Documentation for Emulex® OneConnect® Adapters

User Guide

Version 11.2
December 30, 2016

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

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

Copyright © 2003-2017 Broadcom. All Rights Reserved.

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

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.

Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Table of Contents

Documentation for Emulex OneConnect Adapters User Guide 1

Emulex Software Kit Migration User Guide 27

Chapter 1: Introduction 30

1.1 Supported Adapters 30
1.2 Considerations 30
1.2.1 Inbox Drivers 31
1.2.2 Management Applications 31
1.2.3 Firmware 31
1.3 Driver Kit Designations and Application Kit Names 31
1.4 Additional Information 32

Chapter 2: Installing Drivers 33

2.1 Installing Split Out-of-Box Drivers over Inbox Drivers 33
2.1.1 Windows 33
2.1.2 Linux 33
2.1.3 VMware Drivers 35
2.2 Installing Split Out-of-Box Drivers over Existing 11.1 or Earlier Out-of-Box Drivers 37
2.2.1 Windows 37
2.2.2 Linux 38
2.2.3 VMware 40
2.2.4 Solaris 42

Chapter 3: Reverting Drivers and Applications to Previous Versions 45

3.1 Windows 45
3.2 Linux 47
3.3 VMware 51
3.4 Solaris 53

Chapter 4: Technical Tips 57

4.1 General 57
4.1.1 Applications Specific to LightPulse or OneConnect 57
4.1.2 Multiple ASICs 57
4.1.3 Reverting to a Prior Driver Version 57
4.1.4 Boot Code Functionality 58
4.2 Windows 58
4.2.1 Windows OneInstall 58
4.2.2 Downgrading from a LightPulse 11.2 Kit with LPe16000-series Adapters 58
4.2.3 Downgrading from a LightPulse 11.2 Kit to an 11.0 Kit with LPe12000-Series Adapters 59
4.3 Linux 59
4.3.1 Unloading the lpfc and the brcmfcoe Drivers 59
4.3.2 Installing a pre-11.2 Version of OneCommand Manager 59
4.4 VMware 60
4.4.1 VMware Inbox Driver Installation over 11.2 Drivers 60
4.5 Solaris 61
4.5.1 Running the Install Script of a pre-11.2 Kit Does Not Uninstall the Broadcom OneCommand CNA Manager 61
4.6 Dell Applications Downgrade 61

Appendix A: Driver Names 63

Appendix B: Quick Lookup Guide 64

B.1 Quick Lookup Guide, LightPulse and OneConnect 64
B.1.1 LightPulse Adapters and Kits 65

B.1.2 OneConnect Adapters and Kits 71

Appendix C: Special Instructions 77
C.1 Linux 77
C.2 VMware 77

Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide 82

Chapter 1: Introduction 84
1.1 Abbreviations 84

Chapter 2: Installing and Uninstalling OneCommand CNA Manager Application Components 88
2.1 Installing the OneCommand CNA Manager Application 88
2.1.1 In Windows 88
2.1.2 In Linux 90
2.1.3 In Solaris 92
2.1.4 In VMware 94
2.2 Installing the OneCommand CNA Manager Application Web Launch Interface 94
2.2.1 Requirements 94
2.2.2 Installation 96
2.3 Uninstalling the OneCommand CNA Manager Application 97
2.4 Uninstalling the OneCommand CNA Manager Application Web Launch Interface Only 98

Chapter 3: Starting and Stopping the OneCommand CNA Manager Application 99
3.1 In Windows 99
3.2 In Linux and Solaris 99
3.3 Starting the OneCommand CNA Manager Application Web Launch Interface 100
3.3.1 Managing Files When Running the OneCommand CNA Manager Application Web Launch Interface 100

Chapter 4: Using the OneCommand CNA Manager Application 101
4.1 The OneCommand CNA Manager Application Window Element Definitions 101
4.1.1 Menu Bar 102
4.1.2 Toolbar 102
4.1.3 Discovery-Tree 103
4.1.4 Property Tabs 105
4.1.5 Status Bar 105
4.2 Using OneCommand CNA Manager Secure Management 105
4.2.1 OneCommand CNA Manager Secure Management Configuration Requirements 107
4.3 Changing Management and Read-Only Mode 107
4.3.1 Management Host 108
4.4 Using CIM (Windows Only) 109

Chapter 5: Configuring Discovery 110
5.1 Discovery Using the TCP/IP Access Protocol 110
5.1.1 Hosts File 110
5.1.2 Adding a Single Host 112
5.1.3 Adding a Range of Hosts (IPv4 Only) 113
5.1.4 Removing Hosts 114
5.2 Configuring Discovery and Default CIM Credentials 115
5.3 Viewing Discovery Information 116

Chapter 6: Managing Hosts 118
6.1 Viewing Host Information 118
6.2 Viewing Host Grouping Information 119
6.3 Grouping Hosts 120
6.3.1 Managing Host Groups 121
6.3.2 Creating a Host Group 122

6.3.3 Deleting a Host Group 122
6.3.4 Adding a Host to a Host Group 123
6.3.5 Removing a Host from a Host Group 123
6.3.6 Restoring a Host Group 123
6.3.7 Restoring all Host Groups 123
6.3.8 Exporting Host Grouping Configurations 123
6.4 Searching for Hosts in the Discovery-Tree 124

Chapter 7: Managing Adapters and Ports 125

7.1 Viewing Adapter Information 125
7.1.1 Viewing Channel Management Information (OCe11000-Series Adapters Only) 127
7.1.2 Viewing ASIC Information 130
7.1.3 Viewing Multi-ASIC Information 131
7.1.4 Viewing and Installing Licenses for OCe11000-Series Adapters 133
7.1.5 Changing Personalities on OneConnect OCe11000-Series Adapters 135
7.1.6 Viewing Firmware Information 136
7.1.7 Viewing Physical Port Information 138
7.1.8 Viewing PHY Data (OneConnect 10GBASE-T series Adapters Only) 139
7.1.9 Viewing Transceiver Information 141
7.2 OCe14000-Series Adapter Configuration 142
7.2.1 Configuring Single Personalities 144
7.2.2 Custom Configurations 146
7.2.3 UMC Configuration View 149
7.2.4 Lenovo System X Multichannel Configuration View 150
7.2.5 Link Reconfigurability 156
7.2.6 Dell NPar Configuration View 161
7.2.7 Configuring RoCE on NPar Adapters 164
7.2.8 Removing a Port 170
7.2.9 Configuring HP Flex Adapters 174

Chapter 8: Managing Protocols 177

8.1 NIC 177
8.1.1 Viewing NIC Function Information 177
8.1.2 Viewing NIC VPD Information 179
8.1.3 Enabling and Disabling SR-IOV on NIC Functions (on OCe11000-series Adapters) 180
8.1.4 Enabling and Disabling VEPA on NIC SR-IOV Functions (OCe14000-series Adapters Only) 181
8.1.5 Guest Operating System Discovery and Management from the Base Host Operating System 182
8.1.6 Viewing NIC PCI Registers 184
8.2 FCoE 184
8.2.1 Viewing FCoE Function Information 184
8.2.2 Viewing FCoE Statistics 186
8.2.3 Viewing FCoE Virtual Port Information 188
8.2.4 Creating and Deleting FCoE Virtual Ports 189
8.2.5 Viewing FCoE Fabric Information 192
8.2.6 Viewing FCoE VPD Information 193
8.2.7 Viewing FCoE Maintenance Information 194
8.2.8 Viewing FCoE Target Information 196
8.2.9 Viewing FCoE LUN Information 197
8.2.10 Viewing FCoE Target Mapping 198
8.2.11 Using Automapping and Persistent Binding (Windows Only) 200
8.2.12 Masking and Unmasking LUNs (Windows Only) 202
8.2.13 Changing FCoE WWPN and WWNN 203
8.2.14 Changing FCoE Names 206
8.2.15 Resetting FCoE Functions 207
8.2.16 Configuring FCoE Driver Parameters 208
8.2.17 Creating a Batch Mode Driver Parameters File 213
8.2.18 Configuring FCoE Initialization Protocol (FIP) 215
8.2.19 Configuring Boot from an FCoE SAN 216

8.2.20 Configuring Advanced Settings (Boot from SAN) ... 219
8.3 iSCSI ... 222
8.3.1 Viewing iSCSI Port Information ... 222
8.3.2 Viewing iSCSI Network Information ... 224
8.3.3 Viewing iSCSI PCI Registers ... 226
8.3.4 Viewing iSCSI VPD Information ... 227
8.3.5 Viewing iSCSI Statistics ... 228
8.3.6 Configuring iSCSI Target Discovery ... 230
8.3.7 Viewing iSCSI Target Information ... 238
8.3.8 Viewing iSCSI LUN Information ... 239
8.3.9 Configuring iSCSI Initiator Login Options ... 240

Chapter 9: Updating Adapter Firmware ... 243

9.1 Updating Firmware for a Single Adapter ... 243
9.2 Updating Firmware for Multiple Adapters ... 245

Chapter 10: DCB Configuration ... 247

10.1 Configuring DCB Parameters for NIC Only Adapter Ports ... 247
10.1.1 Configuring DCB Parameters for NIC+iSCSI Adapter Ports ... 251
10.1.2 Configuring DCB Parameters for NIC+FCoE Adapter Ports ... 255
10.1.3 Configuring DCB Parameters for NIC+FCoE+iSCSI Adapter Ports ... 258
10.1.4 Configuring DCB Parameters for RoCE Adapter Ports ... 262

Chapter 11: Exporting SAN Information ... 267

11.1 Creating a SAN Report ... 267

Chapter 12: Diagnostics ... 268

12.1 Loopback Test Combinations ... 269
12.1.1 Error Action ... 269
12.1.2 Test Cycles ... 269
12.1.3 Test Pattern ... 269
12.1.4 Test Status ... 269
12.2 Using Beaconing ... 270
12.3 Running TDR Tests (10GBASE-T Adapters Only) ... 271
12.4 Saving the Log File ... 272
12.5 Creating Diagnostic Dumps ... 272

Chapter 13: Troubleshooting ... 275

13.1 General Situations ... 275
13.2 Emulex Driver for Linux and OneCommand CNA Manager Application Situations ... 277
13.3 Emulex Driver for Solaris and OneCommand CNA Manager Application Situations ... 280
13.4 VPorts and OneCommand CNA Manager Application Situations ... 280

Appendix A: License Notices ... 281

A.1 Secure Hash Algorithm (SHA-1) Notice ... 281
A.2 OpenPegasus Licensing Notice ... 281
A.3 OpenSSL Notice ... 282
A.4 Java Notice ... 285

Emulex OneCommand CNA Manager Command Line Interface for OneConnect Adapters User Guide ... 287

Chapter 1: Introduction ... 289

1.1 Abbreviations ... 289
1.2 OneCommand CNA Manager Secure Management ... 292
1.2.1 OneCommand CNA Manager Secure Management Configuration Requirements ... 293
1.3 Secure Management Installation ... 293
1.3.1 Linux and Solaris ... 294

1.3.2 Windows 294
1.4 Setting Secure Management Mode for Linux and Solaris 294
1.4.1 Using OneCommand CNA Manager with Secure Management Enabled 295

Chapter 2: Installing and Uninstalling the CLI 296

2.1 Linux 296
2.1.1 Citrix 296
2.1.2 Installing in Linux without an Existing OneCommand CLI Kit 296
2.1.3 Installing in Linux with an Existing OneCommand CLI Kit 297
2.1.4 Uninstalling in Linux 298
2.1.5 Uninstalling Older HBAnyware Kits in Linux 299
2.2 Solaris 299
2.2.1 Installing in Solaris 299
2.2.2 Uninstalling in Solaris 300
2.3 VMware ESXi 301
2.4 Windows 301
2.4.1 Installing in Windows by Attended Installation 301
2.4.2 Installing in Windows by Unattended Installation 301
2.4.3 Uninstalling in Windows 303
2.5 Starting and Stopping Daemon Processes for Linux and Solaris Installations 303

Chapter 3: Updating to the OneCommand CNA Manager Application Enterprise Kit 304

3.1 Linux 304
3.2 Solaris 304
3.3 Windows 304

Chapter 4: CLI Client Command Usage 305

4.1 Overview 305
4.1.1 CLI in Read-Only Mode 305
4.2 BrcmHbaCmd Syntax Usage 305
4.3 Secure Management CLI Interface 306
4.3.1 Device Management Using the Secure Management Interface 306
4.3.2 Syntax Rules for the Secure Management Interface 306
4.4 CIM Client Interface 307
4.4.1 Device Management Using the CIM Interface 307
4.4.2 Syntax Rules for the CIM Interface 307

Chapter 5: CLI Client Command Descriptions 309

5.1 Help 317
5.2 Adapter License Management Commands 317
5.2.1 InstallAdapterLicense 318
5.2.2 ShowAdapterLicenseFeatures 318
5.2.3 ShowLicenseAdapterID 319
5.3 Attributes Commands 319
5.3.1 HbaAttributes 319
5.3.2 PortAttributes 320
5.3.3 PortStatistics 321
5.3.4 ServerAttributes 322
5.3.5 SetPhyPortSpeed 322
5.3.6 SetPortEnabled 323
5.4 Boot Commands 324
5.4.1 EnableBootCode 324
5.4.2 GetBootParams 325
5.4.3 SetBootParam 325
5.5 Channel Management Commands 326
5.5.1 CMGetParams 326
5.5.2 CMMode 329
5.5.3 CMSetBW 330

5.5.4 CMSetLPVID 331
5.6 DCB Commands 332
5.6.1 GetDCBParams 332
5.6.2 GetPGInfo 333
5.6.3 SetCnaPGBW 333
5.6.4 SetDCBParam 333
5.6.5 SetDCBPriority 335
5.7 Diagnostic Commands 336
5.7.1 GetBeacon 336
5.7.2 GetXcvrData 337
5.7.3 LoopBackTest 337
5.7.4 PciData 338
5.7.5 SetBeacon 339
5.7.6 SetCableNVP 339
5.7.7 TDRTest 339
5.8 Driver Parameter Commands 340
5.8.1 DriverConfig 340
5.8.2 GetDriverParams 341
5.8.3 GetDriverParamsGlobal 341
5.8.4 SaveConfig 341
5.8.5 SetDriverParam 342
5.8.6 SetDriverParamDefaults 342
5.9 Dump Commands 343
5.9.1 DeleteDumpFiles 343
5.9.2 Dump 343
5.9.3 GetDumpDirectory 344
5.9.4 GetDumpFile 344
5.9.5 GetDumpFileNames 345
5.9.6 GetRetentionCount 345
5.9.7 SetDumpDirectory 345
5.9.8 SetRetentionCount 346
5.10 FCoE Commands 346
5.10.1 GetFCFInfo 346
5.10.2 GetFIPParams 347
5.10.3 SetFIPParam 347
5.11 iSCSI Commands 348
5.11.1 AddARPTableEntry 349
5.11.2 AddiSNSServer 349
5.11.3 AddRouteTableEntry 349
5.11.4 AddTarget 350
5.11.5 AddTargetPortal 351
5.11.6 CleariSNSServer 352
5.11.7 DelARPTableEntry 352
5.11.8 DeleteiSNSServer 353
5.11.9 DelRouteTableEntry 353
5.11.10 DiscoveriSNSServer 353
5.11.11 EraseiSCSIConfig 354
5.11.12 ExportiSCSI 354
5.11.13 GetInitiatorProperties 354
5.11.14 GetiSCSILuns 355
5.11.15 GetiSCSIPortStats 355
5.11.16 GetNetworkConfiguration 355
5.11.17 GetSessionInfo 356
5.11.18 ImportiSCSI 356
5.11.19 iSCSIPing 357
5.11.20 ListSessions 357
5.11.21 RemoveTarget 358
5.11.22 RemoveTargetPortal 358
5.11.23 SetBootTargetSession 358

5.11.24 SetInitiatorProperties . . . 359
5.11.25 SetiSCSIBoot . . . 360
5.11.26 SetNetworkConfiguration . . . 360
5.11.27 SetTargetLoginProperties . . . 361
5.11.28 SetTargetProperties . . . 362
5.11.29 SetTPLoginProperties . . . 363
5.11.30 ShowARPTable . . . 364
5.11.31 ShowiSNSServer . . . 364
5.11.32 ShowRouteTable . . . 364
5.11.33 ShowTarget . . . 365
5.11.34 ShowTargetPortal . . . 365
5.11.35 TargetLogin . . . 365
5.11.36 TargetLogout . . . 366
5.11.37 UpdateiSNSServer . . . 367
5.12 LUN Masking Commands . . . 367
5.12.1 GetLunList . . . 368
5.12.2 GetLunUnMaskByHBA . . . 368
5.12.3 GetLunUnMaskByTarget . . . 368
5.12.4 RescanLuns . . . 369
5.12.5 SetLunMask . . . 369
5.13 Miscellaneous Commands . . . 370
5.13.1 AddHost . . . 370
5.13.2 CnaClearEventLog . . . 370
5.13.3 CnaGetEventLog . . . 371
5.13.4 Download . . . 371
5.13.5 ExportSANInfo . . . 372
5.13.6 GetCimCred . . . 372
5.13.7 GetQoSInfo . . . 372
5.13.8 GetVPD . . . 373
5.13.9 ListHBAs . . . 373
5.13.10 ListVFuncs . . . 374
5.13.11 RemoveHost . . . 374
5.13.12 Reset . . . 375
5.13.13 SetCimCred . . . 375
5.13.14 SRIOVEnable . . . 376
5.13.15 TargetMapping . . . 377
5.13.16 VEPAEnable . . . 377
5.13.17 Version . . . 378
5.14 Persistent Binding Commands . . . 378
5.14.1 AllNodeInfo . . . 379
5.14.2 BindingCapabilities . . . 379
5.14.3 BindingSupport . . . 379
5.14.4 PersistentBinding . . . 379
5.14.5 RemoveAllPersistentBinding . . . 380
5.14.6 RemovePersistentBinding . . . 380
5.14.7 SetBindingSupport . . . 381
5.14.8 SetPersistentBinding . . . 381
5.15 Personality Change Commands . . . 382
5.15.1 ChangePersonality . . . 382
5.15.2 ShowPersonalities . . . 383
5.16 Profile Management Commands . . . 384
5.16.1 GetAdapterPortConfig . . . 384
5.16.2 GetLinkConfig . . . 392
5.16.3 ListProfiles . . . 393
5.16.4 SetAdapterPortConfig (for OCe14000-Series Adapters) . . . 393
5.16.5 Lenovo System X Multichannel for OCe14000-Series Adapters . . . 402
5.16.6 Dell NPar Configurations . . . 402
5.16.7 SetLinkConfig . . . 413
5.17 UMC Commands . . . 413

5.17.1 UmcEnable ... 414
5.17.2 UmcGetParams ... 414
5.17.3 UmcSetBW ... 416
5.17.4 UmcSetLPVID ... 417
5.18 vPort Commands ... 418
5.18.1 CreateVPort ... 418
5.18.2 DeleteVPort ... 418
5.18.3 ListVPorts ... 419
5.18.4 VPortTargets ... 419
5.19 WWN Management Commands ... 419
5.19.1 ChangeWWN ... 419
5.19.2 GetWWNCap ... 420
5.19.3 ReadWWN ... 420
5.19.4 RestoreWWN ... 421

Appendix A: OneCommand CNA Manager Error Messages ... 422

Appendix B: License Notices ... 425
B.1 Secure Hash Algorithm (SHA-1) Notice ... 425
B.2 OpenPegasus Licensing Notice ... 425
B.3 OpenSSL Notice ... 426

Emulex Brcmflash and LpCfg Utilities for OneConnect Adapters User Guide ... 430

Chapter 1: Introduction ... 432
1.1 Brcmflash Offline Kit ... 432
1.2 Brcmflash Utility ... 432
1.3 Command Syntax ... 433
1.4 Supported Platforms ... 433
1.5 Abbreviations ... 434

Chapter 2: Installation ... 435
2.1 Platform Prerequisites ... 435
2.1.1 Linux ... 435
2.1.2 FreeBSD ... 435
2.2 Installing the Brcmflash Offline Kit for Linux ... 435
2.2.1 Updating ... 436
2.2.2 Uninstalling ... 436
2.3 Installing the Brcmflash Offline Kit for Windows PE ... 436
2.3.1 Updating ... 437
2.3.2 Uninstalling ... 437
2.4 Extracting the Driver Files from the Brcmflash Offline Kit for Windows ... 437
2.5 Extracting the Brcmflash Standalone Kit for Linux, Windows, and VMware ... 437
2.6 Installing the Brcmflash Offline Kit for FreeBSD ... 438
2.6.1 Updating ... 438
2.6.2 Uninstalling ... 439
2.7 Firmware and Boot Code ... 439

Chapter 3: Brcmflash Utility Command Line Interface ... 440
3.1 Auto-Discovery (/auto) ... 440
3.2 Downgrade (/downgrade or /g) ... 441
3.2.1 FCoE Adapter Download Summary ... 441
3.2.2 iSCSI and NIC-only Adapter Summary ... 441
3.3 Force Firmware and Boot Code (/f) ... 442
3.3.1 FCoE Adapter Download Summary ... 442
3.3.2 iSCSI and NIC-only Adapter Summary ... 442
3.4 Force Firmware (/ff) ... 443
3.4.1 FCoE Adapter Download Summary ... 443

3.4.2 iSCSI and NIC-only Adapter Summary ... 443
3.5 Firmware Matrix Directory Change (/fmd) ... 444
3.6 Firmware Flash Override (/fo) ... 444
3.7 Help (/h or /?) ... 444
3.8 Image Directory Change (/id) ... 444
3.9 Image Version Display (/iv) ... 445
3.9.1 Log(/log) ... 445
3.10 Preview (/p) ... 445
3.10.1 FCoE Adapter Download Summary ... 446
3.10.2 iSCSI and NIC-only Adapter Download Summary ... 446
3.11 Query (/q) ... 446
3.11.1 FCoE Adapter Status Summary ... 446
3.11.2 iSCSI and NIC Adapter Status Summary ... 446
3.12 Ramdrive (/ramdrive) - (Windows Offline Adapter Management Utility only) ... 447
3.13 Rewrite (/rewrite or /e) ... 447
3.13.1 FCoE Adapter Download Summary ... 447
3.13.2 iSCSI and NIC-only Adapter Download Summary ... 447
3.14 Silent (/s) ... 448
3.15 Discover Adapters Using SysFS (/sysfs) ... 448
3.16 Update (/update) ... 449
3.16.1 FCoE Adapter Download Summary ... 449
3.16.2 iSCSI and NIC Adapter Download Summary ... 449
3.17 Verbose (/v) ... 450
3.17.1 FCoE Adapter Download Summary ... 450
3.17.2 iSCSI and NIC-only Adapter Download Summary ... 450
3.18 VPD (/vpd) ... 450
3.19 XML Output (/xml) ... 451

Chapter 4: LpCfg Utility Command Line Interface ... 452

4.1 Running the LpCfg Utility from the Command Prompt ... 452
4.2 Running the LpCfg Utility Included with the Brcmflash Standalone Kit ... 453
4.2.1 Windows ... 453
4.2.2 Linux and VMware ... 454
4.3 Supported Commands ... 455
4.3.1 Help Command and General Adapter Management Commands ... 457
4.3.2 Viewing Adapter Information ... 458
4.3.3 Firmware and Boot Code Download Commands ... 461
4.3.4 World Wide Name Commands ... 461
4.3.5 Boot Code Commands ... 464
4.3.6 Configuration Commands ... 466
4.3.7 Personality Commands ... 467
4.3.8 Diagnostic Tests ... 467
4.3.9 Using Script Files ... 470
4.3.10 VLAN Commands ... 475

Chapter 5: Status Messages and Error Codes ... 477

5.1 Brcmflash Utility Status Messages ... 477
5.2 LpCfg Error Codes ... 478

Chapter 6: Troubleshooting ... 481

6.1 Unsupported Driver ... 481

Emulex Drivers for Windows for OneConnect Adapters Quick Installation Guide ... 483

Chapter 1: Introduction ... 485

Chapter 2: Driver Installation Options ... 486

2.1 OneInstall Installer ... 487

2.1.1 Loading the OneInstall Package Using Interactive Mode ... 487
2.1.2 Loading the OneInstall Package Using Silent Mode ... 488
2.2 Driver Kit Installer ... 489
2.2.1 Loading the Driver Kit ... 489
2.3 AutoPilot Installer ... 490
2.3.1 Starting Installers from a Command Prompt or Script ... 491
2.3.2 Running a Software Installation Interactively ... 491
2.3.3 Hardware-First Installation or Driver Update ... 492
2.3.4 Software-First Installation ... 493
2.3.5 Text-Only Driver Installation ... 493
2.4 Installing the Emulex iSCSI Crash Dump Driver (OCe14000-Series Adapters Only) ... 494
2.4.1 Supported Configurations ... 494

Chapter 3: Installation Failure ... 496

3.1 AutoPilot Installer Failure ... 496
3.2 OneInstall Installer Failure ... 496

Emulex Drivers for Windows for OneConnect Adapters User Guide ... 498

Chapter 1: Introduction ... 500

1.1 Driver Information ... 500
1.2 Abbreviations ... 500

Chapter 2: Installation ... 504

2.1 OneInstall Installer ... 505
2.1.1 Loading the OneInstall Package Using Interactive Mode ... 505
2.1.2 Loading the OneInstall Package Using Silent Mode ... 505
2.2 Driver Kit Installer ... 507
2.2.1 Loading the Driver Kit ... 507
2.3 AutoPilot Installer ... 507
2.3.1 Starting Installers from a Command Prompt or Script ... 508
2.3.2 Running a Software Installation Interactively ... 509
2.3.3 Hardware-First Installation or Driver Update ... 510
2.3.4 Software-First Installation ... 511
2.3.5 Text-Only Driver Installation ... 511
2.3.6 Unattended Driver Installation ... 511
2.3.7 Installation Failure ... 512
2.4 Manually Installing or Updating the Emulex Protocol Drivers ... 513
2.5 Installing the Emulex PLUS (ElxPlus) Driver for the First Time ... 513
2.5.1 Updating the Emulex PLUS (ElxPlus) Driver ... 514
2.5.2 Installing or Updating the FCoE Storport Miniport Driver ... 514
2.5.3 Installing or Updating the iSCSI Driver ... 515
2.5.4 Installing or Updating the NIC Driver ... 515
2.5.5 Supported Configurations ... 517
2.5.6 Adding Emulex Out-of-Box Drivers to a Nano Server Virtual Hard Disk ... 518
2.5.7 Adding Emulex OOB Drivers to running Nano Server with Pnputil.exe ... 518
2.6 Removing Emulex Driver Kits and Drivers ... 519
2.6.1 Uninstalling Emulex Driver Kits ... 519
2.6.2 Uninstalling the Emulex Drivers ... 520

Chapter 3: Configuration ... 522

3.1 FCoE Driver Configuration ... 522
3.1.1 Configuring FCoE Driver Parameters ... 522
3.2 NIC Driver Configuration ... 526
3.2.1 Configuring NIC Driver Options ... 526
3.2.2 Configuring Windows Server NIC Driver Parameters ... 537
3.2.3 Using OCCFG for Windows NIC Driver Options ... 542
3.2.4 Extracting the occfg.exe Utility on Nano Server ... 548

3.2.5 Using SR-IOV with Emulex Devices ... 549
3.2.6 Configuring NVGRE for the OCe14000-series Adapters ... 557
3.2.7 Configuring RoCE for the OCe14000-Series Adapters ... 564
3.2.8 Configuring UMC ... 579
3.2.9 Network Driver Performance Tuning ... 581
3.3 iSCSI Driver Configuration ... 588
3.3.1 Configuring iSCSI Driver Options ... 588
3.3.2 Interrupt Moderation Policy Settings ... 590
3.3.3 Creating Non-Bootable Targets ... 591
3.3.4 Maximum Transmission Unit (MTU) for iSCSI Connections ... 593

Chapter 4: Troubleshooting ... 595

4.1 General Troubleshooting ... 595
4.2 Troubleshooting the FCoE Driver ... 596
4.2.1 Event Trace Messages ... 596
4.3 Troubleshooting the NIC Drivers ... 609
4.3.1 Monitoring TCP Offloads ... 610
4.3.2 TCP Offload Failure ... 611
4.4 Troubleshooting the iSCSI Driver ... 611
4.4.1 Troubleshooting the Cisco Nexus Switch Configuration ... 611
4.4.2 iSCSI Driver Troubleshooting ... 612

Appendix A: Error and Event Log Information ... 614

A.1 FC/FCoE Error and Event Logs ... 614
A.1.1 Viewing the FC/FCoE Error Log ... 614
A.1.2 Viewing the FCoE Event Log ... 619
A.2 NIC Error and Event Logs ... 623
A.2.1 Viewing the NIC Error Log ... 623
A.2.2 RoCE Event Log ... 623
A.2.3 NIC Event Log ... 624
A.3 iSCSI Error and Event Log ... 627
A.3.1 Viewing the iSCSI Error and Event Log on Windows Server 2008 ... 627
A.3.2 iSCSI Error Log on Windows Server 2008 ... 629
A.3.3 Viewing the iSCSI Error Log on Windows 7, Windows 8, Windows 8.1, Windows Server 2008 R2, Windows Server 2012, and Window Server 2012 R2 ... 632

Appendix B: Configuring iSCSI through DHCP ... 639

B.1 Dynamic Host Configuration Protocol (DHCP) Recommendations ... 639
B.2 Vendor-Specific Option 43 ... 639
B.2.1 Format of Vendor-Specific Option 43 ... 639
B.2.2 Description of Mandatory and Optional Parameters ... 639

Appendix C: Port Speed Specifications ... 641

C.1 Negotiating Speed on a Mezzanine Card ... 641

Appendix D: AutoPilot Installer Command Line and Configuration File Parameters ... 642

D.1 AParg Driver Kit Parameter and Appending to the BrcmAPInstall.exe File ... 642
D.2 AutoPilot Installer Syntax ... 643
D.2.1 Path Specifiers ... 643
D.2.2 Configuration File Location ... 643
D.2.3 Software Configuration Parameters ... 644
D.3 AutoPilot Configuration File ... 646
D.3.1 Using the Windows Environment Variable (%ProgramFiles%) ... 646
D.3.2 Configuration Identification [AUTOPILOT.ID] ... 646
D.3.3 Software Configuration [AUTOPILOT.CONFIG] ... 646
D.3.4 Configuration Prompts/Vendor-Specific Questions [STORPORT.CONFIGURATION] ... 646
D.3.5 QFE Checks [STORPORT.QFES] ... 648
D.3.6 Setting Up FC Driver Parameters [STORPORT.PARAMS] ... 648

D.3.7 Setting Up System Parameters [SYSTEM.PARAMS] 648
D.4 AutoPilot Installer Exit Codes 649
D.5 AutoPilot Installer Installation Reports 649
D.6 Command Script Example 650

Appendix E: RoCE Switch Support 652
E.1 DCBX-Enabled Switch Connection PFC Mode 652
E.1.1 Switch Configuration for PFC Priority 5 652
E.1.2 Host—Client Configuration 653

Appendix F: License Notices 656
F.1 Secure Hash Algorithm (SHA-1) Notice 656

Emulex Drivers for Linux for OneConnect Adapters User Guide 658

Chapter 1: Introduction 660
1.1 Overview 660
1.2 Abbreviations 660

Chapter 2: Installing and Uninstalling 664
2.1 General Installation Requirements 664
2.2 Binary RPM FCoE Driver Kit 664
2.2.1 Installing the Binary RPM FCoE Driver Kit 664
2.2.2 Uninstalling the Binary RPM FCoE Driver Kit 665
2.3 Ethernet Driver Kit 665
2.3.1 Installing the Ethernet Driver Kit 665
2.3.2 Uninstalling the Ethernet Driver Kit 666
2.4 RoCE Driver for the OCe14000-Series Adapters 666
2.4.1 Installing the Infiniband Stack 666
2.4.2 Uninstalling the RoCE Driver 669
2.5 iSCSI Driver Kit 669
2.5.1 Compiling the iSCSI Driver 669
2.5.2 Installing the iSCSI Driver Kit 670
2.5.3 Uninstalling the iSCSI Driver Kit 670
2.6 Booting from a Nonzero LUN Attached to a Emulex FCoE Adapter 671
2.7 Emulex OneCommand CNA Manager Application 671
2.8 Updating the Adapter Firmware 671
2.8.1 Determining the Firmware Version 671
2.8.2 Updating the Firmware Manually 671

Chapter 3: Configuration 673
3.1 FCoE Driver Configuration 673
3.1.1 FCoE Driver Parameters 673
3.1.2 Configuring FCoE Driver Parameters 675
3.1.3 Creating a New Ramdisk Image 678
3.1.4 Dynamically Recognizing LUNs and Targets (Using scan) 678
3.1.5 Persistent Naming 678
3.2 Ethernet Driver Configuration 682
3.2.1 Ethernet Driver Configuration Parameters 682
3.2.2 VLAN Support on UMC 682
3.2.3 Support for Ethtool set-channels and get-channels Commands 682
3.2.4 Transmit/Receive Queue Counts 683
3.2.5 Support for Ethtool set-rxfh/get-rxfh Commands 684
3.2.6 Support for Showing Onboard Die Temperature 684
3.2.7 SR-IOV Configuration 685
3.2.8 Bonding Considerations 693
3.2.9 ARI Considerations for OCe14000-Series Adapters 693
3.2.10 NIC Partitioning (NPar) Configuration (Dell Only) 694

3.3 RoCE Configuration for OCe14000-Series Adapters ... 695
3.3.1 Basic Configuration ... 695
3.3.2 Pinning Interrupts to a CPU Core ... 698
3.3.3 MTU Configuration ... 698
3.3.4 QoS Configuration ... 698
3.3.5 Advanced Applications ... 702
3.3.6 Installing and Using DAPL ... 704
3.4 iSCSI Driver Configuration ... 705
3.4.1 Open-iSCSI Support ... 705
3.4.2 Discovering and Adding iSCSI Targets ... 705
3.4.3 iscsiadm Commands for Configuring the Target ... 707
3.4.4 iSNS Discovery Using iscsiadm ... 708
3.5 VPort Configuration ... 709
3.5.1 VPort Configuration Prerequisites ... 709
3.5.2 Creating, Deleting, and Displaying VPorts ... 709
3.5.3 VPort sysfs ... 710
3.5.4 Monitoring VPorts with fc_vport (FCoE Drivers) ... 713
3.5.5 VPort Configuration Limits ... 714
3.6 FCoE Driver Performance Tuning ... 714
3.6.1 Overview ... 714
3.7 Network Performance Tuning ... 716
3.7.1 Memory Bandwidth Considerations ... 716
3.7.2 Network Memory Limits ... 717
3.7.3 TCP Segmentation Offload (TSO) ... 717
3.7.4 Flow Control ... 717
3.7.5 RX Frame Coalescing ... 718
3.7.6 Maximum Transmission Unit (MTU) ... 719
3.7.7 Interrupt Coalescing ... 719
3.7.8 Receive-Side Scaling (RSS) ... 719
3.7.9 Analyzing Performance Issues ... 720

Chapter 4: Troubleshooting ... 721

4.1 Hardware Situations and Resolutions ... 721
4.2 FCoE Driver Situations and their Resolutions ... 722
4.3 Ethernet Driver Situations and their Resolutions ... 724
4.4 iSCSI Driver Situations and their Resolutions ... 725
4.5 Log Messages ... 725
4.5.1 FCoE Driver Log Messages ... 725
4.5.2 Ethernet Driver Log Messages ... 792
4.5.3 RoCE Error Log Messages ... 797
4.5.4 iSCSI Driver Log Messages ... 800

Appendix A: iSCSI Driver ... 802

A.1 iSCSI Error Handling ... 802
A.1.1 Error Handling Using ETO and LDTO Parameters ... 802
A.1.2 Error Handling Under MPIO and Cluster Configurations ... 802
A.2 iSCSI Driver Log Messages ... 803
A.2.1 Retrieving iSCSI Driver Error Log Messages ... 803
A.2.2 iSCSI Driver Error Log Messages and their Descriptions ... 803

Appendix B: Configuring iSCSI Through DHCP ... 805

B.1 IP Address Reservation ... 805
B.2 DHCP Option 43 (Vendor-Specific Information) ... 805
B.2.1 DHCP Option 43 Parameter Descriptions ... 805
B.2.2 DHCP Option 43 Format ... 806

Appendix C: OneConnect 10GbE Adapter Port Speed Specifications ... 808

C.1 Port Speed Negotiation on Non-Mezzanine Cards ... 808

C.2 Port Speed on Mezzanine Cards . . . 808

Appendix D: Updating Ethernet Firmware . . . 809
D.1 Ethtool Support for Secure Firmware Update . . . 809

Appendix E: Ethtool -S Option Statistics . . . 810

Appendix F: brcmfcoe Driver BlockGuard Functionality . . . 812
F.1 Overview . . . 812
F.2 Enabling BlockGuard . . . 813
F.3 SCSI Command Support – SCSI Operation Codes . . . 813
F.4 brcmfcoe Driver Module Parameters . . . 814
F.4.1 lpfc_prot_mask . . . 814
F.4.2 lpfc_prot_guard . . . 815

Appendix G: Setting Up and Configuring VXLAN . . . 816
G.1 Setting up VXLAN Networks Across Two Systems . . . 816

Appendix H: License Notices . . . 818
H.1 Open-iSCSI GPLv2 Notice . . . 818

Emulex Poll Mode Driver for OneConnect Adapters User Guide . . . 824

Chapter 1: Overview . . . 826
1.1 Purpose . . . 826
1.2 Features . . . 826
1.3 References . . . 827
1.4 Abbreviations . . . 827

Chapter 2: Configuring the Emulex PMD . . . 828
2.1 Basic Setup . . . 828
2.1.1 Download the Required Files . . . 828
2.1.2 Flash the Firmware . . . 828
2.1.3 Install Linux Ethernet (be2net) Source Driver Files . . . 828
2.2 Build and Copy the Binaries . . . 829
2.2.1 Build the DPDK Binaries . . . 829
2.2.2 Build the Emulex SURF and PMD Binaries . . . 829
2.2.3 Copy the Binaries . . . 829
2.3 Parameters for the dpdk_surf.ko Module . . . 830
2.3.1 dpdk_oce_trace <int> . . . 830
2.3.2 lldp_mode <int> . . . 830
2.3.3 max_dpdk_oce_port <int> . . . 830
2.3.4 max_rxq_per_port <int> . . . 830
2.4 Creating the load_drivers and run_testpmd Scripts . . . 831
2.4.1 Creating the load_drivers Script . . . 831
2.4.2 Creating the run_testpmd Script . . . 833
2.5 Running a Test . . . 833

Chapter 3: Limitations . . . 837
3.1 Multichannel Limitations . . . 837
3.2 SR-IOV Limitations . . . 837
3.3 RHEL 7.2 with DPDK 2.0 Limitations . . . 837

Chapter 4: Tuning and Operational Considerations . . . 838
4.1 Tunable Parameters Using DPDK Applications . . . 838
4.2 Recommended RX/TX Queues . . . 838
4.3 Receive and Transmit Tunings . . . 838
4.3.1 Tuning the Receive Function of the Emulex PMD . . . 838

4.3.2 Tuning the Transmit Function of the Emulex PMD 838
4.4 Disabling VLAN Tag Stripping 840
4.5 Disabling LLDP 840

Emulex Drivers for VMware ESXi for OneConnect Adapters User Guide 842

Chapter 1: Introduction 844

1.1 ESXi Compatibility 844
1.2 Supported Guest Operating Systems 844
1.3 Abbreviations 846

Chapter 2: Installation 849

2.1 Overview 849
2.1.1 Migrating Driver Parameters When Upgrading to ESXi 5.5 or Later 849
2.1.2 Preparing to Upgrade to ESXi 5.5 or Later 850
2.2 Installing the FCoE Driver 851
2.2.1 Installing the FCoE Driver and Management Software 851
2.2.2 Uninstalling the FCoE Driver 851
2.2.3 Installing the Native mode FCoE Driver esxcli Plug-in 851
2.3 Installing the NIC Driver 852
2.3.1 Installing the NIC Driver and Management Software 853
2.3.2 Uninstalling the NIC Driver 853
2.3.3 Installing the Native Mode Ethernet Driver esxcli Plug-in 853
2.3.4 Virtualization Support 854
2.3.5 Obtaining Information on the Installed NIC Driver 854
2.4 Installing the iSCSI Driver 854
2.4.1 Installing the iSCSI Driver and Management Software 854
2.5 Updating Drivers with VMware Offline Bundle Media 855

Chapter 3: Configuration 856

3.1 ESXi Command Line Tool Transition 856
3.1.1 FCoE Driver Examples 856
3.1.2 Ethernet Driver Examples 856
3.2 FCoE Driver Configuration 857
3.2.1 FCoE Driver Parameters Configuration Methods 858
3.2.2 Dynamically Adding LUNs 859
3.2.3 Dynamically Adding Targets 860
3.2.4 FCoE Driver Configuration Parameters 860
3.2.5 Creating an FC Remote Boot Disk 865
3.2.6 Managing Devices through the CIM Interface 865
3.2.7 Installing the Emulex CIM Provider 866
3.2.8 Creating, Deleting, and Displaying VPorts 866
3.2.9 Configuring VVols in ESXi 6.0 and 6.5 866
3.3 NIC Driver Configuration 867
3.3.1 NetQueue Support 868
3.3.2 Receive-Side Scaling 868
3.3.3 How an ESXi Server Creates and Names Interfaces 870
3.3.4 Enabling SR-IOV 871
3.3.5 Configuring Emulex Universal Multichannel 872
3.3.6 Configuring VLANs 873
3.3.7 Configuring NPar (Dell Only) 874
3.3.8 Performance Tuning 875
3.4 Emulex Drivers for VMware NSX 6.2 876
3.4.1 Module Parameters 876
3.4.2 Functionality 876
3.4.3 Usage 877
3.4.4 Limitations 878
3.4.5 Limitations Exclusive to OCe11000-Series Adapters 878

3.5 iSCSI Driver Configuration ... 878
3.5.1 Configuring iSCSI Driver Options ... 878
3.5.2 Interrupt Moderation Policy Settings ... 880
3.5.3 Configuring LDTO and ETO on ESXi Server ... 881
3.5.4 Multipath I/O Support ... 882
3.5.5 Reading the Driver Statistics for a Specified Port ... 883
3.5.6 Setting the Log Level ... 883

Chapter 4: Troubleshooting ... 884

4.1 Troubleshooting the FCoE Driver ... 884
4.1.1 FCoE Driver Log Messages ... 886
4.1.2 Extended Link Service Events (0100 to 0199) ... 888
4.1.3 Link Discovery Events (0200 to 0299) ... 891
4.1.4 Mailbox Events (0300 to 0339) ... 897
4.1.5 Initialization Events (0400 to 0599) ... 905
4.1.6 FCP Traffic History (0700 to 0799) ... 916
4.1.7 Node Table Events (0900 to 0999) ... 917
4.1.8 Security Events (1000 to 1099) ... 918
4.1.9 Miscellaneous and FCoE Events (1200 to 1299) ... 922
4.1.10 Link Events (1300 to 1399) ... 925
4.1.11 Port Setup Events (1400 to 1499) ... 926
4.1.12 IOCTL Events (1600 to 1699) ... 930
4.1.13 VPort Events (1800 to 1832) ... 930
4.1.14 Extended Link Service Events (1833 to 2800) ... 934
4.1.15 New Events ... 951
4.2 Troubleshooting the NIC Driver ... 963
4.2.1 ESXi Server NIC Event Log Entries ... 964
4.2.2 NIC Adapter Firmware Error ... 965
4.2.3 NIC Informational Log Groups ... 965
4.3 Troubleshooting the iSCSI Driver ... 967
4.3.1 Retrieving iSCSI Error Log Codes ... 967
4.3.2 iSCSI Error Log Code Entries ... 967
4.3.3 Additional iSCSI Driver Messages ... 969
4.4 Native Mode NIC Driver Troubleshooting Support ... 971
4.5 Native Mode NIC Driver Support for ethtool Commands ... 972

Appendix A: Configuring iSCSI Through DHCP ... 974

A.1 Recommendations ... 974
A.1.1 Vendor-Specific Option 43 ... 974

Emulex CIM Provider Package for OneConnect Adapters Installation Guide ... 976

Chapter 1: Introduction ... 978

1.1 Supported Platforms ... 978
1.2 Supported CIM Provider Profiles ... 978
1.3 Abbreviations ... 979

Chapter 2: Installing the Emulex CIM Provider ... 980

2.1 Installing the VIB in the VMware Operating System ... 980
2.2 Using Adapters ... 980
2.2.1 Updating Firmware on Broadcom Emulex Adapters ... 980
2.2.2 Enabling Logs and Collecting Symptoms ... 981
2.2.3 Generating Provider Logs if Listing Them Fails ... 982
2.2.4 Running Diagnostic Tests ... 982
2.2.5 Discovering vNICs ... 982

Chapter 3: Troubleshooting ... 984

Emulex OneCommand CNA Manager for VMware vCenter for OneConnect Adapters User Guide 986

Chapter 1: Introduction 988

1.1 Compatibility 988
1.2 Abbreviations 989

Chapter 2: Installing and Enabling OneCommand CNA Manager for VMware vCenter 993

2.1 Best Installation Practices 993
2.2 Hardware Requirements 993
2.3 Software Requirements 993
2.4 Installing OneCommand CNA Manager for VMware vCenter 994
2.4.1 Verifying the OneCommand CNA Manager for VMware vCenter Installation 996
2.5 Activating and Enabling ESXi Management 998
2.5.1 Requirements 998
2.5.2 Lockdown Mode Feature 998
2.5.3 Activating OneCommand CNA Manager for VMware vCenter in the Web Client 998
2.5.4 Enabling OneCommand CNA Manager for VMware for vCSA 999
2.5.5 Enabling OneCommand CNA Manager for VMware vCenter Web Client 999
2.6 Enabling and Disabling OneCommand CNA Manager for VMware vCenter with the Plug-in Manager 1000
2.7 Registering and Unregistering OneCommand CNA Manager for VMware vCenter 1001
2.8 Uninstalling OneCommand CNA Manager for VMware vCenter 1001
2.9 Reinstalling OneCommand CNA Manager for VMware vCenter 1002

Chapter 3: Using OneCommand CNA Manager for VMware vCenter 1004

3.1 Accessing with VMware Web Client 6.0 and 6.5 1004
3.2 Accessing with VMware Web Client Version 5.5 1004
3.3 Accessing the vSphere Console 1005
3.3.1 vSphere Console Views 1005
3.3.2 OneCommand CNA Manager for VMware vCenter Window Elements 1006

Chapter 4: Managing Clusters and Hosts 1009

4.1 Managing Clusters 1009
4.1.1 Viewing Hosts in a Cluster 1009
4.1.2 Viewing Adapters in a Cluster 1010
4.1.3 Viewing Physical Ports in a Cluster (Host-centric mode) 1011
4.1.4 Viewing Virtual Ports in a Cluster (Host-centric mode) 1013
4.1.5 Viewing Physical Port Information in a Cluster (Fabric- centric mode) 1014
4.2 Managing Hosts 1015
4.2.1 Viewing Host Information for a Single Host 1016
4.2.2 Viewing Driver Parameters of All Adapters in a Host 1017
4.2.3 Viewing Firmware Information for All Adapters in a Host 1018

Chapter 5: Managing Adapters and Ports 1020

5.1 OneConnect Adapters 1020
5.1.1 Viewing OneConnect Adapter Information 1020
5.1.2 Viewing Firmware Information 1021
5.1.3 Viewing Physical Ports 1022
5.1.4 Viewing PHY Data for 10GBASE–T Adapters 1025
5.1.5 Viewing Port Transceiver Information 1026
5.2 OCe11100-Series Adapters 1027
5.2.1 Viewing Licenses 1027
5.2.2 Changing Personalities 1028
5.2.3 Configuring SR-IOV 1028
5.2.4 Managing Port Channels Using Channel Management 1029
5.3 OCe14000–Series Adapters 1033
5.3.1 Configuring OCe14000–Series Adapters 1033
5.3.2 Setting Up a UMC Configuration 1043
5.3.3 Setting Up an Lenovo System X Multichannel Configuration 1044

5.3.4 Setting Dell NPar Bandwidth Settings 1049
5.3.5 Setting Up HP Flex Configurations 1053
5.3.6 Removing a Port 1054

Chapter 6: Managing Protocols 1056

6.1 NIC 1056
6.1.1 Viewing NIC Port Information 1056
6.1.2 Viewing Virtual Functions on a Host Operating System 1059
6.1.3 Viewing PCI Registers for a NIC Port 1060
6.1.4 Viewing NIC Port Vital Product Data (VPD) 1060
6.2 FCoE 1061
6.2.1 Viewing FCoE Port Details 1061
6.2.2 Viewing FC/FCoE Port Statistics 1063
6.2.3 Viewing PCI Registers 1064
6.2.4 Viewing FCoE Port Maintenance and Firmware Information 1065
6.2.5 Configuring FCoE Port Driver Parameters 1069
6.2.6 Viewing FCoE Port Vital Product Data (VPD) 1071
6.2.7 Configuring FIP 1072
6.2.8 Viewing FCoE Port Transceiver Information 1073
6.2.9 Viewing FC/FCoE Target Information 1075
6.2.10 Viewing FCoE LUN Information 1075
6.3 iSCSI 1076
6.3.1 Viewing iSCSI Port Information 1076
6.3.2 Viewing PCI Registers for an iSCSI Port 1078
6.3.3 Viewing and Changing iSCSI ISNS Server Information 1079
6.3.4 Viewing iSCSI Port VPD 1080
6.3.5 Viewing iSCSI Initiator Statistics 1081
6.3.6 Viewing iSCSI Initiator Login Options 1083
6.3.7 Viewing iSCSI Target Discovery Information 1085
6.3.8 Viewing iSCSI Target Information 1086
6.3.9 Viewing iSCSI LUN Information 1089

Chapter 7: Updating Firmware 1091

7.1 Updating Firmware on a Single Adapter 1091
7.2 Updating Firmware on Multiple Adapters in a Host 1093
7.3 Performing a Batch Firmware Update in Cluster View 1095
7.3.1 Jobs Window 1095

Chapter 8: DCB Configuration 1097

8.1 Configuring DCB Parameters for NIC + FCoE Adapter Ports 1097
8.2 Configuring DCB Parameters for NIC + iSCSI Adapter Ports 1100
8.3 Configuring DCB Parameters for NIC + FCoE + iSCSI Adapter Ports 1104

Chapter 9: Exporting SAN Information in Cluster View 1108

9.1 Capturing SAN Information in XML or CSV Format 1110
9.2 Considerations When Exporting SAN Information in a Cluster View 1110

Chapter 10: Diagnostics 1111

10.1 Running Loopback Tests 1111
10.2 Using Beaconing 1112
10.3 Running a TDR Test (10GBASE–T–Series Adapters Only) 1112
10.4 Setting Up Diagnostic Test Options 1112
10.4.1 Setting Up a Test Failure Error Action 1113
10.4.2 Setting Up Test Cycles 1113
10.4.3 Setting Up a Test Pattern 1113
10.5 Test Status 1113
10.6 Saving the Log File 1113
10.7 Creating OneConnect Adapter Port Diagnostic Dumps 1114

10.7.1 Viewing Diagnostic Dump Files 1114

Chapter 11: Generating and Installing Secured Certificates 1117

11.1 Secure Sockets Layer (SSL) Certificate 1117
11.1.1 Generating an SSL Certificate 1117
11.1.2 Generating a Self-Signed Certificate 1117
11.2 Generating a Certificate Signing Request (CSR) 1118
11.2.1 Generating a CSR for a Server Using the Java Tool 1119
11.2.2 Generating and Validating a CSR 1119
11.2.3 Getting an SSL Certificate 1120
11.2.4 Installing the SSL into the Web Server 1120

Chapter 12: Troubleshooting 1122

12.1 Security 1122
12.1.1 Accepting the Blocked Content 1122
12.1.2 Installing a Security Certificate 1123
12.1.3 Issues and Resolutions 1125

Chapter 13: Using OneCommand CNA Manager for VMware vCenter Command Line Interface 1127

13.1 brcmvcpcmd Syntax Usage 1127
13.2 Help Commands 1128
13.2.1 help (single command) 1128
13.2.2 help (group) 1128
13.3 CLI Command Reference Tables 1128
13.4 Group Commands and CLI Command Descriptions 1134
13.4.1 General Group Commands 1134
13.4.2 AdapterLicense Commands 1136
13.4.3 Adapter Management Commands 1137
13.4.4 Attribute Commands 1140
13.4.5 Boot Command 1145
13.4.6 Channel Management Commands 1145
13.4.7 Cluster Commands 1149
13.4.8 DCB Commands 1151
13.4.9 Diagnostic Commands 1155
13.4.10 Driver Parameter Commands 1158
13.4.11 Dump Commands 1161
13.4.12 FCoE Commands 1164
13.4.13 Firmware Commands 1166
13.4.14 iSCSI Commands 1167
13.4.15 Link Reconfigurability Commands 1171
13.4.16 Personality Commands 1172
13.4.17 SR-IOV Commands 1173
13.4.18 VEPA Commands 1174
13.4.19 Virtual Machines Commands 1175
13.4.20 WWWN Management Commands 1176
13.5 Viewing Audit Logs Using the CLI Command 1177
13.5.1 Syntax 1178
13.5.2 Parameter 1178
13.5.3 List of Supported Event Names 1178
13.5.4 Example 1178

Appendix A: License Notices 1179

A.1 VI Java SDK 1179
A.2 Java Notice 1179

Emulex Boot for NIC, iSCSI, and FCoE Protocols User Guide 1182

Chapter 1: Introduction 1184

1.1 Abbreviation 1185

Chapter 2: Configuring the PXE Boot for NIC on OCe11100-Series Adapters 1188

2.1 Pre-OS 1188
2.2 PXE Boot Process 1188
2.2.1 Remotely Installing with the PXE for Windows Server 2008, 2008R2, 2012, 2012 R2, and Windows Server 2016 1189
2.3 Remotely Installing with the PXE for Linux and Citrix Servers 1190
2.4 Using the PXESelect Utility 1190
2.4.1 Navigating the PXESelect Utility 1190
2.4.2 Running the PXESelect Utility 1191
2.4.3 Setting Up a PXE Bootable Network 1191
2.4.4 Configuring Ports If UMC Is Disabled or Not Supported 1192
2.4.5 Configuring Universal Multichannel 1195
2.4.6 Configuring Multichannel for Lenovo System x Adapters 1197
2.4.7 Advanced Mode Support 1200
2.4.8 Personality Option 1201
2.5 PXE Boot Parameters Default Values 1201

Chapter 3: Configuring PXE Boot for NIC on OCe14000-Series Adapters 1203

3.1 Pre-OS 1203
3.2 PXE Boot Process 1203
3.3 Remotely Installing with PXE for Windows Server 2008, 2008 R2, 2012, 2012 R2, and 2016 1204
3.4 Remotely Installing with PXE for Linux and Citrix Servers 1205
3.5 Using the PXESelect Utility 1205
3.5.1 Navigating the PXESelect Utility 1205
3.5.2 Running the PXESelect Utility 1206
3.5.3 Setting Up a PXE Bootable Network 1206
3.5.4 Configuring Ports If UMC Is Disabled or Not Supported 1208
3.5.5 Configuring Universal Multichannel 1215
3.5.6 Configuring Multichannel for Lenovo System x Adapters 1220
3.6 PXE Boot Parameters Default Values 1224

Chapter 4: Configuring Boot from SAN for the FCoE Protocol 1226

4.1 Windows Server 1226
4.1.1 Configuring Boot from SAN on Windows (x64) 1226
4.1.2 Configuring Boot from SAN on Windows (UEFI) 1227
4.1.3 Installing a New Windows UEFI-Aware Operating System on a UEFI-Based x64 Server 1227
4.1.4 Directing a UEFI-Based Server to a Windows Server Operating System Image Already Installed on the SAN 1228
4.2 Installing Windows Server on a Boot Disk (x64) 1228
4.3 Linux, Citrix, and VMware Operating Systems 1229
4.3.1 Configuring Boot from SAN on Linux, Citrix, or VMware (x86 and x64) 1229
4.3.2 Configuring Boot from SAN on Linux or VMware (UEFI) 1229
4.4 Solaris Operating System 1229
4.4.1 Configuring Boot from SAN on Solaris (x86 and x64) 1230
4.4.2 Determining LUNs to Select for Boot from SAN 1230

Chapter 5: Configuring x86/x64 Platforms for the iSCSI Protocol 1231

5.1 Constructing a Basic iSCSI SAN 1231
5.2 Managing an iSCSI Session 1231
5.2.1 Logging in to an iSCSI Session 1231
5.2.2 Security 1232
5.3 Configuring Boot from SAN for iSCSI 1232
5.3.1 Setting Up Boot from SAN for iSCSI 1232
5.3.2 Booting from SAN for iSCSI MPIO 1233

Chapter 6: Using the FCoE Boot BIOS Utility for x86 and x64 Architectures 1236

6.1 Navigating the FCoE BIOS Utility 1236
6.2 Starting the FCoE BIOS Utility 1236

6.3 Enabling an Adapter to Boot from SAN 1238
6.4 Scanning for Target Devices 1238
6.5 Configuring Boot Devices 1239
6.6 Configuring FCoE FCF CEE Parameters 1243
6.7 Configuring Advanced Adapter Parameters 1244
6.7.1 Changing the Port Login Retry Timer 1245
6.7.2 Enabling or Disabling the Spinup Delay 1246
6.7.3 Setting Auto Scan 1247
6.7.4 Enabling or Disabling EDD 3.0 1248
6.7.5 Enabling or Disabling the Start Unit Command 1249
6.7.6 Enabling or Disabling the Environment Variable 1250
6.7.7 Enabling or Disabling Auto Boot Sector 1251
6.8 Using Multipath Boot from SAN 1252
6.9 Resetting to Default Values 1253

Chapter 7: Updating and Enabling Boot Code 1255

Chapter 8: Configuring iSCSI Boot Support with the iSCSISelect Utility 1256

8.1 Navigating the iSCSISelect Utility 1256
8.2 Setting Up a Basic iSCSI Boot Configuration 1256
8.3 Booting the System 1258

Chapter 9: Configuring and Managing the iSCSI Initiator with the iSCSISelect Utility 1259

9.1 Selecting an Adapter 1259
9.2 Viewing the Controller Properties 1260
9.2.1 Enabling Boot Support 1260
9.2.2 Enabling IPv4 Boot Target Discovery using DHCP 1260
9.2.3 Enabling the MPIO Port 1260
9.2.4 Configuring the iSCSI Initiator Name 1261
9.3 Configuring Network Properties 1262
9.3.1 Selecting the IP Version 1262
9.3.2 Configuring VLAN ID and VLAN Priority 1262
9.3.3 Configuring an IP Address 1263
9.4 Identifying a Port 1266
9.5 Erasing the Current Configuration 1267

Chapter 10: Configuring and Managing iSCSI Targets with the iSCSISelect Utility 1268

10.1 Adding iSCSI Targets 1268
10.1.1 Using the SendTargets Discovery to Add an iSCSI Target 1268
10.1.2 Manually Adding an iSCSI Target 1270
10.2 Managing an iSCSI Target 1271
10.2.1 Viewing Target Information 1271
10.2.2 Editing a Target Configuration 1271
10.2.3 Viewing Advanced Target Properties 1275
10.2.4 Configuring LUN Settings 1275
10.2.5 Logging In to and Out of a Configured Target 1276
10.2.6 Removing a Configured Target 1276
10.2.7 Booting the System 1276
10.3 Discovering Targets through DHCP for iSCSI 1277
10.4 Enabling DHCP Discovery through the iSCSISelect Utility 1278

Chapter 11: Configuring UEFI for Ethernet 1279

11.1 Navigating the Emulex NIC Configuration Utility 1279
11.2 Starting the Emulex NIC Configuration Utility 1280
11.2.1 Advanced Mode Support 1282
11.2.2 Configuring Boot Mode 1283
11.2.3 Configuring iBFT 1284
11.2.4 Configuring Personality 1290

11.2.5 Link Reconfiguration 1293
11.2.6 Configuring Virtual LAN Settings 1294
11.2.7 Configuring the Adapter If UMC Is Disabled or Not Supported 1297
11.2.8 Configuring Universal Multichannel 1301
11.2.9 Configuring Multichannel for Lenovo System x Adapters 1306
11.2.10 Port Management 1311
11.2.11 Feature on Demand 1312
11.3 Downloading Firmware and Boot Code 1313
11.4 Identifying a Port 1317
11.5 Erasing Ports and Adapter Configurations 1317
11.6 UEFI NIC Diagnostics 1318
11.6.1 EFI_DRIVER_DIAGNOSTICS_PROTOCOL 1318

Chapter 12: Configuring UEFI for iSCSI 1320

12.1 Navigating the Emulex iSCSI Configuration Utility 1320
12.2 Starting the Emulex iSCSI Configuration Utility 1320
12.3 Configuring MPIO 1322
12.4 Configuring Boot Support 1322
12.5 Configuring TCP ACK Delay 1322
12.6 Viewing the Controller Properties 1323
12.7 Configuring the Network 1324
12.7.1 Configuring the IP Version 1325
12.7.2 Configuring an IPv4 Address 1325
12.7.3 Configuring an IPv6 Address 1327
12.7.4 Configuring VLAN ID and Priority 1329
12.8 Updating Firmware 1331
12.9 Adding and Configuring Targets 1331
12.9.1 Discovering and Adding Boot Targets through DHCP 1331
12.9.2 Using the SendTargets Discovery to Add an iSCSI Target 1332
12.9.3 Manually Adding, Discovering, and Managing Boot Targets 1333
12.9.4 Setting a Boot Target 1334
12.9.5 Selecting an Authentication Method 1335
12.9.6 Pinging a Target 1336
12.9.7 Viewing Advanced Properties 1336
12.9.8 Logging In to or Out of a Target 1337
12.9.9 Deleting a Target 1337
12.9.10 Configuring LUNs 1337
12.9.11 iSNS Configuration 1338
12.10 Erasing the Configuration 1341

Chapter 13: Configuring UEFI for FCoE 1343

13.1 Navigating the Emulex FCoE Configuration Utility 1343
13.2 Starting the Emulex FCoE Configuration Utility 1343
13.3 Updating Firmware and Boot Code 1345
13.4 Enabling an Adapter to Boot from SAN 1345
13.5 Configuring CEE FCF Parameters 1346
13.6 Scanning for Fibre Devices 1348
13.7 Adding Boot Devices 1348
13.8 Deleting Boot Devices 1350
13.9 Changing Boot Device Order 1351
13.10 Configuring Adapter Parameters 1353
13.10.1 Changing the Port Login Retry Timer 1353
13.10.2 Changing the Maximum LUNs per Target 1354
13.10.3 Changing Boot Target Scan Method 1355
13.10.4 Changing Device Discovery Delay 1356
13.11 Resetting Emulex Adapters to Default Values 1358
13.12 Displaying Adapter Information 1359
13.13 Configuring Legacy Only Settings 1360

13.13.1 Enabling or Disabling the Spinup Delay 1361
13.13.2 Enabling or Disabling EDD 3.0 1361
13.13.3 Enabling or Disabling the Start Unit Command 1362
13.13.4 Enabling or Disabling the Environment Variable 1362
13.13.5 Enabling or Disabling Auto Boot Sector 1362
13.13.6 Configuring the Auto Scan Setting 1362
13.14 Requesting a Reset or Reconnect 1363

Chapter 14: Troubleshooting 1365

14.1 Troubleshooting for the NIC Protocol 1365
14.1.1 Issues During the PXE Boot Process 1365
14.1.2 PXE Error Codes 1366
14.2 Troubleshooting for the iSCSI Protocol 1368
14.3 Troubleshooting for the FCoE Protocol 1370

Appendix A: Configuring iSCSI through a DHCP Server Using Vendor-Specific Option 43 1371

A.1 Format of Vendor-Specific Option 43 1371
A.1.1 Description of Parameters 1371
A.1.2 Example 1372

Appendix B: Example for Installing and Configuring Linux or Citrix for PXE Boot and UEFIBoot 1373

B.1 Linux and Citrix PXE Server Remote Installation Procedure 1373
B.1.1 PXE Server 1374
B.1.2 Copying Files to the TFTP Server with the Graphical Version of the Network Booting Tool 1376

Appendix C: Examples for Configuring and Booting UEFI NIC 1381

C.1 UEFI NIC Server Configuration Script for SLES11 SPx 1381
C.1.1 NFS Server Configuration Script 1381
C.1.2 DHCP Server Configuration 1382
C.1.3 TFTP Server Setup 1383
C.2 UEFI NIC Server Configuration Script for RHEL 6.x 1384
C.2.1 Setup Information 1384
C.2.2 NFS Configuration 1384
C.2.3 TFTP Configuration 1385
C.2.4 PXE Configuration 1385
C.2.5 Configuring DHCP 1386
C.3 Configuring the Lenovo System x HS22 Blade with e-Raptor MEZZ 1387
C.4 UEFI NIC IPv6 PXE Setup Procedure for SLES11 SP2 1395
C.4.1 Setup Information 1395
C.4.2 PXE Server Setup 1395
C.4.3 Configure DHCPD6 1396
C.4.4 Configure TFTP 1396
C.4.5 PXE Client Setup 1401

Appendix D: Example for Installing and Booting UEFI FCoE 1410

Appendix E: Loading and Unloading UEFIBoot from the UEFI Shell 1412

E.1 Loading UEFIBoot from the UEFI Shell 1412
E.2 Unloading UEFIBoot from the UEFI Shell 1412

Appendix F: Dell UEFI 1413

F.1 Accessing the Main Configuration Page 1413
F.2 Main Configuration Page 1413
F.2.1 Firmware Image Properties Page 1415
F.2.2 NIC Configuration Page 1416
F.2.3 iSCSI Configuration Page 1418
F.2.4 FCoE Configuration Page 1431
F.2.5 Device Level Configuration Page 1438
F.2.6 NIC Partitioning Configuration Page 1440

F.3 Secure Firmware Download 1444

Appendix G: Configuring PXE Boot for NIC on Dell OCe14000-Based Systems 1445

G.1 Running the PXESelect Utility 1445
G.2 Navigating the PXESelect Utility 1445
G.3 Setting Up a PXE Bootable Network 1446
G.3.1 Configuring NPar 1448
G.3.2 Selecting the Protocol 1448
G.3.3 Configuring Boot Options 1450
14.3.1 Configuring Port Options 1452
G.4 PXE Boot Parameters Default Values 1454

Emulex OneConnect Adapter Installation Guide 1457

Chapter 1: Introduction 1459

1.1 Adapter Identification 1459

Chapter 2: Installation 1460

2.1 Changing the Bracket 1460
2.2 Installing the Adapter 1462
2.3 Attaching Devices 1462
2.4 Applying Power 1464
2.5 Viewing the LEDs 1464

Emulex® Software Kit Migration

User Guide

Version 11.2
December 30, 2016

pub-005841

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

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

Copyright © 2016-2017 Broadcom. All Rights Reserved.

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

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.

Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Chapter 1: Introduction

Beginning with software release 11.2, LightPulse® adapters and OneConnect® adapters have independent software kits.

This document provides special instructions and considerations for using the driver kits for LightPulse and OneConnect adapters.

Special cases include those in which pre-11.2 (original) drivers and applications are replaced by the new 11.2 drivers and applications, and cases in which inbox drivers are replaced by the new 11.2 out-of-box (OOB) drivers.

NOTES

- Detailed instructions for installing LightPulse kits are provided in the driver and applications user guide for LightPulse adapters.
- Detailed instructions for installing OneConnect kits are provided in the driver and applications user guide for OneConnect adapters.

The installation instructions in the user guides apply to "fresh" installations of the LightPulse or OneConnect drivers, in which the pre-11.2 drivers were never installed. They also apply going forward, after the pre-11.2 drivers have been replaced by either LightPulse or OneConnect drivers.

1.1 Supported Adapters

LightPulse software kits support the following adapters:

- LPe12000 series (Fibre Channel [FC] protocol)
- LPe16202/OCe15100 adapters (FC or Network Interface Card [NIC] +Fibre Channel over Ethernet [FCoE] protocol, depending on the small form factor pluggables [SFPs] installed)
- LPe16000 series (FC protocol)
- LPe31000 series (FC protocol)
- LPe32000 series (FC protocol)

OneConnect software kits support the following adapters:

- OCe14000 series (NIC, Internet small computer system interface [iSCSI], and FCoE protocols)
- OCe11000 series (NIC, iSCSI, and FCoE protocols)

1.2 Considerations

If you have both LightPulse and OneConnect adapters in your system and you are upgrading to software release 11.2, you must install both kits before rebooting your system.

If you are downgrading from software release 11.2, you must first uninstall the 11.2 drivers and applications before rebooting your system.

Existing driver configurations are maintained during the upgrade to software release 11.2.

1.2.1 Inbox Drivers

- Mixing Linux or ESXi inbox drivers with out-of-box drivers is not a supported configuration. If inbox drivers are used, make sure to use inbox drivers for all protocols on all LightPulse and OneConnect adapters in the system.
- Management applications will work with the inbox drivers.

1.2.2 Management Applications

- Install LightPulse management applications from the LightPulse software kit to manage LightPulse adapters.
- Install OneConnect management applications from the OneConnect software kit to manage OneConnect adapters.
- Both sets of management applications can run on the same system at the same time.
- Management applications prior to the 11.2 software release cannot be run with the 11.2 LightPulse or OneConnect management tools on the same system. For instructions on reverting applications to a version earlier than 11.2, see Chapter 3 Reverting Drivers and Applications to Previous Versions.

1.2.3 Firmware

A firmware update to version 11.2 is required for all OCe14000-series adapters, OCe11000-series adapters, and LPe16202/OCe15100 adapters (in NIC+FCoE mode) in the system for Unified Extensible Firmware Interface (UEFI) boot code to operate properly.

1.3 Driver Kit Designations and Application Kit Names

The following table provides the original and new designations of the driver kits for each operating system.

Table 1 Driver Kit Designations

Kit Names	Windows	Linux	ESXi	Solaris
Original FC kit	elxdrvr-fc	elx-lpfc	lpfc	N/A
— LightPulse FC kit	elxdrvr-fc	elx-lpfc	lpfc	N/A
— OneConnect FC kit	N/A	N/A	N/A	N/A
Original FCoE kit	elxdrvr-fcoe	elx-lpfc	lpfc	elxfc
— LightPulse FCoE kit	elxdrvr-fcoe	elx-lpfc (same kit for FC)	lpfc	lpfcoe
— OneConnect FCoE kit	brcmdrvr-fcoe	brcmfcoe	brcmfcoe	elxfc
Original NIC kit	elxdrvr-nic	be2net	elxnet	elxnic
— LightPulse NIC kit	elxdrvr-nic	lpnic	lpnic	lpnic
— OneConnect NIC kit	brcmdrvr-nic	be2net	elxnet	elxnic
Original iSCSI kit	elxdrvr-iscsi	be2iscsi	be2iscsi	N/A
— LightPulse iSCSI kit	N/A	N/A	N/A	N/A
— OneConnect iSCSI kit	brcmdrvr-iscsi	be2iscsi	be2iscsi	N/A

The following table provides the new names of the management applications:

Table 2 Application Kit Names

LightPulse Name	OneConnect Name	Description
OCManager	BrcmOCManager	Management GUI
HbaCmd	BrcmHbaCmd	Management CLI
OneCapture	BrcmOneCapture	Dump capture
Elxflash	brcmflash	Firmware download
linlpcfg	brcmlinlpcfg	Linux management tool
winlpcfg	brcmwinlpcfg	Windows management tool
elxcli-fc elxlpfccli lpniccli lpniccli	elxcli-fcoe brcmfcoecli elxcli elxnetcli	esxcli for ESXi 5.5 FC or FCoE esxcli for ESXi 6.0/6.5 FC or FCoE esxcli for ESXi 5.5 NIC esxcli for ESXi 6.0/6.5 NIC
ELXOCM-VMware-vCenter	BRCMOCM-VMware-vCenter	Management GUI and CLI for VMware vCenter
esx-provider-emulex	esx-provider-brcm	CIM Provider

1.4 Additional Information

Appendix A: provides information about original and new names of driver files, which are embedded in the driver kits.

Appendix B: provides illustrations that show the 11.2 kit names, driver names, operating systems, and applications associated with the LightPulse and OneConnect software kits.

Appendix C: contains special instructions on situations that might arise during installation.

Chapter 2: Installing Drivers

This section provides instructions for installing and uninstalling the drivers in special situations.

2.1 Installing Split Out-of-Box Drivers over Inbox Drivers

2.1.1 Windows

No special procedures are needed to install the new out-of-box drivers over existing inbox drivers on Windows. Follow the installation instructions in the appropriate manual:

- For LightPulse adapters, follow the instructions in the *Emulex Drivers for Windows for LightPulse Adapters User Guide*, located on the Broadcom website at www.broadcom.com.
- For OneConnect adapters, follow the instructions in the *Emulex Drivers for Windows for OneConnect Adapters User Guide*, located on the Broadcom website at www.broadcom.com.

2.1.2 Linux

2.1.2.1 LightPulse FC and FCoE Driver with Only LightPulse Adapters Installed

NOTE The existing inbox lpfc driver cannot be removed; it can only be superseded by an out-of-box driver. When a new out-of-box lpfc driver is installed, it is placed in either the `updates` folder or the `extras` folder. The Linux kernel then checks both of these folders. If an lpfc driver is found, that driver is loaded into memory. Otherwise, the inbox lpfc driver is loaded.

To install the LightPulse FC and FCoE driver with only LightPulse adapters installed in the system, do the following:

1. Issue the `tar xvf elx-lpfc-dd-<`*`Linux distribution version`*`>-<`*`driver version`*`>.tar.gz` command.
2. Do one of the following:
   — Run the `elx_lpfc_install.sh` installation script.
   — Issue the `rpm -ivh elx-lpfc-<`*`driver version`*`>_<`*`Linux distribution versio`*`n>.rpm` command.

2.1.2.2 OneConnect FCoE Driver with Only OneConnect Adapters Installed

To install the OneConnect FCoE driver with only OneConnect adapters installed in the system, do the following:

1. Issue the `tar xvf brcmfcoe-dd-<`*`Linux distribution version`*`><`*`driver version`*`>.tar.gz` command.
2. Do one of the following:
   — Run the `brcmfcoe_install.sh` installation script.
   — Issue the `rpm -ivh brcmfcoe-<`*`driver version`*`>_<`*`Linux distribution version`*`>.rpm` command.

The OneConnect FCoE driver, brcmfcoe, attaches to OneConnect adapters only. Because the inbox lpfc driver also supports these adapters, the OneConnect out-of-box FCoE driver cannot load if the inbox lpfc driver has already been loaded. Therefore, following the previous steps will cause the LightPulse driver to be blacklisted to prevent it from loading.

2.1.2.3 LightPulse FC and FCoE Driver and OneConnect FCoE Driver with a Combination of LightPulse and OneConnect Adapters Installed

To install both the LightPulse FC and FCoE driver and the OneConnect FCoE driver on a system that has both LightPulse and OneConnect adapters installed, perform the following steps:

1. For the LightPulse driver:
   a. Issue the `tar xvf elx-lpfc-dd-<Linux distribution version>-<driver version>.tar.gz` command.
   b. Do one of the following:
      - Run the `elx_lpfc_install.sh` installation script.
      - Issue the `rpm -ivh elx-lpfc-<driver version>_<Linux distribution version>.rpm` command.
2. For the OneConnect driver:
   a. Issue the `tar xvf brcmfcoe-dd-<Linux distribution version><driver version>.tar.gz` command.
   b. Do one of the following:
      - Run the `brcmfcoe_install.sh` installation script.
      - Issue the `rpm -ivh brcmfcoe-<driver version>_<Linux distribution version>.rpm` command.

      NOTE During the brcmfcoe driver installation, if LightPulse adapters are detected, the lpfc driver is not blacklisted. Instead, a message prompts you to install the out-of-box 11.2 or later lpfc driver to ensure the operation of both LightPulse and OneConnect adapters on the system.
3. Reboot the system.

2.1.2.4 LightPulse NIC Driver with Only LightPulse Adapters Installed

NOTE The existing inbox be2net driver cannot be removed; it can only be superseded by an out-of-box driver. When a new out-of-box be2net driver is installed, it is placed in either the `updates` folder or the `extras` folder. The Linux kernel then checks both of these folders. If a be2net driver is found, that driver is loaded into memory. Otherwise, the inbox be2net driver is loaded.

To install the LightPulse NIC driver with only LightPulse adapters installed in the system, do the following:

1. Issue the `tar xvf elx-lpnic-dd-<Linux distribution version>-<driver version>.tar.gz` command.
2. Do one of the following:
   — Run the `elx_lpnic_install.sh` installation script.
   — Issue the `rpm -ivh elx-lpnic-<driver version>_<Linux distribution version>.rpm` command.

The LightPulse NIC driver, lpnic, attaches to LightPulse adapters only. Because the inbox be2net driver also supports these adapters, the OneConnect out-of-box be2net driver cannot load if the inbox be2net driver has already been loaded. Therefore, following the previous steps will cause the OneConnect driver to be blacklisted to prevent it from loading.

2.1.2.5 OneConnect NIC Driver with Only OneConnect Adapters Installed

To install the OneConnect NIC driver with only OneConnect adapters installed in the system, do the following:

1. Issue the `tar xvf elx-be2net-dd-<Linux distribution version><driver version>.tar.gz` command.
2. Do one of the following:
   — Run the `elx_be2net_install.sh` installation script.
   — Issue the `rpm -ivh elx-be2net-<driver version>_<Linux distribution version>.rpm` command.

2.1.2.6 LightPulse NIC Driver and OneConnect NIC Driver with a Combination of LightPulse and OneConnect Adapters Installed

To install both the LightPulse NIC driver and the OneConnect NIC driver on a system that has both LightPulse and OneConnect adapters installed, perform the following steps:

1. For the OneConnect driver:
   a. Issue the `tar xvf elx-be2net-dd-<Linux distribution version><driver version>.tar.gz` command.
   b. Do one of the following:
      - Run the `elx_be2net_install.sh` installation script.
      - Issue the `rpm -ivh elx-be2net-<driver version>_<Linux distribution version>.rpm` command.
2. For the LightPulse driver:
   a. Issue the `tar xvf elx-lpnic-dd-<Linux distribution version>-<driver version>.tar.gz` command.
   b. Do one of the following:
      - Run the `elx_lpnic_install.sh` installation script.
      - Issue the `rpm -ivh elx-lpnic-<driver version>_<Linux distribution version>.rpm` command.

      NOTE During the lpnic driver installation, if OneConnect adapters are detected, the be2net driver is not blacklisted. Instead, a message prompts you to install the out-of-box 11.2 or later be2net driver to ensure the operation of both LightPulse and OneConnect adapters on the system.
3. Reboot the system.

2.1.3 VMware Drivers

2.1.3.1 LightPulse FC and FCoE Driver with Only LightPulse Adapters Installed

This driver maintains the lpfc driver binary name so it will update the existing inbox driver.

To install this driver, type:

```
esxcli software vib install –d <path to driver kit>/Emulex-FCoE-FC-lpfc-<driver version>-offline-bundle-<VMware#>.zip

reboot
```

When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Section C.2, VMware.

2.1.3.2 OneConnect FCoE Driver with Only OneConnect Adapters Installed

This driver introduces the brcmfcoe driver binary name. This driver only supports OneConnect adapters. Because this driver has a different binary name, it will not affect the existing inbox driver on installation.

To install this driver, first remove the lpfc driver, and then install the brcmfcoe driver. This action prevents lpfc from claiming its known PCI IDs in the server.

1. Remove the lpfc driver by typing:

   ```
   esxcli software vib remove –n lpfc
   ```

2. Install the brcmfcoe driver by typing:

   ```
   esxcli software vib install –d <path to driver kit>/brcm-FCoE-brcmfcoe-<driver version>-offline-bundle-<VMware#>.zip


   reboot
   ```

3. When the server has rebooted, check the vmhba enumeration to driver binary as follows:

   ```
   esxcli storage core adapter list
   ```

For special instructions, see Section C.2, VMware.

2.1.3.3 LightPulse FC and FCoE Driver and OneConnect FCoE Driver with a Combination of LightPulse and OneConnect Adapters Installed

If your server configuration has both LightPulse and OneConnect adapters, you must install both of the LightPulse and OneConnect OOB drivers. Installing both drivers provides complete coverage for all adapter families.

NOTE Because your configuration uses LightPulse and OneConnect adapters, the inbox lpfc driver is updated with an OOB LightPulse driver to acquire the LightPulse adapters, and then the OOB OneConnect driver is added to acquire the OneConnect adapters.

1. To install these drivers, type the following:

   ```
   esxcli software vib install –d <path to driver kit>/Emulex-FCoE-FC-lpfc-<driver version>-offline-bundle-<VMware#>.zip
   esxcli software vib install –d <path to driver kit>/brcm-FCoE-brcmfcoe-<driver version>-offline-bundle-<VMware#>.zip
   ```

2. When the server has rebooted, check the vmhba enumeration to driver binary as follows:

   ```
   esxcli storage core adapter list
   ```

For special instructions, see Section C.2, VMware.

2.1.3.4 LightPulse NIC Driver with Only LightPulse Adapters Installed

This driver introduces the lpnic driver binary name. This driver supports only LightPulse adapters. Because this driver has a different binary name, it will not affect the existing inbox driver on installation.

To install this driver, first remove the elxnet driver, and then install the lpnic driver. This action prevents elxnet from claiming its known PCI IDs in the server.

1. Remove the elxnet driver by typing:

   ```
   esxcli software vib remove –n elxnet
   ```

2. Install the lpnic driver by typing:

   ```
   esxcli software vib install –d <path to driver kit>/esx-<OS version>-lpnic-<driver version>.zip
   ```

```
reboot
```

3. When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Section C.2, VMware.

2.1.3.5 OneConnect NIC Driver with Only OneConnect Adapters Installed

This driver maintains the elxnet driver binary name so it will update the existing inbox driver.

To install this driver, type:

```
esxcli software vib install –d <path to driver kit>/esx-<OS
version>-elxnet-<driver version>.zip

reboot
```

When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Section C.2, VMware.

2.1.3.6 LightPulse NIC Driver and OneConnect NIC Driver with a Combination of LightPulse and OneConnect Adapters

If your server configuration has both LightPulse and OneConnect adapters, you must install both of the LightPulse and OneConnect OOB drivers. Installing both drivers provides complete coverage for all adapter families.

NOTE Because your configuration uses LightPulse and OneConnect adapters, the inbox elxnet driver is updated with an OOB OneConnect driver to acquire the OneConnect adapters, and then the OOB LightPulse driver is added to acquire the LightPulse adapters.

1. To install these drivers, type the following:

```
esxcli software vib install –d <path to driver kit>/esx-<OS
version>-elxnet-<driver version>.zip
esxcli software vib install –d <path to driver kit>/esx-<OS
version>-lpnic-<driver version>.zip
```

2. When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Section C.2, VMware.

2.2 Installing Split Out-of-Box Drivers over Existing 11.1 or Earlier Out-of-Box Drivers

2.2.1 Windows

No special procedures are needed to install the new out-of-box drivers over existing out-of-box drivers on Windows. Follow the installation instructions in the appropriate manual:

- For LightPulse adapters, follow the instructions in the *Emulex Drivers for Windows for LightPulse Adapters User Guide*, located on the Broadcom website at www.broadcom.com.

- For OneConnect adapters, follow the instructions in the *Emulex Drivers for Windows for OneConnect Adapters User Guide*, located on the Broadcom website at www.broadcom.com.

2.2.2 Linux

2.2.2.1 LightPulse FC and FCoE Driver with Only LightPulse Adapters Installed

NOTE An existing out-of-box lpfc driver cannot be removed; it can only be superseded by another out-of-box driver. When a new out-of-box lpfc driver is installed, it is placed in the `updates` or the `extras` folder. The Linux kernel then checks both of these folders. If a new lpfc driver is found, that driver is loaded into memory. Otherwise, the older lpfc driver is loaded.

To install the LightPulse FC and FCoE driver with only LightPulse adapters installed in the system, do the following:

1. Issue the `tar xvf elx-lpfc-dd-<`*`Linux distribution version`*`>-<`*`driver version`*`>.tar.gz` command.
2. Do one of the following:
   — Run the `elx_lpfc_install.sh` installation script.
   — Issue the `rpm -ivh elx-lpfc-<`*`driver version`*`>_<`*`Linux distribution version`*`>.rpm` command.

2.2.2.2 OneConnect FCoE Driver with Only OneConnect Adapters Installed

To install the OneConnect FCoE driver with only OneConnect adapters installed in the system, do the following:

1. Issue the `tar xvf brcmfcoe-dd-<`*`Linux distribution version`*`><`*`driver version`*`>.tar.gz` command.
2. Do one of the following:
   — Run the `brcmfcoe_install.sh` installation script.
   — Issue the `rpm -ivh brcmfcoe-<`*`driver version`*`>_<`*`Linux distribution version`*`>.rpm` command.

The OneConnect FCoE driver, brcmfcoe, attaches to only OneConnect adapters. Because the pre-11.2 lpfc driver also supports these adapters, the 11.2 or later OneConnect out-of-box FCoE driver cannot load if the pre-11.2 lpfc driver has already been loaded. Therefore, following the previous steps will cause the 11.2 or later LightPulse driver to be blacklisted to prevent it from loading.

2.2.2.3 LightPulse FC and FCoE Driver and OneConnect FCoE Driver with a Combination of LightPulse and OneConnect Adapters Installed

To install both the LightPulse FC and FCoE driver and the OneConnect FCoE driver on a system that has both LightPulse and OneConnect adapters installed, perform the following steps:

1. For the LightPulse driver:
   a. Issue the `tar xvf elx-lpfc-dd-<`*`Linux distribution version`*`>-<`*`driver version`*`>.tar.gz` command.
   b. Do one of the following:
      - Run the `elx_lpfc_install.sh` installation script.
      - Issue the `rpm -ivh elx-lpfc-<`*`driver version`*`>_<`*`Linux distribution version`*`>.rpm` command.

2. For the OneConnect driver:
   a. Issue the `tar xvf brcmfcoe-dd-<Linux distribution version><driver version>.tar.gz` command.
   b. Do one of the following:
      - Run the `brcmfcoe_install.sh` installation script.
      - Issue the `rpm -ivh brcmfcoe-<driver version>_<Linux distribution version>.rpm` command.

      NOTE During the 11.2 or later brcmfcoe driver installation, if LightPulse adapters are detected, the pre-11.2 lpfc driver is not blacklisted. Instead, a message prompts you to install the out-of-box 11.2 or later lpfc driver to ensure the operation of both LightPulse and OneConnect adapters on the system.
3. Reboot the system.

2.2.2.4 LightPulse NIC Driver with Only LightPulse Adapters Installed

NOTE An existing out-of-box be2net driver cannot be removed; it can only be superseded by an out-of-box driver. When a new out-of-box be2net driver is installed, it is placed in either the `updates` folder or the `extras` folder. The Linux kernel then checks both of these folders. If a be2net driver is found, that driver is loaded into memory. Otherwise, the inbox be2net driver is loaded.

To install the LightPulse NIC driver with only LightPulse adapters installed in the system, do the following:

1. Issue the `tar xvf elx-lpnic-dd-<Linux distribution version>-<driver version>.tar.gz` command.
2. Do one of the following:
   — Run the `elx_lpnic_install.sh` installation script.
   — Issue the `rpm -ivh elx-lpnic-<driver version>_<Linux distribution version>.rpm` command.

The LightPulse NIC driver, lpnic, attaches to only LightPulse adapters. Because the pre-11.2 be2net driver also supports these adapters, the OneConnect out-of-box 11.2 or later be2net driver cannot load if the pre-11.2 be2net driver has already been loaded. Therefore, following the previous steps will cause the OneConnect out-of-box 11.2 or later driver to be blacklisted to prevent it from loading.

2.2.2.5 OneConnect NIC Driver with Only OneConnect Adapters Installed

To install the OneConnect NIC driver with only OneConnect adapters installed in the system, do the following:

1. Issue the `tar xvf elx-be2net-dd-<Linux distribution version><driver version>.tar.gz` command.
2. Do one of the following:
   — Run the `elx_be2net_install.sh` installation script.
   — Issue the `rpm -ivh elx-be2net-<driver version>_<Linux distribution version>.rpm` command.

2.2.2.6 LightPulse NIC Driver and OneConnect NIC Driver with a Combination of LightPulse and OneConnect Adapters Installed

To install both the LightPulse NIC driver and the OneConnect NIC driver on a system that has both LightPulse and OneConnect adapters installed, perform the following steps:

1. For the OneConnect driver:
   a. Issue the `tar xvf elx-be2net-dd-<Linux distribution version><driver version>.tar.gz` command.
   b. Do one of the following:
      - Run the `elx_be2net_install.sh` installation script.
      - Issue the `rpm -ivh elx-be2net-<driver version>_<Linux distribution version>.rpm` command.
2. For the LightPulse driver:
   a. Issue the `tar xvf elx-lpnic-dd-<Linux distribution version>-<driver version>.tar.gz` command.
   b. Do one of the following:
      - Run the `elx_lpnic_install.sh` installation script.
      - Issue the `rpm -ivh elx-lpnic-<driver version>_<Linux distribution version>.rpm` command.

      NOTE During the 11.2 or later lpnic driver installation, if OneConnect adapters are detected, the pre-11.2 out-of-box be2net driver is not blacklisted. Instead, a message prompts you to install the out-of-box 11.2 or later be2net driver to ensure the operation of both LightPulse and OneConnect adapters on the system.
3. Reboot the system.

2.2.3 VMware

2.2.3.1 LightPulse FC and FCoE Driver with Only LightPulse Adapters Installed

This driver maintains the lpfc driver binary name so it will update the out-of-box driver.

To install this driver, type:

```
esxcli software vib install –d <path to driver kit>/Emulex-FCoE-FC-lpfc-<driver version>-offline-bundle-<VMware#>.zip

reboot
```

When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Section C.2, VMware.

2.2.3.2 OneConnect FCoE Driver with Only OneConnect Adapters Installed

This driver introduces the brcmfcoe driver binary name. This driver only supports OneConnect adapters. Because this driver has a different binary name, it will not affect the existing out-of-box driver on installation.

To install this driver, first remove the lpfc driver, and then install the brcmfcoe driver. This prevents lpfc from claiming its known PCI IDs in the server.

1. Remove the lpfc driver by typing:

```
esxcli software vib remove –n lpfc
```

2. Install the brcmfcoe driver by typing:

```
esxcli software vib install –d <path to driver kit>/brcm-FCoE-brcmfcoe-<driver version>-offline-bundle-<VMware#>.zip

reboot
```

3. When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Section C.2, VMware.

2.2.3.3 LightPulse FC and FCoE Driver and OneConnect FCoE Driver with a Combination of LightPulse and OneConnect Adapters Installed

If your server configuration has both LightPulse and OneConnect adapters, you need to install both of the LightPulse and OneConnect OOB drivers. Installing both drivers provides complete coverage for all adapter families.

NOTE Because your configuration uses LightPulse and OneConnect adapters, the existing out-of-box lpfc driver is updated with an out-of-box LightPulse driver to acquire the LightPulse adapters, and then the out-of-box OneConnect driver is added to acquire the OneConnect adapters.

1. To install these drivers, type the following:

```
esxcli software vib install –d <path to driver kit>/Emulex-FCoE-FC-lpfc-<driver version>-offline-bundle-<VMware#>.zip
esxcli software vib install –d <path to driver kit>/brcm-FCoE-brcmfcoe-<driver version>-offline-bundle-<VMware#>.zip

reboot
```

2. When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Appendix C: Section C.2, VMware.

2.2.3.4 LightPulse NIC Driver with Only LightPulse Adapters Installed

This driver introduces the lpnic driver binary name. This driver supports only LightPulse adapters. Because this driver has a different binary name, it will not affect the existing out-of-box driver on installation.

To install this driver, first remove the elxnet driver, and then install the lpnic driver. This action prevents elxnet from claiming its known PCI IDs in the server.

1. Remove the elxnet driver by typing:

```
esxcli software vib remove –n elxnet
```

2. Install the lpnic driver by typing:

```
esxcli software vib install -d <path to driver kit>/esx-<OS
version>-lpnic-<driver version>.zip

reboot
```

3. When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Section C.2, VMware.

2.2.3.5 OneConnect NIC Driver with Only OneConnect Adapters Installed

This driver maintains the elxnet driver binary name so it will update the existing out-of-box driver.

To install this driver, type:

```
esxcli software vib install -d <path to driver kit>/esx-<OS
version>-elxnet-<driver version>.zip

reboot
```

When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Section C.2, VMware.

2.2.3.6 LightPulse NIC Driver and OneConnect NIC Driver with a Combination of LightPulse and OneConnect Adapters Installed

If your server configuration has both LightPulse and OneConnect adapters, you must install both of the LightPulse and OneConnect OOB drivers. Installing both drivers provides complete coverage for all adapter families.

NOTE Because your configuration uses LightPulse and OneConnect adapters, the existing out-of-box elxnet driver is updated with an out-of-box OneConnect driver to acquire the OneConnect adapters. Then the out-of-box LightPulse driver is added to acquire the LightPulse adapters.

1. To install these drivers, type the following:

```
esxcli software vib install -d <path to driver kit>/esx-<OS
version>-elxnet-<driver version>.zip
esxcli software vib install -d <path to driver kit>/esx-<OS
version>-lpnic-<driver version>.zip
```

2. When the server has rebooted, check the vmhba enumeration to driver binary as follows:

```
esxcli storage core adapter list
```

For special instructions, see Section C.2, VMware.

2.2.4 Solaris

2.2.4.1 LightPulse FCoE Driver with Only LightPulse Adapters Installed

Install the LightPulse out-of-box driver over the combined out-of-box driver.

1. Uninstall the combined out-of-box driver (elxfc) if it is present:

```
/elxfc_remove
```

2. Install the lpfcoe driver for the LightPulse adapter.

   Untar the driver kit:

   ```
   tar –xf lpfcoe_kit-<driver version>-<OS version>-s11-i386.tar
   ```

   Install the driver:

   ```
   /lpfcoe_install
   ```

2.2.4.2 OneConnect FCoE Driver with Only OneConnect Adapters Installed

Install the elxfc driver for the OneConnect adapter.

Untar the driver kit:

```
tar –xf elxfc_kit-<driver version>-<OS version>-s11-i386.tar
```

Install the driver:

```
/elxfc_install
```

2.2.4.3 LightPulse FC and FCoE Driver and OneConnect FCoE Driver with a Combination of LightPulse and OneConnect Adapters Installed

Install the LightPulse out-of-box driver over the combined out-of-box driver:

1. Uninstall the combined out-of-box driver (elxfc) if it is present:

   ```
   /elxfc_remove
   ```

2. Install the lpfcoe driver for the LightPulse adapter:
   a. Untar the driver kit:

      ```
      tar –xf lpfcoe_kit-<driver version>-<OS version>-i386.tar
      ```

   b. Install the driver:

      ```
      /lpfcoe_install
      ```

3. Install the elxfc driver for the OneConnect adapter.
   a. Untar the driver kit:

      ```
      tar –xf elxfc_kit-<driver version>-<OS version>-i386.tar
      ```

   b. Install the driver:

      ```
      /elxfc_install
      ```

2.2.4.4 LightPulse NIC Driver with Only LightPulse Adapters Installed

Install the LightPulse out-of-box driver over the combined out-of-box driver.

1. Uninstall the combined out-of-box driver (elxnic) if it is present:

   ```
   /elxnic_remove
   ```

2. Install the lpnic driver for the LightPulse adapter.

   Untar the driver kit:

   ```
   tar –xf lpnic_kit-<driver version>-<OS version>-i386.tar
   ```

   Install the driver:

   ```
   /lpnic_install
   ```

2.2.4.5 OneConnect NIC Driver with Only OneConnect Adapters Installed

Install the elxnic driver for the OneConnect adapter.

Untar the driver kit:

```
tar –xf elxnic_kit-<driver version>-<OS version>-i386.tar
```

Install the driver:

```
/elxnic_install
```

2.2.4.6 LightPulse NIC Driver and OneConnect NIC Driver with a Combination of LightPulse and OneConnect Adapters Installed

Install the LightPulse out-of-box driver over the combined out-of-box driver:

1. Uninstall the combined out-of-box driver (elxfc) if it is present:

   ```
   /elxfc_remove
   ```

2. Install the lpfcoe driver for the LightPulse adapter:
   a. Untar the driver kit:

      ```
      tar –xf lpfcoe_kit-<driver version>-<OS version>-i386.tar
      ```

   b. Install the driver:

      ```
      /lpfcoe_install
      ```

3. Install the elxfc driver for the OneConnect adapter.
   a. Untar the driver kit:

      ```
      tar –xf elxfc_kit-<driver version>-<OS version>-i386.tar
      ```

   b. Install the driver:

      ```
      /elxfc_install
      ```

Chapter 3: Reverting Drivers and Applications to Previous Versions

This section describes the downgrade process if you have 11.2 driver and OneCommand Manager for LightPulse or OneCommand CNA Manager for OneConnect application kits installed and you want to revert to an earlier version. This applies to both LightPulse and OneConnect 11.2 kits.

NOTE You must uninstall *all* 11.2 components before installing pre-11.2 components. Mixing 11.2 and pre-11.2 components will cause unexpected results, such as adapters not being discovered by drivers and management applications.

3.1 Windows

You must uninstall all 11.2 Windows components (drivers and management applications) before installing a pre-11.2 Windows kit.

To uninstall all 11.2 Windows components, do the following:

1. Launch the Control Panel from the start menu and select **Programs and Features**.
2. Click all of the following kits that appear in the list of installed programs.

   NOTE All of the 11.2 driver and management application kits must be removed. Any combination of pre-11.2 and 11.2 kits is not supported.

   — LightPulse driver kits:
   - Emulex/FC kit – 11.2.<*version*>
   - Emulex/FCoE kit – 11.2.<*version*>
   - Emulex/NIC kit – 11.2.<*version*>

   — OneConnect Driver kits:
   - Broadcom/FCoE kit – 11.2.<*version*>
   - Broadcom/NIC kit – 11.2.<*version*>
   - Broadcom/iSCSI kit – 11.2.<*version*>

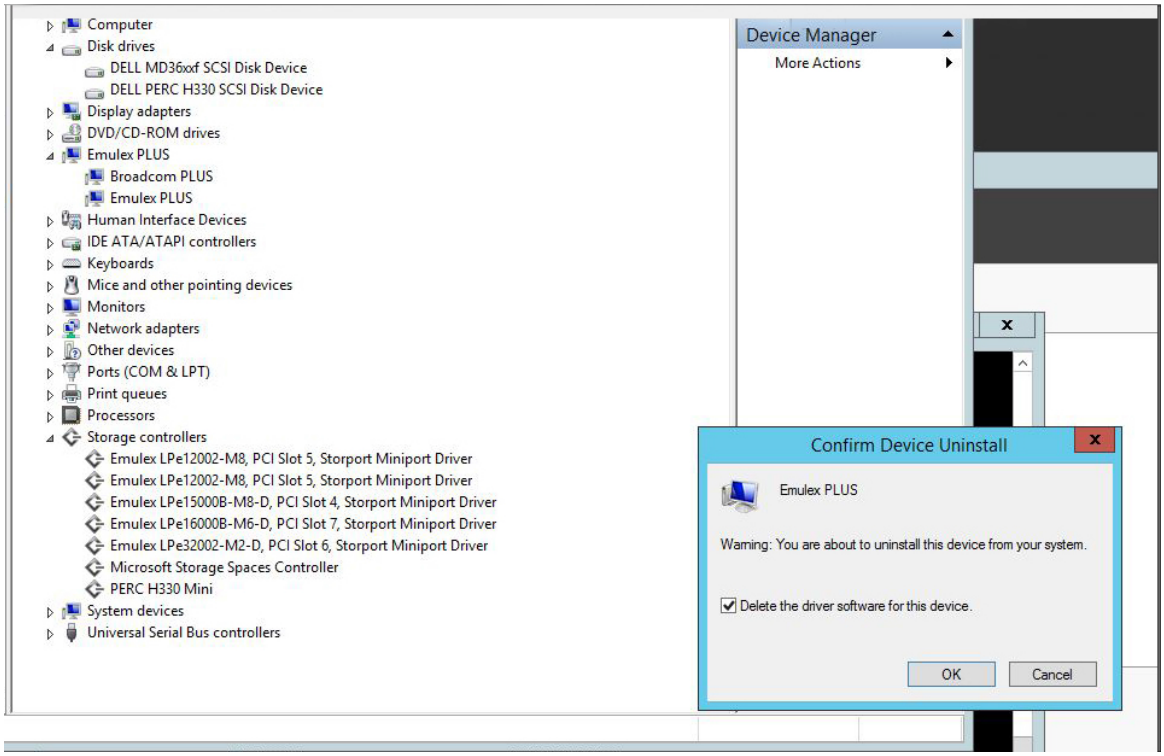
   — LightPulse OneCommand Manager kits:
   - Emulex OCManager Enterprise 11.2.<*version*>
   - Emulex OCManager CLI 11.2.<*version*>

   — OneConnect OneCommand CNA Manager kits:
   - Broadcom OCManager Enterprise 11.2.<*version*>
   - Broadcom OCManager CLI 11.2.<*version*>
3. Select **uninstall**. The selected components are uninstalled.

After the uninstallation process is complete, verify that the 11.2 drivers were uninstalled:

1. From the Computer Management screen, open the Device Manager (see Figure 1).

Figure 1 Windows Device Manager

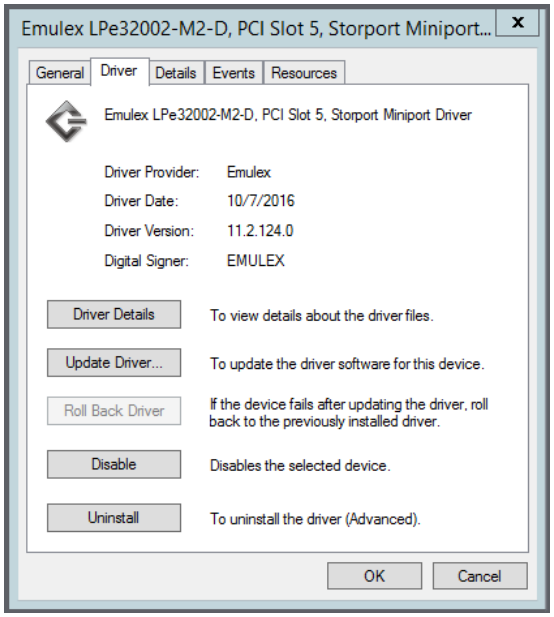


a. Click **Storage controllers**, and then click on an adapter. A window displays the current driver version for the selected adapter. If the 11.2 kit was uninstalled correctly, the driver version shown should not be 11.2. If the 11.2 kit was not uninstalled, the driver version will still be 11.2 (see the example in Figure 2). If the driver version is 11.2, click on **Uninstall** to remove the driver.

b. Click **Emulex PLUS**.
   - Click **Broadcom PLUS** and then select an adapter. A window displays the current driver version for the selected adapter. If the driver version is 11.2, click **Uninstall** to remove the driver.
   - Click **Emulex PLUS** and then select an adapter. A window displays the current driver version for the selected adapter. If the driver version is 11.2, click **Uninstall** to remove the driver.

c. Click **Network adapters** and then select an adapter. A window displays the current driver version for the selected adapter. If the driver version is 11.2, click **Uninstall** to remove the driver.

Figure 2 Windows Driver Version



NOTE When the 11.2 drivers have been uninstalled, if you recheck an adapter, you might see an inbox driver displayed.

2. Verify that the OneCommand Manager kits were uninstalled:
   a. From the Start menu, launch the Control Panel.
   b. Select **Programs and Features**.
   c. Verify that the OneCommand Manager applications have been removed.

   If the 11.2 kits were not uninstalled, perform the uninstall procedure again (see steps 1 through 3).

After all 11.2 Windows components have been uninstalled, install the pre-11.2 drivers and applications as instructed in the *Emulex Drivers for Windows User Guide,* the *OneCommand Manager Application User Guide,* and the *OneCommand Manager Command Line Interface User Guide* that accompanied the kit you are installing.

3.2 Linux

You must uninstall all 11.2 Linux components (drivers and management applications) before installing a pre-11.2 Linux kit.

To uninstall the 11.2 Linux FC and FCoE drivers for LightPulse adapters:

NOTE You must run the uninstall script that shipped with the 11.2 driver kit.

