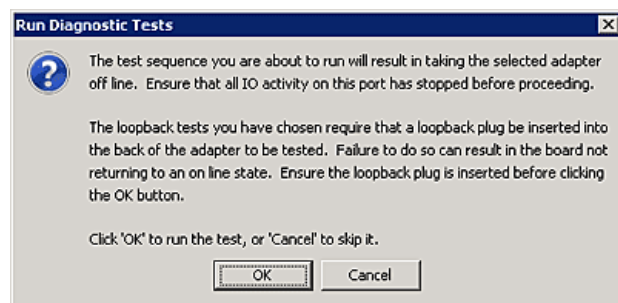
**Note:** You must insert a loopback plug in the selected adapter before running an External Loopback test. Also, you must ensure that the NIC function of the port goes to a link up state. See “Troubleshooting” on page 255 if the NIC link fails to come up.

3. Click **Start**. The following warning appears (Figure 13-15 on page 250):



**Figure 13-15** Run Diagnostic Tests Warning

4. Click **OK**. If you choose to run an External Loopback test, the following window appears (Figure 13-16 on page 250):



**Figure 13-16** Advanced Diagnostic Tests Warning Window for External Loopback

5. Click **OK**. The 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 Status section of the dialog box. Click **Show Test Log** to view and 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 may take some time before they stop.

## 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. (Not available in read-only mode.)

When you enable beaconing for LightPulse adapters, the two LEDs blink rapidly in unison for 24 seconds, after which the LEDs report the adapter health status for 8 seconds. When the 8 seconds are up, the adapter returns to beaconing mode. This cycle repeats indefinitely until you disable beaconing or you reset the adapter.

When you enable beaconing for OneConnect adapters, the two LEDs blink rapidly in unison until you disable beaconing.

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

To enable or disable beaconing:

1. From the discovery-tree (Figure 4-3 on page 40), select the adapter port whose LEDs you want to set.

2. Select the **Diagnostics** tab (Figure 13-14 on page 253) and click **Beacon On** or **Beacon Off**.

**Note:** On OCe11100-series adapters, beaconing can only be done on one port at a time. If you enable beaconing on a port, you cannot enable beaconing on another port until you disable beaconing on the previous port.

## Running TDR Tests (10GBASE-T Adapters Only)

**Note:** The TDR test is not supported by the CIM interface.

The TDR test is a cable diagnostic that can tell you the length of the cable, whether or not the cable has any defects (open/short), and the distance to the defect if one exists. It works by sending a signal down the cable and measuring its reflection.

NVP is a property of the cable that must be known for the TDR test to accurately calculate the cable length (or distance to defect).

**Note:** The TDR test requires that the cable be 'down' and quiet. It is best if the cable is not terminated. If it is terminated, the link partner must not be active during the test.

You can run TDR tests from the Diagnostics Tab when you select a 10GBASE - T adapter in the discovery-tree (Figure 4-3 on page 40).

To run a TDR test:

1. From the discovery-tree (Figure 4-3 on page 40), select the 10GBASE-T adapter on which you want to run the test.
2. Select the **Diagnostics** tab (Figure 13-14 on page 248).
3. Assign the NVP and click **Apply**.
4. Click **Start TDR Test**.

## Saving the Log File

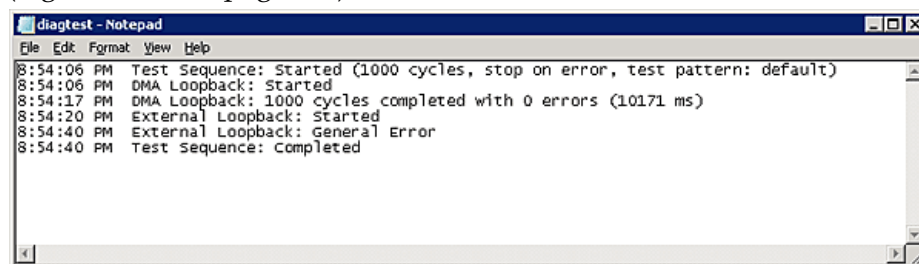
You can save the test log to a log file for later viewing or printing. When new data is written to a saved file, the data is appended to the end of the file. Each entry has a two-line header that contains the identifier of the adapter being tested and the date and time of the test. Over time, the data accumulates to form a chronological history of the diagnostics performed on the adapter. (Not available in read-only mode.)