1. Log in as root.
2. If possible, exit all applications that use FC-attached drives, then unmount the drives. If you cannot exit all applications that use FC-attached drives, the uninstall script works properly, but you must reboot after the uninstall is complete.
3. Run the `elx_lpfc_install.sh` script with the `--uninstall` option:

   ```
   ./elx_lpfc_install.sh --uninstall
   ```
4. Verify that the 11.2 drivers were uninstalled.
   a. To determine whether a driver is installed, type the following at the command prompt:

```
lsmod | grep lpfc
```

If a driver is installed, output similar to the following appears:

```
lpfc  693873 0
scsi_transport_fc 64159 1 lpfc
crc tl0dif 12563 2 sd_mod,lpfc
scsi mod 240575 17
iscsi_tcp,libiscsi,scsi_transport_iscsi,st,sr_mod,usb_storage,sg,sd_m
od,scsi_dh_
alua,scsi_dh_emc,scsi_dh_rdac,scsi_dh_hp_sw,scsi_dh,lpfc,scsi_transpo
rt_fc,scsi
tgt,hpsa
```

If a driver is not installed, typing `lsmod | grep lpfc` only displays the command prompt, indicating that the uninstall was successful.

b. If a driver is installed, determine its version by typing

```
/sys/module/lpfc/version
```

Output similar to the following appears:

```
0:11.1.123.456
```

If the version is 11.2.*x.x*, the driver uninstall was not successful. Perform the uninstall procedure again.

To uninstall the 11.2 Linux NIC driver for LightPulse adapters:

NOTE You must run the uninstall script that shipped with the 11.2 driver kit.

1. Log in as root.
2. If possible, exit all applications that use Ethernet-attached drives, then unmount the drives. If you cannot exit all applications that use Ethernet-attached drives, the uninstall script works properly, but you must reboot after the uninstall is complete.
3. Run the `elx_lpnic_install.sh` script with the `--uninstall` option:

```
./elx_lpnic_install.sh --uninstall
```

4. Verify that the 11.2 drivers were uninstalled.
   a. To determine whether a driver is installed, type the following at the command prompt:

```
lsmod | grep lpnic
```

If a driver is installed, output similar to the following appears:

```
lpnic 693873 0
scsi_transport_fc 64159 1 lpnic
crc tl0dif 12563 2 sd_mod,lpfc
scsi mod 240575 17
iscsi_tcp,libiscsi,scsi_transport_iscsi,st,sr_mod,usb_storage,sg,sd_m
od,scsi_dh_
alua,scsi_dh_emc,scsi_dh_rdac,scsi_dh_hp_sw,scsi_dh,lpfc,scsi_transpo
rt_fc,scsi
tgt,hpsa
```

If a driver is not installed, typing `lsmod | grep lpnic` only displays the command prompt, indicating that the uninstall was successful.

b. If a driver is installed, determine its version by typing

```
/sys/module/lpnic/version
```

Output similar to the following appears:

```
0:11.1.123.456
```

If the version is 11.2.*x.x*, the driver uninstall was not successful. Perform the uninstall procedure again.

To uninstall the 11.2 Linux FCoE driver for OneConnect adapters:

NOTE You must run the uninstall script that shipped with the 11.2 driver kit.

1. Log in as root.
2. If possible, exit all applications that use FCoE-attached drives, then unmount the drives. If you cannot exit all applications that use FCoE-attached drives, the uninstall script works properly, but you must reboot after the uninstall is complete.
3. Run the `brcmfcoe_install.sh` script with the *--uninstall* option:

   ```
   ./brcmfcoe_install.sh --uninstall
   ```

4. Verify that the 11.2 drivers were uninstalled.
   a. To determine whether a driver is installed, type the following at the command prompt:

      ```
      lsmod | grep brcmfcoe
      ```

      If a driver is installed, output similar to the following appears:

      ```
      brcmfcoe 693873 0
      scsi_transport_fc 64159 1 brcmfcoe
      crc tl0dif 12563 2 sd_mod,brcmfcoe
      scsi mod 240575 17
      iscsi_tcp,libiscsi,scsi_transport_iscsi,st,sr_mod,usb_storage,sg,sd_m
      od,scsi_dh_
      alua,scsi_dh_emc,scsi_dh_rdac,scsi_dh_hp_sw,scsi_dh,lpfc,scsi_transpo
      rt_fc,scsi
      tgt,hpsa
      ```

      If a driver is not installed, typing `lsmod | grep brcmfcoe` displays only the command prompt, indicating that the uninstall was successful.
   b. If a driver is installed, determine its version by typing:

      ```
      /sys/module/brcmfcoe/version
      ```

      Output similar to the following appears:

      ```
      0:11.1.123.456
      ```

      If the version is 11.2.*x.x*, the driver uninstall was not successful. Perform the uninstall procedure again.

To uninstall the 11.2 Linux NIC drivers for OneConnect adapters:

NOTE You must run the uninstall script that shipped with the 11.2 driver kit.

1. Log in as root.
2. If possible, exit all applications that use Ethernet-attached drives, then unmount the drives. If you cannot exit all applications that use Ethernet-attached drives, the uninstall script works properly, but you must reboot after the uninstall is complete.
3. Run the `elx_be2net_install.sh` script with the `--uninstall` option:

   ```
   ./elx_be2net_install.sh --uninstall
   ```

4. Verify that the 11.2 drivers were uninstalled.
   a. To determine whether a driver is installed, type the following at the command prompt:

      ```
      lsmod | grep elx_be2net
      ```

      If a driver is installed, output similar to the following appears:

      ```
      elx_be2net 693873 0
      scsi_transport_fc 64159 1 elx_be2net
      crc tl0dif 12563 2 sd_mod,elx_be2net
      scsi mod 240575 17
      ```

```
iscsi_tcp,libiscsi,scsi_transport_iscsi,st,sr_mod,usb_storage,sg,sd_m
od,scsi_dh_
alua,scsi_dh_emc,scsi_dh_rdac,scsi_dh_hp_sw,scsi_dh,lpfc,scsi_transpo
rt_fc,scsi
tgt,hpsa
```

If a driver is not installed, typing `lsmod | grep be2net` displays only the command prompt, indicating that the uninstall was successful.

b. If a driver is installed, determine its version by typing:

```
/sys/module/be2net/version
```

Output similar to the following appears:

```
0:11.1.123.456
```

If the version is 11.2.*x.x*, the driver uninstall was not successful. Perform the uninstall procedure again.

To uninstall the LightPulse 11.2 OneCommand Manager enterprise kit:

1. Log on as root.
2. Change to the `elxocm-<`*`platform`*`>-<`*`version`*`>` installation directory.
3. Type the following:

```
./uninstall
```

4. Verify that the 11.2 OneCommand Manager enterprise kit was uninstalled:
   a. Log on as root.
   b. Run the following command from the command prompt:

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

The following message is displayed if the OneCommand Manager was successfully uninstalled:

```
Bash: /usr/sbin/ocmanager/ocmanager no such a file or directory
```

To uninstall the LightPulse 11.2 OneCommand Manager core kit:

1. Log on as root.
2. Do one of the following:
   — Run the `uninstall_ocmanager.sh` script located in `/usr/sbin/hbanyware/scripts`.
   — Run the `uninstall.sh` script located in the installation tar file.
3. Verify that the 11.2 OneCommand Manager core kit was uninstalled:
   a. Log on as root.
   b. Run the following command from the command prompt:

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

The following message is displayed if the OneCommand Manager core kit was successfully uninstalled:

```
Bash: /usr/sbin/ocmanager/ocmanager no such a file or directory
```

To uninstall the OneConnect 11.2 OneCommand CNA Manager enterprise kit:

1. Log on as root.
2. Change to the `brcmocm-<`*`platform`*`>-<`*`version`*`>` installation directory.
3. Type the following:

```
./uninstall
```

4. Verify that the 11.2 OneCommand CNA Manager enterprise kit was uninstalled:
   a. Log on as root.
   b. Run the following command:

```
/usr/sbin/brcmocmanager/brcmocmanager
```

The following message is displayed if the OneCommand CNA Manager was successfully uninstalled:

```
Bash: /usr/sbin/brcmocmanager/brcmocmanager no such a file or
directory
```

To uninstall the OneConnect 11.2 OneCommand CNA Manager core kit:

1. Log on as root.
2. Do one of the following:
   — Run the `uninstall_brcmocmanager.sh` script located in `/usr/sbin/brcmocmanager/scripts`.
   — Run the `uninstall.sh` script located in the installation tar file.
3. Verify that the 11.2 OneCommand CNA Manager core kit was uninstalled:
   a. Log on as root.
   b. Run the following command from the command prompt:

      ```
      /usr/sbin/brcmocmanager/brcmocmanager
      ```

      The following message is displayed if the OneCommand CNA Manager core kit was successfully uninstalled:

      ```
      Bash: /usr/sbin/brcmocmanager/brcmocmanager no such a file or
      directory
      ```

After all 11.2 Linux driver and OneCommand Manager kits have been uninstalled, install the pre-11.2 drivers and applications as instructed in the *Emulex Drivers for Linux User Guide,* the *OneCommand Manager Application User Guide,* and the *OneCommand Manager Command Line Interface User Guide* that accompanied the kit you are installing.

3.3 VMware

You must uninstall all 11.2 VMware components (drivers, CIM Provider, and OneCommand Manager for VMware vCenter) before installing a pre-11.2 VMware kit.

To uninstall the 11.2 drivers, see the VMware Patch Download page for instructions.

To uninstall the CIM Provider for LightPulse adapters:

1. Type the following:

   ```
   esxcli software vib remove
   --vibname=vmware-esx-provider-emulex-cim-provider-<OS
   version>-<11.2.version>-<VMware#>.vib
   ```

2. To verify that the 11.2 CIM Provider was uninstalled, type the following using the CLI:

   ```
   esxcli software vib remove –n <vibname>
   ```

   For example:

   ```
   esxcli software vib remove –n Emulex-cim-provider
   ```

To uninstall the CIM Provider using Windows:

From the Control Panel, click **Programs and Features**, right-click the CIM Provider, and click **Uninstall**.

To uninstall the CIM Provider for OneConnect adapters:

Type the following:

```
esxcli software vib remove
--vibname=vmware-esx-provider-brcm-cim-provider-<OS
version>-<11.2.version>-<VMware#>.vib
```

To verify that the 11.2 CIM Provider was uninstalled, type the following using the CLI:

```
esxcli software vib remove –n <vibname>
```

For example:

```
esxcli software vib remove –n brcm-cim-provider
```

To uninstall the CIM Provider using Windows:

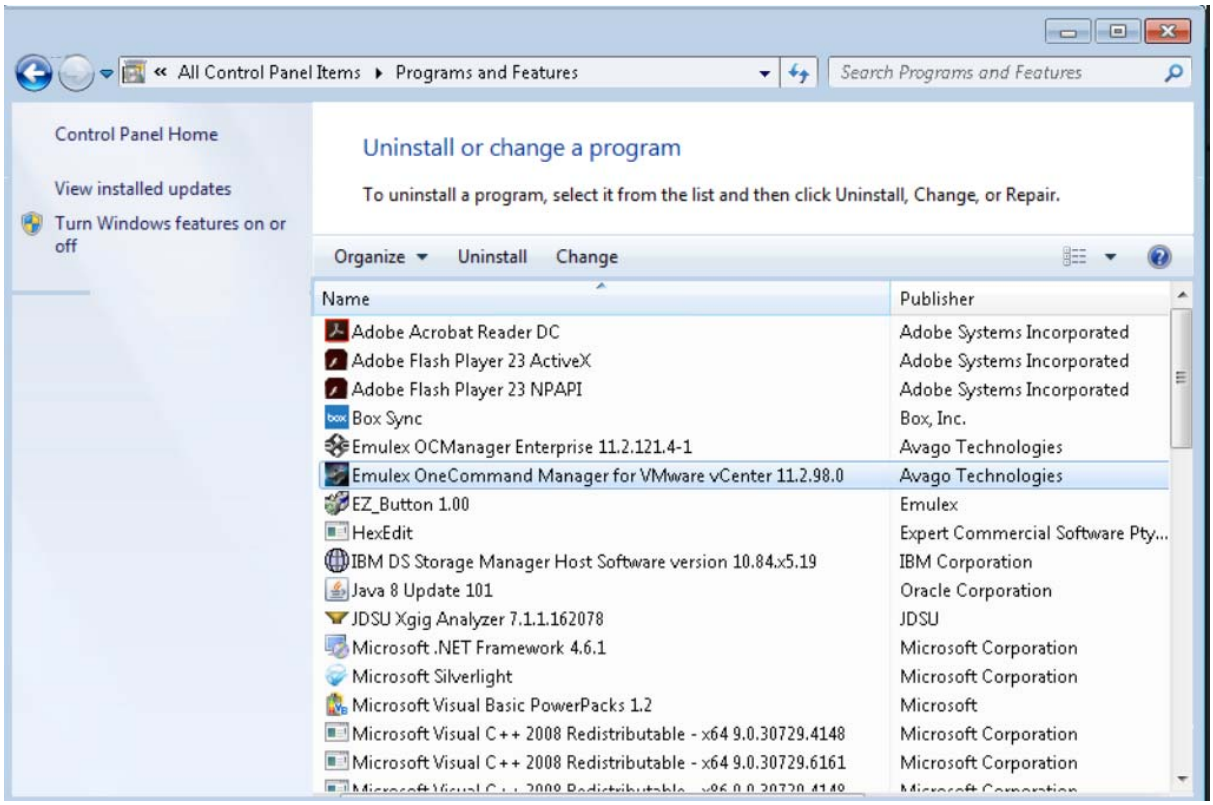
From the Control Panel, click **Programs and Features**, right-click the CIM Provider, and click **Uninstall**.

To uninstall the OneCommand Manager for VMware vCenter plug-in for LightPulse Adapters:

1. Unregister the plug-in with the registered VMware vCenter Server:
   a. Launch the **OCM for VMware vCenter Registration** shortcut located on the desktop.
   b. Provide the details for the registered vCenter server.
   c. Click on **Unregister**.

   Repeat step a to step c if the plug-in is registered with multiple vCenter servers.
2. Uninstall the plug-in server by opening the Control Panel and uninstalling the following kit:
   — Emulex OneCommand Manager for VMware vCenter 11.2.<*version*> (see Figure 3)

Figure 3 Uninstalling the 11.2 OneCommand Manager for VMware vCenter



To verify that the 11.2 OneCommand Manager for VMware vCenter was uninstalled, type the following using the CLI:

```
esxcli software vib remove –n <vibname>
```

For example:

```
esxcli software vib remove –n ELXOCM-VMware-vCenter
```

To uninstall the OneCommand CNA Manager for VMware vCenter plug-in for OneConnect Adapters:

1. Unregister the plug-in with the registered VMware vCenter Server:
   a. Launch the **Broadcom OCM CNA for VMware vCenter Registration** shortcut located on the desktop.
   b. Provide the details for the registered vCenter server.
   c. Click **Unregister**.

   Repeat step a to step c if the plug-in is registered with multiple vCenter servers.
2. Uninstall the plug-in server by opening the Control Panel and uninstalling the following kit:
   — Broadcom OneCommand CNA Manager for VMware vCenter 11.2.<*version*>

   To verify that the 11.2 OneCommand CNA Manager for VMware vCenter was uninstalled, type the following using the CLI:

   ```
   esxcli software vib remove –n <vibname>
   ```

   For example:

   ```
   esxcli software vib remove –n BRCMOCM-VMware-vCenter
   ```

After the 11.2 drivers, CIM Provider, and OneCommand Manager for VMware vCenter kits have been uninstalled, install the pre-11.2 drivers, CIM Provider, and OneCommand Manager for VMware vCenter kits as instructed in the *Emulex Drivers for VMware User Guide, CIM Provider Package Installation Guide,* and the *OneCommand Manager for VMware vCenter User Guide* that accompanied the kit you are installing.

3.4 Solaris

You must uninstall all 11.2 Solaris components (drivers and either OneCommand Manager for Solaris or FCA utilities) before installing a pre-11.2 Solaris kit.

To uninstall the 11.2 Solaris FCoE driver for LightPulse adapters:

1. Log in as root.
2. Remove the driver:
   — For Solaris 10: type `pkgrm lpfcoe`
   — For Solaris 11: type one of the following:
   - `./lpfcoe_remove`
   - `pkg uninstall lpfcoe`

   If the uninstall was successful, the following messages are displayed:

   ```
   Removal of <lpfcoe> was successful.
   <Cleaning directory>
   <lpfcoe_remove complete>
   <uninstall complete>
   ```
3. Reboot the system by typing:

   ```
   reboot
   ```

To uninstall the 11.2 Solaris NIC driver for LightPulse adapters:

1. Log in as root.
2. Remove the driver:
   - For Solaris 10: type `pkgrm lpnic`
   - For Solaris 11: type one of the following:
     - `./lpnic_remove`
     - `pkg uninstall lpnic`

   If the uninstall was successful, the following messages are displayed:

   ```
   Removal of <lpnic> was successful.
   <Cleaning directory>
   <lpnic_remove complete>
   <uninstall complete>
   ```
3. Reboot the system by typing:

   ```
   reboot
   ```

To uninstall the 11.2 Solaris FCoE driver for OneConnect adapters:

1. Log in as root.
2. Remove the driver:
   - For Solaris 10: type `pkgrm elxfc`
   - For Solaris 11: type one of the following:
     - `./elxfc_remove`
     - `pkg uninstall elxfc`

   If the uninstall was successful, the following messages are displayed:

   ```
   Removal of <elxfc> was successful.
   <Cleaning directory>
   <elxfc_remove complete>
   <uninstall complete>
   ```
3. Reboot the system by typing:

   ```
   reboot
   ```

To uninstall the 11.2 Solaris NIC driver for OneConnect adapters:

1. Log in as root.
2. Remove the driver:
   - For Solaris 10: type `pkgrm elxnic`
   - For Solaris 11: type one of the following:
     - `elxnic_remove`
     - `pkg uninstall elxnic`

   If the uninstall was successful, the following messages are displayed:

   ```
   Removal of <elxnic> was successful.
   <Cleaning directory>
   <elxnic_remove complete>
   <uninstall complete>
   ```
3. Reboot the system by typing

   ```
   reboot
   ```

To uninstall the LightPulse 11.2 OneCommand Manager enterprise kit:

1. Log in as root.
2. Run the following script:

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

   If the uninstallation is successful, the following messages are displayed:

   ```
   Removal of <ELXocm> was successful.
   <Cleaning directory>
   <ELXocm_remove complete>
   <uninstall complete>
   ```

To uninstall the LightPulse 11.2 OneCommand Manager core kit:

1. Log on as root.
2. Do one of the following:
   - Run `/opt/ELXocm/scripts/uninstall`
   - Run the uninstall script located in the installation tar file.
   - Enter the command `pkgrm ELXocmcore`.

   If the uninstallation is successful, the following messages are displayed:

   ```
   Removal of <ELXocmcore> was successful.
   <Cleaning directory>
   <ELXocmcore_remove complete>
   <uninstall complete>
   ```

To uninstall the OneConnect 11.2 OneCommand CNA Manager enterprise kit:

1. Log on as root.
2. Run the following script:

   ```
   /opt/BRCMocm/scripts/uninstall
   ```

   If the uninstallation is successful, the following messages are displayed:

   ```
   Removal of <BRCMocm> was successful.
   <Cleaning directory>
   <BRCMocm_remove complete>
   <uninstall complete>
   ```

To uninstall the OneConnect 11.2 OneCommand CNA Manager core kit:

1. Log on as root.
2. Do one of the following:
   - Run `/opt/brcmocm/scripts/uninstall`.
   - Run the uninstall script located in the installation tar file.
   - Enter the command `pkgrm brcmocmcore` command.

   If the uninstallation is successful, the following messages are displayed:

   ```
   Removal of <brcmocmcore> was successful.
   <Cleaning directory>
   <brcmocmcore_remove complete>
   <uninstall complete>
   ```

To uninstall the LightPulse 11.2 emlxu kit:

1. Log on as root.
2. Do one of the following:
   — For Solaris 10: type one of the following:
      - `emlxu_remove`
      - `pkgrm EMLXemlxu`
   — For Solaris 11: type one of the following:
      - `emlxu_remove`
      - `pkg uninstall emlxu`

   If the uninstallation was successful, the following messages are displayed:

```
Removal of <EMLXemlxu> was successful.
<Cleaning directory>
<emlxu_remove complete>
<uninstall complete>
```

To uninstall the OneConnect 11.2 brcmxu kit:

1. Log on as root.
2. Do one of the following:
   — For Solaris 10: type one of the following:
      - `brcmxu_remove`
      - `pkgrm brcmxu`
   — For Solaris 11: type one of the following:
      - `brcmxu_remove`
      - `pkg uninstall brcmxu`

   If the uninstallation was successful, the following messages are displayed:

```
Removal of <brcmxu> was successful.
<Cleaning directory>
<brcmxu_remove complete>
<uninstall complete>
```

After all 11.2 Solaris drivers, OneCommand Manager kits, and FCA utilities kits have been uninstalled, install the pre-11.2 drivers and applications as instructed in the *Emulex Drivers for Solaris User Guide,* the *OneCommand Manager Application User Guide,* the *OneCommand Manager Command Line Interface User Guide,* and the *FCA Utilities for Solaris User Guide* that accompanied the kit you are installing.

Chapter 4: Technical Tips

This section provides information on situations that might be encountered during installation.

4.1 General

4.1.1 Applications Specific to LightPulse or OneConnect

Applications that are installed for LightPulse adapters will not recognize or see any OneConnect adapters if they are also installed in the system. The same is true for applications that are installed for OneConnect adapters; the application will not recognize or see any LightPulse adapters that are installed in the system.

This situation applies to the following applications:

- OneCommand Manager and OneCommand CNA Manager
- Elxflash stand alone utility
- Elxflash and Brcmflash stand alone utilities
- ESX CIM Provider
- OneCommand Manager for VMware vCenter and OneCommand CNA Manager for VMware vCenter

NOTE For the Elxflash stand alone utility, when the `elxflash /fc` command is used in environments that include both OneConnect and LightPulse adapters, Elxflash displays only LightPulse FC adapters. OneConnect adapters will not be displayed.

When the `elxflash /ucna` command is used, Elxflash displays only LightPulse FCoE/NIC adapters. OneConnect adapters will not be displayed.

4.1.2 Multiple ASICs

If an adapter has multiple ASICs, and one ASIC is LightPulse and the other is OneConnect, the OneCommand Manager for LightPulse Adapters ignores the OneConnect ASIC, and the OneCommand CNA Manager for OneConnect Adapters ignores the LightPulse ASIC.

4.1.3 Reverting to a Prior Driver Version

4.1.3.1 Linux and Windows, NIC and iSCSI Protocols

If you attempt to install an 11.1 or earlier driver without first uninstalling the 11.2 or later driver, a message will direct you to uninstall the 11.2 or later driver first.

4.1.3.2 Downgrading from Version 11.0 and Higher to 10.6 is Unsuccessful when NPar and SR-IOV are Both Enabled (Dell)

If SR-IOV and nPAR are both enabled on the adapter, a firmware downgrade from 11.0 and higher to version 10.6 will be unsuccessful. Version 10.6 does not support the concurrent use of SR-IOV and nPAR.

Disable either SR-IOV or nPAR, and attempt the firmware downgrade again.

4.1.3.3 Downgrading from an 11.2 Kit to a pre-11.2 Kit with the Secure Firmware Jumper Installed in the Adapter

If you attempt to downgrade 11.2 firmware to a pre-11.2 on an adapter with the secure firmware jumper installed using UEFI, the following error message is displayed:

```
Incorrect firmware file or no supported device.
```

Workaround

Remove the secure firmware jumper on the adapter to allow firmware downgrading.

4.1.4 Boot Code Functionality

If 11.2 boot code functionality is missing, make sure the firmware for all adapters in the system is at an 11.2 version.

4.2 Windows

4.2.1 Windows OneInstall

The OneInstall kit upgrade fails if the OneCommand Manager application is open and running.You must close the OneCommand Manager application before starting the upgrade.

4.2.2 Downgrading from a LightPulse 11.2 Kit with LPe16000-series Adapters

If you are downgrading from the LightPulse 11.2 kit and the LPe16000-series adapters BIOS is downgraded to a version prior to 10.6.193.22, the system will not boot in a boot from SAN (BFS) environment. This can be avoided if you do not downgrade to versions earlier than 10.6.193.22 in Windows, either locally or BFS.

- Windows BFS environment:

  Before booting to the operating system after a firmware downgrade, set the boot topology to point-to-point in the BIOS. The system will be able to boot, but the non-boot port will still show that the link is down.
- Windows local boot environment:

  When the system is not in a BFS environment, the adapter link will be down. However, if you are using an 8G Brocade switch, the topology can be changed to loop-only mode, and the system can boot.

4.2.3 Downgrading from a LightPulse 11.2 Kit to an 11.0 Kit with LPe12000-Series Adapters

If you attempt to downgrade from an 11.2 kit to a pre-11.2 kit without first uninstalling the 11.2 kit, the following popup message is displayed.

Figure 4 Windows Security Message



Before performing the 11.2 downgrade, uninstall the 11.2 kit.

4.3 Linux

4.3.1 Unloading the lpfc and the brcmfcoe Drivers

Unloading the lpfc and the brcmfcoe drivers fail if the OneCommand Manager application is not stopped first.

- If the BrcmOCManager brcmhbamgrd is running, the LightPulse lpfc driver cannot be unloaded. Stop the BrcmOCManager brcmhbamgrd by typing the following command:

  ```
  /usr/sbin/brcmocmanager/stop_brcmocmanager
  ```

- If the OCManager elxhbamgrd is running, the OneConnect brcmfcoe driver cannot be unloaded. Stop the OCManager elxhbamgrd by typing the following command:

  ```
  /usr/sbin/ocmanager/stop_ocmanager
  ```

The lpfc and brcmfcoe drivers can now be unloaded.

4.3.2 Installing a pre-11.2 Version of OneCommand Manager

Running a pre-11.2 `ocm_install.sh` script to install a pre-11.2 version of OneCommand Manager for Linux fails.

To install the pre-11.2 OneCommand Manager, first uninstall all of the 11.2 Linux drivers. See Section 3.2, Linux for instructions on uninstalling the drivers.

4.4 VMware

4.4.1 VMware Inbox Driver Installation over 11.2 Drivers

The pre-11.2 driver supports all FC and FCoE adapters, and the OOB LightPulse driver named **lpfc** supports only LightPulse FC and FCoE adapters. The OOB OneConnect driver named **brcmfcoe** supports only OneConnect FCoE adapters.

- Installing the pre-11.2 driver when only the OOB LightPulse driver is installed:
  - The `vib upgrade` command to replace the OOB LightPulse driver with the pre-11.2 driver will not work because the 11.2 drivers are always later versions or more recent than the pre-11.2 driver.
  - Remove the existing OOB LightPulse lpfc driver before installing the pre-11.2 driver:

    ```
    esxcli software vib remove –n lpfc
    esxcli software vib install –d <path to
    offline-bundle/lpfc-offline-bundle.zip>
    ```

  - vmhba renumbering is seen. Manually remove old entries that changed from esx.conf to vmhba numbers (see Section C.2, VMware for instructions on how to remove old entries).
- Installing the pre-11.2 driver when only an OOB OneConnect driver is installed:
  - OneConnect brcmfcoe does not share a common name with the pre-11.2 driver, so the upgrade vib command is not needed.
  - Remove the existing OOB OneConnect brcmfcoe driver before installing the pre-11.2 driver:

    ```
    esxcli software vib remove –n  brcmfcoe
    esxcli software vib install –d <path to
    offline-bundle/lpfc-offline-bundle.zip>
    ```

  - vmhba renumbering is not usually seen if only brcmfcoe was previously installed.
- Installing the pre-11.2 driver when both the OOB LightPulse and OneConnect drivers are present:
  - Remove the existing 11.2 OneConnect brcmfcoe driver and 11.2 LightPulse driver (if any, in any order) before installing the pre-11.2 driver:

    ```
    esxcli software vib remove –n brcmfcoe
    esxcli software vib remove –n lpfc
    esxcli software vib install –d <path to
    offline-bundle/lpfc-offline-bundle.zip>
    ```

  - vmhba renumbering is seen. Manually remove old entries that changed from esx.conf to vmhba numbers (see Section C.2, VMware for instructions on how to remove old entries).
- If the pre-11.2 driver is installed before the 11.2 drivers are removed:
  - The pre-11.2 driver will not install if LightPulse lpfc is present using either `vib install` or `vib upgrade` (the `esxcli vib install` command states that lpfc is already present, or there is nothing to upgrade).
  - If only brcmfcoe is present, the pre-11.2 driver will install using the `vib install` command.

4.4.1.1 11.2 Drivers Installation

The server has no inbox or OOB pre-11.2 driver installed, or the pre-11.2 driver has been removed and the system has been rebooted before the 11.2 drivers are installed.

- Both of the 11.2 drivers are installed at the same time:
  - Use the `vib install` command to install the 11.2 drivers, and then reboot:

    ```
    esxcli software vib install –d <path to
    offline-bundle/<lpfc/brcmfcoe>-offline-bundle.zip>
    ```
  - Brcmfcoe claims the lower vmhba numbering for its adapters.
- LightPulse lpfc is installed first and the system was rebooted before the installation of brcmfcoe.
  - This situation can force ESXi to load lpfc first, and then brcmfcoe, so the vmhba numbering is different than when both lpfc and brcmfcoe are installed at the same time; or, when brcmfcoe is installed first and then lpfc is installed after a reboot.
  - LightPulse lpfc claims its HBAs first, so it gets lower vmhba numbering for its adapters.
- Brcmfcoe is installed, the system is rebooted, and then LightPulse lpfc is installed.
  - This situation can be the same as when both of the 11.2 drivers are installed together before a reboot.

NOTE In Boot from SAN configurations, make sure that on reboot at least one driver will be installed that claims the adapter that is configured to boot from SAN. For OneConnect adapters, at least an FCoE driver and a NIC driver must exist on reboot for Boot from SAN setups.

4.5 Solaris

4.5.1 Running the Install Script of a pre-11.2 Kit Does Not Uninstall the Broadcom OneCommand CNA Manager

If the 11.2 OneCommand CNA Manager kits have not been removed and you run the `install.sh` of a pre-11.2 kit, the installation script detects the Emulex OneCommand CNA Manager and prompts you to uninstall the kit. However, the pre-11.2 script does not detect Broadcom OneCommand CNA Manager, and so it does not uninstall it. After the pre-11.2 kit is installed, it will manage only LightPulse adapters.

4.6 Dell Applications Downgrade

LightPulse and OneConnect DUP/MUP application packages do not support automatic downgrading from 11.2 to a prior version. If this is attempted, the following message is displayed:

```
A later version of [OCM Version] is already installed on this machine. The
setup cannot continue.
```

In some situations, when you attempt to downgrade from an 11.2.xx kit to a prior version, the popup message will not be displayed, and the installation of the prior version will succeed. This does not remove the prior versions, but it will install the prior version kit alongside the others. Functionality of the OneCommand Manager application and the drivers being used by adapters in the system is inconsistent.

To downgrade to version 11.1 and earlier from version 11.2 and later, all kits (LightPulse and OneConnect) need to be uninstalled prior to installing the older kit version.

To uninstall 11.2.*xx.xx* kits for both LightPulse and OneConnect:

1. Run the `setup.exe` file that was used to install the current kit on the host.
2. The installation wizard will display an options dialog. Select the **Remove** option.
3. Repeat steps one and two for the remaining LightPulse or OneConnect kits.

NOTE The kits can also be uninstalled by using the Control Panel.

1. From the Start menu, launch the Control Panel.
2. Select **Programs and Features**.
3. Right-click on the kit you want to remove and select **Uninstall**.

After removing the LightPulse or OneConnect 11.2.*xx.xx* kits from the system, an older kit can be installed.

Appendix A: Driver Names

To ensure independent delivery of the LightPulse and OneConnect software kits, the following changes were made:

- Drivers that support both divisions' products have been separated into two drivers with unique names.
- The original names of LightPulse Fibre Channel drivers and utilities have been retained.
- The original names of OneConnect NIC and iSCSI drivers have been retained.

The following have been given new names:

- OneConnect utilities
- OneConnect Linux and VMware FCoE drivers
- LightPulse NIC drivers
- LightPulse Windows and Solaris FCoE drivers

The following table provides the original and new driver names for each operating system.

Table 3 Original and New Driver Names

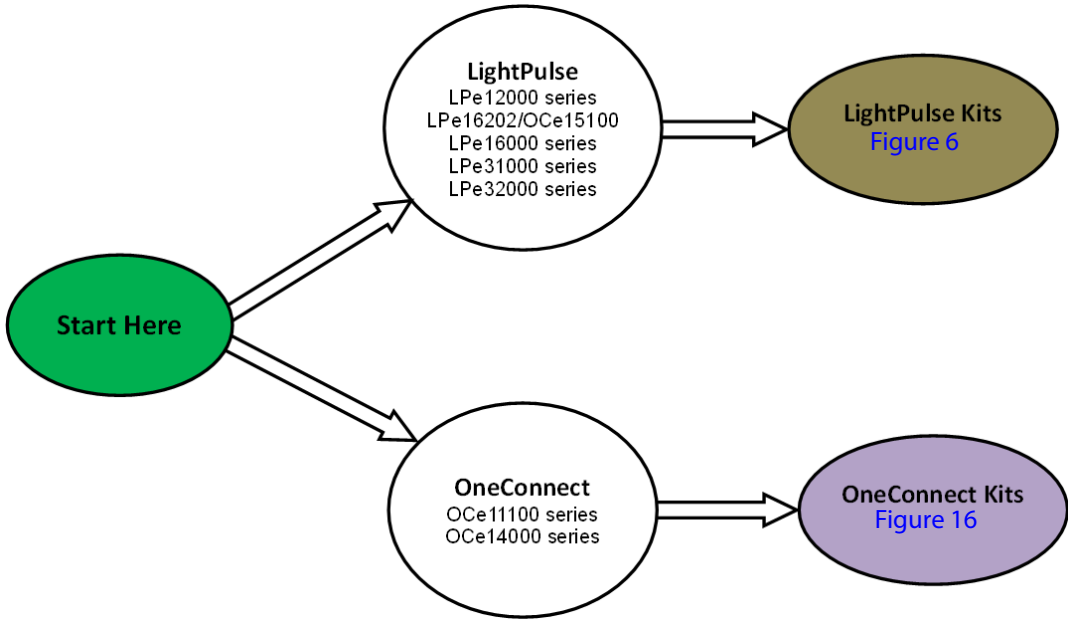
Operating System Family	Original Name	LightPulse Name	OneConnect Name
Windows	elxplus	elxplus	brcmplus
Windows	elxfc	elxfc	N/A
Windows	elxcna	lpfcoe	elxcna
Windows	ocnd63	lpnic	ocnd63
Windows	ocnd64	lpnic	ocnd64
Windows	ocnd65	lpnic	ocnd65
Windows	be2nd62	lpnic	be2nd62
Windows	be2nd6x	lpnic	be2nd6x
Windows	be2iscsi	N/A	be2iscsi
Windows	lancutil	lancutil	N/A
Windows	lanc6ut	lanc6ut	N/A
Linux	lpfc	lpfc	brcmfcoe
Linux	be2net	lpnic	be2net
Linux	be2iscsi	N/A	be2iscsi
Linux	ocrdma	N/A	ocrdma
VMware	lpfc	lpfc	brcmfcoe
VMware	elxlpfccli	elxlpfccli	brcmfcoecli
VMware	elxnet	lpnic	elxnet
VMware	elxnetcli	lpniccli	elxnetcli
VMware	be2iscsi	N/A	be2iscsi
VMware	elxiscsi	N/A	elxiscsi
Solaris	elxfc	lpfcoe	elxfc
Solaris (inbox)	emlxs	emlxs	N/A
Solaris	elxnic	lpnic	elxnic

Appendix B: Quick Lookup Guide

B.1 Quick Lookup Guide, LightPulse and OneConnect

The following illustrations provide a guide to the split kits names, driver names, operating systems, and applications associated with the LightPulse and OneConnect kits. Figure 5 shows an overview.

Figure 5 Overview

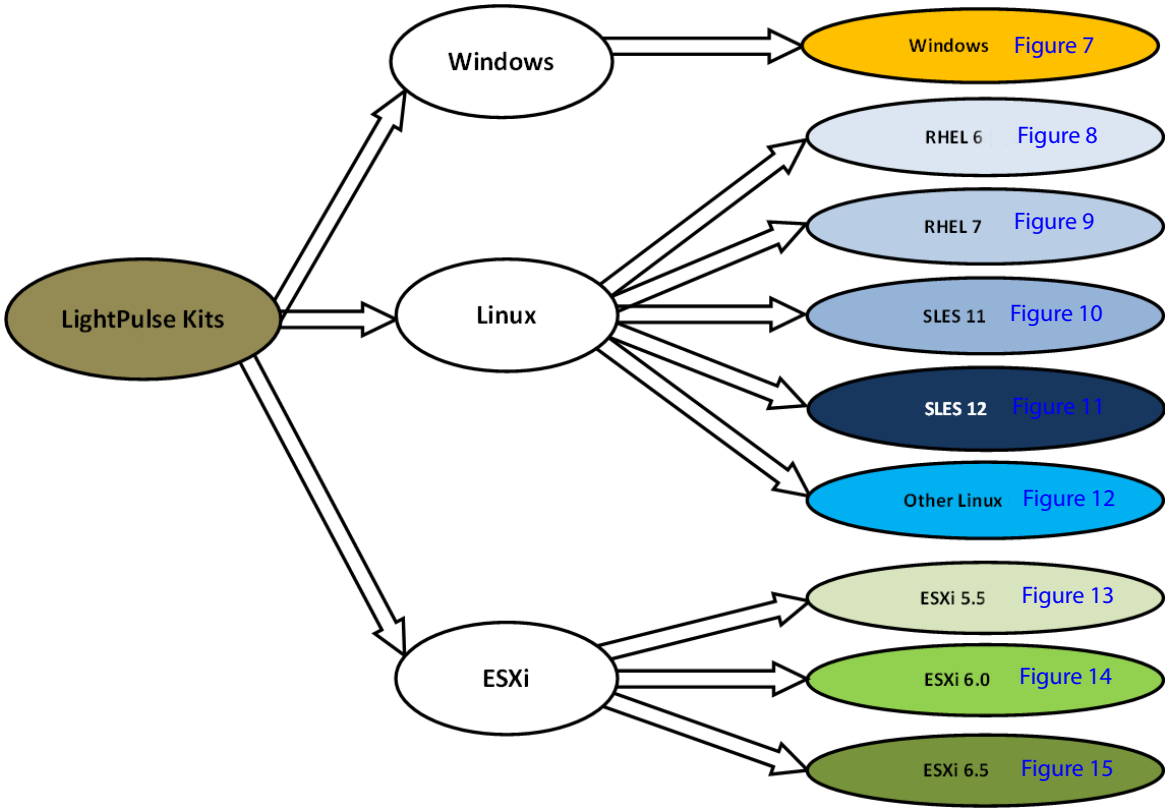


For LightPulse kits, start at flow A (Figure 6).

For OneConnect kits, start at flow K (Figure 16).

B.1.1 LightPulse Adapters and Kits

Figure 6 LightPulse Adapters and Kits



LightPulse	
Windows	Figure 7
RHEL 6	Figure 8
RHEL 7	Figure 9
SLES 11	Figure 10
SLES 12	Figure 11
Other Linux	Figure 12
ESXi 5.5	Figure 13
ESXi 6.0	Figure 14
ESXi 6.5	Figure 15

Figure 7 LightPulse Windows

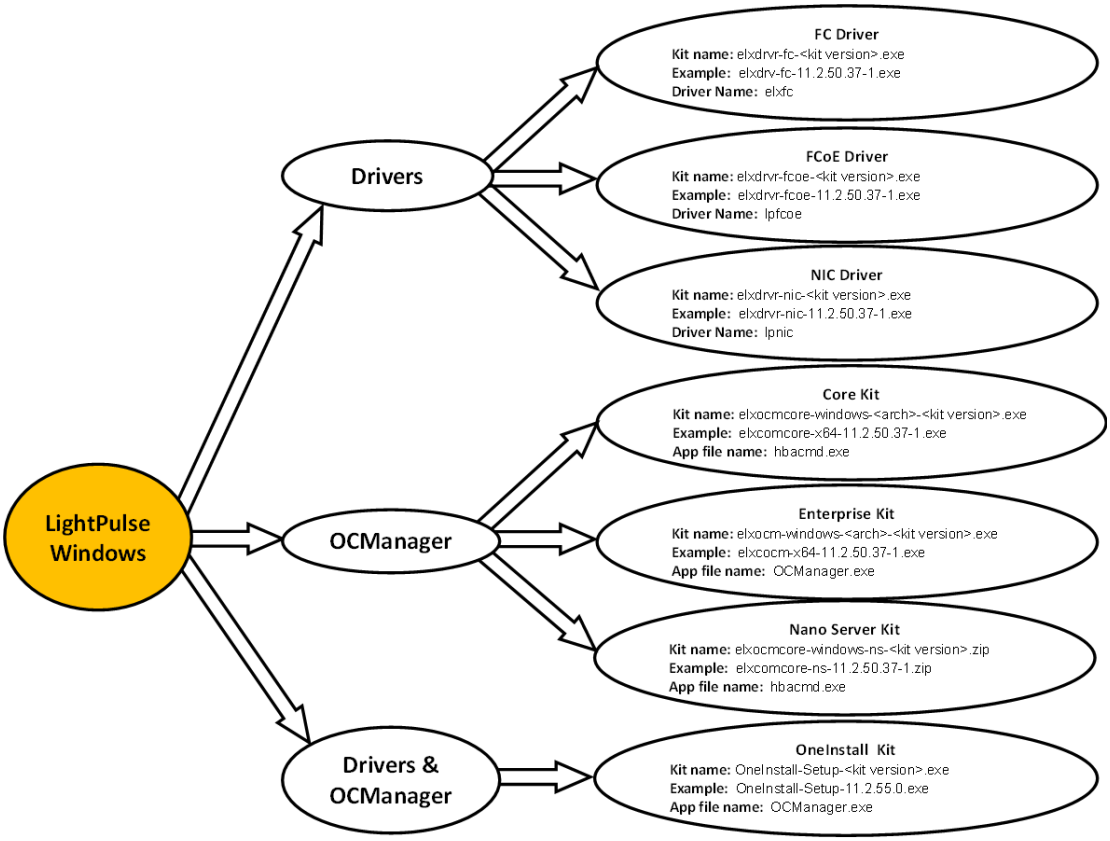


Figure 8 LightPulse Linux RHEL 6

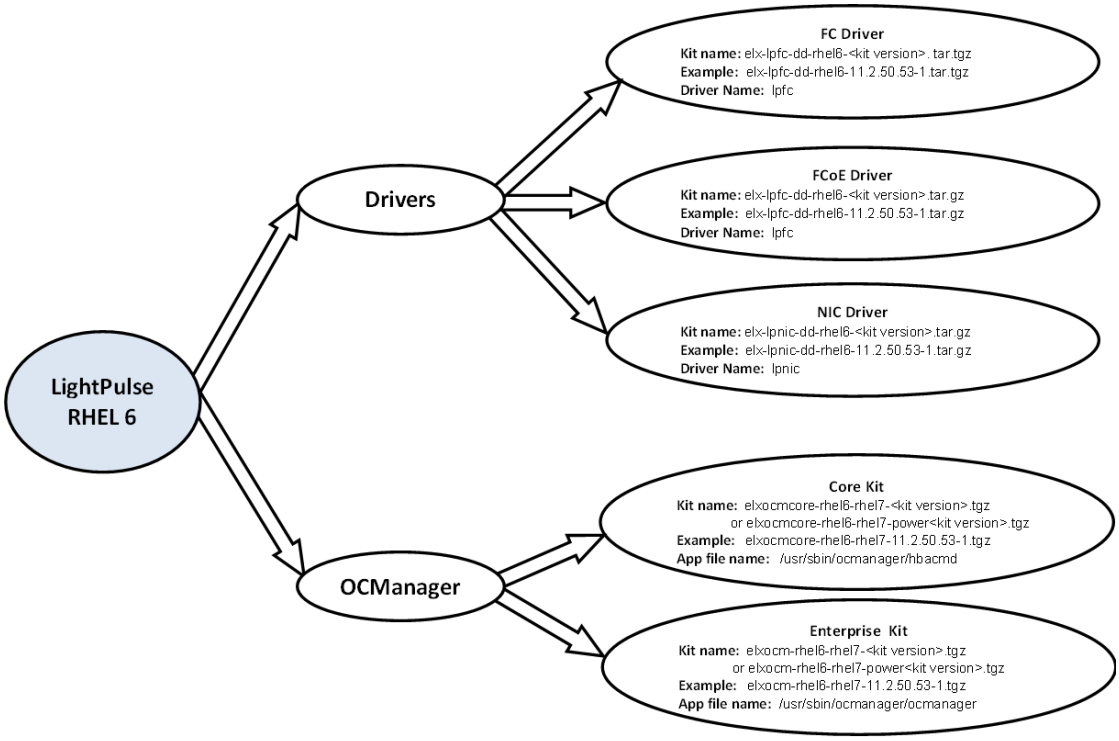


Figure 9 LightPulse RHEL 7

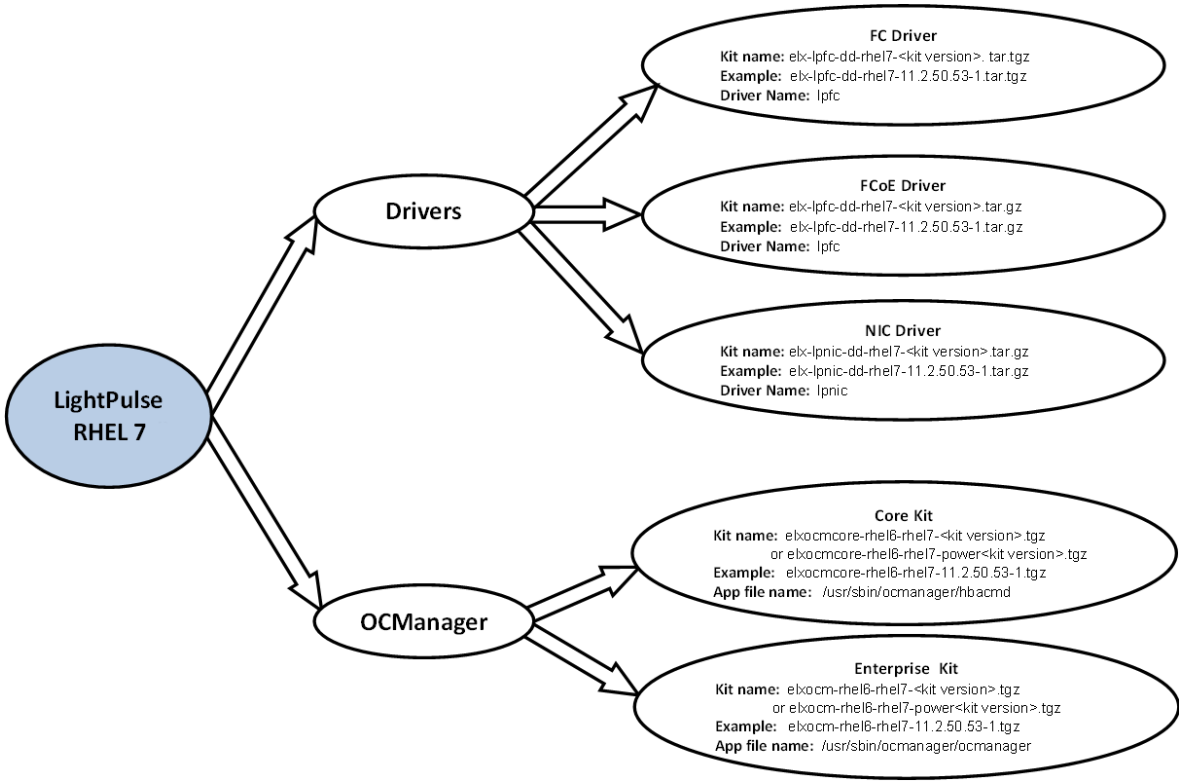


Figure 10 LightPulse SLES 11

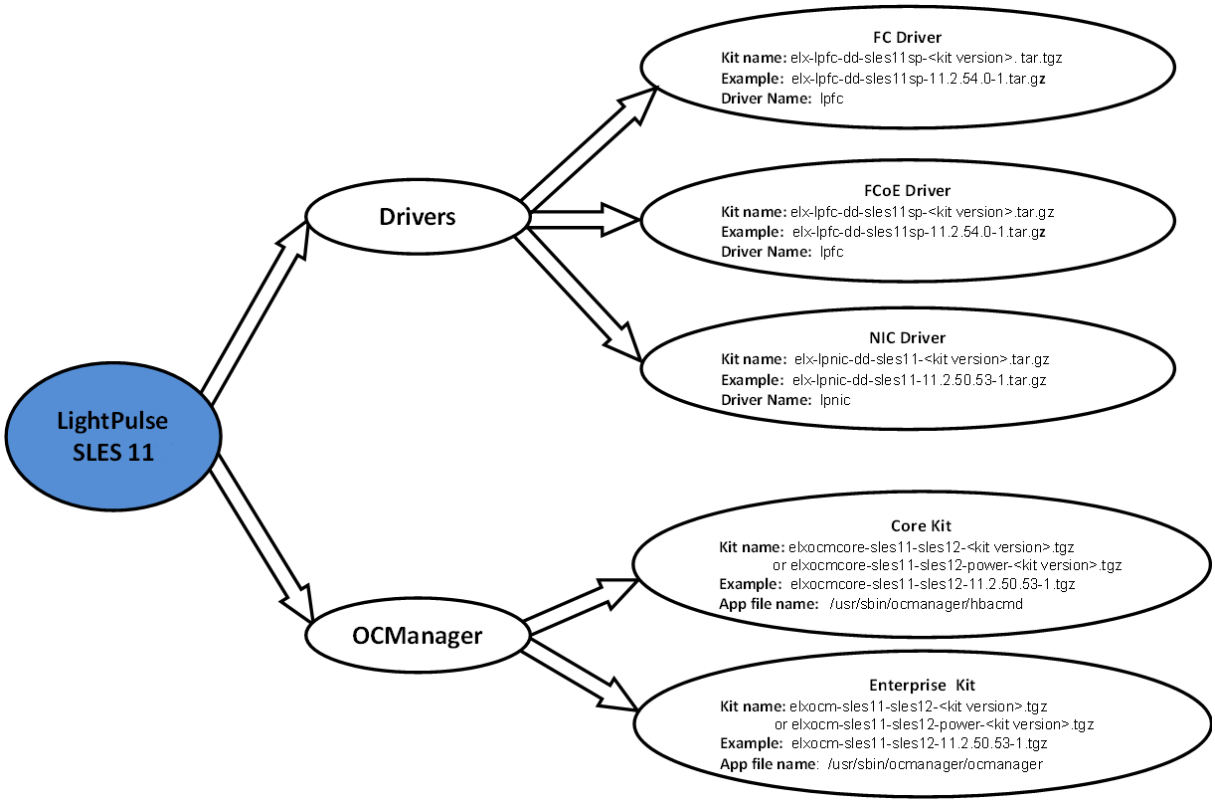


Figure 11 LightPulse SLES 12

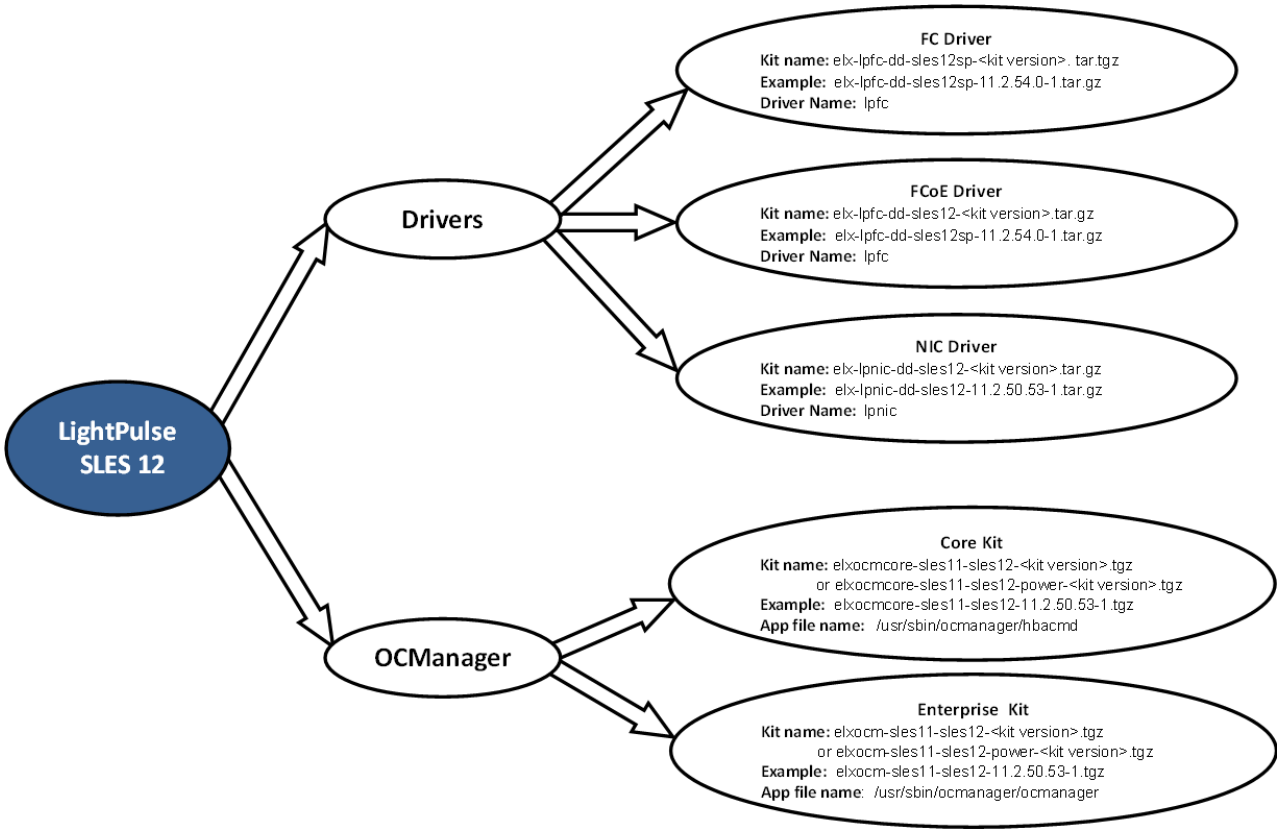


Figure 12 LightPulse Other Linux

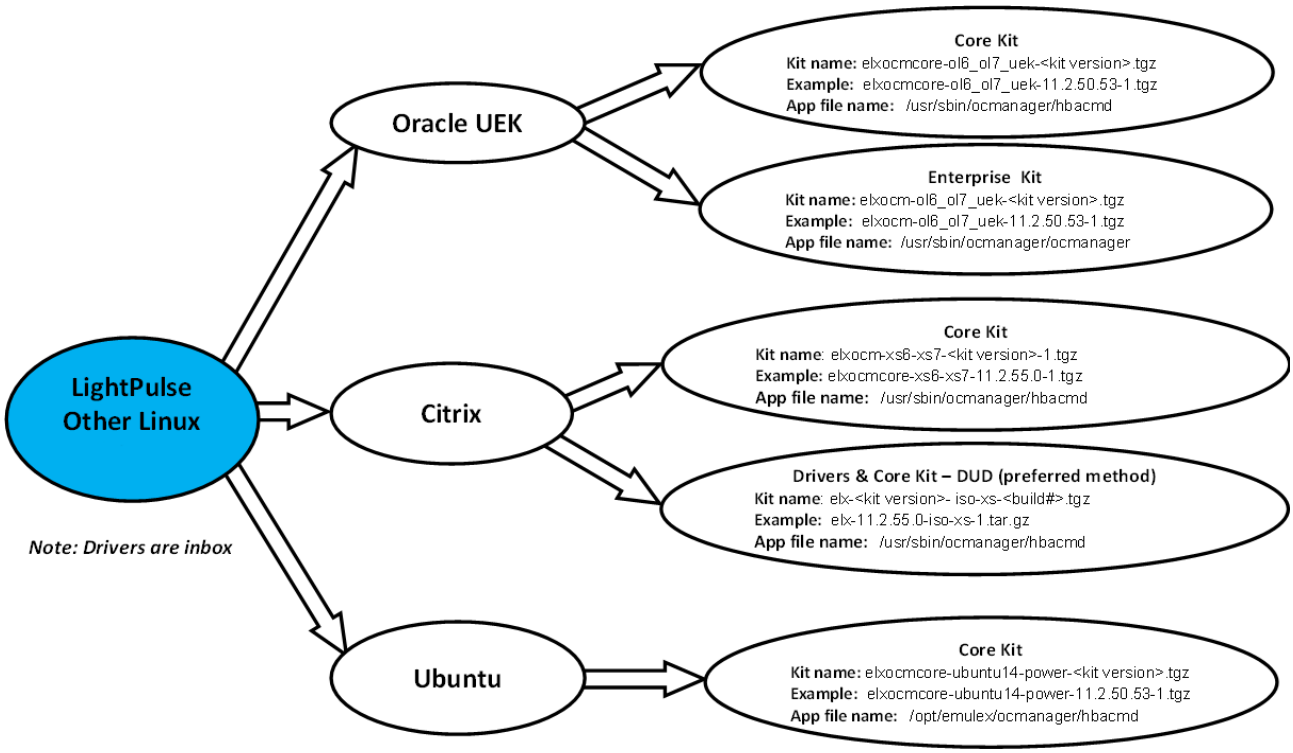


Figure 13 LightPulse ESX 5.5

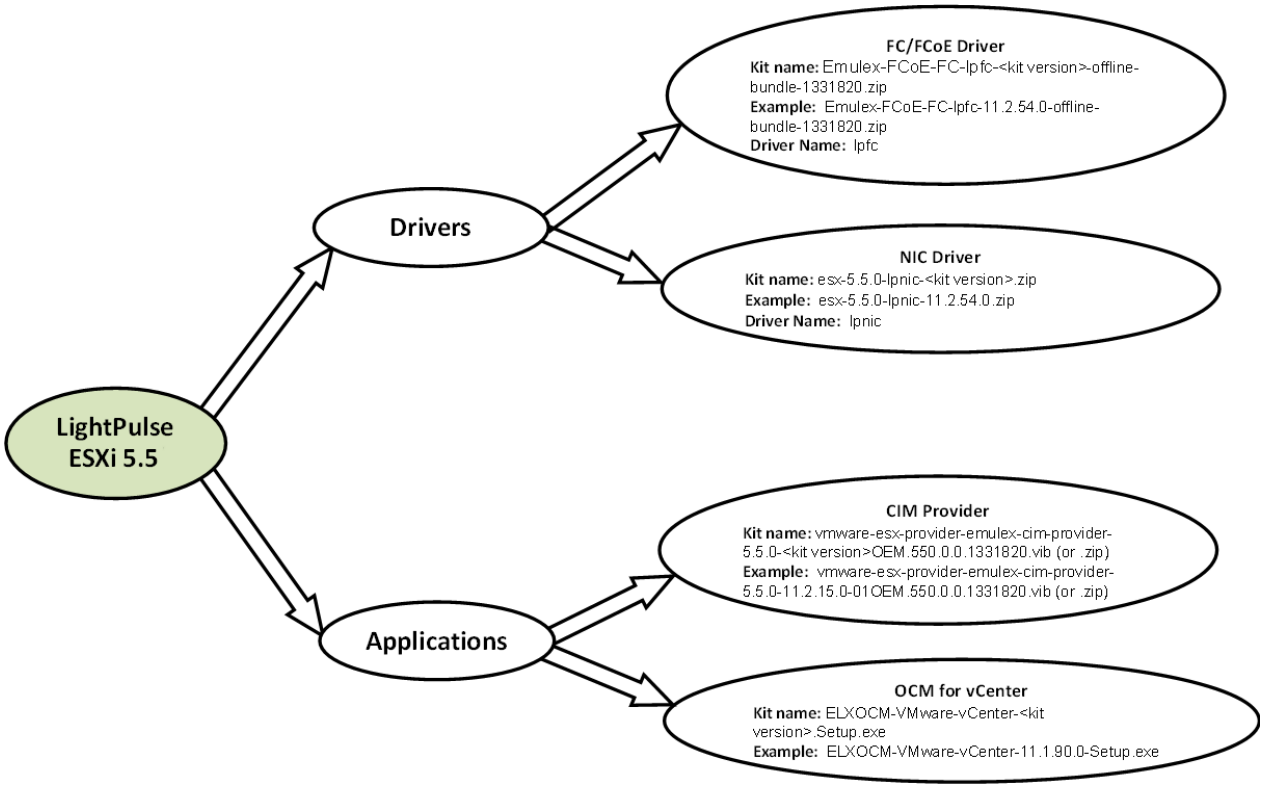


Figure 14 LightPulse ESXi 6.0

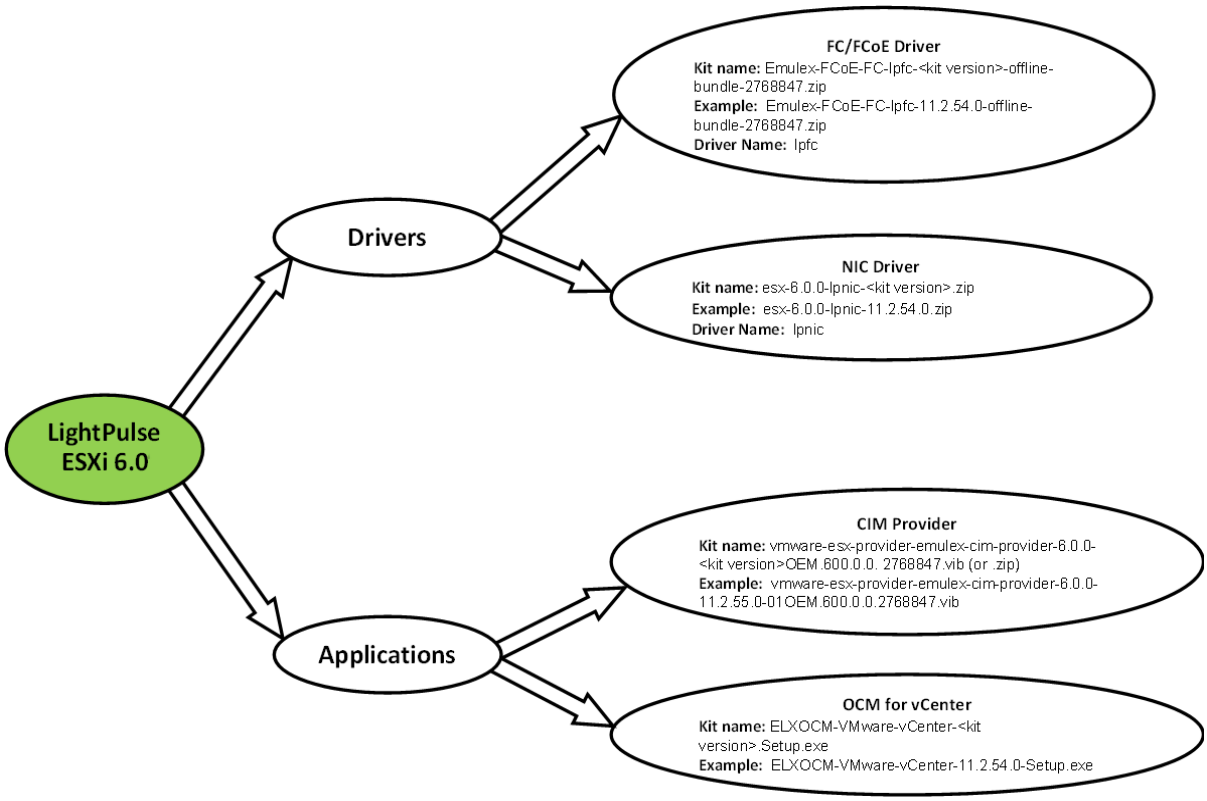
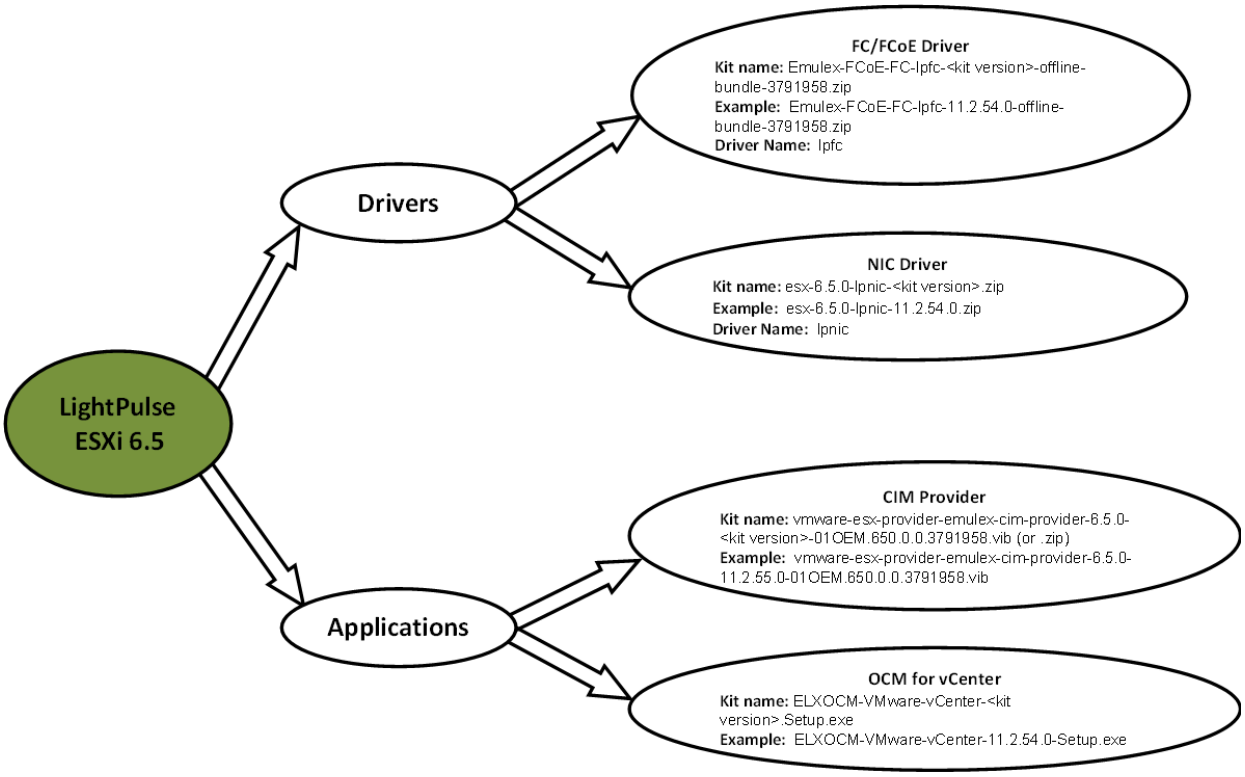
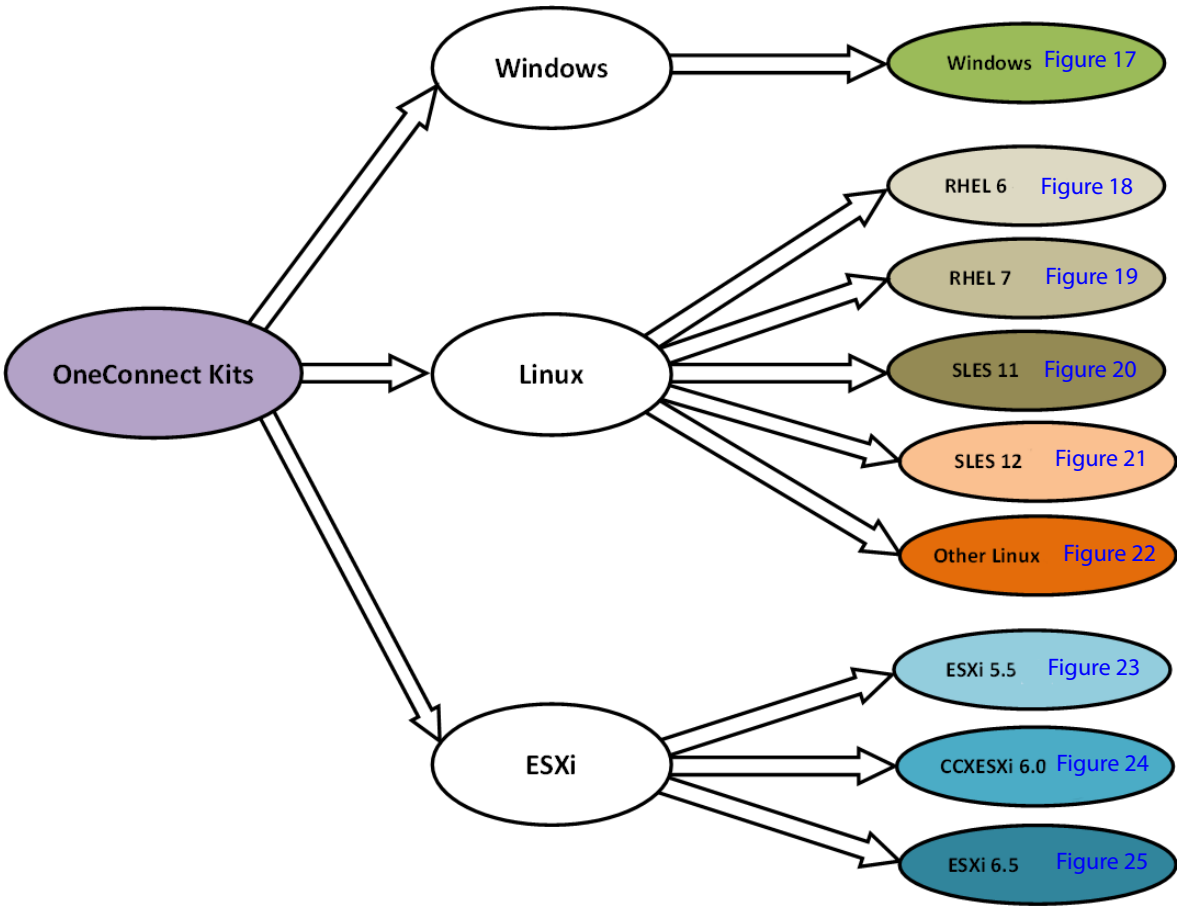


Figure 15 LightPulse ESXi 6.5



B.1.2 OneConnect Adapters and Kits

Figure 16 OneConnect Kits



OneConnect	
Windows	Figure 17
RHEL 6	Figure 18
RHEL 7	Figure 19
SLES 11	Figure 20
SLES 12	Figure 21
Other Linux	Figure 22
ESXi 5.5	Figure 23
ESXi 6.0	Figure 24
ESXi 6.5	Figure 25

Figure 17 OneConnect Windows

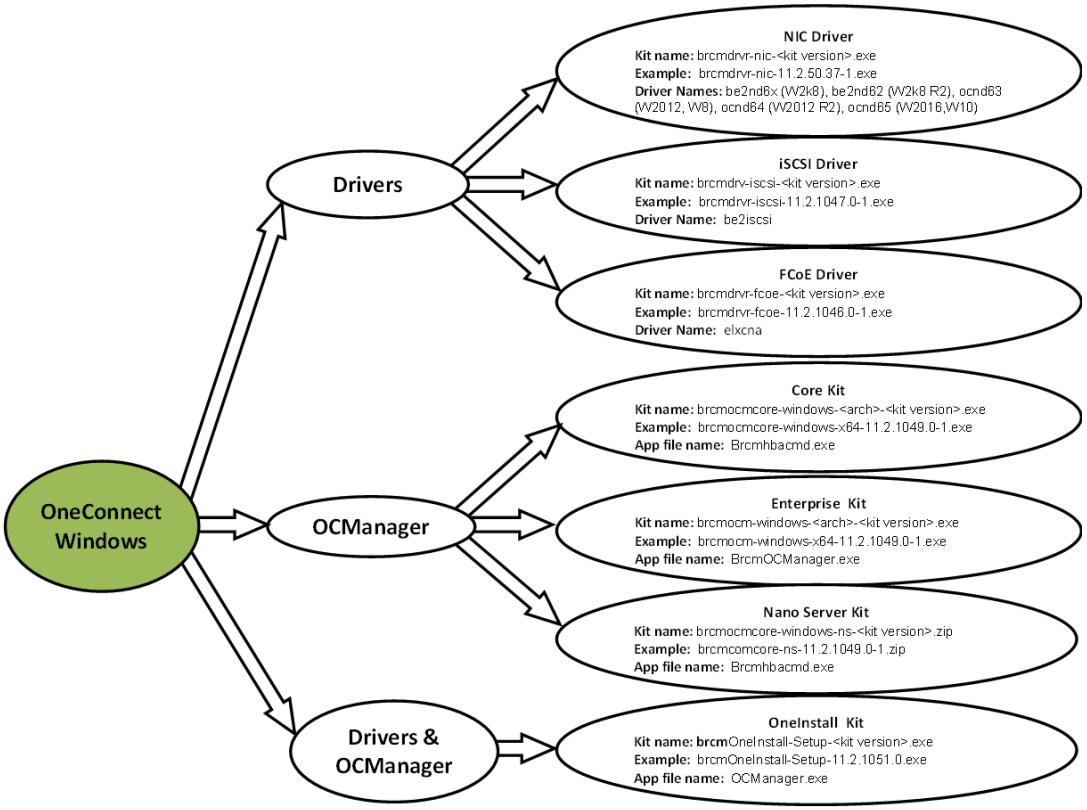


Figure 18 OneConnect RHEL 6

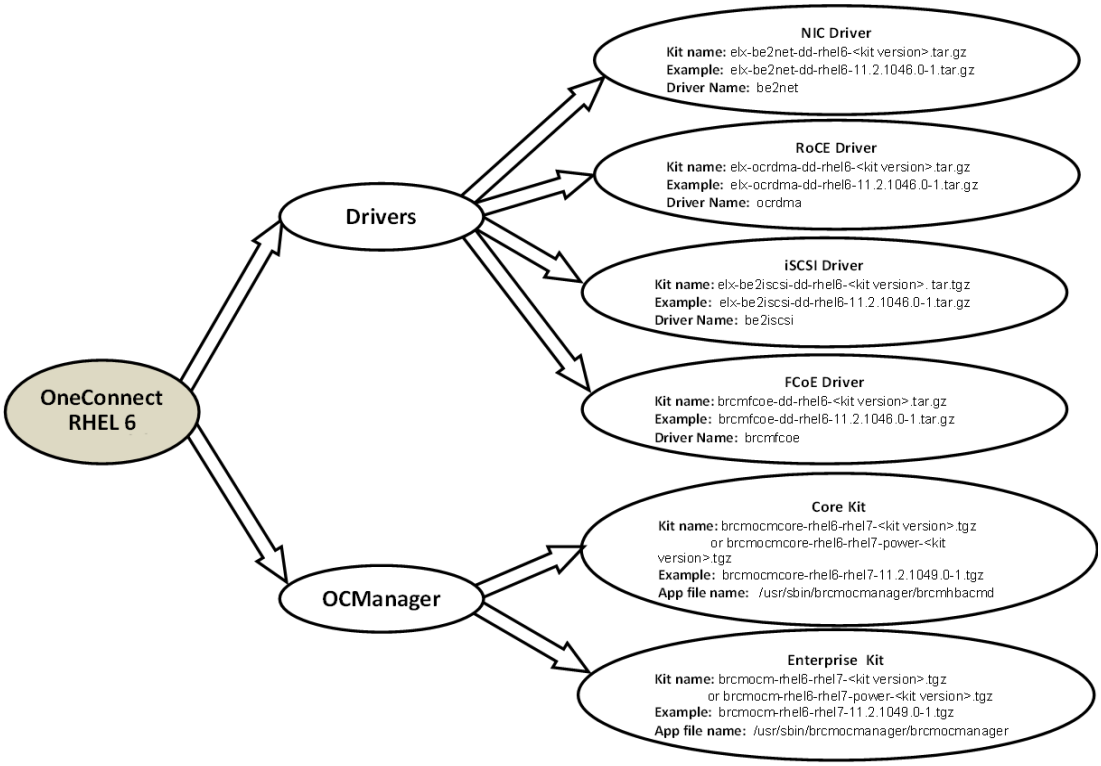


Figure 19 OneConnect RHEL 7

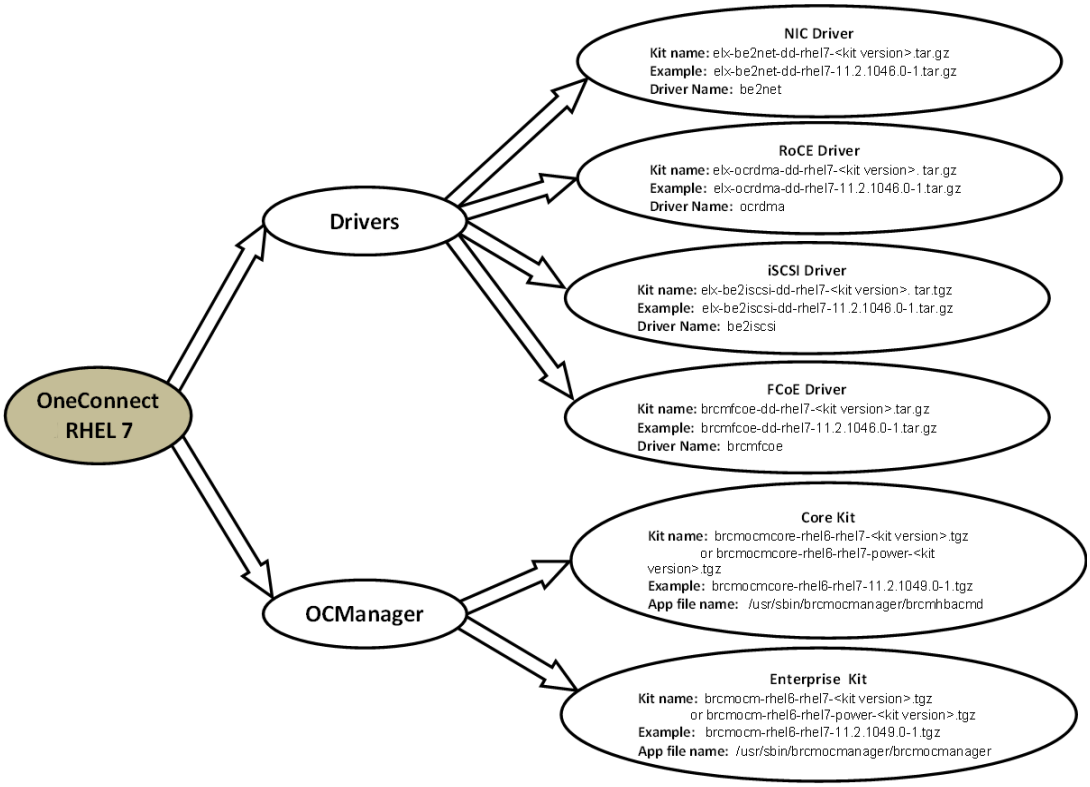


Figure 20 OneConnect SLES 11

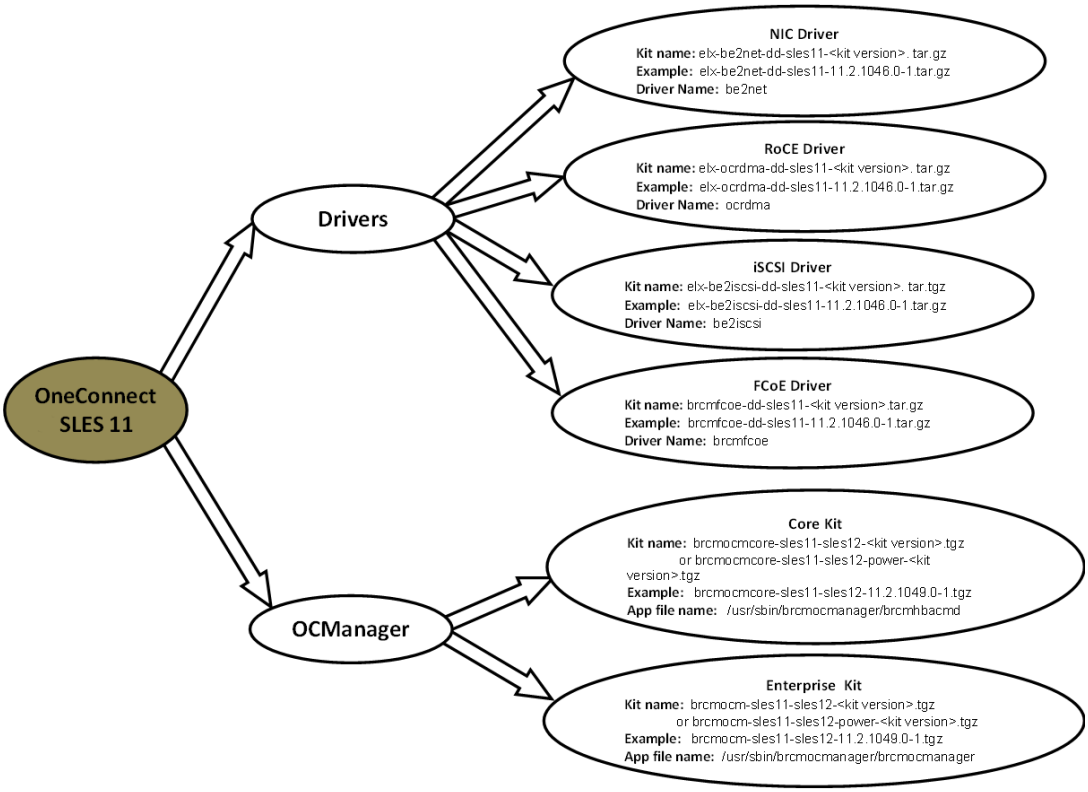


Figure 21 OneConnect SLES 12

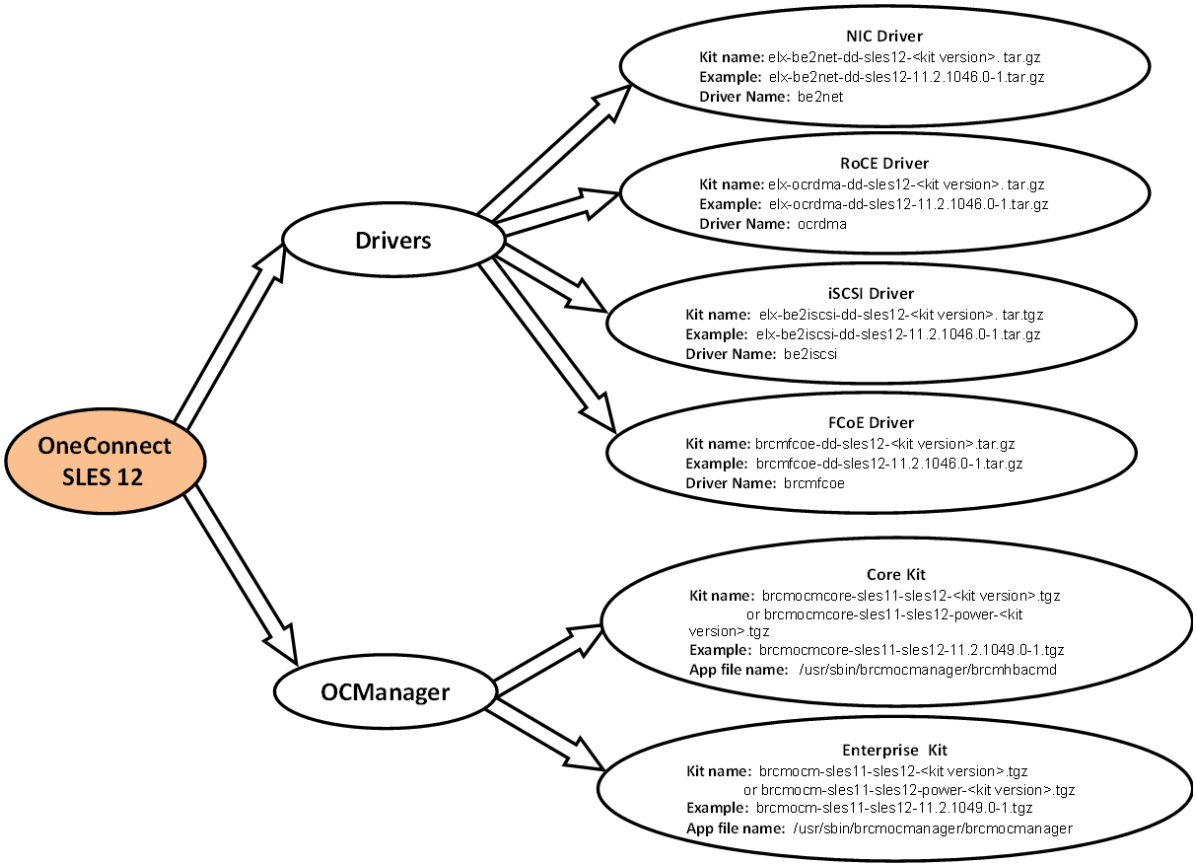


Figure 22 OneConnect Other Linux

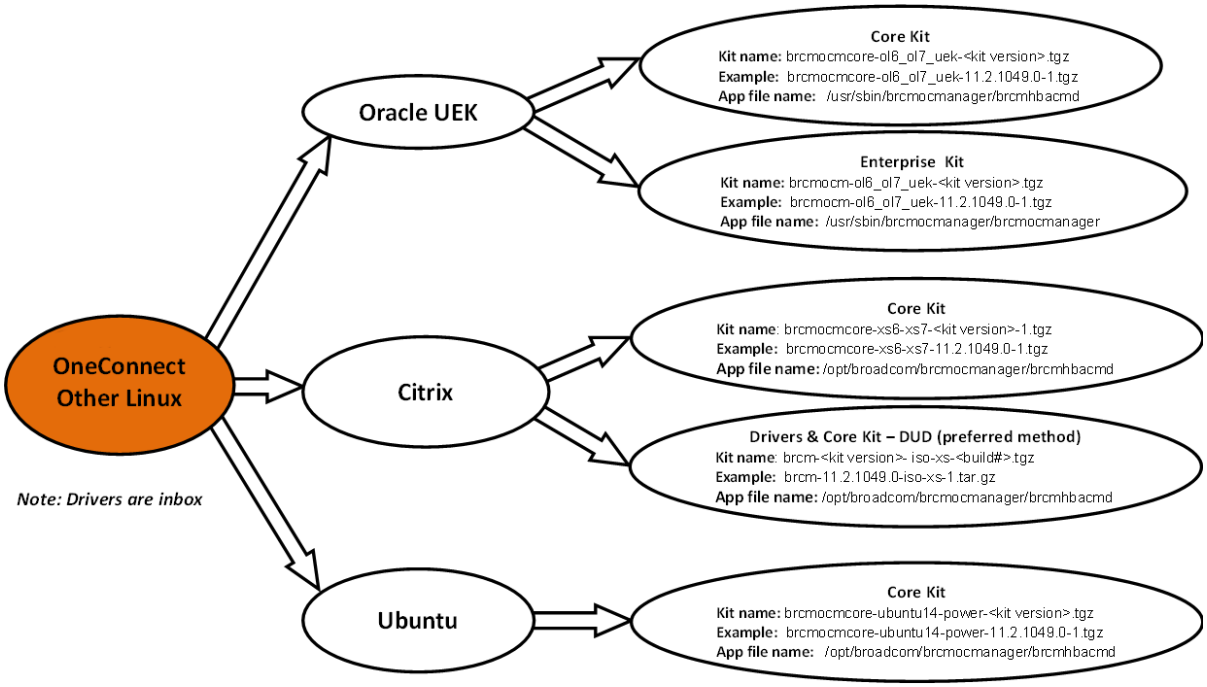


Figure 23 OneConnect ESXi 5.5

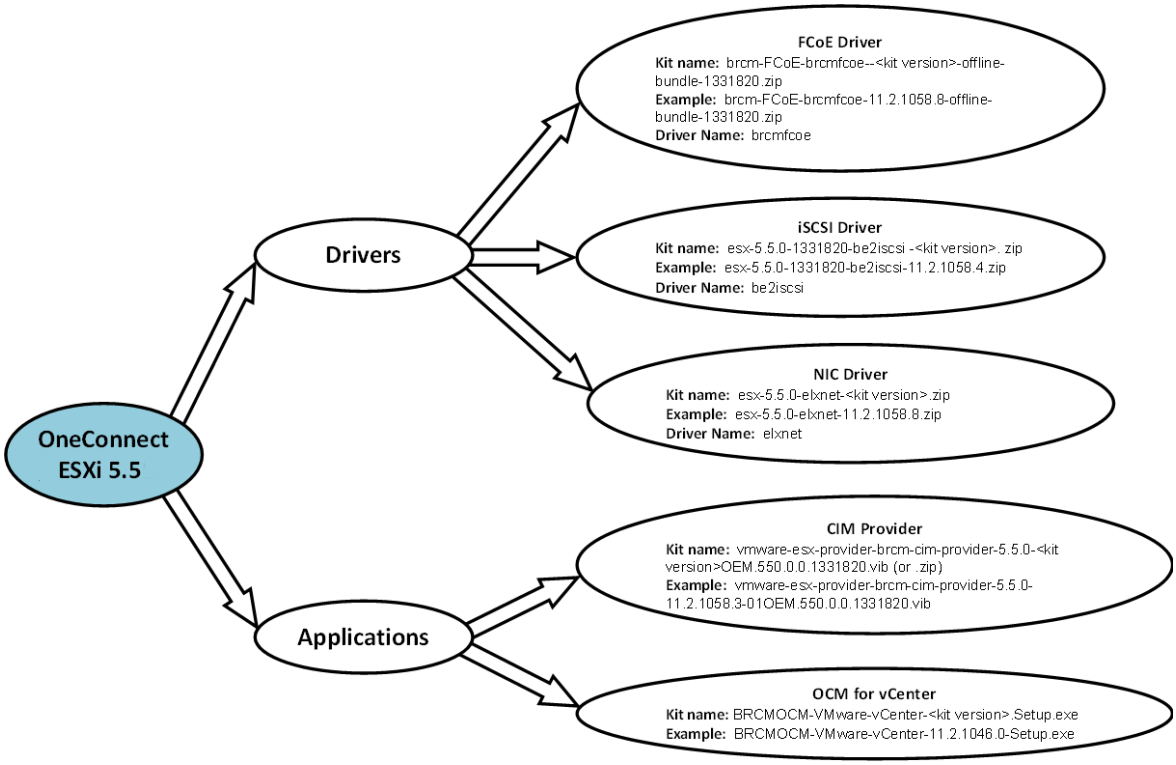


Figure 24 OneConnect ESXi 6.0

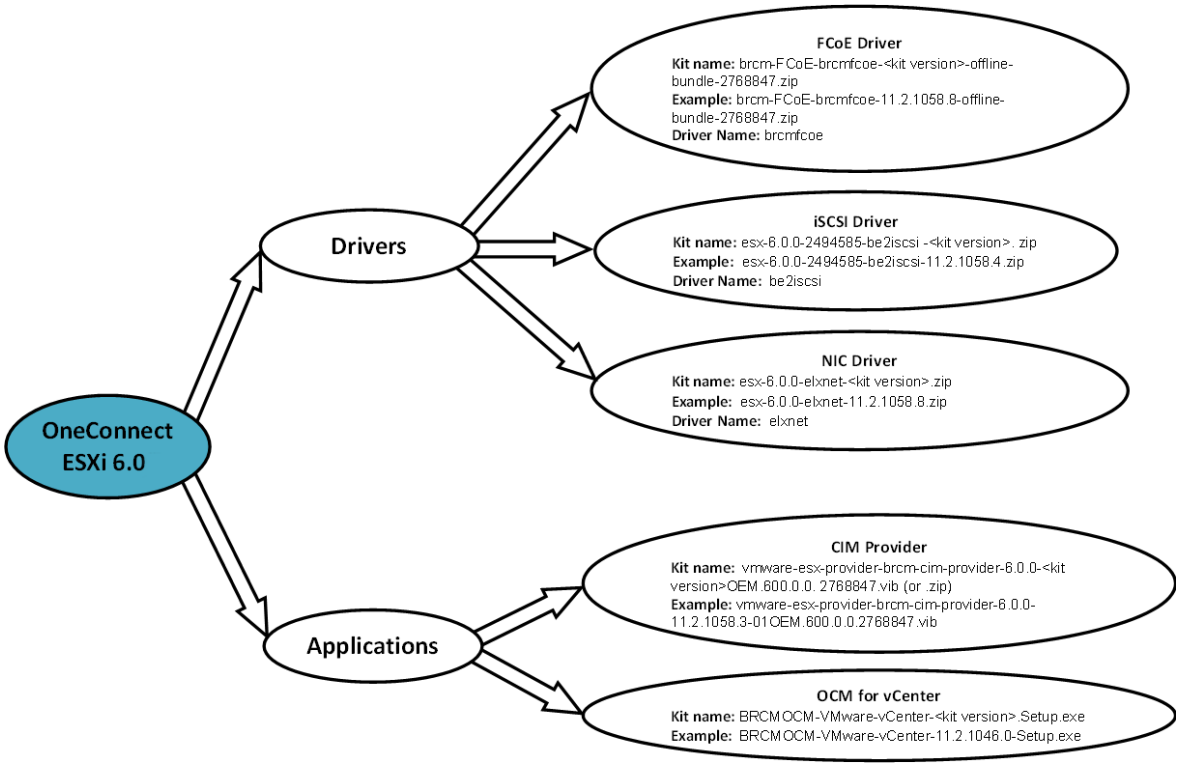
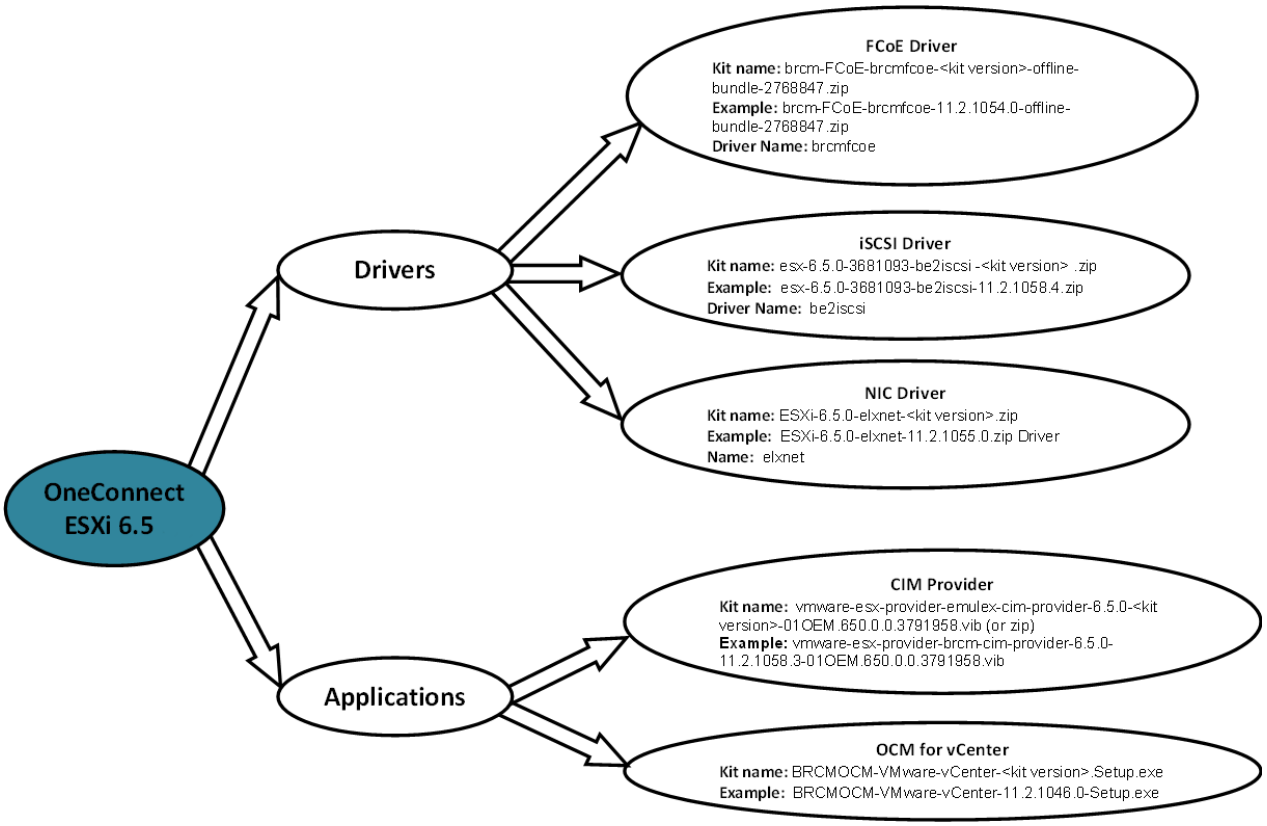


Figure 25 OneConnect ESXi 6.5



Appendix C: Special Instructions

C.1 Linux

Uninstalling the brcmfcoe driver using the brcmfcoe_install.sh script (`brcmfcoe_install.sh -u`) will remove the brcmfcoe driver and unblacklist the lpfc driver.

C.2 VMware

If you need to restore the inbox driver, you can do this in two ways:

- Reboot your server with a local console connected to the server. As POST is completing, watch for instructions to press `shift R`. This allows the administrator to restore the original pristine kernel.

  When the server has rebooted, check the vmhba enumeration to driver binary as follows:

  ```
  esxcli storage core adapter list
  ```

- VMware provides a **depot** zip file with the original inbox drivers. The original lpfc driver is included in this depot zip file. Remove the OOB driver(s) and install the original lpfc driver bundle by typing:

  ```
  esxcli software vib remove –n lpfc (if necessary)
  esxcli software vib remove –n brcmfcoe (if necessary)
  esxcli software vib install -d <path to inbox bundle zip file>
  reboot
  ```

  When the server has rebooted, check the vmhba enumeration to driver binary as follows:

  ```
  esxcli storage core adapter list
  ```

All installation instructions state:

`“When the server has rebooted, check the vmhba enumeration…”.`

The ESXi device manager modifies vmhba numbering because the split driver installation modifies the PCI ID ownership to potentially different driver names. For more information about renumbering, see the following example:

Example

Prior to split kit installation:

NOTE The vmhba instances are vmhba2 to 9. In particular, vmhba2 to 5 are FC and vmhba6 to 9 are FCoE.

```
root@lucic:~] esxcli storage core adapter list
vmhba2  lpfc  link-up  fc.20000090fa021916: (0000:08:00.0) Emulex Corporation Emulex
                       10000090fa021916     LightPulse LPe16000 PCIe Fibre Channel
                                            Adapter
vmhba3  lpfc  link-up  fc.20000090fa021917: (0000:08:00.1) Emulex Corporation Emulex
                       10000090fa021917     LightPulse LPe16000 PCIe Fibre Channel
                                            Adapter
vmhba4  lpfc  link-up  fc.20000000c9f3f9b6: (0000:06:00.0) Emulex Corporation Emulex
                       10000000c9f3f9b6     LPe12000 8Gb PCIe Fibre Channel Adapter
vmhba5  lpfc  link-up  fc.20000000c9f3f9b7: (0000:06:00.1) Emulex Corporation Emulex
                       10000000c9f3f9b7     LPe12000 8Gb PCIe Fibre Channel Adapter
```

```
vmhba6   lpfc   link-up   fc.20000000c9ef3dab: (0000:09:00.2) ServerEngines Corporation
                          10000000c9ef3dab     Emulex OneConnect OCe11100 FCoE Initiator
vmhba7   lpfc   link-up   fc.20000000c9ef3daf: (0000:09:00.3) ServerEngines Corporation
                          10000000c9ef3daf     Emulex OneConnect OCe11100 FCoE Initiator
vmhba8   lpfc   link-n/a  fc.20000090fa5d2f29: (0000:23:00.2) Emulex Corporation Emulex
                          10000090fa5d2f29     OneConnect OCe14000, FCoE Initiator
vmhba9   lpfc   link-n/a  fc.20000090fa5d2f31: (0000:23:00.3) Emulex Corporation Emulex
                          10000090fa5d2f31     OneConnect OCe14000, FCoE Initiator
```

Install the new OOB LightPulse/OneConnect driver bundles:

```
[root@lucic:~] esxcli software vib install -d
Emulex-FCoE-FC-lpfc-11.2.50.60-offline-bundle-2768847.zip
```

Installation result:

Message: `The update completed successfully, but the system needs to be rebooted for the changes to be effective.`

```
Reboot Required: true
VIBs Installed: EMU_bootbank_lpfc_11.2.50.60-1OEM.600.0.0.2768847
VIBs Removed: Emulex_bootbank_lpfc_11.1.166.0-1OEM.550.0.0.1331820
VIBs Skipped:
[root@lucic:~] esxcli software vib install -d
brcm-FCoE-brcmfcoe-11.2.41.63-offline-bundle-2768847.zip
```

Installation result:

Message: `The update completed successfully, but the system needs to be rebooted for the changes to be effective.`

```
Reboot Required: true
VIBs Installed: EMU_bootbank_brcmfcoe_11.2.41.63-1OEM.600.0.0.2768847
VIBs Removed:
VIBs Skipped:
[root@lucic:~] reboot
```

After installing OOB LightPulse/OneConnect drivers:

NOTE vmhba2 to vmhba5 are still FC. However, vmhba6 to9 have been renumbered to vmhba64 to vmhba67. The re-enumeration of the vmhbas is current because the paths are still visible.

```
root@lucic:~] esxcli storage core adapter list
HBA Name Driver    Link      UID                  Capabilities        Description
                   State
vmhba2   lpfc      link-up   fc.20000090fa021916: Second Level Lun ID Emulex Corporation
                             10000090fa021916     (0000:08:00.0)      Emulex LightPulse
                                                                      LPe16000 PCIe Fibre
                                                                      Channel Adapter
vmhba3   lpfc      link-up   fc.20000090fa021917: Second Level Lun ID Emulex Corporation
                             10000090fa021917     (0000:08:00.1)      Emulex LightPulse
                                                                      LPe16000 PCIe Fibre
                                                                      Channel Adapter
vmhba64  brcmfcoe  link-up   fc.20000000c9ef3dab: Second Level Lun ID ServerEngines
                             10000000c9ef3dab     (0000:09:00.2)      Corporation Emulex
                                                                      OneConnect OCe11100
                                                                      FCoE Initiator
```

```
vmhba4   lpfc      link-up   fc.20000000c9f3f9b6: Second Level Lun ID  Emulex Corporation
                             10000000c9f3f9b6     (0000:06:00.0)      Emulex LPe12000 8Gb
                                                                      PCIe Fibre Channel
                                                                      Adapter
vmhba65  brcmfcoe  link-up   fc.20000000c9ef3daf: Second Level Lun ID  ServerEngines
                             10000000c9ef3daf     (0000:09:00.3)      Corporation Emulex
                                                                      OneConnect OCe11100
                                                                      FCoE Initiator
vmhba5   lpfc      link-up   fc.20000000c9f3f9b7: Second Level Lun ID  Emulex Corporation
                             10000000c9f3f9b7     (0000:06:00.1)      Emulex LPe12000 8Gb
                                                                      PCIe Fibre Channel
                                                                      Adapter
vmhba66  brcmfcoe  link-n/a  fc.20000090fa5d2f29: Second Level Lun ID  Emulex Corporation
                             10000090fa5d2f29     (0000:23:00.2)      Emulex OneConnect
                                                                      OCe14000, FCoE
                                                                      Initiator
vmhba67  brcmfcoe  link-n/a  fc.20000090fa5d2f31: Second Level Lun ID  Emulex Corporation
                             10000090fa5d2f31     (0000:23:00.3)      Emulex OneConnect
                                                                      OCe14000, FCoE
                                                                      Initiator
```

```
[root@lucic:~] esxcli storage core path list | grep -e vmhba2 -e vmhba64
Runtime Name:vmhba64:C0:T0:L0
Adapter:vmhba64
Runtime Name:vmhba2:C0:T0:L0
Adapter:vmhba2
Runtime Name:vmhba2:C0:T0:L1
Adapter:vmhba2
Runtime Name:vmhba2:C0:T0:L2
Adapter:vmhba2
Runtime Name:vmhba2:C0:T0:L3
Adapter:vmhba2
Runtime Name:vmhba64:C0:T1:L0
Adapter:vmhba64
```

This renumbering can be repaired by editing the `/etc/vmware/esx.conf` file:

1. Change the directory to `/etc/vmware` and save the existing `esx.conf` to `esx.conf.old`.
2. Open the `esx.conf` file for editing. Search for the vmhba6 -> vmhba9 and vmhba64 -> vmhba67 instances (following this example), delete those lines, save the file, and reboot.