The default location is:

- In Windows: the OneCommand Manager application install directory on your local drive
- In Solaris: /opt/ocmanager/Dump
- In Linux: /var/opt/emulex/ocmanager/logs
- 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 appears below

(Figure 13-17 on page 252):



**Figure 13-17** Example of a DiagTest.log Window

To save the log file:

1. After running a test from the Diagnostic tab (Figure 13-14 on page 248), click **Save Test Log**. The Diagnostic Test Log dialog box appears (Figure 13-17 on page 252). 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 to file**.

## Creating Diagnostic Dumps

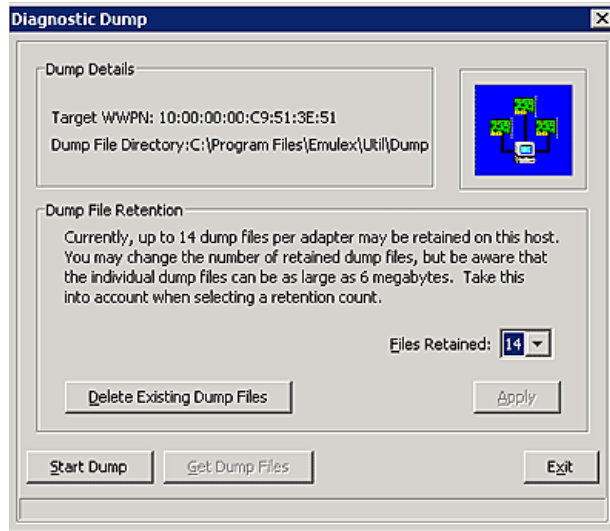
The diagnostic dump capability enables you to create an EFD “dump” file for a selected adapter. Dump files contain various information such as firmware version, driver version, and so on, that is particularly useful when troubleshooting an adapter. You can also retrieve dump files from remote hosts. (Not available in read-only mode.)

To start a diagnostic dump:

1. From the discovery-tree (Figure 4-3 on page 40), select an adapter whose diagnostic information you want to dump.
2. Select the **Firmware** tab (Figure 7-14 on page 84) and click **Diagnostic Dump**. The Diagnostic Dump dialog box appears (Figure 13-18 on page 253).

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 you want to retain using the Files Retained counter. Click **Delete Existing Dump Files** to remove existing dump files for the selected adapter from your system.



**Figure 13-18** Diagnostic Dump Dialog Box

4. Click **Start Dump**. Dump files are created. Where these files are created depends upon your operating system:

**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.

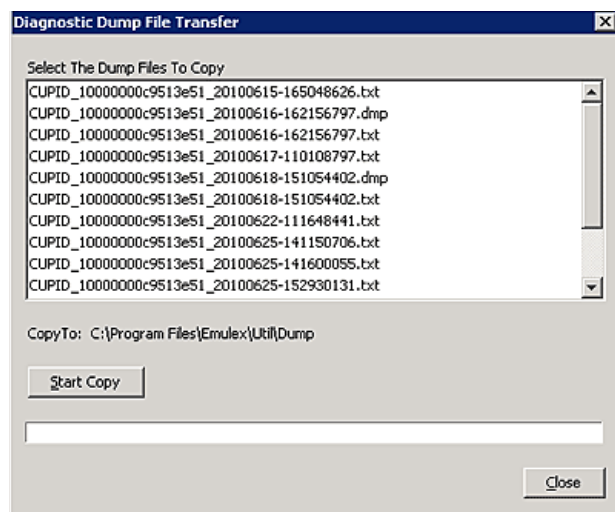
- Windows – %ProgramFiles%\Util\Dump\
- Solaris – /opt/ELXocm/Dump
- Linux – /var/log/emulex/ocmanager/Dump
- VMware – a dump directory you create under /vmfs/volumes.