After the `esx.conf` file has been fixed, the output looks like the following:

```
[root@lucic:~] esxcli storage core adapter list
HBA Name Driver    Link      UID                  Capabilities        Description
                   State
vmhba2   lpfc      link-up   fc.20000090fa021916: Second Level Lun ID  Emulex Corporation
                             10000090fa021916     (0000:08:00.0)      Emulex LightPulse
                                                                      LPe16000 PCIe Fibre
                                                                      Channel Adapter
vmhba3   lpfc      link-up   fc.20000090fa021917: Second Level Lun ID  Emulex Corporation
                             10000090fa021917     (0000:08:00.1)      Emulex LightPulse
                                                                      LPe16000 PCIe Fibre
                                                                      Channel Adapter
```

```
vmhba4   lpfc      link-up   fc.20000000c9f3f9b6: Second Level Lun ID  Emulex Corporation
                             10000000c9f3f9b6     (0000:06:00.0)      Emulex LPe12000 8Gb
                                                                      PCIe Fibre Channel
                                                                      Adapter
vmhba5   lpfc      link-up   fc.20000000c9f3f9b7: Second Level Lun ID  Emulex Corporation
                             10000000c9f3f9b7     (0000:06:00.1)      Emulex LPe12000 8Gb
                                                                      PCIe Fibre Channel
                                                                      Adapter
vmhba6   brcmfcoe  link-up   fc.20000000c9ef3dab: Second Level Lun ID  ServerEngines
                             10000000c9ef3dab     (0000:09:00.2)      Corporation Emulex
                                                                      OneConnect OCe11100
                                                                      FCoE Initiator
vmhba7   brcmfcoe  link-up   fc.20000000c9ef3daf: Second Level Lun ID  ServerEngines
                             10000000c9ef3daf     (0000:09:00.3)      Corporation Emulex
                                                                      OneConnect OCe11100
                                                                      FCoE Initiator
vmhba8   brcmfcoe  link-n/a  fc.20000090fa5d2f29: Second Level Lun ID  Emulex Corporation
                             10000090fa5d2f29     (0000:23:00.2)      Emulex OneConnect
                                                                      OCe14000, FCoE
                                                                      Initiator
vmhba9   brcmfcoe  link-n/a  fc.20000090fa5d2f31: Second Level Lun ID  Emulex Corporation
                             10000090fa5d2f31     (0000:23:00.3)      Emulex OneConnect
                                                                      OCe14000, FCoE
                                                                      Initiator
```

Helpful VMware KB articles can be found by using the following link:

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2091560


BROADCOM

Emulex® OneCommand® CNA Manager Application for OneConnect Adapters

User Guide

Version 11.2
December 30, 2016

OCM-APP-OCA-UG112

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

Broadcom, the pulse logo, Connecting everything, Avago Technologies, Avago, the A logo, Emulex, OneCommand, and OneConnect are among the trademarks of Broadcom in the United States, certain other countries and/or the EU.

Copyright © 2003-2017 Broadcom. All Rights Reserved.

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

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.

Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Chapter 1: Introduction

The Emulex OneCommand CNA 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 CNA Manager application contains a graphical user interface (GUI) and a command line interface (CLI). Refer to the *Emulex OneCommand Manager Command Line Interface for OneConnect Adapters User Guide* for information about installing and using the CLI.

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

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

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

1.1 Abbreviations

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
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
FCoE	Fibre Channel over Ethernet
FIP	FCoE Initialization Protocol
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
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

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
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
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
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 Multichannel
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

Chapter 2: Installing and Uninstalling OneCommand CNA Manager Application Components

This section describes installing and uninstalling the OneCommand CNA Manager application.

2.1 Installing the OneCommand CNA Manager Application

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

2.1.1 In Windows

The OneCommand CNA Manager application can be installed two ways:

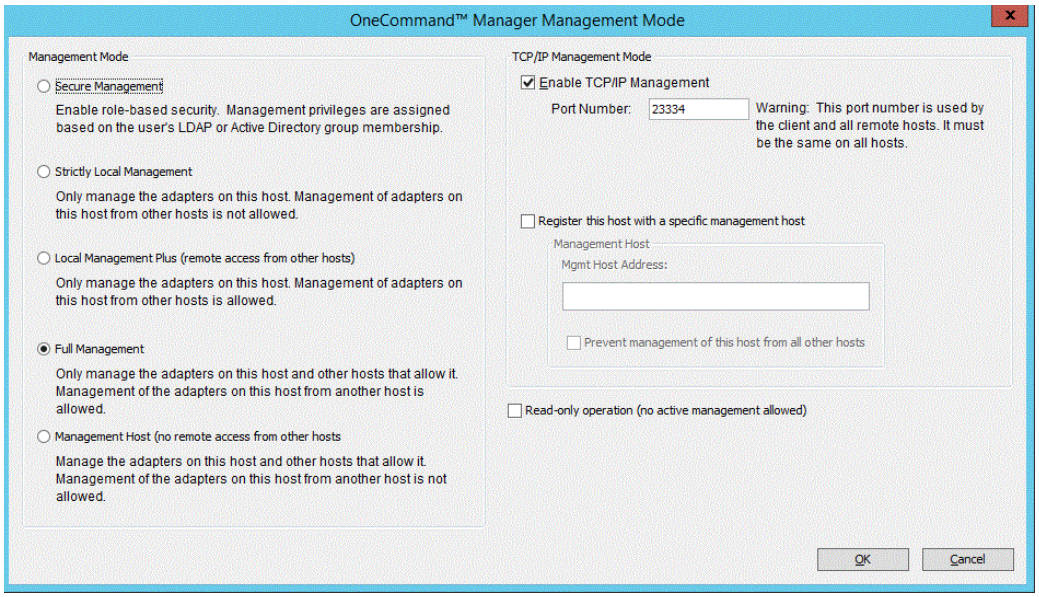
- Attended installation using the GUI.
- Unattended installation using the command line.

2.1.1.1 Attended Installation in Windows

To install the OneCommand CNA Manager application in Windows, perform the following steps:

1. Download the x64 or x86 OneCommand CNA Manager Enterprise Kit installation file from the Documents and Downloads area of http://www.broadcom.com.
2. Navigate to the directory to which you downloaded the file.
3. Double-click the `brcmocm-windows-<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 CNA Manager application files, the Management Mode dialog box appears (Figure 1).

Figure 1 Management Mode Dialog Box



The Management Mode dialog box (Figure 1) enables you to select Secure Management to assign the desired user privileges, or you can choose one of the other management modes. See Section 4.2, Using OneCommand CNA Manager Secure Management or Section 4.3, Changing Management and Read-Only Mode for more information. Choose the management type you want and click **OK**.

5. Check or clear the Enable TCP/IP Management checkbox to enable or disable remote management over TCP/IP. You can also change the TCP/IP port used (23334 is the IANA registered port for Emulex).
6. 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.

2.1.1.2 Unattended Installation in Windows

To install the OneCommand CNA Manager application in Windows, perform the following steps:

1. Download the x64 or x86 OneCommand CNA Manager Enterprise Kit installation file to your system from the Documents and Downloads area of http://www.broadcom.com.
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 Section 4.2, Using OneCommand CNA Manager Secure Management for more information.

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

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

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

   ```
   brcmocm-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 following example. Section 4.3, Changing Management and Read-Only Mode for more information.

   NOTE If you enabled Secure Management in step 3 and entered an mmode value, it results in a 'conflicting parameters' error.

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

   ```
   brcmocm-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 CNA 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 23334 is used.

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

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

2.1.2 In Linux

NOTE The OneCommand CNA Manager application GUI is not supported on Citrix XenServer; however, the OneCommand CNA Manager application CLI is supported. Refer to the *Emulex OneCommand CNA Manager Command Line Interface for OneConnect Adapters User Guide* for Citrix instructions.

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

- The appropriate driver version for your operating system. See the Documents and Downloads area of http://www.broadcom.com for the latest drivers.

  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.

2.1.2.1 Attended Installation in Linux

To install the OneCommand CNA Manager application, or to update an existing installation, in Linux, perform the following steps:

1. Log on as root.
2. Download the utilities from the Documents and Downloads area of http://www.broadcom.com.
3. Copy the OneCommand `brcmocm-<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 6 and RHEL 7, type the following:

```
tar zxvf brcmocm-rhel6-rhel7-<apps_ver>-<rel>.tgz
```

— For SLES 11 and SLES 12, type the following:

```
tar zxvf brcmocm-sles11-sles12-<apps_ver>-<rel>.tgz
```

6. Change to the brcmocm directory created in step 3.

— For RHEL 6 and RHEL 7, type the following:

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

— For SLES 11 and SLES12, type the following:

```
cd brcmocm-sles11-sles12-<apps_ver>-<rel>
```

NOTE Prior to installation, OneCommand CNA Manager application groups must be configured on the LDAP network or the local host machine for Secure Management operation. See Section 4.2.1, OneCommand CNA Manager Secure Management Configuration Requirements for configuration instructions.

7. Run the install script. Type the following:

```
./install.sh
```

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

```
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 23334).

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 23334.)

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 CNA Manager application to operate in read-only mode. Read-only mode prevents users from performing some operations such as resetting adapters, updating an adapter's firmware and changing adapter driver properties and bindings. It only affects the local OneCommand CNA 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.

2.1.2.2 Unattended Installation in Linux

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

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

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

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

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

```
#./install.sh –q2
```

NOTE The Management Mode default for unattended installation is Local Management Plus.

2.1.2.3 Updating an Installation in Linux

The OneCommand CNA Manager application supports the following update paths:

- You can update from an earlier Core Kit to a later Enterprise Kit.
- You can update from an earlier Enterprise Kit to a later Enterprise Kit.

See Section 2.1.2.1, Attended Installation in Linux or Section 2.1.2.2, Unattended Installation in Linux for instructions.

2.1.3 In Solaris

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

- The latest Solaris out-of-box FCoE driver for FCoE management.
- The latest Solaris out-of-box NIC driver for NIC management.

NOTE The NIC driver must be installed and reporting all NIC ports. Otherwise, the OneCommand CNA Manager application cannot manage the CNAs.

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

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

   ```
   tar xvf brcmocm-solaris-<version>.tar
   ```
3. Change to the newly created `brcmocm-solaris-<version>` directory:

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

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

4. Run the install script to begin installation. If the HBAnyware utility, OneCommand CNA Manager Core or OneCommand CNA 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 CNA:

   ```
   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 23334.

   If you select options 2 or 3, you are prompted for the management host address. Leaving the field blank means none.

   You can enter an IP address or host name. If you enter a management host address, you are prompted to exclude management of this host from any other host.

   If you select option 4, management of remote hosts is automatically selected and you are prompted to enter the TCP/IP port number to use. Leaving the field blank defaults to 23334.

   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 CNA Manager application to operate in read-only mode. Read-only mode prevents users from performing some operations such as resetting adapters, updating an adapter's firmware and changing adapter driver properties and bindings. It only affects the local OneCommand CNA 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.

2.1.4 In VMware

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

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

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

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-brcm-cim-provider-<version>.zip  --maintenance-mode
```

3. Reboot the system.

2.2 Installing the OneCommand CNA Manager Application Web Launch Interface

Considerations

- The OneCommand CNA 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 CNA Manager Secure Management mode is not supported for OneCommand CNA Manager application Web Launch clients.

In addition to the driver and OneCommand CNA Manager application, the prerequisites in Section 2.2.1, Requirements must be met before you install the Web Launch Interface.

2.2.1 Requirements

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

2.2.1.1 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.

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

To enable the HTTP port, perform the following steps:

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 CNA Manager JRE installation path. For example:

   ```
   C:\Program Files\broadcom\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 CNA Manager JRE installation path. For example:

   ```
   C:\Program Files\broadcom\util\JRE\bin\rmiregistry.exe.
   ```
1. Click **OK**.
2. Click **OK** to apply the new firewall settings.

To add the MIME type, perform the following steps:

1. Start 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` to the File name extension field.
9. Add `application/x-java-jnlp-file` to the MIME type field.
10. Click **OK**.

2.2.1.2 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 8.0 or later for Intel based systems (x86)
  - 32-bit Java 8.0 or later for x86-64 systems
  - 32-bit Java 8.0 or later for SLES 10 (ppc64)
  - 32-bit Java 8.0 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/brcmocm/README_WEBLAUNCH.txt (Solaris) or /usr/sbin/brcmocmanager/README_WEBLAUNCH.txt (Linux) for more setup information.

2.2.2 Installation

To install the OneCommand CNA Manager application Web Launch Interface, perform the following steps:

- In Windows:

  Click **Programs>Broadcom>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 CNA Manager directory.
   - In Solaris:
     ```
     cd /opt/brcmocm
     ```
   - In Linux:
     ```
     cd /usr/sbin/brcmocmanager
     ```
3. Run the install script. Type the following:
   ```
   ./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.
7. After you have entered the necessary information, you are notified when the installation of the OneCommand CNA Manager application Web Launch package is complete. The Web Launch configuration files are created and Web Launch Service automatically starts.
8. To verify the installation, locate another client, open a web browser window, and enter the following URL:
   ```
   http://IP_ADDR:PORT_NUM/Brcmocmanager.jnlp
   ```
   where IP_ADDR is the IP address of the host on which you installed the OneCommand CNA Manager application Web Launch service, and PORT_NUM is the TCP port number of the listening host's web server. The standard OneCommand CNA Manager application user interface appears.

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

2.2.2.1 Java Exception Site List Maintenance for Java 8

The OneCommand CNA Manager Web Launch application now uses the Java 8 runtime. Because of the strict security requirements of Java 8, you must enter the URI of the site that is hosting the Web Launch Server into a security

exception site list maintained on the browser host. If this is not done, you will be unable to run the Web Launch application from your browser.

To add a Web Launch server site to the browser's security exception site list, perform the following steps:

1. Make sure that Java 8 is installed on the browser host.
2. Locate the Java Control Panel for Java 8 on the browser host. The location of the Java Control Panel executable file depends on the operating system type.

   In Windows:

   Navigate to the location where the Java 8 runtime has been installed. For example:

   ```
   C:\Program Files\Java\jre1.8.x.
   ```

   The Java Control Panel executable is located under the `bin` subdirectory, and is named `javacpl.exe`. For example:

   ```
   C:\Program Files\Java\jre1.8.x\bin\javacpl.exe
   ```

   In Solaris and Linux:

   Navigate to the location where the Java 8 runtime has been installed. The Java Control Panel executable file is located under the `bin` subdirectory of the Java runtime installation location, and is named `ControlPanel`. For example, If the runtime is installed in `/opt/jre1.8.x`, the full path to this executable files is:

   ```
   /opt/jre1.8.x/bin/ControlPanel
   ```
3. Run the Java Control Panel program, and then click on the **Security** tab. The current Exception Site List is displayed.
4. Click on the **Edit Site List** button to the right of the Exception Site List.
5. Click the **Add** button, and then enter the full URI of the OneCommand CNA Manager Web Launch application. The format of this URI is:

   ```
   http://IP-ADDR:PORT_NUM/Brcmocmanager.jnlp
   ```

   `IP-ADDR` is the IP address of the host on which you installed the OneCommand CNA Manager application Web Launch service, and `PORT_NUM` is the TCP port number that the web server is listening on. For example, if the Web Launch service is installed on a server with the IP address of 3.4.5.6, and the TCP port number on which the web server is listening is 5555, you must enter the following site exception:

   ```
   http://3.4.5.6:5555/Brcmocmanager.jnlp
   ```
6. Click **Apply** to save the setting, and then exit the application.

2.3 Uninstalling the OneCommand CNA Manager Application

To uninstall the OneCommand CNA Manager application and OneCommand CNA Manager application Web Launch Interface, perform the following steps:

- In Windows:

1. Select **Start>Control Panel>Programs>Uninstall a Program**.
2. Select **Broadcom OCManager Enterprise [version]** and click **Remove** or **Uninstall**.

- In Linux:

1. Log on as root.
2. Change to the `brcmocm-<platform>-<version>` installation directory.

3. Type the following:

```
./uninstall
```

- In Solaris:

1. Log on as root.
2. Run the OneCommand CNA Manager application uninstall script:

```
/opt/brcmocm/scripts/uninstall
```

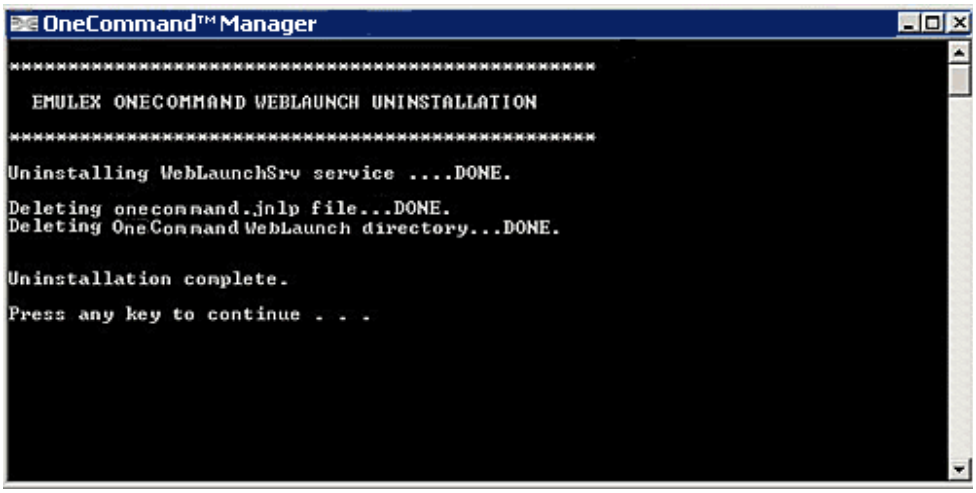
2.4 Uninstalling the OneCommand CNA Manager Application Web Launch Interface Only

To uninstall the OneCommand CNA Manager application Web Launch, but leave the OneCommand CNA Manager application installed, perform the following steps:

- In Windows:

1. Select **Start>Programs>Broadcom>OCManager WebLaunch Uninstall**. The OneCommand CNA Manager application Web Launch Uninstallation Screen appears (Figure 2).

Figure 2 OneCommand CNA Manager Application Web Launch Uninstallation Screen



2. The OneCommand CNA Manager application Web Launch Interface is removed. Press any key to continue.

- In Linux and Solaris:

1. Log on as root.
2. Run the uninstallation script:
   - In Linux:

     ```
     /usr/sbin/brcmocmanager/wsuninstall
     ```

   - In Solaris:

     ```
     /opt/brcmocm/wsuninstall
     ```

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

Chapter 3: Starting and Stopping the OneCommand CNA Manager Application

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

3.1 In Windows

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

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

3.2 In Linux and Solaris

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

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

`brcmhbamgrd` starts at system boot time. `brcmdiscoveryd` starts whenever the OneCommand CNA Manager application GUI process first runs on the host machine.

Additionally, if the Web Launch component of OneCommand CNA 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.

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

- Linux: /usr/sbin/brcmocmanager
- Solaris: /opt/brcmocm

  Example usage on Linux:

  ```
  # /usr/sbin/brcmocmanager/brcmocmanager
  ```

  Example usage on Solaris:

  ```
  # /opt/brcmocm/start_brcmocmanager
  ```

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

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

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

Example usage on Linux:

```
# /usr/sbin/brcmocmanager/stop_brcmocmanager -n
```

Example usage on Solaris:

```
# /opt/brcmocm/stop_brcmocmanager -n
```

NOTE The -n option stops only the OneCommand CNA Manager application and associated elxdiscoveryd daemon. When the `stop_brcmocmanager` script is run without options, the script stops the OneCommand CNA Manager application and all associated daemons.

3.3 Starting the OneCommand CNA Manager Application Web Launch Interface

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

Only the OneCommand CNA Manager application GUI is exported to the requesting client. All adapter detection 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.

NOTE Client operating systems do not support Web Launch, it is only supported on server systems.

To start the OneCommand CNA Manager application with your web browser, perform the following steps:

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 Section 4.2, Using OneCommand CNA Manager Secure Management 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 CNA Manager application Web Launch Interface software installed and running.

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

   where `IP_ADDR` is the IP address of the host on which you installed the OneCommand CNA 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 CNA Manager application user interface is displayed.

3.3.1 Managing Files When Running the OneCommand CNA Manager Application Web Launch Interface

When running the OneCommand CNA Manager application Web Launch Interface, all OneCommand CNA 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.

Chapter 4: Using the OneCommand CNA Manager Application

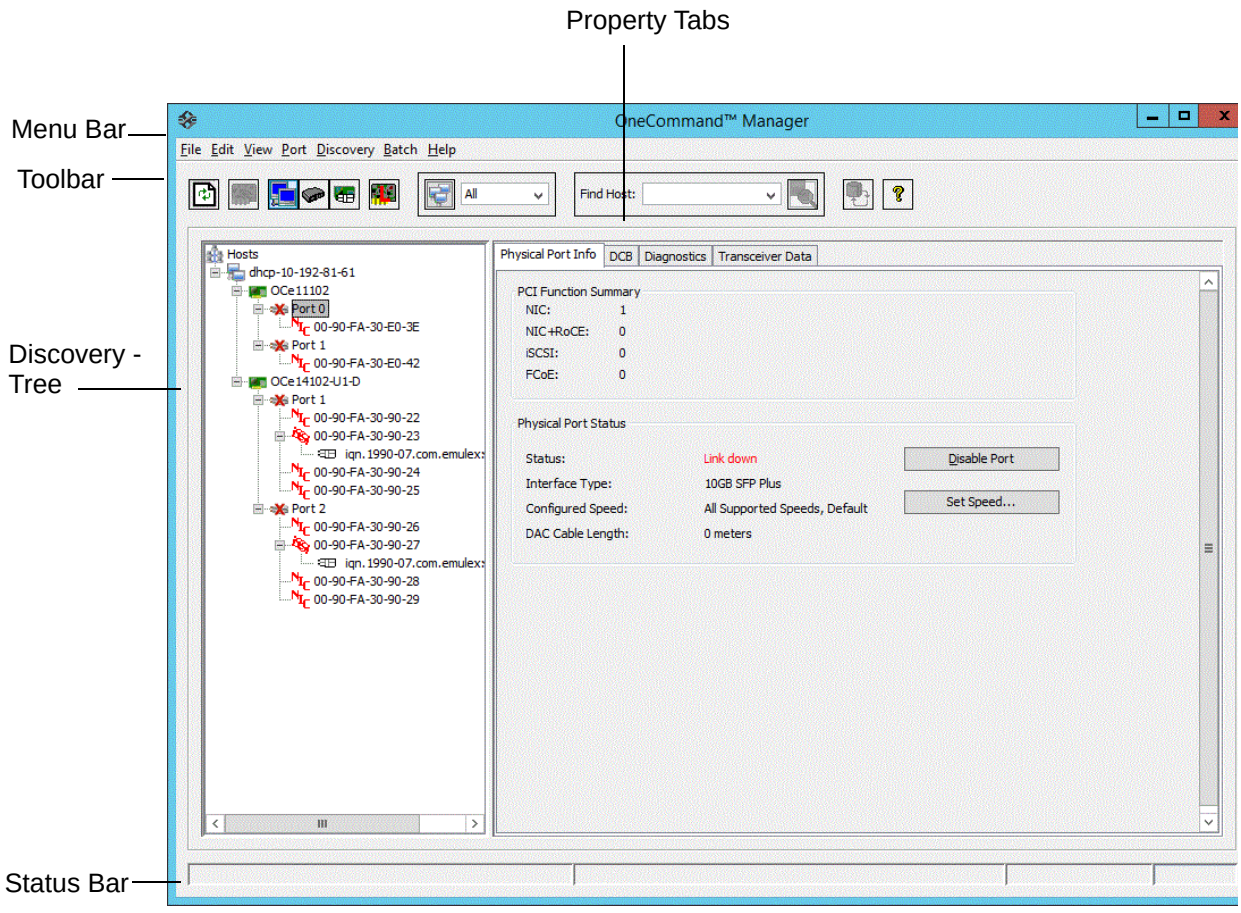
To properly view the OneCommand CNA Manager application, make sure that your system meets the following display requirements:

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

4.1 The OneCommand CNA Manager Application Window Element Definitions

The OneCommand CNA Manager application window (Figure 3) contains five basic components: the menu bar, the toolbar, the discovery-tree, the property tabs, and the status bar.

Figure 3 OneCommand CNA 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 CNA Manager application interface matches those of the remote server. When accessing a remote server running an older version of the OneCommand CNA Manager application, capabilities that are not supported by the server's older version of the OneCommand CNA Manager application are unavailable.
- In some instances, the type of information displayed and available functionality is determined by the operating system in use.

4.1.1 Menu Bar

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

4.1.2 Toolbar

The toolbar (Figure 4) 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 Toolbar



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

4.1.2.1 Toolbar Buttons

The toolbar buttons perform the following tasks:



Discovery Refresh button
- Initiates a discovery refresh cycle.



Reset 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 CNA Manager application displays elements in ascending order.



Host View button (default)
Displays the host system.

NOTE You cannot change host names using the OneCommand CNA 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 iSCSI 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

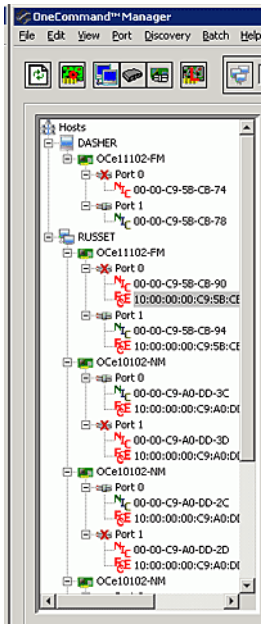
- Displays the OneCommand CNA Manager application's online help.
- Displays the About OneCommand CNA Manager dialog box. The dialog box displays version information including RMAPI, Discovery, DFClib, MILI Library Version (Windows), and Remote Management Agent Version (Windows). It also enables you to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

4.1.3 Discovery-Tree

The discovery-tree (Figure 5) has icons that represent discovered hosts, adapters, ports, virtual ports, fabrics, targets, and LUNs.

Using the View menu, the OneCommand CNA 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 5 Discovery-Tree



4.1.3.1 Discovery-Tree Icons

Discovery-tree icons represent the following:



The local host.



Other hosts connected to the system.



A green adapter icon with black descriptive text represents an online adapter. Blue text represents an adapter port that 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 Section 5.2, Configuring Discovery and Default CIM Credentials for more information about discovery settings.



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

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



The 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.

NOTE RoCE is available as Technical Preview only. Do not use RoCE in a production environment.



The ASIC node icon, only displayed for dual ASIC adapters, represents each ASIC on the adapter. Each ASIC is managed independently. The ASIC node format **ASIC bus#-sub-adapter#** represents the PCI bus number and the sub-adapter number, which is a concatenation of the discovered port numbers for the ASIC. For example, **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 Media Exchanger icon represents connections to individual media exchangers. A media exchanger is a jukebox-type device that is capable of swapping various media device instances (such as records or CDs) in and out.



The Tape LUN icon represents LUNs that are tape devices.



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



The Switch icon represents connections to the switch.

4.1.3.2 Expanding or Collapsing the Discovery-Tree View

You can use the Expand/Collapse capability on the View menu to change the way discovered elements are displayed. By selecting one of the four levels, the discovery-tree (Figure 5) is expanded or collapsed to that level. You can choose Hosts/Fabrics (depending on the view), adapters, ports, PCI functions, and targets.

4.1.4 Property Tabs

The property tabs display configuration, statistical, and status information for network elements (Figure 3). 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 5).

4.1.5 Status Bar

The status bar is located near the bottom of the OneCommand CNA Manager application window (Figure 3). The status bar displays messages about OneCommand CNA Manager application functions, such as **Discovery in progress** or the progress when performing an **Export SAN Info** operation.

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

4.2 Using OneCommand CNA Manager Secure Management

OneCommand CNA 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 CNA Manager application. If users are authenticated, they can only perform the functions allowed by the OneCommand CNA Manager user group to which they belong. If the systems are running in an LDAP or Active Directory domain, the OneCommand CNA Manager application authenticates users with those defined in that domain. For Linux and Solaris systems, this is accomplished using PAM.

NOTE OneCommand CNA Manager Secure Management is not supported on VMware hosts nor is it supported on OneCommand CNA Manager application Web Launch clients.

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

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

Table 1 Secure Management User Privileges

Group Name	OneCommand CNA 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 CNA Manager application unless they are also members of the ocmadmin group. Otherwise, if secure management is enabled, root users or administrators can only manage local adapters (similar to the ocmlocaladmin users).

Remote management operations between two machines are allowed or denied depending on the OneCommand CNA 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 2 Active Commands: Machines on Same Domain

	Remote Server (Secure)	Remote Server (Not Secure)
Client (Secure)	Allowed	Denied [1]
Client (Not Secure)	Denied	Allowed

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

Table 3 Active Commands: Machines on Different Domain

	Remote Server (Secure)	Remote Server (Not Secure)
Client (Secure)	Denied[1]	Denied [2]
Client (Not Secure)	Denied	Allowed

1. Allowed if the username and password are the same on both domains.
2. To inform you of an unsecured server that you may want to secure.

Table 4 Passive Commands: Machines on Any Domain

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

4.2.1 OneCommand CNA Manager Secure Management Configuration Requirements

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

1. Authentication – On Linux and Solaris systems, this is accomplished using the PAM interface and must be configured as follows:
   - For Solaris systems, place the correct setting in the **auth** section of `/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 CNA Manager Secure Management must be able to access the OneCommand CNA 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 CNA Manager group names are listed with their member users, the machine is ready to use OneCommand CNA Manager secure management.
3. For Solaris systems, you must use **useradd -G** groupname for authentication to work. You cannot use a lowercase **g**.

4.3 Changing Management and Read-Only Mode

NOTE This functionality is 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 Section 4.2, Using OneCommand CNA Manager Secure Management for details.
- Strictly Local Management – This setting allows management of adapters on this host. Management of adapters on this host from other hosts is not allowed.
- Local Management Plus – This setting 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 23334. 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 CNA Manager application discovery server and management agent to use the new settings. If the servers cannot be stopped and restarted, you are prompted to reboot the host for the new TCP/IP management settings to take effect.

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

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

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

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

4.3.1 Management Host

The OneCommand CNA 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.

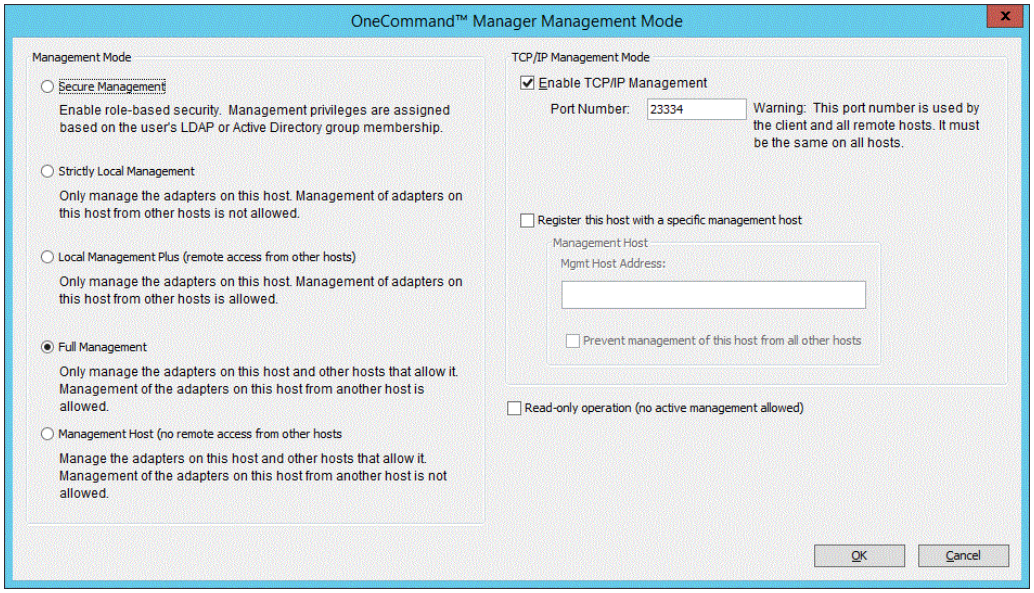
To change management/read-only mode:

NOTE After making changes, you must restart the OneCommand CNA 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 6).

Figure 6 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/brcmocm/set_operating_mode
   ```

2. Choose the management type and read-only mode you want.

- In Linux

1. Stop the OneCommand CNA Manager application.
2. Run the following script:

   ```
   /usr/sbin/brcmocmanager/set_operating_mode
   ```

3. Choose the management type and read-only mode you want.

4.4 Using CIM (Windows Only)

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

To manage the adapters on a VMware host using the OneCommand CNA Manager application, you must install the Emulex CIM Provider on the VMware host.

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

NOTE For VMware hosts, if advanced adapter management capabilities are required (for example, iSCSI Management and port disable), use the OneCommand CNA Manager application for VMware vCenter. For more details, see the *Emulex OneCommand CNA Manager for VMware vCenter for OneConnect Adapters User Guide*.

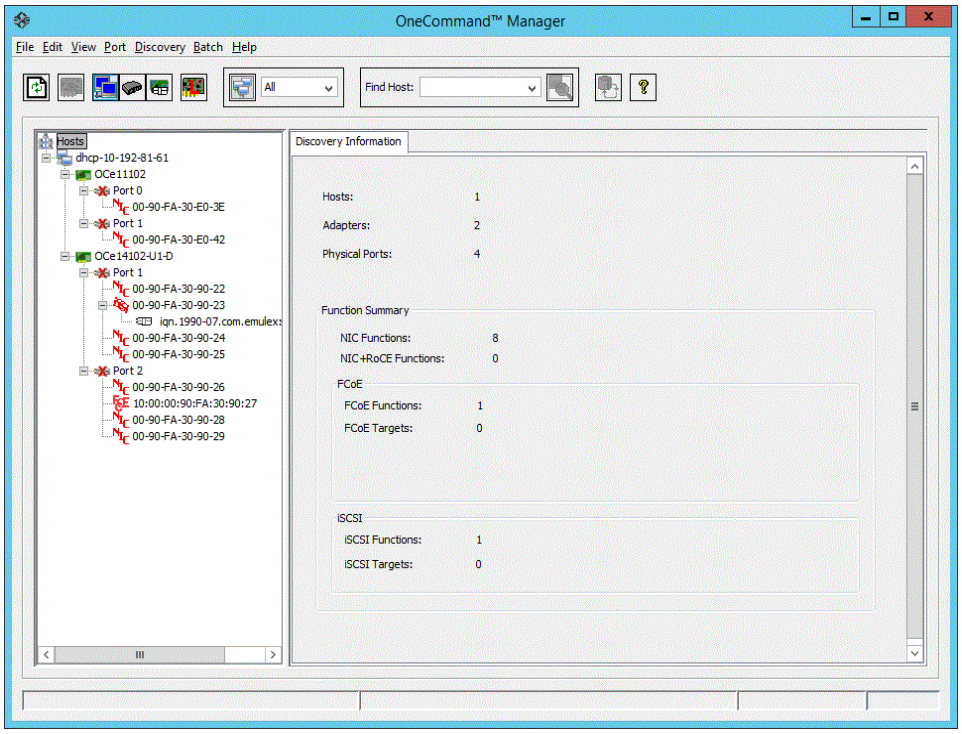
Chapter 5: Configuring Discovery

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

5.1 Discovery Using the TCP/IP Access Protocol

You can discover adapters on IPv4 and IPv6 TCP/IP hosts and on hosts configured to support the CIM interface that have the OneCommand CNA Manager application installed (Figure 7). 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 by using their host IP-address or by the name of the host on which they reside.

Figure 7 Discovery Information



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

```
\Program Files\Broadcom\Util\Common\rmserver.exe
```

5.1.1 Hosts File

The TCP/IP discovery function of the OneCommand CNA 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 8). If the host is discovered, the OneCommand CNA 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 8). For each discovered host, the OneCommand CNA 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 CNA Manager application removes its IP address and name from the host file.
- Adding or removing a host using the CLI.

5.1.1.1 Manually Editing the Hosts File

You can open the hosts file with any text editor, modify the contents and save the file. The name of the 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\broadcom\Util`.
   - Solaris – The file is located in the directory `/opt/brcmocm`.
   - Linux – The file is located in the directory `/usr/sbin/brcmocmanager`.
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 CNA 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 10.192.80.24:23334).
   - 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]:23334`
   - Each line in the file can be up to 1023 characters, although this is longer than is typically needed for a host IP address and host name. A line longer than 1023 characters is truncated, possibly causing discovery to not discover some of the hosts.
   - Blank lines are ignored.
3. Save the file.

5.1.1.2 Copying the File

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

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

5.1.2 Adding a Single Host

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

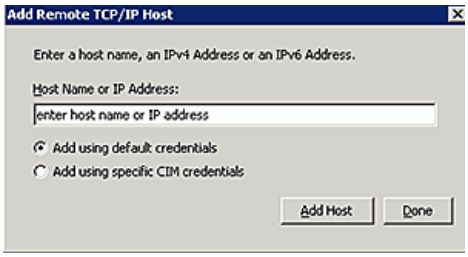
The OneCommand CNA Manager application enables you to specify a single TCP/IP host to manage. You can add a RMAPI host or CIM host using the host name or IP address. If the host is successfully discovered, it is added to the hosts file. If it has not been discovered over FC already, the host and its adapter ports are added to the discovery-tree (Figure 5).

NOTE The OneCommand CNA 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 8).

Figure 8 Add Remote TCP/IP Host Dialog Box



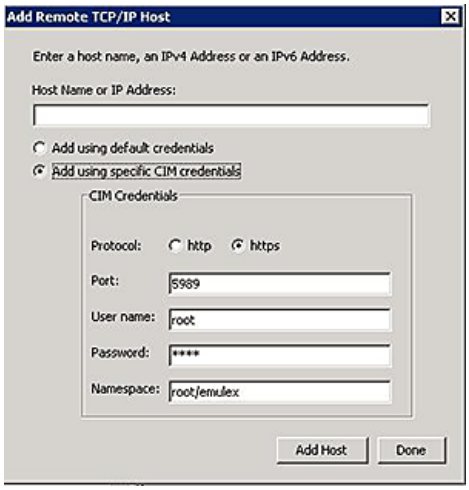
2. Enter the name or the IPv4 or IPv6 address of the host to be added.

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 0s 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 9).

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

Figure 9 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.

5.1.3 Adding a Range of Hosts (IPv4 Only)

NOTE This option is not available in strictly local or local plus management modes.

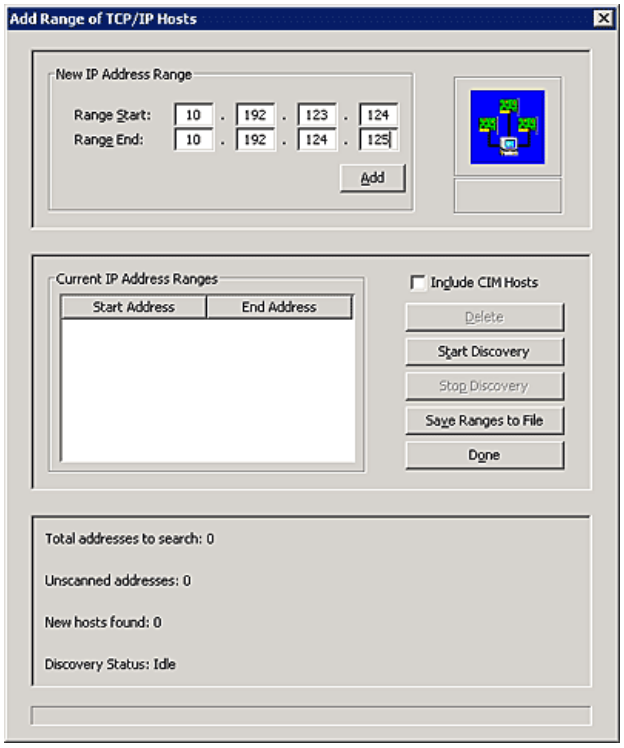
You can find the TCP/IP-accessed manageable hosts by searching a range of IPv4 addresses. The Add Range of TCP/IP Hosts dialog box (Figure 10) enables you to build the initial list of TCP/IP accessed manageable hosts.

NOTE 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.

NOTE The OneCommand CNA Manager application must be installed on all remote hosts.

Figure 10 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 10).
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 5). 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 10).

5.1.4 Removing Hosts

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

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

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.

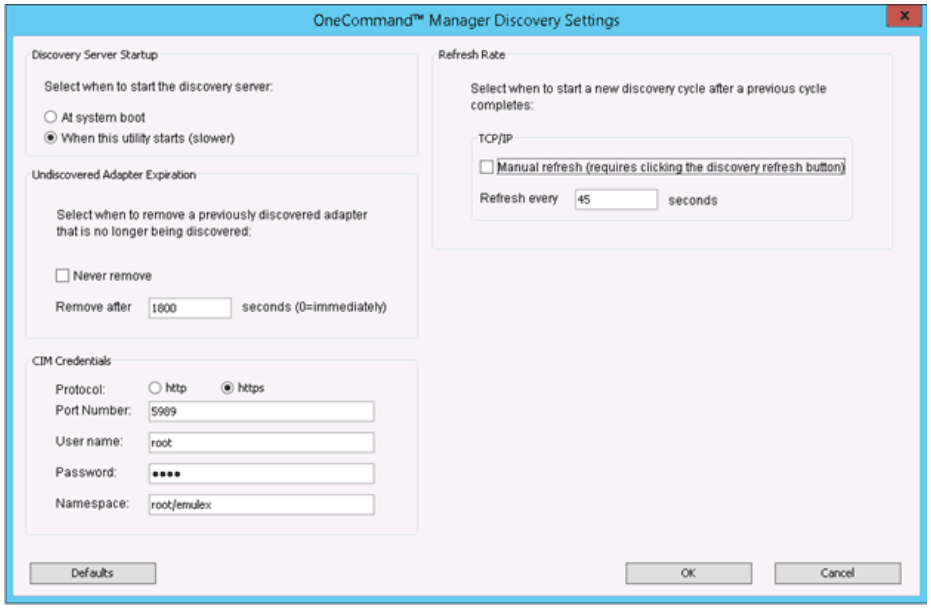
5.2 Configuring Discovery and Default CIM Credentials

Use the OneCommand CNA Manager application Discovery Settings dialog box (Figure 11) 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 5) only once. If the same host name appears for more than one host, the adapters of all these hosts are displayed by the OneCommand CNA Manager application as a single host entry.

Figure 11 Discovery Settings Dialog Box



To configure discovery settings:

1. From the Discovery menu, select **Modify Settings**. The Discovery Settings dialog box appears (Figure 11).
2. Define the discovery properties you want.
3. The CIM credentials area can be used to set the default CIM credentials that are used to connect to all the ESXi hosts that are managed through the CIM interface.

- Protocol: The HTTP or HTTPS protocol can be used to connect to the VMware hosts. The default port numbers used for HTTP and HTTPS are 5988 and 5989 respectively. The port number changes automatically according to the protocol selected. You can also manually change the port number. By default, the HTTP is disabled on sfcb in VMware host, so you must use HTTPS to communicate to the VMware host.
- User name: The user name field contains the username with which to connect to the VMware hosts. By default, this is **root**.
- Password: This password field contains the password of the user name that is used to connect to the VMware host.
- Namespace: Namespace is the namespace of the Emulex provider.

  The default namespace is `root/emulex`.

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

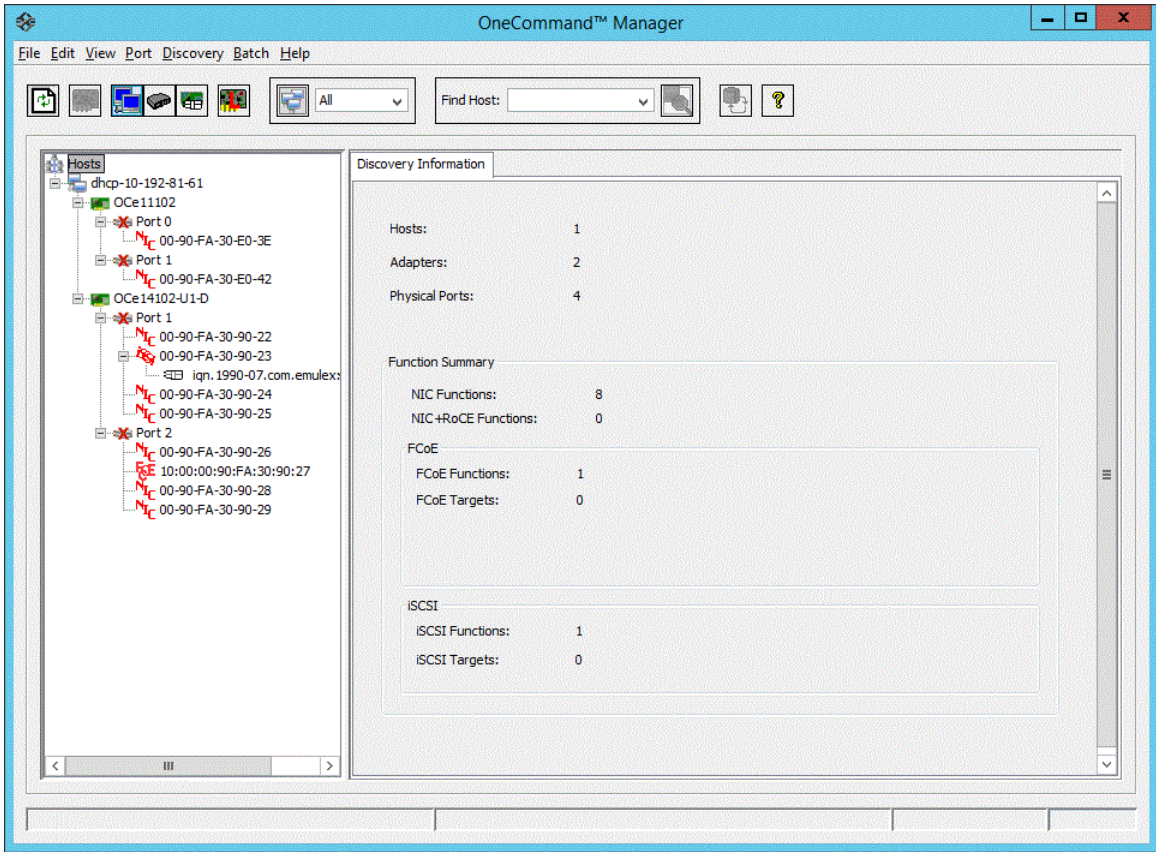
5.3 Viewing Discovery Information

The Discovery Information 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 12 Discovery Information (Host View Selected)



The following Discovery Information fields are displayed:

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

Chapter 6: Managing Hosts

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

6.1 Viewing Host Information

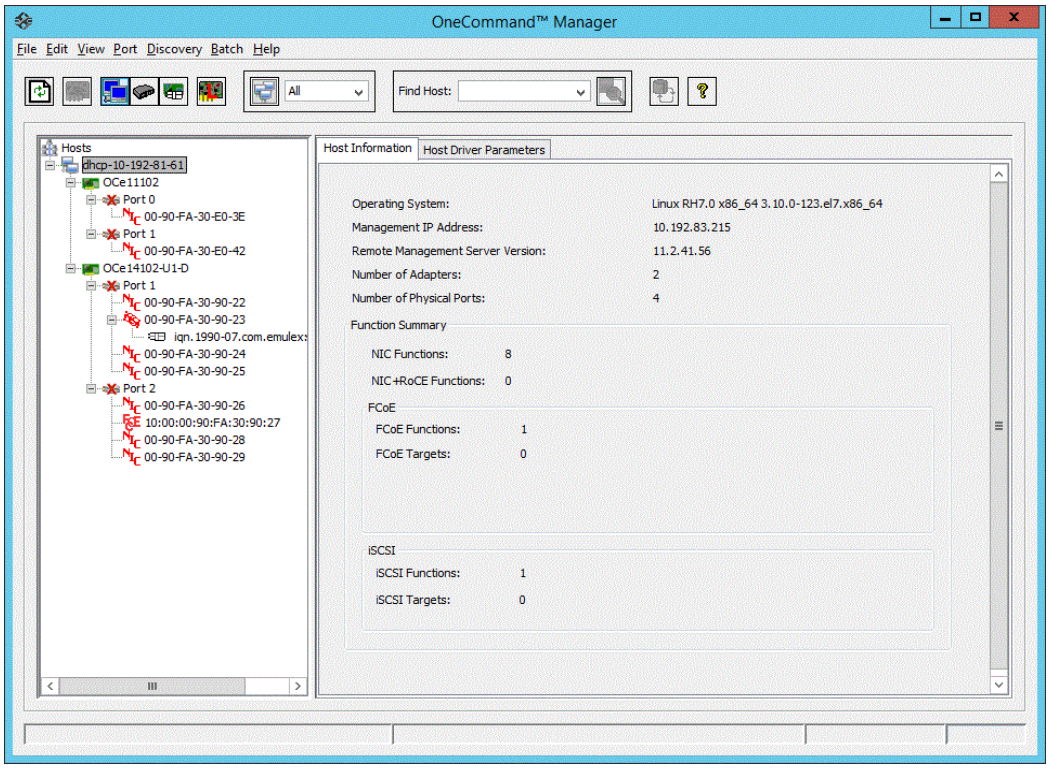
Two tabs show host information: the **Host Information** tab (Figure 13) and the **Host Driver Parameters** tab (Figure 92). The **Host Information** tab is read-only. The **Host Driver Parameters** tab enables you to view and define adapter driver settings for a specific host. See Section 8.2.16.2, The Host Driver Parameters Tab for more information about the **Host Driver Parameters** tab.

To view the **Host Information** and **Host Driver Parameters** tabs:

1. Perform one of the following tasks:
   - 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 13) or the **Host Driver Parameters** tab (Figure 92).

The **Host Information** tab (Figure 13) 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 13 Host Information Tab



The following Host Information fields are displayed:

- Operating System – Details about the installed operating system.

- 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 CNA Manager application.
- Remote Manager Server Version – The version of the OneCommand CNA 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.
  - 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 FCoE functions on this host.
  - VPorts – The number of discovered FCoE 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.

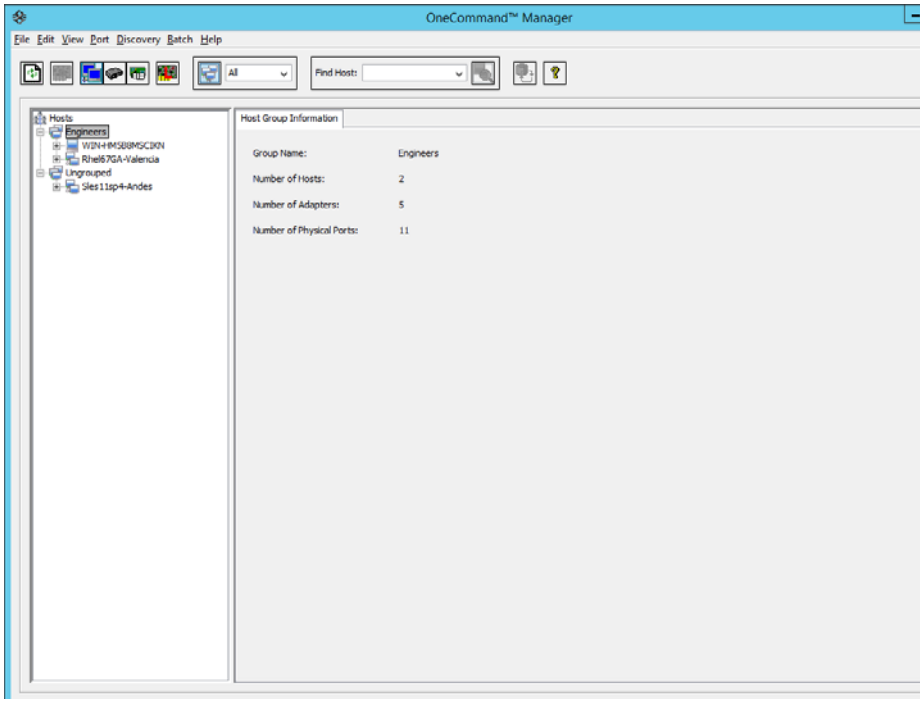
6.2 Viewing Host Grouping Information

The **Host Group Information** tab (Figure 14) displays information about the selected host group, such as the group name and the total number of hosts. See Section 6.3, Grouping Hosts to learn about creating host groups.

NOTE Host grouping is not supported for VMware.

To view host grouping information, from the discovery-tree (Figure 5), select the host group whose information you want to view.

Figure 14 Host Group Information Tab



The following Host Group Information fields are displayed:

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

6.3 Grouping Hosts

The OneCommand CNA 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 5) 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 Section 9.2, Updating Firmware for Multiple Adapters for more information.

NOTE Grouping hosts is not supported on VMware.

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

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

To display all hosts groups:

1. Perform one of the following tasks:
   - 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:

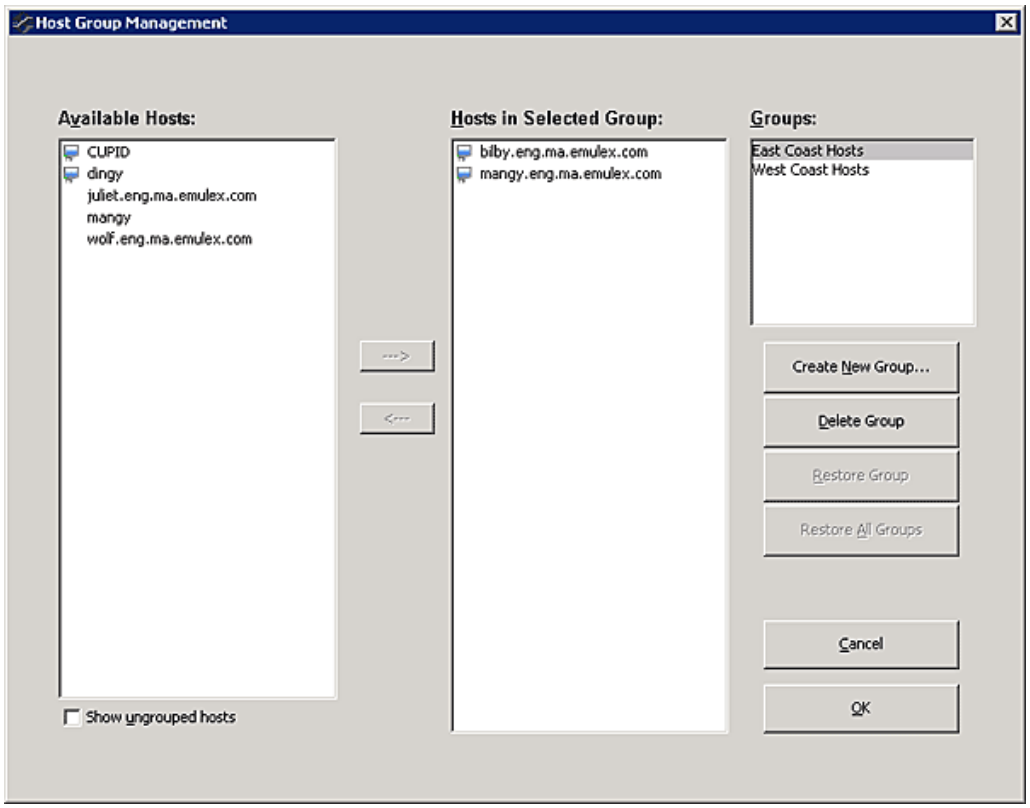
1. Perform one of the following tasks:
   - From the **View** menu, check **Show Groups**.
   - From the toolbar click **Show Host Groups**.
2. From the **Available Host Group** list, choose the group whose hosts you want to view.

6.3.1 Managing Host Groups

Use the Host Group Management dialog box (Figure 15) to create and delete host groups, add and remove hosts, and restore host groups.

NOTE Managing host groups is not supported on VMware.

Figure 15 Host Group Management Dialog Box



6.3.1.1 The following Host Group Management fields are displayed:

- Available Hosts – The list of hosts that can be added to a host group. You can select a host and right-click to see its group assignments.
- Show ungrouped hosts – If 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.

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

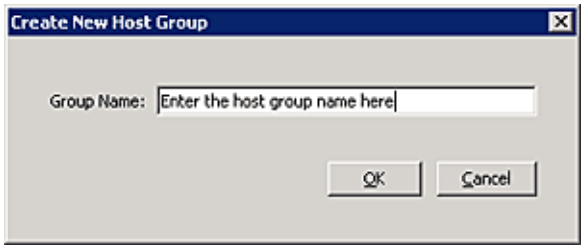
- Indicates the host is currently assigned to a single host group.
- Indicates the host is currently assigned to multiple host groups.

6.3.2 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 15).
2. Click **Create New Group**. The Create New Host Group dialog box is displayed (Figure 16).

Figure 16 Create New Host Group Dialog Box



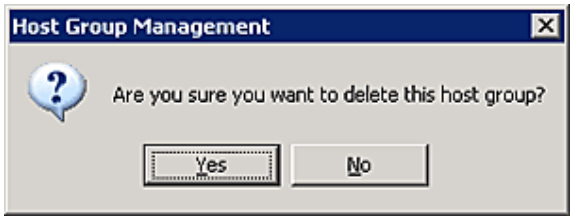
3. Enter the name of the group you want to create and click **OK**. The new group appears in the Groups list on the Host Group Management dialog box.

6.3.3 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 15).
2. From the **Groups** list, select the group you want to delete. The Host Group Management warning dialog box appears (Figure 17).

Figure 17 Host Group Management Warning Dialog Box



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

6.3.4 Adding a Host to a Host Group

To add a host to a group:

1. From the **View** menu, select **Manage Groups**. The Host Group Management dialog box appears (Figure 15).
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 5) displays the new configuration.

6.3.5 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 15).
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 5) displays the new configuration.

6.3.6 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 15).
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.

6.3.7 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 15).
2. Click **Restore All Groups** to return the entire host group configuration to the state that existed when the dialog was opened. All host group assignments are returned to their original configuration. Newly added host groups yet to be committed are removed, and host groups that were deleted are restored.

6.3.8 Exporting Host Grouping Configurations

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

The host groups configuration file locations for the supported platforms are:

- Windows: `<InstallationDriveLetter>:\Program Files\Broadcom\Util\Config\hostgroups`
- Linux: `/usr/sbin/brcmocmanager/config/hostgroups`
- Solaris: `/opt/brcmocm/config/hostgroups`

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

6.4 Searching for Hosts in the Discovery-Tree

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

This capability is especially useful when you are searching for a host in a large installation with hundreds or thousands of hosts. It is also helpful in Fabric view mode because 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.

To search for a host:

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

   The host you are searching for is highlighted in the discovery-tree (Figure 5).

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

Chapter 7: Managing Adapters and Ports

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

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

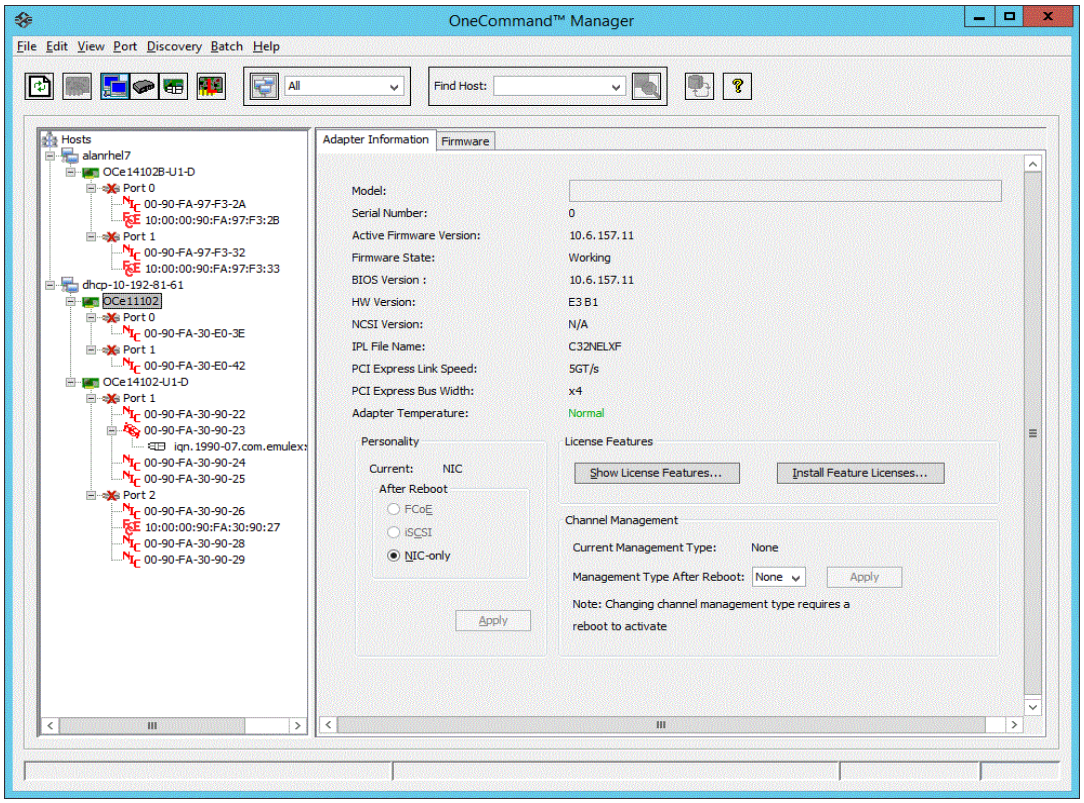
7.1 Viewing Adapter Information

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 Section 7.1.1.1, Configuring UMC Channel Management, Section 7.1.5, Changing Personalities on OneConnect OCe11000-Series Adapters, and Section 7.1.4, Viewing and Installing Licenses for OCe11000-Series Adapters 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 18).

Figure 18 Adapter Information Tab



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

- 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, inoperative 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, inoperative fans, and air conditioning issues that cause high ambient air temperatures.

  After the system overheating issue is resolved and the adapter has cooled down, reboot the system, or if the system supports hot swapping, cycle the power of the adapter slot.
- Personality area (OCe11000-series adapters only):

  NOTE Not available on OneConnect OCe11101-EM/EX or OCe11102-EM/EX adapters.

  - Current – The current personality in use by the adapter.
  - After Reboot – The type of management that is used after a 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.
- License Features area (OCe11000-series adapters only):

  NOTE Not available on OneConnect OCe11101-EM/EX or OCe11102-EM/EX adapters.

  - See Section 7.1.4, Viewing and Installing Licenses for OCe11000-Series Adapters for information.
- Channel Management area (OCe11000-series adapters only):

  NOTE Not available on OneConnect OCe11101-EM/EX or OCe11102-EM/EX adapters.

  - 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.

7.1.1 Viewing Channel Management Information (OCe11000-Series Adapters Only)

The **Channel Management** tab is displayed when an adapter port is selected in the discovery-tree 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.

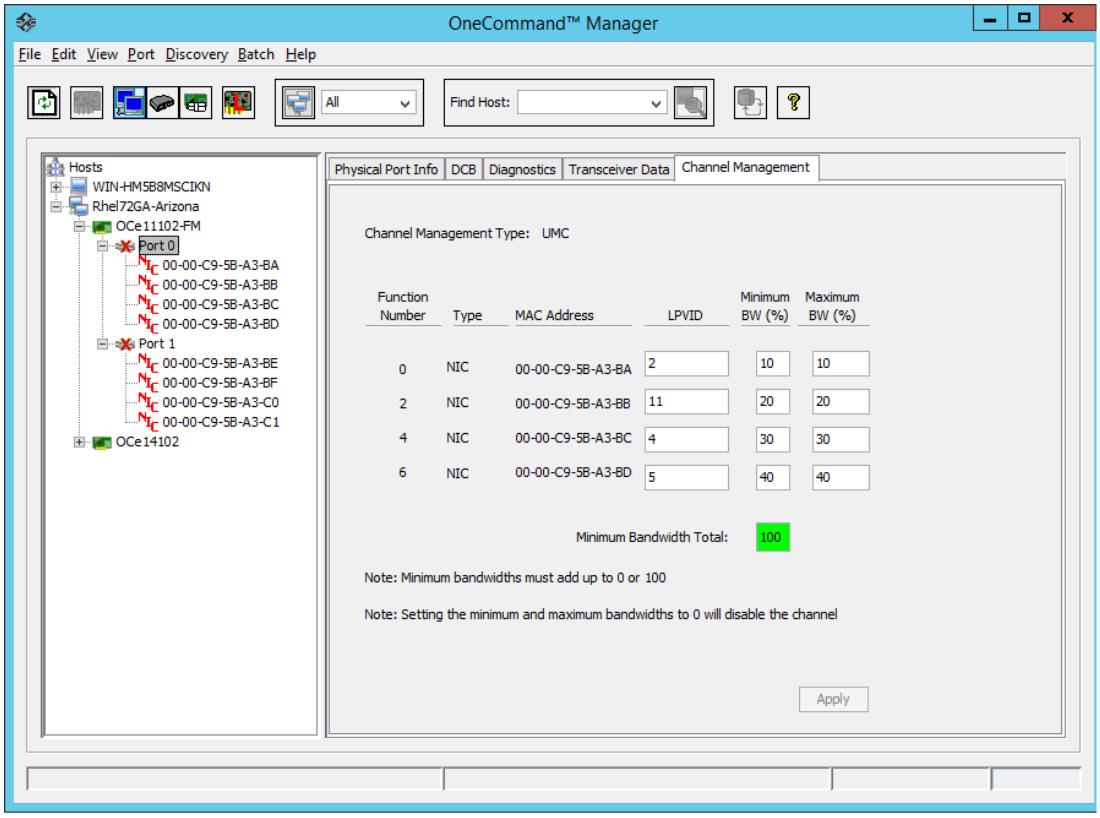
NOTE If Channel Management was enabled on the **Adapter Information** tab without rebooting, the **Channel Management** tab is not displayed.

To view general channel management information:

1. From the discovery-tree (Figure 5), select the adapter port whose channel information you want to view.
2. Select the **Channel Management** tab (Figure 19). 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 19 Example of the Channel Management Tab (UMC Management Depicted)



7.1.1.1 Configuring UMC Channel Management

NOTE Lenovo System X refers to UMC as SIMode.

Using the **Channel Management** tab (Figure 19), 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 Multichannel Reference Manual* for additional information on UMC.

NOTE 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), and reboot and run UMC without further configuration.

For Lenovo System X adapters, UMC mode may be referred to as vNIC2. The OneCommand CNA Manager application displays the Channel Management Type SIMode.

SR-IOV is not supported with UMC.

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 19).

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.

The following UMC **Channel Management** tab fields are displayed:

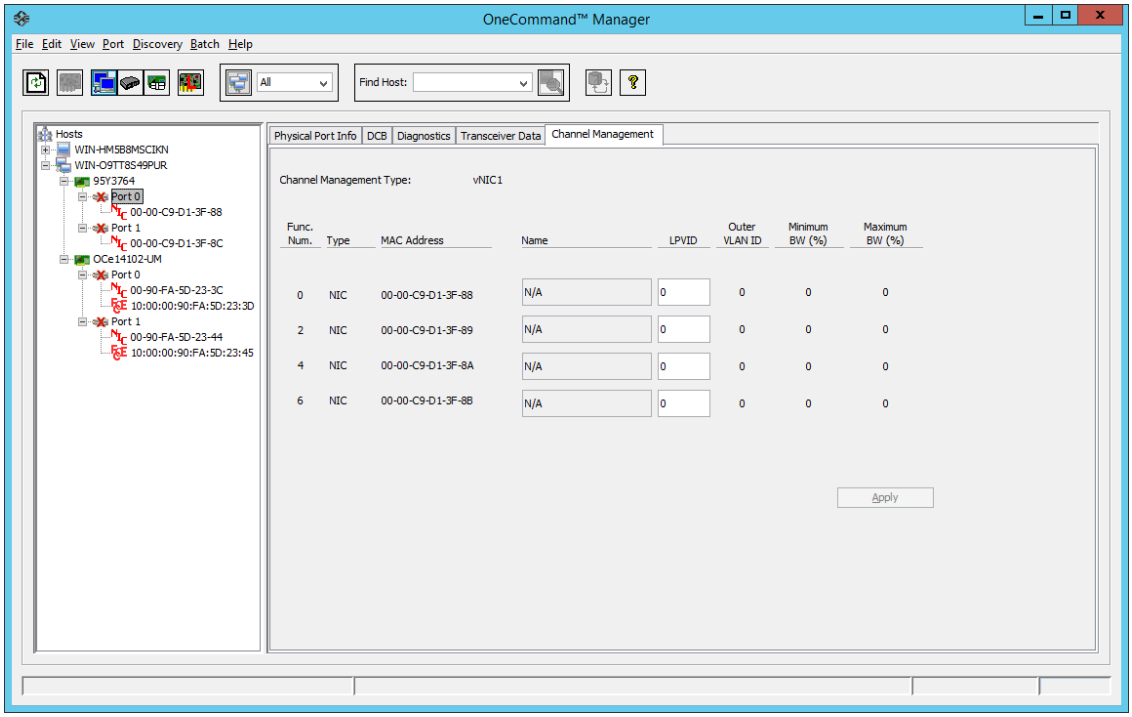
- Function Number (Read-only) – The function number assigned by the system to the channel. Channels 2 and 4 are read-only on OCe11102 adapters.
- Type (Read-only) – The protocol type in use.
- MAC address (Read-only) – The MAC address assigned to the channel.
- LPVID – Each channel can be assigned a 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.

7.1.1.2 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 20) 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 http://www.broadcom.com.

Figure 20 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 20).

The following **Channel Management** tab fields for vNIC1 (Lenovo System X only) are displayed:

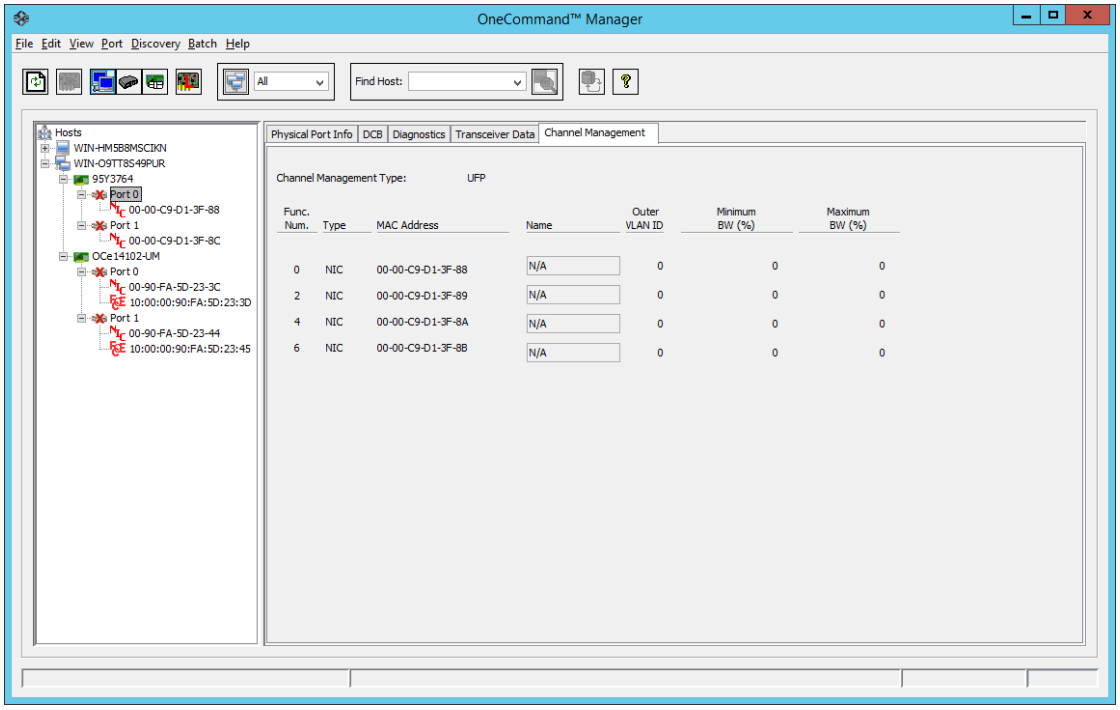
- Type (Read-only) – The protocol type in use.
- Func. Num. (Read-only) – The function number assigned by the system to the channel.
- Type (Read-only) – The protocol type in use by the channel.
- MAC address (Read-only) – The MAC address assigned to the channel.
- Name – The name assigned to the vNIC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

7.1.1.3 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 21) 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 http://www.broadcom.com.

Figure 21 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 21).

The following **Channel Management** tab fields for UFP (Lenovo System X only) are displayed:

- 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.

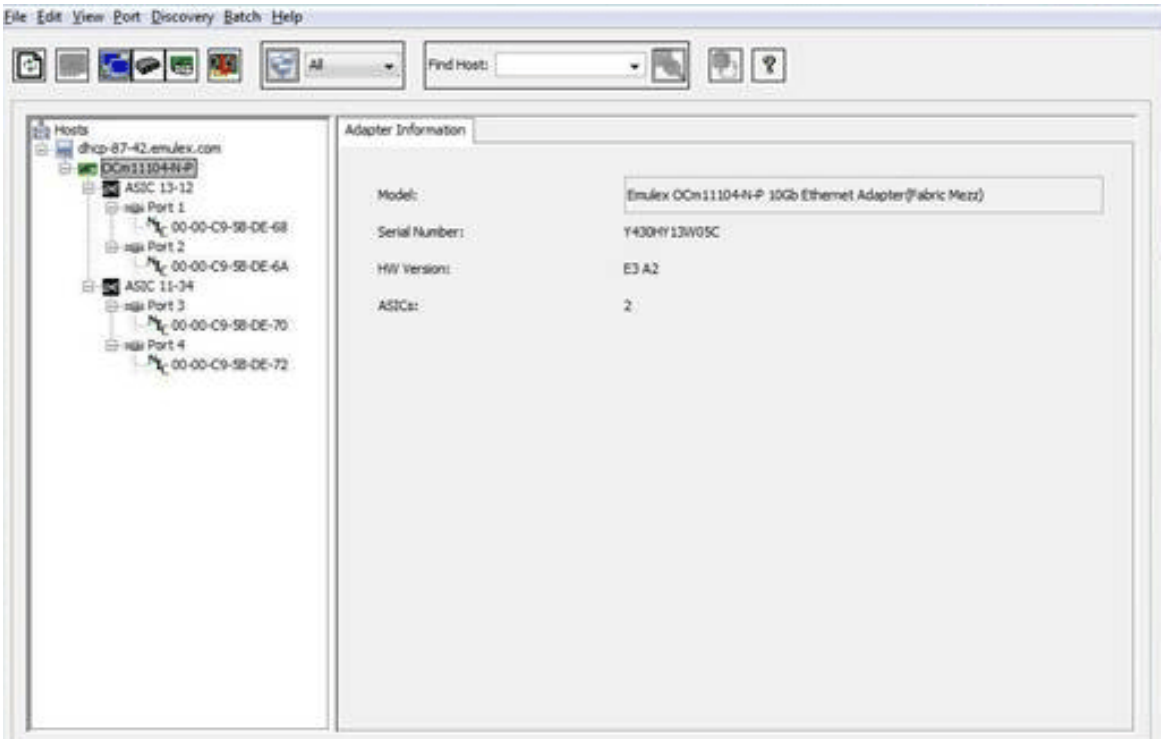
7.1.2 Viewing ASIC Information

When you select a OneConnect multi-ASIC adapter from the discovery-tree, the **ASIC Information** tab contains general attributes associated with the selected dual-ASIC, four-port adapter.

To view general multi-ASIC adapter information:

1. Select **Host** view.
2. Select a multi-ASIC adapter in the discovery-tree. The **Adapter Information** tab is displayed (Figure 22).

Figure 22 Multi-ASIC Adapter Information Tab



The following multi-ASIC **Adapter Information** tab fields are displayed:

- 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.

7.1.3 Viewing Multi-ASIC Information

When you select one ASIC on a multi-ASIC adapter from the discovery-tree, the **ASIC Information** tab appears. It contains general attributes associated with the selected ASIC.

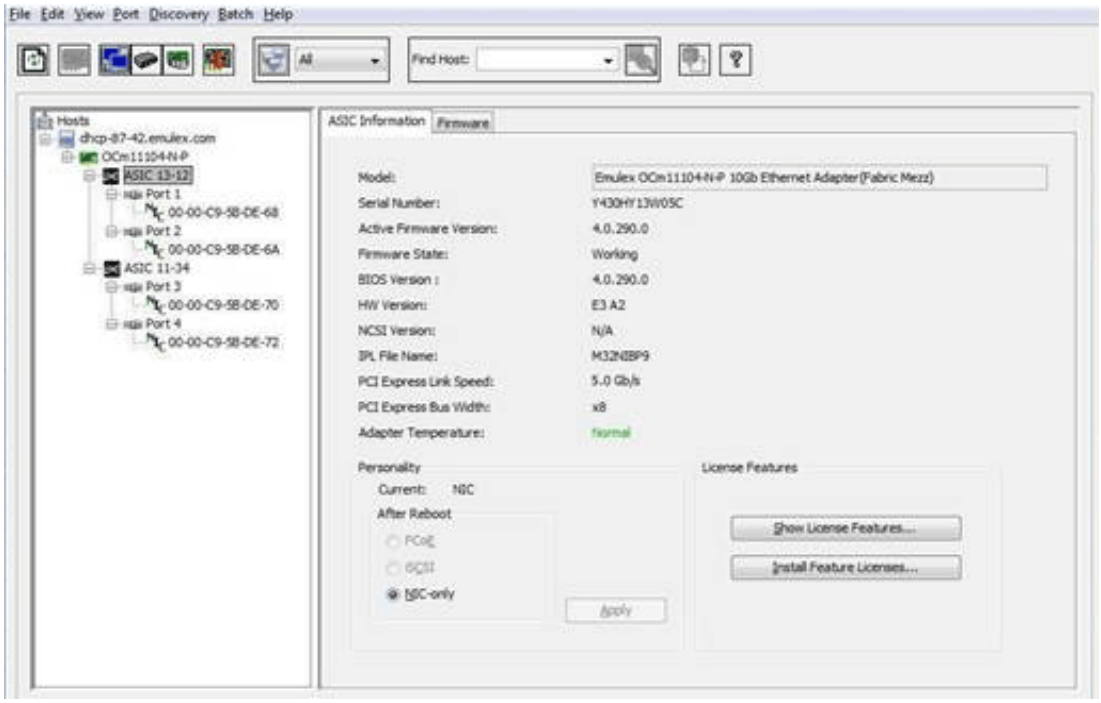
You can also use this tab to view and enable licenses. See Section 7.1.4, Viewing and Installing Licenses for OCe11000-Series Adapters, for more information.

To view general ASIC information:

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

   The **ASIC Information** tab is displayed (Figure 22).

Figure 23 ASIC Information Tab



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

- 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.
  - 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, inoperative 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, inoperative fans, and air conditioning issues that cause high ambient air temperatures.

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

- Personality area (OCe11000-series 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.

The system must be rebooted for your selection to take effect.

- License Features area (OCe11000-series adapters only):
  - Show License Features button – Click to show available licenses. See Section 7.1.4, Viewing and Installing Licenses for OCe11000-Series Adapters for more information.
  - Install License Features button – Click to install licenses.

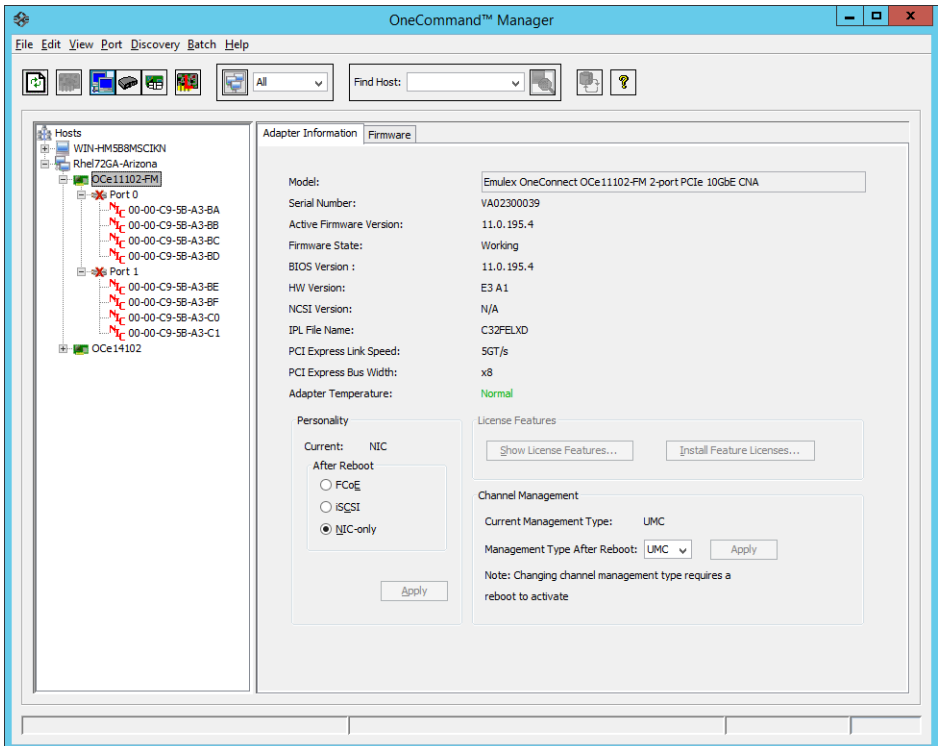
7.1.4 Viewing and Installing Licenses for OCe11000-Series Adapters

The OneCommand CNA Manager application allows you to view available licenses and install licenses to enable capabilities such as FCoE or iSCSI personalities on OCe11000-series adapters without having to re-wire the adapter.

Using the **Adapter Information** tab (Figure 24), you can view the licenses that 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 24 OneConnect OCe11000-Series Adapter Information Tab



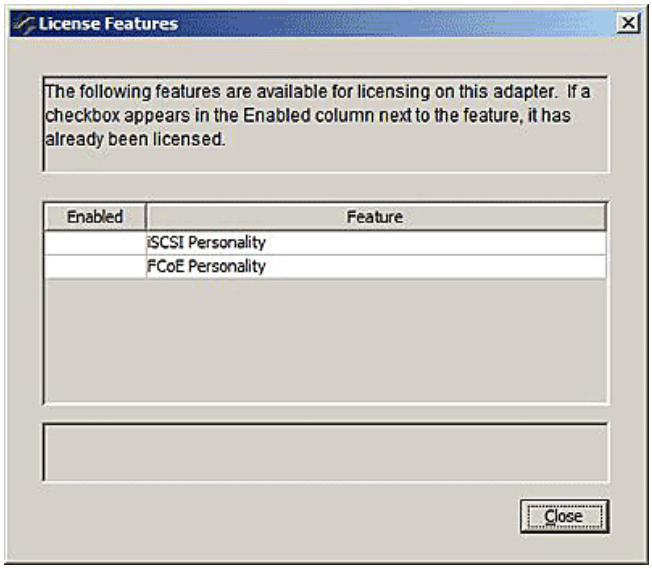
7.1.4.1 Viewing Licenses

To view the available licenses for OCe11000-series adapters:

1. From the discovery-tree, select the OCe11000-series adapter whose licenses you want to view. The **Adapter Information** tab is displayed (Figure 24).
2. On the **Adapter Information** tab, click **Show License Features**. The License Features window appears (Figure 25). 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 25 Licensed Features Window

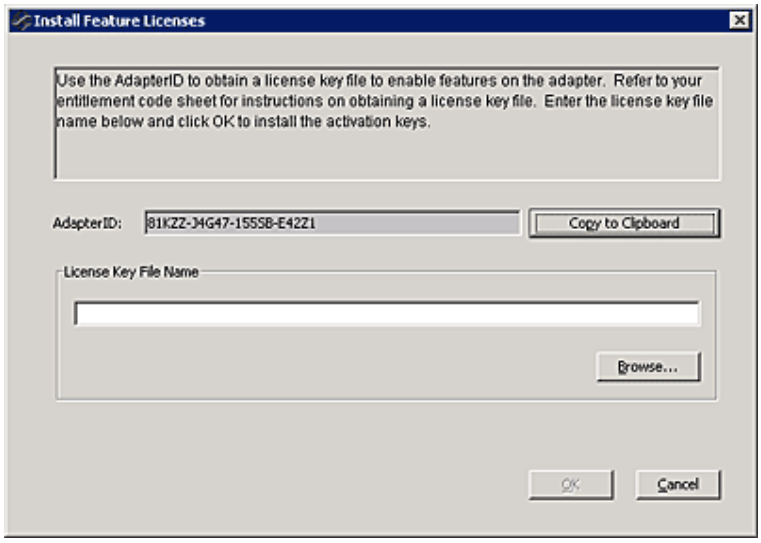


7.1.4.2 Installing Licenses

To install licenses for OCe11000-series adapters:

1. From the discovery-tree (Figure 5), select the OCe11000-series adapter whose licenses you want to install. The **Adapter Information** tab is displayed (Figure 24).
2. From the **Adapter Information** tab, click **Install Feature Licenses**. The Install Feature Licenses dialog box appears displaying the AdapterID (Figure 26).

Figure 26 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 26), 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 prompts you to confirm that you want to install the licenses. Click **OK**.
7. A dialog notifies you that the installation was successful or explains why it failed. Click **OK**.

7.1.5 Changing Personalities on OneConnect OCe11000-Series Adapters

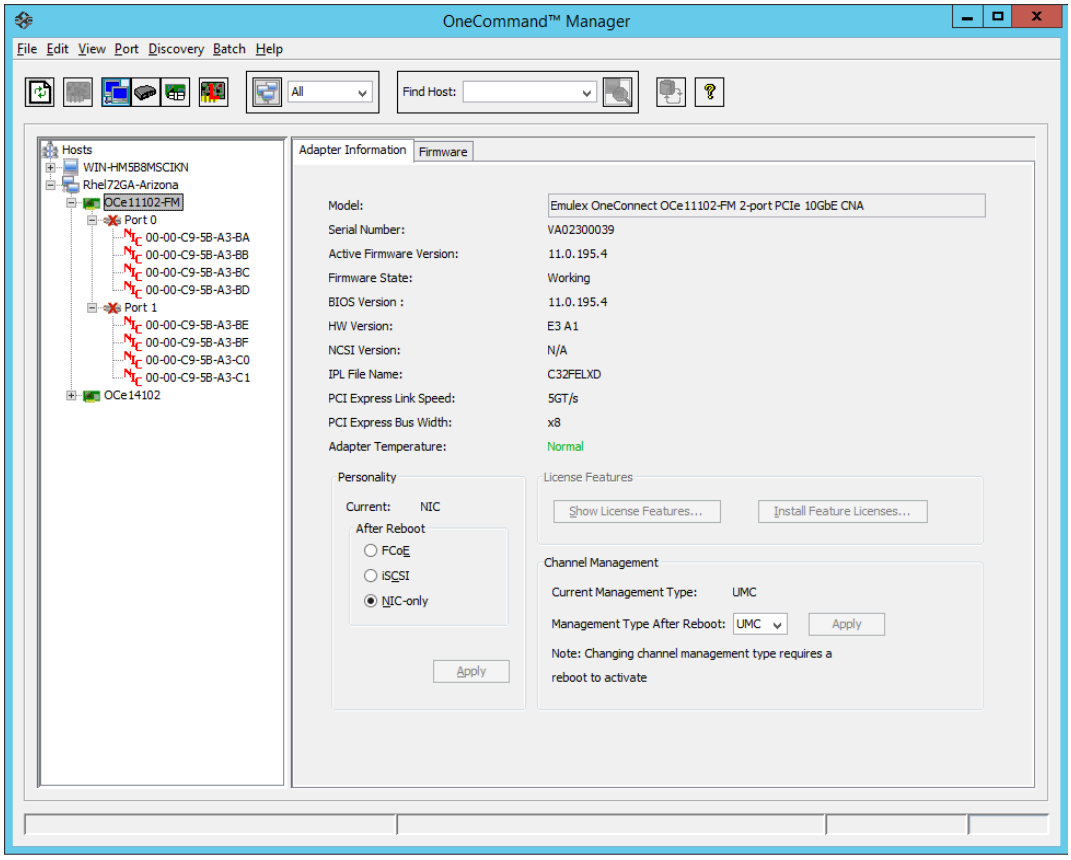
The OneCommand CNA Manager application enables you to change the personality or protocols running on OneConnect OCe11000-series adapters.

When you change the personality of the adapter and reboot the host, the adapter starts running the new personality or protocol. The personalities that OCe11000-series adapters currently run are NIC-Only, NIC + FCoE, and NIC + iSCSI. In some cases, the adapters are preconfigured to support multiple personalities. In other cases, you must install a license key before the adapter can support multiple personalities. See Section 7.1.4, Viewing and Installing Licenses for OCe11000-Series Adapters 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 27) to make personality changes.

Figure 27 OneConnect OCe11102 Adapter Information Tab



To change the personality of OCe11000-series adapters:

1. From the discovery-tree, select the OCe11000-series adapter whose personality you want to change. The **Adapter Information** tab is displayed (Figure 27).
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.

7.1.6 Viewing Firmware Information

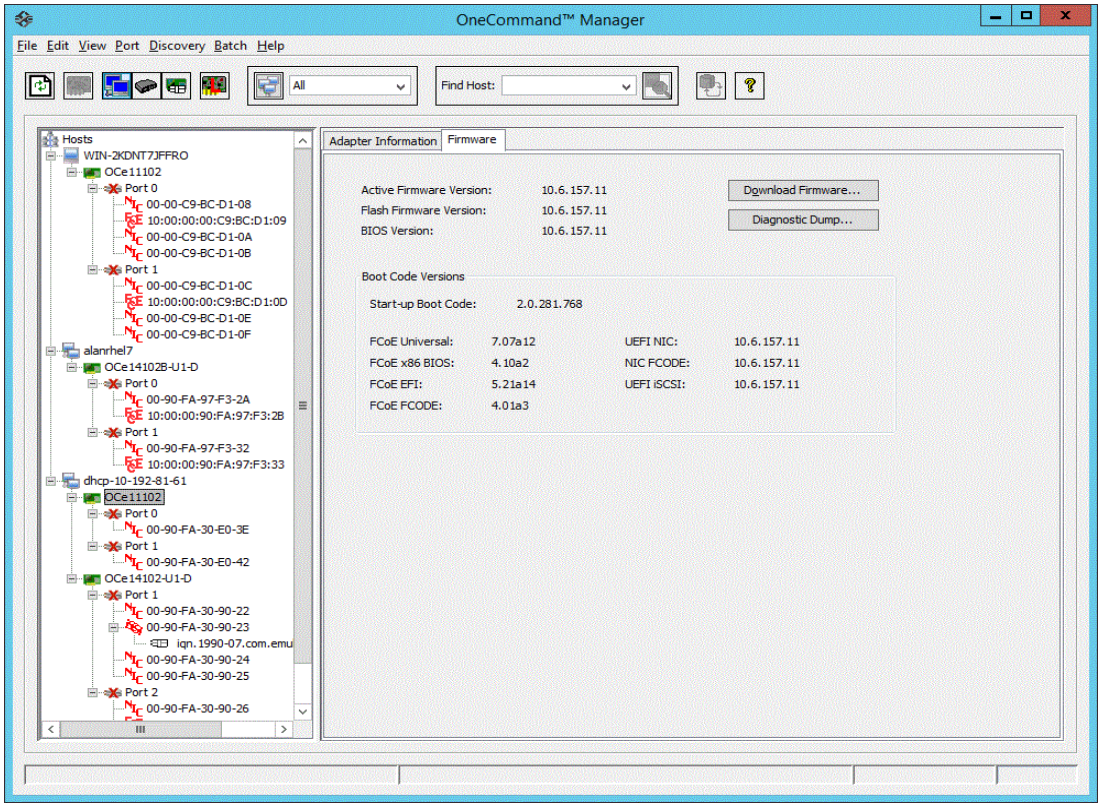
OneConnect adapter firmware is maintained on an adapter-specific basis. Use the Firmware tab (Figure 28) to download firmware and create diagnostic dumps for the selected adapter.

To view 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 5).
3. Select the **Firmware** tab (Figure 28).

Figure 28 OneConnect Firmware Tab



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

- 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.
  - 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.

To update firmware on the selected adapter, see Section 9, Updating Adapter Firmware.

To create a diagnostic dump for the selected adapter, see Section 12.5, Creating Diagnostic Dumps.

7.1.7 Viewing Physical Port Information

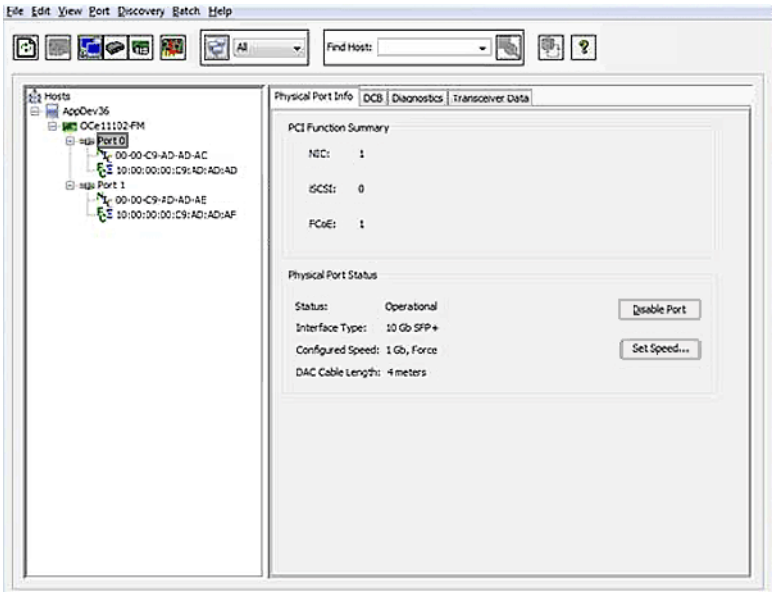
The **Physical Port Info** tab (Figure 29) contains a general summarization of the PCI functions under that physical port and the current physical port status.

OneConnect OCe11000-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 Section 7.1.7.2, Setting Port Speed and DAC Cable Length 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 29).

Figure 29 Physical Port Info Tab



7.1.7.1 Enabling and Disabling OneConnect Physical Ports

Using the **Physical Port Info** tab (Figure 29), 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:

NOTE Disabling a port does not remove the functions from the PCIe bus. It only disables the operation of the port on the adapter. For OCe14000-series adapters, the port can be removed from the host (see Section 7.2.8, Removing a Port for more information).

1. In the discovery-tree (Figure 5), select the physical port you want to enable or disable.
2. Select the **Physical Port Info** tab (Figure 29).
3. Click **Enable Port** or **Disable Port**.

7.1.7.2 Setting Port Speed and DAC Cable Length

The **Physical Port Info** tab enables you to set port speed and DAC cable lengths.

To set the port speed:

1. From the discovery-tree (Figure 5), select the adapter port whose speed you want to change.
2. Click **Set Speed** on the **Physical Port Info** tab (Figure 29). The Change Port Speed dialog box appears (Figure 30).

Figure 30 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 can 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 does 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.

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

The **PHY Data** Tab (Figure 31) displays port-level operational attributes, 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.

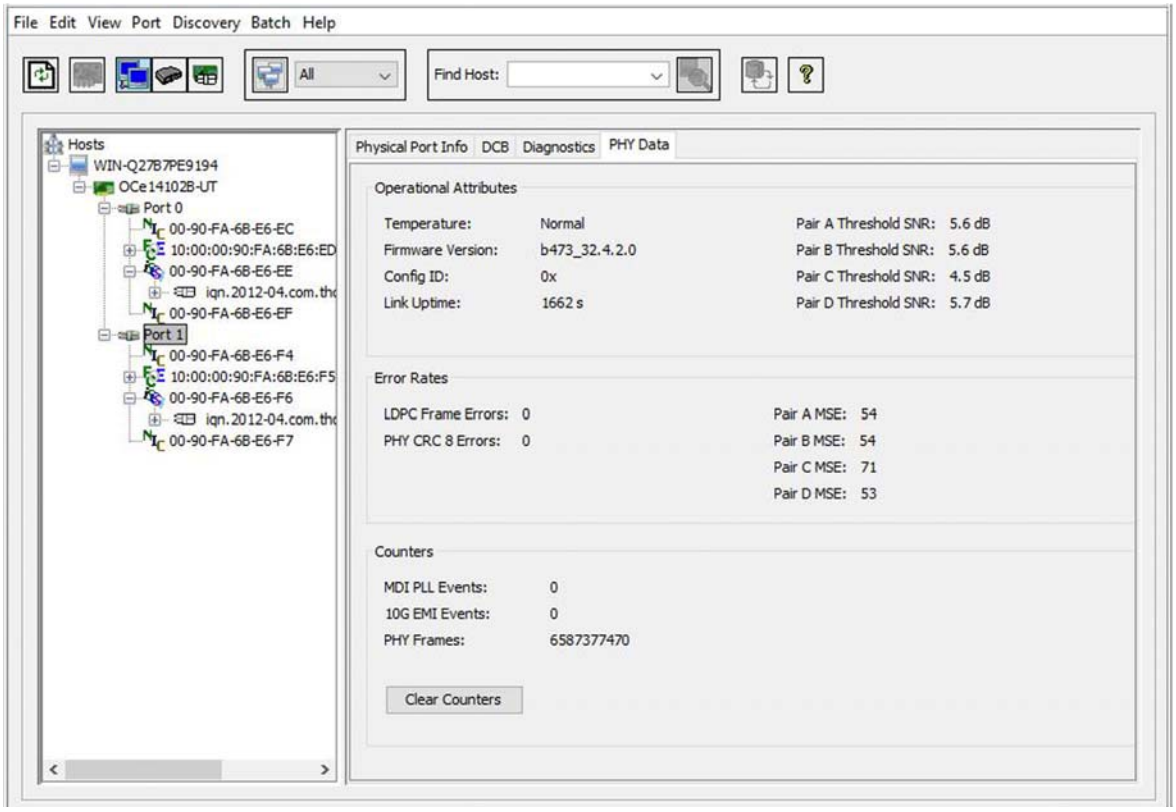
To view OneConnect 10GBASE-T adapter port PHY information:

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 31).

Figure 31 PHY Data Tab



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

- Operational Attributes area:
  - Temperature – The temperature of the selected port.
  - Firmware Version – The PHY firmware version.
  - Config ID – Vendor specific identifier for the PHY firmware version. Value 0 is unavailable for the PHY.
  - Link Uptime – Time in seconds since the link has been up.
- Error Rates area:
  - LDPC Frame Errors – The LDPC counter tracks the number of LDPC frames received by MDI that cannot be corrected (this value should be 0 or close to 0). This counter self-clears at MDI link down.
  - PHY CRC 8 Errors – Number of CRC 8 errors received by the MDI that cannot be corrected (this value should be 0 or close to 0). This counter self-clears at MDI link down.
- 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 to zero. The frame counter is unchanged for OCe11000-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).

7.1.9 Viewing Transceiver Information

When you select an adapter port from the discovery-tree (Figure 5), the **Transceiver Data** tab (Figure 32) 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.

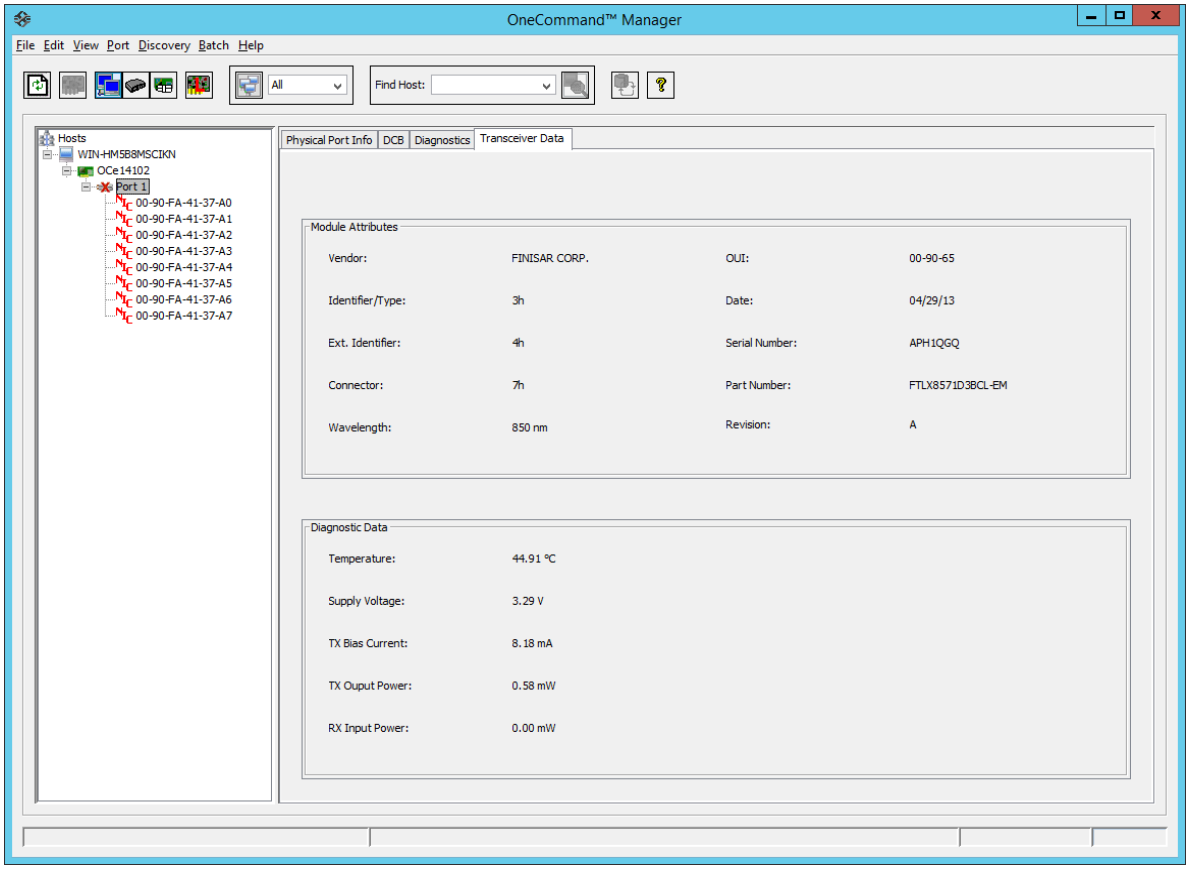
To view transceiver information:

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

   NOTE iSCSI and NIC-Only adapters do not appear in Fabric view.

2. In the discovery-tree (Figure 5), select the OneConnect adapter port whose transceiver information you want to view.
3. Select the **Transceiver Data** tab (Figure 32).

Figure 32 Transceiver Data Tab



7.1.9.1 The following Transceiver Data tab fields are displayed:

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

7.2 OCe14000-Series Adapter Configuration

When you select an OCe14000-series adapter from the discovery-tree (Figure 5), the **Adapter Configuration** tab (Figure 33) 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.
- Multichannel – View the channel properties and change the configurable properties (based on the multichannel type).

The **Adapter Configuration** tab always shows the permanent MAC address for Multichannel and Custom configurations under the Next Boot Configurations panel. To correlate the permanent MAC address with the current MAC address, view the **Port Information** tab. The **Port Information** tab always shows the current (user-defined) MAC address and the permanent MAC address.

For WWPN and WWNN:

- The current FCoE WWPN is displayed in the discovery-tree. If a virtual WWPN is set, it is displayed here; otherwise, the permanent WWPN is displayed.
- Clicking the FCoE function in the discovery-tree displays the **Port Information** tab, where the current WWPN and the current WWNN are displayed.
- If a virtual configuration has not been set, the permanent or factory default setting is shown.

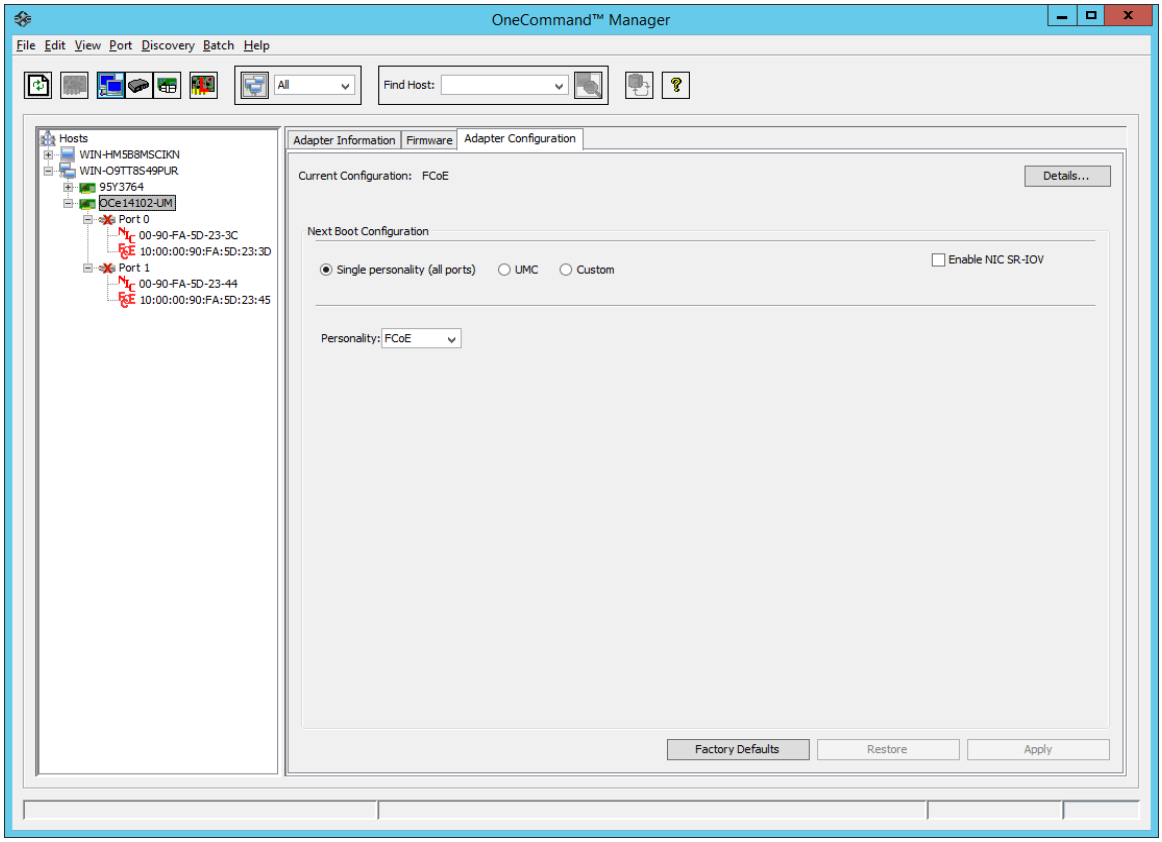
Considerations

- 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 Section 7.2.6, Dell NPar Configuration View for more information.