Two files are created:

- <Hostname\_WWPN\_Date-Time>.efd
- <Hostname\_WWPN\_Date-Time>.txt

5. To obtain remote host dump files and copy them to your local system, click **Get Dump Files**. The Diagnostic Dump File Transfer dialog box appears (Figure 13-19 on page 254).

**Note:** The Get Dump Files button is disabled if a local adapter port is selected.



**Figure 13-19** Diagnostic Dump File Transfer Dialog Box

6. 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 4.

## 14. Troubleshooting

Your system may operate in an unexpected manner due to several circumstances. The Troubleshooting section explains many of these circumstances and offers one or more workarounds for each situation.

### General Situations

Table 14-1 on page 255 lists possible situations and their resolution.

**Table 14-1** General Situations

Situation	Resolution
After installing and starting the OneCommand Manager application, the status bar shows “Initializing discovery engine...”, but after waiting for awhile, nothing is displayed in the discovery-tree.	It is possible the discovery server was not installed properly and therefore is not running. Try uninstalling and re-installing the OneCommand Manager application package.
The Web Launch interface cannot be started. When you attempt to start the OneCommand Manager application Web Launch Interface client interface, you receive an error message stating “Unable to launch OneCommand.”	If the JRE/Web Start version present on your system does not meet the minimum required by the OneCommand Manager application, a temporary copy of the correct Web Start version is downloaded automatically. This is used to open the OneCommand Manager application Web Launch Interface client interface and is then discarded after you terminate your session. On some systems, however, security settings or other factors may prevent this download from completing successfully, resulting in this error. To fix the issue, manually update the JRE on your system to the version required by the OneCommand Manager application.
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 1, or 2GFC device. Only 8 GFC devices are supported on 8GFC adapters and only 16GFC devices are supported on 16GFC adapters.
The other utilities install, but the OneCommand Manager application does not.	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 (see the Installation section of the driver manual). 2) Install the utilities (see the Installation section of the driver manual).
When attempting to start the OneCommand Manager application, the Web browser displays “Emulex Corporation OneCommand Demo of OneCommand WebStart web n.n.n.n...”	The document caching mechanism sometimes behaves erratically if more than one version of Java Runtime is installed on the browser client. Two workarounds are available for this issue: <ul style="list-style-type: none"> <li>Exit the browser and restart it. The OneCommand Manager application Web Launch Interface starts successfully.</li> <li>Uninstall all non-essential versions of the Java Runtime. The OneCommand Manager application Web Launch Interface requires that only a single version of the Java Runtime be installed on the browser client. This single version must be Java 6.0 or later for all platforms.</li> </ul>

**Table 14-1** General Situations (Continued)

Situation	Resolution
In the OneCommand Manager application discovery-tree, multiple CNA FCoE or iSCSI functions are grouped under a single physical port.	Ensure the Emulex NIC driver is loaded and that the operating system sees ALL NIC ports. They do not need to be plumbed or configured; just visible to the operating system.
An operating error occurs when attempting to run the OneCommand Manager application. When you attempt to run the utility, an operating system error may occur. The computer may freeze.	Reboot the system.
Unwanted remote servers appear in the OneCommand Manager application.	<p>To prevent remote servers from appearing in the OneCommand Manager application, do one of the following on the remote systems:</p> <ul style="list-style-type: none"> <li>• In Windows, disable the OneCommand Manager application service.</li> <li>• In Linux, stop the elxhbamgr daemon by running the <code>/usr/sbin/ocmanager/stop_ocmanager</code> script.</li> <li>• In Solaris, stop the elxhbamgr service by issuing the command <code>"svcadm disable elxhbamgr"</code>.</li> </ul> <p><b>Note:</b> Disabling this service or process prevents the local servers from being seen remotely.</p>
<p>If Help--&gt;Contents is selected in the OneCommand Manager application, the online help is not opened in a web browser. The OCManager_Help.htm file may be opened in a text editor (displaying HTML code) or by some other application.</p> <p>This happens when the operating system has associated .HTML files with an application other than a web browser.</p>	<p>On Windows systems, this can be fixed using the following steps:</p> <ol style="list-style-type: none"> <li>1) In Windows Explorer, navigate to the C:\Program Files\Emulex\Util\OCManager\OCManager_help\ directory.</li> <li>2) Right-click on <b>OCManager_Help.htm</b>.</li> <li>3) Select <b>Open With--&gt;Choose default program...</b></li> <li>4) Select a web browser, such as Internet Explorer.</li> <li>5) Check <b>Always use the selected program to open this kind of file</b>.</li> <li>6) Click <b>OK</b>.</li> </ol> <p>On Linux and Solaris, the above steps are very similar, with the OCManager_Help.htm file located at <code>/usr/sbin/ocmanager/ocmanager_help/OCManager_Help.htm</code> and <code>/opt/ELXocm/ocmanager_help/OCManager_Help.htm</code>, respectively.</p>
When using the OneCommand Manager application on a Windows operating system, a reboot message is displayed if the NIC driver properties are changed or the adapter port is disabled.	Ensure that the OneCommand Manager application has been closed prior to making any changes to the NIC driver instance or properties.



## Emulex Driver for Linux and OneCommand Manager Application Situations

Table 14-2 on page 257 lists possible Emulex Driver for Linux situations and their resolution.

**Table 14-2** Emulex Driver for Linux and OneCommand Manager Application Situations

Situation	Resolution
NIC Link fails to come up.	For Emulex CNAs in NIC mode, you may need to properly configure the network interface using system administration utilities.
The OneCommand Manager application software package does install. An error message states that: "inserv Service Elxlpfc has to be enabled for service ElxDiscSrvinserv: exiting now/sbin/ inserv failed exit code 1."	Reinstall the driver with the lpfc-install script.
If a SAN configuration has 256 targets mapped by the FC, FCoE, or FC/FCoE driver, any additional added targets do not get a target ID mapping by the driver and cause target discovery to fail. Removing targets or reinitializing the link does not solve the issue.	Unload and reload the driver to reset available target IDs. Ensure that the SAN configuration is correct prior to reloading the driver. This clears the driver's consistent binding table and free target IDs for new target nodes.
In some cases, after loading a vendor-supplied combined firmware/OpenBoot image you are not able to enable BootBIOS from the lputil Boot BIOS Maintenance menu.	<ol style="list-style-type: none"> <li>1) Download the current OpenBoot only image for your adapter from the Emulex website.</li> <li>2) Load the current OpenBoot only image following steps listed in Updating BootBIOS section of this manual.</li> <li>3) Run lputil, return to the Boot BIOS Maintenance menu.</li> <li>4) Enable BootBIOS.</li> </ol>
rmmod fails to unload FC, FCoE, or FC/FCoE driver module due to ERROR: Module LPFC is in use.  This message can appear when you attempt to remove the driver and there is a Logical Volume Group dependent on the driver.	<p>Make the Logical Volume Group unavailable.</p> <p>Type</p> <pre>lvchange -a n xxxxxxxx</pre> <p>where xxxxxx is the Volume Group Name.</p>
Slow targets or extended link faults on the storage side may result in storage being marked off-line by the mid-layer and remaining off-line (not recovered) when the link faults are corrected.	The 8.2 version of the driver should eliminate this issue. However, if you experience off-line device issues, increase the SCSI command timeout to a value greater than or equal to sixty seconds. Emulex also provides a script which addresses this issue (for 2.6 kernels). To access the lun_change_state.sh script, go to <a href="http://www.emulex.com/files/downloads/linux/tools.html">http://www.emulex.com/files/downloads/linux/tools.html</a> , then click the link to the appropriate driver, and click the Linux tools link.