NOTE RoCE is available as Technical Preview only. Do not use RoCE in a production environment.

- RoCE configurations are not supported with SR-IOV.
- Emulex adapters running iSCSI require the Open-iSCSI driver. Make sure that the appropriate iSCSI Initiator Package for each Linux distribution is installed.
- For RHEL 6.4, RHEL 7, and SLES 11 SP2 and later Linux distributions, Broadcom has replaced the proprietary out-of-box driver with the Open-iSCSI driver (refer to the latest *Emulex Drivers for Linux User Guide* for more information).

Figure 33 OCe14000-series Adapter Configuration Tab (FCoE Selected)



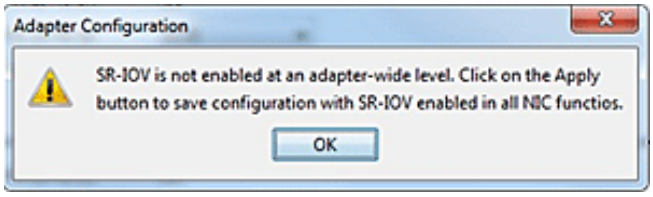
The following OCe14000-series **Adapter Configuration** tab fields are displayed:

- 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-multichannel and non-NIC+RoCE configurations.
- 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 34, prompting you to click the **Apply** button to save the configuration with SR-IOV enabled on all NIC functions.

Figure 34 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), multichannel, or a custom configuration.

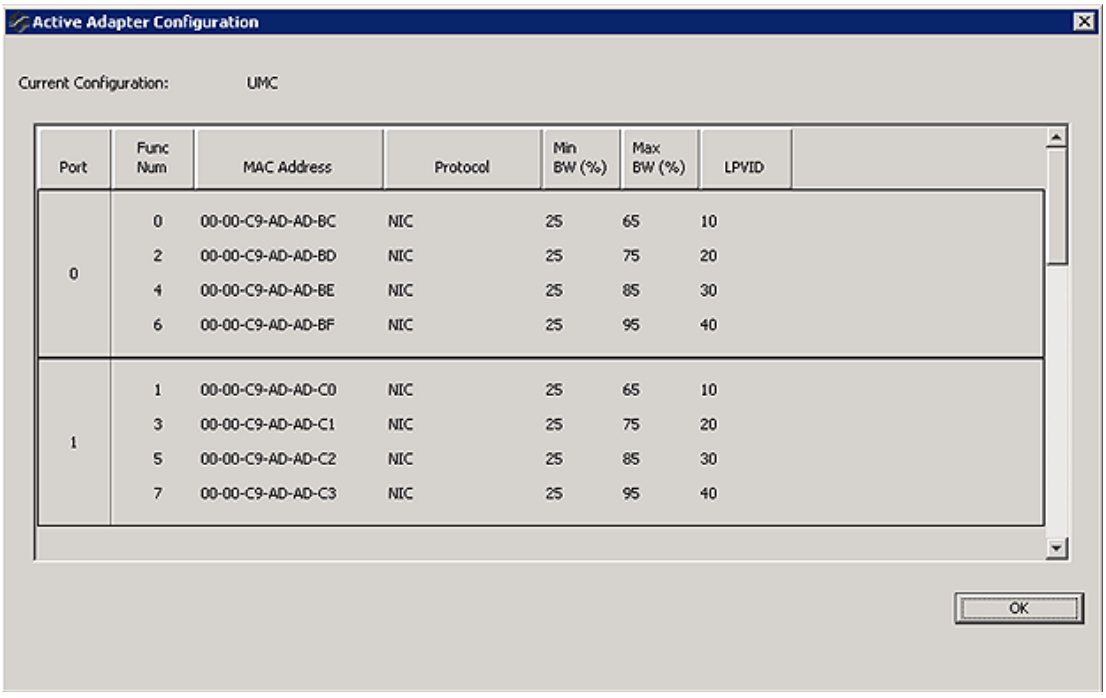
To revert any unsaved changes back to the currently saved settings on the adapter, click **Restore**.

To return the adapter's ports to the factory default profile and settings for that adapter, click **Factory Defaults**. A reboot is required.

NOTE For Dell only: If virtual MAC addresses are used for an adapter, clicking **Factory Defaults** results in resetting all the configured virtual MAC addresses to the factory default MAC addresses.

To display a pop-up window with more information about the current configuration running, click **Details**. The Active Adapter Configuration window (Figure 35) is displayed.

Figure 35 Active Adapter Configuration Example

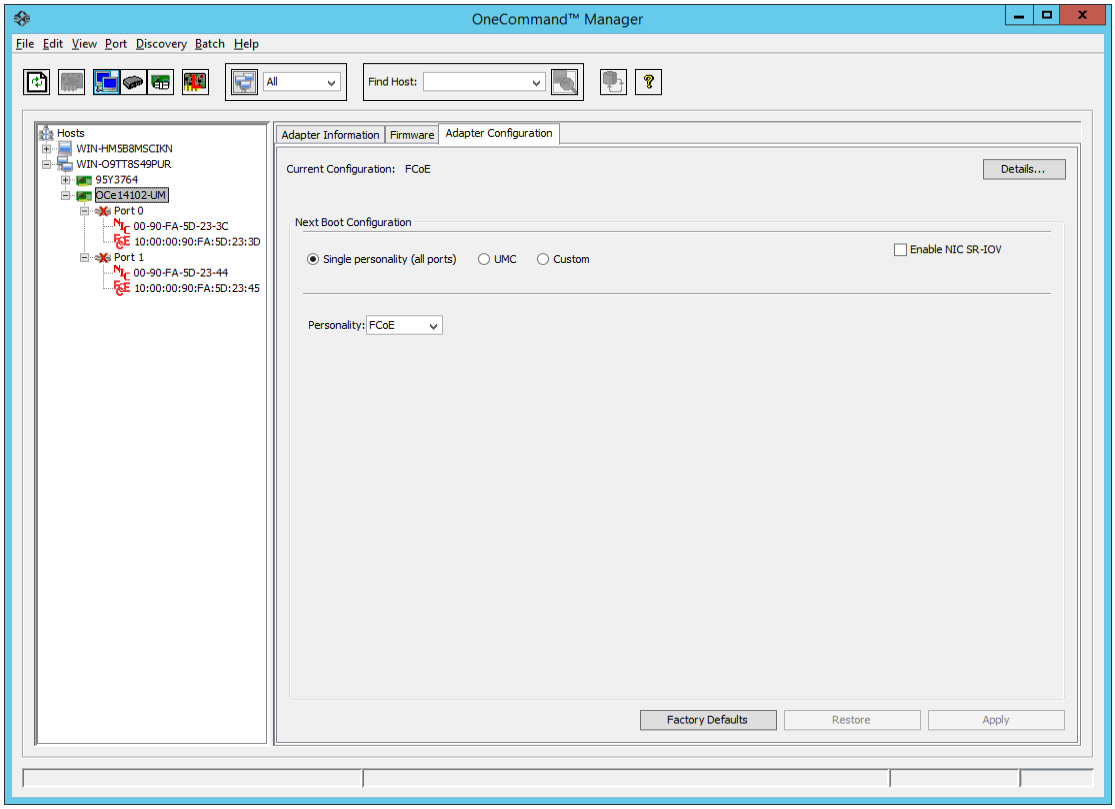


7.2.1 Configuring Single Personalities

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 36.

Figure 36 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 5), select the OCe14000-series adapter whose personality you want to configure.
2. From the **Adapter Configuration** tab (Figure 33), 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 Knowledge Base on http://www.broadcom.com 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.

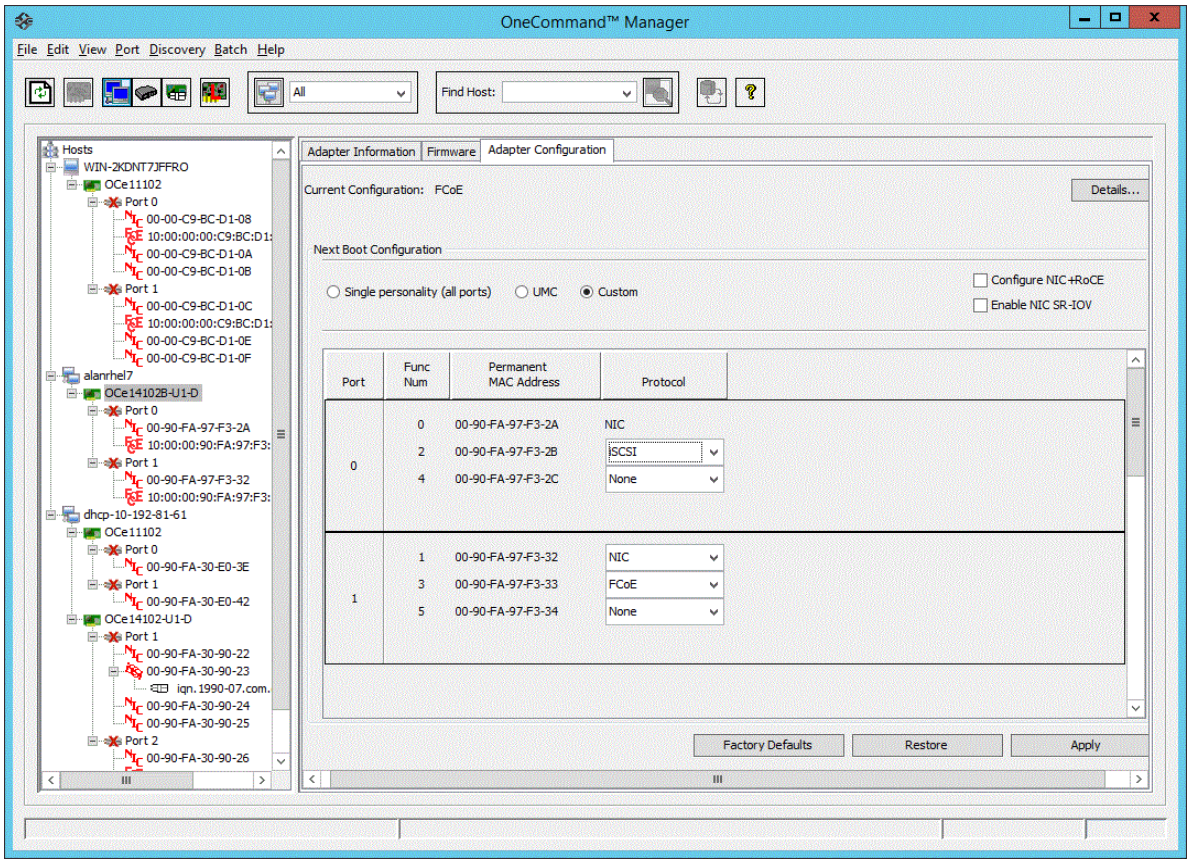
7.2.2 Custom Configurations

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 37.

NOTES

- If you configure all ports the same, you essentially configured a single personality. In this case, the tab automatically reverts to single personality mode.
- NIC+RoCE and SR-IOV are mutually exclusive. If you select the Configure NIC+RoCE checkbox, you cannot select the Enable NIC SR-IOV checkbox because it is disabled. Similarly, if you select the Enable NIC SR-IOV checkbox, you cannot select the Configure NIC+RoCE checkbox because it is disabled.

Figure 37 Custom View



NOTE

Enabling and configuring more than four functions per physical port requires that Custom mode is enabled on the adapter. The options to enable and configure more than four functions per physical port are not available unless Custom mode is selected.

If you enable Custom mode but do not create a Custom configuration, the personality might revert to NIC, iSCSI, or FCoE. If Custom mode is enabled, the following conditions apply:

- If you leave all functions set to NIC, the configured personality is automatically switched from Custom to NIC.
- If you leave the functions configured the same as they would be for iSCSI, the configured personality is automatically switched from Custom to iSCSI.
- If you leave the functions configured the same as they would be for FCoE, the configured personality is automatically switched from Custom to FCoE.

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 33), 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.

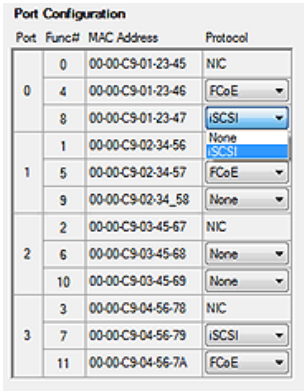
   Or, check the NIC+RoCE checkbox to configure NIC+RoCE or only NIC on each port.

   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.

In custom configuration view, up to three functions per port are configurable as depicted in Figure 37.

When configuring storage, the first function must be set to NIC. There are pull-down menus for selecting a storage protocol to run on the second and third functions. If two storage functions are configured on a port, they must be different storage protocols. 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 38).

Figure 38 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 have 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.

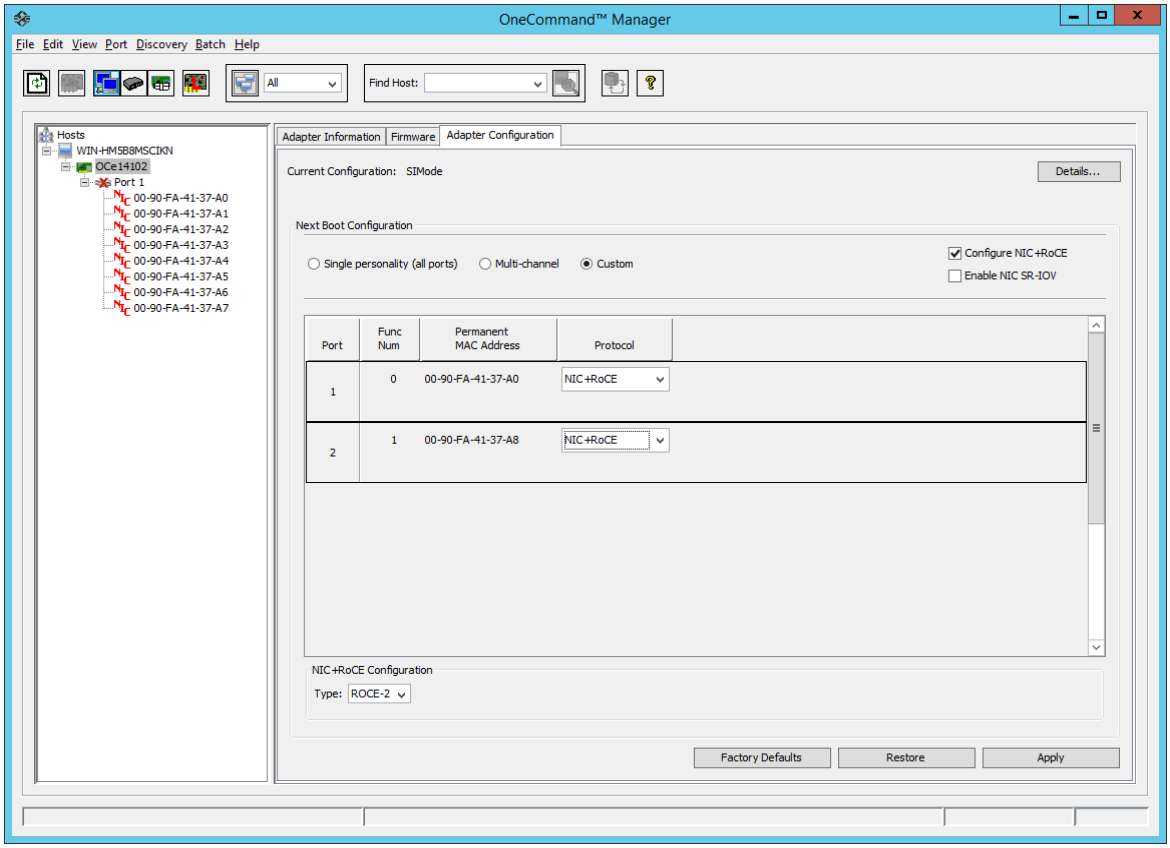
7.2.2.1 Configuring RoCE in a Custom View

In Custom view, NIC+RoCE can be configured on the first function of a port (see Figure 39). 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 Knowledge Base on http://www.broadcom.com for any updated information on additional use.

NOTE RoCE configurations are not supported with SR-IOV.

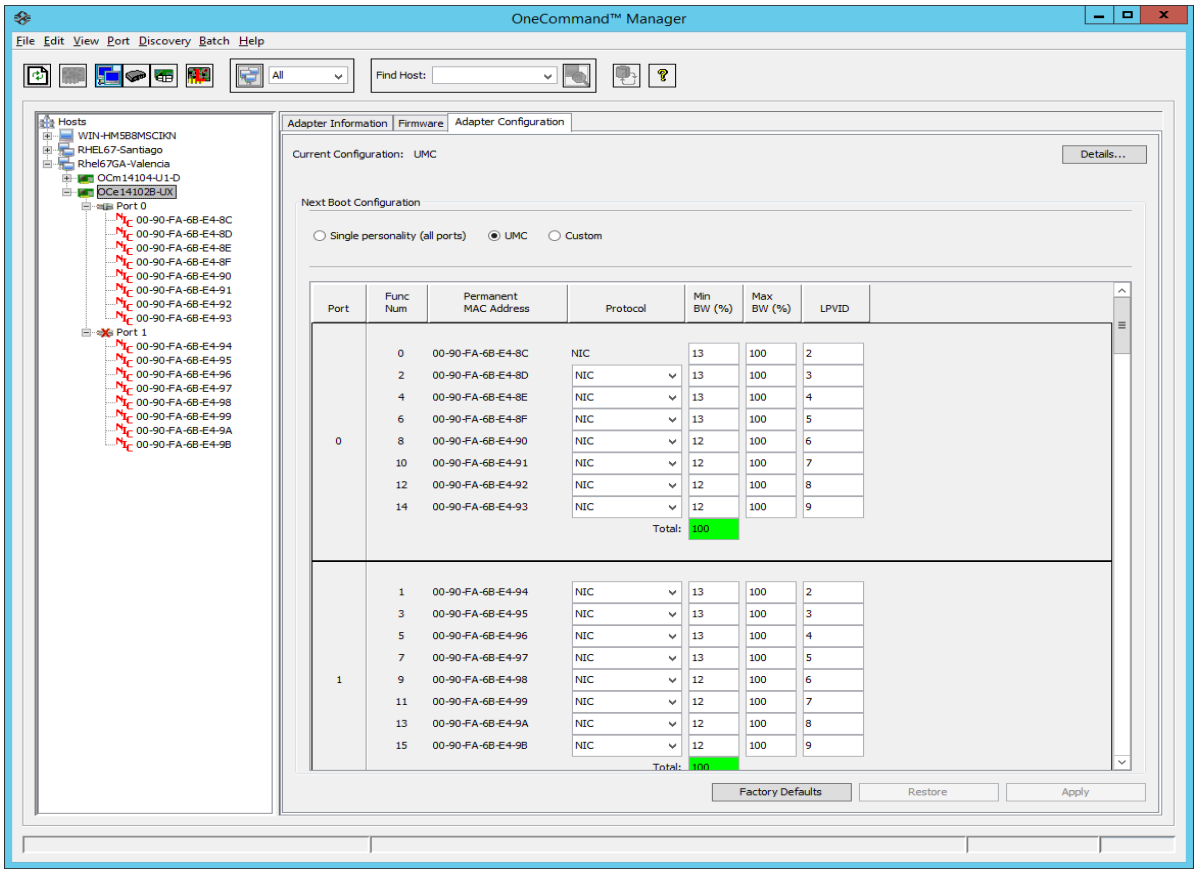
Figure 39 Custom NIC+RoCE Configuration



7.2.3 UMC Configuration View

When UMC is available and the UMC radio button is selected, the **Adapter Configuration** tab looks like Figure 40.

Figure 40 UMC View Adapter Configuration Tab (UMC, 2 Ports, 8 Functions per Port)



For UMC, the OneCommand CNA 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 are present and discovered by the OneCommand CNA 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
- 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 to 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.

7.2.3.1 Default UMC Settings

The default UMC settings prevent the enabling of functions that require ARI, in the event that the host does not support ARI. This prevents you from enabling channels that are not 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 is still 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:

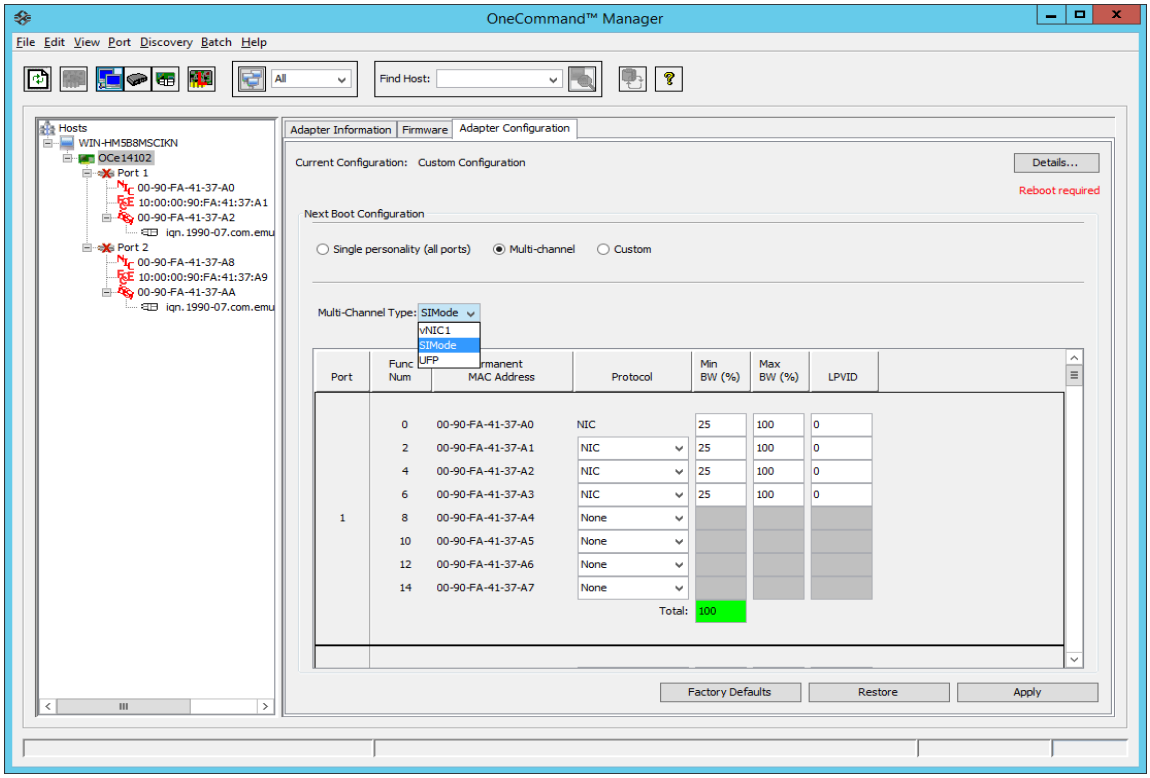
1. All enabled functions (that is, functions not set to None) continue to be enabled running the same protocol before UMC was enabled even if its function number is greater than 7 (only possible on four-port adapters).
2. Any channel with a function number less than 8 is enabled with NIC.
3. Any channel with function number greater than 7 is set to None (that is, channel disabled).
4. The minimum bandwidths are assigned by dividing up 100% as evenly as possible between the enabled channels.
5. The maximum bandwidths for the enabled channels are set to 100.
6. Both the minimum and maximum bandwidth for the channels assigned to None are set to 0.
7. The LPVIDs are 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.

7.2.4 Lenovo System X Multichannel Configuration View

For Lenovo multichannel configuration, there are three different multichannel types available: vNIC, SIMode, and UFP. When you check the Multi-channel radio button, the **Adapter Configuration** tab looks like Figure 41.

Figure 41 Multichannel View (Showing Multi-Channel Type Drop-down)



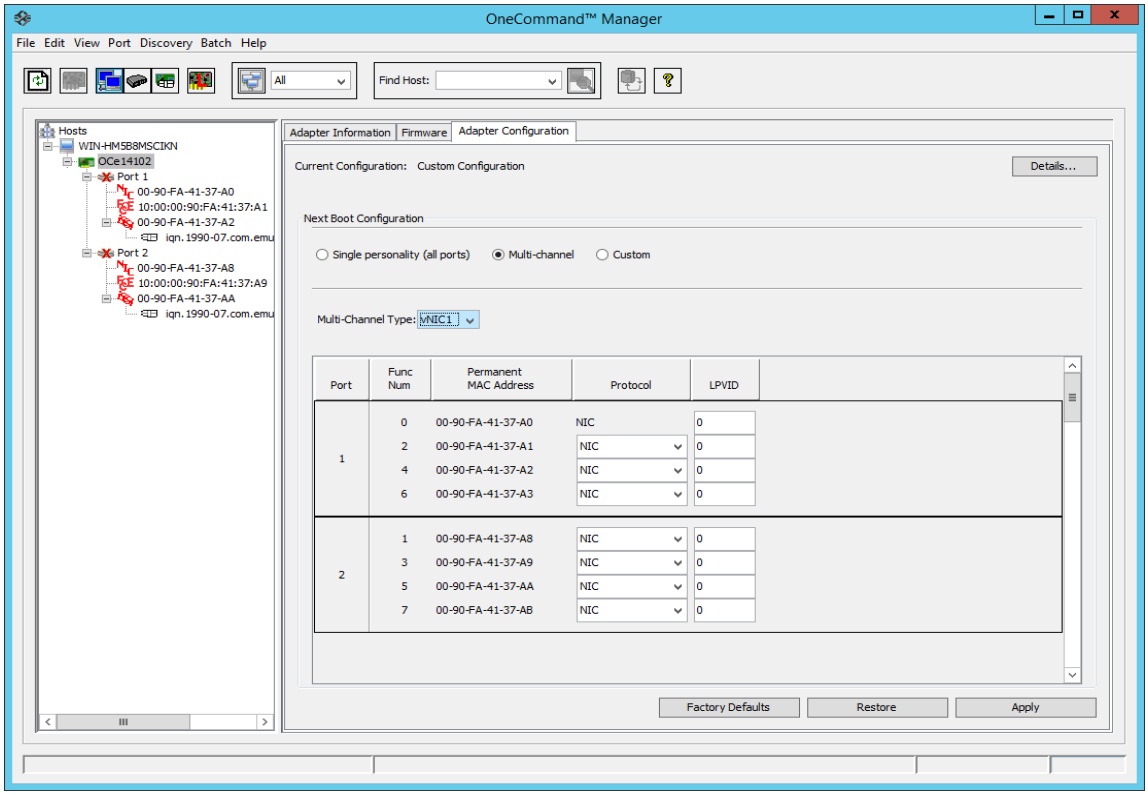
A Multi-Channel Type drop-down menu is available to select the type of multichannel configuration on the adapter. The contents of the Channel Configuration table depend upon the selected multichannel type.

NOTE If Lenovo System X multichannel is enabled, RoCE cannot be configured on any function.

7.2.4.1 vNIC Configuration

When vNIC is selected on the adapter, the **Adapter Configuration** tab looks like Figure 42.

Figure 42 Multichannel vNIC View



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 to 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.

7.2.4.2 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
- 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 to 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 5.

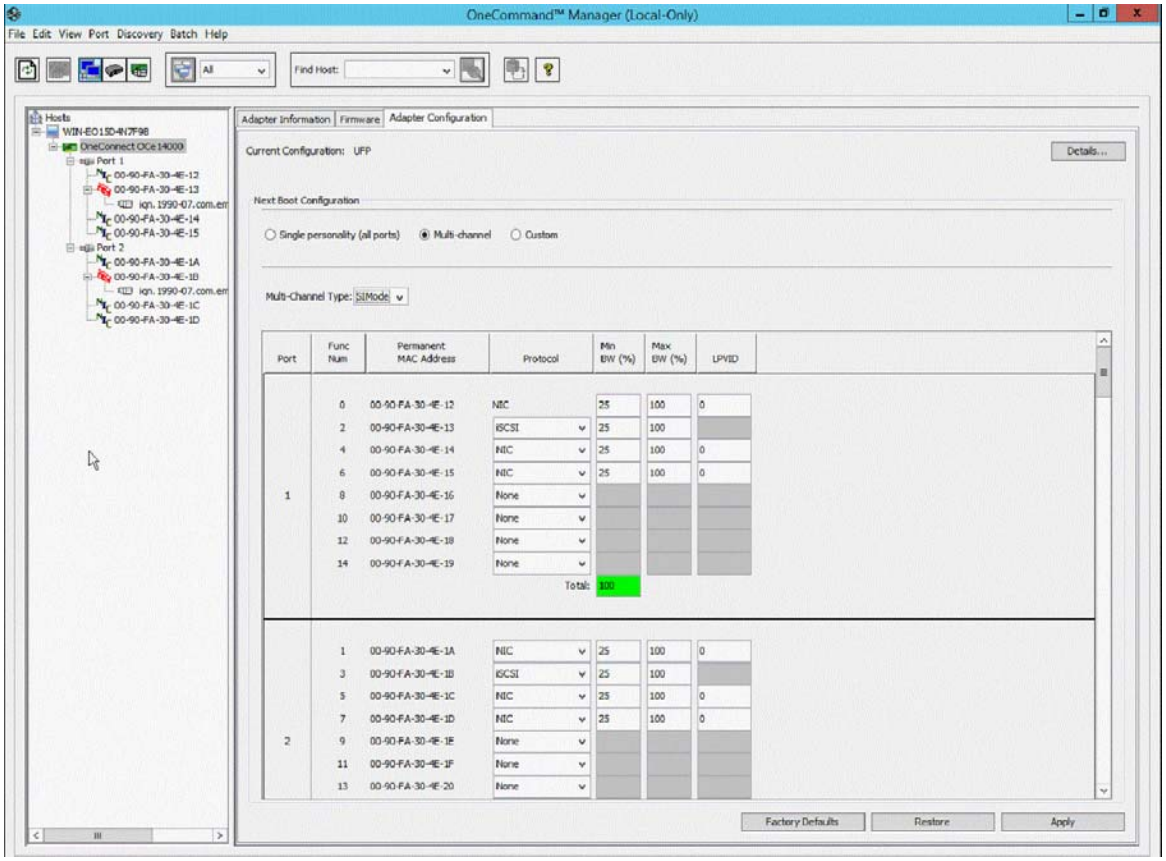
Table 5 vNIC Maximum Number of Functions per Port

Number of Ports	Port Speed	Default Functions per Port	Maximum Functions per Port
1	10Gb/s	4	4
2	10Gb/s	4	4
4	10Gb/s	4	4

7.2.4.3 SIMode Configuration

When SIMode is selected, the **Adapter Configuration** tab looks like Figure 43.

Figure 43 Multichannel SIMode View

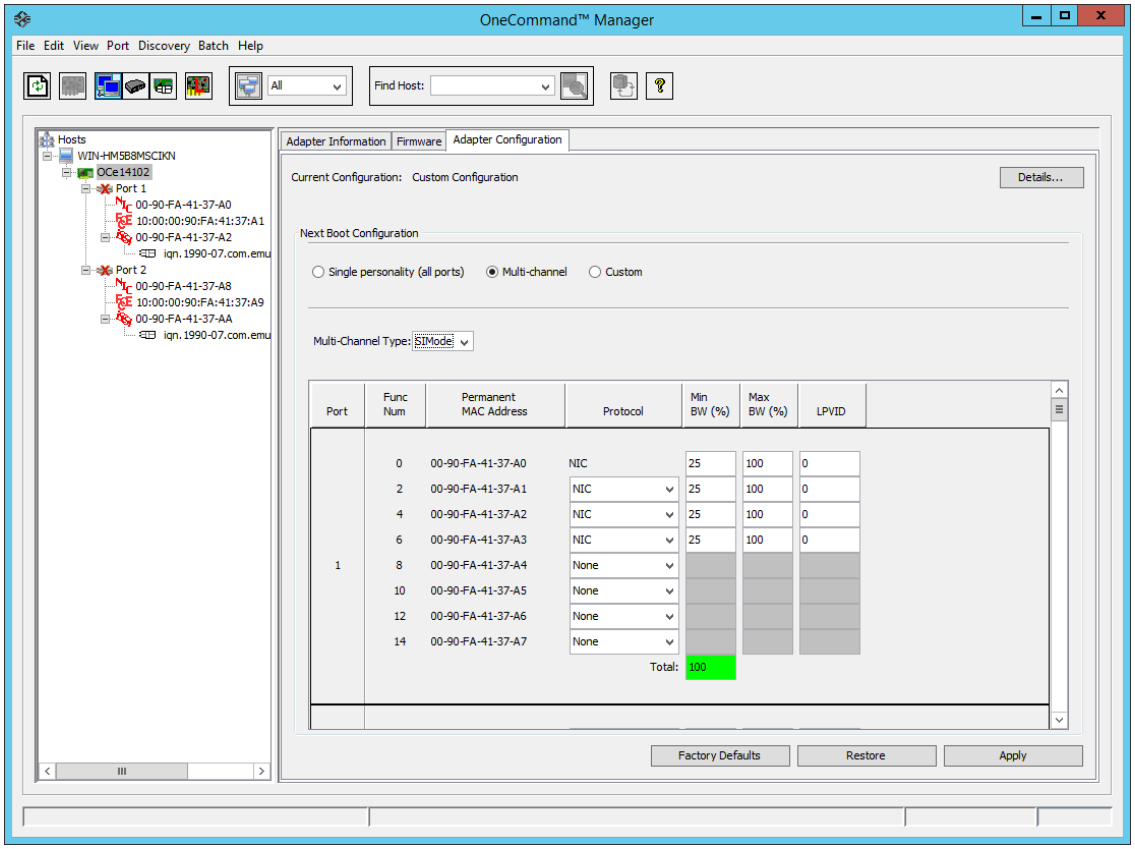


SIMode and UMC configuration are exactly the same. See Section 7.2.3, UMC Configuration View for more information.

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

When switching from a non-multichannel 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 44 shows this configuration.

Figure 44 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 6.

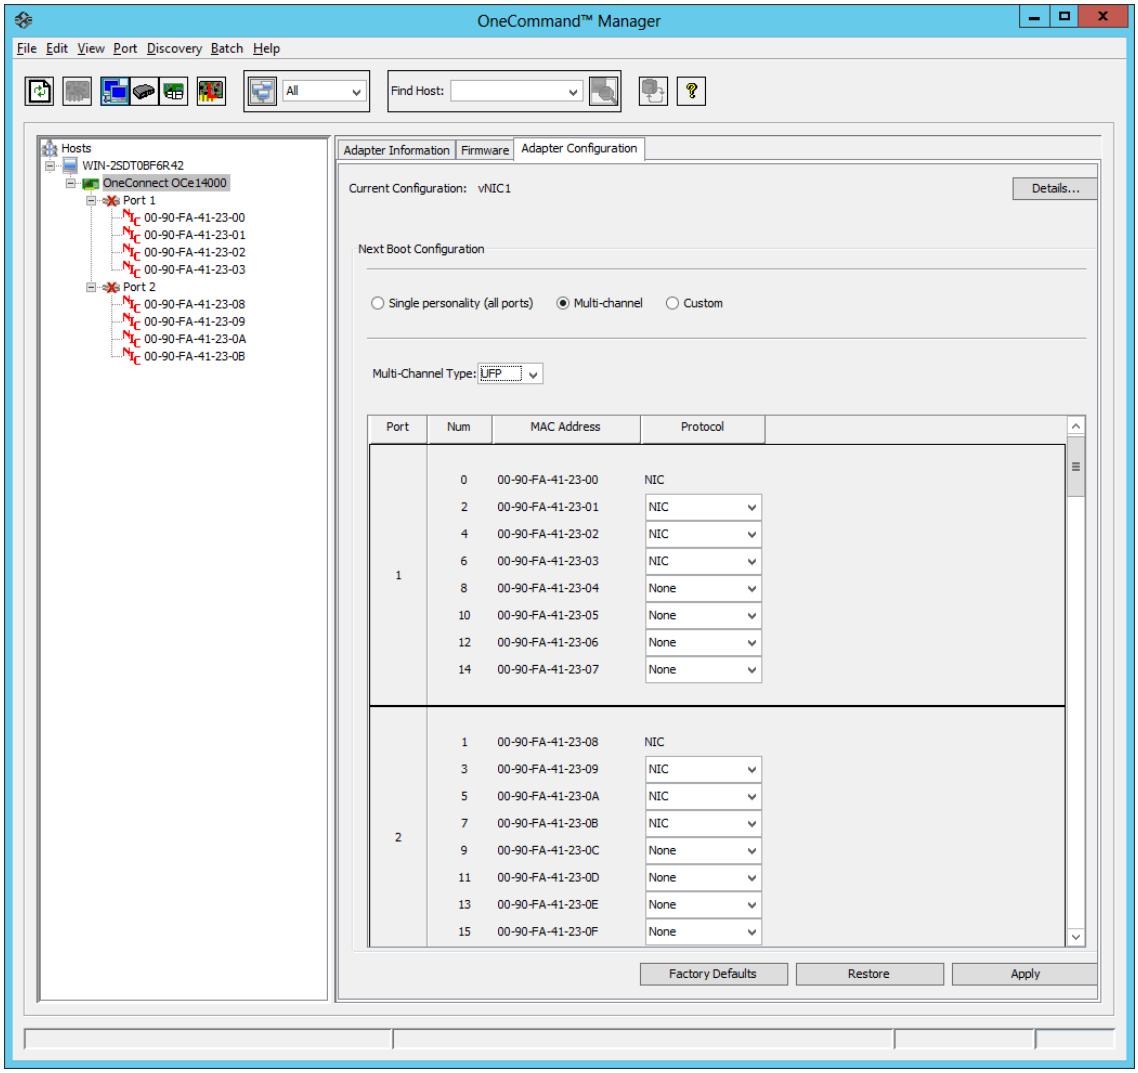
Table 6 SIMode Maximum Number of Functions per Port

Number of Ports	Port Speed	Default Functions per Port	Maximum Functions per Port
1	10Gb/s	4	8
2	10Gb/s	4	8
4	10Gb/s	4	4
1	40Gb/s	4	16

7.2.4.4 UFP Configuration

If UFP is selected, the **Adapter Configuration** tab looks like Figure 45.

Figure 45 Multichannel UFP View



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 four functions per port on one port and six functions per port on another. If you configure six 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 four functions per port. The maximum number of functions per port is shown in Table 7.

Table 7 UFP Maximum Number of Functions per Port

Number of Ports	Port Speed	Default Functions per Port	Maximum Functions per Port
1	10Gb/s	4	8
2	10Gb/s	4	8
4	10Gb/s	4	4
1	40Gb/s	4	16

7.2.4.5 UFP Configuration Rules

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

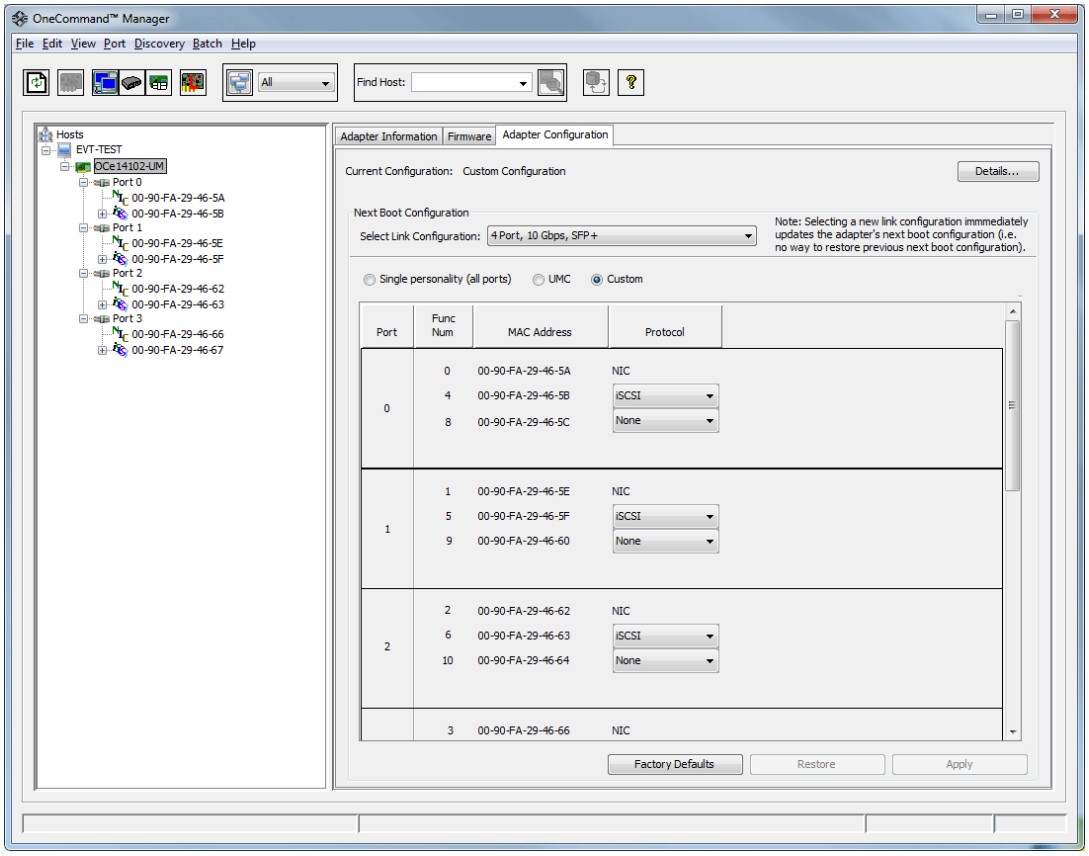
- The first channel is always NIC.
- NIC+RoCE cannot be configured with UFP.
- The second and third channels can be assigned the following:
  - NIC
  - None (channel disabled)
  - Storage protocol
- The fourth and higher channels of a port (if available) can be assigned NIC or None.

7.2.5 Link Reconfigurability

For adapters that support it, Link Reconfigurability allows you to adjust the number of ports and port speeds on the adapter by selecting a specific Link Reconfigurable setting. The adapter provides a list of Link Reconfigurable descriptions (for example, one port, 40Gb/s, SFP+) and Link IDs.

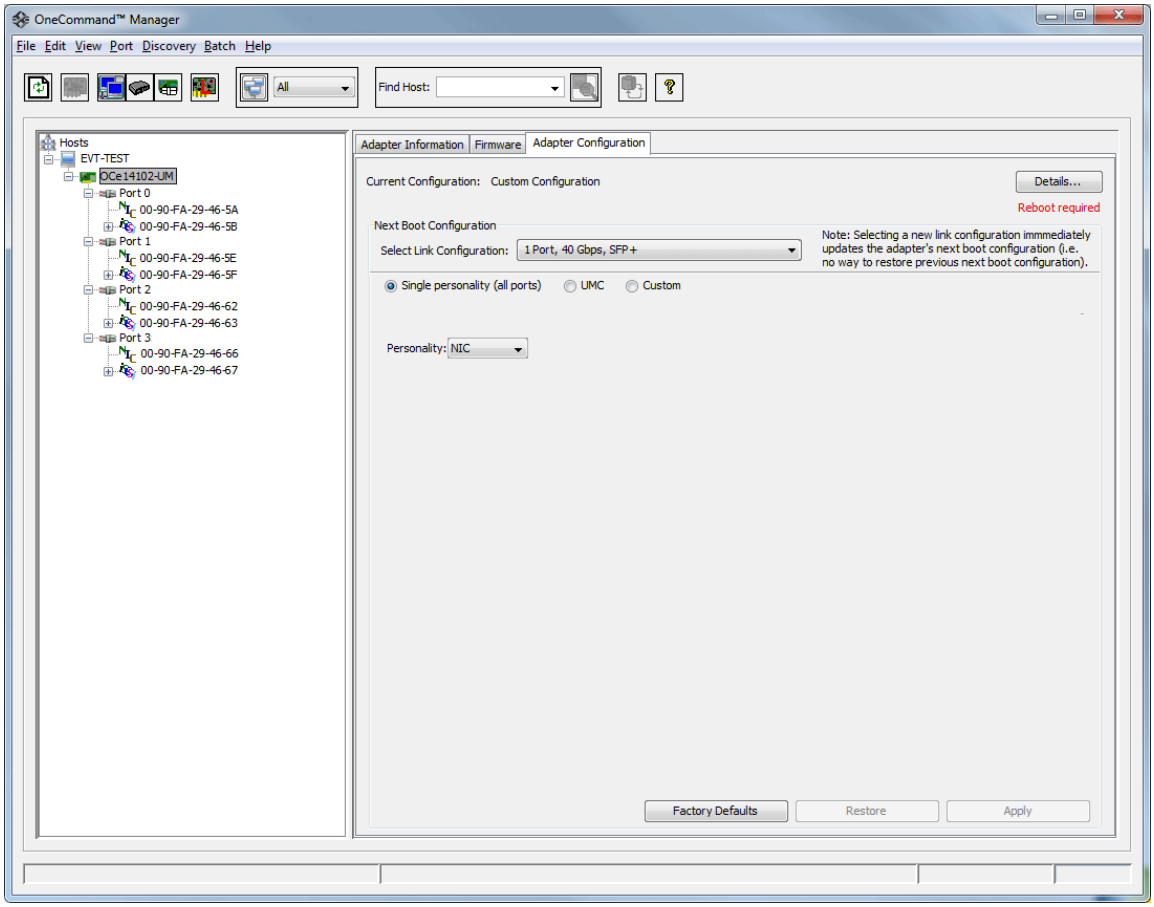
For supported adapters, the Select Link Configuration menu on the **Adapter Information** tab enables you to select the link configuration. An example of this is shown in Figure 46.

Figure 46 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 47.

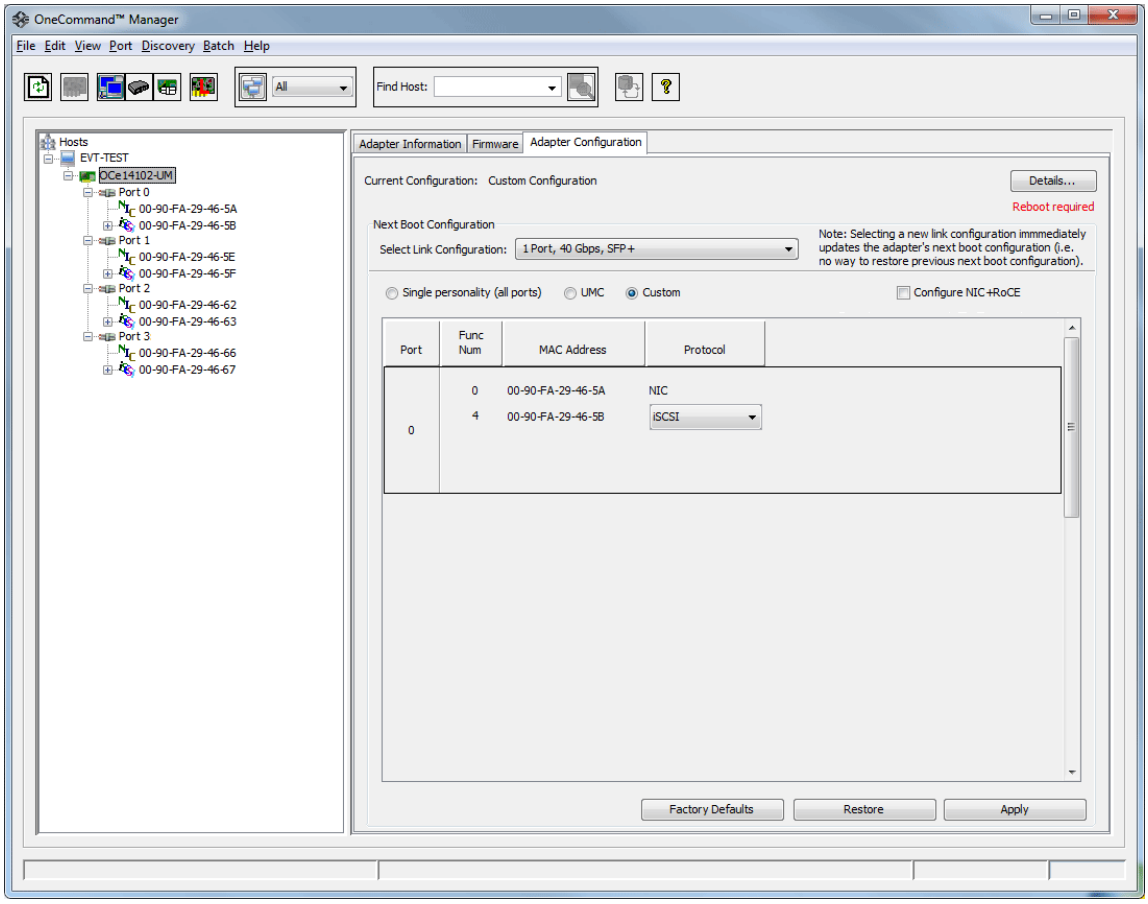
Figure 47 Adapter Configuration Tab After Selecting a New Link Configuration



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 47 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).

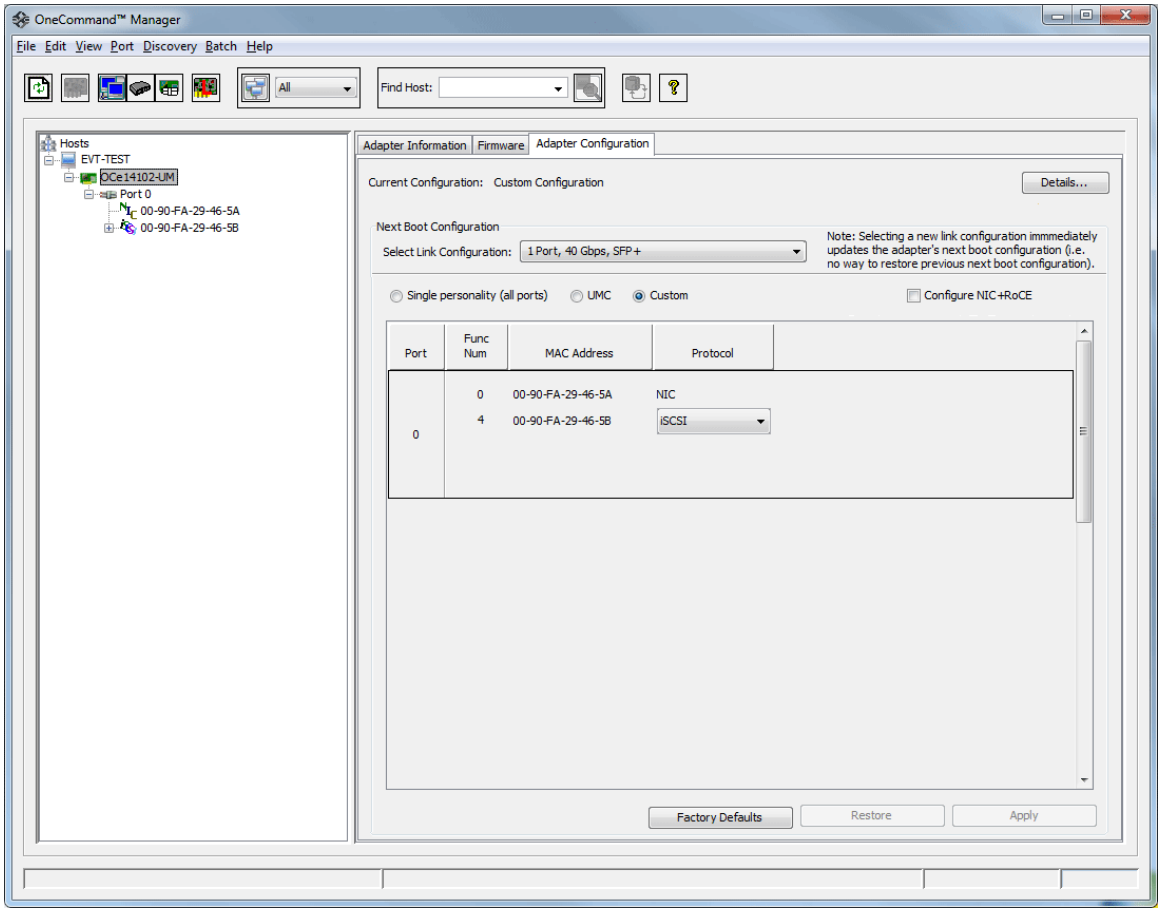
Before rebooting, you can change the default port configuration that was set on the adapter by changing the link configuration. For example, you can change the single port configuration from a NIC-only configuration to NIC+iSCSI as shown in Figure 48.

Figure 48 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 49.

Figure 49 Adapter Configuration Tab After Reboot with Single Port Configuration



7.2.5.1 Adapter Configuration Buttons with Link Reconfigurability

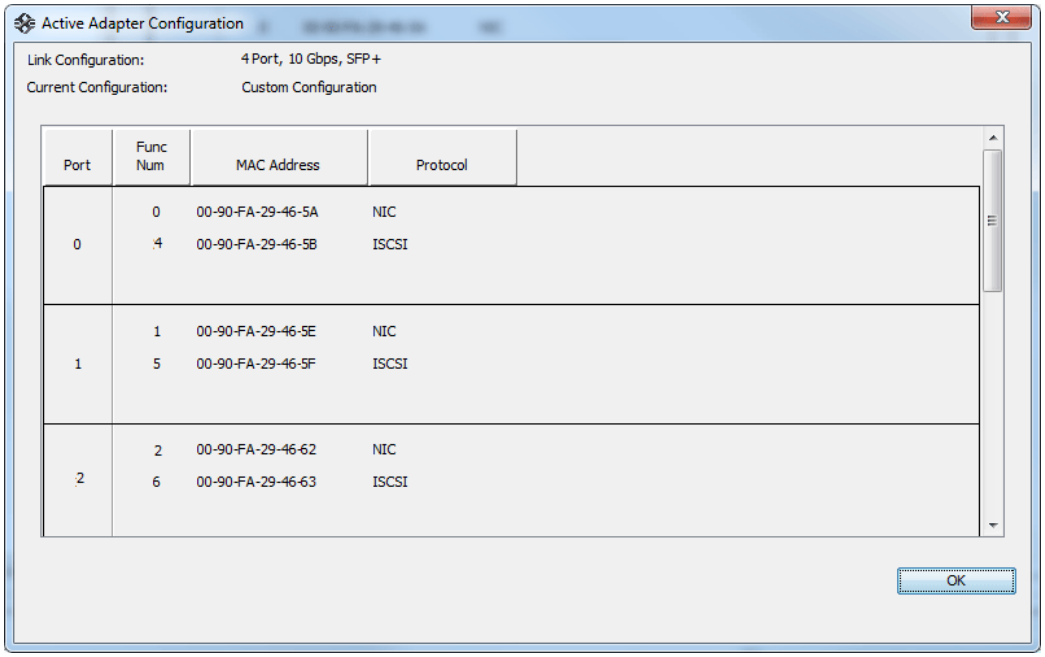
To returns the adapter's ports to their factory default profile and settings for that adapter, click **Factory Defaults**. This also returns the adapter to its default link configuration. For example, if the current adapter's link configuration was for four ports, but the factory default link configuration for the adapter is one port, clicking **Factory Defaults** changes the adapter's link configuration to one port.

To revert any unsaved changes back to the currently saved settings on the adapter, click **Restore**.

NOTE Changing the link configuration immediately updates the adapter's next boot configuration, so 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.

To display a pop-up window with more information about the current configuration running, click **Details**. If Link Reconfigurability is available, the current link configuration description is displayed at the top of the window as shown in Figure 50.

Figure 50 Active Adapter Configuration Window Showing Active Link Configuration



7.2.6 Dell NPar Configuration View

NOTE For NPar, functions are referred to as partitions.

For Dell adapters, the **Adapter Configuration** tab (Figure 51) is different than other adapters. Single personality, multichannel, 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 CNA 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 allocated12%. The maximum bandwidth for all partitions defaults to 100%. You can change the bandwidths before applying the changes. However, the minimum bandwidths must add up to 100%.

Considerations

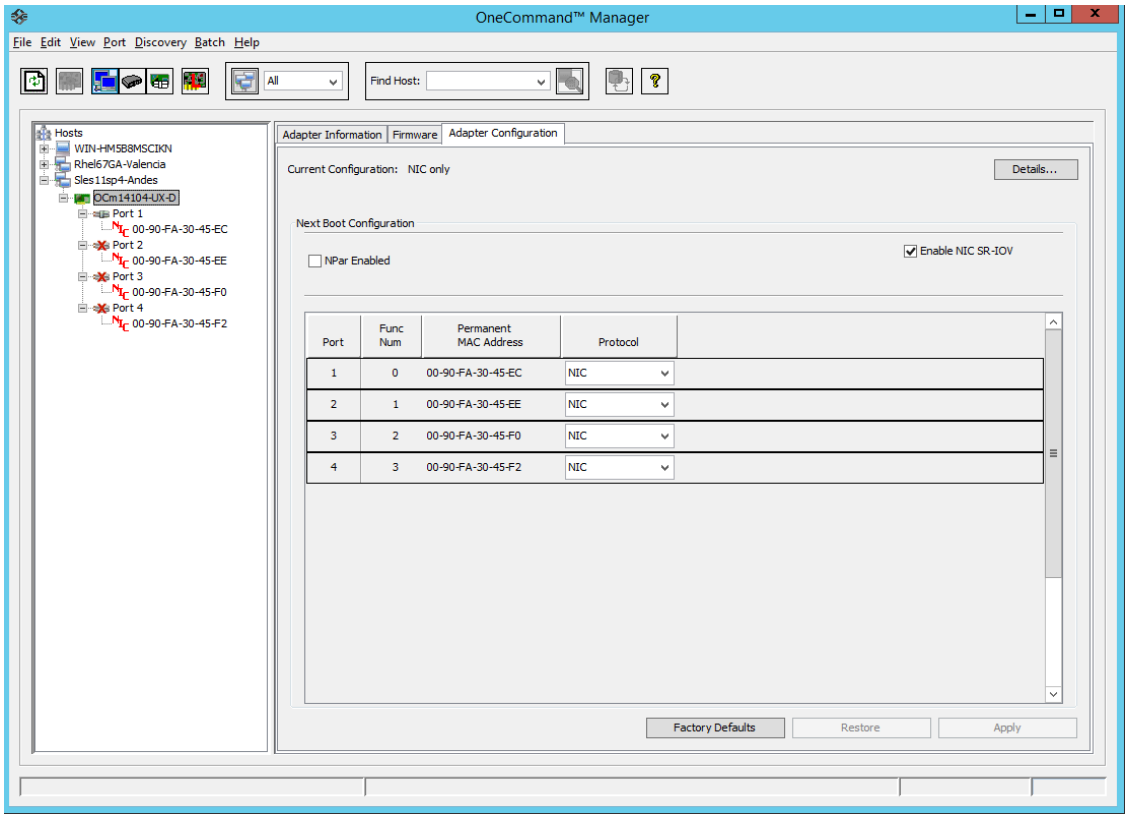
- NParEP must be available to support up to sixteen functions on an adapter. OCe14000-series adapters automatically support NParEP. However, the system's motherboard must support NParEP, it must be enabled in the system BIOS, and the operating system must support NParEP.
- On Linux and VMware systems, 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 are present and discovered by the OneCommand CNA Manager application after a reboot.
- SR-IOV can be enabled (for the first NIC function on each port) when NPar is enabled. Additionally, SR-IOV is enabled at the adapter level instead of the port level. When switching between NPar disabled and enabled, the SR-IOV state of the adapter is preserved.
  - For example, if NPar is disabled and SR-IOV is enabled, when you enable NPar, SR-IOV continues to be enabled.

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 cleared), only a single partition per port is available. When NPAR is disabled and all ports are configured to run NIC functions, the **Adapter Configuration** tab looks like Figure 51.

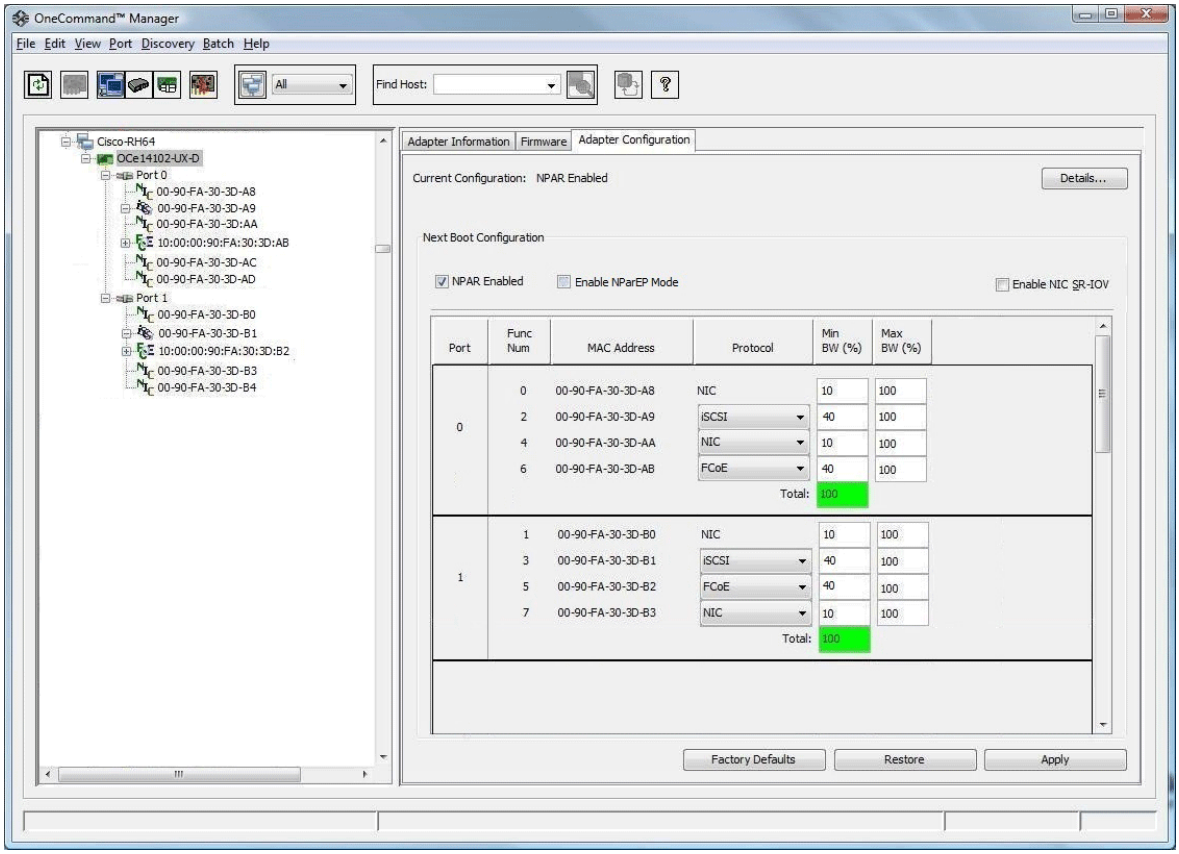
Figure 51 Adapter Configuration Tab for NPar Adapters (NPar Disabled, NIC-only Configuration)



The Enable NIC SR-IOV checkbox can be used to enable SR-IOV on the single NIC function on each port. When the Enable NIC SR-IOV is checked, the NIC-Only Configuration drop-down is displayed, from which the user can select either the NIC or NIC – ETS Disabled profile.

When NPAR is enabled, SR-IOV can be enabled at any time. Also, NParEP Mode has no effect on the ability to enable or disable SR-IOV. If NPar is enabled, the **Adapter Configuration** tab looks like Figure 54.

Figure 52 Adapter Configuration Tab with NPar Enabled (NParEP Mode Disabled)



The Enable NIC SR-IOV checkbox can be checked for any NPar enabled configuration. When it is checked, SR-IOV is only enabled on the first NIC function on each port. All other NIC functions on the port do not have SR-IOV enabled.

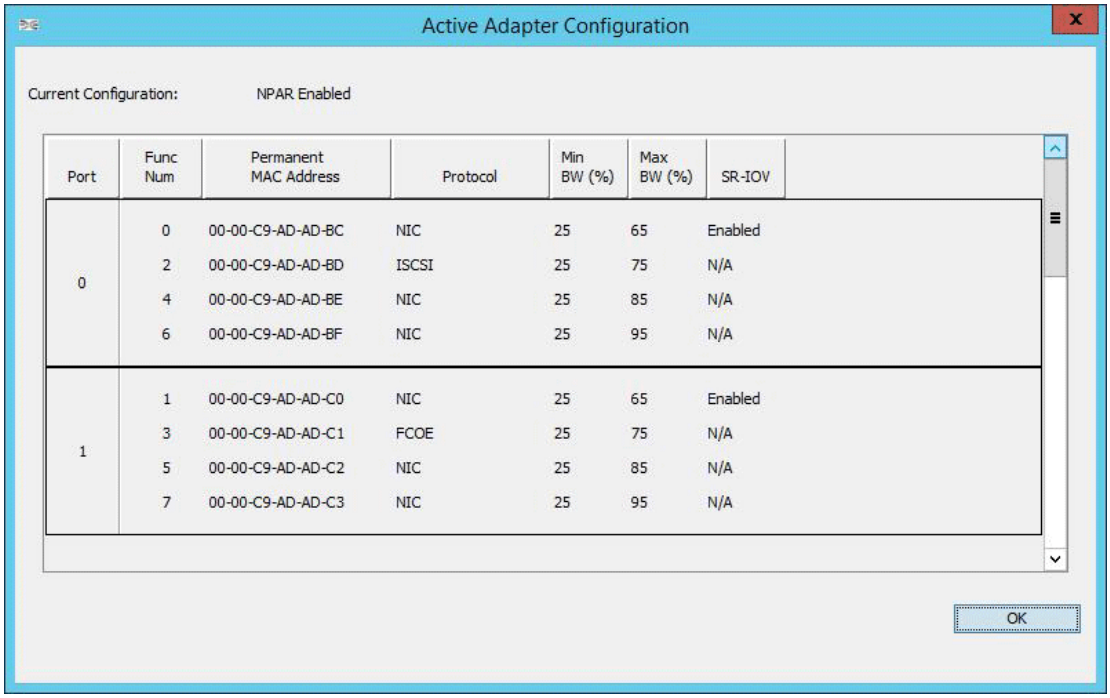
When you switch between SR-IOV enabled and disabled, the port configuration is returned to the default setting, where each partition is reset to the NIC protocol, the minimum bandwidths are evenly divided between the partitions, and the maximum bandwidth is set to 100. This automatically occurs each time the Enable NIC SR-IOV checkbox is checked or unchecked.

> For example, if you have changed any of the partitions to storage protocols or you have updated any of the minimum or maximum bandwidths, these changes are lost when the Enable NIC SR-IOV checkbox is checked or unchecked. This behavior is also true for the Enable NParEP mode checkbox. However, enabling or disabling SR-IOV has no effect on the current NParEP mode setting, which is retained whenever SR-IOV is enabled or disabled. In addition, enabling or disabling NParEP has no effect on the current SR-IOV setting, which is retained whenever NParEP is enabled or disabled.

Unlike the situation where NPar is disabled and a NIC-only configuration is specified, the NIC-Only Configuration Type drop-down is not available. Instead, the NIC profile is selected automatically because the NIC – ETS Disabled profile is not an available option when NPar is enabled.

When NPAR is enabled in the active configuration, the Active Adapter Configuration dialog box includes an extra column labeled SR-IOV. This column contains the active SR-IOV status for each function. The Active Adapter Configuration dialog looks like Figure 53.

Figure 53 Active Adapter Configuration Dialog Box for an NPar Enabled Adapter



For functions that do not support SR-IOV (FCoE/iSCSI and all other NIC functions except the first one on each port), the SR-IOV status is reported as N/A.

7.2.7 Configuring RoCE on NPar Adapters

You can configure NIC+RoCE on NPar adapters if NPar is disabled (Figure 51).

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 Knowledge Base on http://www.broadcom.com for any updated information on additional use.

NOTE RoCE configurations are not supported with SR-IOV.

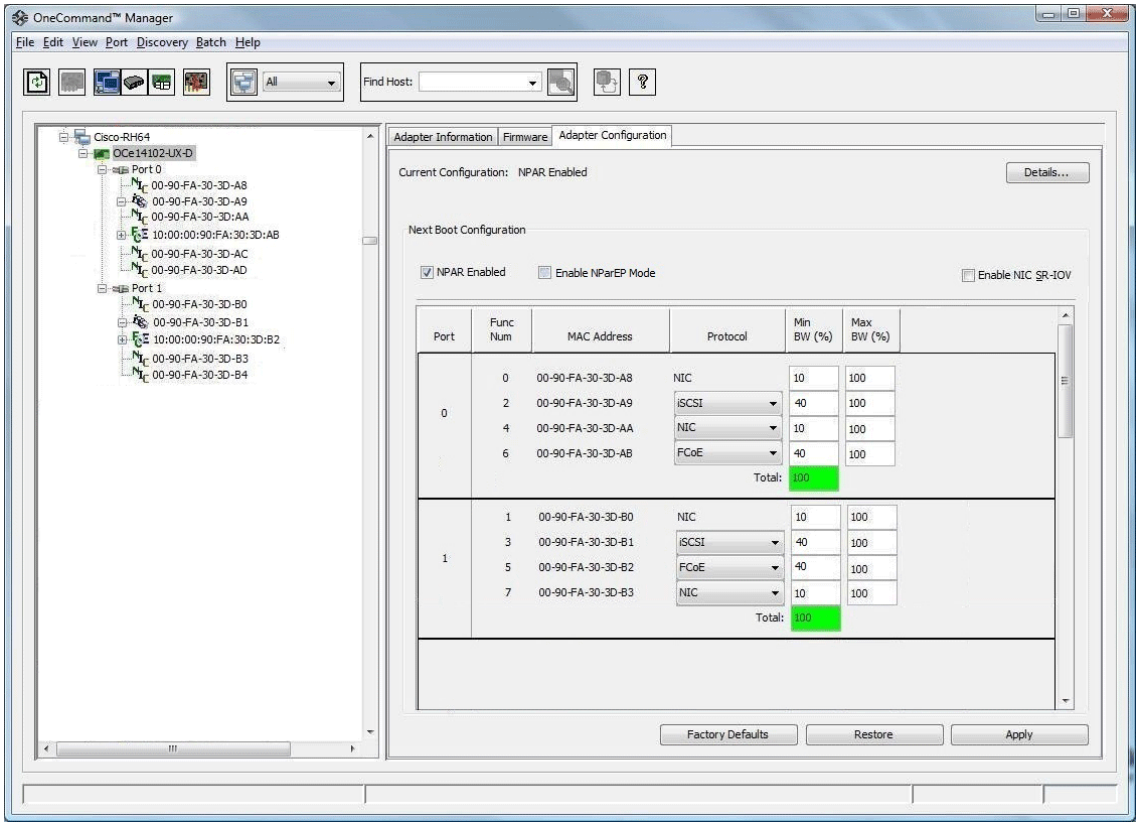
7.2.7.1 Dell NPar Enabled

If NPar is enabled, the **Adapter Configuration** tab looks like Figure 54.

NOTE 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.

Figure 54 Adapter Configuration Tab with NPar Enabled (NParEP Mode Disabled)

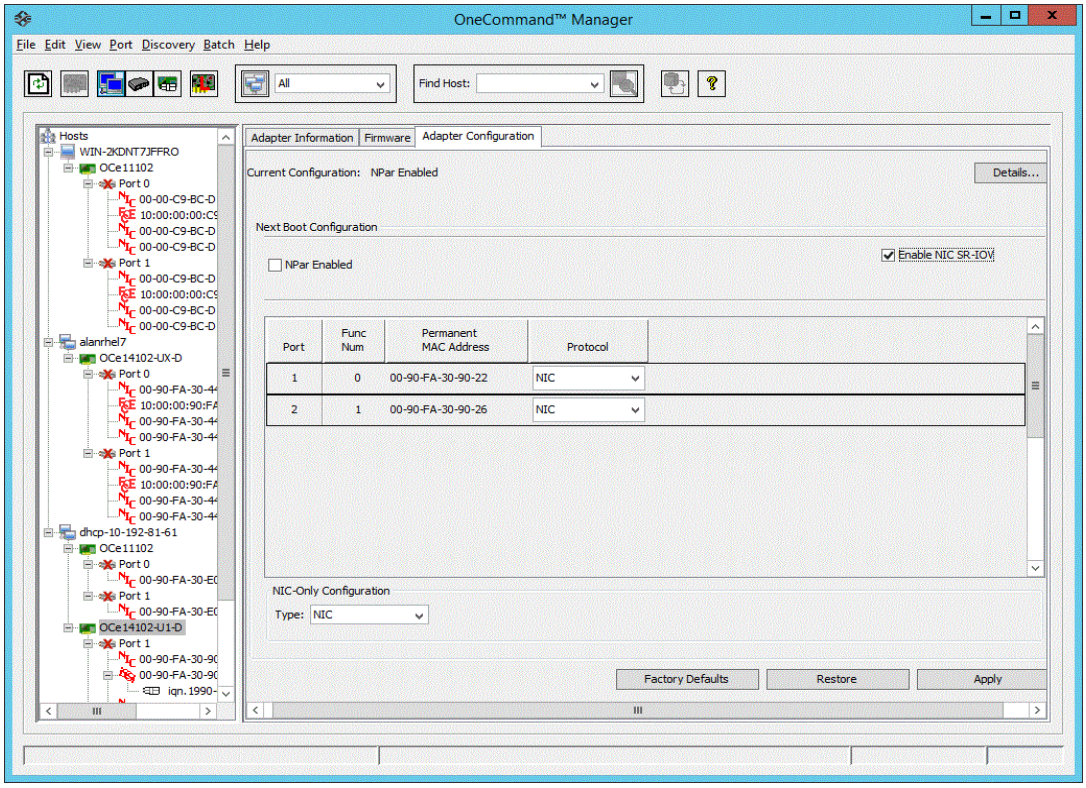


For adapters that support sixteen functions, on two port adapters you can configure up to eight partitions per port. Four-port adapters allow up to four partitions per port. The None setting can be configured on any but the first partition.

NOTE If NPar is enabled, RoCE cannot be configured on any partition.

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. The **Adapter Configuration** tab looks like Figure 55.

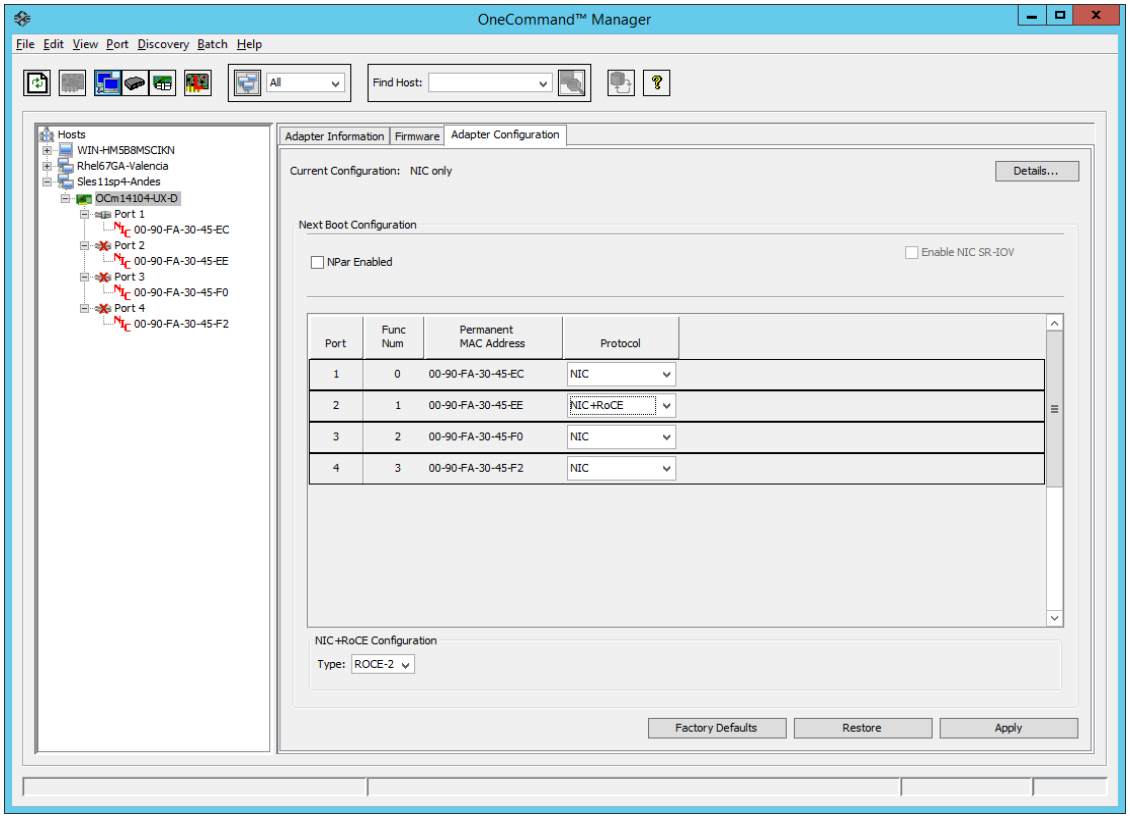
Figure 55 Adapter Configuration Tab with NPar and NIC SR-IOV Enabled



When NPar is disabled and one or more ports are configured for NIC+RoCE functions, the **Adapter Configuration** tab looks like Figure 56.

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 is automatically cleared and disabled.

Figure 56 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 34, 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 CNA 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 35). However, you can enable or disable SR-IOV on the entire adapter on a OneCommand CNA Manager host by checking or clearing the Enable NIC SR-IOV checkbox and clicking the **Apply** button.

7.2.7.2 Dell NPar and NParEP Mode Enabled

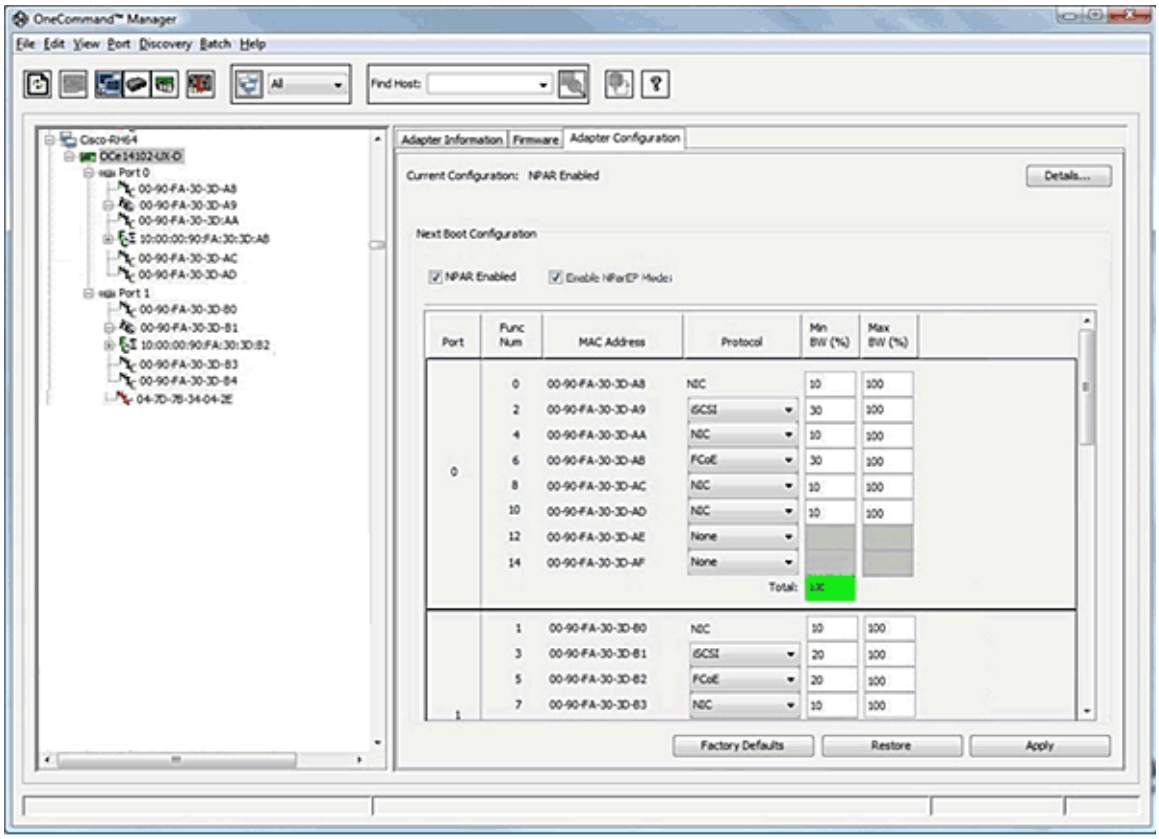
NOTE On Linux and VMware systems, SR-IOV must be enabled on the system BIOS if NParEP is used. See the documentation that accompanied your Dell server for more information.

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 CNA Manager application. If NParEP Mode is enabled, you can view and configure sixteen partitions per adapter.

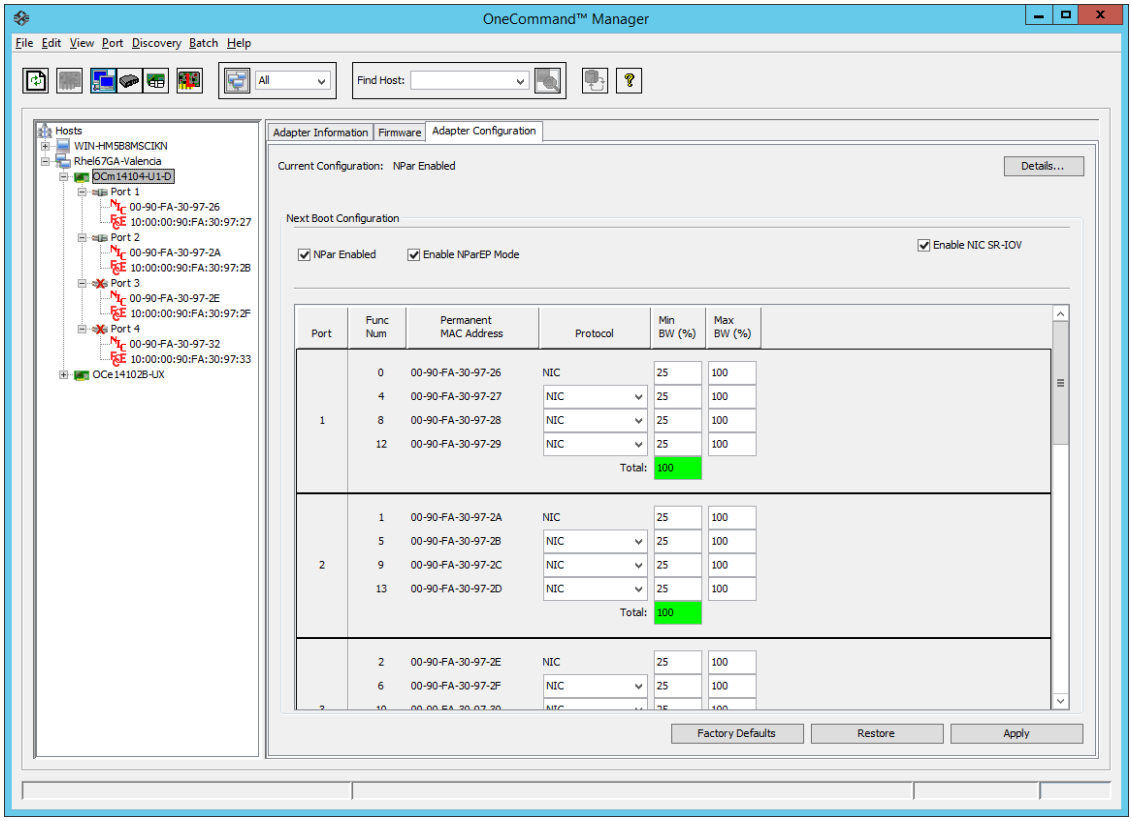
If NParEP Mode is enabled on 2-port adapters, up to eight partitions per port can be configured (Figure 57).

Figure 57 Adapter Configuration Tab with NPar and NParEP Mode Enabled (2-Port configuration)



Four-port adapters allow up four partitions per port to be configured if NParEP Mode is enabled (Figure 58).

Figure 58 Adapter Configuration Tab with NPar and NParEP Mode Enabled (4-Port configuration)



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.

7.2.7.3 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 of amount of data passes 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.

7.2.8 Removing a Port

Port removal allows you to prevent the appearance of the port and its functions on the PCiE bus. It is supported for non-NPar OCe14000-series adapters. A port can only be removed in Multichannel and Custom modes. This option is not available in Single Personality mode, since this type of configuration does not allow configuring ports and functions individually or separately.

NOTE The first port of an adapter cannot be disabled. This is because function 0 (which is always on the first port) cannot be set to None.

To remove a port:

Set the first function to None in all except the first port in Multichannel and Custom modes.

When all functions belonging to a port are set to None, the port is removed and is not discovered after reboot.

To add a port back, perform one of the following:

- Set the first function in the removed port to NIC or NIC+RoCE protocol.
- Configure the adapter to single personality, which causes all the removed ports to be added back after reboot.

7.2.8.1 Custom View Configuration

The Custom view configuration provides a drop-down menu, which includes None as one of the items in the list for all the first functions belonging to ports other than the first port. Function 0 (on the first port) is not able to change its protocol (see Figure 59). By setting the first function on the port to None, all of the other functions on that port are automatically set to None.

For the NIC + RoCE configuration, None is available as a drop-down menu item for functions greater than 0 (see Figure 60).

Figure 59 Port Removal, Custom Configuration

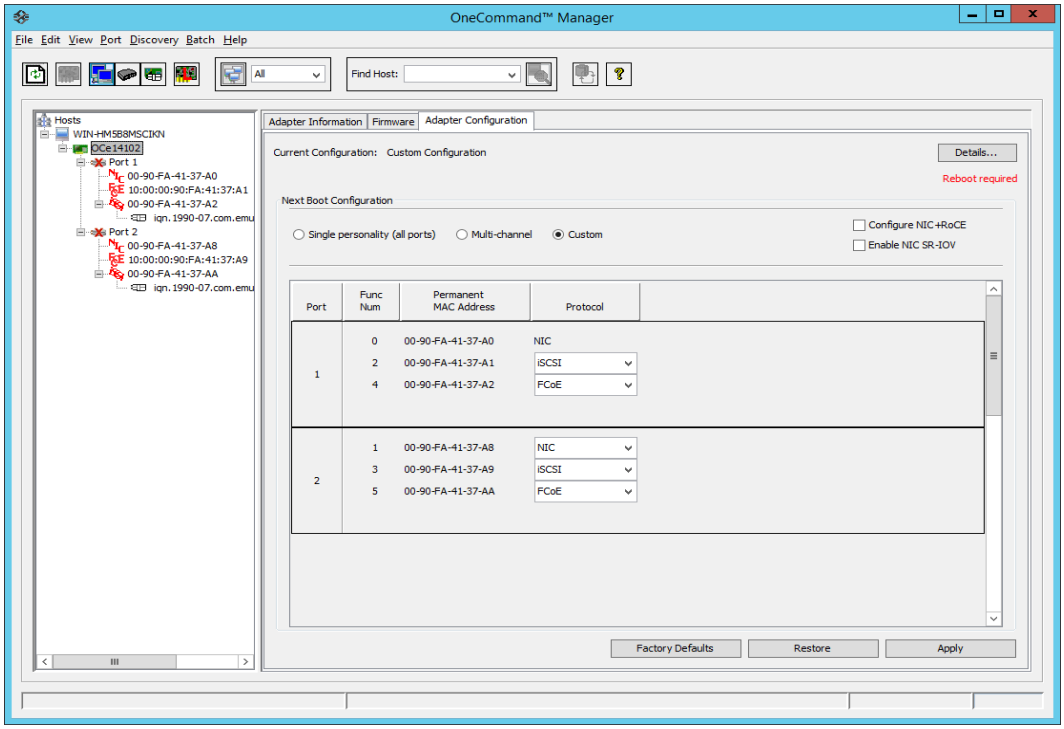
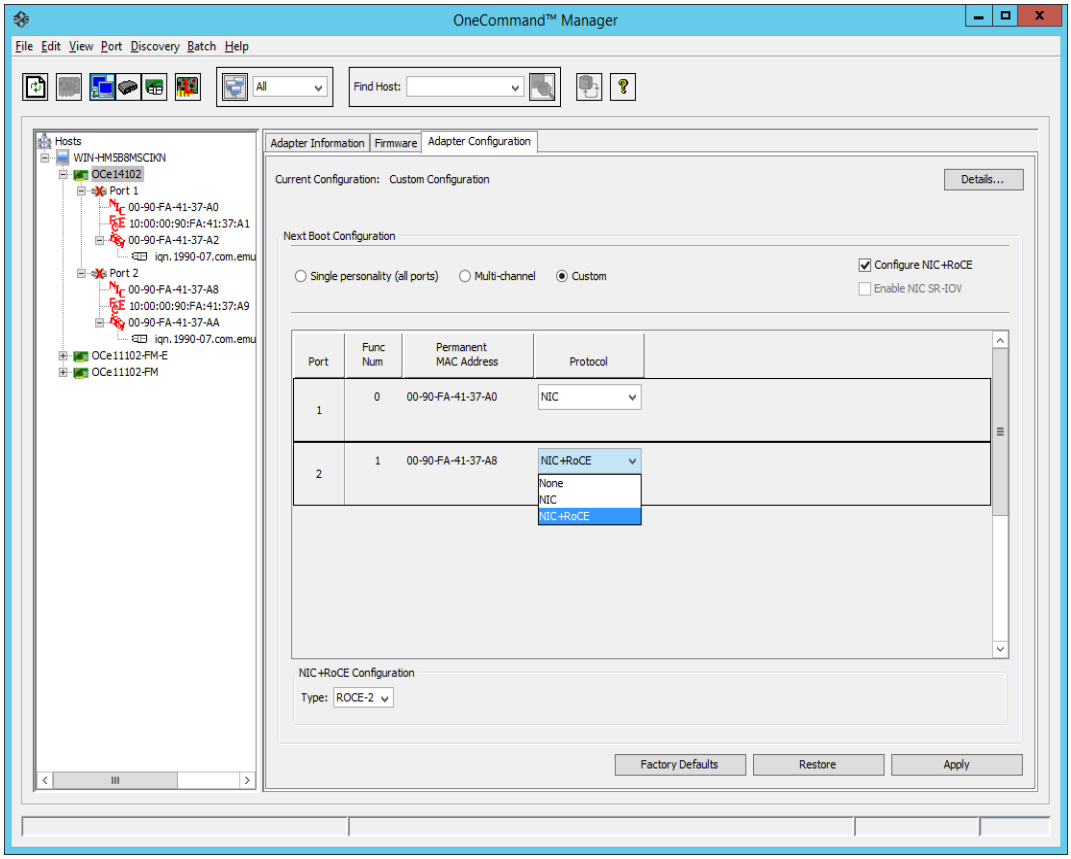


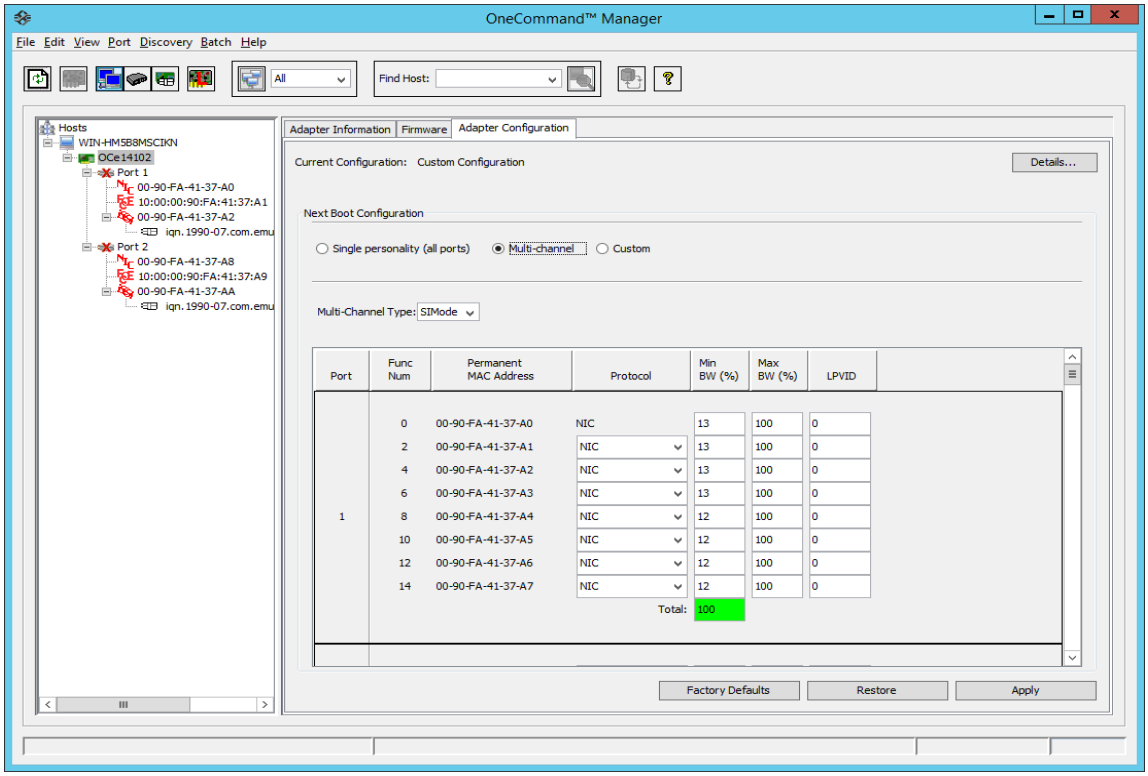
Figure 60 Port Removal, Custom View for the NIC + RoCE Configuration



7.2.8.2 Multichannel Configuration View (UMC, SIMode, vNIC, UFP)

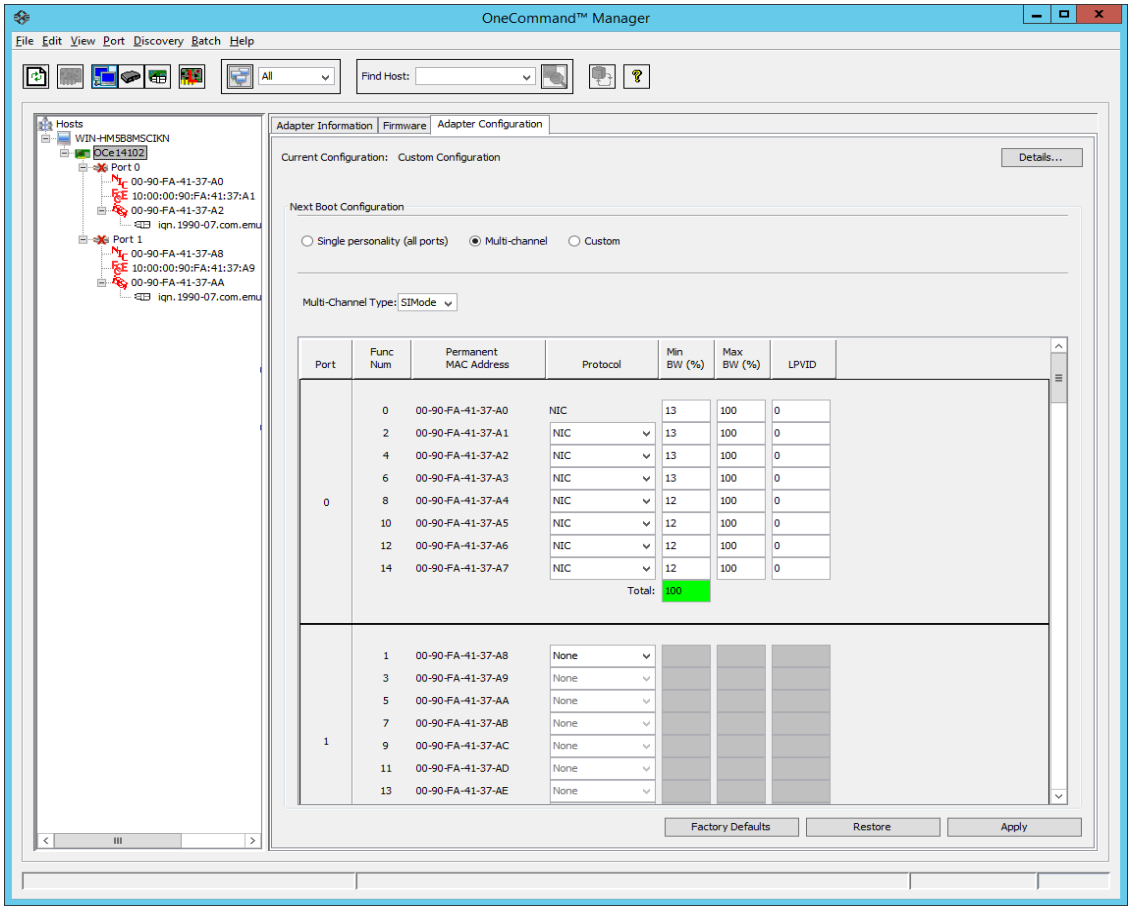
The Multichannel view configuration provides a drop-down menu, which includes None as one of the items in the list, for all the first functions belonging to ports other than the first port. Function 0 (on the first port) is not able to change its protocol (see Figure 61).

Figure 61 Port Removal, Multichannel View



If the first function belonging to ports other than the first port is set to None, then all the remaining functions from that port are forced to None, and the protocol drop-down menus become disabled to prevent changing protocols. The drop-down menus become enabled when the first function is reset to NIC (see Figure 62).

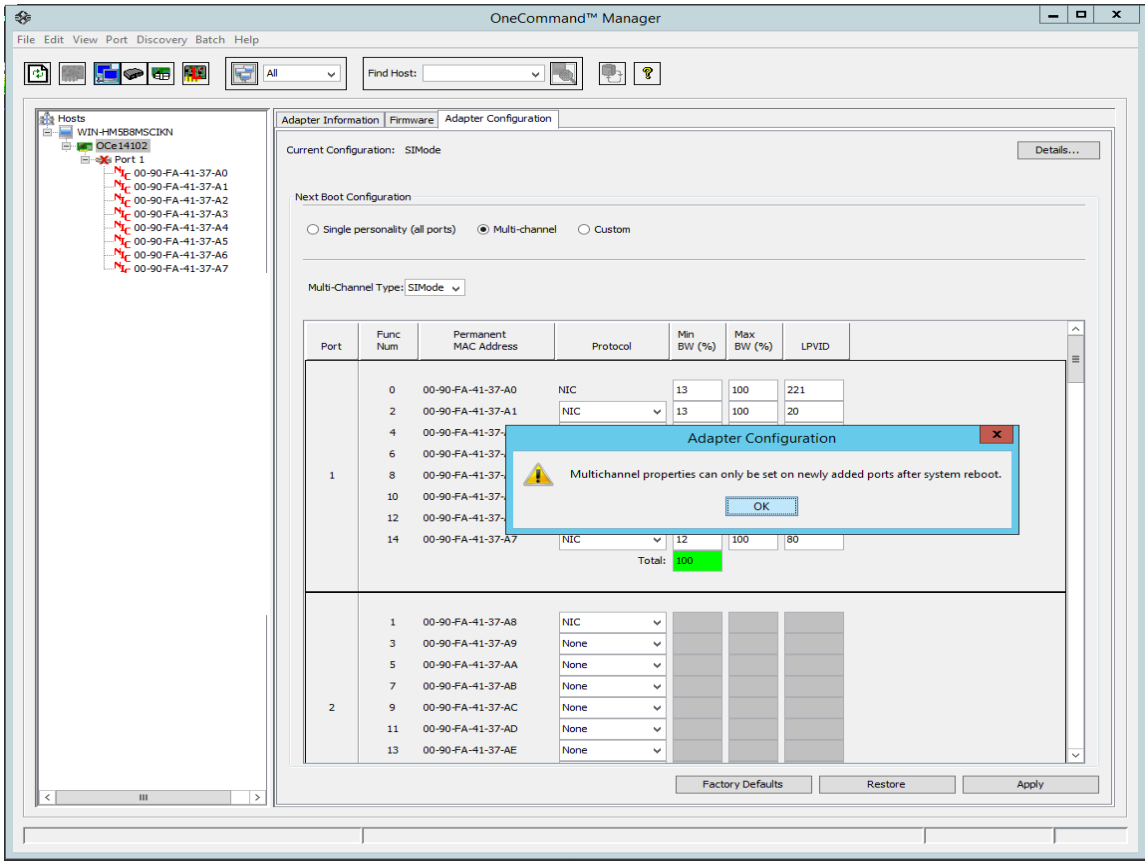
Figure 62 Port Removal, First Function from Port 1 Set to None



7.2.8.3 Adding Ports in Multichannel Mode (UMC, SIMode, vNIC, UFP)

When you add a port while running in multichannel mode, you cannot set any multichannel properties for the port being added until after a system reboot. The Min BW, Max BW, and LPVID (for UMC and SIMode) text field boxes (or any other multichannel property box) remain unavailable. A pop-up box (see Figure 63) is also displayed immediately after clicking **Apply**, indicating that multichannel properties can only be set on this port after a system reboot has been performed.

Figure 63 Port Removal, Re-enabling the Port in Multichannel Mode



NOTE The pop-up message continues to be displayed whenever you select the **Adapter Configuration** tab if a reboot is still required to enable the port.

7.2.9 Configuring HP Flex Adapters

The OneCommand CNA 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 to the Flex mode capability of the adapter.

NOTE Even though the OneCommand CNA 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.

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 shown in Figure 64.

Figure 64 HP Flex Adapter Configuration Tab (Flex Enabled)

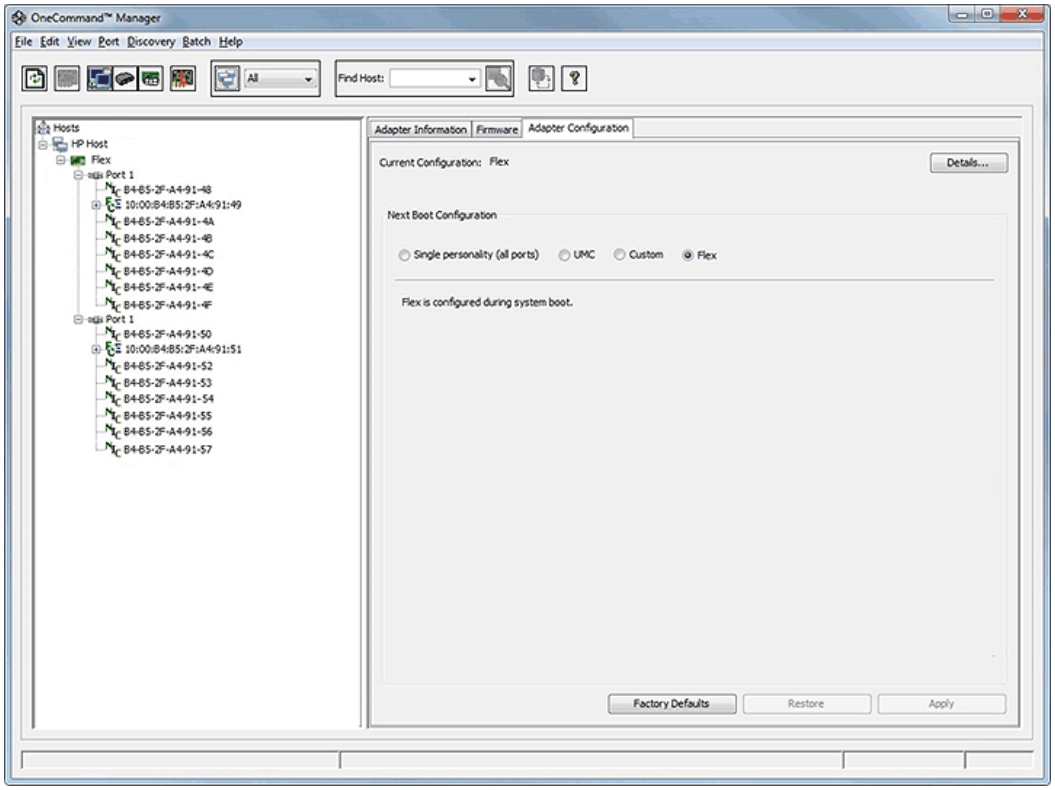
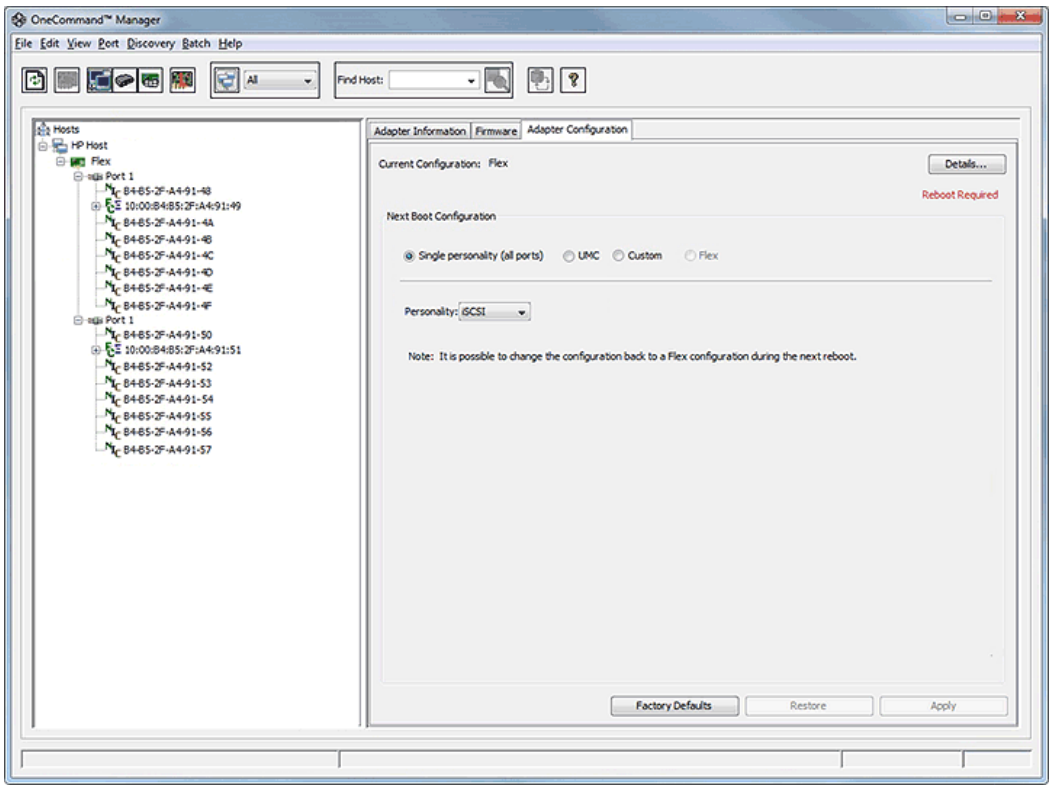


Figure 65 Flex Adapter After Applying an iSCSI Single Personality Change



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 Section 7.2.3, UMC Configuration View for UMC configuration instructions.

See Section 7.2.2, Custom Configurations for instructions on custom configurations.

Chapter 8: Managing Protocols

This section describes how to manage the available protocols.

8.1 NIC

This section describes managing NIC functions.

NOTE If you change NIC driver properties (such as packet size) using the Windows Driver Properties applet while the OneCommand CNA Manager application is running, the OneCommand CNA Manager application may no longer be able to manage OneConnect adapters. As a result, most data fields show n/a, adapter settings are not be configurable, and firmware download is not available. If this occurs, exit and restart the OneCommand CNA Manager application.

8.1.1 Viewing NIC Function Information

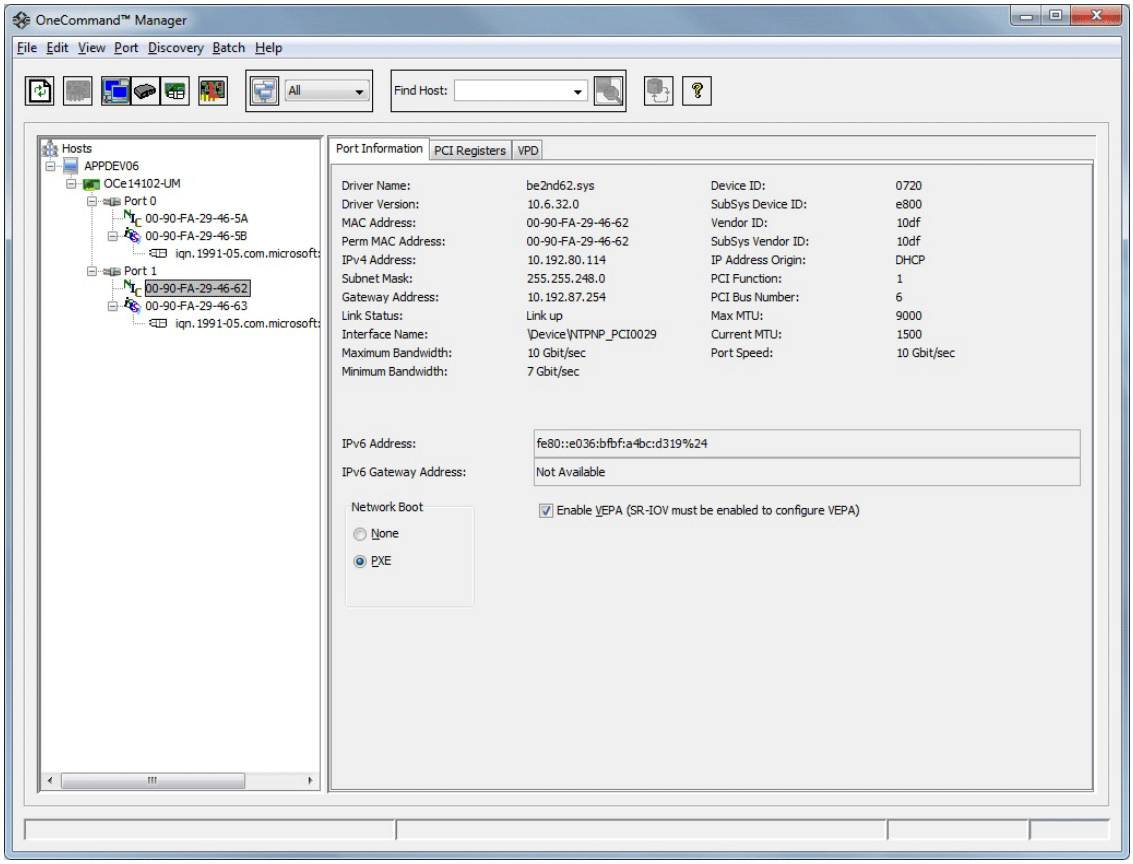
When you select a NIC function from the discovery-tree (Figure 5), the NIC **Port Information** tab (Figure 66) contains general attributes associated with the selected NIC function.

The NIC **Port Information** tab 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 5).
3. Select the **Port Information** tab (Figure 66).

Figure 66 NIC Port Information Tab



The following NIC **Port Information** Tab fields are displayed:

- 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.

NOTE When an OCe11000-series adapter is configured in multichannel disabled mode, the minimum bandwidth for NIC functions will always be shown as zero.

- 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 area:
  - 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.

For information about VEPA, see Section 8.1.4, Enabling and Disabling VEPA on NIC SR-IOV Functions (OCe14000-series Adapters Only).

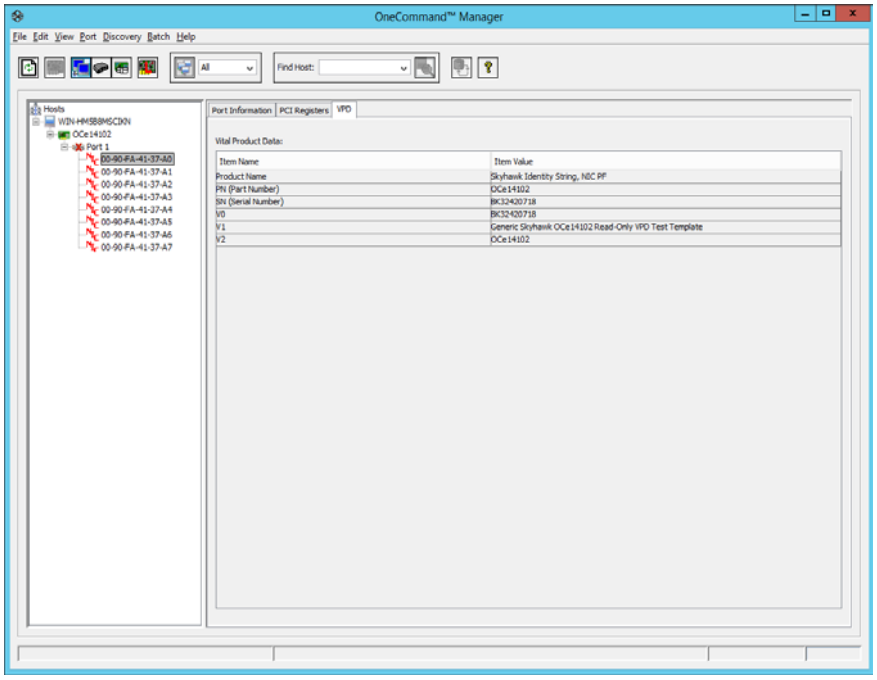
8.1.2 Viewing NIC VPD Information

The VPD tab (Figure 67) displays vital product data, such as the product name, part number, serial number (if available), for the selected NIC function.

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 67).

Figure 67 NIC VPD Tab



The following **VPD** Tab fields are displayed:

- 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 V*x* (where *x* is a number).

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

8.1.3 Enabling and Disabling SR-IOV on NIC Functions (on OCe11000-series Adapters)

For OCe14000-series adapters, refer to Section 7.2, OCe14000-Series Adapter Configuration.

NOTES

- This section only applies when running the OneCommand CNA Manager application on supported SR-IOV adapters and operating system platforms.
- SR-IOV is not supported with UMC.

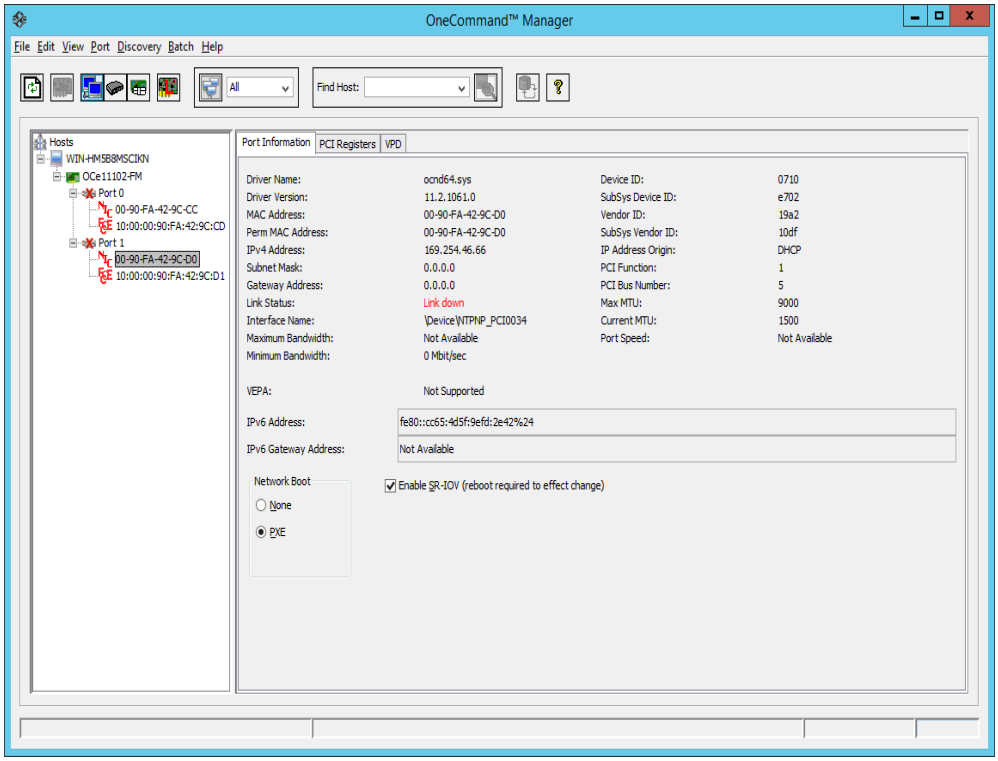
If SR-IOV is available, the NIC **Port Information** tab (Figure 66) allows you to enable SR-IOV on NIC adapter functions of an adapter. If SR-IOV is enabled, multiple VFs can be created on a NIC PF for an adapter port. These VFs are assigned by virtual operating systems such as Microsoft Hyper-V and 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 5) displays the VFs running on the discovered PFs. It also shows the selected VFs 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 5).
3. Select the **Port Information** tab (Figure 68).
4. Check or clear **Enable SR-IOV**.
5. Reboot to change the SR-IOV state.

Figure 68 Port Information Tab, SR-IOV Enabled



8.1.4 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 clear the Enable VEPA checkbox to enable or disable VEPA for the virtual functions on the NIC port. See Figure 66. If SR-IOV is not enabled on the adapter, the VEPA checkbox is disabled.

NOTE

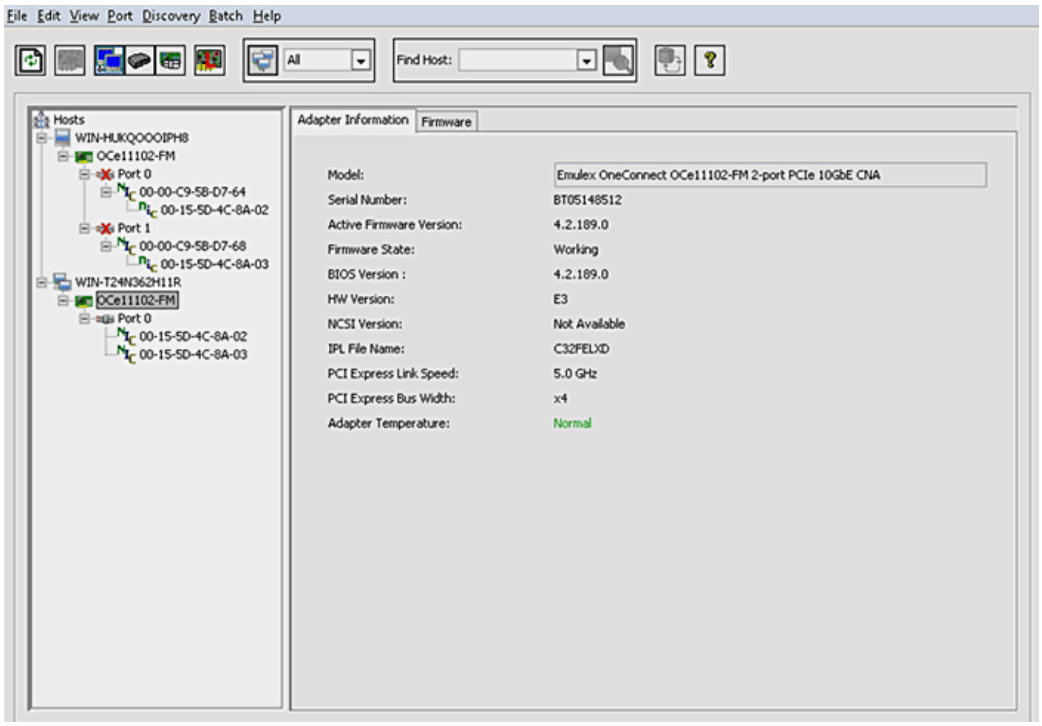
When VEPA is not available for the next boot configuration (for example, SR-IOV is disabled or multichannel is enabled), this checkbox is displayed in a disabled state.

When VEPA is currently enabled and the system is rebooted, if the adapter's configuration after rebooting does not allow VEPA (for example, the boot configuration does not have SR-IOV enabled or has multichannel enabled), VEPA is automatically disabled.

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

When the OneCommand CNA Manager application is installed on a guest operating system, the guest operating system and VF are discovered by the OneCommand CNA Manager application running on the host operating system. The guest operating system host appears as a remote host in the discovery-tree (Figure 69).

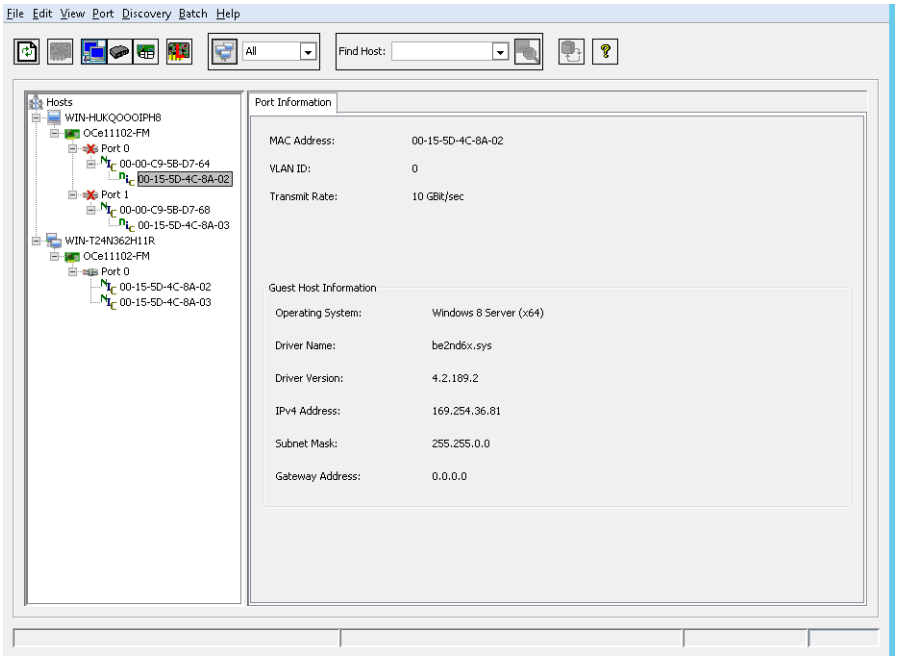
Figure 69 OneCommand CNA Manager Application Running on the Base Host Operating System after Discovering the Guest Host



The NIC function for the discovered guest operating system matches one of the VFs in the base host operating system as seen in Figure 70.

Select the VF in under the base operating system host in the discovery-tree to display the **Port Information** tab (Figure 70). Some of the information displayed in the tab is obtained from the OneCommand CNA Manager remote management agent running on the guest operating system.

Figure 70 VF Selected Showing the Port Information Tab for the Discovered NIC in the Guest Operating System



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

- 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 MB/s.
- Guest Host Information area:

  NOTE The following fields are only available when the OneCommand CNA 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 CNA Manager application running on the base host operating system. See Section 5.1, Discovery Using the TCP/IP Access Protocol.

  - 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.

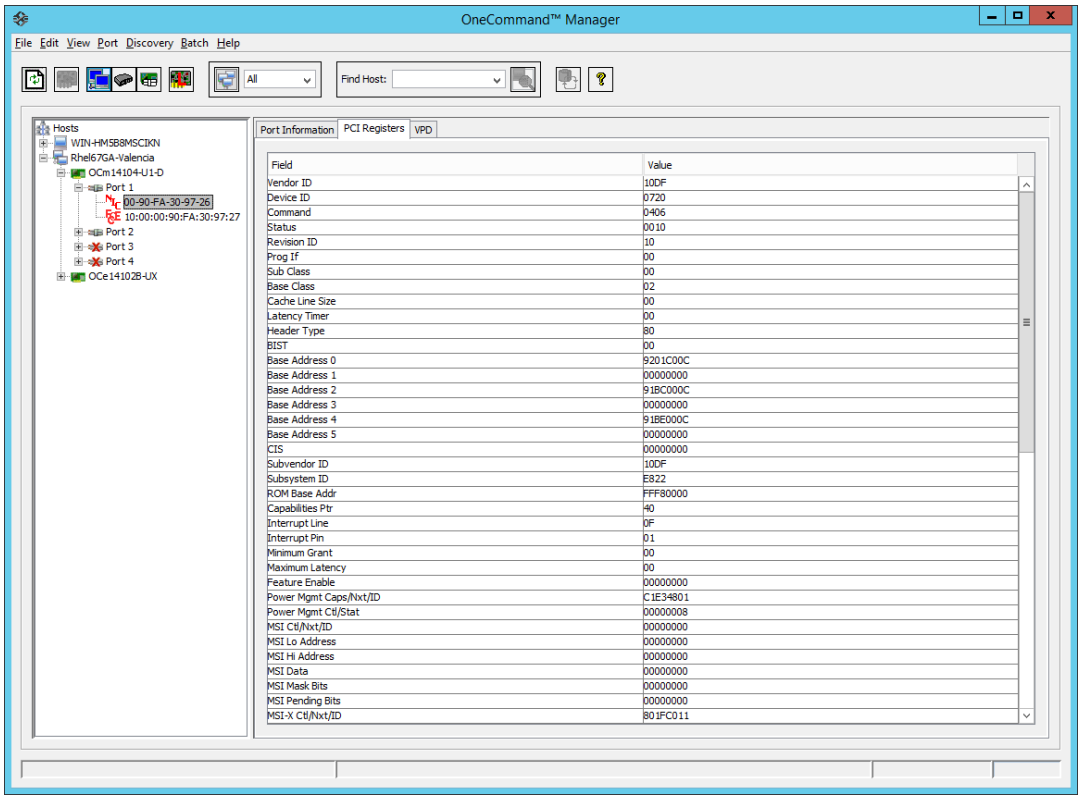
8.1.5.1 Running the OneCommand CNA Manager Application on a Guest Operating System

When the OneCommand CNA 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 CNA 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.

8.1.6 Viewing NIC PCI Registers

The NIC PCI Registers tab (Figure 71) displays base PCI registers.

Figure 71 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 71).

8.2 FCoE

This section describes managing FCoE ports.

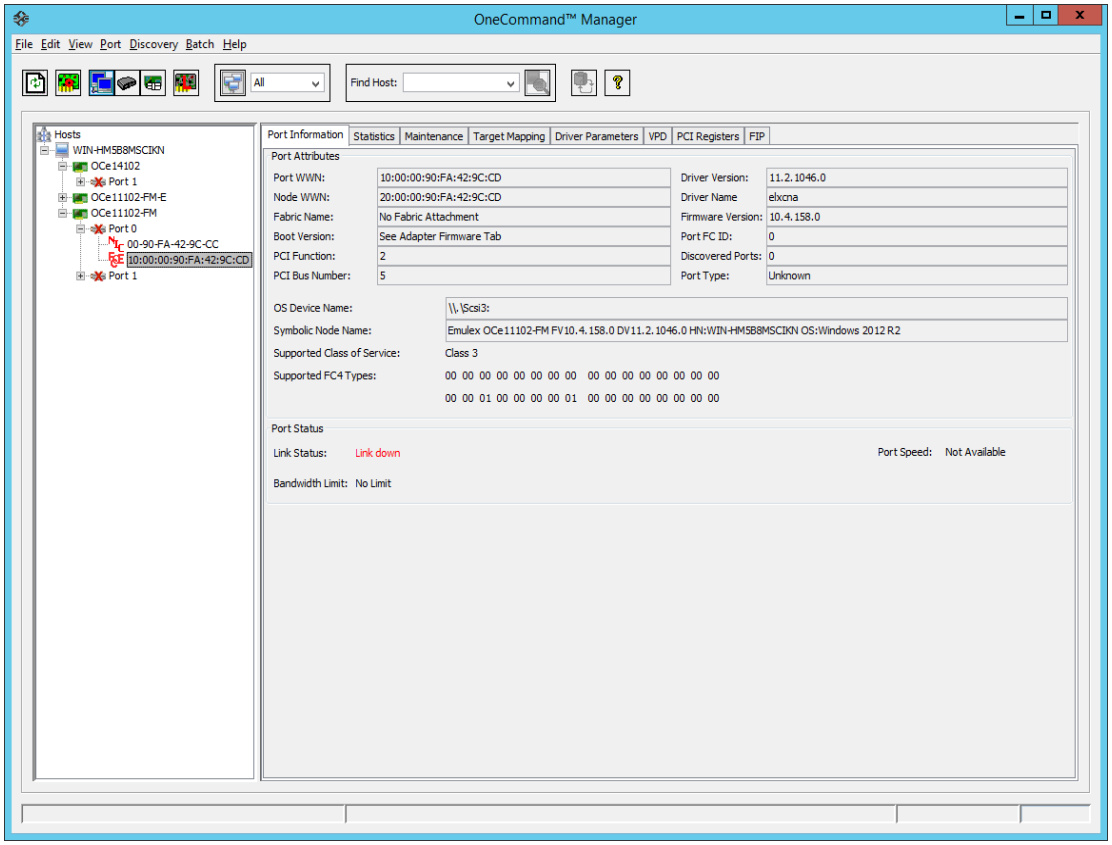
8.2.1 Viewing FCoE Function Information

When you select an FCoE function from the discovery-tree, the **Port Information** tab (Figure 72) 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 72).

Figure 72 FCoE Port Information Tab



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

- Port Attributes area:
  - Port WWN – The WWPN 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:

- 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.

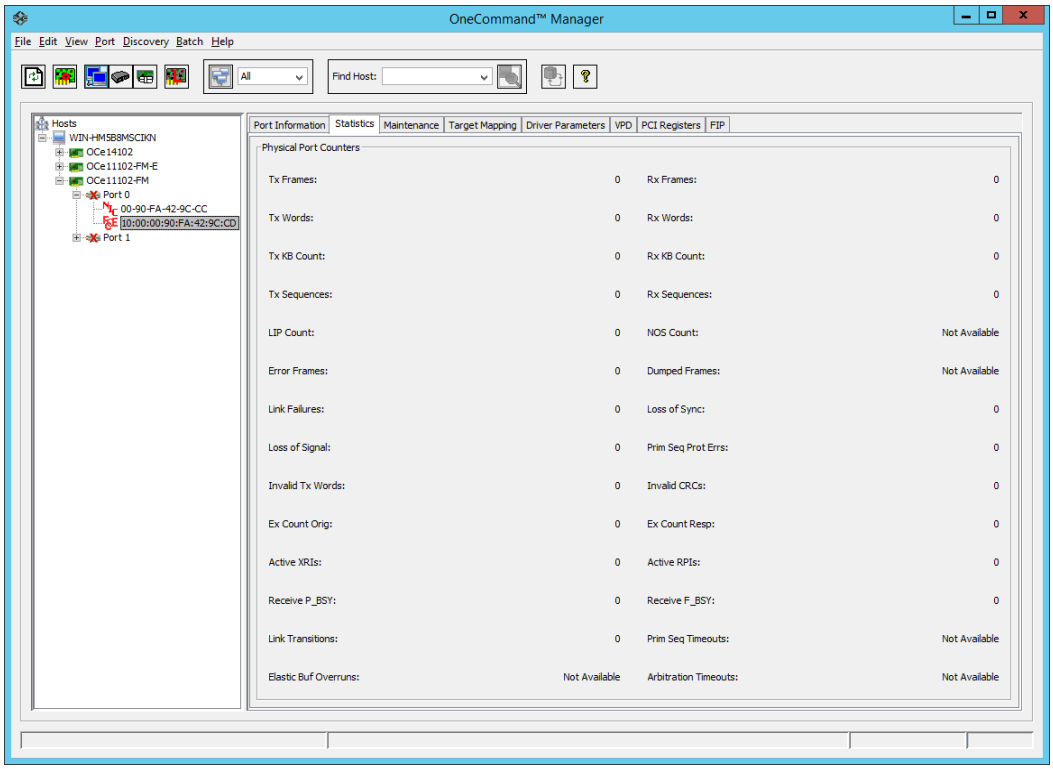
8.2.2 Viewing FCoE Statistics

When you select an FCoE function from the discovery-tree, the **Statistics** tab (Figure 73) provides cumulative totals for various error events and statistics on the port. Some statistics are cleared when the adapter is reset.

To view /FCoE statistics:

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

Figure 73 Statistics Tab



The following Port Statistics tab Fields are displayed:

- Tx Frames – FC frames transmitted by this FCoE function.
- Tx Words – FC words transmitted by this FCoE function.
- Tx KB Count – FC kilobytes transmitted by this FCoE function.
- Tx Sequences – FC sequences transmitted by this FCoE function.
- LIP count – The number of loop initialization primitive (LIP) events that have occurred for the 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.
  -
  - 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 FCoE function.
- Ex Count Orig – The number of FC exchanges originating on this 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 FCoE function.
- Rx Words – The number of FC words received by this FCoE function.
- Rx KB Count – The received kilobyte count by this FCoE function.
- Rx Sequences – The number of FC sequences received by this 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. The NOS count is not currently supported for Emulex Windows drivers or arbitrated loop.
- Dumped Frames – The number of frames that were lost because of a lack of host buffers available. This option is not currently supported for the Storport Miniport driver or the driver for Solaris.
- Loss of Sync – The number of times loss of synchronization has occurred.
- Prim Seq Prot Errs – The primitive sequence protocol error count. This counter is incremented whenever there is any type of protocol error.
- Invalid CRCs – The number of frames received that contain CRC failures.
- Ex Count Resp – The number of FC exchange responses made by this 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).

8.2.3 Viewing FCoE Virtual Port Information

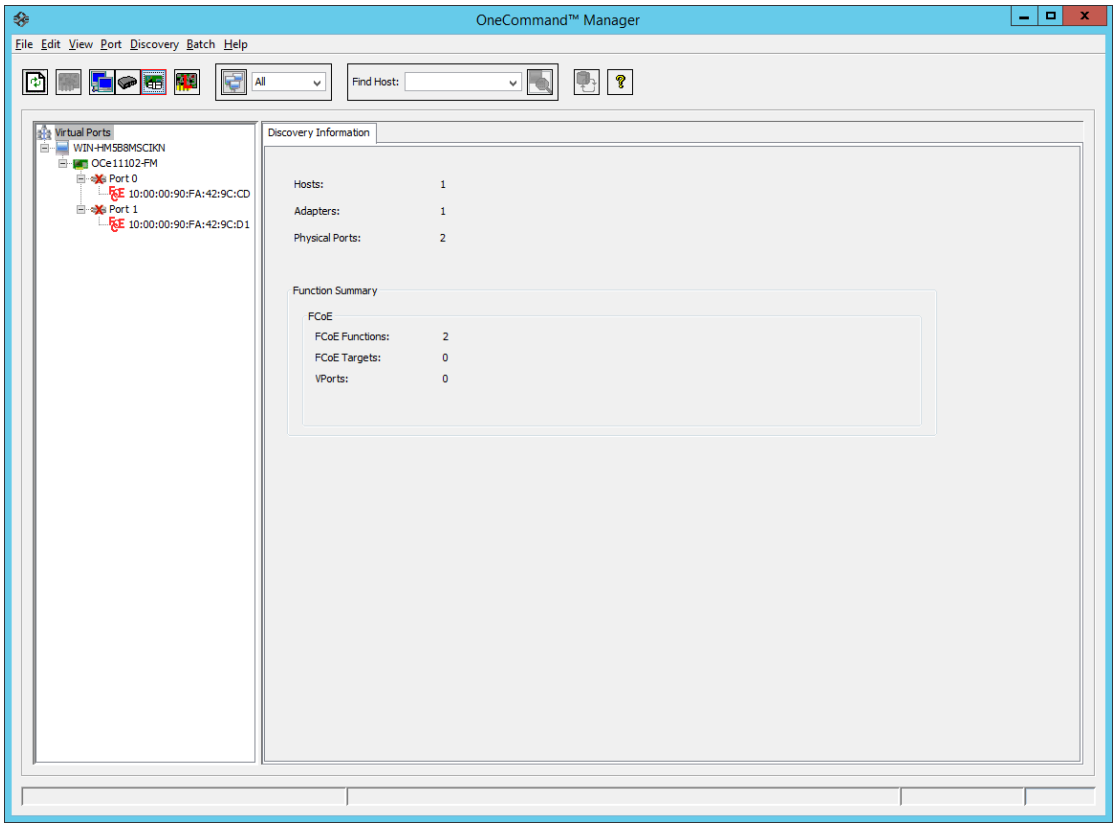
The **Discovery Information** tab (Figure 74) can display information about FCoE virtual ports and their associated targets and LUNs.

To view virtual port information:

1. Perform one of the following tasks:
   - 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 74).

Figure 74 Discovery Information Tab Displaying Virtual Ports Information



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

- Number of Hosts – The total number of hosts 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.

Function Summary

- FCoE
  - Number of FCoE Functions – The total number of FCoE functions discovered in the SAN.
  - Number of FC Targets – The total number of FC targets discovered in the SAN.
  - Number of VPorts – The total number of virtual FCoE ports discovered in the SAN.

8.2.4 Creating and Deleting FCoE Virtual Ports

This section describes how to create and delete virtual ports.

8.2.4.1 Creating Virtual Ports

Using the **Virtual Ports** tab (Figure 75), 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.

NOTES

- The OneCommand CNA 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:

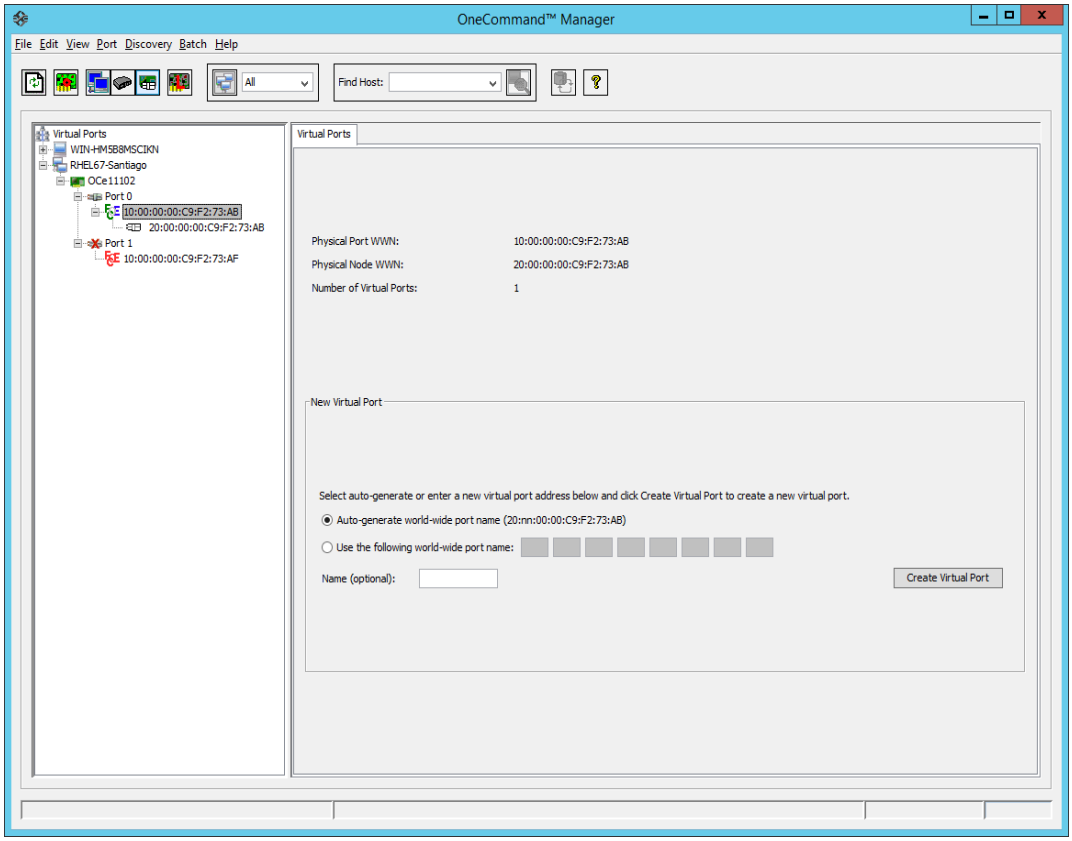
- For Windows: From the **Driver Parameters** tab, highlight enableNPIV, then select **Enable**. On the Storport Miniport system, the SLIMode driver parameter must also be set to 0 or 3.
- For Solaris: enable-npiv
- For Linux: lpfc_enable_npiv

See Section 8.2.16, Configuring FCoE Driver Parameters for more information on enabling driver parameters.

To create a virtual port:

1. Perform one of the following tasks:
   - From the **View** menu, select **Group Adapters by Virtual Ports**.
   - From the toolbar, click  **Group Adapters by Virtual Ports**.
2. From the discovery-tree, select the FCoE function on which you want to create a virtual port. The **Virtual Ports** tab appears (Figure 75).

Figure 75 Virtual Ports Tab



3. Perform one of the following tasks:
   - Check **Auto-generate world wide port name**. The OneCommand CNA Manager application creates the unique WWPN for the new virtual port based on the WWPN of the 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
     ```

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

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

5. Click **Create Virtual Port**. A 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 5) under the physical port where it was created and the Number of Virtual Ports field is updated.

   NOTE The OneCommand CNA 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 5).

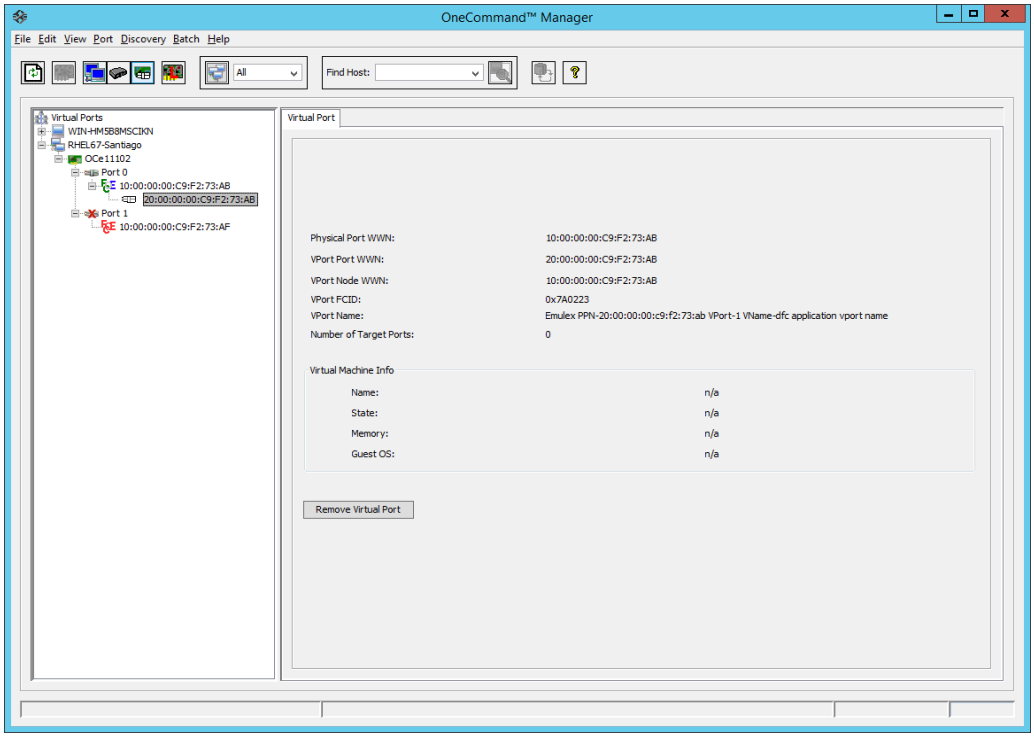
8.2.4.2 Deleting Virtual Ports

NOTE The OneCommand CNA 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:

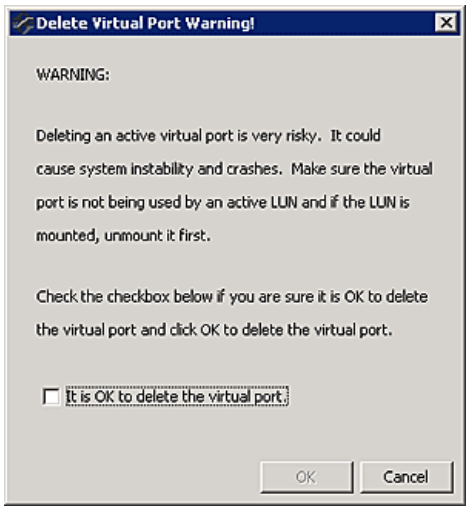
1. Perform one of the following tasks:
   - From the **View** menu, select **Group Adapters by Virtual Ports**.
   - From the toolbar, click **Group Adapters by Virtual Ports**.
2. From the discovery-tree, select the virtual port you want to delete. The Virtual Ports tab appears (Figure 76).

Figure 76 Virtual Port Tab



3. Click **Remove Virtual Port.** The Delete Virtual Port Warning dialog box appears (Figure 77).

Figure 77 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 5).
5. Click **OK**.

8.2.5 Viewing FCoE Fabric Information

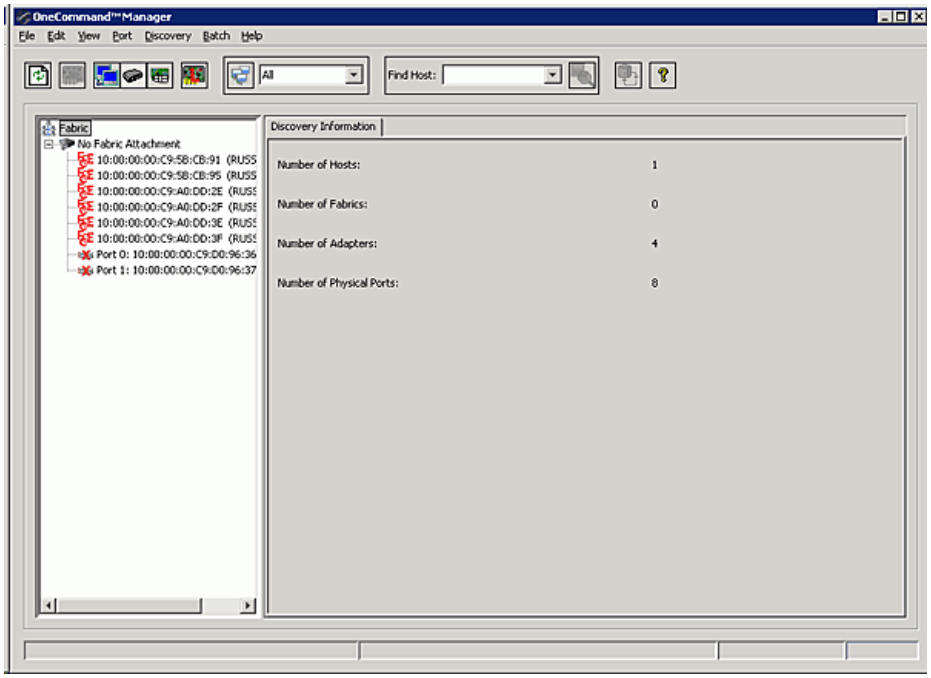
The **Discovery Information** tab (Figure 78) contains information about the selected fabric.

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

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

The **Discovery Information** tab shows information about the fabric (Figure 78).

Figure 78 Fabric Discovery Information



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

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

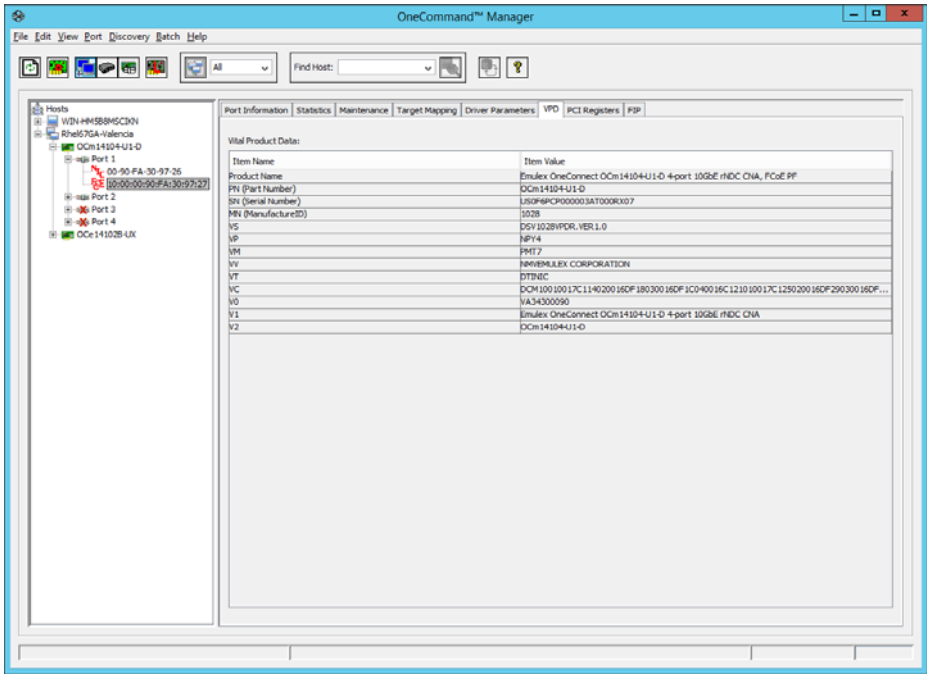
8.2.6 Viewing FCoE VPD Information

The **VPD** tab (Figure 79) displays vital product data (if available) for the selected FCoE 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 FCoE function whose VPD information you want to view.
3. Select the **VPD** tab.

Figure 79 FCoE VPD Tab



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

- Product Name – Product information about the selected 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 V*x* (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.

8.2.7 Viewing FCoE Maintenance Information

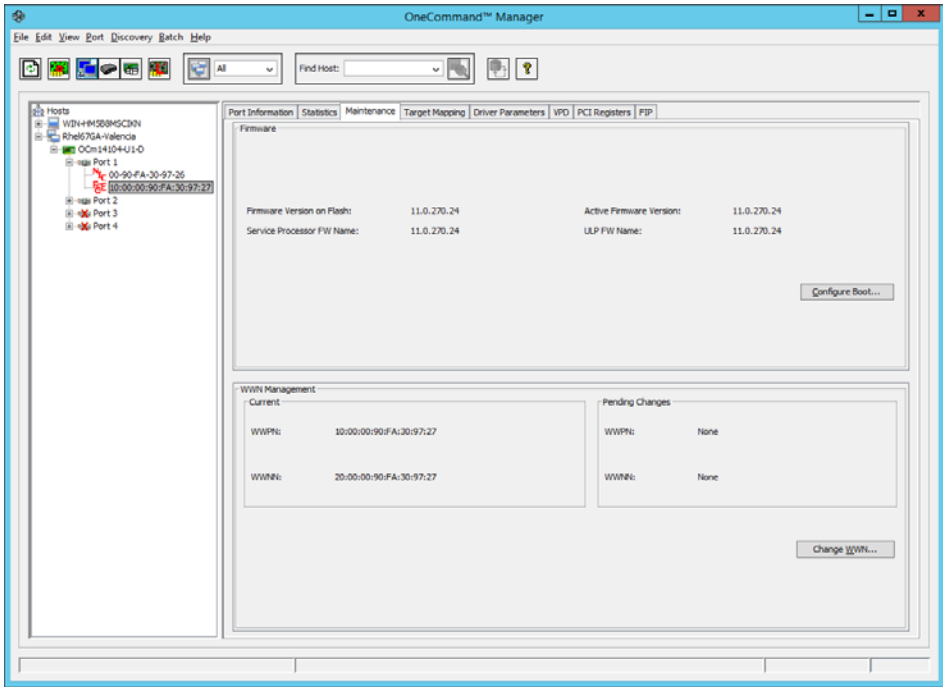
NOTE This option is not available in read-only mode.

Use the **Maintenance** tab (Figure 80) to view firmware information. You can also configure boot from SAN and change WWPN and WWNN information for the selected FCoE function.

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 80 FCoE Maintenance Tab



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

- 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 Section 8.2.13, Changing FCoE WWPN and WWNN 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 Section 8.2.13, Changing FCoE WWPN and WWNN for more information.

To configure boot from SAN, see Section 8.2.19, Configuring Boot from an FCoE SAN (not available on VMware ESXi servers being managed through the CIM interface).

To change the selected FCoE function's WWNN or WWPN, see Section 8.2.13, Changing FCoE WWPN and WWNN.

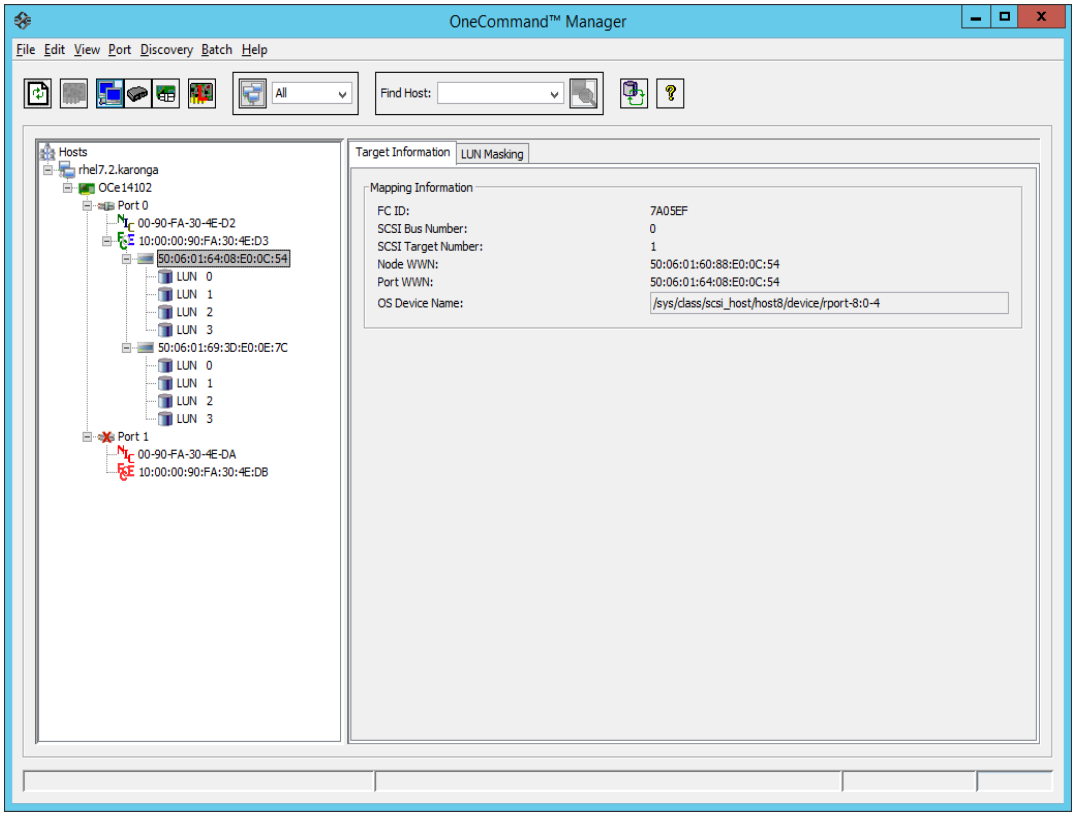
8.2.8 Viewing FCoE Target Information

When you select a target associated with an FCoE adapter from the discovery-tree (Figure 5), the **Target Information** tab (Figure 81) displays information associated with that target.

To view FCoE target information:

1. Select **Host**, **Fabric** or **Virtual Port** view.
2. In the discovery-tree (Figure 5), select the FCoE target whose information you want to view. The **Target Information** tab appears (Figure 81).

Figure 81 Target Information Tab



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

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

8.2.9 Viewing FCoE LUN Information

When you select a LUN associated with an FCoE target from the discovery-tree (Figure 82), 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:

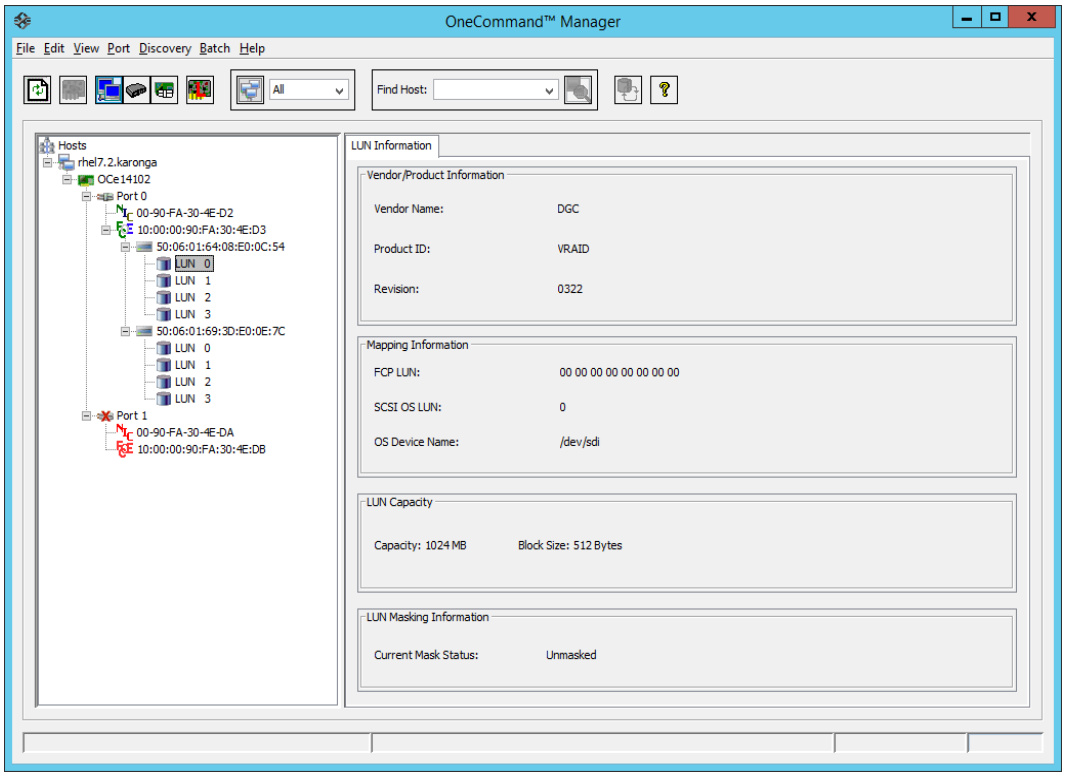
```
/usr/sbin/brcmfcoe/lun_scan all
```

- 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 82).

Figure 82 FCoE LUN Information Tab



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

- Vendor Product Information area:
  - Vendor Name – The name of the vendor of the LUN.

- Product ID – The vendor-specific ID for the LUN.
- Revision – The vendor-specific revision number for the LUN.

- Mapping Information area:
  - FCP LUN – The FC identifier used by the adapter to map to the SCSI OS LUN.
  - SCSI OS LUN – The SCSI identifier used by the operating system to map to the specific LUN.
  - OS Device Name – The name assigned by the operating system to the LUN.
- LUN Capacity area:

  NOTE LUN capacity information is 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 Section 8.2.12, Masking and Unmasking LUNs (Windows Only) for more information on LUN Masking.

8.2.10 Viewing FCoE Target Mapping

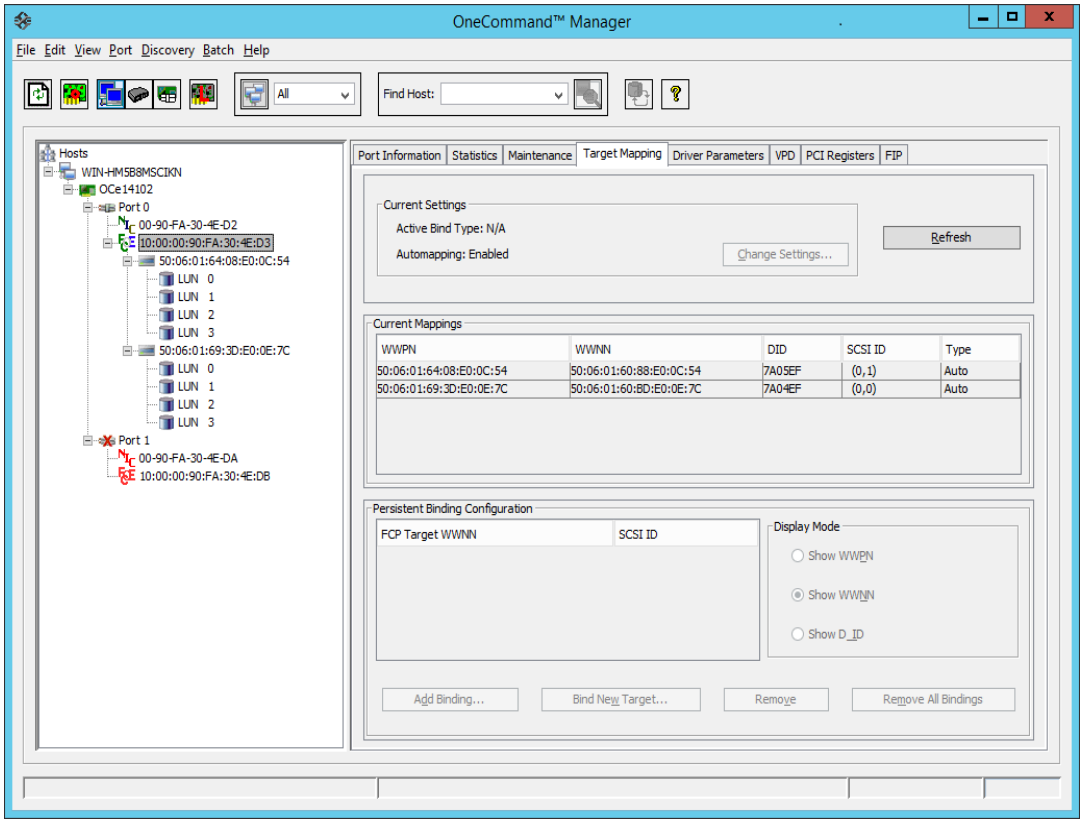
The **Target Mapping** tab (Figure 83) enables you to view current target mapping and to set up persistent binding.

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 FCoE function whose target mapping information you want to view.
3. Select the **Target Mapping** tab (Figure 83).

Figure 83 Target Mapping Tab



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

- 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 FCoE function.
- Persistent Binding Configuration area:
  - This table lists persistent binding information for the selected FCoE function (not available on VMware ESXi servers being managed through the CIM interface).
- Display Mode area:

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

  - Select the method by which you want to display information in the Persistent Binding Configuration table.

For information on changing settings, see Section 8.2.11.1, Changing Automapping Settings.

For information on adding a binding, see Section 8.2.11.2, Adding a Persistent Binding.

For information on binding a new target, see Section 8.2.11.3, Binding a Target that Does Not Appear in the Persistent Binding Table.

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

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

8.2.11 Using Automapping and Persistent Binding (Windows Only)

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

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

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

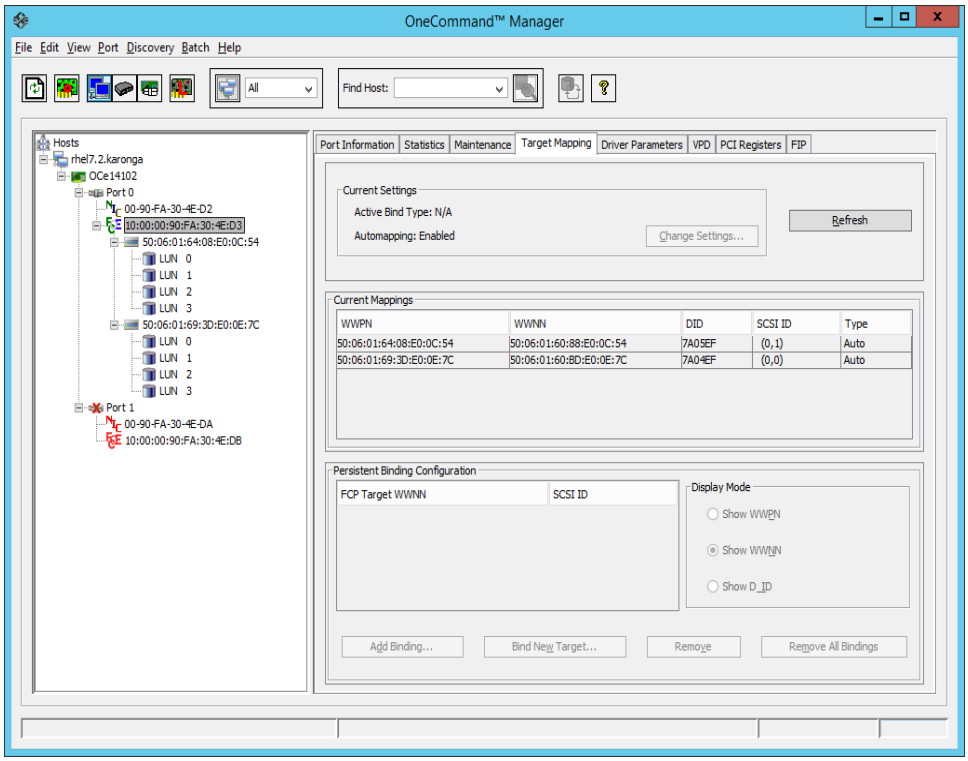
- The SCSI ID (target/bus combination) specified in the binding request must not be mapped to another target. For example, the SCSI ID must not already appear in the Current Mappings table under SCSI ID. If the SCSI ID is already in use, then the binding cannot be made dynamic, and a reboot is required.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

8.2.11.1 Changing Automapping Settings

To change automapping settings:

1. Select **Host** or **Fabric** view.
2. In the discovery-tree, select the FCoE function you want to set up with persistent binding.
3. Select the **Target Mapping** tab (Figure 84). All targets are displayed.

Figure 84 Target Mapping Tab



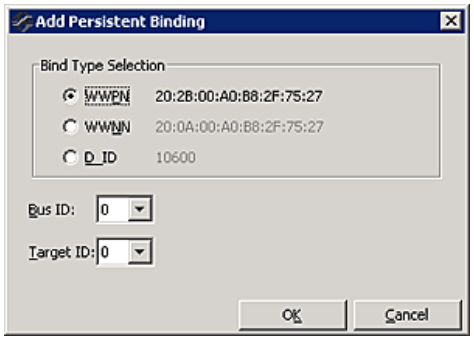
4. Target mappings are displayed by WWPN, WWNN, or D_ID. *PB* indicates mapping from persistent binding, while *Auto* indicates an automapped target. In the Display Mode section, choose the display mode you want to use.
5. If you want to make changes, click **Change Settings**. The Mapped Target Settings 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.

8.2.11.2 Adding a Persistent Binding

To add a persistent binding:

1. Select **Host** or **Fabric** view.
2. In the discovery-tree (Figure 5), select the FCoE function you want to set up with persistent binding.
3. Select the **Target Mapping** tab (Figure 84). 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 85) is displayed.

Figure 85 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.

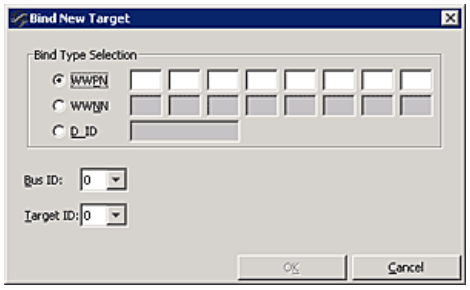
8.2.11.3 Binding a Target that Does Not Appear in the Persistent Binding Table

NOTE It is possible to specify a SCSI bus and target that have already been used on behalf of a different FC target. Attempting to bind a target already in the Persistent Binding table on the Target Mapping tab results in 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 5), select the FCoE function you want to set up with persistent binding.
3. Select the **Target Mapping** tab (Figure 83). All targets are displayed.
4. Click **Bind New Target**. The Bind New Target dialog box is displayed (Figure 86).

Figure 86 Bind New Target Dialog Box



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

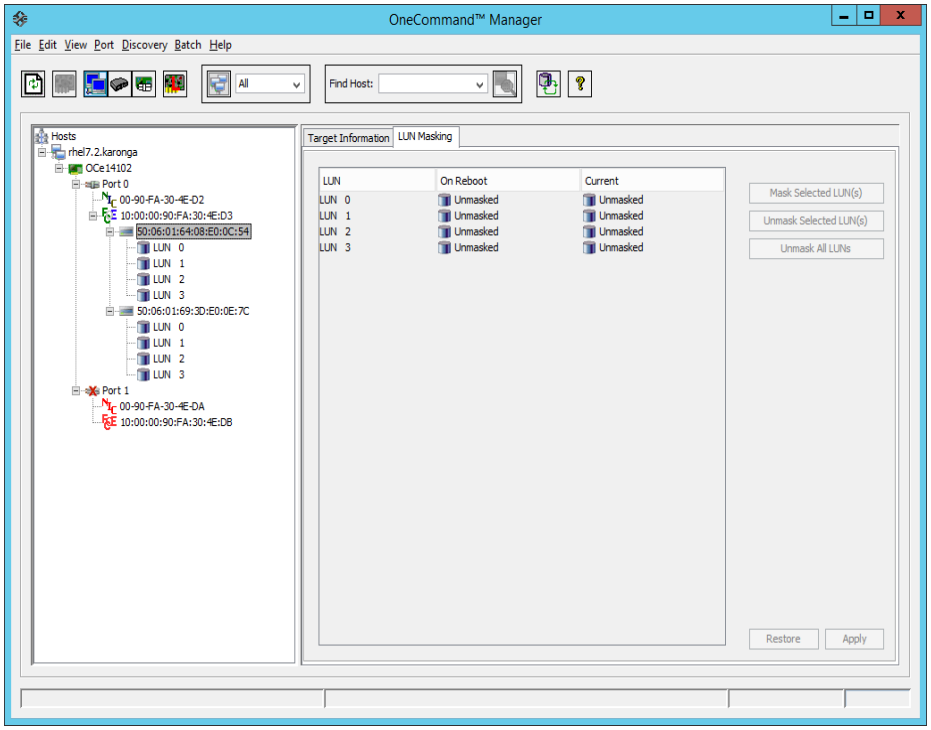
NOTE A target does not appear on the target list if automapping is disabled and the target is not already persistently bound.

8.2.12 Masking and Unmasking LUNs (Windows Only)

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 87) to mask or unmask LUNs at the host level.

NOTE The LUN Masking tab (Figure 87) is not shown in Virtual Port view because LUN masking is not available for virtual ports.

Figure 87 LUN Masking Tab



8.2.12.1 LUN Masking Conventions and Guidelines

LUN icons in the discovery-tree (Figure 5) 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).

The following LUN Masking information is displayed:

- LUN – The FC LUN number.
- On Reboot – The On Reboot column shows the mask configuration currently saved to the configuration file on disk (Solaris) or to the Registry (Windows). 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 5).

To change the mask status of a LUN:

1. Select **Host** view.
2. From the discovery-tree (Figure 5), select the target whose LUN masking state you want to change. A set of LUNs appears below the selected target.
3. Select the **LUN Masking** tab (Figure 87). This tab contains a list of the same set of LUNs that appear below the FCoE target in the discovery-tree (Figure 5).
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. Make 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.

8.2.13 Changing FCoE WWPN and WWNN

The **Maintenance** tab (Figure 88) enables you to change the WWPN and the WWNN of a selected FCoE function. For example, you can use an installed adapter as a standby in case another installed adapter fails. By changing the standby adapter's WWPN or WWNN, it can assume the identity and configuration (for example, driver parameters, persistent binding settings, and so on) of the failed adapter.

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 or a power outage.
- Volatile WWN – A temporary value that is saved in volatile memory on the flash. If volatile WWNs are set, they are used instead of the non-Volatile WWNs.

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

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

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

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

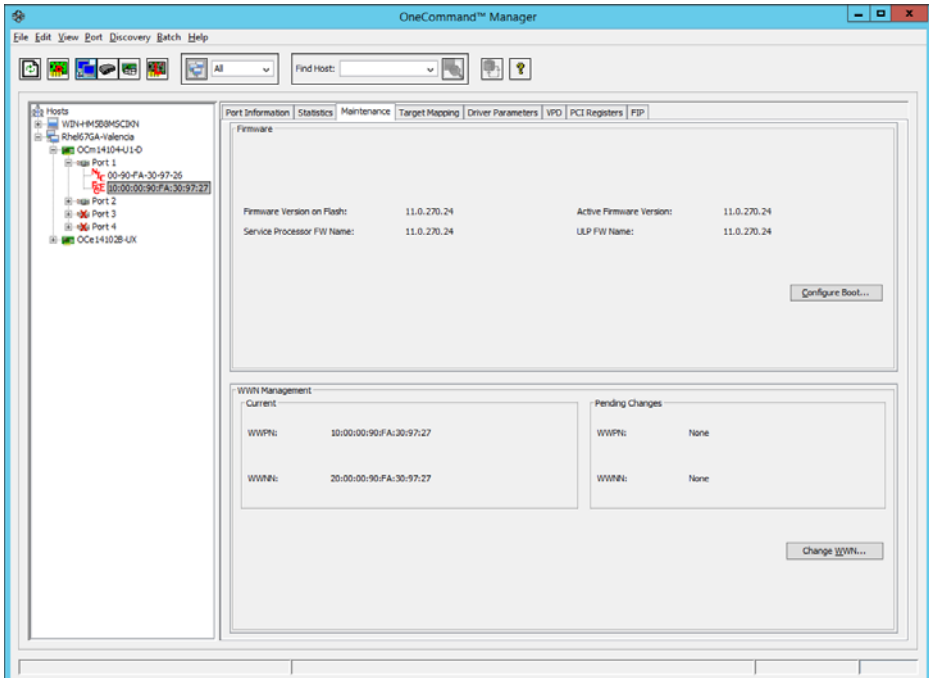
For example, in an environment with CLP or BOFM:

The OneCommand CNA Manager application modifies the WWNN or WWPN. The OneCommand CNA 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 by the command from the BOFM.

To change an FCoE function's WWPN or WWNN:

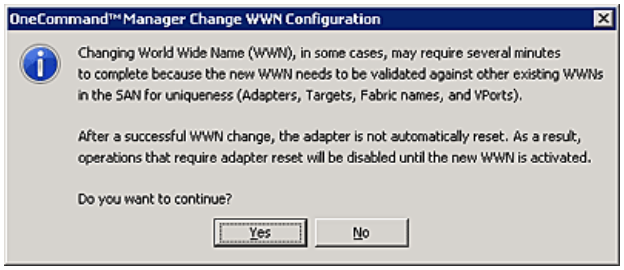
1. Perform one of the following tasks:
   - From the **View** menu, click **Group Adapters by Host Name**.
   - From the toolbar, click  **Group Adapters by Host Name**.
   - From the **Host Grouping** menu, select **Group Adapter by Fabric Names**.
2. In the discovery-tree, select the FCoE function that you want to change.
3. Select the **Maintenance** tab (Figure 88).

Figure 88 Maintenance Tab



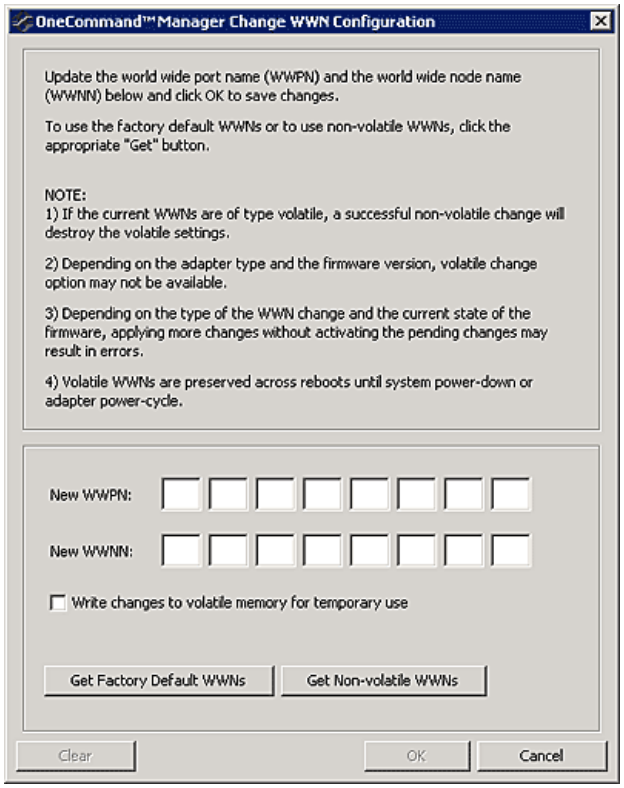
4. Click **Change WWN**. The warning in Figure 89 appears:

Figure 89 Warning About Changing WWN



5. Click **Yes**. The Change World Wide Name Configuration dialog box appears (Figure 90).

Figure 90 Change World Wide Name Configuration Dialog Box



6. Perform one of the following tasks:
   - Enter a new WWPN and WWNN.
   - Click **Get Factory Default WWNs** to load the settings that were assigned to the 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 cleared, the New WWPN and New WWNN settings are saved as Non-Volatile WWNs.

   NOTE If the adapter or firmware does not support Volatile WWNs, the Write changes to volatile memory for temporary use 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.

8.2.14 Changing FCoE Names

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

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

For example, you may want to identify a particular 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 FCoE function's displayed name from the default WWPN, the change occurs in the discovery-tree (Figure 5) only. The function'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 CNA Manager application on another host do not see your name changes.

To change the name of an FCoE function:

1. From the discovery-tree (Figure 5), select the FCoE port that you want to change.
2. Perform one of the following tasks:
   - Select **Edit Name** from the **Port** menu.
   - From the discovery-tree, right-click the port that you want to change and select **Change Name**.
3. Edit the name in the discovery-tree.

To use the FCoE function's default name:

1. From the discovery-tree (Figure 5), select the FCoE function that you want to change.
   - Perform one of the following tasks:
   - Select **Use Default Name** from the **Port** menu.
   - From the discovery-tree, right-click the port that you want to change and select **Restore Default Name**.

8.2.15 Resetting FCoE Functions

You can reset remote and local FCoE functions. This option is 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.

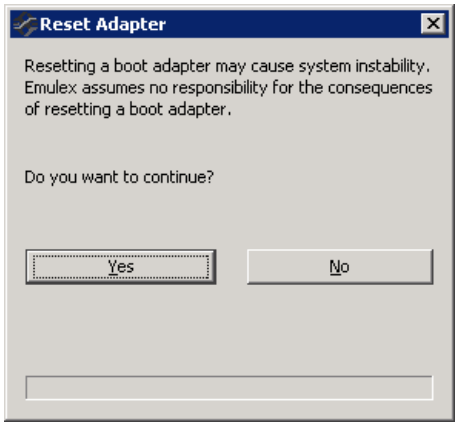
NOTE A reset is necessary only to activate updated driver parameters that require a reset or FIP settings. It does not actually perform an adapter function level reset.

To reset the FCoE function:

1. In the discovery-tree (Figure 5), select the FCoE function you want to reset.
2. Perform one of the following tasks:
   - From the **Port** menu, click **Reset Port**.
   - From the toolbar, click  **Reset**.

   The Reset Adapter warning appears (Figure 91).

Figure 91 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.

8.2.16 Configuring FCoE Driver Parameters

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

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

NOTES

- This section only applies to 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. Refer to the appropriate driver manual for instructions.

8.2.16.1 Activation Requirements

A parameter has one of the following activation requirements:

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

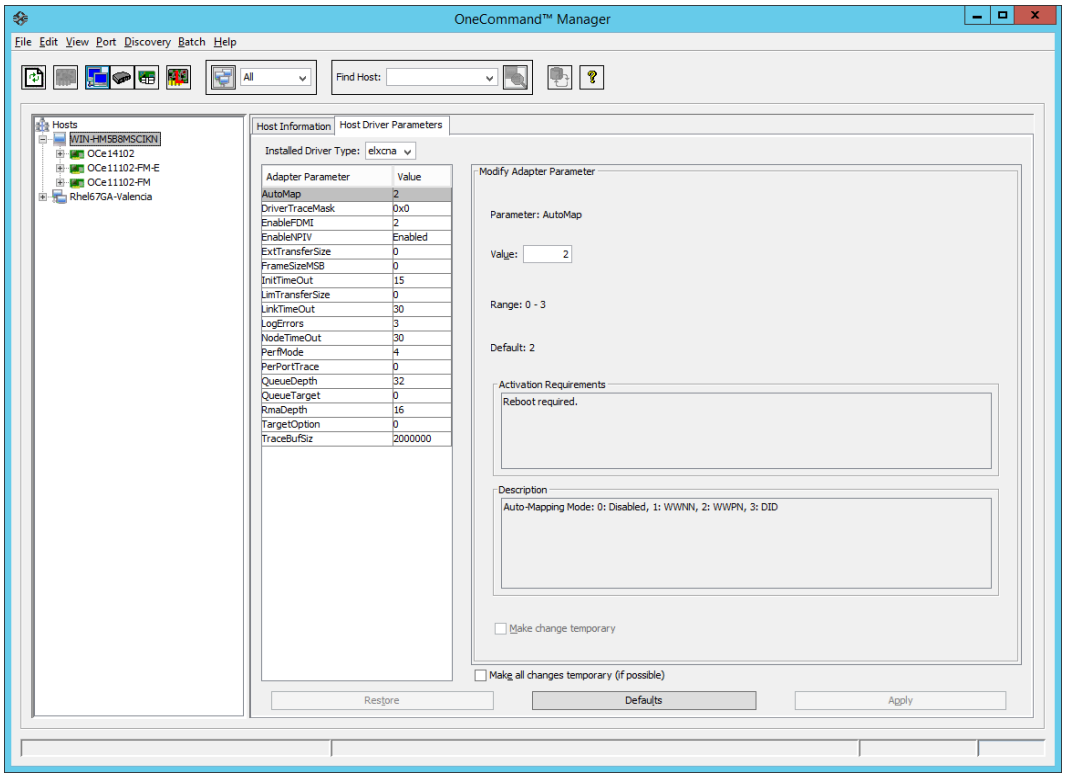
8.2.16.2 The Host Driver Parameters Tab

The **Host Driver Parameters** tab (Figure 92) enables you to view and edit the adapter driver parameter settings contained in a specific host. The host driver parameters are global values, and apply to all adapters in that host unless they are overridden by parameters assigned to a specific adapter using the adapter **Driver Parameters** tab. For each parameter, the tab shows the current value, the range of acceptable values, the default value, and whether the parameter is dynamic. A dynamic parameter allows the change to take effect without resetting the adapter or rebooting the system.

For information on changing parameters for a single adapter, see Section 8.2.16.3, Setting Driver Parameters. For information on changing parameters for the host, see Section 8.2.16.9, Setting Driver Parameters for All FCoE Functions in a Host.

NOTE If there are no discovered FCoE functions, the entire **Host Driver Parameters** tab (Figure 92) is dimmed. This occurs because there are no drivers to which the host driver parameters apply.

Figure 92 Host Driver Parameters Tab



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

- 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.

8.2.16.3 Setting Driver Parameters

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

For example, if you select a host in the discovery-tree (Figure 5), you can globally change the parameters for all FCoE functions in that host. If you select an FCoE function in the discovery-tree, you can change parameters for only that 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 FCoE function to other FCoE functions in the system using the **Driver Parameters** tab, thereby simplifying multiple adapter configuration. See Section 8.2.17, Creating a Batch Mode Driver Parameters File for more information.

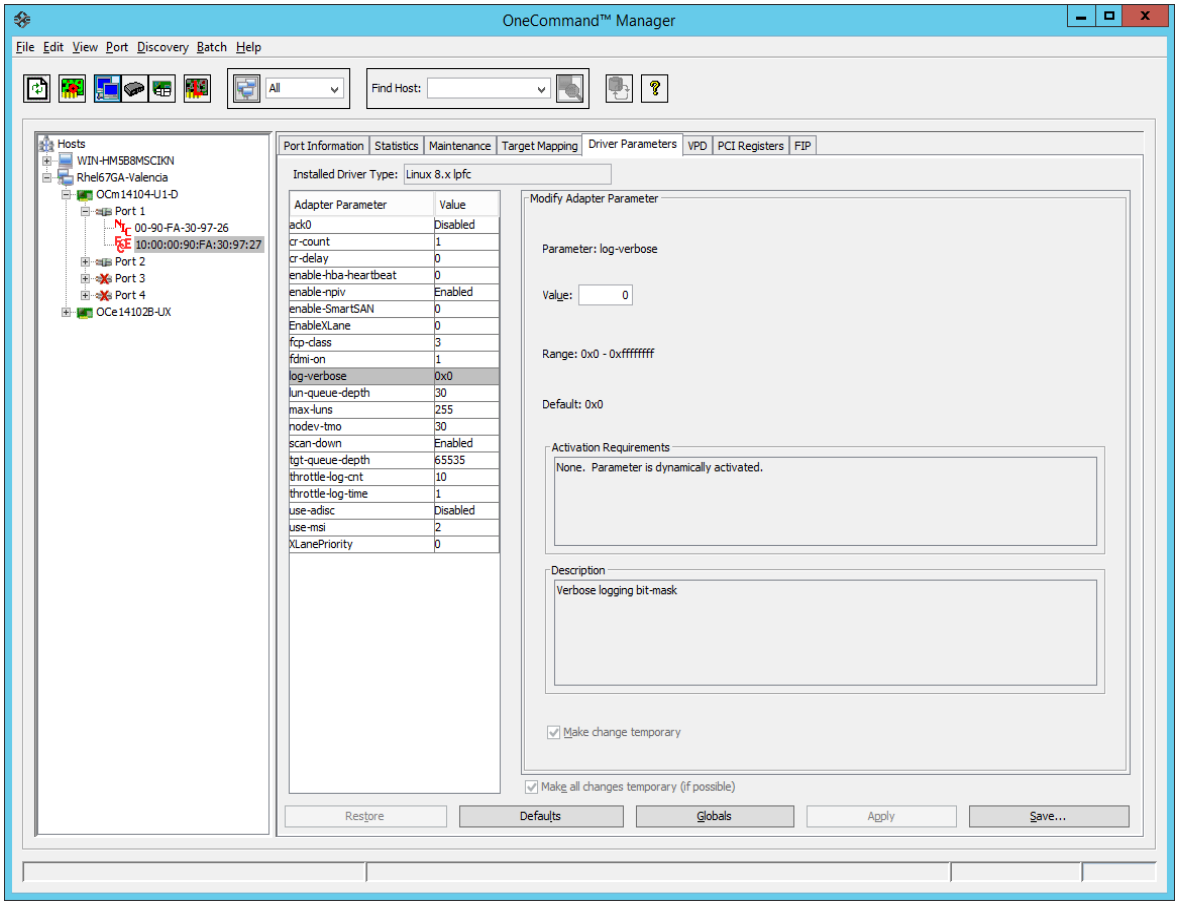
NOTE The Linux 2.6 kernel only supports setting some of the driver parameters for individual FCoE functions. Some driver parameters must be applied to all FCoE functions contained in the host. See the *Emulex Driver for Linux for OneConnect Adapters User Guide* for more information.

8.2.16.4 Setting Driver Parameters for a Single FCoE Function

To change the driver parameters for a single FCoE function:

1. Select **Host** or **Fabric** view.
2. In the discovery-tree (Figure 5), select the FCoE function that you want to change.
3. Select the **Driver Parameters** tab (Figure 93). The parameter values for the selected FCoE functions are displayed.

Figure 93 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**.

8.2.16.5 Restoring All Parameters to Their Earlier Values

If you changed parameters, but did not click **Apply** in the **Driver Parameters** tab (Figure 93) and you want to restore the parameters to their last saved values, click **Restore**.

8.2.16.6 Resetting All Default Values

To reset all parameter values to their default (factory) values, click **Defaults** in the **Driver Parameters** tab (Figure 93).

8.2.16.7 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 5), select the adapter port.
3. Select the **Driver Parameters** tab (Figure 93).
4. Click **Globals**. All parameter values are now the same as the global, or host, values.
5. To apply the global values, click **Apply**.

8.2.16.8 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 93). To apply your configuration changes, click **Apply**.

NOTE OneCommand CNA Manager application Web Launch Interface driver parameters files are saved on the host that the browser was started 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\Broadcom\Util\Broadcom Repository`
- In Linux: `/usr/sbin/brcmocmanager/RMRepository`
- In VMware ESXi: `/tmp/RMRepository`
- In Solaris: `/opt/brcmocm/RMRepository`

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

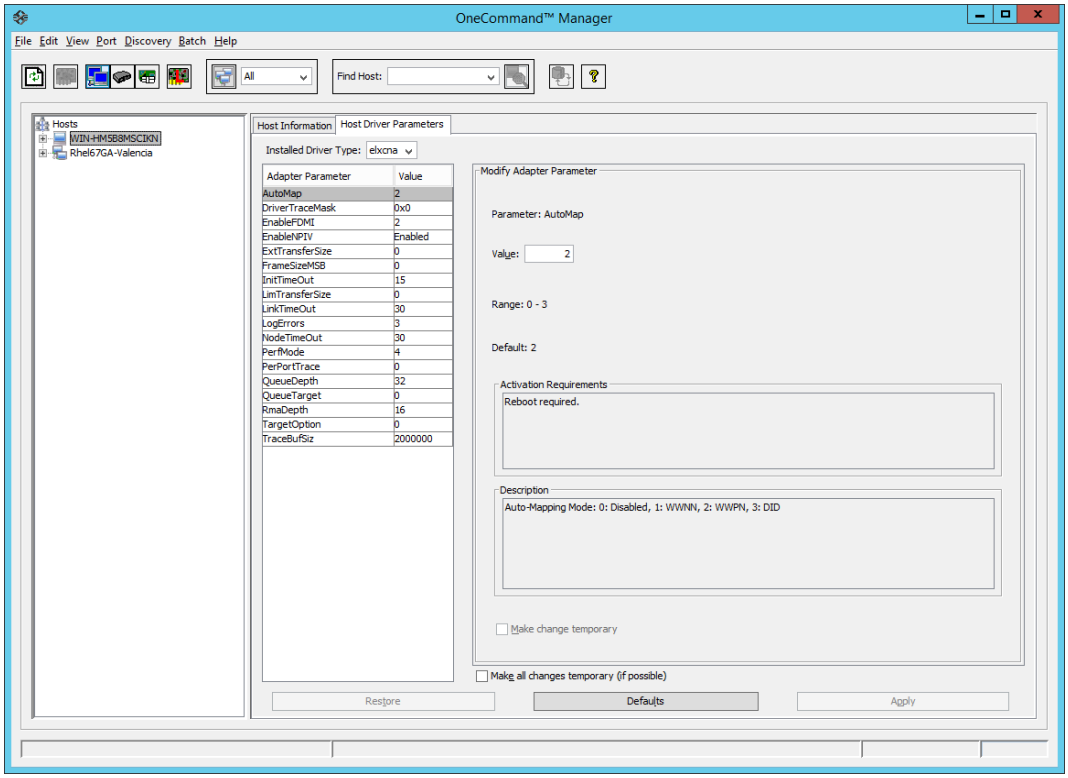
NOTE Host driver parameters and persistent binding settings cannot be saved.

8.2.16.9 Setting Driver Parameters for All FCoE Functions in a Host

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

1. Perform one of the following tasks:
   - 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 94). 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 94 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**.

8.2.16.10 Changing Non-Dynamic Parameter Values

To change non-dynamic parameter values for Linux:

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

   ```
   ./stop_brcmocmanager
   ```

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

   ```
   modprobe -r brcmfcoe
   ```

4. Reload the driver. Type the following command:

   ```
   modprobe brcmfcoe
   ```

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

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

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

   ```
   ./start_brcmocmanager
   ```

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

If DHCHAP authentication is currently employed on Emulex adapters on this machine, you must type the following command to restart the authentication daemon:

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

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

```
./start_weblaunch
```

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

8.2.17 Creating a Batch Mode Driver Parameters File

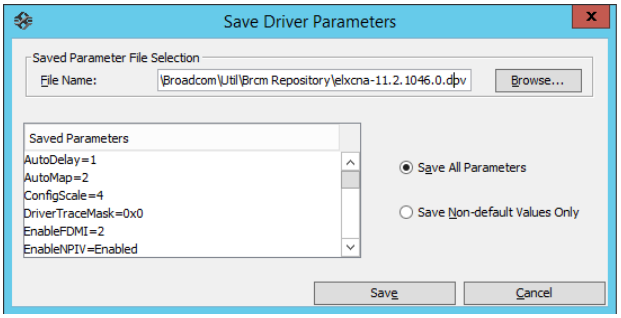
NOTE This option is not available in read-only mode.

You can apply driver parameters for one FCoE function to other 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 CNA Manager application enables you to assign the `.dpv` file parameters to multiple adapters in the system.

To create the `.dpv` file:

1. Select **Host** or **Fabric** view.
2. Select the FCoE function whose parameters you want to apply to other FCoE functions from the discovery-tree (Figure 5).
3. Select the **Driver Parameters** tab (Figure 93).
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 95). You can save the file to a different directory or change its name.

Figure 95 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.

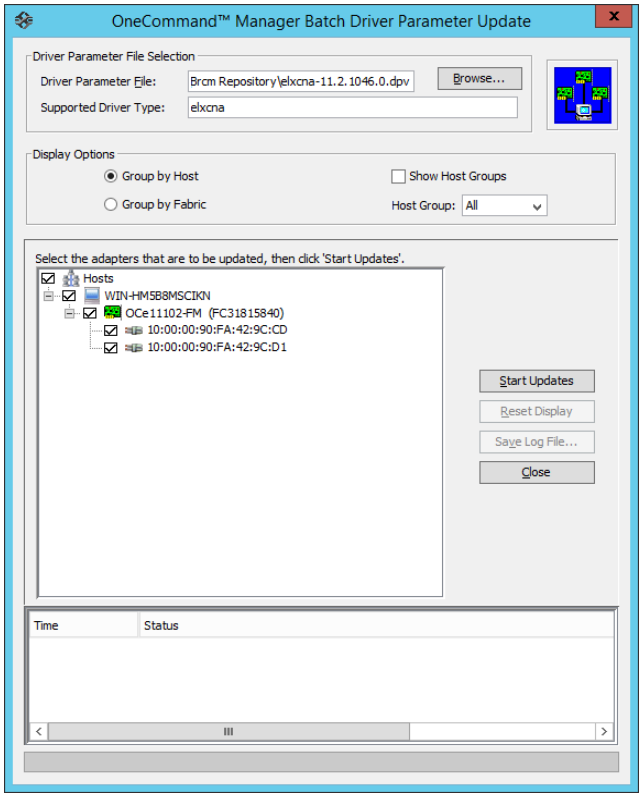
8. Click **Save**.

8.2.17.1 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 5) elements at this time.)
2. When the Batch Driver Parameter Update dialog box appears, click **Browse**.

Figure 96 Batch Driver Parameters Update Dialog Box



3. The Driver Parameter File Selection dialog box appears (Figure 96). Select the file you want to use and click **OK**. A dialog box appears notifying you that the OneCommand CNA Manager application is searching for compatible adapters.

   After compatible 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 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 FCoE functions entries are used to select or clear an entry. Checking an adapter selects or removes all 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 CNA Manager application Batch Driver Parameter Update dialog box (Figure 96) shows the current status of the update. When the update completes, a final summary shows the number of FCoE functions that were successfully processed, and the number of FCoE functions for which one or more parameter updates failed.
5. You can click **Save Log File** to save a report of the update.

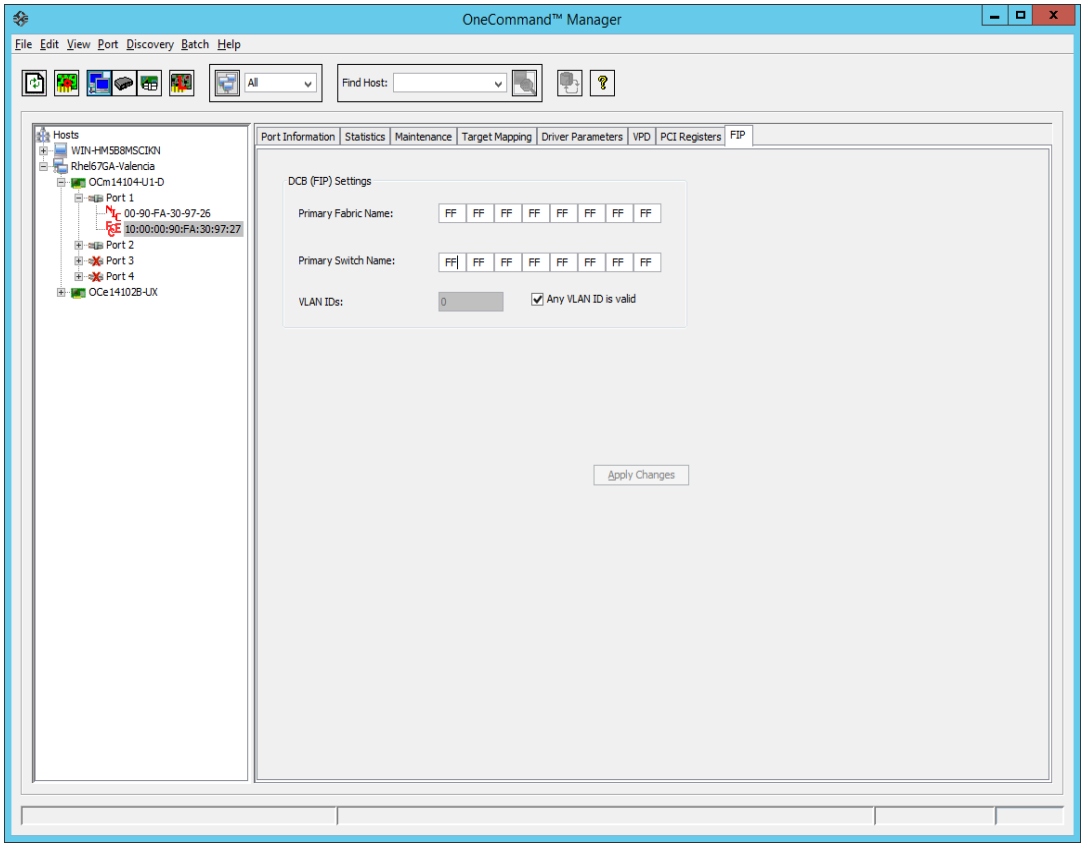
8.2.18 Configuring FCoE Initialization Protocol (FIP)

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 97).
3. Set the parameters you want and click **Apply Changes**.

Figure 97 FIP Tab for FCoE Adapters



The following FIP tab fields are displayed:

- Primary Fabric Name – Indicates the fabric's WWN to connect to. If the Primary Fabric Name is wild (all 0xFFs), then connection to any fabric name is allowed.
- Primary Switch Name – Indicates the switch's WWN to connect to. If the Primary Switch Name is wild (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 to 4095 and supports wild card values if **Any** is checked.
- Any VLAN ID is valid checkbox – If checked, the VLAN ID field of the FCoE forwarder can be any valid value.

8.2.19 Configuring Boot from an FCoE SAN

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

8.2.19.1 Boot Types

Using the Maintenance tab, you can enable, disable, or configure boot from SAN for x86 BootBIOS and EFIBoot.

- x86 BootBIOS works with the existing BIOS on x64 and x86 systems.
- EFIBoot works with x64-based systems and provides 64-bit system boot capability through the use of the EFI Shell.

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

8.2.19.2 Boot Device Parameters

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

- The OneCommand CNA Manager application runs on a running operating system, so you must boot the host to configure boot from SAN with the OneCommand CNA Manager application.
- You must work from a running host that supports the OneCommand CNA Manager application. Often, this host has booted from a direct-attached drive. With the OneCommand CNA 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 CNA Manager application cannot detect this. Without knowing this, you could configure a boot type but not be able to boot from it because the adapter lacks the correct boot code.
- You must know what boot code type the system supports; the OneCommand CNA 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 CNA Manager application that does not support boot from SAN, the Configure Boot button does not appear.

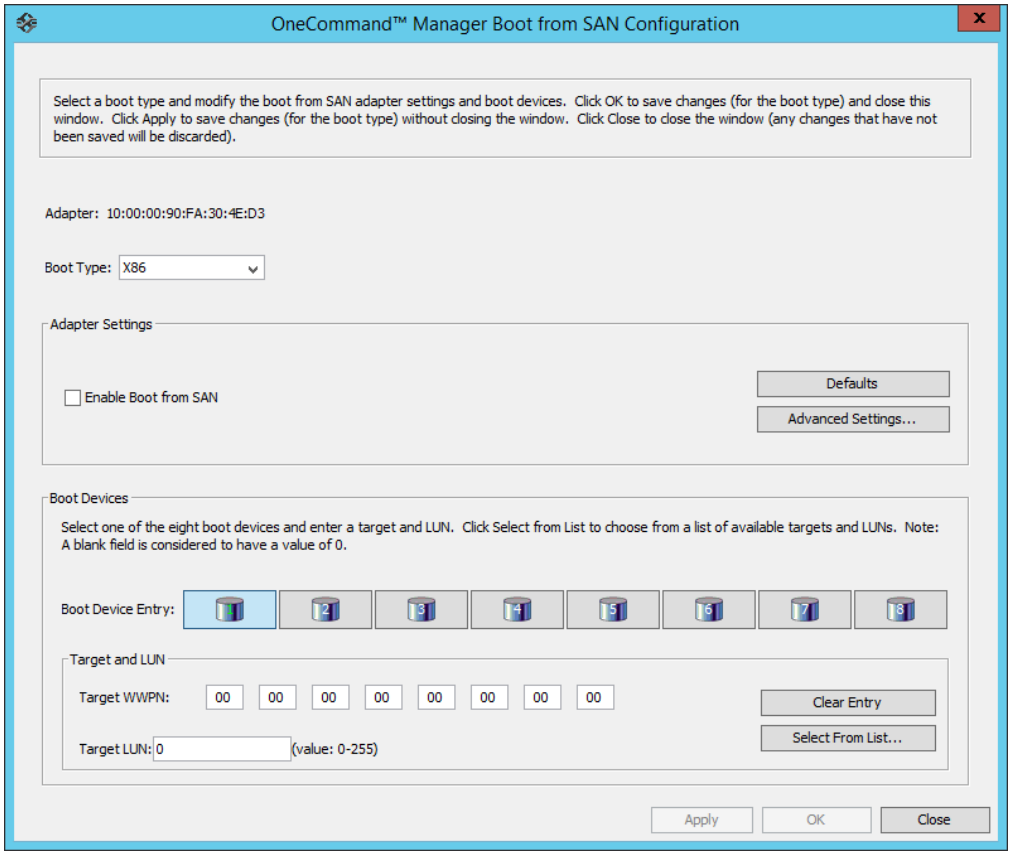
  NOTE You can configure boot from SAN before boot by using the Emulex Boot BIOS setup command line interface that runs during system startup. See the *Boot for NIC, iSCSI, and FCoE Protocols User Guide* for details.

- One of the following 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 5), click the FCoE adapter function on which you want to enable boot from SAN.
3. Select the **Maintenance** tab (Figure 80) and click **Configure Boot**. The Boot from SAN Configuration dialog box appears (Figure 98).

Figure 98 Boot from SAN Configuration Dialog Box



The Boot from SAN Configuration dialog box varies for each boot type. Figure 98 depicts the boot from SAN configuration for the x86 type boot.

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

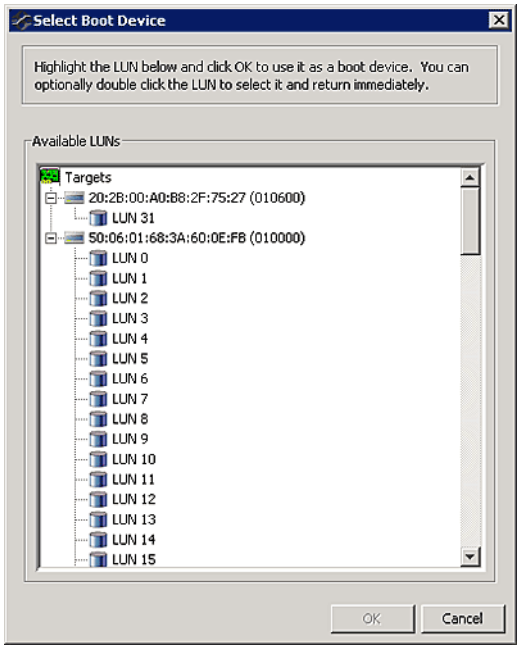
   NOTE If you modified the settings for the current boot type and then change to a new boot type, a message appears telling you to save the current settings before changing to the new boot type.

6. Check **Enable Boot from SAN**.
7. If you want, click **Advanced Settings** to configure autoscan, spinup delay and so on. See Section 8.2.20, Configuring Advanced Settings (Boot from SAN) for more information.
8. Select one or more boot devices.
9. Perform one of the following tasks on the Boot from SAN Configuration dialog box (Figure 98):
   - Select **Target World Wide 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, type in an 8-byte LUN (hexadecimal) 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 zeros.
   - 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 99). 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 99.) The OneCommand CNA Manager application attempts to update the boot parameters. If successful, a window appears with a confirmation message. Click **OK** on this confirmation window.

Figure 99 Select Boot Device Window (for x86 or EFIBoot)



10. On the Boot from SAN Configuration dialog box (Figure 98), 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.

8.2.20 Configuring Advanced Settings (Boot from SAN)

The OneCommand CNA Manager application provides advanced settings for each boot type. From the Boot from SAN Configuration dialog box (Figure 98), 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.

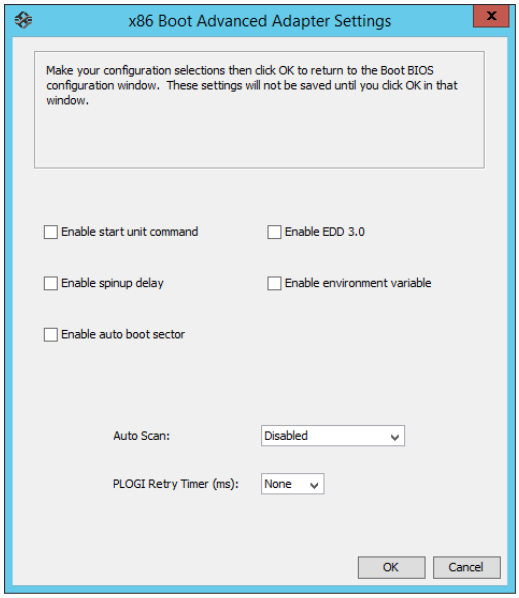
If you make changes you must click **OK** to save the changes and close the dialog box. You can click **Cancel** and close the dialog box without saving the changes.

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.

8.2.20.1 x86 Boot Advanced Adapter Settings Dialog Box

Use the x86 Boot Advanced Adapter Settings dialog box (Figure 100) to configure **advanced settings** for the selected x86 adapter. All checkboxes are cleared (off) by default. All changes require a reboot to activate.

Figure 100 x86 Boot Advanced Adapter Settings Dialog Box



The following x86 Boot Advanced Adapter Settings fields are displayed:

- 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 stop responding 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 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. 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

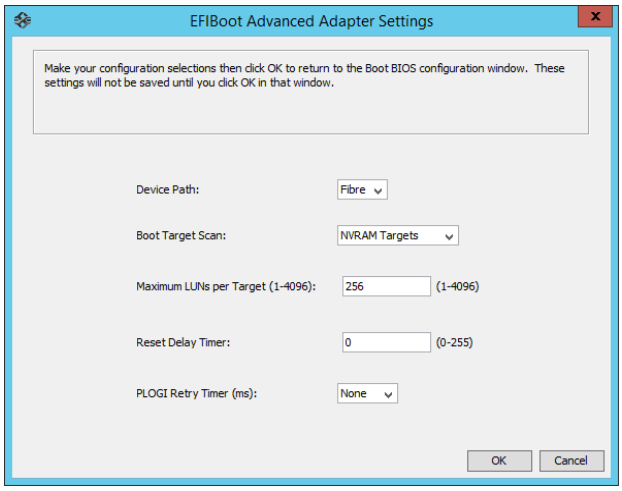
- First non-LUN 0 Device

- 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 back online, the PLOGI retry interval scans the loop to discover this device. This default setting is None (0 ms). Set the Retry Timer to one of the following values:
  - None (default)
  - 50 ms
  - 100 ms
  - 200 ms

8.2.20.2 EFIBoot Advanced Adapter Settings Dialog Box

Use the EFIBoot Advanced Adapter Settings dialog box (Figure 101) to configure the advanced settings for the selected EFIBoot adapter.

Figure 101 EFIBoot Advanced Adapter Settings Dialog Box



The following EFIBoot Advanced Adapter Settings fields are displayed:

- 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 by using 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 FCoE 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

8.3 iSCSI

This section describes configuring iSCSI functions.