**Table 14-2** Emulex Driver for Linux and OneCommand Manager Application Situations (Continued)

Situation	Resolution
Under certain conditions of an I/O load, some targets cannot retire an I/O issued by a Linux initiator within the default timeout of 30 seconds given by the SCSI midlayer. If the situation is not corrected, the initiator-to-target condition deteriorates into abort/recovery storms leading to I/O failures in the block layer. These types of failures are preceded by a SCSI IO error of hex 6000000.	Emulex provides a script which addresses this issue. To access the <code>set_target_timeout.sh</code> script, go to <a href="http://www.emulex.com/files/downloads/linux/tools.html">http://www.emulex.com/files/downloads/linux/tools.html</a> , then click the link to the appropriate driver, and click the Linux tools link.
FC, FCoE, or FC/FCoE driver fails to recognize an adapter and logs “unknown IOCB” messages in the system log during driver load. The adapter is running outdated firmware.	Update adapter firmware to minimum supported version listed in the driver installation guide (or newer).
<code>rmmod</code> of FC, FCoE, or FC/FCoE driver hangs and module reference count is 0.	Due to a small race condition in the kernel it is possible for an <code>rmmod</code> command to hang. Issue the <code>rmmod-w</code> command. If this does not help, reboot the computer.
System panics when booted with a failed adapter installed.	Remove the failed adapter and reboot.
<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 a LightPulse 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. Type <code>./stop_ocmanager</code> Unmount any disks connected to the adapter. Unload the driver. Type <code>rmmod lpfc</code>
Driver install fails. The <code>lpfc-install</code> script fails to install the driver.	The install script may fail for the following reasons: <ul style="list-style-type: none"> <li>• A previous version of the driver is installed. Run the <code>/usr/src/lpfc/lpfc-install --uninstall</code> script and then try to install the driver.</li> <li>• The current driver is already installed.</li> <li>• The kernel source does not match the standard kernel name or you are running a custom kernel.</li> </ul>
<ul style="list-style-type: none"> <li>• “No module <code>lpfc</code> found for kernel” error message. When updating the kernel, the <code>.rpm</code> file generates the following error: “No module <code>lpfc</code> found for kernel <code>KERNELVERSION</code>”.</li> <li>• A recently updated kernel cannot find the ramdisk. After updating the kernel, the kernel cannot find the ramdisk which halts or panics the system.</li> <li>• The driver is not loaded after a system reboot after updating the kernel.</li> </ul>	<p>These three situations may be resolved by updating the kernel. Two ways to install the driver into an updated kernel are available. The method you use depends on whether or not you are updating the driver.</p> <ul style="list-style-type: none"> <li>• Update the kernel using the same version of the driver.</li> <li>• Update the kernel using a new version of the driver.</li> </ul> <p>See the Installation section of the driver manual for these procedures.</p>