8.3.1 Viewing iSCSI Port Information

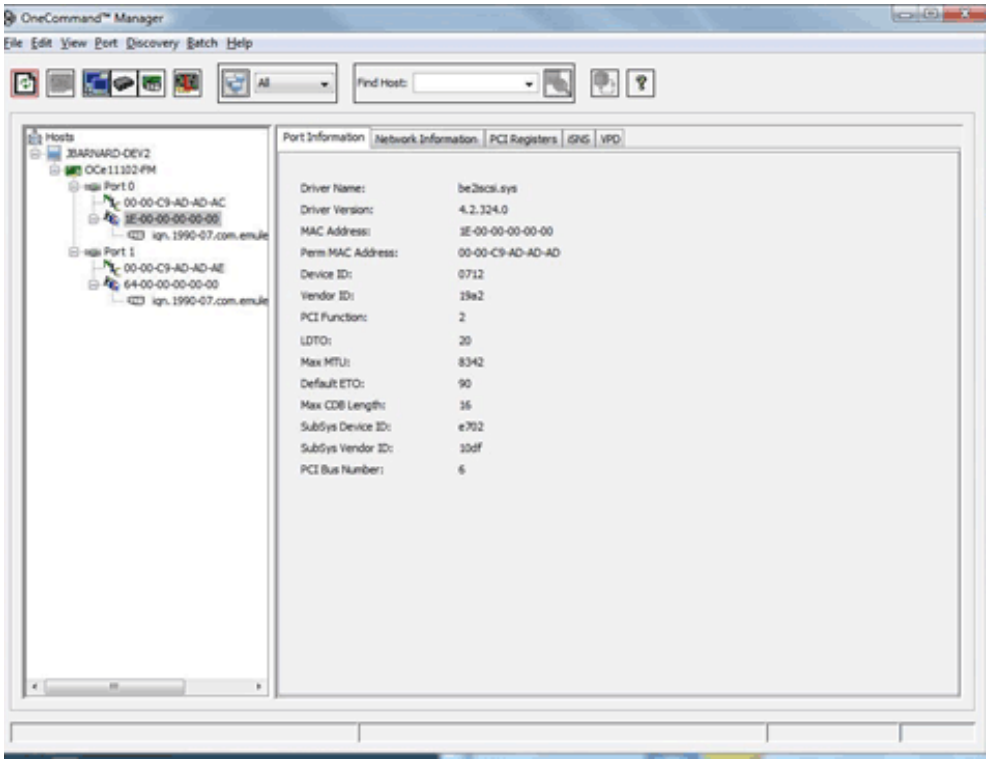
When you select an iSCSI function from the discovery-tree (Figure 5), the iSCSI **Port Information** tab (Figure 102) contains general attributes associated with the selected iSCSI function.

NOTE iSCSI functions are not discovered in the OneCommand CNA 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 5).
3. Select the **iSCSI Port Information** tab (Figure 102).

Figure 102 iSCSI Port Information Tab



8.3.1.1 The following iSCSI Port Information tab fields are displayed:

- 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.

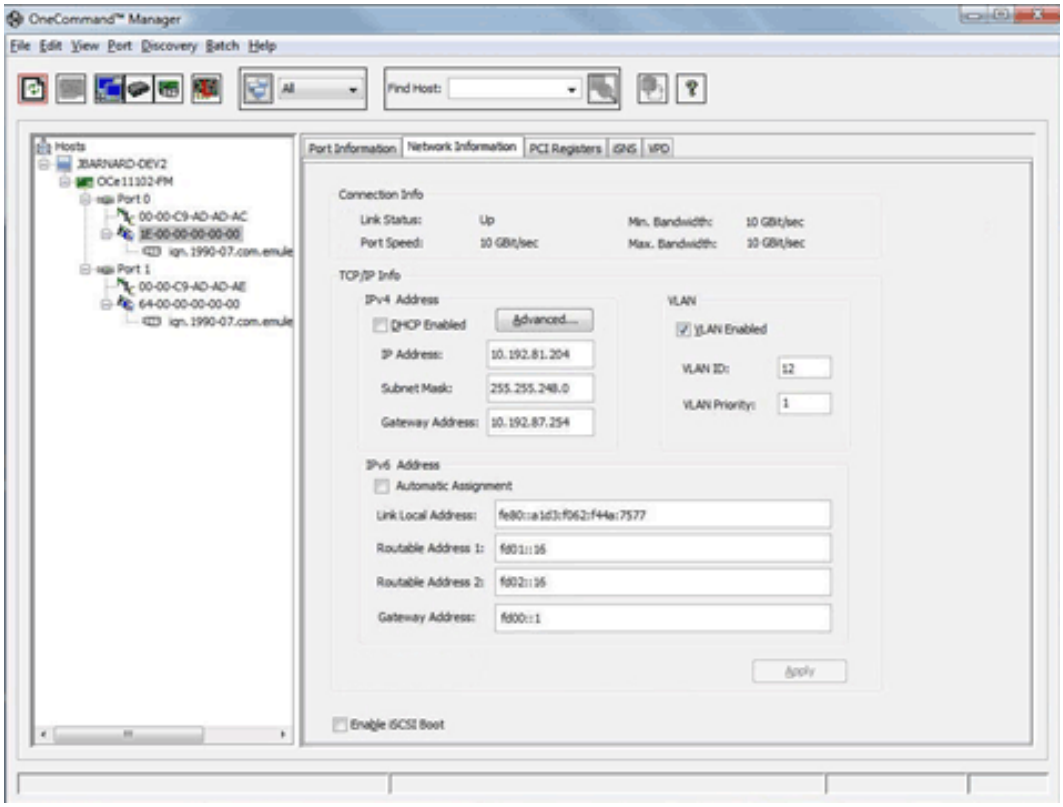
8.3.2 Viewing iSCSI Network Information

The **Network Information** tab (Figure 103) 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 103).

Figure 103 iSCSI Network Information Tab



The following iSCSI **Network Information** tab fields are displayed:

- 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.

    NOTE When an OCe11000-series adapter is configured in multichannel disabled mode, the minimum bandwidth for iSCSI functions will always be shown as zero.
  - 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.

  NOTE Checking DHCP Enabled disables the IP Address, Subnet Mask, and Gateway Address fields.
- 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.

  NOTE 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 fails.
- 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 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 clear this box to enable or disable iSCSI boot on the iSCSI function.

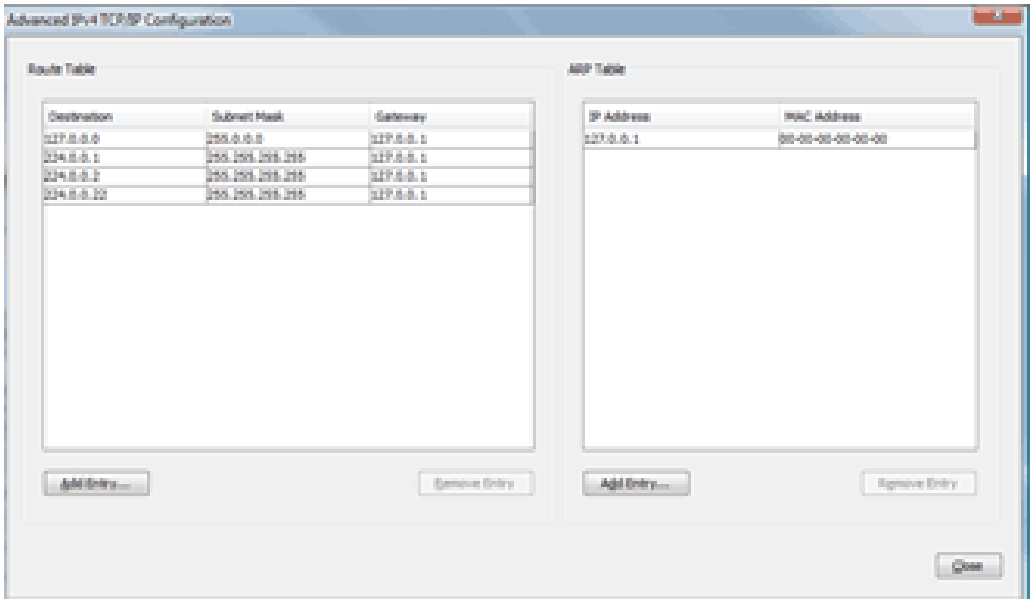
To modify TCP/IP configurations for iSCSI ports:

1. From the discovery-tree (Figure 5), select the iSCSI function whose configuration you want to modify.
2. Select the **Network Information** tab (Figure 103).
3. Make your selections.
4. Click **Apply**.

8.3.2.1 Advanced TCP/IP Configuration

The Advanced TCP/IP Configuration dialog box (Figure 104) enables you to add and remove Route and ARP table entries (applicable to IPv4 only) for the selected iSCSI function.

Figure 104 Advanced TCP/IP Configuration Dialog Box



To add table entries:

1. From the discovery-tree (Figure 5), select the iSCSI function whose configuration you want to modify.
2. Select the **Network Information** tab (Figure 103).
3. Click **Advanced**. The Advanced TCP/IP Configuration dialog box appears (Figure 104).
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.

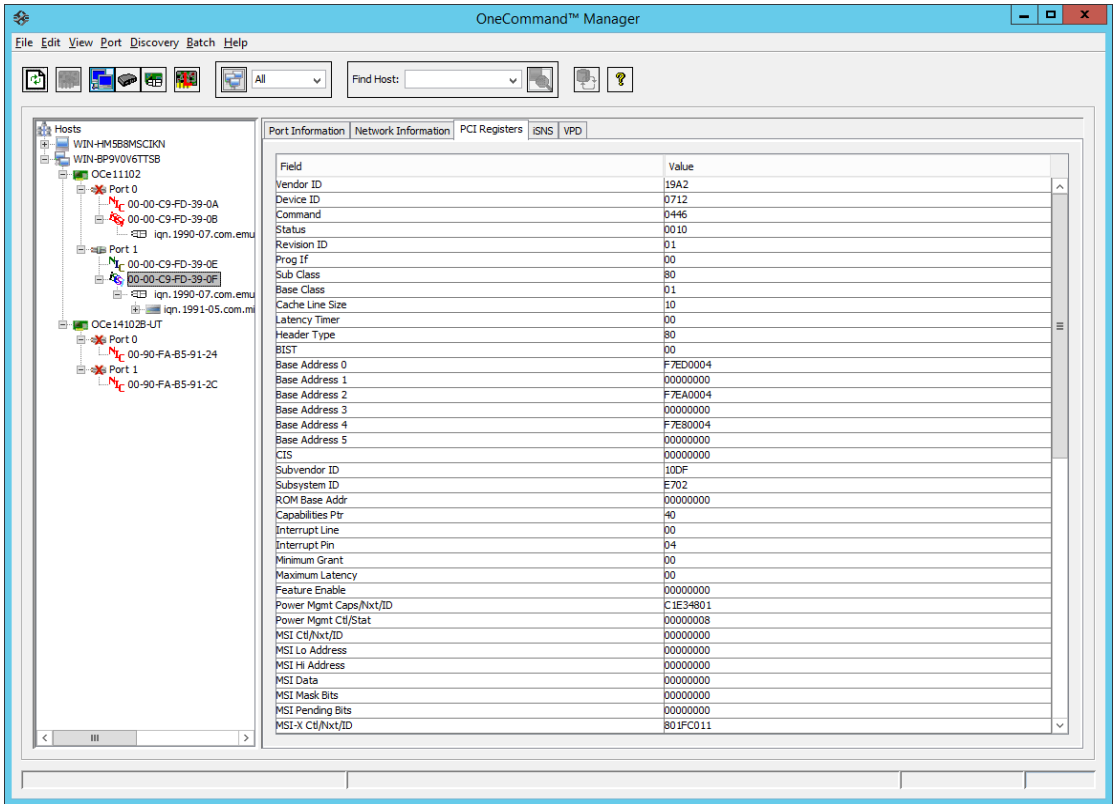
To delete table entries:

1. From the discovery-tree (Figure 5), select the iSCSI function whose configuration you want to modify.
2. Select the **Network Information** tab (Figure 103).
3. Click **Advanced**. The Advanced TCP/IP Configuration dialog box appears (Figure 104).
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.

8.3.3 Viewing iSCSI PCI Registers

The iSCSI PCI Registers tab (Figure 105) displays base PCI registers.

Figure 105 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 105).

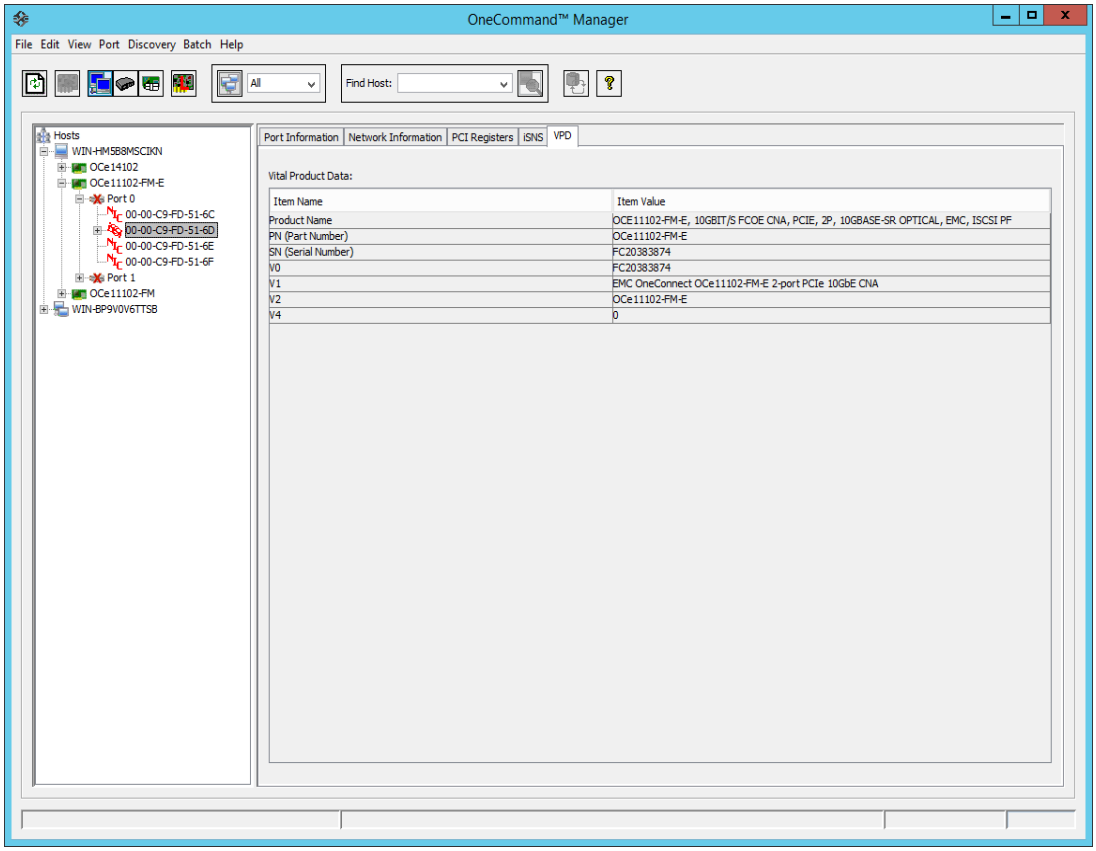
8.3.4 Viewing iSCSI VPD Information

The **VPD** tab (Figure 106) 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 106).

Figure 106 iSCSI VPD Tab



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

- 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 V*x* (where *x* is a number).

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

8.3.5 Viewing iSCSI Statistics

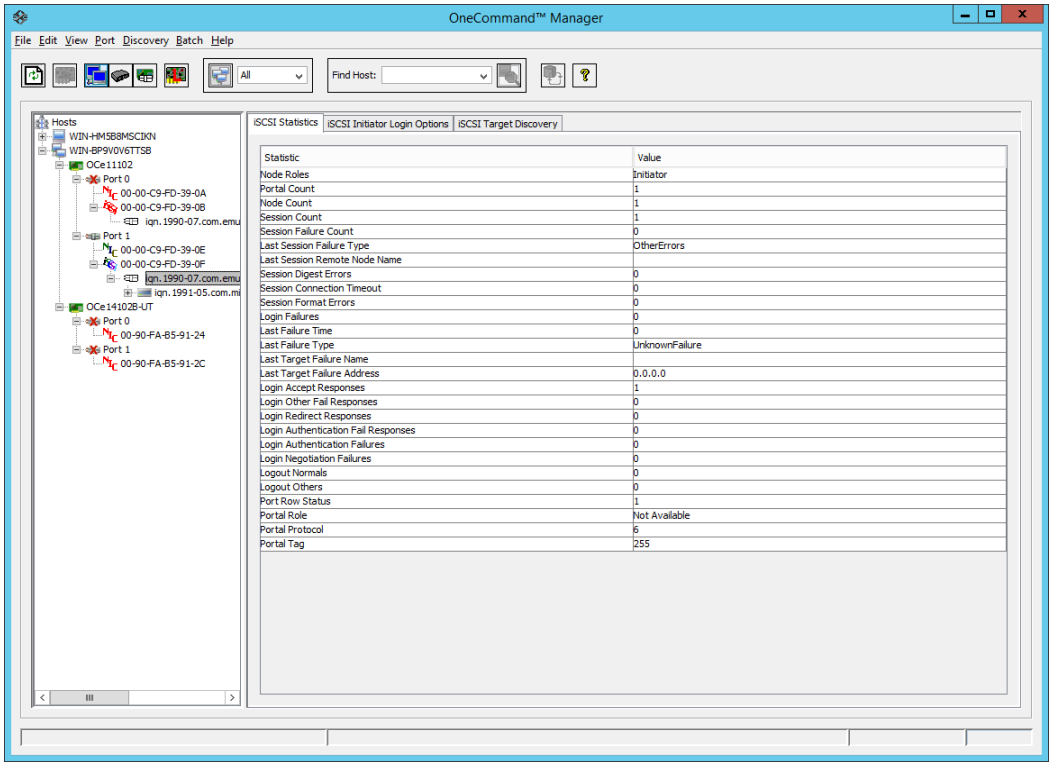
When you select an iSCSI initiator node from the discovery-tree (Figure 5), the **iSCSI Statistics** tab (Figure 107) provides cumulative totals for various error events and statistics on the port.

NOTE iSCSI statistical information is not available from 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 107)

Figure 107 iSCSI Statistics Tab



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

- Node Roles – The node role for this iSCSI initiator.
- Portal Count – The number of rows in the iscsiPortaltypeTable that are currently associated with this iSCSI instance.
- Node Count – The number of rows in the iscsiNodetypeTable that are currently associated with this iSCSI instance.
- Session Count – The number of rows in the iscsiSessiontypeTable 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 caused by a receipt of a PDU containing header or data digest errors.
- Session Connection Timeout – The count of sessions that failed because of a sequence exceeding a time limit.
- Session Format Errors – The count of sessions that failed caused by a 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 halted a login because the target could not be authenticated.
- Login Negotiation Failures – The number of times the initiator has halted 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.

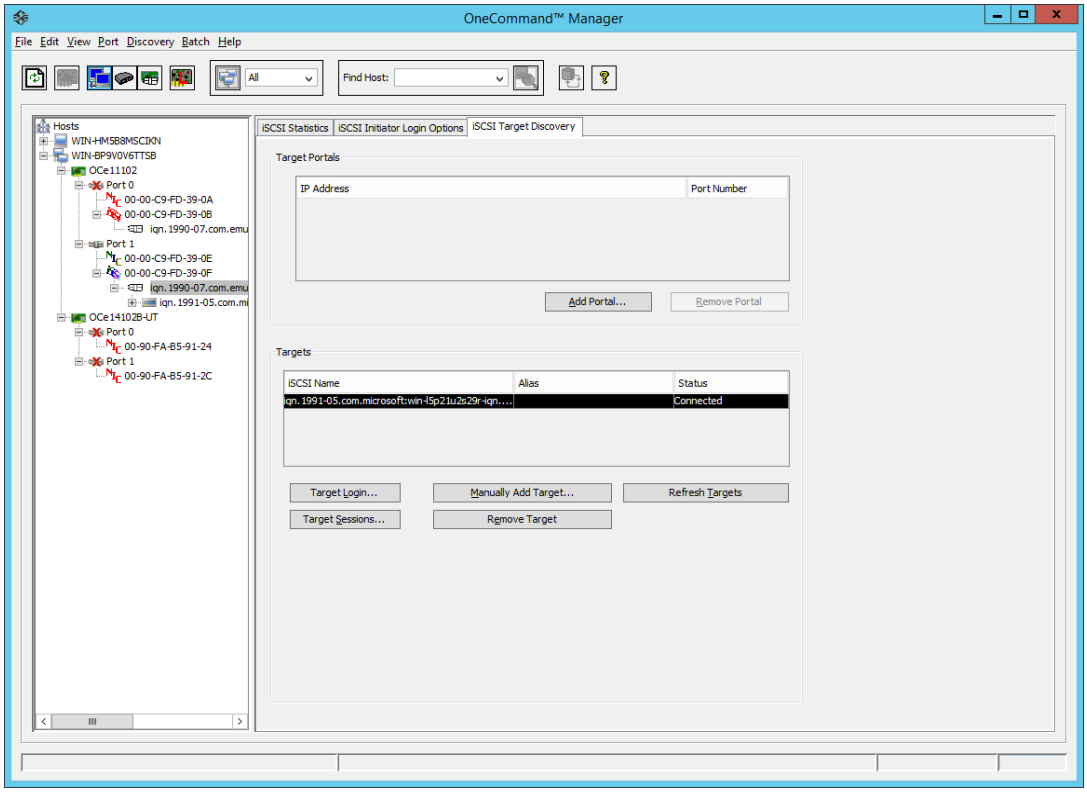
8.3.6 Configuring iSCSI Target Discovery

The **iSCSI Target Discovery** tab allows you to configure iSCSI target discovery related parameters.

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 108).

Figure 108 iSCSI Target Discovery Tab



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

- 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 can 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 by a target portal.
  - The target was found through an iSNS server query.

To add a portal, refer to Section 8.3.6.1, Adding Target Portals.

To remove a portal, refer to Section 8.3.6.2, Removing a Target Portal.

To log in to a target, refer to Section 8.3.6.4, Logging into Targets.

To view active sessions for a selected target, refer to Section 8.3.6.7, Viewing Target Sessions.

To add an iSCSI target manually, refer to Section 8.3.6.5, Manually Adding an iSCSI Target.

To remove an iSCSI target manually, refer to Section 8.3.6.6, Removing Targets.

To manually force a complete rediscovery of the targets, querying all configured iSNS servers and target portals, click **Refresh Targets**.

8.3.6.1 Adding Target Portals

To add a target portal:

1. From the **iSCSI Target Discovery** tab (Figure 108), click **Add Portal**. The Add Target Portal dialog box appears.

2. Enter the server IP address and TCP port number and click **OK**. When a target portal has been successfully added, 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, refer to Section 8.3.9, Configuring iSCSI Initiator Login Options.
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.

8.3.6.2 Removing a Target Portal

To remove a target portal:

1. From the **iSCSI Target Discovery** tab (Figure 108), 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.

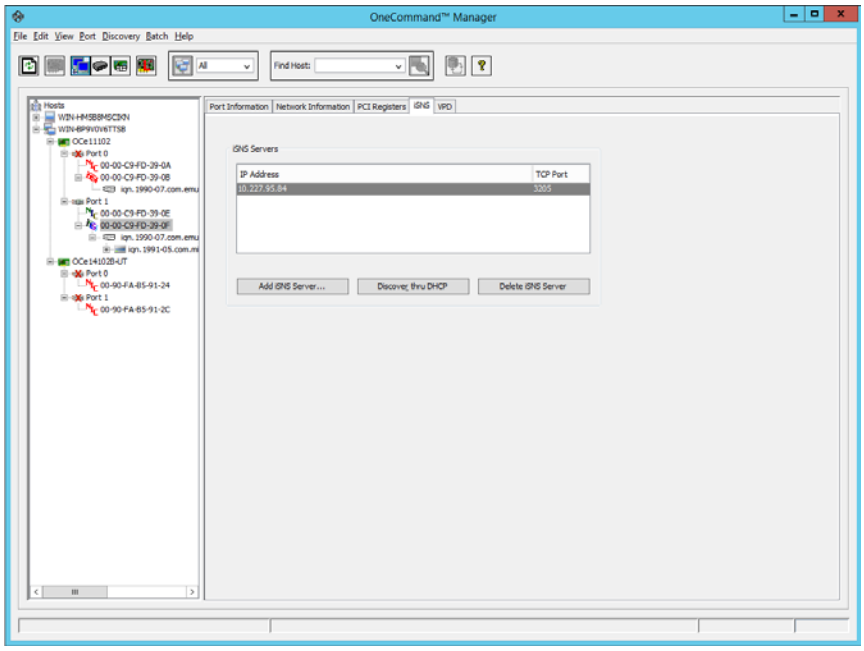
8.3.6.3 Configuring iSNS for iSCSI Target Discovery

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 CNA Manager application.

NOTE On Linux hosts running the Open-iSCSI driver, iSNS is not supported.

Use the **iSCSI iSNS** tab (Figure 109) to configure the iSNS server or to discover the server using DHCP.

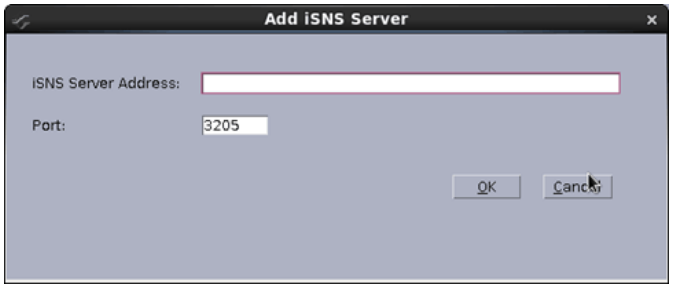
Figure 109 iSCSI iSNS Tab



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 110).

Figure 110 Add iSNS Server Dialog Box



3. Enter the IPv4 or IPv6 address of the iSNS server and the TCP port number.

NOTE On OCe14000-series adapters:

- An IPv4 and a IPv6 iSNS address can be added. For other adapters, only an IPv4 address can be added.
- 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 109). 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:

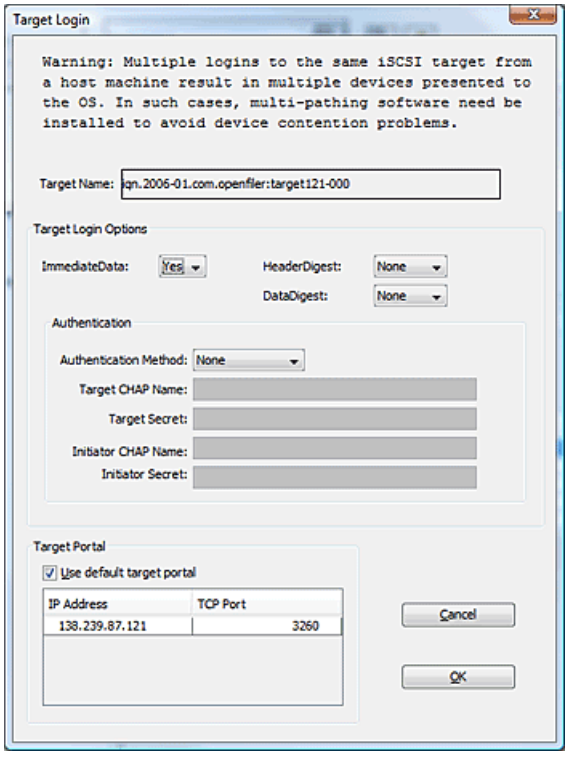
Click **Discover thru DHCP** (Figure 109). If an iSNS server can be discovered through a DCHP server, it is added to the iSNS Server list. After an iSNS is added, its discovered targets should appear in the Targets list of the **iSCSI Target Discovery** tab.

8.3.6.4 Logging into Targets

Only connected targets (targets that are successfully logged in to) are displayed in the discovery-tree (Figure 5). However, the Targets table in the **iSCSI Target Discovery** tab (Figure 108) 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 111 Target Login Dialog Box



To log into a target:

1. From the **iSCSI Target Discovery** tab (Figure 108), select the target from the Targets table.
2. Click **Target Login**. The Target Login dialog box appears (Figure 111). 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 Section 8.3.9, Configuring iSCSI Initiator Login Options.
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.

8.3.6.5 Manually Adding an iSCSI Target

The **iSCSI Target Discovery** tab (Figure 108) enables you to manually add and log into iSCSI targets.

To manually add an iSCSI target:

1. From the **iSCSI Target Discovery** tab (Figure 108), 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 Section 8.3.9, Configuring iSCSI Initiator Login Options.
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 cleared, 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.

8.3.6.6 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 108), select the target you want to remove and click **Remove Target**.

8.3.6.7 Viewing Target Sessions

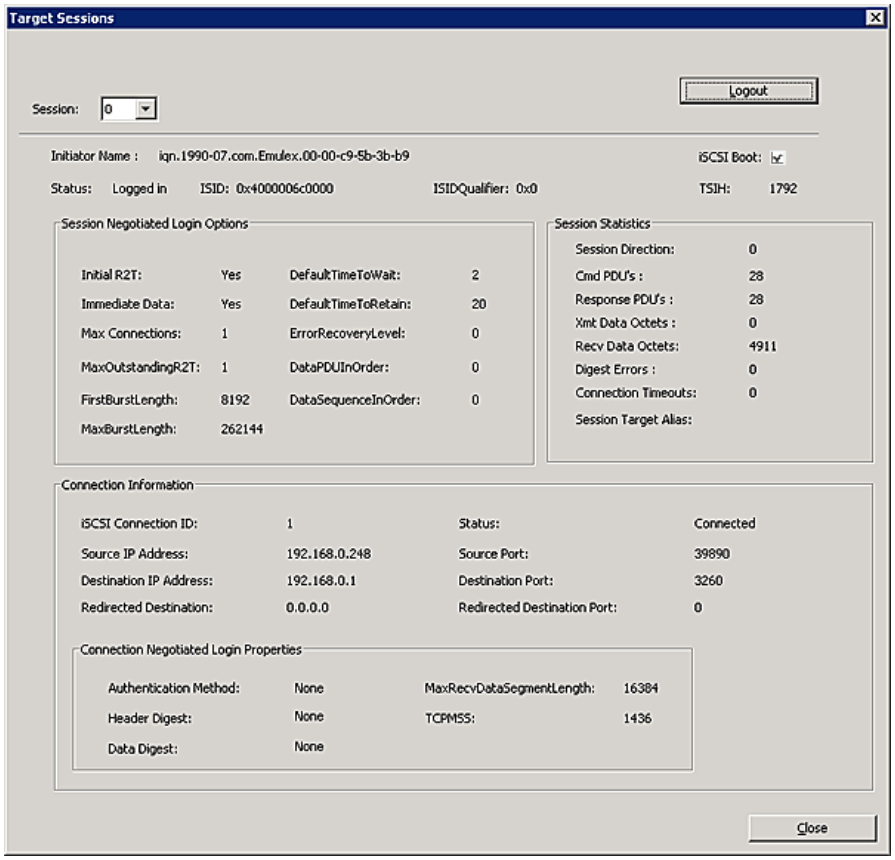
The Target Sessions dialog box (Figure 112) enables you to view sessions for a target.

NOTE Making a target session a boot session for open-iSCSI implementation in the OneCommand CNA Manager application is not supported. Therefore, when running the Emulex open-iSCSI driver, the OneCommand CNA 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 108), select the target whose sessions you want to view and click **Target Sessions**. The Target Sessions dialog box appears (Figure 112).
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.

Figure 112 Target Sessions Dialog Box



The following Target Sessions fields are displayed:

- 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. Clear it to disable boot from this target. You must clear this box to log out from a target. However, even if you clear 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 caused by 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 Properties area:
  - 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.

8.3.6.8 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 108), select the target whose sessions you want to log out of and click **Target Sessions**. The Target Sessions dialog box appears (Figure 112).

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 5). However, the target is still available for login later.

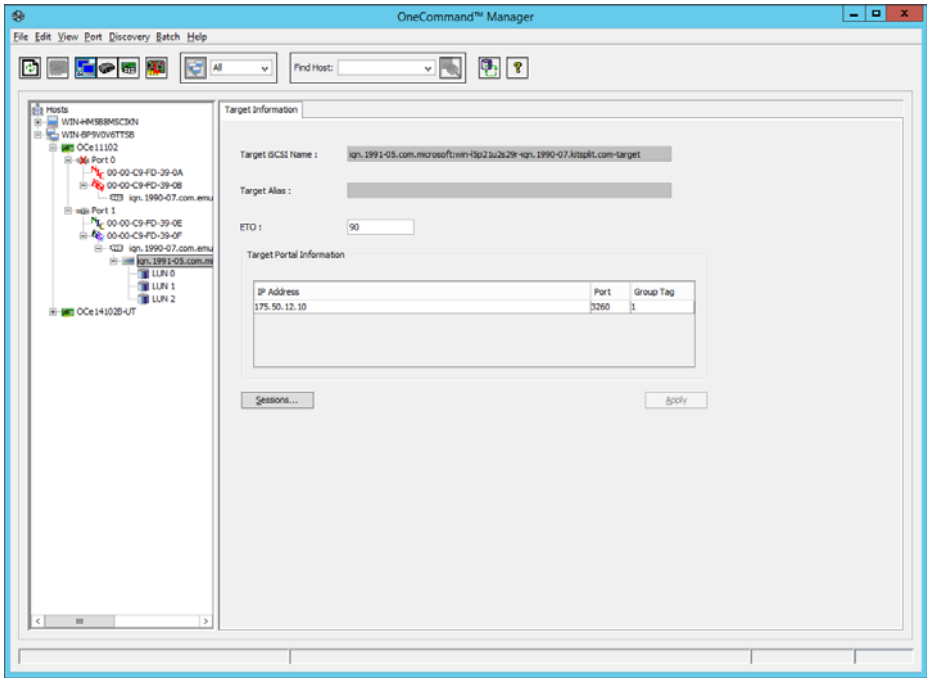
8.3.7 Viewing iSCSI Target Information

When you select a target associated with a iSCSI adapter from the discovery-tree (Figure 5), the **Target Information** tab (Figure 113) 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 113).

Figure 113 Target Information Tab



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

- 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 CNA Manager application.
- ETO – The ETO for the target. The driver makes sure that I/Os are not timed out until this time has expired (from the time the target stopped responding). This value can be changed.
- 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 subgroups must be refreshed.

To view the currently active sessions for the target, see Section 8.3.6.7, Viewing Target Sessions.

8.3.8 Viewing iSCSI LUN Information

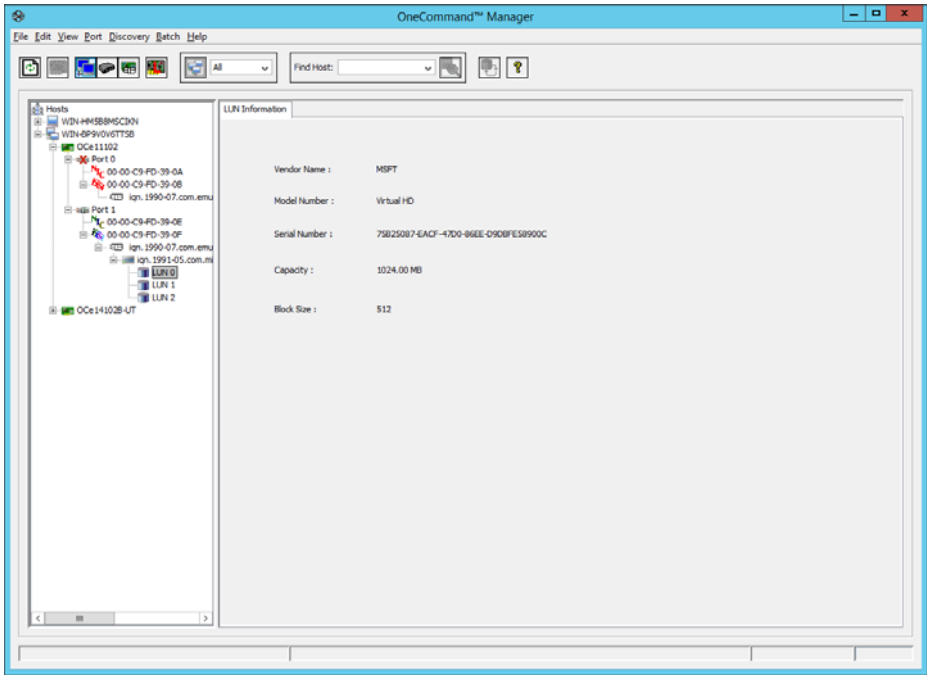
When you select a LUN associated with an iSCSI adapter from the discovery-tree, the **LUN Information** tab (Figure 114) displays information associated with that LUN.

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 114).

Figure 114 LUN Information Tab



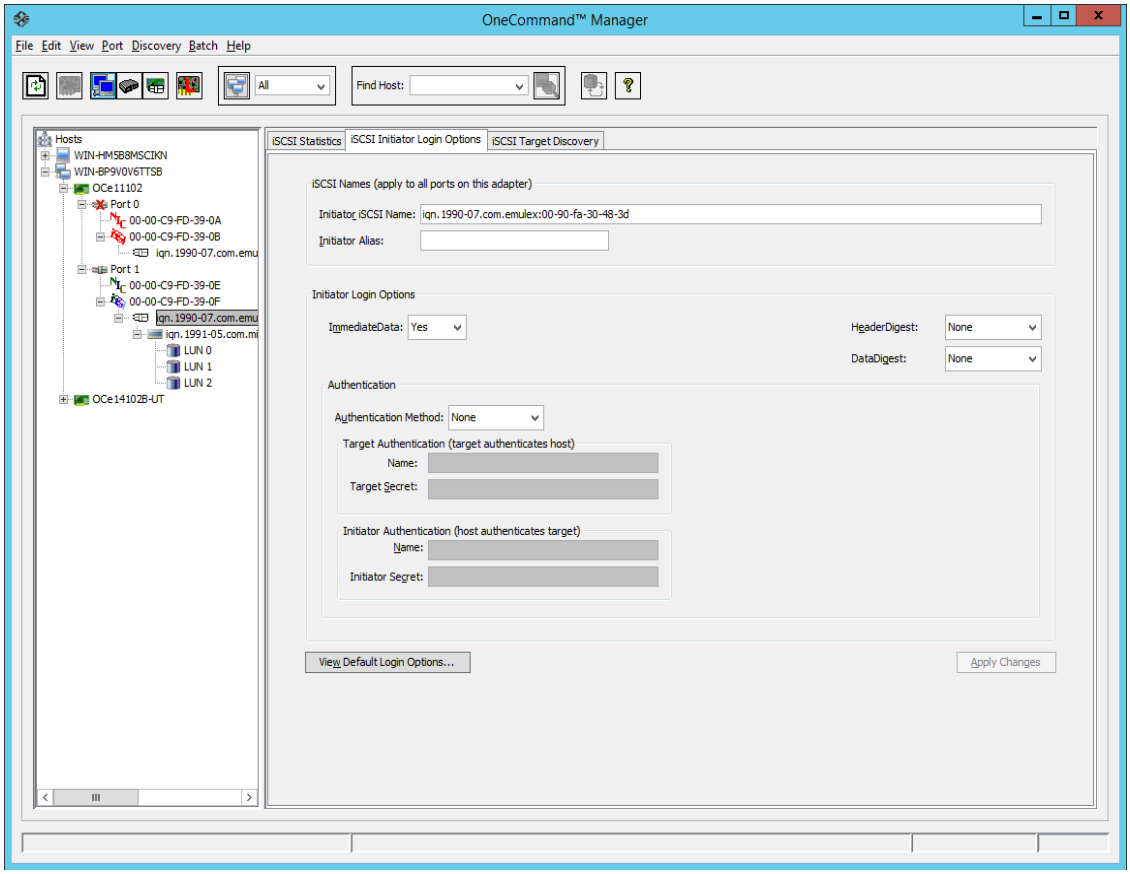
The following **LUN Information** tab fields are displayed:

- 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.

8.3.9 Configuring iSCSI Initiator Login Options

The iSCSI Initiator Login Options dialog box (Figure 115) 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 115 iSCSI Initiator Login Options Tab



The following **iSCSI Initiator Login Options** tab fields are displayed:

- iSCSI Names area:
  - Initiator iSCSI Name – The iSCSI qualifier name of the initiator. The initiator name must be followed by a colon and the port number.
  - 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.

To configure iSCSI port initiator login:

1. In the discovery-tree (Figure 5), select the iSCSI initiator node you want to configure.
2. Select the **iSCSI Initiator Login Options** tab (Figure 115) and make your changes.
3. Click **Apply Changes**.

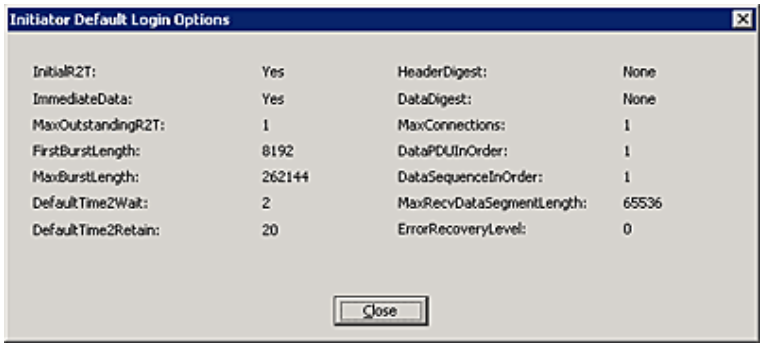
NOTES

- 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:

1. In the discovery-tree (Figure 5), select the iSCSI initiator node whose default login settings you want to view.
2. Select the **iSCSI Initiator Login Options** tab (Figure 115) and click **View Default Login Options**. The Initiator Default Login Options window appears (Figure 116).

Figure 116 Initiator Default Login Options Window



The following Initiator Default Login Options fields are displayed:

- 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 transmits (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.

Chapter 9: Updating Adapter Firmware

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

NOTE If a secure version of firmware (version 11.0 or later) is installed on an OCe14000B-series adapter and you want to update to an earlier unsecured version of firmware, you must remove the secure firmware jumper block before performing the update. See the installation guide for the adapter for more information.

9.1 Updating Firmware for a Single Adapter

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

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

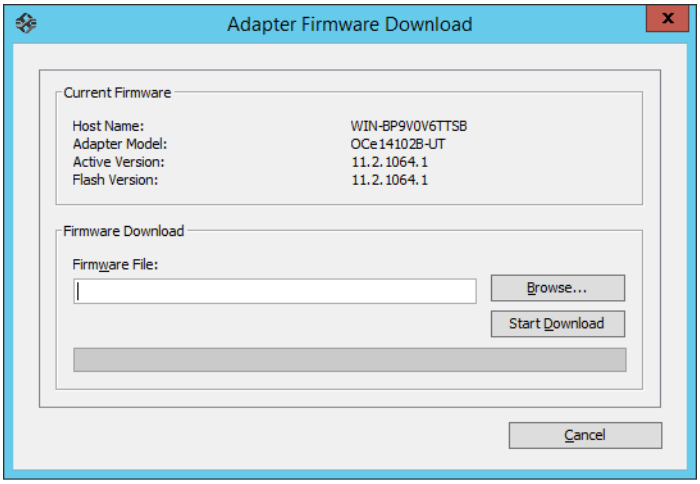
- Download the firmware file from http://www.broadcom.com to a local drive.
- Make sure that the Emulex driver is installed.
- Make sure that the OneCommand CNA Manager application is installed.
- If the adapter is already connected to a boot device, the system must be in a state in which this type of 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.
- 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.)

To update firmware for a single adapter, adapter port, or ASIC:

1. Select **Host** or **Fabric** view.
2. In the discovery-tree (Figure 5), select the adapter, FC port, or ASIC whose firmware you want to update.
3. Select the **Firmware** tab (Figure 28) and click **Download Firmware**. The Firmware Download dialog box appears (Figure 117).

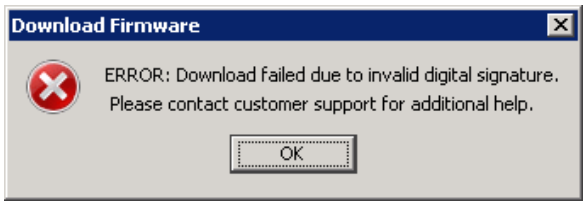
Figure 117 Firmware Download Dialog Box



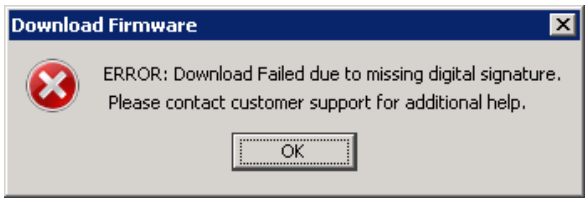
4. Using the Firmware Download dialog box (Figure 117), navigate to the unzipped, extracted image file you want to download. The firmware image may be specified either by entering the image file's full pathname in the Firmware File field or by clicking the **Browse** button.

   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.

NOTE If you attempt to update unauthenticated firmware for a secure OCe14000B-series adapter, the following error message is displayed:



If you attempt to update unsecured firmware for a secure OCe14000B-series adapter, the following error message is displayed:



Contact customer support for more information.

A status bar shows the progress of the download. The adapter in the discovery-tree (Figure 5) is displayed in black text when the update is complete.

NOTE The adapter in the discovery-tree may change to blue during the download. This is normal.

7. Click **Close**. The Firmware tab displays the updated firmware information for the selected adapter.

9.2 Updating Firmware for Multiple Adapters

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

When using the OneCommand CNA Manager application Web Launch Interface, the firmware file must reside on the host where the browser window was started from, not the host that was specified in the web address.

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

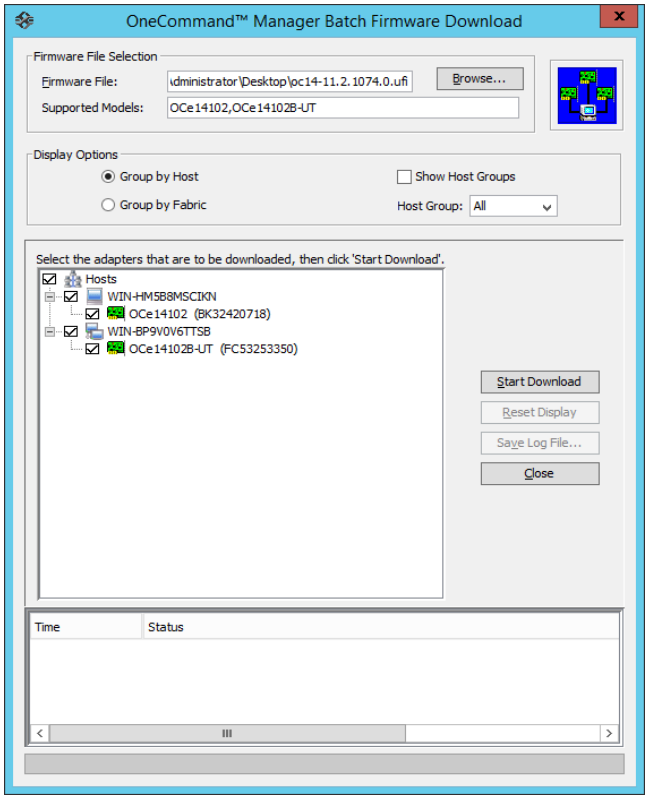
Before you can perform a batch update, the firmware file must be downloaded from http://www.broadcom.com and extracted to a directory on your local drive.

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

To update firmware for multiple adapters:

1. From the **Batch** menu, select **Download Firmware**.
   You do not need to select a particular discovery-tree element for this operation.
2. When the Batch Firmware Download dialog box appears (Figure 118), click **Browse**.

Figure 118 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 CNA 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.

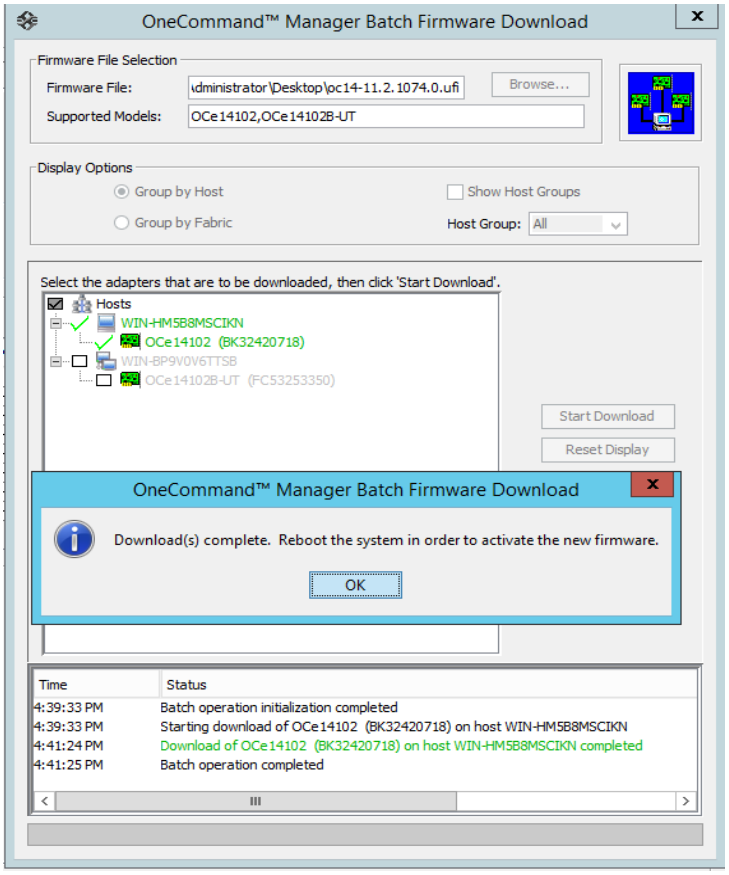
You can also display host groups by checking **Show Host Groups**. To display a particular host group, choose that group from the **Host Group** menu.

Checkboxes next to the 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 119).

Figure 119 Batch Firmware Download Dialog Box, Download Complete



When downloading is finished, you can click **Save Log File** to save a copy of the activity log.

Chapter 10: DCB Configuration

The following section describes how to configure DCB parameters on OneConnect adapter ports.

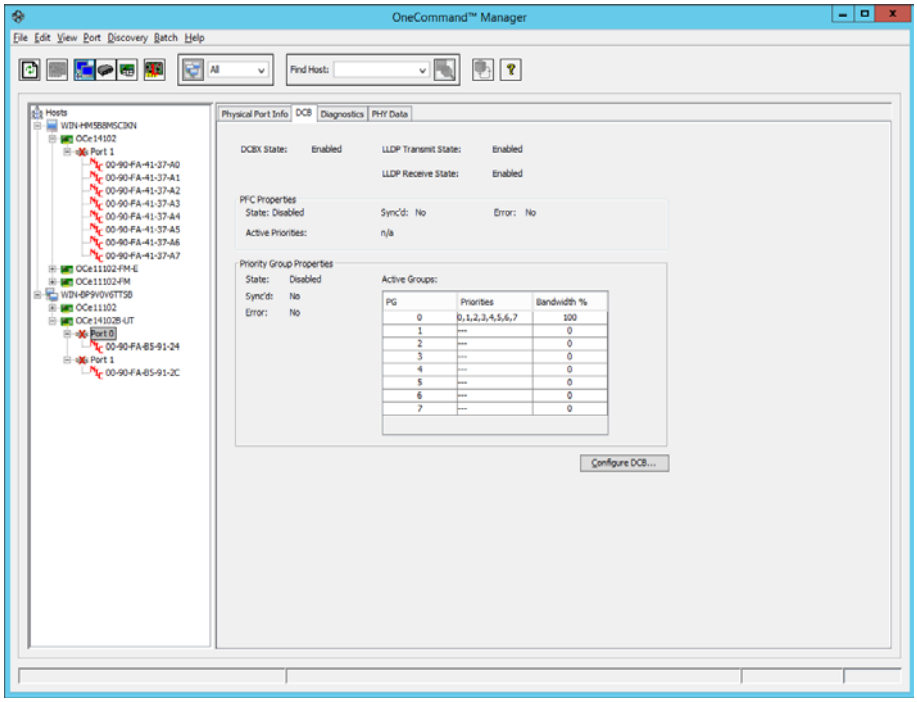
10.1 Configuring DCB Parameters for NIC Only Adapter Ports

The **DCB** tab (Figure 120) 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 120 DCB Tab for NIC Only Adapter Ports (NIC Only Adapter Port Selected)



The following DCB tab fields are displayed:

- 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.

NOTE iSCSI PFC priority 4 is not enabled by default when an OCe14104 adapter is configured for converged mode when configuring DCB parameters for the iSCSI, NIC+iSCSI, and NIC+FCoE+iSCSI ports.

Although you can set one port on an OCe14104 adapter to run multiple protocols, only one priority can be enabled for PCF. The default is to enable FCoE and not iSCSI. iSCSI cannot specify the bandwidth in its own priority group, so the switch must be configured to include it in the default priority group. The adapter does not use the switch settings if the PFC iSCSI priority is manually enabled.

NOTE In NIC-Only mode, PFC is disabled by default. To enable PFC, NIC ETS must be enabled.

To enable NIC ETS:

- On Windows, enable Enhanced Transmission Selection in the driver properties.
- On Linux, load the NIC driver with the tx_prio driver parameter set to 1.

- 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 multichannel 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 area:
  — 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.

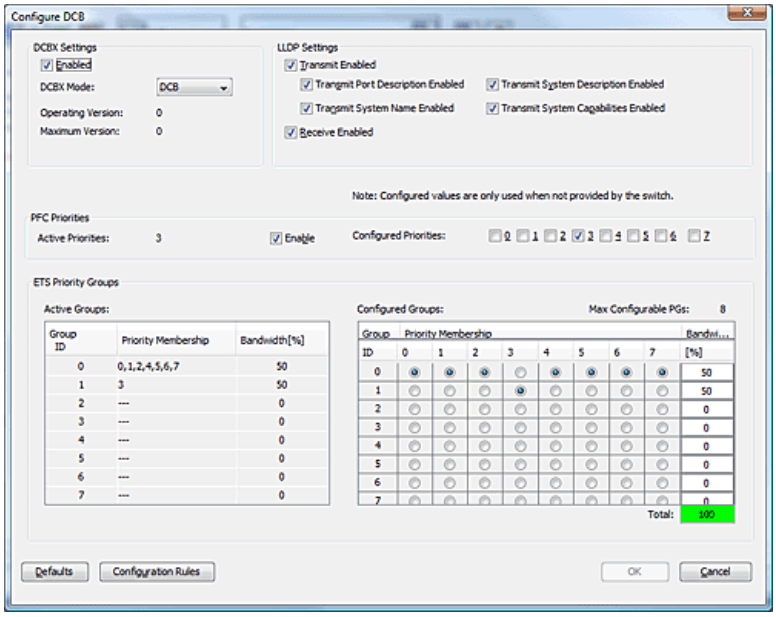
To configure DCB for NIC only adapter ports:

1. From the discovery-tree (Figure 5), select the NIC only adapter port whose DCB properties you want to configure.
2. Select the **DCB** tab (Figure 120).
3. Click **Configure DCB**. The Configure DCB dialog box appears (Figure 121).

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 121 Configure DCB Dialog Box for NIC Only Adapter Ports



The following Configure DCB dialog box fields are displayed:

- 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 or functions 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 checkboxes 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 multichannel 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.

Click to return to the factory settings, click **Defaults**.

To display the **NIC Priority** window, which lists the rules for configuring NIC priorities, click **Configuration Rules**. You must observe the following rules when configuring priority groups for NIC-Only adapter ports:

- Only one PFC priority can be configured.
- The PFC Priority must be assigned to a priority group that has no other priorities.
- Bandwidths of all the priority groups must add up to 100%.

10.1.1 Configuring DCB Parameters for NIC+iSCSI Adapter Ports

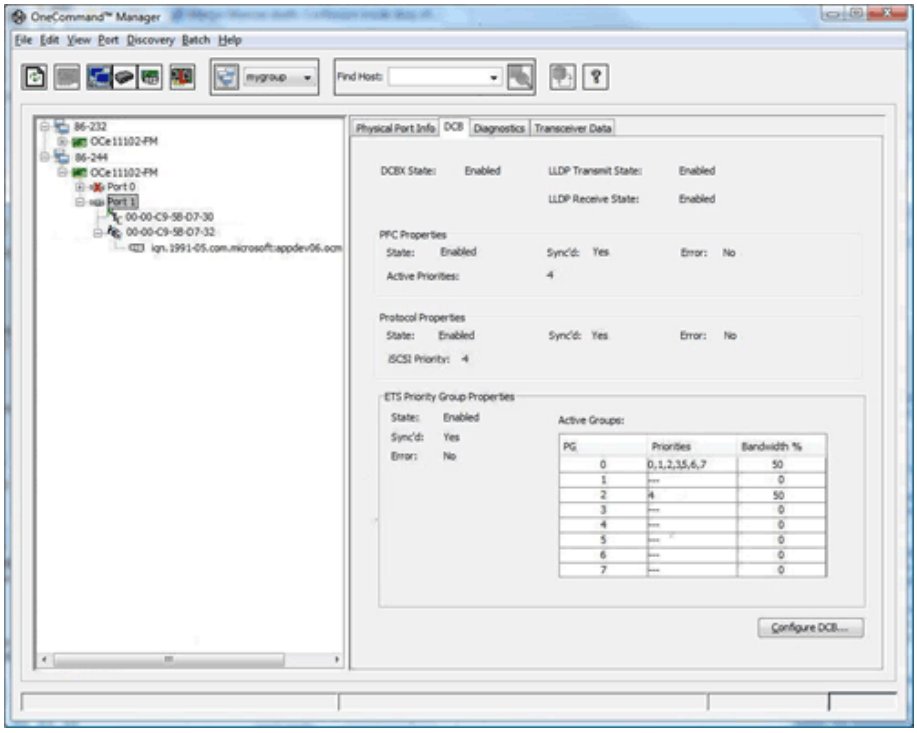
The DCB tab (Figure 122) displays parameters for NIC+iSCSI adapter ports.

NOTE For ports running both iSCSI and FCoE, see Section 10.1.3, Configuring DCB Parameters for NIC+FCoE+iSCSI Adapter Ports.

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 122).

Figure 122 DCB Tab for NIC+iSCSI Adapter Ports (NIC+iSCSI Adapter Port Selected)



The following DCB tab fields are displayed:

- 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 multichannel is enabled on the adapter with the exception of NPar and UFP (if enabled by the switch).

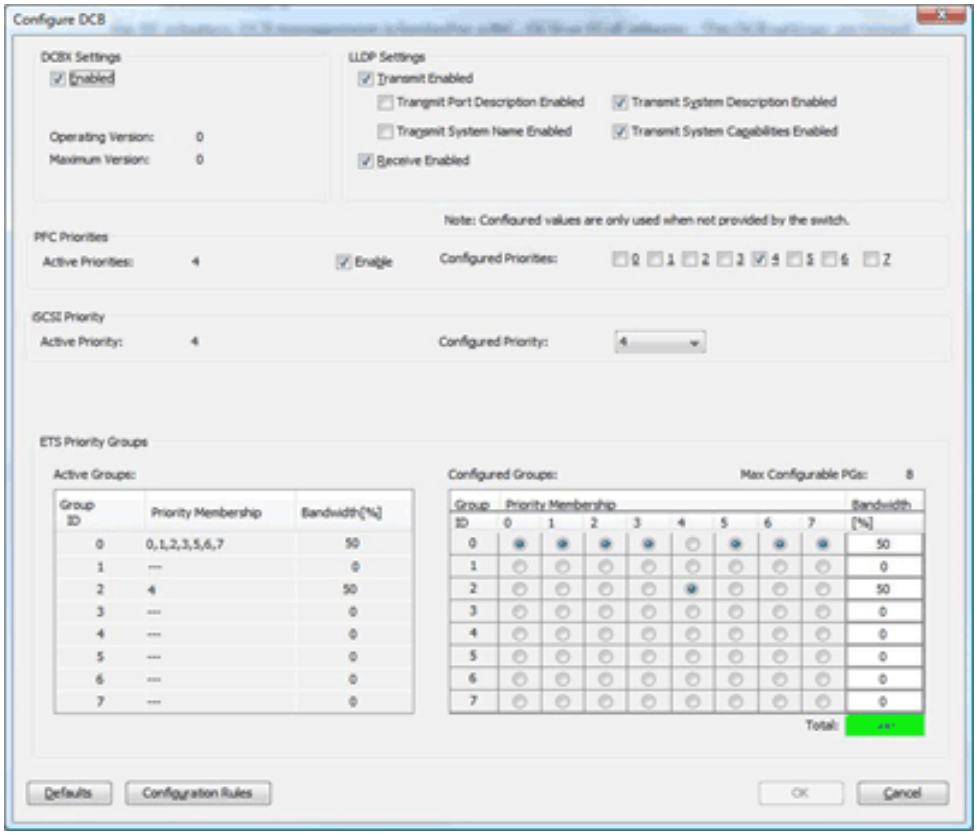
  - State – The current priority group state. It can be enabled or disabled.
  - Sync'd – If yes, the priority groups have been set by the peer. This parameter cannot be set.
  - Error – The error state. This capability indicates whether an error has occurred during the configuration exchange with the peer.
- Active Groups area:
  - 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.

To configure DCB for NIC+iSCSI adapter ports, perform the following steps:

1. From the discovery-tree (Figure 5), select the NIC+iSCSI adapter port whose DCB properties you want to configure.
2. Select the **DCB** tab (Figure 122).
3. Click **Configure DCB**. The Configure DCB dialog box appears (Figure 123).
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 123 Configure DCB Dialog Box for NIC+iSCSI Adapter Ports (DCBX Enabled)



The following Configure DCB dialog box fields are displayed:

- 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 or functions 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 checkboxes 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 multichannel 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.

To return parameters to the default iSCSI DCB settings, click **Defaults**.

To display the iSCSI Priority window, which lists the rules for configuring iSCSI priorities, click **Configuration Rules**. You must observe the following rules when configuring priority groups for iSCSI adapter ports:

- Only one priority can be configured as the iSCSI priority.
- 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.

- The priority group to which the iSCSI priority is assigned must contain no other priorities.
- The additional PFC priority must be assigned to a priority group that has no other priorities.
- Bandwidths of all the priority groups must add up to 100%.

10.1.2 Configuring DCB Parameters for NIC+FCoE Adapter Ports

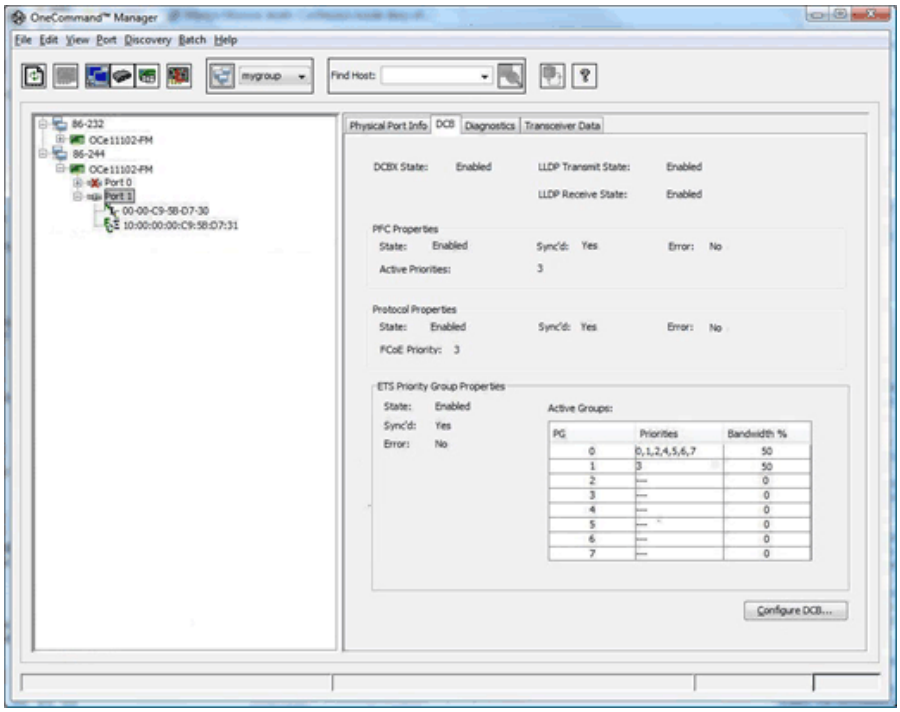
NOTE For ports running both iSCSI and FCoE, see Section 10.1.3, Configuring DCB Parameters for NIC+FCoE+iSCSI Adapter Ports.

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 124).

Figure 124 DCB Tab (NIC+FCoE Adapter Port Selected)



10.1.2.1 The following DCB tab fields are displayed:

- 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 multichannel is enabled on the adapter with the exception of NPar and UFP (if enabled by the switch).

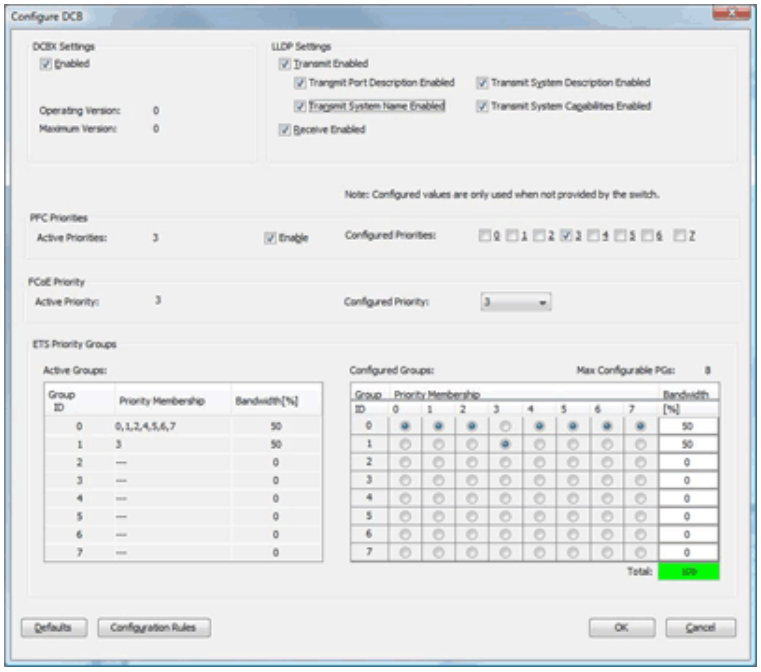
  - State – The priority group state. It can be enabled or disabled.
  - Sync'd – If yes, the priority groups have been set by the peer. This parameter cannot be set.
  - Error – The error state. This capability indicates whether an error has occurred during the configuration exchange with the peer. Error 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.

To configure DCB for NIC+FCoE adapter ports:

1. From the discovery-tree (Figure 5), select the NIC+FCoE adapter port whose DCB properties you want to configure.
2. Select the **DCB** tab (Figure 124).
3. Click **Configure DCB**. The Configure DCB dialog box appears (Figure 125).
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 125 Configure DCB Dialog Box for NIC+FCoE Adapter Ports (DCBX Enabled)



The following Configure DCB dialog box fields are displayed:

- 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 or functions 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 multichannel 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.

To return parameters to default FCoE DCB settings, click **Defaults**.

To display the FCoE Priority window, which lists the rules for configuring FCoE priority group information, click **Configuration Rules**. You must observe the following rules when configuring priority groups for FCoE adapter ports:

- One and only one priority is configured for the FCoE priority.
- 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.
- The priority group to which the FCoE priority is assigned must contain no other priorities.
- The additional PFC priority must be assigned to a priority group which has no other priorities.
- Bandwidths of all the priority groups must add up to 100%.

10.1.3 Configuring DCB Parameters for NIC+FCoE+iSCSI Adapter Ports

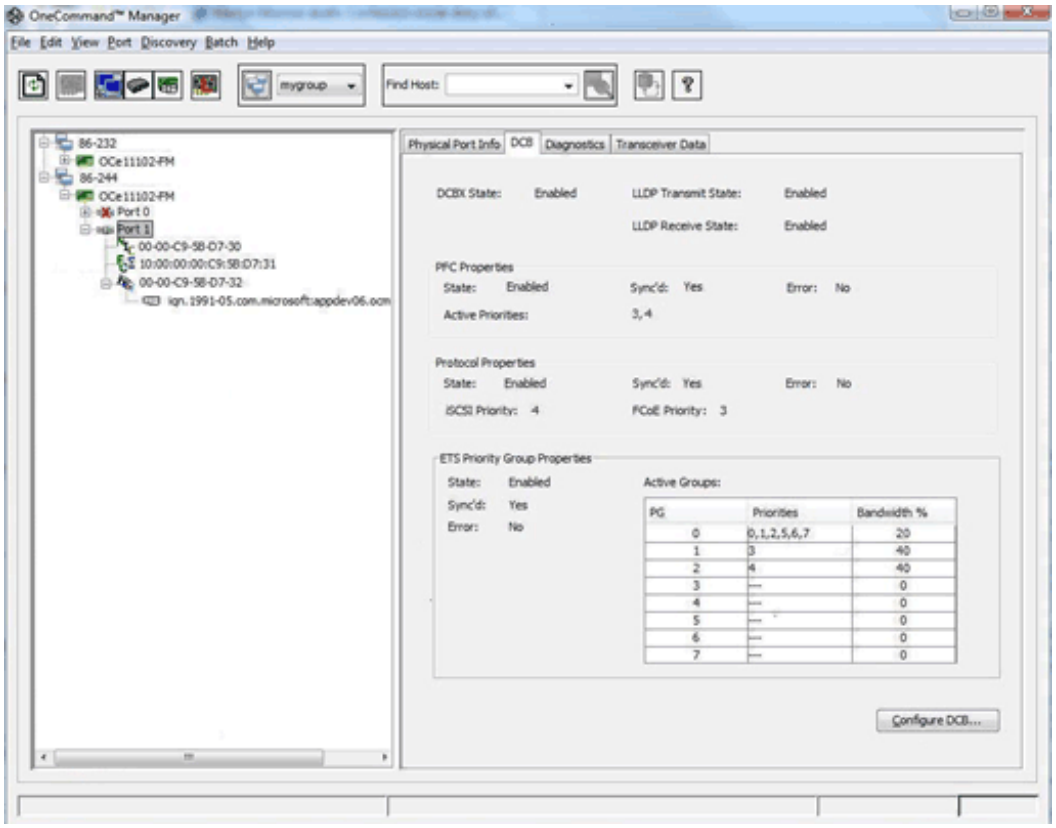
The **DCB** tab displays parameters for NIC+FCoE+iSCSI adapter ports.

NOTE For adapter ports running only FCoE or only iSCSI, refer to Section 10.1.2, Configuring DCB Parameters for NIC+FCoE Adapter Ports and to Section 10.1.1, Configuring DCB Parameters for NIC+iSCSI Adapter Ports.

To view the DCB parameters for NIC+FCoE+iSCSI adapter ports:

1. From the discovery-tree (Figure 5), select the NIC+FCoE+iSCSI adapter port whose DCB properties you want to view.
2. Select the **DCB** tab (Figure 126).

Figure 126 DCB Tab (NIC+FCoE+iSCSI Adapter Port Selected)



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

- 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 multichannel 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.

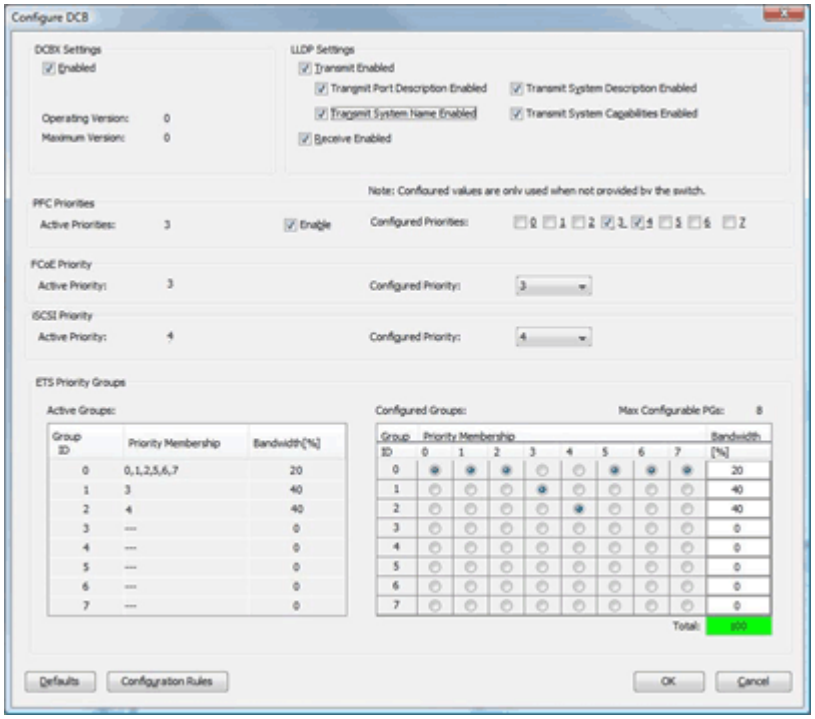
  - Error – The error state. This capability indicates whether an error has occurred during the configuration exchange with the peer. Error is also set to Yes when the Compatible method for the capability fails.
- Active Groups:
  - PG – The priority group number. It can be 0 to 7.
  - Priorities – The 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.

To configure DCB for NIC+FCoE+iSCSI adapter ports:

1. From the discovery-tree (Figure 5), select the NIC+FCoE+iSCSI adapter port whose DCB properties you want to configure.
2. Select the **DCB** tab (Figure 126).
3. Click **Configure DCB**. The Configure DCB dialog box appears (Figure 127).
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 127 Configure DCB Dialog Box for NIC+FCoE+iSCSI Adapter Ports (DCBX Enabled)



The following Configure DCB dialog box fields are displayed:

- 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 or functions 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. Only one PFC priority is supported, so the FCoE or iSCSI priority can be checked, but 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 multichannel 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.

To return parameters to default FCoE /iSCSI DCB settings, click **Defaults**.

To display the iSCSI+FCoE Priority window, which lists the rules for configuring iSCSI+FCoE priority group information, click **Configuration Rules**. You must observe the following rules when configuring priority groups for FCoE adapter ports:

- 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.
- 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.
- The priority group to which the FCoE priority is assigned must contain no other priorities.
- The priority group to which the iSCSI priority is assigned must contain no other priorities.
- The additional PFC priority must be assigned to a priority group that has no other priorities.
- Do not exceed the maximum number of configurable priority groups displayed above the priority groups box when assigning priorities to priority groups.
- Bandwidths of all the priority groups must add up to 100%.

10.1.4 Configuring DCB Parameters for RoCE Adapter Ports

NOTE RoCE is available as Technical Preview only. Do not use RoCE in a production environment.

The **DCB** tab (Figure 128) displays parameters for RoCE adapter ports.

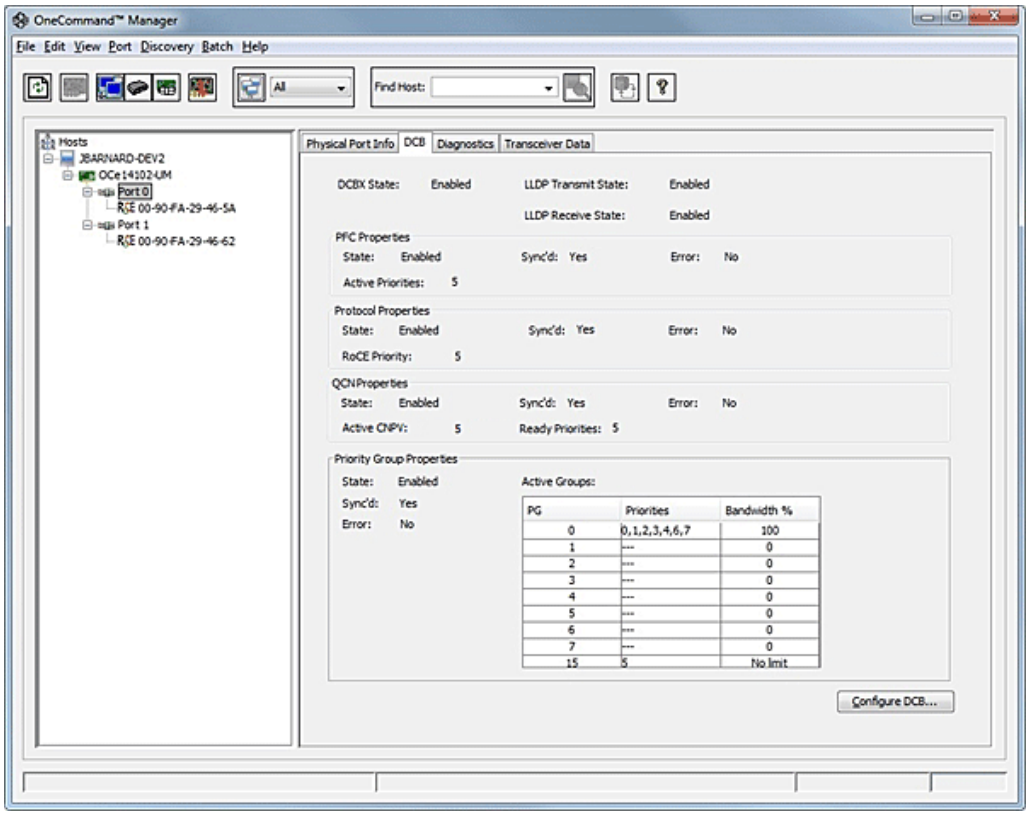
NOTE RoCE configurations are not supported with SR-IOV.

To view the DCB parameters for RoCE adapter ports:

1. From the discovery-tree (Figure 5), select the RoCE adapter port whose DCB properties you want to view.

2. Select the **DCB** tab (Figure 128).

Figure 128 DCB Tab for RoCE Adapter Ports



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

- 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 misconfigured.

- 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. 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.
- 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.

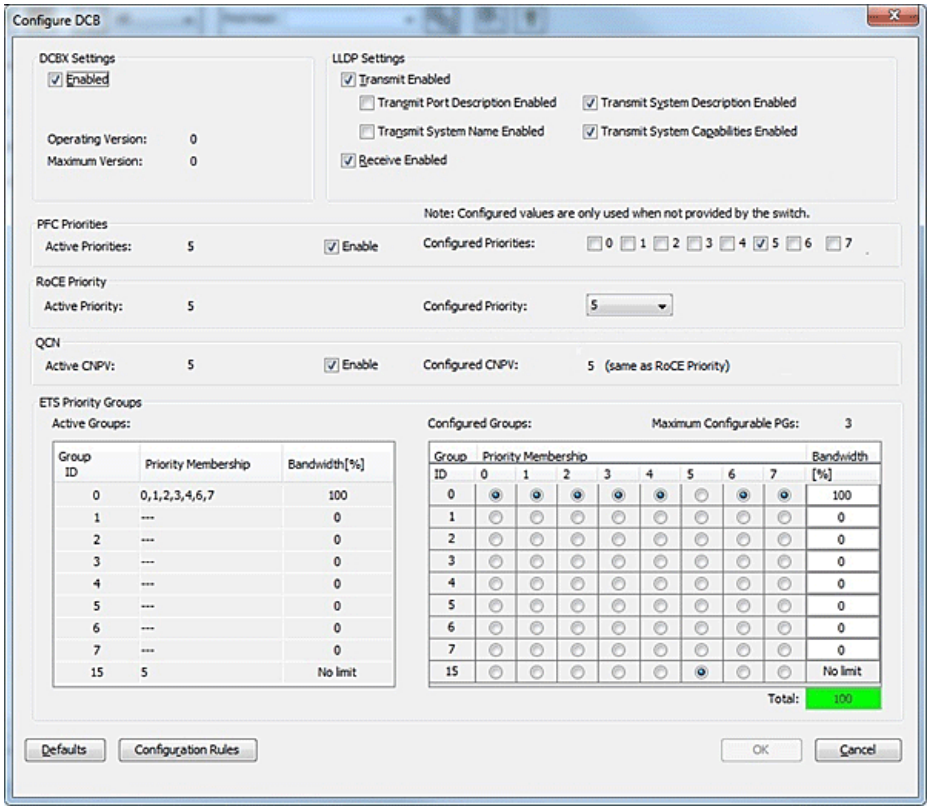
    NOTE PG15 is only available for NIC+RoCE and has no bandwidth limitation.

To configure DCB for RoCE adapter ports:

1. From the discovery-tree (Figure 5), select the RoCE adapter port whose DCB properties you want to configure.
2. Select the **DCB** tab (Figure 128).
3. Click **Configure DCB**. The Configure DCB dialog box appears (Figure 129).
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 129 Configure DCB Dialog Box for RoCE Adapter Ports (DCBX Enabled)



The following Configure DCB dialog box fields are displayed:

- 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 or functions 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 dimmed, 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%.

To return to the factory settings, click **Defaults**.

To display the NIC Priority window, which lists the rules for configuring NIC priorities, click **Configuration Rules**. You must observe the following rules when configuring priority groups for RoCE adapter ports:

- The RoCE priority must also be assigned as a PFC priority.
- Only the RoCE priority can be assigned to PG15.
- 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 must be assigned to a single PG.
- No bandwidth is specified for PG15. The bandwidths for PG0-PG7 must add up to 100.

Chapter 11: Exporting SAN Information

11.1 Creating a SAN Report

The OneCommand CNA 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 CNA 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.

Chapter 12: Diagnostics

This section describes the diagnostic tests that can be run.

NOTES

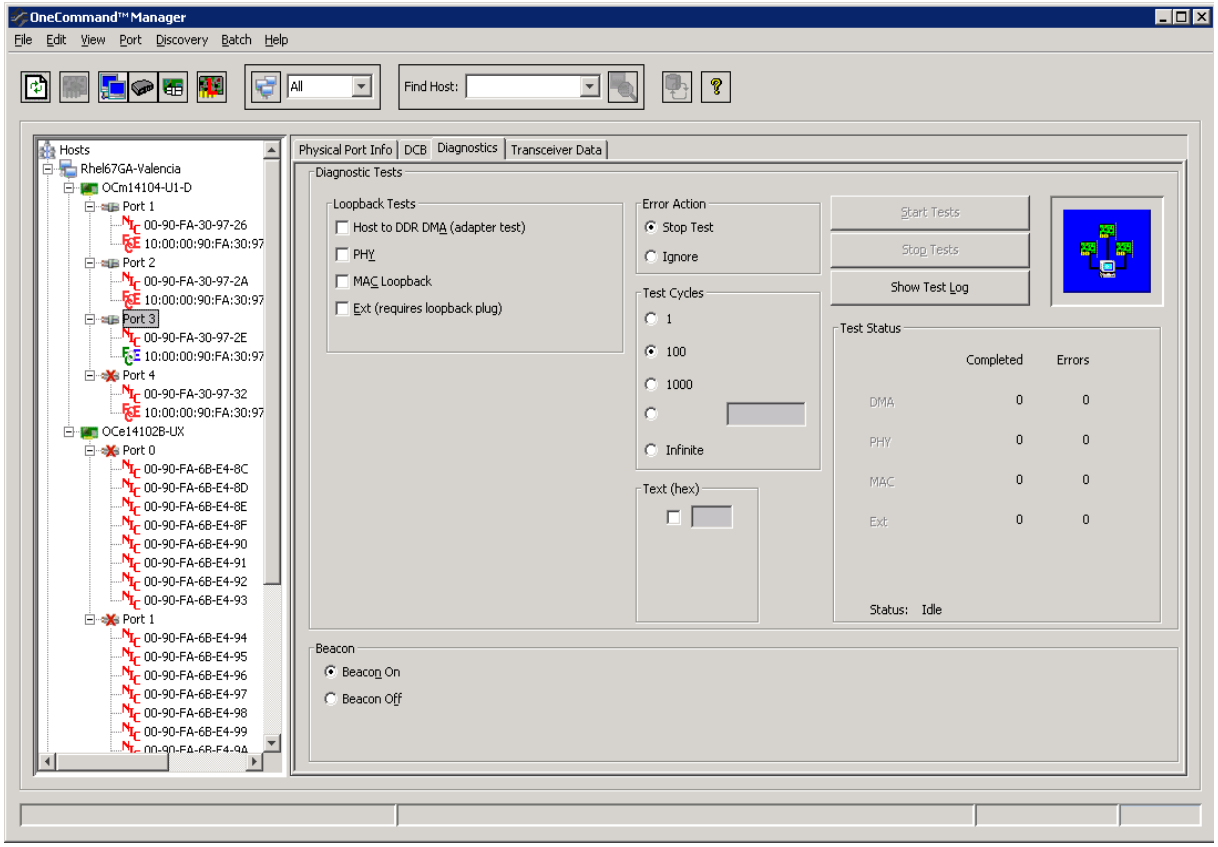
- Diagnostics are not available in read-only mode. See Section 4.3, Changing Management and Read-Only Mode for more information.
- The Echo test is not supported on OCe14000-series and OCe11100-series FCoE adapters.

You can use the **Diagnostics** tab (Figure 130) to perform the following tasks:

- Run the following tests on OneConnect adapters installed in the system:
  - Host to DDR DMA Loopback
  - PHY Loopback
  - MAC Loopback
  - 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 130 Diagnostics Tab (10GBASE-T Adapter Selected)



12.1 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. (This test is 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 by way of a loopback connector. The returned data is subsequently validated for integrity. This test is 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).

12.1.1 Error Action

Error Action enables you to define the actions to be performed in the event of a test failure. Two error action option exists:

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

12.1.2 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.

12.1.3 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.

12.1.4 Test Status

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

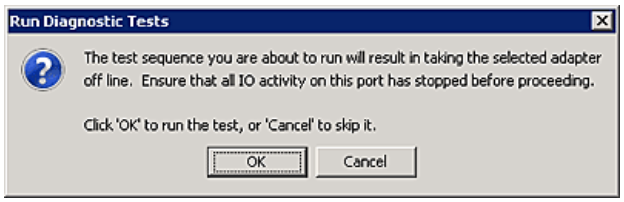
To run loopback tests:

1. From the discovery-tree (Figure 5), select the adapter port on which you want to run the Loopback Test.
2. Select the **Diagnostics** tab (Figure 130). 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. 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 in Section 12.5, Creating Diagnostic Dumps, if the NIC link fails to come up.

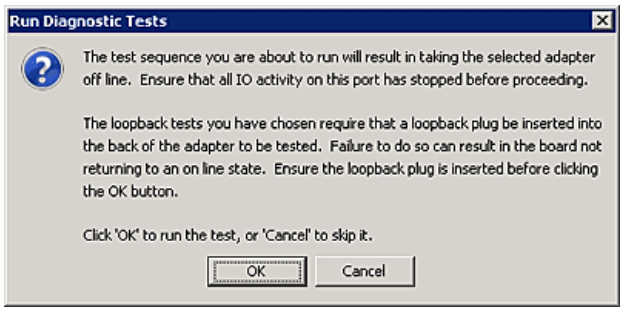
3. Click **Start**. The following warning appears (Figure 131):

Figure 131 Run Diagnostic Tests Warning



4. Click **OK**. If you choose to run an External Loopback test, the following window appears (Figure 132):

Figure 132 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.

12.2 Using Beaconing

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

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.

When you enable beaconing, 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 5), select the adapter port whose LEDs you want to set.

2. Select the **Diagnostics** tab (Figure 130) and click **Beacon On** or **Beacon Off**.

NOTE On OCe11000-series adapters, beaconing can only be used 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.

12.3 Running TDR Tests (10GBASE-T Adapters Only)

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.

NOTE The TDR test is not supported by the CIM interface.

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 5).

To run a TDR test:

1. From the discovery-tree (Figure 5), select the 10GBASE-T adapter on which you want to run the test.
2. Select the **Diagnostics** tab (Figure 130).
3. Assign the NVP and click **Apply**.
4. Click **Start TDR Test**.

12.4 Saving the Log File

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

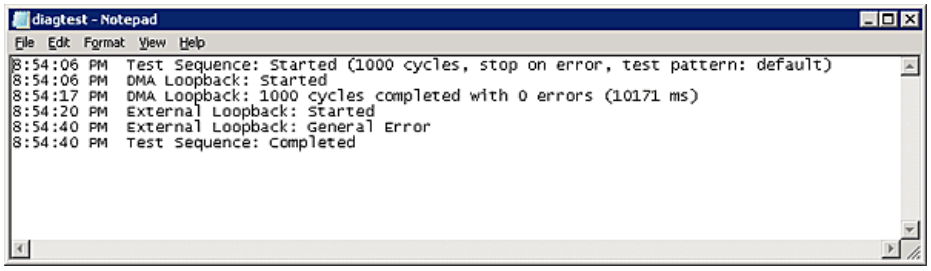
You can save the test log to a log file for later viewing or printing. When new data is written to a saved file, the data is appended to the end of the file. Each entry has a two-line header that contains the identifier of the adapter being tested and the date and time of the test. Over time, the data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

The default location is:

- In Windows: the OneCommand CNA Manager application installation directory on your local drive.
- In Solaris: `/opt/brcmocmanager/Dump`.
- In Linux: `/var/opt/broadcom/brcmocmanager/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 is shown in Figure 133.

Figure 133 Example of a DiagTest.log Window



To save the log file:

1. After running a test from the Diagnostic tab (Figure 130), click **Save Test Log**. The Diagnostic Test Log dialog box appears (Figure 133). 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**.

12.5 Creating Diagnostic Dumps

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

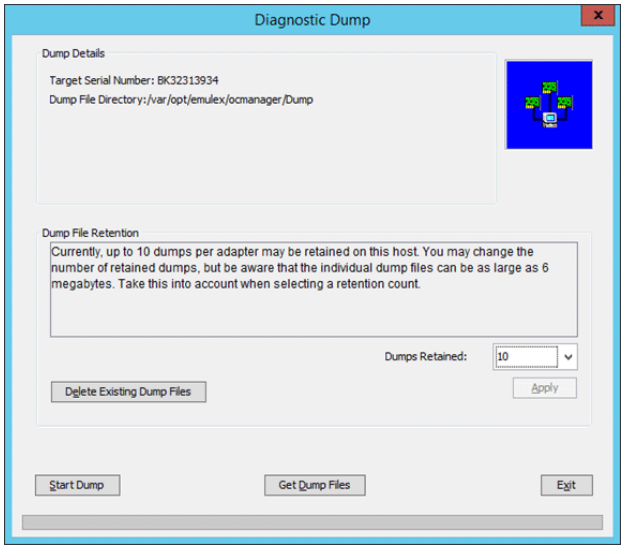
The diagnostic dump capability enables you to create an EFD dump file for a selected adapter. Dump files contain various information such as firmware version, driver version, and so on, that is particularly useful when troubleshooting an adapter. You can also retrieve dump files from remote hosts.

To start a diagnostic dump:

1. From the discovery-tree (Figure 5), select an adapter whose diagnostic information you want to dump.
2. Select the **Firmware** tab (Figure 28) and click **Diagnostic Dump**. The Diagnostic Dump dialog box appears (Figure 134).

   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 134 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/brcmocm/Dump`.
- Linux – `/var/log/broadcom/brcmocmanager/Dump`.
- VMware – a dump directory you create under `/vmfs/volumes`.

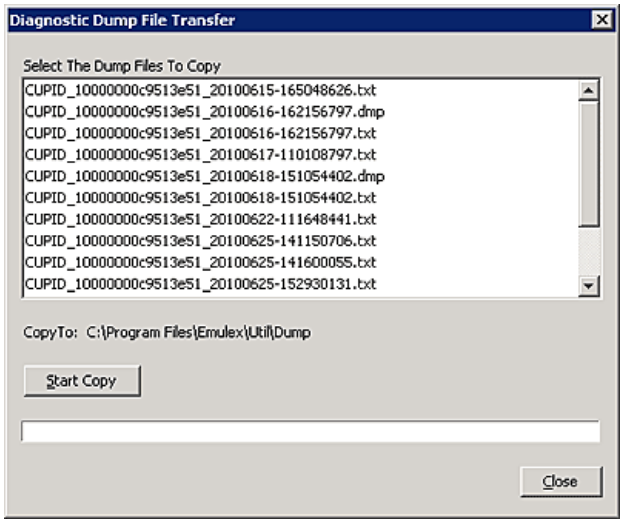
  Two files are created:

  — `<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 135).

   NOTE The Get Dump Files button is disabled if a local adapter port is selected.

Figure 135 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.

Chapter 13: Troubleshooting

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

13.1 General Situations

Table 8 lists possible situations and their resolution.

Table 8 General Situations

Situation	Resolution
After installing and starting the OneCommand CNA 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 CNA Manager application package.
The Web Launch interface cannot be started. When you attempt to start the OneCommand CNA Manager application Web Launch Interface client interface, you receive the following error message: `Unable to launch OneCommand.`	If the JRE/Web Start version present on your system does not meet the minimum required by the OneCommand CNA Manager application, a temporary copy of the correct Web Start version is downloaded automatically. This is used to open the OneCommand CNA 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 CNA Manager application.
The other utilities install, but the OneCommand CNA 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 guide). 2. Install the utilities (see the Installation section of the driver guide).
When attempting to start the OneCommand CNA Manager application, the Web browser displays **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: ■ Close the browser and restart it. The OneCommand CNA Manager application Web Launch Interface starts successfully. ■ Uninstall all nonessential versions of the Java Runtime. The OneCommand CNA 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 8.0 or later for all platforms.
In the OneCommand CNA 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 CNA Manager application. When you attempt to run the utility, an operating system error may occur. The computer may freeze.	Reboot the system.

Table 8 General Situations (Continued)

Situation	Resolution

Unwanted remote servers appear in the OneCommand CNA Manager application.	To prevent remote servers from appearing in the OneCommand CNA Manager application, perform one of the following tasks on the remote systems:
	■ In Windows, disable the OneCommand CNA Manager application service.
	■ In Linux, stop the brcmhbamgr daemon by running the following script:
	`/usr/sbin/brcmocmanager/stop_brcmocmanager`
	■ In Solaris, stop the elxhbamgr service by issuing the following command:
	`svcadm disable brcmhbamgr.`
	NOTE Disabling this service or process prevents the local servers from being seen remotely.

If Help-->Contents is selected in the OneCommand CNA 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.
This happens when the operating system has associated .HTML files with an application other than a web browser.
	On Windows systems, this can be fixed using the following steps:
	1. In Windows Explorer, navigate to the `C:\Program Files\Emulex\Util\OCManager\OCManager_help\` directory.
	2. Right-click on **OCManager_Help.htm**.
	3. Select **Open With-->Choose default program...**
	4. Select a web browser, such as Internet Explorer.
	5. Check **Always use the selected program to open this kind of file**.
	6. Click **OK**.
	On Linux and Solaris, the preceding steps are very similar, with the `brcmocmanager_Help.htm` file located in
	`/usr/sbin/brcmocmanager/brcmocmanager_help/brcmocmanager_Help.htm`
	and
	`/opt/brcmocm/brcmocmanager_help/brcmocmanager_Help.htm`
	respectively.

When using the OneCommand CNA 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 CNA Manager application has been closed prior to making any changes to the NIC driver instance or properties.

13.2 Emulex Driver for Linux and OneCommand CNA Manager Application Situations

Table 9 lists possible Emulex driver for Linux situations and their resolution.

Table 9 Emulex Driver for Linux and OneCommand CNA Manager Application Situations

Situation	Resolution
NIC Link fails to come up.	For Emulex CNAs in NIC mode, you must properly configure the network interface using system administration utilities.
The OneCommand CNA 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 `brcmfcoe-install` script.
If a SAN configuration has 256 targets mapped by the 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.
rmmod fails to unload the FCoE driver module. `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.	Make the Logical Volume Group unavailable. Type the following: `lvchange -a n xxxxxxx` where *xxxxxx* is the Volume Group Name.
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.	If you experience off-line device issues, increase the SCSI command timeout to a value greater than or equal to sixty seconds.
The 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 the latest version.
rmmod of the FCoE driver hangs and the module reference count is 0.	If there is a small race condition in the kernel, it is possible for an rmmod command to stop responding. Issue the rmmod-w command. If this does not help, reboot the computer.
System stops responding when booted with a failed adapter installed.	Remove the failed adapter and reboot.

rmmod 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 CNA Manager application or when the OneCommand CNA Manager application is installed and running.	Stop the OneCommand CNA Manager application before attempting to unload the driver. The script is located in the /usr/sbin/ocmanager directory.
	Type the following:
	`./stop_brcmocmanager`
	Unmount any disks connected to the adapter.
	Unload the driver.
	Type the following:
	`rmmod brcmfcoe`

Driver install fails. The brcmfcoe-`install` script fails to install the driver.	The install script may fail for the following reasons: ■ A previous version of the driver is installed. Run the `/usr/src/brcmfcoe/brcmfcoe-install --uninstall` script and then try to install the driver. ■ The current driver is already installed. ■ The kernel source does not match the standard kernel name or you are running a custom kernel.

Table 9 Emulex Driver for Linux and OneCommand CNA Manager Application Situations (Continued)

Situation	Resolution
■ No module brcmfcoe found for kernel error message. When updating the kernel, the `.rpm` file generates the following error: No module lpfc found for kernel KERNELVERSION. ■ A recently updated kernel cannot find the ramdisk. After updating the kernel, the kernel cannot find the ramdisk that halts the system. ■ The driver is not loaded after a system reboot after updating the kernel.	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. ■ Update the kernel using the same version of the driver. ■ Update the kernel using a new version of the driver. See the Installation section of the driver guide for these procedures.
Driver uninstall fails. The `brcmfcoe_install --uninstall` script fails with an error.	Try the following solutions: ■ Uninstall the OneCommand CNA Manager application by running the `/uninstall` script from the OneCommand CNA Manager application installation directory. ■ Unload the FCoE driver.
`brcmfcoe_install` script exit code.	The `brcmfcoe_install` script contains exit codes that can be useful in diagnosing installation issue. See the `brcmfcoe-install` script for a complete listing of codes and definitions.
The OneCommand CNA Manager application software package does not install. The following error message is displayed: `inserv Service Elxlpfc has to be enabled for service ElxDiscSrvinserv: exiting now/sbin/inserv failed exit code 1.`	Reinstall the driver with the `brcmfcoe_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/brcmfcoe/lun_scan`. SuSE supplies `/bin/rescan-scsi-bus.sh` which can be changed to scan everything.

Table 9 Emulex Driver for Linux and OneCommand CNA Manager Application Situations (Continued)

Situation	Resolution
Cannot see any adapters.	Try the following solutions: ■ Perform an **lsmod** to see if the Emulex drivers are loaded. Look for an error message on the command line stating that the brcmfcoe driver is not loaded. If this is the case, do an **insmod** of the FCoE driver and restart the OneCommand CNA Manager application. ■ Exit the OneCommand CNA Manager application and run the following sripts in this order: 1. `/usr/sbin/brcmocmanager/stop_brcmocmanager` – stops the OneCommand CNA Manager application daemons. 2. `/usr/sbin/brcmocmanager/start_brcmocmanager` – starts the OneCommand CNA Manager application daemons. 3. `/usr/sbin/brcmocmanager/brcmocmanager` – starts the OneCommand CNA Manager application GUI. The adapters should be visible. If they are not visible, reboot your system.
Cannot see new LUNs.	Try the following: 1. Click the **Refresh LUNs** button in the toolbar. 2. Exit the OneCommand CNA Manager application and restart it. If new LUNs are visible, you are finished. If that does not work, try the following: 1. Exit the OneCommand CNA Manager application. 2. Navigate to `/usr/sbin/brcmocmanager`. 3. Run `./stop_brcmocmanager` to stop both the `elxhbamgr` and `brcmdiscovery` processes. 4. Run `./start_brcmocmanager` and `./start_brcmdiscovery` to restart both processes. 5. Start the OneCommand CNA Manager application.

13.3 Emulex Driver for Solaris and OneCommand CNA Manager Application Situations

Table 10 lists possible Emulex driver for Solaris situations and their resolution.

Table 10 Emulex driver for Solaris and OneCommand CNA Manager Application Situations

Situation	Resolution
NIC link fails to come up.	For Emulex CNAs in NIC mode, you must properly configure the network interface using system administration utilities.

13.4 VPorts and OneCommand CNA Manager Application Situations

Table 11 lists possible VPorts situations and their resolution.

Table 11 VPorts and OneCommand CNA Manager Application Situations

Situation	Resolution
vPort creation failure.	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: 1. Start the OneCommand CNA Manager application. 2. Select **View>Group Adapters by Virtual Port** from the Main menu. 3. In the discovery-tree, select the FCoE function on which you want to create a virtual port. 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 Broadcom technical support.
Port not ready.	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: ■ The driver NPIV parameter is disabled. ■ SLI-3 is not being used by a port. ■ The adapter port is out of resources for additional virtual ports. ■ The port is not connected to a fabric. ■ The fabric switch does not support virtual ports. ■ The fabric switch is out of resources for additional virtual ports. ■ The port link state is down.

Appendix A: License Notices

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

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

A.2 OpenPegasus Licensing Notice

Licensed to The Open Group (TOG) under one or more contributor license agreements. Refer to the OpenPegasusNOTICE.txt file distributed with this work for additional information regarding copyright ownership. Each contributor licenses this file to you under the OpenPegasus Open Source License; you may not use this file except in compliance with the License.
Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit

persons to whom the Software is furnished to do so, subject to the following conditions:
The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.
THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

A.3 OpenSSL Notice

This is a copy of the current LICENSE file inside the CVS repository.
LICENSE ISSUES
==========
The OpenSSL toolkit stays under a dual license, i.e. both the conditions of the OpenSSL License and the original SSLeay license apply to the toolkit. below for the actual license texts. Actually both licenses are BSD-style Open Source licenses. In case of any license issues related to OpenSSL please contact openssl-core@openssl.org.
OpenSSL License

/* ==
* Copyright (c) 1998-2011 The OpenSSL Project. All rights reserved.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
*
* 1. Redistributions of source code must retain the above copyright
* notice, this list of conditions and the following disclaimer.
*
* 2. Redistributions in binary form must reproduce the above copyright
* notice, this list of conditions and the following disclaimer in
* the documentation and/or other materials provided with the
* distribution.
*
* 3. All advertising materials mentioning features or use of this
* software must display the following acknowledgment:
* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit. (http://www.openssl.org/)"
*
* 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to
* endorse or promote products derived from this software without
* prior written permission. For written permission, please contact
* openssl-core@openssl.org.

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

Original SSLeay License

/* Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)
* All rights reserved.
*
* This package is an SSL implementation written
* by Eric Young (eay@cryptsoft.com).
* The implementation was written so as to conform with Netscapes SSL.
*
* This library is free for commercial and non-commercial use as long as
* the following conditions are aheared to. The following conditions
* apply to all code found in this distribution, be it the RC4, RSA,
* lhash, DES, etc., code; not just the SSL code. The SSL documentation
* included with this distribution is covered by the same copyright terms
* except that the holder is Tim Hudson (tjh@cryptsoft.com).

```
*
* Copyright remains Eric Young's, and as such any Copyright notices in
* the code are not to be removed.
* If this package is used in a product, Eric Young should be given attribution
* as the author of the parts of the library used.
* This can be in the form of a textual message at program startup or
* in documentation (online or textual) provided with the package.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
* 1. Redistributions of source code must retain the copyright
* notice, this list of conditions and the following disclaimer.
* 2. Redistributions in binary form must reproduce the above copyright
* notice, this list of conditions and the following disclaimer in the
* documentation and/or other materials provided with the distribution.
* 3. All advertising materials mentioning features or use of this software
* must display the following acknowledgement:
* "This product includes cryptographic software written by
* Eric Young (eay@cryptsoft.com)"
* The word 'cryptographic' can be left out if the 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.]
*/
```

A.4 Java Notice

Copyright(c) 2002 Sun Microsystems, Inc. All Rights Reserved.

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

- Redistribution of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistribution in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of Sun Microsystems, Inc. or 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 "AS IS," without a warranty of any kind. ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE HEREBY EXCLUDED. SUN MICROSYSTEMS, INC. ("SUN") AND ITS LICENSORS SHALL NOT BE LIABLE FOR ANY DAMAGES SUFFERED BY LICENSEE AS A RESULT OF USING, MODIFYING OR DISTRIBUTING THIS SOFTWARE OR ITS DERIVATIVES. IN NO EVENT WILL SUN OR ITS LICENSORS BE LIABLE FOR ANY LOST REVENUE, PROFIT OR DATA, OR FOR DIRECT, INDIRECT, SPECIAL, CONSEQUENTIAL, INCIDENTAL OR PUNITIVE DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF THE USE OF OR INABILITY TO USE THIS SOFTWARE, EVEN IF SUN HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

You acknowledge that this software is not designed, licensed or intended for use in the design, construction, operation or maintenance of any nuclear facility.


BROADCOM

Emulex® OneCommand® CNA Manager Command Line Interface for OneConnect® Adapters

User Guide

Version 11.2
December 30, 2016

OCMCLI-OCA-UG112

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

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

Copyright © 2007–2017 by Broadcom. All Rights Reserved.

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

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.

Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Chapter 1: Introduction

The OneCommand® CNA (converged network adapter) Manager command line interface (CLI) is a comprehensive management utility for Emulex® adapters. The CLI provides support for commonly used commands without requiring the installation of the OneCommand CNA Manager graphical user interface (GUI). The OneCommand CNA Manager CLI console application name is `Brcmhbacmd`. At the command line interface, a single operation is performed by entering `brcmhbacmd`, followed by a CLI client command and its possible parameters.

The OneCommand CNA Manager application can be installed on multiple operating systems: Windows, Linux, and Solaris. For VMware ESXi hosts, use the OneCommand CNA Manager application for VMware vCenter. For details, refer to the *Emulex OneCommand CNA Manager for VMware vCenter for OneConnect Adapters User Guide*. You can also manage adapters using the OneCommand CNA Manager CLI on Windows, but you must install and use the appropriate Emulex CIM Provider on those VMware hosts.

NOTE For VMware ESXi hosts, when advanced adapter management capabilities are required (for example, iSCSI Management and port disablement), use the OneCommand CNA Manager for VMware vCenter Server. For more details, refer to the *Emulex OneCommand CNA Manager for VMware vCenter for OneConnect Adapters User Guide*.

This product supports the following Emulex OneConnect converged network Adapters (CNAs):

- OCe11000-series adapters
- OCe14000-series adapters

For supported versions of operating systems and platforms, go to http://www.broadcom.com.

1.1 Abbreviations

API	application programming interface
ARI	alternative routing-ID interpretation
ARP	Address Resolution Protocol
ASIC	application-specific integrated circuit
BIOS	basic input-output system
CHAP	Challenge Handshake Authentication Protocol
CIMOM	CIM Model Object Manager
CIN	chassis internal network
CLI	command line interface
CNA	Converged Network Adapter
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
ETO	extended timeout
FAT	file allocation table
FCF	Fibre Channel over Ethernet Forwarder

FCoE	Fibre Channel over Ethernet
FEC	forward error correction
FIP	FCoE Initialization Protocol
GUI	graphical user interface
HBA	host bus adapter
iBFT	iSCSI boot firmware table
ICMP	Internet Control Message Protocol
IP	internet protocol
ISID	initiator session identifier
iSCSI	Internet Small Computer Systems Interface
ISID	initiator session identifier
iSNS	Internet Storage Name Service
LDAP	Lightweight Directory Access Protocol
LLDP	Link Layer Discovery Protocol
LPVID	logical port VLAN ID
LUN	logical unit number
MAC	Media Access Control
MILI	Management Interface Library
MSI	message signaled interrupt
MTU	maximum transmission unit
NIC	network interface card
NPar	NIC partitioning
NPIV	N_Port_ID Virtualization
NVP	normal velocity of propagation
NVRAM	nonvolatile random access memory
OB	open boot
OS	operating system
PAM	pluggable authentication module
PCI	Peripheral Component Interconnect
PFC	priority flow control
PG	priority group
POST	power-on self-test
PXE	Pre-boot execution Environment
QCN	Quantized Congestion Notification
QoS	quality of service
RoCE	RDMA over Converged Ethernet
RHEL	Red Hat Enterprise Linux
Rx	receive
SAN	storage area network
SCSI	Small Computer Systems Interface
SFCB	Small Footprint CIM Broker
SFP	small form-factor pluggable
SLES	SUSE Linux Enterprise Server

SMB	Server Message Block
SR-IOV	single root I/O virtualization
TCP	Transmission Control Protocol
TDR	time-domain reflectometer
Tx	transmit
UEFI	Unified Extensible Firmware Interface
UFP	Unified Fabric Port
UMC	Universal Multi-Channel
VEPA	virtual Ethernet port aggregator
VLAN	virtual local area network
VLAN ID	VLAN identifier
VM	virtual machine
VPD	vital product data
vPort	virtual port
WWN	World Wide Name
WWNN	World Wide Node Name
WWPN	World Wide Port Name
XML	Extensible Markup Language

1.2 OneCommand CNA Manager Secure Management

OneCommand CNA 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 specify their user name and password to run the OneCommand CNA Manager CLI. When users are authenticated, only they can perform the functions allowed by the OneCommand CNA Manager user group to which they belong. If your systems are running in an LDAP or Active Directory domain, the OneCommand CNA Manager CLI will authenticate the user with those defined in that domain. For Linux and Solaris systems, this is accomplished using PAM.

NOTE OneCommand CNA Manager Secure Management is supported on Linux, Solaris, and Windows, but it is not supported on VMware hosts. For VMware hosts, the CIM credentials are used.

Administrators set up user accounts such that a user belongs to one of the OneCommand CNA Manager user groups. The user groups define the management capabilities for the user. Table 1 defines the OneCommand CNA Manager user groups and each group's management capabilities.

Table 1 Secure Management User Privileges

Group Name	OneCommand CNA 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 can belong to the administrator group or be a root user, they will not have full privileges to run the OneCommand CNA Manager unless they are also a member of the ocmadmin group. Otherwise, when secure management is enabled, a root user or an administrator can only manage local adapters (similar to the ocmlocaladmin user).

Remote management operations between two machines is allowed or denied depending on the OneCommand CNA Manager secure management status of the machines, and the domains to which the machines belong. The following tables (Table 2, Table 3, and Table 4) list the expected behavior for each machine domain condition (assuming appropriate user credentials are used).

Table 2 Active Commands: Machines on Same Domain

	Remote Server (Secure)	Remote Server (Not Secure)
Client (Secure)	Allowed	Denied [1]
Client (Not Secure)	Denied	Allowed

1. Informs you of an unsecured server that you might want to secure.

Table 3 Active Commands: Machines on Different Domain

	Remote Server (Secure)	Remote Server (Not Secure)
Client (Secure)	Denied[1]	Denied [2]
Client (Not Secure)	Denied	Allowed

1. Allowed if the user name and password are the same on both domains.
2. Informs you of an unsecured server that you might want to secure.

Table 4 Passive Commands: Machines on Any Domain

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

1.2.1 OneCommand CNA Manager Secure Management Configuration Requirements

For systems to run in the OneCommand CNA Manager secure management environment, they must be configured to provide the following two capabilities:

- Authentication – On Linux and Solaris, this is accomplished by using the PAM interface and must be configured as follows:
  - On Solaris, place the correct value in the auth section of the `/etc/pam.d/other` file, or its earlier equivalent, `/etc/pam.conf`.

    NOTE For Solaris systems, you must use `useradd -G groupname` for authentication to work. You cannot use a lowercase **g**.
  - On Linux, it is the `/etc/pam.d/passwd` file auth section, or the equivalent.
- User Group Membership – From the host machine, OneCommand CNA Manager Secure Management must be able to access the OneCommand CNA 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 CNA Manager group names are listed with their member users, the system is ready to use OneCommand CNA Manager secure management.

1.3 Secure Management Installation

The enabling or disabling of the Secure Management feature is specified at OneCommand CNA Manager installation time. This can be accomplished either interactively or by using dedicated installation switches on Windows, Linux, and Solaris. On Linux and Solaris, if the OneCommand CNA Manager groups described in the previous section (see Table 1) are not configured on the machine at the time of the OneCommand CNA Manager installation, the installation will fail when the secure management feature is selected.

NOTE Only a user with administrator or root privileges is allowed to enable or disable the secure management feature on a local host machine.

Management mode cannot be used if Secure Management is enabled.

1.3.1 Linux and Solaris

This section describes the Secure Management installation options for Linux and Solaris operating systems.

1.3.1.1 Interactive Installation

Enterprise OneCommand CNA Manager installations performed in interactive mode ask if OneCommand CNA Manager Secure Management mode should be enabled. If the answer is **yes**, the other management mode questions are skipped. If the answer is **no** to the OneCommand CNA Manager Secure Management mode question, then the management mode installation questions follow.

1.3.1.2 Unattended Installation with Install Script Switch Option Support

Enterprise OneCommand CNA Manager installations performed in unattended mode provide a switch option to enable OneCommand CNA Manager Secure Management. If the OneCommand CNA Manager Secure Management switch is not used with the installation, Secure Management is disabled.

1.3.2 Windows

During OneCommand CNA Manager installations performed in interactive mode, you are presented with a management mode window where you can select Secure Management as the management mode.

1.4 Setting Secure Management Mode for Linux and Solaris

To set the secure management mode for the Linux and Solaris operating systems:

1. Log on as root.
2. Set secure management:
   - To set Secure Management Mode for Linux, type the following command:
     ```
     # /usr/sbin/brcmocmanager/set_operating_mode
     ```
   - To set Secure Management Mode for Ubuntu 14, type the following command:
     ```
     # /opt/emulex/brcmocmanager/scripts/set_operating_mode.sh
     ```
   - To set Secure Management Mode for Solaris, type the following command:
     ```
     # /opt/brcmocm/brcmocmanager/set_operating_mode
     ```

   Example

   The following example text is displayed:

   ```
   Do you want to enable Secure Management feature for OneCommand? (s/u)
   The secure management feature requires OneCommand groups be configured on
   the LDAP network or the local host machine to provide for OneCommand
   operation.
   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': s
   ```

1.4.1 Using OneCommand CNA Manager with Secure Management Enabled

To run the OneCommand CNA Manager CLI when Secure Management is enabled, you must include your user name and password each time you type a command.

The syntax for entering your user name and password is the following:

```
brcmhbacmd <m=sec> <u=userid> <p=password> <command>
```

For example

```
>brcmhbacmd m=sec u=jsmith p=password download 00-12-34-56-78-9A
oc11-4.6.96.2.ufi
```

User names and passwords are used to authenticate the commands. After the credentials are authenticated, the OneCommand CNA Manager CLI will determine which one of the four user groups you belong to and will allow command usage as appropriate.

Chapter 2: Installing and Uninstalling the CLI

This chapter details prerequisites and procedures for installing and uninstalling the OneCommand CNA Manager CLI in the following operating systems: Linux, Solaris, and Windows. It also describes the Secure Management capability and the procedure for starting and stopping daemon processes.

2.1 Linux

The following instructions are for installing and uninstalling the OneCommand CNA Manager CLI on Linux operating systems. You can install Linux with or without an existing OneCommand CLI kit. Additionally, you can install the OneCommand CNA Manager CLI for Citrix-based operating systems.

2.1.1 Citrix

Citrix is based on CentOS Linux, however, for the OneCommand CNA Manager CLI, Citrix is more comparable to VMware –, a hypervisor-style server for managing virtual machines. Citrix XenServer 6.5 and 7.0 operating systems require the OneCommand CNA Manager CLI installation.

2.1.2 Installing in Linux without an Existing OneCommand CLI Kit

NOTE For Secure Management, prior to installation, OneCommand CNA Manager groups must be configured on the LDAP network or the local host machine for Secure Management operation. See Section 1.2.1, OneCommand CNA Manager Secure Management Configuration Requirements, for configuration instructions.

2.1.2.1 Linux OneCommand CNA Manager Requirements

For new systems, install the specific Linux driver rpm files before installing the OneCommand CNA Manager CLI.

2.1.2.1.1 Libnl Library

On RHEL 6.x and RHEL 7.x, the OneCommand CNA Manager Core rpm file requires the Libnl library. This library is not installed by default, but it can be obtained from the operating system distribution media.

- For i386 RHEL use the 32-bit libnl library.
- For x86_64 RHEL use the 64-bit libnl library.

2.1.2.1.2 libhbaapi Library

To install the OneCommand CNA Manager CLI in Linux without an existing OneCommand CNA Manager CLI:

1. Copy the applications kit tar file to a directory on the installation machine.
2. Change to the directory where you copied the tar file.
3. Untar the file:

   ```
   tar zxvf brcmocmcore-<supported_os>-<app_ver>-<rel>.tgz
   ```
4. Change to the core kit directory created in step 3.

   ```
   cd brcmocmcore-<supported_os>-<app_ver>-<rel>
   ```
5. Run the `install.sh` script.

   ```
   ./install.sh
   ```

The core kit consists of three or four rpm files for each supported architecture and each supported version of Linux. For example:

- `brcmocmlibhbaapi-*.rpm` (on 64-bit platforms that support 32-bit applications, there are two of these files)
- `brcmocmcore-*.rpm`
- `brcmocmcorelibs-*.rpm`

6. When you are prompted, choose whether to enable Secure Management for OneCommand:

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

If you enter **u**, an additional prompt is given for the management mode:

```
You selected: Secure Management Disabled
Select desired mode of operation for OneCommand CNA Manager:
Enter the number 1, 2, 3, or 4: 1
You selected: 'Local Only Mode'

1 Strictly Local Management:  Only manage the adapters on this host. Management of
                              adapters on this host from other hosts is not allowed.
2 Local Management Plus:      Only manage the adapters on this host. Management of
                              adapters on this host from other hosts is allowed.
3 Full Management:            Manage the adapters on this host and other hosts that
                              allow it. Management of the adapters on this host from
                              another host is allowed.
4 Management Host:            Manage the adapters on this host and other hosts that
                              allow it. Management of the adapters on this host from
                              another host is not allowed.
```

2.1.2.1.3 Unattended Installation

The `install.sh` script can be run in noninteractive (unattended or quiet) mode. Enter the following command to view the syntax:

```
./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.

2.1.3 Installing in Linux with an Existing OneCommand CLI Kit

NOTE The OneCommand CNA Manager core kit cannot be installed if a previous version of the HBAnyware utility is installed.

There are two options for installing the OneCommand CNA Manager CLI on a Linux system with an existing OneCommand CLI kit:

- Updating an existing installation – Preserve existing settings
- Performing a clean install – Overwrite existing settings

2.1.3.1 Updating (Preserving Existing Settings)

To update the OneCommand CNA Manager CLI and preserve settings, you must install the current core kit as detailed in Section 2.1.2, Installing in Linux without an Existing OneCommand CLI Kit. The `.rpm` file handles the configuration file update. The install script executes an rpm file update (`rpm -U *.rpm`) to update the installed version of the core kit to the current version.

NOTE There is no update path from an HBAnyware 4.x or 3.x core kit to a OneCommand CNA Manager core kit. You must uninstall previous versions of the HBAnyware utility before installing a OneCommand CNA Manager core kit. For information on uninstalling older versions of HBAnyware, see Section 2.1.5, Uninstalling Older HBAnyware Kits in Linux.

2.1.3.2 Performing a Clean Install (Removing Existing Settings)

1. Uninstall the existing OneCommand CNA Manager CLI using the uninstall script included in the tar file or in the `/usr/sbin/brcmocmanager/scripts` directory. The configuration files are backed up by rpm with an `.rpmsave` extension.

   For Ubuntu 14, use the uninstall script in the following location:

   ```
   /opt/emulex/brcmocmanager/scripts/uninstall.sh
   ```

   NOTE If an HBAnyware CLI or enterprise kit is installed, follow the procedure in Section 2.1.5, Uninstalling Older HBAnyware Kits in Linux.

2. Install the specific rpm file for your driver for Linux version. For information on installing the rpm file, see Section 2.1.2, Installing in Linux without an Existing OneCommand CLI Kit.

2.1.4 Uninstalling in Linux

To uninstall the OneCommand CNA Manager CLI in Linux:

1. Log on as root.
2. Perform one of the following tasks:
   - Run the `uninstall_brcmocmanager.sh` script located in `/usr/sbin/brcmocmanager/scripts`.
   - Run the `uninstall.sh` script located in the installation tar file.
   - For Ubuntu 14, use the uninstall script in the following location:
     `/opt/broadcom/brcmocmanager/scripts/uninstall.sh`

2.1.5 Uninstalling Older HBAnyware Kits in Linux

2.1.5.1 Uninstalling an Older HBAnyware Core Kit

Run the following command to remove the core kit.

```
rpm -e elxlinuxcorekit
```

2.1.5.1.1 Uninstalling an Older HBAnyware Enterprise Kit

1. Perform one of the following tasks:
   - Run the uninstall script located in `/usr/sbin/hbanyware/scripts` to remove the enterprise kit.
   - Run the uninstall script located in the tar file to remove the enterprise kit.

   If the HBAnyware Security Configurator is installed, you must uninstall it before uninstalling the HBAnyware configuration utility. You must use the uninstall script that shipped with the version of OneCommand Security Configurator you want to remove and proceed to step 2. If the Security Configurator is not installed, proceed to step 3.
2. If the HBAnyware Security Configurator is installed, follow these steps:
   a. Log on as root.
   b. Change to the directory containing the tar file.
   c. Extract the tar file using the `tar -xvf` command.
   d. Change to the newly created directory.
   e. Type the following uninstall script with the ssc parameter specified:
      ```
      ./uninstall ssc
      ```
3. Uninstall the HBAnyware utility and the Application Helper module:
   a. Log on as root.
   b. Change to the directory containing the tar file.
   c. Extract the tar file using the `tar -xvf` command.
   d. Change to the newly created directory.
   e. Uninstall any previously installed versions. Type the following command:
      ```
      ./uninstall
      ```

2.2 Solaris

The following instructions are for installing and uninstalling the OneCommand CNA Manager CLI on Solaris operating systems.

2.2.1 Installing in Solaris

NOTE For Secure Management, prior to installation, OneCommand CNA Manager groups must be configured on the LDAP network or the local host machine for Secure Management operation. See Section 1.2.1, OneCommand CNA Manager Secure Management Configuration Requirements, for configuration instructions.

To install the OneCommand CNA Manager CLI in Solaris:

1. Copy the OneCommand CNA Manager core kit to a temporary directory on the system.
2. Untar the core kit by typing the following command:

   ```
   tar xvf brcmocmcore-solaris-<kit version>.tar
   ```

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

   ```
   cd ./brcmocmcore-solaris-<kit version>/
   ```

4. Run the `install` script and follow the instructions.

   ```
   ./install
   ```

   NOTE The `install` script can also be run in non-interactive (unattended, quiet) mode. Enter the following command to view the syntax:

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

   If any of the following are already present on the system, the `install` script attempts to remove them first:
   - HBAnyware utility
   - OneCommand Manager or OneCommand CNA Manager core kit
   - OneCommand Manager or OneCommand CNA Manager enterprise kit
   - Solaris driver (FCSA) utilities
5. When you are 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 **u** here, an additional prompt is given for the management mode:

   ```
   You selected: Secure Management Disabled
   Select desired mode of operation for OneCommand CNA Manager:
   Enter the number 1, 2, 3, or 4: 1
   You selected: 'Local Only Mode'

   1 Strictly Local Management: Only manage the adapters on this host. Management of
                                adapters on this host from other hosts is not allowed.
   2 Local Management Plus:     Only manage the adapters on this host. Management of
                                adapters on this host from other hosts is allowed.
   3 Full Management:           Manage the adapters on this host and other hosts that
                                allow it. Management of the adapters on this host from
                                another host is allowed.
   4 Management Host:           Manage the adapters on this host and other hosts that
                                allow it. Management of the adapters on this host from
                                another host is not allowed.
   ```

2.2.2 Uninstalling in Solaris

To uninstall the OneCommand CNA Manager CLI in Solaris:

1. Log on as root.
2. Perform one of the following tasks:
   - Run `/opt/brcmocm/scripts/uninstall`.

— Run the `uninstall` script located in the installation tar file.
— Enter the command `pkgrm brcmocmcore`.

NOTE The `uninstall` script can also be run in noninteractive (quiet) mode. Enter the following command to view the syntax:

```
./uninstall --help
```

2.3 VMware ESXi

The OneCommand CNA Manager CLI cannot be run on a VMware ESXi operating system. However, a VMware ESXi host can be accessed remotely from the Windows OneCommand CNA Manager CLI if the Emulex CIM provider is installed on the ESXi host. For instructions on installing Emulex CIM Provider on VMware ESXi operating systems, refer to the *Emulex CIM Provider Package for OneConnect Adapters Installation Guide*.

2.4 Windows

The following instructions are for installing and uninstalling the OneCommand CNA Manager CLI on Windows operating systems. There are two ways to install the OneCommand CNA Manager CLI in Windows:

- Attended installation – You are present during the installation. You are prompted for more information for the installation to continue.
- Unattended installation – You do not need to be present during the installation. Installation will complete on its own. Installation progress can be displayed as an option.

2.4.1 Installing in Windows by Attended Installation

To install the OneCommand CNA Manager CLI, run the `installation.exe` file for a core Windows driver kit that does not include the OneCommand CNA Manager GUI, and follow the installer directions.

Use the following syntax for the installation executable file:

```
brcmocmcore-windows-<arch>-<kit version>.exe
```

- <*arch*> is either x64 or x86.
- <*kit version*> represents the complete kit version.

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

```
brcmocmcore-windows-x64-5.0.2.14-1.exe
```

2.4.2 Installing in Windows by Unattended Installation

To install the OneCommand CNA Manager CLI in Windows unattended:

1. From http://www.broadcom.com, download the x64 or x86 OneCommand CNA Manager Core Kit installation file to your system.

2. Use the following syntax for the installation executable file:

   ```
   brcmocmcore-windows-<arch>-<kit version>.exe <option>
   ```

3. Activate the kit with switch `/q` or `/q2`.
   - The `/q` switch displays progress reports.
   - The `/q2` switch does not display progress reports.
4. Enable Secure Management Mode by adding the `sec=1` argument or disable it by adding `sec=0`. If the sec argument is not entered, Secure Management is disabled by default. See Section 1.2, OneCommand CNA Manager Secure Management, for more information.

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

   ```
   brcmocm-windows-<arch>-kit version>.exe sec=1 /q2
   ```

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

   ```
   brcmocm-windows-<arch>-kit version>.exe sec=0 /q2
   ```

   NOTE There are two management mode defaults for unattended installation:
   - `mmode=3` (Full Management Mode)
   - `achange=1`
5. Select a Management Mode by adding the `mmode` argument, and select the ability to change the Management Mode by adding the `achange` argument with selected values as in the following example.

   NOTE If you enabled Secure Management in step 4 and attempt to enter an `mmode` value, a `conflicting parameters` error can occur.

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

   ```
   brcmocm-windows-x64-5.01.00.10-4.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 will be disabled. If the switch is specified, the value can be a host name or an IP address, which 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 nonmanagement-host user to select whether the OneCommand CNA 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 23333 is used.

If the management host option is selected, you must select the default port number or enter a valid TCP/IP port number on the command line. A value of `0` will not be accepted.

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

2.4.3 Uninstalling in Windows

You can uninstall the OneCommand CNA Manager CLI in Windows in one of two ways:

- Through the Control Panel
- Through the command line

2.4.3.1 Uninstalling through the Control Panel

To uninstall the OneCommand CNA Manager CLI in Windows through the Control Panel:

1. In the Control Panel, select **Programs and Features**.
2. If present, select **Emulex OC CNA Manager CLI** [*version*], and click **Uninstall/Change;** you are prompted to continue. Click **Yes**.

The OneCommand CNA Manager CLI components are removed from the system.

2.4.3.2 Uninstalling through the Command Line

To uninstall the OneCommand CNA Manager CLI in Windows through the command line:

1. Change to the appropriate uninstall directory:
   ```
   cd <Install Location>\broadcom\Util\Uninstall
   ```
2. Type the following command:
   ```
   uninstall_BrcmOCManager_Core.bat
   ```

2.5 Starting and Stopping Daemon Processes for Linux and Solaris Installations

On Linux and Solaris machines, you can stop and start the OneCommand CNA Manager daemon processes using the `stop_brcmocmanager` and `start_brcmocmanager` scripts, respectively. These are found in the following OneCommand CNA Manager installation directories:

- Linux – `/usr/sbin/brcmocmanager`
- Ubuntu 14 – `/opt/broadcom/brcmocmanager/scripts`:
  — `stop_brcmocmanger.sh`
  — `start_brcmocmanager.sh`
- Solaris – `/opt/brcmocm`

The `brcmhbamgrd` daemon process (included with OneCommand CNA Manager CLI) is affected by these scripts. It is a remote management daemon that services requests from OneCommand CNA Manager clients running on remote host machines.

The daemon processes start at system boot time.

Chapter 3: Updating to the OneCommand CNA Manager Application Enterprise Kit

NOTE The full-featured OneCommand CNA Manager application enterprise kit is not supported on Citrix XenServer 6.x, Citrix XenServer 7.x, or VMware ESXi server.

This chapter details procedures for updating the OneCommand CNA Manager CLI to the OneCommand CNA Manager application enterprise kit in Linux, Solaris, and Windows operating systems. An update can be performed only if the version of the OneCommand CNA Manager application enterprise kit is the same or later than the OneCommand CNA Manager CLI version.

NOTE You cannot update a OneCommand CNA Manager CLI with a previous version of the OneCommand CNA Manager application enterprise kit.

3.1 Linux

To update from the OneCommand CNA Manager CLI to the full-featured OneCommand CNA Manager enterprise kit in Linux, run the `install.sh` script of the OneCommand CNA Manager application enterprise kit.

The install script executes an rpm file update (`rpm -U *.rpm`) to update the installed core kit to an enterprise kit.

3.2 Solaris

To update from the OneCommand CNA Manager CLI to the full-featured OneCommand CNA Manager enterprise kit in Solaris:

1. Download the OneCommand CNA Manager enterprise kit to a temporary directory on your system.
2. Untar the OneCommand CNA Manager enterprise kit tar file:

   ```
   tar xvf brcmocm-solaris-<kit version>.tar
   ```
3. Change to the newly created `brcmocm-<kit version>` directory:

   ```
   cd ./brcmocm-solaris-<kit version>/
   ```
4. Run the `install` script and follow the instructions:

   ```
   ./install
   ```

The `install` script can also be run in noninteractive (quiet) mode. To view the syntax, type the following command:

```
/install --help
```

3.3 Windows

To update from the OneCommand CNA Manager CLI to the full-featured OneCommand CNA Manager enterprise kit in Windows:

From the desktop, run the `brcmocm-windows-<kit version>.exe` file that contains the full application kit. Running this executable file removes the OneCommand CNA Manager CLI and installs a full-featured version of the OneCommand CNA Manager application that includes the CLI and the GUI.

Chapter 4: CLI Client Command Usage

The CLI Client component of the OneCommand CNA Manager application provides access to the capabilities of the Remote Management library or the CIM interface from a console command prompt to get the management information.

4.1 Overview

The CLI Client is intended for use in command shells or scripted operations from within shell scripts or batch files. The CLI Client is a console application named `BrmcHbaCmd`. A single operation is performed by typing `brcmhbacmd` at the command line, followed by a CLI client command and its possible parameters. For example:

```
brcmhbacmd [cli options]<command> [parameters]
```

The CLI options are specified for running the CLI commands to remote hosts or with Secure Management.

When the specified operation is completed, the command prompt is displayed. For a majority of commands, the first parameter following the command is the WWPN or MAC address of the port that the command is to act upon.

4.1.1 CLI in Read-Only Mode

The CLI does not allow the execution of some commands if it is configured for read-only mode. The following error message is returned if such a command is attempted:

```
Error: Read-only management mode is currently set on this host. The
requested command is not permitted in this mode.
```

4.2 BrcmHbaCmd Syntax Usage

The following syntax rules and usage apply to the BrcmHbaCmd application:

- Parameters denoted within angle brackets < > are required.
- Parameters denoted within square brackets [ ] are optional.
- For Linux and Solaris, (which are case-sensitive), program names must be in lowercase letters. Therefore, the command line must begin with brcmhbacmd (rather than BrcmHbaCmd). Windows is not case-sensitive, so the program name is not required to be in all lowercase letters.
- To run the command on a remote host, an IP address or host name must be specified using the h option with the following syntax:

  ```
  brcmhbacmd [h=IP_Address[:port] | Hostname[:port]] <command> [parameters]
  ```

  — If the h option is omitted, the command is run on the local host.

  — If the h option is specified, the command is sent to the specified remote host (assuming it is specified correctly, the remote host is up, and the remote host is running the OneCommand CNA Manager remote management agent.

  — The *`:port`* option is optional. If it is omitted, the OneCommand CNA Manager remote management protocol uses the default TCP port. If it is specified, the OneCommand CNA Manager remote management protocol uses the user-specified TCP port.

  — **Examples**

  Using the IP address:

  ```
  brcmhbacmd h=138.239.91.121 ListHBAs
  ```

Using the host name:

```
brcmhbacmd h=cp-hp5670 ListHBAs
```

— The h option is available for all commands except for the AddHost, RemoveHost, and the Version commands.

- For FCoE functions, the WWPN of the adapter must be specified. Where the WWPN is specified, each pair of numbers within the WWPN is separated by colons (:) or spaces ( ). If space separators are used, the entire WWPN must be enclosed in quotation marks (“ ”).

  For example, the following command displays the port attributes for the adapter with the specified WWPN:

  ```
  brcmhbacmd PortAttributes 10:00:00:00:c9:20:20:20
  ```

- For iSCSI and NIC functions, the MAC address must be specified. Where a MAC address is specified, each pair of numbers within the MAC address is separated by a dash (-).

  For example, the following command sets the target properties for the iSCSI port with the specified MAC address with an extended timeout value of 1:

  ```
  brcmhbacmd SetTargetProperties 00-11-22-33-44-55 iscsiTarget 1
  ```

- For NIC functions, only the permanent MAC address is supported for the port address parameter on a BrcmHbaCmd command line.

  Normally, for a NIC function, the functions's permanent MAC address and current MAC address parameters are equal. However, it is possible to set a user-specified (current) MAC address that is different from the permanent MAC address. Also, for some implementations, it is possible to have multiple NIC functions with the same current MAC addresses, but with unique permanent MAC addresses. Therefore, to be sure that the OneCommand CNA Manager application can access the correct function, only the permanent MAC address is supported.

NOTE Both the permanent MAC address and the current MAC address are displayed by using the `ListHBAs` command. See Section 5.13.9, ListHBAs.

4.3 Secure Management CLI Interface

The Secure Management CLI interface is supported by the Linux, Solaris, and Windows operating systems.

NOTE Users with root or Administrator privileges on the local machine will retain full OneCommand CNA Manager CLI configuration capability without the use of credentials (local machine only).

4.3.1 Device Management Using the Secure Management Interface

To run the BrcmHbaCmd CLI client application when the Secure Management feature is enabled, each invocation must include a user name and password. The user name and password options are added to the existing BrcmHbaCmd command in the same way as they are for CIM commands, except the <m=CIM> option is replaced by the <m=sec> option (to distinguish it from a CIM command). For example:

Without Secure Management (or if running as root or administrator):

```
brcmhbacmd  <cmd>
```

With Secure Management (as non-root or non-administrator user):

```
brcmhbacmd  <m=sec> <u=userid> <p=password> <cmd>
```

4.3.2 Syntax Rules for the Secure Management Interface

For the secure management interface, all of the syntax rules in Section 4.2, BrcmHbaCmd Syntax Usage, apply.

Example

In Windows, to download firmware on an adapter managed on a remote host at IP address 192.168.1.122 using the Secure Management interface, run the following command:

```
brcmhbacmd h=192.168.1.122 m=sec u=jsmith p=password download 00-12-34-56-78-9A
oc14-11.2.123.45.ufi
```

4.4 CIM Client Interface

NOTE In Linux and Solaris, you cannot use `brcmHbaCmd` as a CIM client.

4.4.1 Device Management Using the CIM Interface

VMware on the hypervisor-based ESXi platforms use the CIM as the only standard management mechanism for device management.

For VMware ESXi hosts, you can manage adapters using the OneCommand CNA Manager CLI on Windows, but you must install and use the appropriate Emulex CIM Provider on the VMware ESXi host. For installation, refer to the *Emulex CIM Provider Package for OneConnect Adapters Installation Guide*.

NOTE For VMware ESXi hosts, if advanced adapter management capabilities are required (such as iSCSI management), use the OneCommand CNA Manager for VMware vCenter Server. For more details, refer to the *Emulex OneCommand CNA Manager for VMware vCenter for OneConnect Adapters User Guide*.

4.4.2 Syntax Rules for the CIM Interface

For the CIM interface, all the syntax rules in section Section 4.2, BrcmHbaCmd Syntax Usage, apply, except that the h option is required. Additionally, the `m=cim` parameter is required in the command line for getting the data from the ESXi host. For example:

```
brcmhbacmd h=192.168.1.110 m=cim u=root p=password n=root/brcmccx listhbas
```

4.4.2.1 Syntax Options and Setting CIM Credentials

For issuing CIM-based commands, two main syntax options are available.

Option A

```
brcmhbacmd <h=IP_Address[:port]> m=cim [u=userid] [p=password]
[n=root/brcmccx] <command> <WWPN>
```

Option B

```
brcmhbacmd <h=IP_Address[:port]> <m=cim> <command>
```

Before using the option B syntax, you must set the CIM credentials. Perform one of the following tasks:

- Set the default CIM credentials using the `SetCimCred` command (see Section 5.13.13, SetCimCred). This command sets only the CIM credentials. After you have set them, subsequent `BrcmHbaCmd` commands do not require you to specify the CIM credentials on the command line.

  Command syntax:

  ```
  brcmhbacmd SetCimCred <username> <password> <namespace> <portnum>
  ```

- Add the host IP address with CIM credentials using the `AddHost` command.

  Command syntax:

  ```
  brcmhbacmd <m=cim> [u=userid] [p=password] [n=namespace] AddHost
  <IP_Address>
  ```

4.4.2.1.1 Default CIM Credentials

If you specify the command with the CIM method `m=cim` without specifying the CIM credentials (*`userid`*, *`password`*, or *`namespace`*), the default value for the missing CIM credential is obtained in the following order:

1. The information entered using the `addhost` command is looked up.
2. If no values exist, the information entered using the `setcimcred` command is used.
3. If no values exist, the following defaults are used:

   ```
   username=root
   password=root
   namespace=root/brcmccx
   portnum=5988
   ```

4.4.2.2 Example of Using the CIM Interface to Display Adapters

In Windows, to display a list of adapters managed for a specified host using the CIM interface, run the following command:

```
brcmhbacmd h=10.192.113.128 m=cim u=root p=root n=root/brcmccx listhbas
```

For a list of BrcmHbaCmd commands supported through the CIM interface, see Table 6, CLI Client Command Reference.

Chapter 5: CLI Client Command Descriptions

CLI Client commands are organized by command groups. Two tables are presented for your convenience; a table organized by command group and another by alphabetically listing CLI client commands.

The following table shows each command group with a short description and the commands in each group. After you determine the command group of interest, click on the command link and go directly to the command you selected.

Table 5 CLI Client Command Reference Functional Groups

Command Group	Description	Commands
Adapter License Management Commands	This group manages the adapter licensing. You can use these commands to install license keys to enable functionality and list current licensed functionality.	InstallAdapterLicense ShowAdapterLicenseFeatures ShowLicenseAdapterID
Attributes Commands	This group manages the display of adapter, port, server attributes, and port statistics for each adapter specified. You can also set the port speed on OCe11100-series and OCe14000-series adapters.	HbaAttributes ServerAttributes SetPhyPortSpeed SetPortEnabled
Boot Commands	This group manages the commands that enable or disable the network boot for NIC ports or the boot code for FCoE adapter ports. You can also show and change FCoE boot parameters.	EnableBootCode GetBootParams SetBootParam
Channel Management Commands	These commands enable and disable channel management and set channel properties for CNAs. The OCe14000-series adapters support NPar, a form of channel management.	CMGetParams CMMode CMSetBW CMSetLPVID
DCB Commands	These commands display and set the DCB and LLDP parameters for iSCSI, FCoE, and NIC adapter ports.	GetDCBParams GetPGInfo SetCnaPGBW SetDCBParam SetDCBPriority
Diagnostic Commands	This group provides commands that enable you to detect cabling problems, to examine transceiver data, and to flash memory load lists. Additionally, you can run specific diagnostic tests, such as the Loopback test and the POST test.	GetBeacon GetXcvrData LoopBackTest PciData SetBeacon SetCableNVP TDRTest
Driver Parameter Commands	Use the driver parameter commands to show, set, and save the driver parameter values. You can also change the parameters back to factory default values. Driver Parameter commands are supported only for FCoE ports.	DriverConfig GetDriverParams GetDriverParamsGlobal SaveConfig SetDriverParam SetDriverParamDefaults

Table 5 CLI Client Command Reference Functional Groups (Continued)

Command Group	Description	Commands
Dump Commands	Use the diagnostic dump feature to create a dump file for a selected adapter. Dump files contain information such as firmware version, driver version, and operating system information. This information is useful for troubleshooting an adapter, but it is unavailable in read-only mode.	DeleteDumpFiles Dump GetDumpDirectory GetDumpFile GetDumpFileNames GetRetentionCount SetDumpDirectory SetRetentionCount
FCoE Commands	This group of commands manages the FIP parameters and displays the FCF for an adapter in the FCoE mode.	GetFCFInfo GetFIPParams SetFIPParam

Table 5 CLI Client Command Reference Functional Groups (Continued)

Command Group	Description	Commands
iSCSI Commands	The commands in this section support the iSCSI interface. iSCSI commands are supported only on OneConnect iSCSI ports.	AddARPTableEntry AddiSNSServer AddRouteTableEntry AddTarget AddTargetPortal CleariSNSServer DelARPTableEntry DeleteiSNSServer DelRouteTableEntry DiscoveriSNSServer EraseiSCSIConfig ExportiSCSI GetInitiatorProperties GetiSCSILuns GetiSCSIPortStats GetNetworkConfiguration GetSessionInfo ImportiSCSI iSCSIPing ListSessions RemoveTarget RemoveTargetPortal SetBootTargetSession SetInitiatorProperties SetiSCSIBoot SetNetworkConfiguration SetTargetLoginProperties SetTargetProperties SetTPLoginProperties ShowARPTable ShowiSNSServer ShowRouteTable ShowTarget ShowTargetPortal TargetLogin TargetLogout UpdateiSNSServer
LUN Masking Commands	The commands in this group manage LUN masking activities. LUN masking is supported only for FCoE ports.	GetLunList GetLunUnMaskByHBA GetLunUnMaskByTarget RescanLuns SetLunMask

Table 5 CLI Client Command Reference Functional Groups (Continued)

Command Group	Description	Commands
Miscellaneous Commands	This group contains commands that do not belong in other groups.	AddHost CnaClearEventLog CnaGetEventLog Download ExportSANInfo GetCimCred GetQoSInfo GetVPD ListHBAs ListVFuncs Reset SetCimCred SRIOVEnable TargetMapping VEPAEnable Version
Persistent Binding Commands	This group of commands facilitates persistent binding operations. These commands are supported only for FCoE ports.	AllNodeInfo BindingCapabilities BindingSupport PersistentBinding RemoveAllPersistentBinding RemovePersistentBinding SetBindingSupport SetPersistentBinding
Personality Change Commands	This group of commands changes the personality or protocol running on OneConnect adapters.	ChangePersonality ShowPersonalities
Profile Management Commands	This group of commands manages profile configuration for OCe14000-series adapters. You can display active and reboot port configurations for an adapter, list available profile IDs, and configure the function protocol for all ports on OCe14000-series adapters.	GetAdapterPortConfig GetLinkConfig ListProfiles SetAdapterPortConfig (for OCe14000-Series Adapters) SetLinkConfig

Table 5 CLI Client Command Reference Functional Groups (Continued)

Command Group	Description	Commands
UMC Commands	UMC commands allow viewing of the UMC configuration, enabling and disabling of the UMC at the adapter level, and the modification of some of the channel properties. The UMC commands cannot be used to manage other channel management types.	UmcEnable UmcGetParams UmcSetBW UmcSetLPVID
vPort Commands	vPort commands manage virtual ports and functions only on FCoE adapters. In Linux, VPorts do not persist across system reboots.	CreateVPort DeleteVPort ListVPorts VPortTargets
WWN Management Commands	WWN management validates WWNs to avoid WWPN duplication; however, WWNN duplication is acceptable. You might see error and warning messages if a name duplication is detected. Make sure that the activation requirement is fulfilled after each WWN is changed or restored. If pending changes exist, some diagnostic and maintenance features are not available. WWN management commands are available only on FCoE adapters.	ChangeWWN GetWWNCap ReadWWN RestoreWWN

Table 6 lists each command alphabetically and shows the operating system and CIM Interface support for each command. A linked section number for each command is provided for your convenience. A check mark (✓) designates a supported command for a particular operating system and CIM interface.

NOTE For VMware ESXi, two options supprt the CLI:

- Using the OneCommand CNA Manager CLI on Windows with the appropriate Emulex CIM Provider installed on a VMware host. These commands are covered in this section. iSCSI management commands are not supported by this option.
- Using the OneCommand CNA Manager for VMware vCenter command line interface (`brcmvcpcmd`). Although the available commands are listed in Table 6 for your convenience, see the *Emulex OneCommand CNA Manager for VMware vCenter for OneConnect Adapters User Guide* for specific information.

Table 6 CLI Client Command Reference

Command	Linux		Solaris	Windows	CIM Interface Support	brcmvcpcmd	Section
	RHEL, SLES, Ubuntu, and Oracle	Citrix					
`AddARPTableEntry`	✓			✓		✓	5.11.1
`AddHost`	✓	✓	✓	✓	✓		5.13.1
`AddiSNSServer`	✓	✓		✓		✓	5.11.2
`AddRouteTableEntry`	✓	✓		✓		✓	5.11.3
`AddTarget`	✓	✓		✓		✓	5.11.4
`AddTargetPortal`	✓	✓		✓		✓	5.11.5
`AllNodeInfo`	✓	✓	✓	✓	✓	✓	5.14.1
`BindingCapabilities`			✓	✓			5.14.2
`BindingSupport`			✓	✓			5.14.3
`ChangePersonality`	✓	✓	✓	✓	✓	✓	5.15.1

Table 6 CLI Client Command Reference (Continued)

Command	Linux		Solaris	Windows	CIM Interface Support	brcmvcpcmd	Section
	RHEL, SLES, Ubuntu, and Oracle	Citrix					
ChangeWWN	✓	✓	✓	✓	✓	✓	5.19.1
ClearSNSServer	✓	✓		✓		✓	5.11.6
CMGetParams	✓	✓	✓	✓	✓	✓	5.5.1
CMMode	✓	✓	✓	✓	✓	✓	5.5.2
CMSetBW	✓	✓	✓	✓	✓	✓	5.5.3
CMSetLPVID	✓	✓	✓	✓	✓	✓	5.5.3
CnaClearEventLog	✓	✓		✓		✓	5.13.2
CnaGetEventLog	✓	✓		✓		✓	5.13.3
CreateVPort	✓		✓	✓			5.18.1
DelARPTableEntry	✓	✓		✓		✓	5.11.7
DelRouteTableEntry	✓	✓		✓		✓	5.11.9
DeleteDumpFiles	✓	✓	✓	✓	✓	✓	5.9.1
DeleteiSNSServer	✓	✓		✓		✓	5.11.8
DeleteVPort	✓		✓	✓			5.18.2
DiscoveriSNSServer	✓	✓		✓		✓	5.11.10
Download	✓	✓	✓	✓	✓	✓	5.13.4
DriverConfig	✓	✓		✓		✓	5.8.1
Dump	✓	✓	✓	✓	✓	✓	5.9.2
EnableBootCode	✓	✓	✓	✓	✓	✓	5.4.1
EraseiSCSIConfig	✓	✓	✓	✓			5.11.11
ExportSANInfo	✓	✓	✓	✓			5.13.5
ExportiSCSI				✓			5.11.12
GetBeacon	✓	✓	✓	✓	✓	✓	5.7.1
GetAdapterPortConfig	✓	✓	✓	✓	✓	✓	5.16.1
GetBootParams	✓	✓	✓	✓		✓	5.4.2
GetCimCred				✓			5.13.6
GetDCBParams	✓	✓	✓	✓	✓	✓	5.6.1
GetDriverParams	✓	✓	✓	✓	✓	✓	5.8.2
GetDriverParamsGlobal	✓	✓	✓	✓	✓	✓	5.8.3
GetDumpDirectory	✓	✓	✓	✓	✓	✓	5.9.3
GetDumpFile	✓	✓	✓	✓	✓	✓	5.9.4
GetDumpFileNames	✓	✓	✓	✓	✓	✓	5.9.5
GetFCFInfo	✓	✓	✓	✓	✓	✓	5.10.1
GetFIPParams	✓	✓	✓	✓	✓	✓	5.10.2
GetInitiatorProperties	✓	✓		✓		✓	5.11.13
GetiSCSILuns	✓	✓		✓		✓	5.11.14
GetiSCSIPortStats	✓	✓		✓		✓	5.11.15
GetLinkConfig	✓	✓	✓	✓	✓	✓	5.16.2

Table 6 CLI Client Command Reference (Continued)

Command	Linux		Solaris	Windows	CIM Interface Support	brcmvcpcmd	Section
	RHEL, SLES, Ubuntu, and Oracle	Citrix					
GetLunList	✓		✓	✓	✓	✓	5.12.1
GetLunUnMaskByHBA				✓			5.12.2
GetLunUnMaskByTarget				✓			5.12.3
GetNetworkConfiguration	✓	✓		✓		✓	5.11.16
GetPGInfo	✓	✓	✓	✓	✓	✓	5.6.2
GetQosInfo	✓	✓	✓	✓	✓	✓	5.13.7
GetRetentionCount	✓	✓	✓	✓	✓	✓	5.9.6
GetSessionInfo	✓	✓		✓		✓	5.11.17
GetVPD	✓	✓	✓	✓	✓	✓	5.13.8
GetWWNCap	✓	✓	✓	✓	✓	✓	5.19.2
GetXcvrData	✓	✓	✓	✓	✓	✓	5.7.2
HbaAttributes	✓	✓	✓	✓	✓	✓	5.3.1
Help	✓	✓	✓	✓	N/A	✓	5.1
ImportiSCSI				✓			5.11.18
InstallAdapterLicense	✓	✓	✓	✓	✓	✓	5.2.1
iSCSIPing	✓	✓		✓		✓	5.11.19
ListHBAs	✓	✓	✓	✓	✓	✓	5.13.9
ListProfiles	✓	✓	✓	✓	✓	✓	5.16.3
ListSessions	✓	✓		✓		✓	5.11.20
ListVFuncs	✓	✓	✓	✓	✓	✓	5.18.3
ListVPorts	✓	✓	✓	✓	✓		5.18.3
LoopBackTest	✓	✓	✓	✓	✓	✓	5.7.3
PciData	✓	✓	✓	✓	✓		5.7.4
PersistentBinding		✓	✓	✓		✓	5.14.4
PortAttributes	✓	✓	✓	✓	✓		5.3.2
PortStatistics	✓	✓	✓	✓	✓	✓	5.3.3
ReadWWN	✓	✓	✓	✓	✓	✓	5.19.3
RemoveAllPersistentBinding			✓	✓		✓	5.14.5
RemovePersistentBinding			✓	✓			5.14.6
RemoveHost	✓	✓	✓	✓	✓		5.13.10
RemoveTarget	✓	✓		✓			5.11.21
RemoveTargetPortal	✓	✓		✓		✓	5.11.22
RescanLuns	✓		✓	✓		✓	5.12.4
Reset	✓	✓	✓	✓	✓		5.13.12
RestoreWWN	✓	✓	✓	✓	✓	✓	5.19.4
SaveConfig	✓	✓	✓	✓		✓	5.8.4
ServerAttributes	✓	✓	✓	✓	✓	✓	5.3.4
SetAdapterPortConfig	✓	✓	✓	✓	✓	✓	5.16.4

Table 6 CLI Client Command Reference (Continued)

Command	Linux		Solaris	Windows	CIM Interface Support	brcmvcpcmd	Section
	RHEL, SLES, Ubuntu, and Oracle	Citrix					
SetBeacon	✓	✓	✓	✓	✓		5.7.1
SetBindingSupport			✓	✓		✓	5.14.7
SetBootParam	✓	✓	✓	✓			5.4.3
SetBootTargetSession	✓	✓		✓		✓	5.11.23
SetCableNVP	✓	✓	✓	✓		✓	5.7.6
SetCnaPGBW	✓	✓	✓	✓	✓	✓	5.6.3
SetCimCred				✓		✓	5.13.13
SetDCBParam	✓	✓	✓	✓	✓		5.6.4
SetDCBPriority	✓	✓	✓	✓	✓	✓	5.6.5
SetDriverParam	✓	✓	✓	✓	✓	✓	5.8.5
SetDriverParamDefaults	✓	✓	✓	✓		✓	5.8.6
SetDumpDirectory					✓	✓	5.9.7
SetFIPParam	✓	✓	✓	✓	✓	✓	5.10.3
SetInitiatorProperties	✓	✓		✓		✓	5.11.24
SetiSCSIBoot	✓	✓		✓		✓	5.11.25
SetLinkConfig	✓	✓	✓	✓	✓	✓	5.16.7
SetLunMask				✓		✓	5.12.5
SetNetworkConfiguration	✓	✓		✓			5.11.26
SetPersistentBinding			✓	✓			5.14.8
SetPhyPortSpeed	✓	✓	✓	✓	✓	✓	5.3.5
SetPortEnabled	✓	✓	✓	✓	✓	✓	5.3.6
SetRetentionCount	✓	✓	✓	✓	✓	✓	5.9.8
SetTargetLoginProperties	✓	✓		✓		✓	5.11.27
SetTargetProperties	✓	✓		✓		✓	5.11.28
SetTPLoginProperties	✓	✓		✓		✓	5.11.29
ShowAdapterLicenseFeatures	✓	✓	✓	✓	✓	✓	5.2.2
ShowARPTable	✓	✓		✓		✓	5.11.30
ShowiSNSServer	✓	✓		✓		✓	5.11.31
ShowLicenseAdapterID	✓	✓	✓	✓	✓	✓	5.2.3
ShowPersonalities	✓	✓	✓	✓	✓	✓	5.15.2
ShowRouteTable	✓	✓		✓		✓	5.11.32
ShowTarget	✓	✓		✓		✓	5.11.33
ShowTargetPortal	✓	✓		✓		✓	5.11.34
SRIOVEnable	✓	✓	✓	✓	✓	✓	5.13.14
TargetLogin	✓	✓		✓		✓	5.11.35
TargetLogout	✓	✓		✓		✓	5.11.36
TargetMapping	✓	✓	✓	✓	✓	✓	5.13.15
TDRTest	✓	✓	✓	✓		✓	5.7.7

Table 6 CLI Client Command Reference (Continued)

Command	Linux		Solaris	Windows	CIM Interface Support	brcmvcpcmd	Section
	RHEL, SLES, Ubuntu, and Oracle	Citrix					
UmcEnable	✓	✓	✓	✓	✓	✓	5.17.1
UmcGetParams	✓	✓	✓	✓	✓	✓	5.17.2
UmcSetBw	✓	✓	✓	✓	✓	✓	5.17.3
UmcSetLPVID	✓	✓	✓	✓	✓	✓	5.17.4
UpdateiSNSServer	✓	✓		✓		✓	5.11.37
VEPAEnable	✓	✓	✓	✓	✓	✓	5.13.16
Version	✓	✓	✓	✓	✓	✓	5.13.17
VPortTargets	✓	✓	✓	✓	✓	✓	5.18.4

5.1 Help

This command displays command information for the BrcmHbaCmd application. Without using its optional parameters, the Help command lists all the commands in their respective groups. Using the optional parameter GroupName, it lists the commands in a group. Using the optional parameter CmdName, it shows the details for a specific command.

Supported By

Linux, Solaris, and Windows

Syntax

```
Help [GroupName] [CmdName]
```

Parameters

GroupName This optional parameter lists the commands in a particular group.

CmdName This optional parameter shows the details for a particular CLI command.

Examples

This Help command example lists all the commands in their respective groups:

```
brcmhbacmd help
```

This Help command example shows the details for the SetDCBParam command:

```
brcmhbacmd help setdcbparam
```

5.2 Adapter License Management Commands

The Adapter License Management Group manages adapter licensing. Use these commands to install license keys to enable functionality and to list current licensed functionality.

For these commands, the *WWPN* or *MAC* address argument specifies the adapter the command is acting upon. The `BrcmHbaCmd` application uses the WWPN or MAC address to identify the adapter, but this does not mean that the command works successfully on the specified port.

These commands are supported on ESXi systems only if the CIM provider is used on a remote system.

NOTE Adapter License Management commands are supported only on some OCe11100-series adapters. They are not available on OCe11101-EM/EX, OCe11102-EM/EX, or OCe14000-series adapters. If license management commands are attempted with unsupported adapters, the following error message is returned:

```
There are no license features for this adapter
```

5.2.1 InstallAdapterLicense

This command installs the license keys from a license file to enable specific features on the adapter.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
InstallAdapterLicense <WWPN|MAC> <LicenseFile>
```

Parameters

WWPN	The WWPN of an FCoE port.
MAC	The MAC address of a NIC or iSCSI port.
License File	The path to the license key file containing the license keys obtained from the License area of http://www.broadcom.com.

Examples

For non-VMware ESXi hosts:

```
brcmhbacmd InstallAdapterLicense 00-12-34-56-78-9A K:\lf1324.lic
```

For VMware ESXi hosts:

```
brcmhbacmd h=<IP_Address> m=cim u=root p=<password> n=<namespace>
InstallAdapterLicense 00-12-34-56-78-9A K:\lf1324.lic
```

5.2.2 ShowAdapterLicenseFeatures

This command displays the licensed and licensable features and the features that are already licensed. The output is a list of features with an indication of whether the feature has been licensed.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
ShowAdapterLicenseFeatures <WWPN|MAC>
```

Parameters

WWPN	The WWPN of an FCoE port.

MAC The MAC address of a NIC or iSCSI port.

Examples

For non-VMware ESXi hosts:

```
brcmhbacmd ShowAdapterLicenseFeatures  00-12-34-56-78-9A
```

For VMware ESXi hosts:

```
brcmhbacmd h=<IP_Address> m=cim u=root p=<password> n=<namespace>
ShowAdapterLicenseFeatures  00-12-34-56-78-9A
```

5.2.3 ShowLicenseAdapterID

This command returns the adapter ID used for enabling licensed features. The adapter ID and the entitlement code are used to obtain license keys that enable various features on the adapter.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
ShowLicenseAdapterID <WWPN|MAC>
```

Parameters

WWPN The WWPN of an FCoE port.

MAC The MAC address of a NIC or iSCSI port.

Examples

For non-VMware ESXi hosts:

```
brcmhbacmd ShowLicenseAdapterID 00-12-34-56-78-9A
```

For VMware ESXi hosts:

```
brcmhbacmd h=<IP_Address> m=cim u=root p=<password> n=<namespace>
ShowLicenseAdapterID 00-12-34-56-78-9A
```

5.3 Attributes Commands

The Attributes Command group manages the display of adapter, port, function, server attributes, and port statistics for each adapter specified. It also enables you to set the port speed.

5.3.1 HbaAttributes

This command shows a list of all adapter attributes for the adapter. The type of information listed might vary according to the adapter model.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
HbaAttributes <WWPN|MAC>
```

Parameters

WWPN	The WWPN of an FCoE function.
MAC	The MAC address of a NIC function.

Example

```
HBA Attributes for 00-90-fa-30-90-29

Host Name                       : dhcp-10-192-81-61
Manufacturer                    : Emulex Corporation
Serial Number                   : US0CG7YT0000039N000MX03
Model                           : OCe14102-U1-D
Model Desc                      : Emulex OneConnect OCe14102-U1-D 2-port PCIe 10GbE
CNA
HW Version                      : E4 B0
FW Version                      : 11.2.123.456
Vendor Spec ID                  : 10DF
Number of Ports                 : 6
Driver Name                     : be2net.ko
Driver Version                  : 11.2.123.456
Device ID                       : 0720
Operational FW                  : 11.2.123.456
Network Boot                    : Not Available
Available Network Boot Methods: PXE
Boot Version                    : 11.2.123.456
Board Temperature               : Normal
Function Type                   : NIC
Sub Device ID                   : E820
Port Number                     : 2
PCI Bus Number                  : 5
PCI Func Number                 : 7
Sub Vendor ID                   : 10DF
Firmware Status                 : Working
IPL Filename                    : CS2FDEL3
NCSI Version                    : N/A
Start-up Boot Code Version      : 2.0.281.768
FCoE Universal Version          : N/A
FCoE x86 BIOS Version           : 11.2.123.456
FCoE EFI Version                : 11.2.123.456
FCoE FCODE Version              : 11.2.123.456
UEFI NIC Version                : 11.2.123.456
NIC FCODE Version               : 11.2.123.456
UEFI iSCSI Version              : 11.2.123.456
PCI Express Link Speed          : 5GT/s
PCI Express Bus Width           : x4
Flash Firmware Version          : 11.2.123.456
```

5.3.2 PortAttributes

This command shows a list of attributes for the adapter-specified function. The type of information listed might vary according to the adapter model and can include IPv4 and IPv6 addresses.

NOTE If VEPA is supported on your adapter, the VEPA state is displayed as enabled or disabled.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
PortAttributes <WWPN|MAC>
```

Parameters

WWPN	The WWPN of an FCoE function.
MAC	The MAC address of a NIC or iSCSI function.

NOTE For Dell adapters, this command displays virtual WWPNs and WWNNs for each FCoE function, and virtual MAC addresses for each NIC function.

Example

```
brcmhbacmd h=10.192.78.68 portattributes 00-9 0-fa-30-44-84
Port Attributes for 00-90-fa-30-44-84
Port MAC:00-90-fa-30-44-84
Permanent MAC:00-90-fa-30-44-84
Port State:Operational
Interface Type:10GB SFP Plus
Auto-Neg Supported Speeds:Not Supported
Force Supported Speeds:1Gb,10Gb
Configured Port Speed:All Supported Speeds
Port Speed Mode:Default
DAC Cable Length:0 meters
Maximum MTU:9000
Current MTU:1500
Function Type:NIC
Function Port State:Up
Function Port Speed:10 GBit/sec
Interface Name:\Device\NTPNP_PCI0114
Minimum Bandwidth:10 GBit/sec
Maximum Bandwidth:10 GBit/sec
IPv4 Address:196.1.8.1
Subnet Mask:255.255.0.0
IPv4 Gateway Address:0.0.0.0
IPv6 Address: Not Available
IPv6 Gateway Address:Not Available
IP Address Origin:Static
SRIOV State:Enabled
VEPA State:Disabled
```

5.3.3 PortStatistics

This command shows all function statistics for the specified function. The type of information listed may vary according to the adapter model.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
PortStatistics <WWPN|MAC> [clear]
```

Parameters

WWPN	The WWPN of an FCoE function.
MAC	MAC address of NIC function (on 10GBase-T adapters only).
clear	Clear counters for a 10GBASE-T function.

5.3.4 ServerAttributes

This command shows a list of server attributes for the server where the specified function is running. The type of information listed may vary according to the adapter model.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
ServerAttributes <WWPN|MAC>
```

Parameters

WWPN	The WWPN of an FCoE function.
MAC	The MAC address of a NIC or iSCSI function.

5.3.5 SetPhyPortSpeed

This command sets the port speed.

OneConnect adapters have configurable physical port speeds. Depending on the port module or transceiver installed in the physical port, the speed settings can be forced to a specific value, for instance 1 Gb, or to a range of values for auto-negotiation with the switch; for example, 10 Gb/1 Gb/100 Mb. Three values can be configured: port speed mode, speed value(s), and the DAC cable length.

The configurable port speeds are based on the port module type and the mode defined by the port speed *Mode* parameter. For the default port speed Mode, the speed setting is not required.

NOTE 10GBASE-CX4, 10GBASE XFP, and SGMII port module types do not support port speed settings.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
SetPhyPortSpeed <WWPN|MAC> <Mode> [Speed [Length]]
```

Parameters

WWPN	The WWPN of an FCoE function.
MAC	The MAC address of a NIC or iSCSI function.

Mode	The Mode number: 0 = Default 1 = Auto-negotiate; requires the Speed parameter 2 = Force; requires the Speed and Length parameters If the adapter's port speed value and the switch's port speed value conflict, the link will not be brought up.
Speed	The speed string of the PHY port. Some valid string values include 100 Mb, 1 Gb, and 10 Gb. The `PortAttributes` command lists all the valid speeds in Auto-negotiate and Force modes. **Auto-negotiated Speeds** A comma-separated list of available auto-negotiated speeds is displayed by the `PortAttributes` command. For combinations of speeds, each speed is separated by a slash, for example, 10 Gb/1 Gb/100 Mb. If the port does not support auto-negotiated speeds, this property is displayed as `Not Supported`. **Forced Speeds** A comma-separated list of available forced speeds is displayed by the `PortAttributes` command. Combinations of speeds for forced speeds are not available. If the port does not support forced speeds, this property is displayed as `Not Supported`. If the Mode parameter is 1 or 2, the Speed parameter is required. If the Mode parameter is 0, the Speed parameter is ignored.
Length	The length of the DAC cable in meters. Valid values are 0 to 10. A length value of 0 indicates an optical cable. A Length value is required if you are using a 10 Gb SFP and QSFP transceiver interface type. If the Mode parameter is 0, Speed and Length parameters are ignored. If the Mode parameter is 1, the Length parameter is ignored. For an embedded mezzanine adapter linked to an embedded switch on the internal port, the Length value is ignored.

Examples

The following example configures the PHY port to a forced speed of 1 Gb/s with a cable length of 10 meters:

```
brcmhbacmd setphyportspeed 00-00-c9-ad-ad-ac 2 1Gb 10
```

The following example tries to configure the PHY port to a forced speed of 100 Mbps:

```
brcmhbacmd setphyportspeed 00-00-c9-a9-41-88 2 100Mb
```

If the command is successful, the following is displayed:

```
Successfully changed speed settings on port.
```

If the *`Mode`* parameter is 2, this command results in the following error because you must include a value for the *`Length`* parameter:

```
ERROR: <431>: Cable length required for force mode and interface type
```

5.3.6 SetPortEnabled

This command enables or disables a port. When a port is disabled, packets are not transmitted or received on the port.

NOTE Make sure that all I/O traffic on the port is stopped before disabling the port.

NOTE OneConnect adapters do not require a reset when the adapter port is enabled or disabled.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
SetPortEnabled <WWPN|MAC> <PortEnable>
```

Parameters

WWPN	The WWPN of an FCoE function on the port.
MAC	The MAC address of a NIC or iSCSI function on the port.
PortEnabled	The port-enabled state: 0 = Disabled 1 = Enabled

5.4 Boot Commands

The Boot Commands group manages the commands that enable or disable the network boot for NIC ports or the boot code for FCoE ports. You can also show and change FCoE boot parameters.

5.4.1 EnableBootCode

This command enables or disables network boot. If network boot is being enabled, it is necessary to select the specific network boot type. The supported network boot types are PXE and iBFT. iBFT is not supported on all NIC adapter types.

NOTE To enable or disable boot on an iSCSI target, see Section 5.11.25, SetiSCSIBoot.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
EnableBootCode <MAC> <Flag> <NetworkBootMethod>
```

Parameters

MAC	The MAC address of a NIC function.
Flag	D = Disable the boot code. E = Enable the boot code.
NetworkBootMethod	The network boot method to be used by the NIC (PXE or iBFT). The network boot types supported on the specific NIC can be found in the Available Network Boot Methods string displayed by the HbaAttributes command.

Examples

The following example enables iBFT:

```
brcmhbacmd EnableBootCode  00-00-c9-11-22-33 e iBFT
```

The following example disables network boot:

```
brcmhbacmd EnableBootCode  00-00-c9-11-22-33 d
```

5.4.2 GetBootParams

This command shows the FCoE boot parameters. If any arguments are missing or invalid, an error is reported. If all arguments are correct, the data is displayed in tabular form.

Supported By

Linux, Solaris, and Windows

Syntax

```
GetBootParams <WWPN> <Type>
```

Parameters

WWPN	The WWPN of an FCoE function.
Type	X86, OpenBoot, or UEFI.

5.4.3 SetBootParam

This command changes the FCoE boot parameters. You can change function parameters and boot device parameters for x86, OpenBoot, and EFI boot.

- If you change adapter parameters, omit the `BootDev` keyword and value; otherwise, an error is reported.
- If you change boot device parameters for OpenBoot, omit the `BootDev` keyword and value; otherwise, an error is reported.
- For boot device parameters for x86 or UEFI, you must provide the `BootDev` keyword and value.

Supported By

Linux, Solaris, and Windows

Syntax

```
SetBootParam <WWPN> <Type> <Param> <Value1> [BootDev <Value2>]
```

Parameters

WWPN	The WWPN of an FCoE port.
Type	{x86, EFI, OB}
Param	The parameter name.
Value1	The parameter value.
BootDev	The boot device.
Value2	The boot device entry number: {0 to 7}.

Adapter Parameters	Boot Type	Value
DefaultAlpa	All	{ Value }
EnableAdapterBoot	All	{ State } (0=Disable, 1=Enable)
EnableBootFromSan	All	{ State } (0=Disable, 1=Enable)
LinkSpeed	All	{ 0, 1, 2, 4, 8 }
PlogiRetryTimer	All	{ 0, 1, 2, 3 }
Topology	All	{ 0, 1, 2, 3 }
AutoScan	X86	{ 0, 1, 2, 3 }
AutoBootSectorEnable	X86	{ State } (0=Disable, 1=Enable)

EDD30Enable	X86	{ State } (0=Disable, 1=Enable)
EnvVarEnable	X86	{ State } (0=Disable, 1=Enable)
SpinupDelayEnable	X86	{ State } (0=Disable, 1=Enable)
StartUnitCommandEnable	X86	{ State } (0=Disable, 1=Enable)
BootTargetScan	EFI	{ 0, 1, 2 }
DevicePathSelection	EFI	{ 0, 1 }
MaxLunsPerTarget	EFI	{ Value }
ResetDelayTimer	EFI	{ Value }
SfsFlag	OB	{ State } (0=Disable, 1=Enable)
Boot Device Parameters		
D_ID	All	{ Value [BootDev <Value2>] }
LUN	All	{ Value [BootDev <Value2>] }
TargetWwpn	All	{ Value [BootDev <Value2>] }
TargetID	OB	{ Value }

5.5 Channel Management Commands

These commands enable and disable channel management and set channel properties for CNAs.

NOTE OneCommand CNA Manager CLI commands for UMC management are still available for backward compatibility with existing UMC scripts. They cannot be used to manage other channel management types. See Section 5.17, UMC Commands, for information on UMC management.

Each port on each adapter's physical port can be partitioned into isolated channels providing a converged conduit for network and storage traffic. Each channel has its own unique MAC address. Depending on the type of channel management in effect, each channel provides various traffic management and provisioning capabilities, such as enabling and disabling, minimum and maximum bandwidth, and VLAN ID (in UMC for untagged packets, also called the LPVID).

The OneCommand CNA Manager application allows you to enable or disable channel management. In the case of UMC or SIMode, you can set each channel's properties. For the vNIC1 and UFP channel management types, the OneCommand CNA Manager application displays the channel properties, but it does not allow modification, except for vNIC1 where the LPVID (inner VLAN ID) can be modified.

Additionally, OCe14000-series adapters support NPar. Use the `GetAdapterPortConfig` and `SetAdapterPortConfig` command to manage NPar for these adapters. See Section 5.16.1, GetAdapterPortConfig, and Section 5.16.4, SetAdapterPortConfig (for OCe14000-Series Adapters). For Lenovo System X adapters, UMC mode is referred to as SIMode.

NOTE SR-IOV is not supported with channel management.

5.5.1 CMGetParams

This command shows the current channel management configuration for an adapter's physical port.

5.5.1.1 Multichannel

If multichannel is used, this command displays the adapter's active (booted) multichannel state, the configured state, the configured channel management mode (N/A if the configured state is disabled), and the available channel management modes. This is followed by a table showing the specified port's channel properties. The Type column shows the protocol that is running on the channel. The output is different depending upon the multichannel type.

NOTE While Flex is displayed in the output for the CMGetParams command as the multichannel type for HP adapters currently running in Flex mode, you cannot specify Flex for the *mctype* parameter of the SetAdapterPortConfig command. The only *mctype* parameter that can be specified on HP UMC-capable adapters is UMC.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
CMGetParams <MAC | WWPN>
```

Parameters

MAC	MAC address of any NIC or iSCSI function on an adapter port.
WWPN	WWPN of an FCoE function on an adapter port.

Consideration

If you specify a MAC address that belongs to a newly disabled port, an error message indicates that the operation cannot be performed on this port.

For example, if you type the following command:

```
>brcmhbacmd cmgetparams 00-90-FA-41-22-F0
```

The following error message is displayed:

```
ERROR: This port is configured to be removed after reboot. Channel
Management properties are unavailable.
```

Examples

The following table provides locations of application examples.

Example Types	Section
Multichannel disabled	5.5.1.1.1
Multichannel enabled – UMC, NIC-only	5.5.1.1.2
Multichannel enabled – UMC, with storage	5.5.1.1.3
Multichannel disabled – Lenovo System X	5.5.1.1.4
Multichannel enabled – Lenovo System X UFP	5.5.1.1.5
Dell NPar disabled	5.5.1.1.6
Dell NPar enabled – 4 functions per port	5.5.1.1.7

5.5.1.1.1 Multichannel Disabled

```
>brcmhbacmd CMGetParams 00-00-c9-12-34-56
Active Mode:     None
Configured mode: None
Available modes: UMC
```

5.5.1.1.2 Multichannel Enabled – UMC, NIC-only

```
>brcmhbacmd CMGetParams 00-00-c9-12-34-56
Active Mode:     UMC
Configured mode: UMC
Available modes: UMC

Func# Type  MAC Address       LPVID Min BW Max BW
----- ----- ----------------- ----- ------ ------
0     NIC   00-00-c9-12-34-56  2      25     50
1     NIC   00-00-c9-12-34-57  3      25      50
2     NIC   00-00-c9-12-34-58  4      25     50
3     NIC   00-00-c9-12-34-59  5      50     75
```

5.5.1.1.3 Multichannel Enabled – UMC, with Storage

```
>brcmhbacmd CMGetParams 00-00-c9-12-34-56
Active Mode:     UMC
Configured mode: UMC
Available modes: UMC

Func# Type  MAC Address       LPVID Min BW Max BW
----- ----- ----------------- ----- ------ ------
0     NIC   00-00-c9-12-34-56  2      25     50
1     FCoE  n/a                n/a    40     100
2     NIC   00-00-c9-12-34-58  4      25     50
3     NIC   00-00-c9-12-34-59  5      10     20
```

5.5.1.1.4 Multichannel Disabled – Lenovo System X

```
>brcmhbacmd CMGetParams 00-00-c9-12-34-56
Active Mode:     None
Configured mode: None
Available modes: vNIC1, SIMode, UFP
```

5.5.1.1.5 Multichannel Enabled – Lenovo System X UFP

```
>brcmhbacmd CMGetParams 00-00-c9-12-34-56
Active Mode:     UFP
Configured mode: UFP
Available modes: vNIC1, SIMode, UFP

                              Outer
Func# Type  MAC Address       VLAN  Min BW Max BW
----- ----- ----------------- ----- ------ ------
0     NIC   00-00-c9-12-34-56  2      25     50
1     NIC   00-00-c9-12-34-57  3      25     50
2     NIC   00-00-c9-12-34-58  4      25     50
3     NIC   00-00-c9-12-34-59  5      25     75
```

5.5.1.1.6 Dell NPar Disabled

```
>brcmhbacmd CMGetParams 00-00-c9-12-34-56
Active Mode:     None
Configured mode: None
Available modes: NPAR
```

5.5.1.1.7 Dell NPar Enabled – 4 Functions per Port

NOTE Even though function 1 has a minimum bandwidth of 0, some traffic flows through.

This command lists permanent MAC addresses, not virtual addresses.

```
>brcmhbacmd CMGetParams 00-00-c9-12-34-56
Active Mode:     NPAR
Configured mode: NPAR
Available modes: NPAR

Func# Type  MAC Address       Min BW Max BW
----- ----- ----------------- ------ ------
0     NIC   00-00-c9-12-34-56   25     50
1     NIC   00-00-c9-12-34-57   0      50
2     NIC   00-00-c9-12-34-58   25     50
3     NIC   00-00-c9-12-34-59   50     75
```

5.5.2 CMMode

This command enables or disables channel management mode only on OCe11100-series adapters. For UMC configurations, the `UMCEnable` command can still be used to enable UMC.

This command also sets the channel management type at the adapter level. A system reboot is required for the change take effect.

For OCe14000-series adapters, you must set the *`mctype`* parameter in the `SetAdapterPortConfig` command to enable multichannel, including NPar. See Section 5.16.4, SetAdapterPortConfig (for OCe14000-Series Adapters), for more information. If you use the `CMMode` command on an OCe14000-series adapter, an error message indicates that the command is not supported by the firmware or hardware.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
CMMode <WWPN | MAC Address> <None | Type>
```

Parameters

WWPN	WWPN of an FCoE function on the physical port.
MAC	MAC address of any NIC function on the physical adapter port.
None	Specify None to disable channel management or a channel.
Type	Management type (see Section 5.5.1, CMGetParams) to enable channel management.

Example

```
>brcmhbacmd CMMode 00-00-c9-bb-cc-aa None
```

This example disables channel management on an adapter containing a NIC function with a MAC address of 00-00-c9-bb-cc-aa.

5.5.3 CMSetBW

This command sets the minimum and maximum bandwidths for each channel on the physical port. For UMC and SIMode configurations, this command can also disable a channel's logical link by setting both the minimum and maximum bandwidths to `0`. To enable the logical link, specify a nonzero value for the minimum and maximum bandwidth.

If you are using NPar, setting the minimum bandwidth to zero does not bring the logical link down on the NPar function or prevent the NPar function from receiving or transmitting a small amount of network traffic.

The number of channels, the number of channels that need to be specified, and the number of bandwidth combinations that need to be specified depend on the adapter model. See Section 5.5.1, CMGetParams, to determine how many bandwidths need to be specified for a port by looking at the number of functions indicated in the `Func #` column of the output.

The total of the minimum bandwidths for the enabled channels or partitions (NPar) must add up to 100. An exception to this rule is for UMC and SIMODE configurations, when the minimum and maximum bandwidths for all channels are 0, effectively bringing down the logical link on all channels. The maximum bandwidth must be greater than or equal to the minimum bandwidth for that channel or partition (NPar), up to a maximum value of 100.

NOTE This command fails with an error message if the configured multichannel mode is *`vNIC1`*, *`UFP`*, or *`None`*.

Considerations

- If too many or too few minimum and maximum bandwidth combinations are provided, an error is generated.
- For UMC, the `UMCSetBW` command can still be used instead of `CMSetBW`.
- A reboot is not required to change the channel bandwidths when multichannel is currently active (including NPar) on the adapter.

  For NPar, these bandwidths are only in effect when ETS priority group bandwidths are not available. ETS priority group bandwidths are set by the `SetCnaPGBW` command.
- If a channel's protocol is configured to `None`, the minium and maximum bandwidth must be `0`.
- If you specify a MAC address that belongs to a newly disabled port, an error message indicates that the operation cannot be performed on this port.

  For example, if you type the following command:

  ```
  >brcmhbacmd cmSetBW  00-90-FA-41-22-F0 25,100 25,100 25,100 25,100
  ```

  The following message is displayed:

  ```
  ERROR: This port is configured to be removed after reboot. Bandwidth values
  can only be set on ports that will be available after reboot.
  ```

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
CMSetBW <MAC | WWPN> <Min0,Max0> [Min1,Max1]  ...  [MinN,MaxN]
```

Parameters

WWPN	WWPN of an FCoE function on the physical adapter port.
MAC	MAC address of any NIC or iSCSI function on the physical adapter port.
Min0, Max0	Minimum and maximum bandwidths for channel 0.
Min1, Max1	Minimum and maximum bandwidths for channel 1.
MinN, MaxN	Minimum and maximum bandwidths for channel N.

Examples

Down Logical Link on the Third Channel of a 4-Channels-per-Port Adapter

```
>brcmhbacmd CMSetBW 00-00-c9-12-34-56 25,50 50,100 0,0 25,50
```

```
>brcmhbacmd CMGetParams 00-00-c9-12-34-56
```

```
Active mode:     UMC
Configured mode: UMC
Available modes: UMC

Func# Type      MAC Address       LPVID Min BW Max BW
----- -----     ----------------- ----- ------ ------
0     NIC       00-00-c9-12-34-56  2      25     50
2     iSCSI     00-00-c9-12-34-57  n/a    50     100
4     NIC       00-00-c9-12-34-58  3      0      0
6     NIC       00-00-c9-12-34-59  4      25     50
```

5.5.4 CMSetLPVID

This command sets the LPVID values for the UMC and SIMode NIC channels. Use the `CMGetParams` command to determine the number of LPVIDs required. See Section 5.5.1, CMGetParams, for more information. A reboot is not required for these changes to take effect if UMC is enabled.

NOTE If the current multichannel mode is not UMC or SIMode, the `CMSetLPVID` command fails.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
CMSetLPVID <WWPN|MAC> <LPVID0> <LPVID1> ... [LPVIDn]
```

Parameters

WWPN	WWPN of an FCoE function on the physical adapter port.
MAC	MAC address of any NIC or iSCSI function on the physical adapter port.
LPVID0	LPVID for channel 0.
LPVID1	LPVID for channel 1.
LPVIDn	LPVID for channel n.

Considerations

- LPVID values are in the range of 2 to 4094.
- Every NIC channel on a physical port must have a unique LPVID.
- For FCoE and iSCSI channels, `0` must be entered because LPVIDs can only be specified for NIC channels.

- LPVIDs specified for channels with protocols set to None are ignored.
- This command is not supported on 1 Gb ports.
- If channel management is disabled when this command is executed, an error message is displayed.
- If you specify a MAC address that belongs to a newly disabled port, an error message indicates that the operation cannot be performed on this port.

  For example, if you type the following command:

  ```
  >brcmhbacmd cmSetLPVID 00-90-FA-41-22-F0 2 3 4 5 6 7 8 9
  ```

  The following message is displayed:

  ```
  ERROR: This port is configured to be removed after reboot. LPVID values can
  only be set on ports that will be available after reboot.
  ```

Examples

4 NIC Channels

```
>>brcmhbacmd CMSetLPVID 00-00-c9-12-34-56 1001 1002 1003 1004
```

8 Channels, Storage on Second Channel

```
>brcmhbacmd CMSetLPVID 00-00-c9-55-43-21 1001 0 1002 1003 1004 1005 1006 1007
```

5.6 DCB Commands

This command group shows the DCB and LLDP parameters for iSCSI, FCoE, and NIC adapter ports.

DCB commands are not available on OCe11101-EM/EX or OCe11102-EM/EX adapters. If a command is used on an unsupported adapter, the following error message is returned:

```
ERROR: <222>: DCB not available
```

5.6.1 GetDCBParams

This command shows the active and configured DCB and LLDP settings on a port of a OneConnect adapter. The active parameters show what the adapter port is currently running, and the configured parameters show the value to which the adapter port's DCB parameter is set.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetDCBParams <WWPN|MAC>
```

Parameters

WWPN	The WWPN of an FCoE function on the port.
MAC	The MAC address of a NIC or iSCSI function on the port.

Example

```
brcmhbacmd GetDCBParams 00-00-c9-93-2f-d8
```

5.6.2 GetPGInfo

This command shows the ETS priority group bandwidth percentages for the port of a OneConnect adapter. Additionally, this command displays the number of priority groups supported by an adapter.

NOTE On OCe14000-series adapters, if UFP multichannel is active and the PGs are enabled by the switch, the `GetPGInfo` command displays the active PG values.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetPGInfo <WWPN|MAC>
```

Parameters

WWPN	The WWPN address of an FCoE function on the port.
MAC	The MAC address of a NIC or iSCSI function on the port.

Example

```
brcmhbacmd getpginfo 00-00-c9-93-2f-d8
```

5.6.3 SetCnaPGBW

This command sets the ETS priority group bandwidth percentages on a port of a OneConnect adapter according to the following rules:

- Bandwidths (BW0–BW7) for priority groups 0 to 7 (PG0 to PG7) must total 100 (for 100%).
- Bandwidth can be assigned to a priority group that has priorities.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
SetCnaPGBW <WWPN|MAC> <BW0–BW7>
```

Parameters

WWPN	The WWPN of an FCoE function on the port.
MAC	The MAC address of a NIC or ISCI function on the port.
BW0–BW7	The bandwidths allocated for the priority groups 0 to 7.

Example

This command sets the bandwidth of PG0 to 50%, PG1 to 50%, and PG2 to PG7 to 0%.

```
brcmhbacmd SetCnaPGBW 10:00:00:00:c9:3c:f7:88 50 50 0 0 0 0 0 0
```

5.6.4 SetDCBParam

This command configures the DCB and LLDP settings on a OneConnect adapter port. Use the `GetDCBParams` command to obtain valid parameter names for use in this command.

NOTE You cannot set DCBX mode. If you attempt to specify a *dcbxmode* parameter, an error message is displayed.

NOTE RoCE is available only on OCe140000-series adapters. RoCE is available as Technical Preview only. Do not use RoCE in a production environment.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
SetDCBParam <WWPN|MAC> <Param> <Value>
```

Parameters

WWPN	The WWPN of an FCoE function on the port.
MAC	The MAC address of a NIC or iSCSI function on the port.
Param	The parameter name. See DCB Settings for <Param> and <Value> and LLDP Settings for <Param> and <Value>.
Value	The parameter value. See DCB Settings for <Param> and <Value> and LLDP Settings for <Param> and <Value>.

DCB Settings for <Param> and <Value>

<Param>	Description and <Value>
DCBXState	The DCBX protocol state. 0 = Disabled 1 = Enabled
PFCEnable	Flow control in both directions (transmit and receive). 0 = Disabled 1 = Enabled
FCoEPriority	This parameter is only applicable for ports running FCoE. A single priority must be specified. The range of valid values is 0 to 7. Only one priority can be specified for each invocation of this command and must be for a protocol running on the port. If more than one protocol priority can be set, they must be unique values.
iSCSIPriority	This parameter is applicable for ports running iSCSI only. A single priority must be specified. The range of valid values is 0 to 7. Only one priority can be specified for each invocation of this command and must be for a protocol running on the port. If more than one protocol priority can be set, they must be unique values.
RoCEPriority	This parameter is only applicable for ports running RoCE. A single priority must be specified. The range of valid values is 0 to 7. Only one priority can be specified for each invocation of this command and must be for a protocol running on the port. If more than one protocol priority can be set, they must be unique values.
PFCPriority	A list of comma-separated values where multiple PFC priorities are supported. The comma-separated list can contain up to seven values ranging from 0 to 7.
QCNEnable	This parameter is only applicable for ports on OCe14000-series adapters running NIC+RoCE, and only for RoCE traffic. 0 = Disabled 1 = Enabled

defaults	Use to set the DCB parameters (including priority groups) to their default values. For example:

```
brcmhbacmd SetDCBParam <WWPN|MAC> defaults
```

LLDP Settings for <Param> and <Value>

<Param>	Description and <Value>
TxState	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. 0 = Disabled 1 = Enabled
RxState	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. 0 = Disabled 1 = Enabled
TxPortDesc	Transmit Port Description: Provides a description of the port in an alphanumeric format. 0 = Disabled 1 = Enabled
TxSysDesc	Transmit System Description: Provides a description of the network entity in an alphanumeric format. 0 = Disabled 1 = Enabled
TxSysName	Transmit System Name: Provides the system's assigned name in an alphanumeric format. 0 = Disabled 1 = Enabled
TxSysCap	Transmit System Capabilities: 0 = Disabled 1 = Enabled

Example

```
brcmhbacmd SetDCBParam 00-00-c9-3c-f7-88 fcoepriority 3
```

5.6.5 SetDCBPriority

This command sets the PFC priorities and the ETS priority groups priorities. The values must be set according to the following rules:

- The priorities range from 0 to 7.
- A priority (0 to 7) must exist in only one priority group.
- All priorities must appear once in any of the eight (PG0 to PG7) priority groups or if available, PG15.

NOTE For OCe14000-series adapters, you can use the PG15 priority group for only RoCE priority.

- To not specify priorities for a priority group, use a dash (-).
- Any assigned PFC priority must be assigned as the single priority in a priority group (for example, no other priorities allowed in a group assigned the PFC priority).
- Any PG assigned one or more priorities must also be assigned a nonzero bandwidth value (see Section 5.6.3, SetCnaPGBW).

The following rules are specific to FCoE and iSCSI adapters:

- A maximum of two PFC priorities can be assigned.
- If FCoE is running on the port, one of the PFC priorities must match the FCoE priority.
- If iSCSI is running on the port, one of the PFC priorities must match the iSCSI priority.

The following rules are specific to NIC-only adapters:

- Only one PFC priority can be assigned.
- In NIC-Only mode, PFC is disabled by default. To enable PFC, NIC ETS must be enabled.

  To enable NIC ETS:

  - In Windows, enable `Enhanced Transmission Selection` in the driver properties (for example, in the Device Manager property page for the NIC driver).
  - In Linux, load the NIC driver with the `tx_prio` driver parameter set to 1.

The following rules are specific to RoCE adapters:

- The RoCE priority must be assigned to the PFC priority and as the only priority to Priority Groups PG0 to PG7, or to PG15.
- All non-RoCE priorities must be assigned to only one priority group other than the RoCE priority group.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
SetDcbPriority <WWPN|MAC> <PFC> <PG0> <PG1> <PG2> <PG3> <PG4> <PG5> <PG6> <PG7>
[PG15]
```

Parameters

WWPN	The WWPN of an FCoE function on a port.
MAC	The MAC address of a NIC or iSCSI function on a port.
PFC	The PFC priority that is a comma-separated list of up to eight values, ranging from 0 to 7.
PG0-PG7	Priority group membership that is a comma-separated list of priorities ranging from 0 to 7. Each set of priorities for a group must be separated by a space. All priorities (0 to 7) must be assigned to a PGID.
PG15	PG15 group membership (only required on RoCE ports)

Example

```
brcmhbacmd h=10.192.203.151 m=cim SetDCBPriority 10:00:00:00:c9:3c:f7:88 3
0,1,2,4,5,6,7 3 0 0 0 0 0 0
```

5.7 Diagnostic Commands

The Diagnostic Commands group provides commands that enable you to detect cabling problems, to examine transceiver data, and to flash memory load lists. Additionally, you can run specific diagnostic tests, such as the Loopback test.

5.7.1 GetBeacon

This command shows the current beacon state (either on or off).

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

GetBeacon <WWPN|MAC>

Parameters

WWPN	The WWPN of an FCoE function on the port.
MAC	The MAC address of a NIC or iSCSI function on the port.

5.7.2 GetXcvrData

This command shows transceiver data for a port on an adapter.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetXcvrData <WWPN|MAC> [Type]
```

Parameters

WWPN	The WWPN of an FCoE function on a port.
MAC	The MAC address of a NIC or iSCSI function on a port.
Type	This optional parameter indicates the type of SFP data to display: 1 = Formatted SFS data (default) 2 = Raw SFS data (not supported by Windows + CIM Provider on a VMware host)

Example

```
brcmhbacmd GetXcvrData 00-00-c9-93-2f-d6
```

5.7.3 LoopBackTest

This command runs one of the loopback tests available on the adapter port specified by the WWPN or MAC address.

NOTE PHY diagnostics are not supported on mezzanine cards and blade network daughter cards because they do not contain PHYs.

NOTE The external loopback test is not supported for OCe11102-xT adapters and OCe14102-UT adapters.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
LoopBackTest <WWPN 1 MAC> <Type> <Count> <StopOnError> [Pattern]
```

Parameters

WWPN	The WWPN of an FCoE function on a port.
MAC	The MAC address of a NIC or iSCSI function on a port.

Type	The type of loopback test to run: — 2 = External loopback test (requires loopback plug) — 3 = DMA loopback test — 4 = PHY loopback test — 5 = MAC loopback test
Count	Number of times to run the test. Possible values are 1 to 99,999. To run the test infinitely, use 0.
StopOnError	Checks whether the test must be halted on error. 0 = No halt 1 = Halt
Pattern	An optional parameter that specifies 1–8 hexadecimal bytes to use for loopback data (for example, 1a2b3c4d).

Example

```
brcmhbacmd LoopBackTest 00-00-c9-93-2f 4 120 0
```

5.7.4 PciData

This command shows the PCI configuration data (if available).

The PCI registers displayed are specific to the function referenced in the OneCommand CNA Manager CLI. For example, if you specify the WWPN for the FCoE function, the PCI registers for that FCoE function are returned. If you specify the MAC address for the NIC function on that same physical port, the PCI registers for that NIC function are returned.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
PciData <WWPN|MAC>
```

Parameters

WWPN	The WWPN of an FCoE function.
MAC	The MAC address of a NIC or iSCSI function.

Example

```
brcmhbacmd PciData 00-00-c9-93-2f-d6
The example output:
Vendor ID:        0x19A2       Device ID:          0x0700
Command:          0x0406       Status:             0x0010
Revision ID:      0x02         Prog If:            0x00
Subclass:         0x00         Base Class:         0x02
Cache Line Size:  0x10         Latency Timer:      0x00
Header Type:      0x80         Built In Self Test: 0x00
Base Address 0:   0x00000000   Base Address 1:     0xDF478000
Base Address 2:   0xDF480004   Base Address 3:     0x00000000
Base Address 4:   0xDF4A0004   Base Address 5:     0x00000000
CIS:              0x00000000   SubVendor ID:       0x10DF
SubSystem ID:     0xE622       ROM Base Address:   0x00000000
Interrupt Line:   0x00         Interrupt Pin:      0x01
Minimum Grant:    0x00         Maximum Latency:    0x00
Capabilities Ptr: 0x40
```

5.7.5 SetBeacon

This command turns the beacon on or off on the adapter port.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
SetBeacon <WWPN | MAC> <BeaconState>
```

Parameters

WWPN	The WWPN of an FCoE function on the port.
MAC	The MAC address of a NIC or iSCSI function on the port.
BeaconState	Indicates the state of the beacon. 0 = Off 1 = On

5.7.6 SetCableNVP

This command sets the NVP, which is required for the `TDRTest` command, for the cable that connects to the physical port associated with the WWPN or MAC.

NOTE This command supports only OCe11100-series and OCe14000-series 10GBASE-T adapters.

Supported By

Linux, Solaris, and Windows

Syntax

```
SetCableNVP <WWPN|MAC> <NVP>
```

Parameters

WWPN	The WWPN of an FCoE function on the port.
MAC	The MAC address of a NIC or iSCSI function on the port.
NVP	A percentage value between 1 and 100. Consult your cable documentation to obtain the proper NVP value.

5.7.7 TDRTest

NOTE This command only supports OneConnect adapters that use coax copper cables, such as the 10GBASE-T adapter.

The TDR test attempts to determine whether any cable faults are compromising the integrity of the link. The test requires that the cable is free of disturbances, is down, and is quiet. It is best if the cable is not terminated; however, if it is terminated, the link partner must not be active during the test.

For each twisted pair cable (labeled A, B, C, or D):

- If a fault cannot be detected, the test output displays `OK` and an estimated cable length (in meters), if possible. If the length cannot be determined, the estimated length is displayed as invalid.
- If one of two faults (a short or an open connection) is detected, the test output displays the fault type (`Shorted` or `Open`) and the distance to the fault (in meters).

Supported By

Linux, Solaris, and Windows

Syntax

```
TDRTest <MAC_Address>
```

Parameters

MAC_Address	The MAC address of the NIC or iSCSI port.

Example

```
brcmhbacmd TDRTest 00-90-FA-27-A1-70
```

The example output:

```
Pair A: OK. Cable Length Estimation: 50m.
Pair B: OK. Cable Length Estimation: invalid.
Pair C: Open. Distance to Fault: 38m.
Pair D: Shorted. Distance to Fault: 36m.
```

5.8 Driver Parameter Commands

The Driver Parameter Commands group controls the driver parameter values. You can also change the parameters back to factory default values.

NOTE Driver Parameter commands are supported only for FCoE ports.
The `DriverConfig` and `SetDriverParamDefaults` commands are not supported for Solaris.

Considerations

- Driver parameters set to temporary or global values (using the `T` and `G` flags, respectively) must be read using the `GetDriverParams` command to view the current value of the parameter. The `GetDriverParamsGlobal` command returns only permanently set driver parameter values.

  Additionally, if temporary and global values are set for one or more driver parameters, the `SaveConfig` command must be run with the `N` flag (using the `N` flag is analogous to using the `GetDriverParams` command) to force the driver parameter values for the specified adapter to be saved. Inaccurate values can be saved if the `G` flag is used for this command.
- The list of available driver parameters that can be configured are different depending on the operating system.

5.8.1 DriverConfig

This command sets all driver parameters to the values in the `.dpv` file type. The `.dpv` file's driver type must match the driver type of the host operating system adapter.

Supported By

Linux and Windows

Syntax

```
DriverConfig <WWPN> <FileName> <Flag>
```

Parameters

WWPN	The WWPN of an FCoE function.
FileName	The name of the `.dpv` file, which is stored in the Emulex Repository directory. For Linux, this is `/opt/broadcom/brcmocmanager/RMRepository`. For Windows, this is `\Program Files\Broadcom\Util\EmulexRepository`.
Flag	G = Make the change global (all FCoE functions on this host). N = Make the change non-global (function-specific).

5.8.2 GetDriverParams

This command shows the name and values of each parameter.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetDriverParams <WWPN>
```

Parameters

WWPN	The WWPN of an FCoE function.

5.8.3 GetDriverParamsGlobal

This command shows the name and the global value of each driver parameter.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetDriverParamsGlobal <WWPN>
```

Parameters

WWPN	The WWPN of an FCoE function.

5.8.4 SaveConfig

This command saves the specified adapter's driver parameters to a file. The resulting file contains a list of driver parameter definitions in ASCII file format with definitions delimited by a comma. Each definition has the following syntax:

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

The command saves either the values of the global set, or those specific to the adapter in the Emulex Repository directory. For Linux, this is `/opt/broadcom/brcmocmanager/RMRepository`. For Windows, this is `\Program Files\Broadcom\Util\EmulexRepository`.

Supported By

Linux, Solaris, and Windows

Syntax

```
SaveConfig <WWPN> <FileName> <Flag>
```

Parameters

WWPN	The WWPN of an FCoE function.
FileName	Name of the file that contains the driver parameters list.
Flag	G = Save the global parameter set. N = Save the local (function-specific) parameter set.

5.8.5 SetDriverParam

This command changes a driver parameter and designates the scope of the change.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
SetDriverParam <WWPN> <Flag1> <Flag2> <Param> <Value>
```

Parameters

WWPN	The WWPN of an FCoE function.
Flag1	L = Make change local for only this function. G = Make change global (all FCoE functions on this host).
Flag2	P = Make change permanent (persists across reboot). For Linux, to make a permanent change that persists across reboots, you must set Flag1 to G (Global). T = Make change temporary.
Param	Name of the parameter to modify.
Value	New parameter value, decimal or hexadecimal (0xNNN).

Example

To enable dynamic target mode:

```
brcmhbacmd SetDriverParam 10:00:00:00:c9:ff:ff:ff L P
```

To disable dynamic target mode, set the flag to 0.

5.8.6 SetDriverParamDefaults

This command changes all values to the default for the adapter.

Supported By

Linux, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
SetDriverParamDefaults <WWPN> <Flag1> <Flag2>
```

Parameters

WWPN	The WWPN of an FCoE function.
Flag1	L = Make change local for only this function. G = Make change global (applies to all FCoE functions on this host).
Flag2	P = Make change permanent (the change persists across reboot). T = Make change temporary.

5.9 Dump Commands

The diagnostic dump feature enables you to create a dump file for a selected adapter. Dump files contain information such as firmware version, driver version, and operating system information. This information is useful when your are troubleshooting an adapter, but it is unavailable in read-only mode.

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

The dump files created are text files (`.txt` extension) and binary files. The extension for binary files depends on the dump type:

- Enhanced FAT dump – `.edf` extension
- Core dump – `.core` extension

5.9.1 DeleteDumpFiles

This command deletes all diagnostic dump files for an adapter.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
DeleteDumpFiles <WWPN|MAC>
```

Parameters

WWPN	The WWPN of an FCoE function on the adapter.
MAC	The MAC address of a NIC or iSCSI port function on the adapter.

5.9.2 Dump

This command creates a diagnostic dump file in the `BrcmHbaCmd` dump file directory.

NOTE In some cases, a core dump can be performed on an inoperative adapter. To view inoperative adapters on the local host, use the `ListHBAs down` command. See Section 5.13.9, ListHBAs.

If the core optional parameter is not specified, an enhanced FAT dump is performed by default.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
Dump <WWPN|MAC> [core [Options]]
```

Parameters

WWPN	The WWPN of an FC or FCoE port.
MAC	The MAC address of a NIC or iSCSI port.
core	Perform a core dump on an adapter (local host only).
Options	For available core dump options, contact Broadcom® technical support.

5.9.3 GetDumpDirectory

This command shows the dump file directory for the adapters in the host.

NOTE The dump directory can be set only on VMware ESXi hosts.
The dump directory applies to all adapters in the server. A separate dump directory for each adapter does not exist.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetDumpDirectory [WWPN|MAC]
```

Parameters

WWPN	Obsolete; ignored if specified.
MAC	Obsolete; ignored if specified.

5.9.4 GetDumpFile

This command gets the dump file. This command gets the user-specified dump file to the local client's dump directory. The dump directory (local and remote) is named `Dump`. The dump files are copied from the dump directory of the remote host to the dump directory of the local host. Therefore, if the remote host option is not specified (h=*`IP_Address`*`[:port]`), this command returns an error because the source and destination directories are the same.

Dump directory:

- Windows – *`SystemDrive_Letter`*`:\Program Files\broadcom\Util\Dump`
- Linux – `/var/opt/broadcom/brcmocmanager/Dump`
- Solaris – `/opt/brcmocm/Dump`
- VMware ESXi – The dump directory set using the `SetDumpDirectory` command.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetDumpFile <h=IP_Address[:port]>[WWPN|MAC] <filename>
```

Parameters

WWPN	Obsolete; ignored if specified.
MAC	Obsolete; ignored if specified.
filename	The name of the dump file to be copied from the remote host.

Example

```
brcmhbacmd h=10.192.193.154  GetDumpFile
BG-HBANYWARE-15_10000000c97d1314_20100120-032820421.dmp
```

5.9.5 GetDumpFileNames

This command gets the names of the files in the host's dump directory.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetDumpFileNames [WWPN|MAC]
```

Parameters

WWPN	Obsolete; ignored if specified.
MAC	Obsolete; ignored if specified.

Example

```
brcmhbacmd GetDumpFileNames
```

5.9.6 GetRetentionCount

This command shows the maximum number of diagnostic dump files to keep.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetRetentionCount [WWPN|MAC]
```

Parameters

WWPN	Obsolete; ignored if specified.
MAC	Obsolete; ignored if specified.

5.9.7 SetDumpDirectory

This command sets the dump directory (valid only on VMware ESXi hosts).

Supported By

Windows + CIM Provider on a VMware host

To use the `SetDumpDirectory` command, you must have a directory (which must be a Storage partition) mapped under `/vmfs/volumes` where the files are dumped. This directory points to the internal hard disk or an external storage area and can also be mapped using the vSphere Client utility from VMware.

The application checks for the dump directory and creates the dump files in that location.

In a remote environment, you can use the `SetDumpDirectory` command from a host running any operating system (including Linux, Solaris, and Windows), but only to a remote host that is running VMware ESXi.

NOTE The dump directory applies to all adapters in the server. A separate dump directory for each adapter does not exist.

Syntax

```
SetDumpDirectory <DumpDirectoryName>
```

Parameters

`DumpDirectoryName`	The directory under `/vmfs/volumes` that you created to store the dump files.

Example

This example shows the dump directory set to `/vmfs/volumes/ocm-datastore`:

```
brcmhbacmd h=10.192.203.173 m=cim u=root p=Swamiji001 n=root/brcmccx
SetDumpDirectory 10:00:00:00:c9:61:f2:64 ocm-datastore
```

5.9.8 SetRetentionCount

This command specifies the maximum number of diagnostic dump files for the adapter. When the count reaches the limit, the next dump operation deletes the oldest file.

NOTE The retention count applies to all adapters in the server.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
SetRetentionCount [WWPN|MAC] <Value>
```

Parameters

WWPN	Obsolete; ignored if specified.
MAC	Obsolete; ignored if specified.
Value	The number of dump files to retain.

Example

```
brcmhbacmd SetRetentionCount 6
```

5.10 FCoE Commands

The FCoE Commands group manages the FIP parameters and displays the FCF for an FCoE function.

5.10.1 GetFCFInfo

This command shows the FCF information of the FCoE function.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetFCFInfo <WWPN>
```

Parameters

WWPN The WWPN of an FCoE function.

Example

```
brcmhbacmd GetFCFInfo 10:00:00:00:c9:3c:f7:88
Number of FCFs:   1
Active FCFs:      1
Entry 0:
      State:            1
      Priority:         133
      Fabric Name:      10:00:00:05:1E:0C:54:49
      Switch Name:      10:00:00:05:1E:0C:54:49
      MAC:              00:05:9B:71:3D:71
      FC Map:           0x0EFC00
      VLAN IDs:
      LKA Period:       8
```

5.10.2 GetFIPParams

This command gets the FIP parameters of an FCoE function.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetFIPParams <WWPN>
```

Parameters

WWPN The WWPN of an FCoE function.

Example

```
brcmhbacmd h=10.231.140.83 getfipparams 10:00:00:00:c9:bc:a9:31
Param Description       Param Name  Value
-----------------       ----------  -----
Primary Fabric Name     pfabric     FF:FF:FF:FF:FF:FF:FF:FF
Primary Switch Name     pswitch     FF:FF:FF:FF:FF:FF:FF:FF
DCB VLan ID             vlanid      Any VLAN ID is valid
```

5.10.3 SetFIPParam

This command sets the FIP parameters of an FCoE function.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
SetFIPParam <WWPN> <Param> <Value>
```

Parameters

WWPN — The WWPN of an FCoE function.

Param — The FIP parameter name:
- pfabric
- pswitch
- vlanid
- fcmap
- cinvlanid

Value — The value based on the FIP parameter name:
- pfabric: 8-byte fabric name (format XX:XX:XX:XX:XX:XX:XX:XX)
- pswitch: 8-byte switch name (format XX:XX:XX:XX:XX:XX:XX:XX)
- vlanid: 2-byte VLAN ID [0–4095] or *any* for any VLANID
- fcmap: 3-byte FC_map, 0x0EFCxx
- cinvlanid: 2-byte VLAN_ID [0–4095]

Example

```
brcmhbacmd SetFIPParam 10:00:00:00:c9:5b:3a:6d fcmap 0x0efc99
```

5.11 iSCSI Commands

The iSCSI Commands group supports the iSCSI interface in the CLI.

NOTE iSCSI commands are supported only on OneConnect iSCSI functions. Only OCe14000-series adapters support IPv6 addresses.

The MAC address <MAC_Address> of the iSCSI port must be passed to each command as the first argument.

Some commands require values to be set in a format similar to `option_name=value`. Type the full option name or the abbreviated option name (shown in Table 7) and enter the value. The abbreviations are not case sensitive.

Table 7 Option Names

Option Name	Abbreviation	Example
Auth	au	au=1
DataDigest	dd	dd=1
DHCP	dh	dh=1
HeaderDigest	hd	hd=1
ImmediateData	id	id=1
Initiator_Alias	ia	ia=”initiator_alias”
Initiator_Name	in	in=”initiator_name”
Priority	pr	pr=1
VLAN_ENABLED	ve	ve=1
VLAN_ID	vi	vi=1

5.11.1 AddARPTableEntry

This command adds an ARP table entry.

Supported By

Linux and Windows

Syntax

```
AddARPTableEntry <MAC_Address> <Dest_MAC_Address> <Dest_IP_Address>
```

Parameters

MAC_Address	The MAC address of the iSCSI function.
Dest_MAC_Address	The destination MAC address to add to the ARP table.
Dest_IP_Address	The destination IP address to add to the ARP table.

5.11.2 AddiSNSServer

This command adds a new iSNS server to the existing set of iSNS servers. It accepts either an IPv4 or IPv6 server address.

NOTE For OCe11000-series adapters, only one iSNS server can be configured. For OCe14000-series adapters, up to four iSNS servers can be configured.

Supported By

Linux and Windows

Syntax

```
AddiSNSServer <MAC_Address> <Server_IP> <Port>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
Server_IP	IP address of the iSNS server to configure.
Port	Port number of the iSNS server to configure (value: 1024 to 65535). iSNS is not supported on iSCSI functions running the Open iSCSI driver. In this case, adding an iSNS server will not detect any iSCSI targets.

5.11.3 AddRouteTableEntry

This command adds a new route table entry to the route table of the specified function.

Supported By

Linux and Windows

Syntax

```
AddRouteTableEntry <MAC_Address> <Dest_IP_Address> <Subnet_Mask> <Gateway>
```

Parameters

MAC_Address	MAC address of an iSCSI function.
Dest_IP_Address	Destination IP address to add to the route table.
Subnet_Mask	Subnet mask to add to the route table.
Gateway	Gateway to add to the route table.

5.11.4 AddTarget

This command adds a target to the list of targets seen by the initiator and logs into the target after it has been successfully created. This command requires that you specify a valid IPv4 or IPv6 target IP <*Target_IP*>, port number <*Port*>, and iSCSI name <*iscsi_target_name*>. If you do not specify the remaining options, these options are set to their default values.