**Table 14-2** Emulex Driver for Linux and OneCommand Manager Application Situations (Continued)

Situation	Resolution
Driver uninstall fails. The lpfc-install --uninstall script fails with an error.	<p>Try the following solutions:</p> <ul style="list-style-type: none"> <li>Uninstall the OneCommand Manager application by running the ./uninstall script from the OneCommand Manager application installation directory.</li> <li>Unmount all FC disk drives.</li> <li>Unload the FC, FCoE, or FC/FCoE driver.</li> </ul>
lpfc-install script exit code.	The lpfc-install script contains exit codes that can be useful in diagnosing installation issue. See the lpfc-install script for a complete listing of codes and definitions.
The OneCommand Manager application software package does not install. An error message states that: “inserv Service Elxlpfc has to be enabled for service ElxDiscSrvinserv: exiting now/sbin/ inserv failed exit code 1.”	Reinstall the driver with the lpfc-install script.
The Linux SCSI subsystem only sees eight LUNs when more are present.	Some SCSI drivers do not scan past eight LUNs when the target reports as a SCSI-2 device. Force a SCSI bus scan with /usr/sbin/lpfc/lun_scan. SuSE supplies /bin/rescan-scsi-bus.sh which can be changed to scan everything.
Cannot see any adapters.	<p>Try the following solutions:</p> <ul style="list-style-type: none"> <li>Perform an 'lsmod' to see if the Emulex drivers are loaded. Look for an error message on the command line stating the LPFC driver is not loaded. If this is the case, do an insmod of the FC, FCoE, or FC/FCoE driver and re-launch the OneCommand Manager application.</li> <li>Exit the OneCommand Manager application and run the following scripts in this order: <ol style="list-style-type: none"> <li>1.) /usr/sbin/ocmanager/stop_ocmanager - stops the OneCommand Manager application daemons</li> <li>2.) /usr/sbin/ocmanager/start_ocmanager - starts the OneCommand Manager application daemons</li> <li>3.) /usr/sbin/ocmanager/ocmanager - starts the OneCommand Manager application GUI</li> </ol> </li> </ul> <p>The adapters should be visible. If they are not visible, reboot your system.</p>
Cannot see new LUNs.	<p>Try the following:</p> <ol style="list-style-type: none"> <li>1) Click the <b>Refresh LUNs</b> button in the toolbar.</li> <li>2) Exit the OneCommand Manager application and restart it. If new LUNs are visible, you are finished.</li> </ol> <p>If that does not work, try the following:</p> <ol style="list-style-type: none"> <li>1) Exit the OneCommand Manager application.</li> <li>2) Navigate to /usr/sbin/ocmanager.</li> <li>3) Run ./stop_ocmanager to stop both the elxhamgr and elxdiscov processes.</li> <li>4) Run ./start_ocmanager and ./start_elxdiscov to restart both processes.</li> <li>5) Start the OneCommand Manager application.</li> </ol>

## Emulex Driver for Solaris and OneCommand Manager Application Situations

Table 14-3 on page 260 lists possible Emulex Driver for Solaris situations and their resolution.

**Table 14-3** Emulex Driver for Solaris and OneCommand Manager Application Situations

Situation	Resolution
NIC link fails to come up.	For Emulex CNAs in NIC mode, you may need to properly configure the network interface using system administration utilities.

## VPorts and OneCommand Manager Application Situations

Table 14-4 on page 260 lists possible VPorts situations and their resolution.

**Table 14-4** VPorts and OneCommand Manager Application Situations

Situation	Resolution
vPort creation failure.	<p>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. This may be the issue. For a detailed list of unsatisfied conditions:</p> <ol style="list-style-type: none"> <li>1) Start the OneCommand Manager application.</li> <li>2) Select <b>View&gt;Group Adapters by Virtual Port</b> from the Main menu.</li> <li>3) In the discovery-tree, select the FC/FCoE function on which you would like to create a virtual port.</li> <li>4) The Virtual Ports tab should contain a list of unsatisfied conditions (if any) that are preventing a virtual port from being created. If there are no unsatisfied conditions, yet vPort creation still fails, contact Emulex technical support.</li> </ol>
Port not ready.	<p>The controls in the New Virtual Port box of the Virtual Port window are replaced by a list of reasons why VPorts cannot be created. The reasons can be one or more of the following:</p> <ul style="list-style-type: none"> <li>• The driver NPIV parameter is disabled.</li> <li>• SLI-3 is not being used by a port.</li> <li>• The adapter port is out of resources for additional virtual ports.</li> <li>• The port is not connected to a fabric.</li> <li>• The fabric switch does not support virtual ports.</li> <li>• The fabric switch is out of resources for additional virtual ports.</li> <li>• The port link state is down.</li> </ul>

## Appendix A. License Notices

### 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.
 */
```

## 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.

## 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. 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](mailto: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 rouines 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.]  
\*/