If you set the authentication method <*Auth*> to a value other than 0, you must set additional parameters. Each string must be enclosed in quotation marks to avoid mishandling by the Windows, Linux, Solaris, or VMware shell's parser.

- If you set the authentication method to One-Way CHAP (<*Auth*>=1), you must also specify the *TargetCHAPName* and *TargetSecret*. For example:

```
brcmhbacmd AddTarget 00-11-22-33-44-55 192.168.1.1 8000 iscsitarget Auth=1
"TgtCHAPName" "TargetSecret1"
```

- If you set the authentication method to Mutual CHAP (<Auth>=2), you must specify all four values. For example:

```
brcmhbacmd AddTarget 00-11-22-33-44-55 192.168.1.1 8000 iscsitarget Auth=1
"TgtCHAPName" "TargtSecret1" "InitCHAPName" "InitialSecret1"
```

Supported By

Linux and Windows

Syntax

```
AddTarget <MAC_Address> <Target_IP> <Port> <iscsi_target_name>
[ImmediateData=<0|1>] [HeaderDigest=<0|1>] [DataDigest=<0|1>] [Boot=<0|1>]
[Login=<0|1>] [Auth=<0|1|2> "TgtCHAPName" "TgtSecret" "InitCHAPName"
"InitSecret"]
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
Target_IP	IP address of the target portal. You can specify an IPv4 or IPv6 IP address.
Port	Port number of the target portal (value: 102 to 65535).
iscsi_target_name	Target's iSCSI name enclosed in quotation marks (string length: 11 to 255).
ImmediateData	0 = No 1 = Yes (default)
HeaderDigest	0 = None (default) 1 = CRC32C
DataDigest	0 = None (default) 1 = CRC32C
Boot	This optional parameter specifies if the added target is a boot device: 0 = Added target is not a boot device 1 = Added target is a boot device

Login	This optional parameter specifies whether to log in to the target after it has been added: 0 = Do not log in to the target 1 = Specify log in to the target If the Login parameter is omitted, the default is 1.
Auth	0 = None (default) 1 = One-Way CHAP 2 = Mutual CHAP
TgtCHAPNam	Target CHAP name enclosed in quotation marks (string length: 1 to 256). Required when Auth =1 or 2.
TgtSecret	Target Secret enclosed in quotation marks (string length: 12 to 16). Required when Auth =1 or 2.
InitCHAPName	Initiator CHAP name enclosed in quotation marks (string length: 1 to 256). Required when Auth =2.
InitSecret	Initiator Secret enclosed in quotation marks (string length: 12 to 16). Required when Auth =2.

5.11.5 AddTargetPortal

This command adds a new SendTarget Portal for the initiator and runs a target detection after the SendTarget Portal is created. This command requires that you specify a valid IPv4 or IPv6 portal IP address <*Target_IP*> and a valid port number <*Port*>. If you do not specify the remaining options, these options are set to their default values.

If you set the authentication method <*Auth*> to a value other than 0, you must set additional parameters. Each string must be enclosed in quotation marks to avoid mishandling by the Windows, Linux, Solaris, or VMware shell's parser.

- If you set the authentication method to One-Way CHAP (<*Auth*>=1), you must also specify the Target CHAP Name and Target Secret. For example:

```
brcmhbacmd AddTargetPortal 00-11-22-33-44-55 10.0.0.1 8000 Auth=1
"TgtCHAPName" "TargetSecret1"
```

- If you set the authentication method to Mutual CHAP (<*Auth*>=2), you must specify all four values.

  For example:

```
brcmhbacmd AddTargetPortal 00-11-22-33-44-55 10.0.0.1 8000 Auth=2
"TgtChapName" "TargetSecret1" "InitCHAPName" "InitialSecret1"
```

You must specify either the TSIH value or the ISID qualifier. If you specify the ISID qualifier, you must also specify the Target's ID address.

Supported By

Linux and Windows

Syntax

```
AddTargetPortal <MAC_Address> <Target_IP> <Port> [ImmediateData=<0|1>]
[HeaderDigest=<0|1>] [DataDigest=<0|1>] [Auth=<0|1|2> "TgtCHAPName" "TgtSecret"
"InitCHAPName" "InitSecret"]
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
Target_IP	IP address of the target portal. You can specify an IPv4 or IPv6 IP address.
Port	Port number of the target portal (value: 1024 to 65535).
ImmediateData	0 = No 1 = Yes (default)

HeaderDigest	0 = None (default) 1 = CRC32C
DataDigest	0 = None (default) 1 = CRC32C
Auth	0 = None (default) 1 = One-Way CHAP 2 = Mutual CHAP
TgtCHAPNam	Target CHAP name enclosed in quotation marks (string length: 1 to 256). Required when Auth =1 or 2.
TgtSecret	Target Secret enclosed in quotation marks (string length: 12 to 16). Required when Auth =1 or 2.
InitCHAPName	Initiator CHAP name enclosed in quotation marks (string length: 1 to 256). Required when Auth =2.
InitSecret	Initiator Secret enclosed in quotation marks (string length: 12 to 16). Required when Auth =2.

5.11.6 ClearISNSServer

This command clears the configured iSNS server and disables iSNS target detection. If no iSNS server is currently configured, or if two or more iSNS servers are defined, any attempt to use this command returns an error.

NOTE This command works only if one iSNS server is defined.

Supported By

Linux and Windows

Syntax

```
ClearISNSServer <MAC_Address>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.

5.11.7 DelARPTableEntry

This command removes an ARP table entry.

Supported By

Linux and Windows

Syntax

```
DelARPTableEntry <MAC_Address> <Dest_MAC_Address> <Dest_IP_Address>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
Dest_MAC_Address	The destination MAC address to remove from the ARP table.
Dest_IP_Address	The destination IP address to remove from the ARP table.

5.11.8 DeleteiSNSServer

This command deletes an iSNS server from the current iSNS server list.

Supported By

Linux and Windows

Syntax

```
DeleteiSNSServer <MAC_Address> <IP_Address>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
IP_Address	The IPv4 or IPv6 IP address of the iSNS server.

5.11.9 DelRouteTableEntry

This command removes a route table entry from the specified function.

Supported By

Linux and Windows

Syntax

```
DelRouteTableEntry <MAC_Address> <Dest_IP_Address> <Subnet_Mask> <Gateway>
```

Parameters

MAC_Address	MAC address of an iSCSI function.
Dest_IP_Address	Destination IP address to delete from the route table.
Subnet_Mask	Subnet Mask to delete from the route table.
Gateway	Gateway to delete from the route table.

5.11.10 DiscoveriSNSServer

This command detects an iSNS server address through DHCP. If the DHCP server returns an iSNS server address, it replaces all manually configured iSNS servers and can be viewed using the `ShowiSNSServer` command.

NOTE OCe11000-series adapters can detect one iSNS server, and OCe14000-series adapters can detect up to four iSNS servers.

Supported By

Linux and Windows

Syntax

```
DiscoveriSNSServer <MAC_Address>
```

Parameters

MAC_Address	The MAC address of an iSCSI function. iSNS is not supported on iSCSI functions running the Open iSCSI driver. In this case, adding an iSNS server will not detect any iSCSI targets.

5.11.11 EraseiSCSIConfig

This command resets the iSCSI function configuration for OneConnect adapters to factory defaults.

Supported By

Linux and Windows

Syntax

```
EraseiSCSIConfig <MAC address>
```

Parameters

MAC_Address	The MAC address of an iSCSI function for which the configuration is to be erased.

5.11.12 ExportiSCSI

This command outputs the iSCSI function and target information in XML format. The output can be redirected to a specified file, and the file can be used for the `ImportiSCSI` command. This command works only on the local host.

Supported By

Windows

Syntax

```
ExportiSCSI
```

Example

In this example, the command exports all the iSCSI targets of all the adapters on the host. The output is redirected to the file `targets.xml`.

```
brcmhbacmd ExportiSCSI > targets.xml
```

5.11.13 GetInitiatorProperties

This command shows all the initiator login options for the specified port.

These properties are set as the target portal's login properties to be used when detecting the targets on the target portal. The detected targets inherit these properties.

Supported By

Linux and Windows

Syntax

```
GetInitiatorProperties <MAC_Address>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.

5.11.14 GetiSCSILuns

This command shows all the LUNs and their information for a specified target. The command gathers the information from the iSCSI target indicated by the <*iscsi_target_name*> parameter.

Supported By

Linux and Windows

Syntax

```
GetiSCSILuns <MAC_Address> <iscsi_target_name>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
iscsi_target_name	Target's iSCSI name enclosed in quotation marks (string length: 11 to 255)

5.11.15 GetiSCSIPortStats

This command shows all the iSCSI statistics for a specified function.

Supported By

Linux and Windows

Syntax

```
GetiSCSIPortStats <MAC_Address>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.

5.11.16 GetNetworkConfiguration

This command lists a port's TCP/IP information for IPv4 and IPv6 protocols for the iSCSI function.

Supported By

Linux and Windows

Syntax

```
GetNetworkConfiguration <MAC_Address>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.

Example

```
TCP/IP Configuration for 00-00-c9-ad-ad-b1:

VLAN Enabled:          No
VLAN ID:               0
Priority:              0

IPv4 Configuration:
DHCP Enabled:          Yes
IP Address:            10.192.81.204
Subnet Mask:           255.255.248.0
Gateway:               10.192.87.254
IPv6 Configuration:
Automatic Assignment:  No
Link Local Address:    fe80::a1d3:f062:f44a:7577
IP Address 1:          fd01::16
IP Address 2:          fd02::16
Gateway:               fd00::1
```

5.11.17 GetSessionInfo

This command lists all session information for a specified session.

You must specify the <*iscsi_target_name*> and either the <*TSIH*> of the session or the session's ISID Qualifier <*ISID_Qual*> and the target's IPv4 or IPv6 IP address <*Target_IP*>. These parameters tell the command to gather the information from the specified target and session. You can find the TSIH and ISID qualifier by running the `ListSessions` command.

Supported By

Linux and Windows

Syntax

```
GetSessionInfo <MAC_Address> <iscsi_target_name> <TSIH | <ISID_Qual Target_IP>>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
iscsi_target_name	Target's iSCSI name enclosed in quotation marks (string length: 11 to 255).
TSIH	TSIH value of the session (value: 1 to 65535).
ISID_Qual	ISID qualifier of the session (value: 0 to 65535).
Target_IP	The Target's IP address. You can specify an IPv4 or IPv6 IP address.

5.11.18 ImportiSCSI

This command imports iSCSI function configuration and targets from an XML file to the iSCSI functions on the local host. The XML file is created by the `ExportiSCSI` command.

Supported By

Windows

Syntax

```
ImportiSCSI <Import_File> [clean]
```

Parameters

Import_File	The name of the XML file containing import information that was generated by the ExportiSCSI command (see Section 5.11.12, ExportiSCSI).
clean	Erases entire iSCSI configuration on all iSCSI ports before importing targets.

NOTE This command fails if the system is booted from an iSCSI target.

Example

In this example, the command imports the iSCSI targets found in the targets.xml file to the iSCSI ports found on the host, and erases the iSCSI configuration (including targets) of all iSCSI ports on these host before importing the targets.

```
brcmhbacmd ImportiSCSI targets.xml clean
```

The example output:

```
All existing targets have been removed.
Added target iqn.2006-01.com.openfiler:target121-000.on.port.00-00-c9-be-1a-24
Added target iqn.2006-01.com.openfiler:target121-001.on.port.00-00-c9-be-1a-24
Added target iqn.2006-01.com.openfiler:target122-000.on.port.00-00-c9-2f-45-1b
Added target iqn.2006-01.com.openfiler:target122-001.on.port.00-00-c9-2f-45-1b
```

5.11.19 iSCSIPing

This command issues ICMP echo requests to an iSCSI target.

Supported By

Linux and Windows

Syntax

```
iSCSIPing <MAC_Address> <IP_Address>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
IP_Address	IP address of target to send ICMP echo request. You can specify an IPv4 or IPv6 IP address.

5.11.20 ListSessions

This command lists all the sessions on a specified target. The command gathers the information from the iSCSI target indicated by the <*iscsi_target_name*> parameter.

Supported By

Linux and Windows

Syntax

```
ListSessions <MAC_Address> <iscsi_target_name>
```

Parameters

MAC_Address	The MAC address of the an iSCSI function.
iscsi_target_name	Target's iSCSI name enclosed in quotation marks. The string length is 11 to 255.

5.11.21 RemoveTarget

This command removes the target with the specified iSCSI target name <*iscsi_target_name*>.

Supported By

Linux and Windows

Syntax

```
RemoveTarget <MAC_Address> <iscsi_target_name>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
iscsi_target_name	Target's iSCSI name enclosed in quotation marks. The string length is 11 to 255.

5.11.22 RemoveTargetPortal

This command removes the Target Portal containing the IPv4 or IPv6 target IP <*Target_IP*> and the port <*Port*> from the list of target portals for the specified initiator.

Supported By

Linux and Windows

Syntax

```
RemoveTargetPortal <MAC_Address> <Target_IP> <Port>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
Target_IP	The target portal's IP address. You can specify an IPv4 or IPv6 target IP address.
Port	The port number of the target portal. The possible values are 1024 to 65535.

5.11.23 SetBootTargetSession

This command enables and disables a iSCSI target's session as a boot session. If a session is enabled as a boot session and the system reboots, it attempts to boot from that target.

Depending on whether the target is logged in, there are two different ways to specify the session. If the target is logged in, use the <*TSIH*> parameter. If the target is not logged in, use the <*ISID_Qual Target_IP*> parameter. You can specify an IPv4 or IPv6 target IP address.

Supported By

Linux and Windows

Syntax

```
SetBootTargetSession <MAC Address> <Target> <TSIH | <ISID_Qual Target_IP>> <0|1>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
Target	Specifies the iSCSI name of the desired iSCSI target.
TSIH	TSIH value of the session. The possible values are 1 to 65535.
ISID_Qual Target_IP	ISID qualifier of the session. The possible values are 0 to 65535. You can specify an IPv4 or IPv6 Target IP address that the target is using depending on the adapter type.
0|1	Specifies the state of the BootProperty for the target: 0 = Disabled 1 = Enabled

5.11.24 SetInitiatorProperties

This command sets the initiator properties for the specified iSCSI function. It allows you to specify an initiator name <*Initiator_Name*> and an initiator alias <*Initiator_Alias*>. If you opt not to specify these fields, a default iSCSI name is assigned.

Except for the <*Initiator_Name*> and <*Initiator_Alias*> properties, these properties are set as the target portal's login properties to be used when detecting the targets on the target portal. The targets inherit the target portal's properties when they are detected. The detected target's login properties can be changed using the `SetTargetProperties` command.

If you set the authentication method <*Auth*> to a value other than `0`, you must set additional parameters. Each string must be enclosed in quotation marks to avoid mishandling by the Windows, Linux, Solaris, or VMware shell's parser. Additionally, these properties are used for iSNS target detection to set the detected target's login properties.

- If you set the authentication method to One-Way CHAP (<*Auth*>=1), you must also specify the *TgtCHAPName* and *TgtSecret*, which are used by the target to authenticate the initiator. For example:

  ```
  brcmhbacmd SetInitiatorProperties 00-11-22-33-44-55 Auth=1 "TgtChapName"
  "TargetSecret1"
  ```

- If you set the authentication method to Mutual CHAP (<*Auth*>=2), you need to specify the *TgtCHAPName* and *TgtSecret* and the *InitCHAPName* and *InitSecret*. The *InitCHAPName* and *InitSecret* are used for the initiator to authenticate the target. For example:

  ```
  brcmhbacmd SetInitiatorProperties 00-11-22-33-44-55 Auth=2 "TgtChapName"
  "TargetSecret1" "InitCHAPName" "InitialSecret1"
  ```

Supported By

Linux and Windows

Syntax

```
SetInitiatorProperties <MAC_Address> [Initiator_Name="initiator_name"]
[Initiator_Alias="initiator_alias"] [ImmediateData=<0|1>] [HeaderDigest=<0|1>]
[DataDigest=<0|1>] [Auth=<0|1|2> "TgtCHAPName" "TgtSecret" "InitCHAPName"
"InitSecret"]
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
Initiator_Name	Initiator iSCSI name enclosed in quotation marks (string length: 1 to 224).
Initiator_Alias	Initiator iSCSI alias enclosed in quotation marks (string length: 0 to 32).
ImmediateData	0 = No 1 = Yes (default)

HeaderDigest	0 = None (default) 1 = CRC32C
DataDigest	0 = None (default) 1 = CRC32C
Auth	0 = None (default) 1 = One-Way CHAP 2 = Mutual CHAP
TgtCHAPNam	Target CHAP name enclosed in quotation marks (string length: 1 to 256). Required when Auth =1 or 2.
TgtSecret	Target Secret enclosed in quotation marks (string length: 12 to 16). Required when Auth =1 or 2.
InitCHAPName	Initiator CHAP name enclosed in quotation marks (string length: 1 to 256). Required when Auth =2.
InitSecret	Initiator Secret enclosed in quotation marks (string length: 12 to 16). Required when Auth =2.

5.11.25 SetiSCSIBoot

This command defines whether a specific iSCSI function's boot ROM is active.

Supported By

Linux and Windows

Syntax

```
SetiSCSIBoot <MAC_Address> <0|1>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
0|1	Specifies the iSCSI boot ROM state: 0 = Disabled 1 = Enabled

5.11.26 SetNetworkConfiguration

This command sets the TCP/IP configuration on a specified port. The required fields for this command depend on the values set for <*DHCP*> and <*VLAN_Enabled*>.

Supported By

Linux and Windows

Syntax

```
SetNetworkConfiguration <MAC address> VLAN_Enabled=<0|1>[VLAN_ID=<0-4094>
Priority=<0-7>]DHCP=<0|1> [<IPv4_Address> <Subnet> [IPv4_Gateway]] [AA=<0|1>
[<LL=IPv6_Address> [RA1=IPv6_Address RA2=IPv6_Address] [GW6=IPv6_Address]]]
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
VLAN_Enabled	0 = Disabled 1 = Enabled

VLAN_ID	VLAN ID of the interface. The possible values are 0 to 4095. This value is specified only if VLAN is enabled.
Priority	VLAN priority of the interface. The possible values are 0 to 7. This value is specified only if VLAN is enabled.
DHCP	Dynamic Host Configuration Protocol for automatic IPv4 address assignment. 0 = Disabled 1 = Enabled If the DHCP parameter is disabled, you must specify its parameters. For example: IPv4_Address, Subnet, and IPv4_Gateway.
IPv4_Address	IPv4 address of initiator port. For example: 10.192.1.1. IPv4_Address is required if DHCP is disabled.
Subnet	Subnet mask of initiator port. For example: 255.255.255.0. Subnet is required if DHCP is disabled.
IPv4_Gateway	IPv4 gateway of initiator port. For example: 10.192.1.1. IPv4_Gateway is optional if DHCP is disabled.
AA	Automatic IPv6 address assignment. 0 = Disabled 1 = Enabled If the AA (automatic assignment) parameter is disabled, you must specify its parameters. For example, LL: RA1, RA2 and GW6.
LL	Link local IPv6 address. LL is required if AA is disabled, and it is ignored if AA is enabled.
RA1	Routable IPv6 address 1. RA1 is ignored if AA is enabled.
RA2	Routable IPv6 address 2. RA2 is ignored if AA is enabled.
GW6	IPv6 gateway address. GW6 is ignored if AA is enabled.

Considerations

- *VLAN_ID* and *Priority* are required only if *VLAN_Enabled* is enabled; otherwise, these values must be omitted.
- *IPv4_Address* and *Subnet* are required only if *DHCP* is disabled; otherwise these values must be omitted.
- At a minimum, the *DHCP* or *AA* parameter must be specified, and if required, their associated parameters as well. If either the *DHCP* or *AA* parameter is not specified, an error results.
- The *AA*, *LL*, *RA1*, *RA2*, and *GW6* parameters are valid only on OCe14000-series adapters.

5.11.27 SetTargetLoginProperties

This command sets the login and authentication properties associated with a specific target. This command requires that you specify a valid iSCSI target name <*iscsi_target_name*>. If you do not specify some of the remaining properties, these options are set to their default values. However, if no properties are changed, an error is generated. You must change at least one property for this command to be returned successfully.

If you set the authentication method <*Auth*> to a value other than 0, you must set additional parameters. Each string must be enclosed in quotation marks to avoid mishandling by the Windows, Linux, Solaris, or VMware shell's parser.

- If you set the authentication method to One-Way CHAP (<*Auth*>=1), you must also specify the *TgtCHAPName* and *TgtSecret*, which is used by the target to authenticate the initiator. For example:

```
brcmhbacmd SetTargetLoginProperties 00-11-22-33-44-55 iscsitarget Auth=1
"TgtCHAPName" "TargetSecret1"
```

- If you set the authentication method to Mutual CHAP (<*Auth*>=2), you must specify the *TgtCHAPName* and *TgtSecret*, and the *InitCHAPName* and *InitSecret*. The *InitCHAPName* and *InitSecret* are used for the initiator to authenticate the target.

  For example:

```
brcmhbacmd SetTargetLoginProperties 00-11-22-33-44-55 iscsitarget Auth=2
"TgtChapName" "TargetSecret1" "InitCHAPName" "InitialSecret1"
```

Supported By

Linux and Windows

Syntax

```
SetTargetLoginProperties <MAC_Address> <iscsi_target_name> [ImmediateData=<0|1>]
[HeaderDigest=<0|1>] [DataDigest=<0|1>] [Auth=<0|1|2> "TgtCHAPName" "TgtSecret"
"InitCHAPName" "InitSecret"]
```

Parameters

MAC_Address	The MAC address of an iSCSI port.
iscsi_target_name	Target's iSCSI name enclosed in quotation marks (string length: 11 to 255).
ImmediateData	0 = No 1 = Yes (default)
HeaderDigest	0 = None (default) 1 = CRC32C
DataDigest	0 = None (default) 1 = CRC32C
Auth	0 = None (default) 1 = One-Way CHAP 2 = Mutual CHAP
TgtCHAPNam	Target CHAP name enclosed in quotation marks (string length: 1 to 256). Required if Auth =1 or 2.
TgtSecret	Target Secret enclosed in quotation marks (string length: 12 to 16). Required if Auth =1 or 2.
InitCHAPName	Initiator CHAP name enclosed in quotation marks (string length: 1 to 256). Required if Auth =2.
InitSecret	Initiator Secret enclosed in quotation marks (string length: 12 to 16). Required if Auth =2.

5.11.28 SetTargetProperties

This command sets the ETO value of an iSCSI target.

Supported By

Linux and Windows

Syntax

```
SetTargetProperties <MAC_Address> <iscsi_target_name> <ETO>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
iscsi_target_name	Target's iSCSI name enclosed in quotation marks (string length: 11 to 255).

ETO The extended time out option for the target:
- For Windows, valid values are 0 to 3600
- For Linux and Solaris, valid values are 0 to 30

5.11.29 SetTPLoginProperties

This command sets a target portal's login properties. This command requires that you specify a valid IPv4 or IPv6 Target IP <*Target_IP*> and Port <*Port*>. However, if you specify no options other than the *Target_IP* and *Port*, no changes are made. You must change at least one of the optional parameters for this command to make any changes to the target portal's login properties.

These properties are used when detecting the targets on the target portal. The targets inherit the target portal's properties when they are detected. Targets already detected do not inherit the updated properties, only newly detected targets inherit the properties.

If you set the authentication method <*Auth*> to a value other than 0, you must set additional parameters. Each string must be enclosed in quotation marks to avoid mishandling by the Windows, Linux, Solaris, or VMware shell's parser.

- If you set the authentication method to One-Way CHAP (<*Auth*>=1), you must also specify the *TgtCHAPName* and *TgtSecret*. For example:

```
brcmhbacmd SetTPLoginProperties 00-11-22-33-44-55 10.192.1.1 5050 Auth=1
"TgtChapName" "TargetSecret1"
```

- If you set the authentication method to Mutual CHAP (<*Auth*>=2), you need to specify the *TgtCHAPName* and *TgtSecret*, and the *InitCHAPName* and *InitSecret*. The *InitCHAPName* and *InitSecret* are used for the initiator to authenticate the target. For example:

```
brcmhbacmd SetTPLoginProperties 00-11-22-33-44-55 10.192.1.1 5050 Auth=2
"TgtChapName" "TargetSecret1" "InitCHAPName" "InitialSecret1"
```

Supported By

Linux and Windows

Syntax

```
SetTPLoginProperties <MAC_Address> <Target_IP> <Port> [ImmediateData=<0|1>]
[HeaderDigest=<0|1>] [DataDigest=<0|1>] [Auth=<0|1|2> TgtCHAPName TgtSecret
InitCHAPName InitSecret]
```

Parameters

MAC_Address	The MAC address of an iSCSI port.
Target_IP	The IP address of the target portal. You can specify an IPv4 or IPv6 Target IP address.
Port	The port number of the target portal (value: 1024 to 65535).
ImmediateData	0 = No 1 = Yes (default)
HeaderDigest	0 = None (default) 1 = CRC32C
DataDigest	0 = None (default) 1 = CRC32C
Auth	0 = None (default) 1 = One-Way CHAP 2 = Mutual CHAP
TgtCHAPNam	The Target CHAP name enclosed in quotation marks (string length: 1 to 256). Required if Auth =1 or 2.

TgtSecret	The Target Secret enclosed in quotation marks (string length: 12 to 16). Required if Auth =1 or 2.
InitCHAPName	The Initiator CHAP name enclosed in quotation marks (string length: 1 to 256). Required if Auth =2.
InitSecret	The Initiator Secret enclosed in quotation marks (string length: 12 to 16). Required if Auth =2.

5.11.30 ShowARPTable

This command shows the current ARP table for the specified iSCSI function.

Supported By

Linux and Windows

Syntax

```
ShowARPTable <MAC_Address>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.

5.11.31 ShowiSNSServer

This command shows the currently configured Internet Storage Name Server. OCe11000-series adapters support only one iSNS server. OCe14000-series adapters support up to four iSNS servers, and they include IPv6 addresses in addition to IPv4 addresses.

Supported By

Linux and Windows

Syntax

```
ShowiSNSServer <MAC_Address>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.

5.11.32 ShowRouteTable

This command shows the route table for an iSCSI function.

Supported By

Linux and Windows

Syntax

```
ShowRouteTable <MAC_Address>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.

Example

```
brcmhbacmd ShowRouteTable 00-00-c9-a0-ce-77
```

5.11.33 ShowTarget

This command shows the properties for a specified target or all targets for an iSCSI function. If you do not specify the iSCSI target name, all targets and their associated properties are returned.

Supported By

Linux and Windows

Syntax

```
ShowTarget <MAC_Address> [iscsi_target_name | refreshtargets]
```

NOTE Only a single command option can be specified with this command. That is, you can specify only *iscsi_target_name* or *refreshtargets*.

Parameters

MAC_Address	The MAC address of an iSCSI function.
iscsi_target_name	iSCSI target name of a specific target. If not specified, all targets for the iSCSI function are displayed.
refreshtargets	Refresh all targets before displaying the information.

5.11.34 ShowTargetPortal

This command shows the properties for a specified Target Portal. If the <Target_IP> and <Port> are not specified, all Target Portals and their associated properties return.

Supported By

Linux and Windows

Syntax

```
ShowTargetPortal <MAC_Address> [<Target_IP> <Port>]
```

Parameters

MAC_Address	The MAC address of an iSCSI port.
Target_IP	IP address of the target portal. You can specify an IPv4 or IPv 6 address.
Port	Port number of the target portal.

5.11.35 TargetLogin

This command logs in to a target. The iSCSI target name <*iscsi_target_name*> is the only mandatory option. The <*target_portal*> and <*port*> information are optional, and if they are not provided, a default target portal is used. If you do not specify the remaining options, these options are set to their default values.

If you set the authentication method <*Auth*> to a value other than 0, you must set additional parameters. Each string must be enclosed in quotation marks to avoid mishandling by the Windows, Linux, Solaris, or VMware shell's parser.

- If you set the authentication method to One-Way CHAP (<*Auth*>=1), you must also specify the *TgtCHAPName* and *TgtSecret*, which is used by the target t authenticate the initiator. For example:

```
brcmhbacmd TargetLogin 00-11-22-33-44-55 iscsitarget Auth=1 "TgtChapName"
"TargetSecret1"
```

- If you set the authentication method to Mutual CHAP (<*Auth*>=2), you need to specify the *TgtCHAPName* and *TgtSecret*, and the *InitCHAPName* and *InitSecret*. The *InitCHAPName* and *InitSecret* are used for the initiator to authenticate the target. For example:

```
brcmhbacmd TargetLogin 00-11-22-33-44-55 iscsitarget Auth=2 "TgtChapName"
"TargetSecret1" "InitCHAPName" "InitialSecret1"
```

Supported By

Linux and Windows

Syntax

```
TargetLogin <MAC address> <Target Name> [<Target IP> <Port>]
[HeaderDigest=<0|1>][ImmediateData=<0|1>][DataDigest=<0|1>][Auth=<0|1|2>
["TgtCHAPName"] ["TgtSecret"] ["InitCHAPName"] ["InitSecret"]]
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
target_name	The target's iSCSI name enclosed in quotation marks (string length: 11 to 255 characters).
Target IP	Target IP Address. You can specify an IPv4 or IPv6 Target IP address.
Port	The port number of the target portal. The possible values are 1024 to 65535.
HeaderDigest	0 = None (default) 1 = Uses CRC32C checksum
ImmediateData	0 = No 1 = Yes (default) for appending solicited data to a command.
DataDigest	0 = None (default) 1 = CRC32C checksum
Auth	0 = None (default) 1 = One-Way CHAP 2 = Mutual CHAP
TgtCHAPName	The Target CHAP name enclosed in quotation marks (string length: 1 to 255 characters). The Target CHAP name is required it you set Auth to 1 or 2.
TgtSecret	The Target Secret enclosed in quotation marks (string length: 12 to 16 characters). The Target CHAP name is required it you set Auth to 1 or 2.
InitCHAPName	The Initiator CHAP name enclosed in quotation marks (string length: 1 to 255 characters). The Target CHAP name is required it you set Auth to 2.
InitSecret	The Initiator Secret enclosed in quotation marks (string length: 12 to 16 characters). The Target CHAP name is required it you set Auth to 2.

5.11.36 TargetLogout

This command logs out of an iSCSI target session.

You must specify the iSCSI target name <*iscsi_target_name*> and either the TSIH <*TSIH*> of the session, or the session's ISID Qualifier <*ISID_Qual*> and the target's IP address <*Target_IP*>.

Supported By

Linux and Windows

Syntax

```
TargetLogout <MAC_Address> <iscsi_target_name> <TSIH | <ISID_Qual Target_IP>>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
iscsi_target_name	The target's iSCSI name enclosed in quotation marks (string length: 11 to 255).
TSIH	The TSIH value of the session. The possible values are 1 to 65535.
ISID_Qual	The ISID qualifier of the session. The possible values are 0 to 65535.
Target_IP	The target's IP address.

5.11.37 UpdateiSNSServer

This command updates a configured iSNS server and accepts IPv4 or IPv6 iSNS server addresses. This command requires the server IP <*Server_IP*> and port number <*Port*> of the iSNS server to be available to respond to the iSNS requests. If a single iSNS server is configured, this command replaces the single iSNS server IP address with a new IP address. If no iSNS server is configured, this command adds a new iSNS server.

NOTE If two or more iSNS servers are already configured, any attempt to use this command results in an error.

Supported By

Linux and Windows

Syntax

```
UpdateiSNSServer <MAC_Address> <Server_IP> <Port>
```

Parameters

MAC_Address	The MAC address of an iSCSI function.
Server_IP	IP address of the iSNS server to configure. You can specify an IPv4 or IPv6 server IP address.
Port	Port number of the iSNS server to configure (value: 1024 to 65535).

5.12 LUN Masking Commands

The LUN Masking Commands group manage LUN masking activities. LUN Masking commands are supported only for FCoE functions.

NOTE Linux does not support LUN masking commands.

NOTE Solaris and Windows + CIM Provider on a VMware host do not support the following commands:

- `GetLunUnMaskbyHBA`
- `GetLunUnMaskbyTarget`
- `RescanLuns`
- `SetLunMask`

5.12.1 GetLunList

This command queries for the presence of any masked LUNs.

Supported By

Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetLunList <HBA WWPN> <Target WWPN> <Option>
```

Parameters

HBA WWPN	The WWPN of an FCoE function on the adapter.
Target WWPN	The WWPN of the target.
Option	0 = Get information from the driver 1 = Get information from the configuration

5.12.2 GetLunUnMaskByHBA

This command queries for the presence of any unmasked LUNs by FCoE functions.

Supported By

Windows

Syntax

```
GetLunUnMaskByHBA <HBA WWPN> <Option>
```

Parameters

HBA WWPN	The WWPN of an FCoE port.
Option	0 = Get information from the driver 1 = Get information from the configuration

5.12.3 GetLunUnMaskByTarget

This command queries for any unmasked LUNs by target.

Supported By

Windows

Syntax

```
GetLunUnMaskByTarget <HBA WWPN> <Target WWPN> <Option>
```

Parameters

HBA WWPN	The WWPN of an FCoE function.
Target WWPN	The WWPN of the target.
Option	0 = Get information from the driver 1 = Get information from the configuration

5.12.4 RescanLuns

This command rescans LUNs to find any new LUNs.

Supported By

Windows

Syntax

```
RescanLuns <HBA WWPN> <Target WWPN>
```

Parameters

HBA WWPN	The WWPN of an FCoE function.
Target WWPN	The WWPN of the target.

5.12.5 SetLunMask

This command masks the specified LUNs.

Supported By

Windows

Syntax

```
SetLunMask <HBA WWPN> <Target WWPN> <Option> <Lun> <LunCount> <MaskOp>
```

Parameters

HBA WWPN	The WWPN of an FCoE function.
Target WWPN	The WWPN of the target.
Option	0 = Get information from the driver 1 = Get information from the configuration (make persistent) 2 = Send information to both
Lun	The starting LUN number.
LunCount	The number of LUNs.
MaskOp	A = Mask LUN B = Clear unmask target level C = Clear unmask HBA level D = Unmask LUN E = Unmask target level F = Unmask HBA level

5.13 Miscellaneous Commands

Commands in the Miscellaneous Command group do not fit in other groups. See specific commands for adapter limitations.

5.13.1 AddHost

This command adds a host to the hosts file for remote TCP/IP management in the OneCommand CNA Manager application. The adapters for these hosts are also presented by the `ListHBAs` command (see Section 5.13.9, ListHBAs).

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

To add non-VMware ESXi hosts:

```
AddHost <hostname|IP_address>[:Port_Number]
HostName:          Name of host to add to hosts file
IP_address:        IP address (IPv4 or IPv6) of host to add to hosts file
                   Example IPv4: 10.192.80.102
                   Example IPv6: fe80::6445:80e9:9878:a527
Port_Number:       Optional IP port number to access remote host
                   Example: 10.192.80.102:9876
                   Note: When specifying IPv6 address with Port_Number, it must be
                   enclosed in [].
                   Example: [fe80::6445:80e9:9878:a527]:9876
```

- An attempt is made to contact the host to confirm remote access before adding it to the host list. If the attempt fails, the host is not added.
- The h option (for specifying an optional IP address or host name) after `brcmhbacmd` is not available for the `AddHost` command.

To add VMware ESXi hosts to Windows using the OneCommand CNA Manager application:

```
m=cim [u=<username>] [p=<password>] [n=<namespace>] AddHost  <IP_Address>
```

If the *username*, *password*, and *namespace* are not specified, see Section 4.4.2.1.1, Default CIM Credentials.

Parameters

host_address	The IP address (using the IPv4 or IPv6 format) or the host name.

5.13.2 CnaClearEventLog

This command clears the event log for the adapter specified by the WWPN or MAC address

Supported By

Linux and Windows

Syntax

```
CnaClearEventLog <WWPN|MAC>
```

Parameters

WWPN	The WWPN of an FCoE function on the adapter.
MAC	The MAC address of a NIC or iSCSI function on the adapter.

5.13.3 CnaGetEventLog

This command shows the adapter event log for the adapter specified by the WWPN or MAC address.

NOTE This command is supported only for OneConnect adapters.

Supported By

Linux and Windows

Syntax

```
CnaGetEventLog <WWPN|MAC>
```

Parameters

WWPN	The WWPN of an FCoE function on the adapter.
MAC	The MAC address of a NIC or iSCSI function on the adapter.

5.13.4 Download

This command downloads a firmware image to the port function or adapter specified by the WWPN or MAC address.

Considerations

- If a secure version of firmware (version 11.0 or later) is installed on an OCe14000B-series adapter and you want to update to an earlier unsecured version of firmware, you must remove the secure firmware jumper block before performing the update. See the installation manual for the adapter for more information.
- If you attempt to update unauthenticated firmware for a secure OCe14000B-series, adapter, the following error message is displayed:

  ```
  ERROR: Download Failed due to invalid firmware digital signature. Please
  contact customer support for additional help.
  ERROR: <203>: Failed validating firmware digital signature
  ```

- If you attempt to update unsecured firmware for a secure OCe14000B-series adapter, the following error message is displayed:

  ```
  ERROR: Download Failed due to missing digital signature in firmware file.
  Please contact customer support for additional help.
  ERROR: <209>: Firmware digital signature missing
  ```

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
Download <WWPN|MAC> <FileName>
```

Parameters

WWPN	The WWPN of an FCoE function on the adapter.
MAC	The MAC address of a NIC or iSCSI function on the adapter.
FileName	The name and location of the firmware image (any file accessible to the CLI client).

5.13.5 ExportSANInfo

For reporting purposes, this command captures the SAN information in `.xml` for XML-formatted files and `.csv` for CSV-formatted files.

NOTE This command can take a long time on large SAN configurations because of the large amount of information that must be obtained and reported. The output can also be redirected to a file if required.

Supported By

Linux, Solaris, and Windows

Syntax

```
ExportSANInfo [format]
```

NOTE The h option (for specifying an optional IP address or host name) after `brcmhbacmd` is not available for the `ExportSANInfo` command.

Parameters

format — An optional parameter that specifies the format of the adapter information:
- csv
- xml

Leaving the format blank shows the data in xml format (default).

5.13.6 GetCimCred

This command shows the default credentials set for the CIM client.

NOTE The password is encrypted.

Supported By

Windows

Syntax

```
GetCimCred
```

Parameters

None.

5.13.7 GetQoSInfo

This command shows the QoS information for a specified NIC function if multichannel support is enabled for the port on which the NIC function exists.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetQoSInfo <MAC_Address>
```

Parameters

MAC The MAC address of a NIC function.

Example

```
brcmhbacmd GetQoSInfo 00-00-c9-93-20f-d6
```

5.13.8 GetVPD

This command shows the port's VPD.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetVPD <WWPN|MAC>
```

Parameters

WWPN The WWPN of an FCoE function.

MAC The MAC address of a NIC or iSCSI function.

5.13.9 ListHBAs

This command shows a list of the manageable Broadcom Emulex adapters found by local discovery. For a NIC-only or iSCSI adapter, the MAC address is displayed rather than the port WWN. The node WWN and fabric WWN are not displayed. The type of information listed may vary according to the adapter model.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
ListHBAs [local] [m=model] [pt=type] [down]
```

Parameters

local	Displays only local adapters. For Dell adapters, this parameter displays virtual WWPNs and WWNNs for each FCoE function and virtual MAC addresses for each NIC function. It does not display permanent settings.
m=model	Model filter. Append * to the end of the model name for a wildcard match. For example: OCe14*
pt=type	The port type filter. Valid types are NIC, RoCE, iSCSI, and FCoE.
down	Displays only the NIC functions of adapters on the local system in which the adapter's ARM processor has stopped. This parameter detects adapters that need to have a core dump collected in addition to adapters that might not respond to commands from the OneCommand CNA Manager CLI or application.

5.13.10 ListVFuncs

This command lists the virtual functions (SR-IOV) on a specified NIC function.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
ListVFuncs <MAC>
```

Parameters

MAC The MAC address of a NIC function.

Example

```
brcmhbacmd ListVFuncs 00-00-c9-12-34-56
```

The example output:

```
Virtual Functions for 00-00-c9-12-34-56:

MAC Address    : 00-00-c9-12-34-ab
VLADID         : 10
Transmit Rate  : 100 Mbit/sec.

MAC Address    : 00-00-c9-12-34-cd
VLADID         : 10
Transmit Rate  : 100 Mbit/sec.

MAC Address    : 00-00-c9-12-34-ef
VLADID         : 10
Transmit Rate  : 1 bbit/sec.

MAC Address    : 00-00-c9-13-34-01
VLADID         : 20
Transmit Rate  : 1 Gbit/sec.
```

5.13.11 RemoveHost

This command removes a host from the hosts file use for TCP/IP management in the OneCommand CNA Manager application GUI. The <*host_address*> can be an IP address that uses the IPv4 or IPv6 format, or it can be a host name.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

For the remote management interface:

```
RemoveHost host_address
```

For VMware ESXi using the CIM interface:

```
m=cim RemoveHost <IP_Address>
```

NOTE The h option (for specifying an optional IP address or host name) after `brcmhbacmd` is not available for the `RemoveHost` command.

Parameters

host_address	The host to remove.
IP_Address	The IP address of the host to remove.

5.13.12 Reset

This command resets the FCoE function. A reset can require several seconds to complete, especially for remote devices. When the reset is completed, the system command prompt is displayed.

NOTE This command applies only to FCoE functions.

This command only resets the driver to update changed driver parameters that require a driver reset. It does not cause a hardware reset of the FCoE function.

Supported By

Linux, Solaris, Windows, and VMware ESXi from a Windows remote host

Syntax

```
Reset <WWPN>
```

Parameters

WWPN	The WWPN of an FCoE function.

5.13.13 SetCimCred

This command sets the default CIM credentials. You must specify all four credentials: *`username`*, *`password`*, *`namespace`*, and *`portnum`*. Default credentials are used if any credential is not in the `brcmhbacmd` command argument. After the default credentials for a host are set, any other command can be issued by specifying `m=cim`.

Supported By

Windows

Syntax

```
SetCimCred <username> <password> <namespace> <portnum>
```

NOTE Use this command to set only the CIM credentials. After this is finished, subsequent `brcmhbacmd` commands do not require you to specify the CIM credentials in the command line.

Parameters

username	The login user ID of the VMware ESXi.
password	The login password of the VMware ESXi.
namespace	The namespace where the Emulex CIM provider is registered in the SFCB CIMOM of VMware ESXi, specifically `root/brcmccx`.
portnum	The port number of the SFCB that CIMOM is listening to, that is, 5988 (HTTP) or 5989 (HTTPS).

5.13.14 SRIOVEnable

For OCe11000-series adapters, this command enables/disables SR-IOV only on the specified NIC function. For OCe14000-series adapters, this command enables or disables SR-IOV on all NIC functions on the adapter. This command can be used to enable or disable SR-IOV with the NPar or NParEP mode enabled.

This command is not available on OCe11101-EM/EX or OCe11102-EM/EX adapters. The following error is returned:

```
ERROR: <251>: Hardware or firmware does not support command.
```

SRIOVEnable returns an error if channel management is enabled.

For OCe14000-series adapters, it is preferable to use the SetAdapterPortConfig command (see Section 5.16.4, SetAdapterPortConfig (for OCe14000-Series Adapters)) to enable or disable SR-IOV.

NOTE SR-IOV is not supported with RoCE configurations.

Supported By

Linux, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
SRIOVEnable <MAC> <0|1>
```

Parameters

MAC	The MAC address of a NIC function.
0|1	0 = Disables SR-IOV 1 = Enables SR-IOV

Examples

The following command enables SR-IOV on a NIC function with MAC address 00-00-c9-12-34-56:

```
brcmhbacmd SRIOVEnable 00-00-c9-12-34-56 1
```

The following command enables or disables SR-IOV with NPar or NParEP mode enabled:

```
>brcmhbacmd SriovEnable 00-90-fa-41-22-f0 1

SR-IOV successfully enabled in all ports. A system reboot is required.
```

NOTE You can use the GetAdapterPortConfig command to display the SR-IOV state.

```
brcmhbacmd getAdapterPortConfig 00-90-fa-41-22-f0

Number of Ports             : 2
Maximum FCoE Functions      : 1
Maximum iSCSI Functions     : 1
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : Yes
ROCE Functions Allowed      : Yes
RebootRequired              : Yes

Active Configuration
  Profile ID  : 16
  MC Type     : None
  SR-IOV State: Disabled
```

```
  Port Assignments:
    Port 1: NIC
    Port 2: NIC

Reboot Configuration
  Profile ID  : 16
  MC Type     : NPar
  SR-IOV State: Enabled
  NParEP Mode : Disabled
  Port Assignments:
    Port 1 (p0): NIC, NIC, NIC, NIC
    Port 2 (p1): NIC, NIC, NIC, NIC

Note:  (pN) labels are the port parameter names to use in the SetAdapterPortConfig
command (e.g. p0=NIC,FCoE p1=NIC,iSCSI).
```

5.13.15 TargetMapping

This command shows a list of mapped targets and the LUNs for an FCoE function on a port.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`TargetMapping <WWPN>`

Parameters

WWPN	The WWPN of an FCoE adapter.

5.13.16 VEPAEnable

This command enables or disables VEPA Management. VEPA Management is supported only for OCe14000-series NIC ports that have SR-IOV enabled. For all other board types and port types, this feature is undefined, and this command returns a `Hardware Does Not Support` error. In addition, when the next boot configuration does not have SR-IOV enabled or if multichannel is enabled, the following error message is displayed:

```
ERROR: <223>: VEPA is only settable on NIC ports when the reboot
configuration is non-multichannel and SR-IOV is enabled.
```

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

`VEPAEnable <MAC> <0|1>`

Parameters

MAC	The MAC address of the NIC function.
0|1	The VEPA state: 0 = Disabled 1 = Enabled The current VEPA state for the NIC function is displayed by the PortAttributes command. If the current or next boot configuration does not allow VEPA, the VEPA state is displayed as N/A.

NOTE If VEPA is currently enabled, when the system is rebooted with either SR-IOV disabled or multichannel enabled, VEPA is automatically disabled.

5.13.17 Version

This command shows the current version of the OneCommand CNA Manager CLI Client.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

For the remote management interface:

`Version`

NOTE The h option (for specifying an optional IP address or host name) after `brcmhbacmd` is not available for the `Version` command.

For VMware ESXi using the CIM interface:

`h=<IP address> m=cim Version`

Parameters

None

5.14 Persistent Binding Commands

The Persistent Binding Commands group facilitates persistent binding operations.

In a remote environment, you can perform persistent bindings operations from a host running any operating system (including Linux or VMware ESXi), but only to a remote host that is running Windows or Solaris.

For a binding to take effect immediately (that is, `SetPersistentBinding parameter`: `Scope` = `I` or B), the <*`SCSIBus`*> and <*`SCSITarget`*> parameters must match the SCSI bus and SCSI target to which the FCoE target is already automapped. If automapping is disabled, the binding takes effect immediately if the FCoE target is not already persistently bound, and the specified <*`SCSIBus`*> and <*`SCSITarget`*> parameters are available to be persistently bound. Also, the <*`BindType`*> parameter must match the currently active bind type. Otherwise, you are notified that you must reboot the system to cause the persistent binding to become active.

These commands are supported only on FCoE ports.

The following persistent binding commands are not supported on Linux or on Windows + CIM Provider on a VMware host:

- `BindingCapabilities`
- `BindingSupport`
- `PersistentBinding`
- `RemoveAllPersistentBinding`
- `RemovePersistentBinding`
- `SetPersistentBinding`
- `SetBindingSupport`

5.14.1 AllNodeInfo

This command shows target node information for each target accessible by the adapter.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
AllNodeInfo <WWPN>
```

Parameters

WWPN	The WWPN of an FCoE function.

5.14.2 BindingCapabilities

This command shows the binding capabilities of the adapter. If a binding is configured, it is maintained across reboots.

Supported By

Solaris and Windows

Syntax

```
BindingCapabilities <WWPN>
```

Parameters

WWPN	The WWPN of an FCoE function.

5.14.3 BindingSupport

This command shows the binding support for an FCoE function.

Supported By

Solaris and Windows

Syntax

```
BindingSupport <WWPN> <Source>
```

Parameters

WWPN	The WWPN of an FCoE function.
Source	C = Configuration support L = Live support

5.14.4 PersistentBinding

This command specifies the set of persistent binding information (configuration or live state) that is requested.

Supported By

Solaris and Windows

Syntax

```
PersistentBinding <WWPN> <Source>
```

Parameters

WWPN	The WWPN of an FCoE function.
Source	C = Configuration support L = Live support

5.14.5 RemoveAllPersistentBinding

This command removes all persisting bindings for an FCoE function.

Supported By

Solaris and Windows

Syntax

```
RemoveAllPersistentBinding <WWPN>
```

Parameters

WWPN	The WWPN of an FCoE function.

5.14.6 RemovePersistentBinding

This command removes persistent binding between an FCoE target and a SCSI bus and target. The binding to be removed can be to a target WWPN, a target WWNN, or a target D_ID.

Supported By

Solaris and Windows

Syntax

```
RemovePersistentBinding <WWPN> <BindType> <ID> <SCSIBus> <SCSITarget>
```

Parameters

WWPN	The WWPN of an FCoE function.
BindType	P = Remove binding by WWPN N = Remove binding by WWNN D = Remove binding by D_ID
ID	The type of ID based on <BindType>: — Target WWPN if <BindType> = P — Target WWNN if <BindType> = N Target D_ID if <BindType> = D
SCSIBus	The bus number of the SCSI device.
SCSITarget	The target number of the SCSI device.

5.14.7 SetBindingSupport

This command enables and sets the binding support for an FCoE function.

Supported By

Solaris and Windows

Syntax

```
SetBindingSupport <WWPN> <BindFlag>
```

Parameters

WWPN	The WWPN of an FCoE function.
BindFlag	The type of binding support for the adapter: D = Binding by D_ID (not available for Windows driver) P = Binding by WWPN N = Binding by WWNN (not available for Windows driver) A = Binding by automap (not available for Windows driver) DA = Binding by D_ID and automap PA = Binding by WWPN and automap NA = Binding by WWNN and automap

5.14.8 SetPersistentBinding

This command sets a persistent binding between an FCoE target and a SCSI bus target. The binding can be to a target WWPN, target WWNN, or target D_ID.

Supported By

Solaris and Windows

Syntax

```
SetPersistentBinding <WWPN> <Scope> <BindType> <TargetId> <SCSIBus> <SCSITarget>
```

Parameters

WWPN	The WWPN of an FCoE function.
Scope	P = Permanent binding (survives reboot) I = Immediate binding B = Binding is both permanent and immediate
BindType	P = Enable binding by WWPN N = Enable binding by WWNN D = Enable binding by D_ID
TargetId	If BindType = P, Target WWPN If BindType = N, Target WWNN If BindType = D, Target D_ID
SCSIBus	The bus number of the SCSI device.
SCSITarget	The target number of the SCSI device.

5.15 Personality Change Commands

The OneCommand CNA Manager application enables you to change the personality or protocol running on OneConnect adapters. If you change the personality of the adapter and reboot the host, the adapter starts running the new personality. The personalities that the OneConnect adapters currently run are NIC-only, NIC + FCoE, NIC + iSCSI, and NIC + RoCE for OCe14000-series adapters.

For RoCE-1, check the Knowledge Base on http://www.broadcom.com for any updated information on additional use cases for the RoCE-1 profile. Choose the RoCE-2 profile for Windows SMB Direct, or for Linux iSER, NFS, or VM Migration support.

NOTE It is possible to install one (or more) driver kits for the current personality profile, then change the personality profile and no longer have the drivers necessary to run the adapter. If you change personality profiles, you must install the appropriate drivers. The appropriate drivers are available at http://www.broadcom.com.

5.15.1 ChangePersonality

This command changes the personality on the adapter. It is the only way to change the personality of an OCe11000-series adapter. For an OCe14000-series adapter, `ChangePersonality` is an efficient way to set all ports on the adapter to run the same protocols.

However, the OCe14000-series adapter is capable of running different protocols on different ports. If you want to set different protocols on different ports, use the `SetAdapterPortConfig` command. See Section 5.16.4, SetAdapterPortConfig (for OCe14000-Series Adapters) for more information.

After a successful personality change, a system reboot is required.

NOTE RoCE is available only on OCe14000-series adapters.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
ChangePersonality <WWPN|MAC> <type> [Profile_ID]
```

Parameters

WWPN	The WWPN of an FCoE function on the adapter.
MAC	The MAC address of a NIC, NIC+RoCE, or iSCSI function on the adapter.
type	The personality type to which the adapter is being changed. The available values are shown when you are using the `ShowPersonalities` command. See Section 5.15.2, ShowPersonalities.
Profile_ID	Optionally specifies the profile ID number to used to identify the personality. This number can be obtained from Section 5.16.3, ListProfiles by looking at the Active Profile ID field. Profile_ID is only required in the following instances: — When specifying a NIC-only personality — When specifying a NIC+RoCE personality — When more than one NIC profile is available Selecting a profile that cannot be used with the specified personality results in an error.

Examples

Changing to FCoE personality:

```
brcmhbacmd ChangePersonality 00-12-34-56-78-9A fcoe
```

Changing to FCoE personality:

```
brcmhbacmd changepersonality 00-00-c9-12-34-56 fcoe
```

Changing to NIC + RoCE personality with Profile ID 21:

```
brcmhbacmd ChangePersonality 00-00-c9-12-34-56 nic+roce 21
brcmhbacmd ShowPersonalities 00-00-c9-12-34-56
Adapter Personalities:
NIC
NIC+RoCE (configured)
iSCSI
FCoE (active)
```

5.15.2 ShowPersonalities

This command displays the list of single personalities available on an adapter. The personality type is displayed as either NIC, NIC + RoCE, iSCSI, or FCoE.

NOTE If the active configuration is not a single personality (only on OCe14000-series), the active personality indicators are not displayed.

NOTE The available personalities are adapter dependent.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
ShowPersonalities <WWPN|MAC>
```

Parameters

WWPN	The WWPN of an FCoE function on the adapter.
MAC	The MAC address of a NIC, NIC + RoCE or iSCSI function on the adapter.

Example

```
>brcmhbacmd showpersonalities 00-00-c9-12-34-56
```

```
Adapter Personalities:
NIC
NIC+RoCE
iSCSI
FCoE (active) (configured)
```

5.16 Profile Management Commands

The Profile Management Commands group manages profile configuration for OCe14000-series adapters. You can display active and reboot port configurations for an adapter, list available profile IDs, and configure the function protocol for all ports on OCe14000-series adapters.

5.16.1 GetAdapterPortConfig

This command displays the active and reboot adapter port configurations of an OCe14000-series adapter. This includes the total number of ports, the maximum number of FCoE and iSCSI functions that can be defined, and whether mixed mode and concurrent mode storage are supported by the adapter.

NOTE All adapters support concurrent mode; therefore, up to two storage protocols can be configured on a port.

Additionally, it shows whether RoCE is supported on the adapter and whether a reboot is currently required to activate the new configuration.

The number of functions displayed per port depends on the current configuration of the adapter, such as custom mode, multichannel, or NPar.

In the information displayed for the reboot configuration, each of the ports displays the `pN` parameter name in parenthesis. The `pN` parameter is used in the `SetAdapterPortConfig` command to change the port's protocol assignments. This is useful when the physical port numbering on the adapter does not start at 0.

For multichannel and Dell NPar configurations, this command displays only the port protocol assignments. For the entire multi channel configuration, use the Section 5.5.1, CMGetParams.

NOTE If ARI is present on a host, this command may show more partitions per port than are actually running on the host.

If `Flex` is displayed in the output for the `GetAdapterPortConfig` command as the multi channel type for HP adapters currently running in Flex mode, you cannot specify `Flex` for the *`mctype`* parameter of the `SetAdapterPortConfig` command. The only *`mctype`* parameter that can be specified on HP UMC-capable adapters is UMC.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetAdapterPortConfig <MAC|WWPN>[table]
```

Parameters

MAC	MAC address of any iSCSI, or NIC, NIC + RoCE function on the adapter WWPN.
WWPN	WWPN of any FCoE function on the adapter.
table	The table option displays the output in tabular form, including all function and available protocols for each function.

Examples

The following table provides locations of application examples.

Example Types	Section
Tabular output	5.16.1.1
2-port adapter, with the table option	5.16.1.2
4-port adapter	5.16.1.3
2-port adapter, port number starting at 1	5.16.1.4
UMC	5.16.1.5
NIC+RoCE	5.16.1.6
Default output format with a disabled port	5.16.1.7
Tabular output format with a disabled port	5.16.1.8
Default output with port added back	5.16.1.9
Tabular output with port added back	5.16.1.10

5.16.1.1 Tabular Output

The adapter configuration information is displayed in a table format when the table (tab or tabular) option is specified.

```
>brcmhbacmd getadapterportconfig 00-90-FA-41-22-F0 table

Number of Ports             : 2
Maximum FCoE Functions      : 1
Maximum iSCSI Functions     : 1
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : Yes
ROCE Functions Allowed      : Yes
RebootRequired              : Yes

Active Configuration
  Profile ID  : 16
  MC Type     : None
  SR-IOV State: Enabled

  Port#  Func#  Protocol
  -----  -----  --------
    1       0   NIC
    2       1   NIC

Reboot Configuration
  Profile ID  : 16
  MC Type     : UMC  (Available MC Types: UMC)
  SR-IOV State: Disabled


  Port#  Func#  Protocol  Available Protocols for Function
  -----  -----  --------  ----------------------------------
  1 (p0)    0   NIC       NIC, NIC+RoCE
            2   NIC       none, NIC, iSCSI, FCoE
            4   NIC       none, NIC, iSCSI, FCoE
```

```
              6   NIC       none, NIC, iSCSI, FCoE
              8   none      none, NIC
             10   none      none, NIC
             12   none      none, NIC
             14   none      none, NIC
  2 (p1)     1   NIC       none, NIC, NIC+RoCE
              3   NIC       none, NIC, iSCSI, FCoE
              5   NIC       none, NIC, iSCSI, FCoE
              7   NIC       none, NIC, iSCSI, FCoE
              9   none      none, NIC
             11   none      none, NIC
             13   none      none, NIC
             15   none      none, NIC
```

Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig command (e.g. p0=NIC,FCoE p1=NIC,iSCSI).

5.16.1.2 2-port Adapter, with the Table Option

```
>brcmhbacmd getadapterportconfig 00-00-c9-12-34-56
Number of Ports             : 2
Maximum FCoE Functions      : 1
Maximum iSCSI Functions     : 1
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : Yes
ROCE Functions Allowed      : Yes
RebootRequired              : No

Active Configuration
  Profile ID  : 37
  MCType      : None
  SR-IOV State: disabled


  Port #  Func # Protocol
 ------  ------  --------
    0       0     NIC
            2     iSCSI
            4     FCoE
    1       1     NIC
            3     iSCSI
            5     FCoE

Reboot Configuration
  Profile ID  : 37
  MC Type     : None
  SR-IOV State: Disabled (Available Types: UMC)

  Port #  Func #  Protocol  Available Protocols for Function
  ------  ------  --------  ----------------------------------
  0 (p0)    0     NIC       NIC, NIC+RoCE
            2     iSCSI     NIC, iSCSI, FCoE
            4     FCoE      NIC, iSCSI, FCoE
            6     n/a       NIC
```

```
              8     n/a       NIC
              10    n/a       NIC
              12    n/a       NIC
              14    n/a       NIC
  1 (p1)      1     NIC       NIC, NIC+RoCE
              3     iSCSI     NIC, iSCSI, FCoE
              5     FCoE      NIC, iSCSI, FCoE
              7     n/a       NIC
              9     n/a       NIC
              11    n/a       NIC
              13    n/a       NIC
              15    n/a       NIC
```

Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).

5.16.1.3 4-Port Adapter

```
>brcmhbacmd getadapterportconfig 00-00-c9-12-34-56
Number of Ports             : 4
Maximum FCoE Functions      : 4
Maximum iSCSI Functions     : 4
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : No
ROCE Functions Allowed      : Yes
RebootRequired              : No

Active Configuration
  Profile ID  : 37
  MCType      : None
  SR-IOV State: disabled
  Port Assignments:
    Port 0: NIC,iSCSI
    Port 1: NIC,iSCSI
    Port 2: NIC,FCoE
    Port 3: NIC,FCoE

Reboot Configuration
  Profile ID  : 37
  MCType      : None
  SR-IOV State: disabled
  Port Assignments:
    Port 0 (p0): NIC,iSCSI
    Port 1 (p1): NIC,iSCSI
    Port 2 (p2): NIC,FCoE
    Port 3 (P3): NIC,FCoE
```

Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).

5.16.1.4 2-Port Adapter, Port Number Starting at 1

```
>brcmhbacmd getadapterportconfig 00-00-c9-12-34-56
Number of Ports             : 2
Maximum FCoE Functions      : 4
Maximum iSCSI Functions     : 4
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : Yes
ROCE Functions Allowed      : Yes
RebootRequired              : Yes

Active Configuration
  Profile ID  : 37
  MCType      : None
  SR-IOV State: disabled
  Port Assignments:
    Port 1: NIC,iSCSI,FCoE
    Port 2: NIC,iSCSI,FCoE

Reboot Configuration
  Profile ID  : 23
  MCType      : None
  SR-IOV State: disabled
  Port Assignments:
    Port 1 (p0): NIC,FCoE,None
    Port 2 (p1): NIC,FCoE,None
```

Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).

5.16.1.5 UMC

```
>brcmhbacmd getadapterportconfig 00-00-c9-12-34-56
Number of Ports             : 4
Maximum FCoE Functions      : 4
Maximum iSCSI Functions     : 4
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : Yes
ROCE Functions Allowed      : Yes
RebootRequired              : No

Active Configuration
  Profile ID  : 37
  MC Type     : UMC
  SR-IOV State: disabled
  Port Assignments:
    Port 0: NIC,iSCSI,FCoE,NIC
    Port 1: NIC,iSCSI,FCoE,NIC
    Port 2: NIC,FCoE,NIC,NIC
    Port 3: NIC,iSCSI,NIC,NIC

Reboot Configuration
  Profile ID  : 37
  MC Type     : UMC
  SR-IOV State: disabled
```

```
  Port Assignments:
    Port 0 (p0): NIC,iSCSI,FCoE,NIC
    Port 1 (p1): NIC,iSCSI,FCoE,NIC
    Port 2 (p2): NIC,FCoE,NIC,NIC
    Port 3 (p3): NIC,iSCSI,NIC,NIC

Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig
command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).
```

5.16.1.6 NIC+RoCE

```
>brcmhbacmd getadapterportconfig 00-00-c9-12-34-56
Number of Ports             : 4
Maximum FCoE Functions      : 4
Maximum iSCSI Functions     : 4
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : Yes
ROCE Functions Allowed      : Yes
RebootRequired              : No

Active Configuration
  Profile ID  : 21
  MC Type     : None
  SR-IOV State: disabled
  Port Assignments:
    Port 0: NIC+RoCE,None,None
    Port 1: NIC,None,None
    Port 2: NIC+RoCE,None,None
    Port 3: NIC,None,None

Reboot Configuration
  Profile ID  : 21
  MC Type     : None
  SR-IOV State: disabled
  Port Assignments:
    Port 0 (p0): NIC+RoCE,None,None
    Port 1 (p1): NIC,None,None
    Port 2 (p2): NIC+RoCE,None,None
    Port 3 (p3): NIC,None,None

Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig
command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).
```

5.16.1.7 Default Output Format with a Disabled Port

If you have removed a port, the `GetAdapterPortConfig` output no longer displays a port number or port name for the removed port, since this information becomes unavailable. The port number is replaced with the word `Removed` in the default output or `N/A` in the tabular output.

Example

```
Number of Ports             : 2
Maximum FCoE Functions      : 1
Maximum iSCSI Functions     : 1
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : Yes
ROCE Functions Allowed      : Yes
```

```
RebootRequired              : No

Active Configuration
  Profile ID  : 23
  MC Type     : None
  SR-IOV State: disabled
  Port Assignments:
    Port 1: NIC, FCoE, None
    Removed: (p1) : None, None, None

Reboot Configuration
  Profile ID  : 23
  MC Type     : None
  SR-IOV State: disabled
  Port Assignments:
    Port 1 (p0): NIC, FCoE, None
    Removed: None, None, None

Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig
command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).
```

5.16.1.8 Tabular Output Format with a Disabled Port

```
>brcmhbacmd setadapterportconfig 00-90-FA-41-22-F0

Active Configuration
  Profile ID  : 23
  MC Type     : None
  SR-IOV State: disabled

Port#    Func#   Protocol
 ------  ------  --------
    1       0     NIC
            2     FCoE
            4     None
    N/A     1     None
            3     None
            5     None

Reboot Configuration
  Profile ID  : 23
  MC Type     : None  (Available MC Types: vNIC1, SIMode, UFP)
  SR-IOV State: Disabled

  Port #  Func #  Protocol  Available Protocols for Function
  ------  ------  --------  ----------------------------------
  1 (p0)    0     NIC       NIC, NIC+RoCE
            2     iSCSI     None, NIC, iSCSI, FCoE
            4     FCoE      None, NIC, iSCSI, FCoE
            6     None      None, NIC
            8     None      None, NIC
            10    None      None, NIC
```

```
            12    None      None, NIC
            14    None      None, NIC
  N/A (p1)   1    None      None, NIC, NIC+RoCE
```

5.16.1.9 Default Output with Port Added Back

```
>brcmhbacmd getadapterportconfig 00-90-FA-41-22-F0

Number of Ports             : 2
Maximum FCoE Functions      : 1
Maximum iSCSI Functions     : 1
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : Yes
ROCE Functions Allowed      : Yes
RebootRequired              : No

Active Configuration
  Profile ID  : 23
  MC Type     : None
  SR-IOV State: Disabled
  Port Assignments:
    Port 1  : NIC, FCoE, None
    Disabled: None, None, None

Reboot Configuration
  Profile ID  : 23
  MC Type     : None
  SR-IOV State: Disabled
  Port Assignments:
    Port 1 (p0): NIC, FCoE, None
    Added (p1) : NIC, FCoE, None

Note:  (pN) labels are the port parameter names to use in the SetAdapterPortConfig
command (e.g. p0=NIC,FCoE p1=NIC,iSCSI).
```

5.16.1.10 Tabular Output with Port Added Back

```
>brcmhbacmd setadapterportconfig 00-90-FA-41-22-F0

Active Configuration
  Profile ID  : 23
  MC Type     : None
  SR-IOV State: Disabled

  Port#     Func#  Protocol
  -------   -----  --------
    1          0   NIC
               2   FCoE
               4   None
   N/A         1   None
               3   None
               5   None
```

```
Reboot Configuration
  Profile ID  : 23
  MC Type     : None  (Available MC Types: vNIC1, SIMode, UFP)
  SR-IOV State: Disabled

  Port#     Func#  Protocol  Available Protocols for Function
  -------   -----  --------  ---------------------------------
  1 (p0)        0   NIC       NIC, NIC+RoCE
                2   FCoE      None, NIC, iSCSI, FCoE
                4   None      None, NIC, iSCSI, FCoE
                6   None      None, NIC
                8   None      None, NIC
               10   None      None, NIC
               12   None      None, NIC
               14   None      None, NIC
  N/A (p1)      1   None      None, NIC, NIC+RoCE

Note: the port number for re-enabled ports will be available after system reboot.
```

5.16.2 GetLinkConfig

This command displays the active and configured link configuration IDs and lists the available link configurations (descriptions).

NOTE The link reconfiguration feature is available on only some OCe14000-series adapters.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetLinkConfig <MAC|WWPN>
```

Parameters

WWPN	The WWPN of any FCoE function on the adapter.
MAC	The MAC address of any NIC+RoCE or iSCSI function on the adapter.

Example

```
OCM: brcmhbacmd GetLinkConfig 00-90-fa-41-28-f0

Active Link ID:     19
Configured Link ID: 19

Available Link Configurations

ID Description
-- --------------------------------
19   1x40Gbps QSFP+ OCe14401-UX CNA
20   4x10Gbps SFP+ OCe14401-UX CNA
```

5.16.3 ListProfiles

This command shows a list of the available profile IDs and their descriptions for an adapter, including the currently active and configured profile IDs. This information helps to determine the appropriate profile ID to specify to the `ChangePersonality` and `SetAdapterPortConfig` commands required for RoCE and NIC-only configurations. The specific profiles listed may vary according to the adapter model.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
ListProfiles <MAC|WWPN>
```

Parameters

MAC	MAC address of any NIC, NIC+RoCE or iSCSI function on the adapter.
WWPN	WWPN of any FCoE function on the adapter.

Example

NOTE Different adapter models will have a different set of profiles.

```
>brcmhbacmd ListProfiles 00-00-c9-12-34-56
Active Profile ID:      23
Configured Profile ID:  23

Available Profiles
ID     Description
--  ---------------------------------

16    NIC
17    ISCSI initiator, no DIF
19    FCOE initiator + Target, with DIF
20    RoCE-1
21    RoCE-2
23    FCOE initiator + Target, no DIF
36    Mixed Storage, 25% ISCSI initiator only, 75% FCOE, no DIF
37    Mixed Storage, 50% ISCSI initiator only, 50% FCOE, no DIF
38    Mixed Storage, 75% ISCSI initiator only, 25% FCOE, no DIF
```

5.16.4 SetAdapterPortConfig (for OCe14000-Series Adapters)

This command can do the following:

- Configure the function protocol assignments for all ports on an OCe14000-series adapter.
- Restore the adapter to its factory default profile and settings.
- Enable or disable SR-IOV on the entire adapter.
- Remove a port when all functions are set to None, even if multichannel is enabled.

NOTE This command fails under the following conditions:

- Three storage protocols on a port are specified.
- Two of the same storage protocols on a port are specified.

- The number of ports and function protocol assignments specified in this command depend on the available ports and the multichannel and NPar states. A reboot is required to activate the new port configuration.

To support up to 16 functions on an adapter, ARI must be available on the system and the following conditions must be met:

- The system hardware, such as the motherboard and BIOS, must support ARI.
- ARI must be enabled in the system BIOS.
- The operating system must support ARI, such as Windows Server 2012 and later.
- The management tools you use must support ARI, such as OneCommand CNA Manager.

If these conditions are not met, although you may configure all 16 functions, only eight functions will be present and discovered by the OneCommand CNA Manager application after a reboot.

NOTE SR-IOV is not supported with RoCE or with multichannel configurations. For a Linux or VMware operating system, SR-IOV must be enabled on the system BIOS when NParEP is used. See the documentation that accompanied your Dell server for more information.

You cannot specify `Flex` for the *`mctype`* parameter of the `SetAdapterPortConfig` command. The only *`mctype`* parameter that can be specified on HP UMC-capable adapters is UMC, even though `Flex` is displayed in the output for the `GetAdapterPortConfig` and the `CMGetParams` commands as the multichannel type for HP adapters currently running in `Flex` mode.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
SetAdapterPortConfig <MAC|WWPN> defaults
```

or

```
SetAdapterPortConfig <MAC | WWPN> p0=fcnCfg [p1=fcnCfg] [p2=fcnCfg] [p3=fcnCfg]
[pid=ProfileID] [mctype=Type] [NParEpMode=State][sriov=SriovState]
```

Parameters

MAC	MAC address of any iSCSI, NIC or NIC+RoCE function on the adapter.
WWPN	WWPN of any FCoE function on the adapter.
defaults	Sets the adapter port configuration back to its factory default configuration. No additional parameters should be specified with the default parameter. For Dell adapters, this command resets virtual addresses to their factory defaults. A system reboot is required to activate the new settings.
p0=fcnCfg	First port's function
p1=fcnCfg	Second port's function (required for 2- and 4-port adapters)
p2=fcnCfg	Third port's function (required for 4-port adapters)
p3=fcnCfg	Fourth port's function (required for 4-port adapters)
fcnCfg	:: f0[,f1[,...fM]] where: — f0 – First function protocol on port (must be NIC or NIC+RoCE). — f1 – Second function protocol on port (optional). — f*m* – *M*th function protocol on port (optional). The number of functions per port (the value of *M*) is specific to the model and multichannel type.

pid=ProfileID — This parameter is required in the following instances:

- When specifying a NIC-only personality
- When specifying a NIC+RoCE personality
- When more than one NIC profile is available

Otherwise, this parameter must be omitted.

mctype=Type — Optional multichannel type specification. This parameter is required to enable multichannel. Specify None to disable multichannel. If you do not specify an mctype, no change to the configured multichannel type is made. Multichannel types are listed as follows:

Restrictions	Options: Type =
None	None
None	UMC (this option is not available for Lenovo System X or Dell)
Lenovo System X	vNIC1
Lenovo System X	SIMode
Lenovo System X	UFP
Dell	NPar

NParEpMode=State — NParEP Mode state (Dell only). This parameter can only be specified when mctype=NPar. See Section 5.16.6.3, NParEP Mode for more information.

0 = Disable, 1 = Enable sriov = SriovState (0 = disabled, 1 = enabled)

SriovState — SR-IOV state (0 = disabled, 1 = enabled)

Considerations

- The protocols specified by the f0–fm parameters are the names displayed by the `ShowProperties` command.
- If you are using `SIMODE` or `UFP` with OCe-14000-series adapters, the maximum number of channels is eight per port for 2-port adapters and four channels per port for 4-port adapters. If you are using `vNIC1`, the maximum number of channels is four channels per port for 2-port and 4-port adapters.
- The following rules apply when using this command with non-NPar multichannel enabled:
  - Do not add spaces before or after the equal sign (=) and commas (,) in the port specification. For example,

    ```
    p0 = nic, iscsi
    ```

    would cause an error.
  - The only choice for f0 is NIC or NIC+RoCE. Choices for f1 through f*m* are `None`, `iSCSI`, `FCoE` or `NIC`. `NIC` for f1 through f*m* can be specified only for multichannel configurations. Except for `NPar`, `iSCSI` and `FCoE` are limited to f1 and f2.
  - NIC+RoCE can be specified only as the f0 parameter with no other f parameters specified, and it cannot be specified with multichannel configurations. If NIC+RoCE is specified for any of the port functions, all other port function specifications must be NIC or NIC+RoCE. For example, a single NIC or NIC+RoCE function per port is all that is allowed. Also, NIC+RoCE can only be specified when *`mctype`*`=none`.
  - The number of functions (f) specified depends on the adapter model and configuration type; for example, non-multichannel, multichannel, RoCE, and NPar. If an incorrect number of functions is specified, an error is generated.
  - The *`pid`*`=ProfileID` option is required for NIC+RoCE configurations. If it is specified and the profile does not work with the specified port functions, an error is generated. When using a NIC-only profile, if more than one NIC profile exists, you must specify it. To determine how many NIC-only profiles are available, use the `ListProfiles` command. See Section 5.16.3, ListProfiles for more information.
  - f0 must always be specified. However, if the rest of the functions on a port can be set to None, they do not have to be specified. For example, p0=NIC, iSCSI p1=NIC, None could be specified as `p0=NIC, iSCSI p1=NIC`.
  - Storage protocols can be configured on f1 or f2 or both.

- If *mctype* is specified, it must be an available multichannel type, or it must be set to None to disable multichannel. If it is not specified, the currently configured multichannel type (including None) remains unchanged.
- If the functions being set to None are the last functions on a port, they do not need to be specified. For example, p0=NIC, iSCSI, None can be shortened to p0=nic,iscsi.
- NParEP mode can be specified only if the *mctype* parameter is set to NPar or if the currently configured multichannel type is NPar.
- NParEP support is available only on Dell 13G or newer systems.
- When disabling NPar (*mctype*=none), NParEP mode is automatically disabled.

Examples

The following table provides the locations of application examples.

Example Types	Section
One storage protocol plus SR-IOV	5.16.4.1
Two storage protocols – 2 ports	5.16.4.2
RoCE	5.16.4.3
Enabling NPar with SR-IOV	5.16.4.4
Removing a port	5.16.4.6

5.16.4.1 One Storage Protocol Plus SR-IOV

```
>brcmhbacmd SetAdapterPortConfig 00-00-c9-12-34-56 p0=nic,iscsi p1=nic,fcoe
p2=nic,fcoe p3=nic sriov=1

>brcmhbacmd GetAdapterPortConfig 00-00-c9-12-34-56
Number of Ports              : 4
Maximum FCoE Functions       : 4
Maximum iSCSI Functions      : 4
Mixed Mode I/O Allowed       : Yes
Concurrent Mode I/O Allowed  : No
ROCE Functions Allowed       : Yes
RebootRequired               : Yes

Active Configuration
 Profile ID  : 37
 MCType      : None
 SR-IOV State: Disabled
 Port Assignments:
   Port 0: NIC,iSCSI
   Port 1: NIC,FCoE
   Port 2: NIC,iSCSI
   Port 3: NIC,FCOE

Reboot Configuration
  Profile ID : 37
  MCType     : None SR-IOV State: Enabled
  Port Assignments:
    Port 0 (P0): NIC,iSCSI
    Port 1 (P1): NIC,FCoE
    Port 2 (P2): NIC,FCoE
    Port 3 (P3): NIC,None
```

```
Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig
command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).
```

5.16.4.2 Two Storage Protocols – 2 Ports

```
>brcmhbacmd SetAdapterPortConfig 00-00-c9-12-34-56 p0=nic,iscsi,fcoe p1=nic,fcoe

>brcmhbacmd GetAdapterPortConfig 00-00-c9-12-34-56
Number of Ports             : 2
Maximum FCoE Functions      : 4
Maximum iSCSI Functions     : 4
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : Yes
ROCE Functions Allowed      : Yes
RebootRequired              : Yes

Active Configuration
  Profile ID  : 33
  MCType      : None
  SR-IOV State: Disabled
  Port Assignments:
    Port 1: NIC,iSCSI,FCoE
    Port 1: NIC,iSCSI,FCoE

Reboot Configuration
  Profile ID  : 33
  MCType      : None
  SR-IOV State: Disabled
  Port Assignments:
    Port 1 (p0): NIC,iSCSI,FCoE
    Port 2 (p1): NIC,FCoE,None

Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig
command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).
```

5.16.4.3 RoCE

```
>brcmhbacmd SetAdapterPortConfig 00-00-c9-12-34-56 p0=nic+roce p1=nic p2=nic
p3=nic+roce pid=21
>brcmhbacmd GetAdapterPortConfig 00-00-c9-12-34-56
Number of Ports             : 4
Maximum FCoE Functions      : 4
Maximum iSCSI Functions     : 4
Mixed Mode I/O Allowed      : Yes
Concurrent I/O Mode Allowed : Yes
ROCE Functions Allowed      : Yes
RebootRequired              : Yes
Active Configuration
  Profile ID  : 37
  MCType      : None
  SR-IOV State: Disabled
  Port Assignments:
    Port 0: NIC,iSCSI,FCoE
    Port 1: NIC,FCoE,None
    Port 2: NIC,iSCSI,None
```

```
    Port 3: NIC,None,None

Reboot Configuration
  Profile ID  : 21
  MCType      : None
  SR-IOV State: Disabled
  Port Assignments:
    Port 0 (p0): NIC+RoCE,None,None
    Port 1 (p1): NIC,None,None
    Port 2 (p2): NIC,None,None
    Port 3 (p3): NIC+RoCE,None,None

Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig
command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).
```

5.16.4.4 Enabling NPar with SR-IOV

```
>brcmhbacmd setadapterportconfig 00-90-fa-26-14-2a p0=nic,iscsi,fcoe,nic p1=nic,
iscsi,none,nic p2=nic,iscsi,nic,fcoe p3=nic,nic,iscsi,none mctype=npar
nparepmode=1 sriov=1

Successfully changed port configuration. A reboot is required to complete changes.
```

NOTE You can use the `GetAdapterPortConfig` command to display the SR-IOV state.

5.16.4.5 Multichannel Configurations[

For additional information about universal multichannel, refer to the *Emulex Universal Multichannel Reference Manual*.

ARI must be available to support sixteen functions on an adapter. OCe14000-series adapters automatically support ARI. However, the system's motherboard must support ARI, ARI 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 detected by the OneCommand CNA Manager application after a reboot.

5.16.4.5.1 Enabling and Disabling Multichannel

The *`mctype`* parameter in the `SetAdapterPortConfig` command enables or disables multichannel on an adapter. This parameter is required to enable multichannel.

Disable multichannel on an adapter by setting *`mctype`* to `None`. When multichannel is enabled, the additional function protocols must be specified for the extra channels that are not available when multichannel is disabled.

NOTE The `CMMode` command that enables or disables multichannel on OCe11100-series adapters cannot be used on OCe14000-series adapters. If it is used, an error message is generated. Use the `SetAdapterPortConfig` command instead.

Examples

Enable Multichannel for UMC:

```
>brcmhbacmd SetAdapterPortConfig 00-00-c9-12-34-56 p0=nic,iscsi,nic,nic
p1=nic,fcoe,nic,nic p2=nic,fcoe,nic,nic p3=nic,iscsi,nic,nic mctype=umc
```

Disable Multichannel:

```
>brcmhbacmd SetAdapterPortConfig 00-00-c9-12-34-56 p0=nic,iscsi p1=nic,fcoe
p2=nic,fcoe p3=nic,iscsi mctype=none
```

5.16.4.5.2 Showing Multichannel Configuration

The `CMGetParams` command shows the current multichannel configuration for an adapter's physical port. See Section 5.5.1, CMGetParams for more information.

5.16.4.5.3 Setting Channel Protocols using SetAdapterPortConfig

Use the `SetAdapterPortConfig` command to configure the protocols running on the channels or functions.

The following rules apply when using `SetAdapterPortConfig` with non-NPar multichannel enabled:

- The first function must always be NIC; other functions can be None.
- If NIC+RoCE is specified on any port, all other ports must be either NIC or NIC+RoCE.
- Storage functions cannot be specified on adapters running NIC+RoCE.
- RoCE cannot be specified for any multichannel configurations.
- This command fails under the following conditions:
  - Three storage protocols are specified on a port.
  - The f2 and f3 storage protocols are the same.
- The optional *`mctype`* parameter can specify the multichannel type, or it can be set to `None` to disable multichannel. If the *`mctype`* parameter is not specified, the currently configured multichannel type is not changed.

NOTE You must reboot your system after running this command to run the new protocols assigned to the channels.

Examples

The following table provides the locations of application examples.

Example Types	Section
Enable UMC on 4-port, 4 channels per port	5.16.4.5.3.1
Enable UMC on 2-port, 8 channels per port	5.16.4.5.3.2

5.16.4.5.3.1 Enable UMC on 4-Port, 4 Channels per Port

```
>brcmhbacmd SetAdapterPortConfig 00-00-c9-12-34-56 p0=nic,iscsi,nic,nic
p1=nic,fcoe,nic,nic p2=nic,fcoe,nic,nic p3=nic,iscsi,nic,nic mctype=umc

Successfully changed port configuration. A reboot is required to complete changes.

>brcmhbacmd GetAdapterPortConfig 00-00-c9-12-34-56
Number of Ports             : 4
Maximum FCoE Functions      : 4
Maximum iSCSI Functions     : 4
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : No
ROCE Functions Allowed      : Yes
RebootRequired              : Yes

Active Configuration
  Profile ID  : 37
  MC Type     : UMC
  SR-IOV State: Disabled
  Port Assignments:
    Port 0: NIC,NIC,NIC,NIC
```

```
    Port 1: NIC,NIC,NIC,NIC
    Port 2: NIC,NIC,NIC,NIC
    Port 3: NIC,NIC,NIC,NIC

Reboot Configuration
  Profile ID  : 37
  MC Type     : UMC
  SR-IOV State: Disabled
  Port Assignments:
    Port 0 (p0): NIC,iSCSI,NIC,NIC
    Port 1 (p1): NIC,FCoE,NIC,NIC
    Port 2 (p2): NIC,FCoE,NIC,NIC
    Port 3 (p3): NIC,iSCSI,NIC,NIC


Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig
command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).
```

5.16.4.5.3.2 Enable UMC on 2-Port, 8 Channels per Port

```
>brcmhbacmd SetAdapterPortConfig 00-00-c9-12-34-56
p0=nic,iscsi,fcoe,nic,nic,nic,nic,nic p

1= nic,fcoe,iscsi,nic,nic,nic,nic,nic mctype=UMC
Successfully changed port configuration. A reboot is required to complete changes.

>brcmhbacmd getadapterportconfig 00-00-c9-12-34-56
Number of Ports             : 4
Maximum FCoE Functions      : 4
Maximum iSCSI Functions     : 4
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : Yes
ROCE Functions Allowed      : Yes
RebootRequired              : Yes

Active Configuration
  Profile ID  : 37
  MC Type     : None
  SR-IOV State: Disabled
  Port Assignments:
    Port 0: NIC,iSCSI,FCoE
    Port 1: NIC,FCoE,iSCSI

Reboot Configuration
  Profile ID  : 37
  MC Type     : UMC
  SR-IOV State: Disabled
  Port Assignments:
    Port 0 (p0): NIC,iSCSI,FCoE,NIC,NIC,NIC,NIC,NIC
    Port 1 (p1): NIC,FCoE,iSCSI,NIC,NIC,NIC,NIC,NIC

Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig
command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).
```

5.16.4.5.4 Setting Multichannel Bandwidths

See Section 5.5.3, CMSetBW for more information.

5.16.4.6 Removing a Port

NOTE Removed ports are not discovered by the `ListHba` command after a system reboot.

Channel management operations cannot be performed on a newly disabled port.

The following conditions generate error or warning messages (see Table 9 for a list of messages):

- Adding a port while enabling multichannel mode
- Attempting to remove a port on an unsupported adapter
- Setting the first function from the first port to none
- Setting the first function from a port other the first port to none

To remove a port on supported adapters, type the following command:

```
brcmhbacmd setAdapterPortConfig 00-90-fa-41-22-f0 p0=NIC p1=none pid=16
```

The following output is displayed:

```
Adapter configuration successfully set to new configuration.
System reboot required to activate it.
```

If you set the first function of a port other than the first port to `None`, you must either also set the rest of the functions from that port to `None`, or do not specify any other function (which will cause the remaining functions to be set to `None`). If you attempt to set any other function from this port to a protocol other than `None`, the attempt fails and the following error message is generated:

```
>brcmhbacmd setadapterportconfig 00-90-FA-41-22-F0 p0=nic,fcoe p1=none,fcoe

ERROR: If f0 is set to None, all functions on the port must be set to NONE.
ERROR: <4>: Invalid Argument
```

5.16.4.7 Adding a Port While Enabling Multichannel Mode

If you add a port back while also enabling multichannel mode, a warning message notifies you that multichannel properties can be set on this port only after a system reboot has been performed.

For example, if you type the following command:

```
brcmhbacmd setadapterportconfig 00-90-FA-41-22-F0 p0=nic,fcoe,nic,nic
p1=nic,fcoe,nic,nic mctype=umc
```

The following messages are displayed:

```
Adapter configuration successfully set to new configuration.
System reboot required to activate it.

The multichannel type has been changed. Be sure to use the additional
multichannel commands to set the other multichannel properties in order to
complete the configuration.

Multichannel properties can only be set on newly added ports after system
reboot.
```

5.16.4.8 Adding a Port While Operating in Multichannel Mode

If you add a port back while the adapter is currently running in `Multichannel` mode, a warning message notifies you that multichannel properties can be set on this port only after a system reboot has been performed.

For example, if you type the following command:

```
brcmhbacmd setadapterportconfig 00-90-FA-41-22-F0 p0=nic,fcoe,nic,nic
p1=nic,fcoe,nic,nic
```

The following messages are displayed:

```
Adapter configuration successfully set to new configuration.
System reboot required to activate it.

Multichannel properties can only be set on newly added ports after system
reboot.
```

5.16.5 Lenovo System X Multichannel for OCe14000-Series Adapters

For Lenovo System X multichannel configuration, three multichannel types are available: vNIC, SIMode, and UFP.

For vNIC, a maximum of four channels per port can be configured.

For SIMode and UFP, the maximum number of channels per port is shown in Table 8.

Table 8 SIMode and UFP Maximum number of functions per port

Number of Ports	Port Speed	Default Functions per Port	Maximum Functions per Port
1	10Gb/s	4	8
2	10Gb/s	4	8
4	10Gb/s	4	4
1	40Gb/s	4	16

5.16.6 Dell NPar Configurations

NOTE NParEP must be available to support 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. If these conditions are not met, although you are required to configure all sixteen functions, only eight functions will be present and detected by the OneCommand CNA Manager application after a reboot.

5.16.6.1 Enabling and Disabling NPar

Enable NPar by setting the *mctype* parameter in the `SetAdapterPortConfig` command to `NPar`. When you enable NPar, you must also specify the additional function protocols for the extra channels, or they will be set to `None`, meaning that there is no protocol running on the function.

Disable NPar by setting *mctype* to `None`.

NOTE The `CMMode` command that enables or disables multichannel on OCe11100-series adapters cannot be used to enable or disable NPar. Using the `CMMode` command to enable or disable NPar displays an

error message indicating that the command is not supported by the firmware or hardware.

When NPar is disabled, NParEP mode is automatically disabled as well. It is not necessary to specify the *nparepmode* parameter on the command line when disabling NPar (*mctype*=none), and an error is generated if this is attempted.

After enabling NPar with the SetAdapterPortConfig command, default bandwidths are automatically assigned to each partition with an assigned protocol. See Section 5.16.6.7, Default NPar Bandwidths for more information.

5.16.6.2 Showing NPar Configuration

The GetAdatperPortConfig command and the CMGetParams command show different aspects of the NPar configuration.

5.16.6.2.1 Showing NPar Configuration Using GetAdapterPortConfig

The GetAdatperPortConfig command displays all of the adapter's ports and the functions running on each port. It also shows the state of NParEP mode. See Section 5.16.1, GetAdapterPortConfig for more information.

Examples Using GetAdatperPortConfig

The following table provides the locations of application examples.

Example Types	Section
NPar disabled – 2 ports, NParEP mode disabled	5.16.6.2.1.1
NPar enabled – 2 ports, NParEP mode enabled on reboot	5.16.6.2.1.2

5.16.6.2.1.1 NPar Disabled – 2 Ports, NParEP Mode Disabled

```
>brcmhbacmd getadapterportconfig 00-00-c9-12-34-56
Number of Ports             : 4
Maximum FCoE Functions      : 4
Maximum iSCSI Functions     : 4
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : Yes
ROCE Functions Allowed      : Yes
RebootRequired              : No

Active Configuration
  Profile ID  : 16
  MC Type     : None
  SR-IOV State: Enabled
  Port Assignments:
    Port 1: NIC
    Port 2: NIC

Reboot Configuration
  Profile ID  : 16
  MC Type     : None
  SR-IOV State: Enabled
  Port Assignments:
    Port 1 (p0): NIC
    Port 2 (p1): NIC
```

```
Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig
command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).
```

5.16.6.2.1.2 NPar Enabled – 2 Ports, NParEP Mode Enabled on Reboot

```
>brcmhbacmd getadapterportconfig 00-00-c9-12-34-56
Number of Ports             : 4
Maximum FCoE Functions      : 4
Maximum iSCSI Functions     : 4
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : Yes
ROCE Functions Allowed      : Yes
RebootRequired              : Yes

Active Configuration
  Profile ID  : 37
  MC Type     : NPAR
  SR-IOV State: Disabled
  Port Assignments:
    Port 1: NIC,iSCSI,FCoE,NIC
    Port 2: NIC,iSCSI,FCoE,NIC

Reboot Configuration
  Profile ID  : 37
  MC Type     : NPAR
  SR-IOV State: Disabled
  NParEP Mode : Enabled
  Port Assignments:
    Port 1 (p0): NIC,iSCSI,NIC,NIC,NIC,NIC,NIC,NIC
    Port 2 (p1): NIC,iSCSI,FCoE,NIC,NIC,NONE,NONE,NONE


Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig
command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).
```

5.16.6.2.2 Showing NPar Configuration Using CMGetParams

The `CMGetParams` command shows the current NPar configuration for an adapter's physical port. When NPar is enabled, `CMGetParams` shows the function information such as the protocol and bandwidths. See Section 5.5.1, CMGetParams for more information.

When you are using NPar, this command displays the adapter's active (booted) mode as `NPar` or `None`, the configured mode as `NPar` or `None`, and the available management mode, which can only be `NPar`. This is followed by a table showing the specified port's functions and function properties. When `NPar` is active, the `Type` column shows the protocol that is running on each function.

Examples Using CMGetParams

The following table provides the locations of application examples.

Example Types	Section
NPar disabled	5.16.6.2.2.1
NPar enabled – 4 partitions per port	5.16.6.2.2.2
NPar enabled – 8 partitions per port	5.16.6.2.2.3

5.16.6.2.2.1 NPar Disabled

```
>brcmhbacmd cmgetparams 00-00-c9-12-34-56
Active Mode:     None
Configured mode: None
Available modes: NPAR
```

5.16.6.2.2.2 NPar Enabled – 4 Partitions per Port

```
>brcmhbacmd cmgetparams 00-00-c9-12-34-56
Active Mode:     NPAR
Configured mode: NPAR
Available modes: NPAR

Func# Type  MAC Address       Min BW Max BW
----- ----- ----------------- ------ ------
0     NIC   00-00-c9-12-34-56   20     75
4     iSCSI 00-00-c9-12-34-57   60     100
8     FCoE  00-00-c9-12-34-58   40     100
12    NONE  00-00-c9-12-34-59   0      0
```

5.16.6.2.2.3 NPar Enabled – 8 Partitions per Port

```
>brcmhbacmd cmgetparams 00-00-c9-12-34-56
Active Mode:     NPAR
Configured mode: NPAR
Available modes: NPAR

Func# Type  MAC Address       Min BW Max BW
----- ----- ----------------- ------ ------
0     NIC   00-00-c9-12-34-56   10     50
2     iSCSI 00-00-c9-12-34-57   20     100
4     NIC   00-00-c9-12-34-58   10     50
6     FCoE  00-00-c9-12-34-59   30     100
8     NIC   00-00-c9-12-34-5A   10     50
10    NIC   00-00-c9-12-34-5B   10     50
12    NIC   00-00-c9-12-34-5C   10     50
14    NIC   00-00-c9-12-34-5C   0      50
```

5.16.6.3 NParEP Mode

For NPar adapters, you can change the total number of partitions displayed and configured on the adapter to 8 if NParEP mode is disabled, or to as many as16 if NParEP mode is enabled. Use the *nparepmode* parameter in the SetAdapterPortConfig command to enable or disable this property.

NOTE If NPar is disabled, NParEP mode is automatically disabled as well. It is not necessary to specify the nparepmode parameter on the command line when disabling NPar by setting the *mctype* parameter to None.

After enabling or disabling NParEP mode with the SetAdapterPortConfig command, default bandwidths are automatically assigned to each partition with an assigned protocol. See Section 5.16.6.7, Default NPar Bandwidths for more information.

Syntax

```
SetAdapterPortConfig <MAC | WWPN> p0=fcnCfg [p1=fcnCfg] [p2=fcnCfg] [p3=fcnCfg]
[pid=ProfileID] [mctype=Type] [NParEpMode=State][sriov=SriovState]
```

5.16.6.4 Parameters

MAC	MAC address of any iSCSI, NIC or NIC+RoCE function on the adapter.
WWPN	WWPN of any FCoE function on the adapter
defaults	Sets the adapter port configuration back its factory default configuration. No additional parameters should be specified with the defaults parameter.
p0=fcnCfg	First port's function
p1=fcnCfg	Second port's function (required for 2- and 4-port adapters)
p2=fcnCfg	Third port's function (required for 4-port adapters)
p3=fcnCfg	Fourth port's function (required for 4-port adapters)
fcnCfg	:: f0[,f1[,...f*M*]] where: — f0 – First function protocol on port (must be NIC or NIC+RoCE). — f1 – Second function protocol on port (optional). — f*m* – *M*th function protocol on port (optional). The number of functions per port (the value of *M*) are model and multichannel-type specific.
pid=ProfileID	This parameter is required when specifying a NIC-only or NIC+RoCE personality; otherwise, this parameter must be omitted.
mctype=Type	Optional multichannel type specification. This parameter is required to enable multichannel. Specify None to disable multichannel or NPar if you are using the NParEP mode. If you do not specify an mctype, no change to the configured multichannel type is made.
NParEpMode=State	NParEP Mode state (Dell only). This parameter can only be specified when mctype=NPar. See Section 5.16.6.3, NParEP Mode for more information. 1 = enables the NParEP Mode to display and configure 16 functions. 0 = disables the NParEP Mode to display and configure only 8 functions. sriov = SriovState (0 = disabled, 1= enabled

5.16.6.5 Setting NPar Function Protocols

When NPar is enabled, the protocols for additional functions on each port can be specified depending on the board model, the number of ports available on the adapter, and the NParEPMode setting. If protocols for additional functions are not specified, they are set to `None`. In other words, no protocol is running on the function. The `SetAdapterPortConfig` command configures the protocols running on each function.

After NPar mode is enabled with the `SetAdapterPortConfig` command, default bandwidths are automatically assigned to each partition with an assigned protocol. See Section 5.16.6.7, Default NPar Bandwidths for more information.

The following rules apply when using the `SetAdapterPortConfig` command with multichannel enabled to configure NPar:

- The first function on a port is always NIC. NIC+RoCE is not available when NPar is enabled.
- Storage can be configured only on the second, third, or fourth partitions. A total of two different storage protocols can be configured on a port.
- Any partition other than the first partition can be configured as `None`; that is, for no protocol.
- A reboot is required after executing this command to run the new protocols assigned to the functions.

Examples

The following table provides the locations of application examples.

Example Types	Section
4 ports – NParEP mode disabled, 2 functions per port	5.16.6.5.1
2 ports – NParEP mode disabled, 4 functions per port	5.16.6.5.2
4 ports – NParEPMode enabled, 4 functions per port	5.16.6.5.3

5.16.6.5.1 4 Ports – NParEP Mode Disabled, 2 Functions per Port

```
>brcmhbacmd setadapterportconfig 00-00-c9-12-34-56 p0=nic,iscsi p1=nic,fcoe
p2=nic,fcoe p3=nic,iscsi mctype=npar nparepmode=0
Successfully changed port configuration. A reboot is required to complete changes.

>brcmhbacmd getadapterportconfig 00-00-c9-12-34-56
Number of Ports             : 4
Maximum FCoE Functions      : 4
Maximum iSCSI Functions     : 4
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : Yes
ROCE Functions Allowed      : Yes
RebootRequired              : Yes

Active Configuration
  Profile ID  : 16
  MC Type     : NPAR
  Port Assignments:
    Port 1: NIC,NIC
    Port 2: NIC,NIC
    Port 3: NIC,NIC
    Port 4: NIC,NIC

Reboot Configuration
  Profile ID  : 37
  MC Type     : NPAR
  NParEP Mode : Disabled
  Port Assignments:
    Port 1 (p0): NIC,iSCSI
    Port 2 (p1): NIC,FCoE
    Port 3 (p2): NIC,FCoE
    Port 4 (p3): NIC,iSCSI


Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig
command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).
```

5.16.6.5.2 2 Ports – NParEP Mode-Disabled, 4 Functions per Port

```
>brcmhbacmd setadapterportconfig 00-00-c9-12-34-56 p0=nic,iscsi,nic,fcoe
p1=nic,fcoe,iscsi,nic mctype=npar nparepmode=0

Successfully changed port configuration. A reboot is required to complete changes.
```

```
>brcmhbacmd getadapterportconfig 00-00-c9-12-34-56
Number of Ports             : 2
Maximum FCoE Functions      : 2
Maximum iSCSI Functions     : 2
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : No
ROCE Functions Allowed      : Yes
RebootRequired              : Yes

Active Configuration
  Profile ID  : 16
  MC Type     : None
  Port Assignments:
    Port 1: NIC
    Port 2: NIC

Reboot Configuration
  Profile ID  : 33
  MC Type     : NPAR
  NParEP Mode : Disabled
  Port Assignments:
    Port 1 (p0): NIC,iSCSI,NIC,FCoE
    Port 2 (p1): NIC,FCoE,iSCSI,NIC
```

Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).

5.16.6.5.3 4 Ports – NParEPMode-Enabled, 4 Functions per Port

```
>brcmhbacmd setadapterportconfig 00-00-c9-12-34-56 p0=nic,iscsi,fcoe,nic
p1=nic,iscsi,none,nic p2=nic,iscsi,nic,fcoe p3=nic,nic,iscsi,none mctype=npar
nparepmode=1

Successfully changed port configuration. A reboot is required to complete changes.
>brcmhbacmd getadapterportconfig 00-00-c9-12-34-56
Number of Ports             : 4
Maximum FCoE Functions      : 4
Maximum iSCSI Functions     : 4
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : Yes
ROCE Functions Allowed      : Yes
RebootRequired              : Yes

Active Configuration
  Profile ID  : 33
  MC Type     : NPAR
  SR-IOV State: Disabled
  Port Assignments:
    Port 1: NIC,NIC,NIC,NIC
    Port 2: NIC,NIC,NIC,NIC
    Port 3: NIC,iSCSI,NIC,FCoE
    Port 4: NIC,iSCSI,FCoE,None
```

```
Reboot Configuration
  Profile ID  : 33
  MC Type     : NPAR
  SR-IOV State: Disabled
  NParEP Mode : Enabled
  Port Assignments:
    Port 1 (p0): NIC,iSCSI,FCoE,NIC
    Port 2 (p1): NIC,iSCSI,None,NIC
    Port 3 (p2): NIC,iSCSI,NIC,FCoE
    Port 4 (p3): NIC,NIC,iSCSI,None


Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig
command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).
```

5.16.6.6 Setting NPar Bandwidths

After enabling NPar with the `SetAdapterPortConfig` command, default bandwidths are set for each partition that is enabled. Use the `CMSetBW` command to configure new bandwidths if the default bandwidths are not acceptable. See Section 5.5.3, CMSetBW for more information.

NOTE A partition is referred to as a channel in the `CMSetBW` help.

The number of bandwidth combinations that need to be specified depend on the adapter model. See Section 5.16.1, GetAdapterPortConfig to determine how many bandwidths need to be specified for a port by looking at the number of functions indicated for each port number under `Reboot Configuration`. For example, if `Reboot Configuration` shows four functions for each port number, four sets of minimum and maximum bandwidths must be specified. If too many or too few minimum and maximum bandwidth combinations are provided, an error is generated.

The total of the minimum bandwidths for the enabled functions, such as partitions assigned protocols, must add up to 100.Although you can set the minimum bandwidths on some enabled functions to 0, you cannot set them all to 0.

The maximum bandwidth of a partition must be greater than or equal to the minimum bandwidth for that partition, up to a maximum of 100. The minimum and maximum bandwidths on disabled functions, those configured as `None`, must be 0.

NOTE A reboot is not required for bandwidth changes to take effect when channel management is enabled or when NPar is currently running on the adapter.

Setting the minimum and maximum bandwidths to zero does not bring the logical link down on the NPar function, or prevent receiving or transmitting packets on that function, because a small amount of network traffic still gets through.

5.16.6.7 Default NPar Bandwidths

After NParEP or NPar are enabled or disabled, the default bandwidths that are automatically assigned to each partition with an assigned protocol have minimum and maximum bandwidths. The default minimum bandwidths are set by dividing 100 by the number of partitions with assigned protocols to run on them. This can cause an uneven distribution if 100 is divided evenly by the number of partitions, meaning that the minimum bandwidth of some partitions may be one greater than others. The maximum bandwidth defaults to 100 for all partitions with assigned protocols.

NOTE If default bandwidths are not required, use the `CMSetBW` command to configure the minimum and maximum bandwidths. See Section 5.5.3, CMSetBW for more information.

Examples

The following table provides the locations of application examples.

Example Types	Section
4 partitions per port	5.16.6.7.1
4 partitions per port – third & fourth partition protocols = None	5.16.6.7.1
8 partitions per port – 2-port adapter, NParEP mode = Enabled, last channel MinBandwith = 0	5.16.6.7.1

5.16.6.7.1 4 Partitions per Port

```
>brcmhbacmd CMSetBW 00-00-c9-12-34-56 15,50 30,100 40,100 15,50
>brcmhbacmd CMGetParams 00-00-c9-12-34-56
Active mode:     NPAR
Configured mode: NPAR
Available modes: NPAR

Func# Type  MAC Address       Min BW Max BW
----- ----- ----------------- ------ ------
0     NIC   00-00-c9-12-34-56   15     50
1     iSCSI 00-00-c9-12-34-57   30     100
2     FCoE  00-00-c9-12-34-58   40     100
3     NIC   00-00-c9-12-34-59   15     50
```

5.16.6.7.2 4 Partitions per Port – Third and Fourth Partition Protocol = None

```
>brcmhbacmd CMSetBW 00-00-c9-12-34-56 50,100 60,100 0,0 0,0

>brcmhbacmd CMGetParams 00-00-c9-12-34-56
Active mode:     NPAR
Configured mode: NPAR
Available modes: NPAR

Func# Type  MAC Address       Min BW Max BW
----- ----- ----------------- ------ ------
0     NIC   00-00-c9-12-34-56   40     100
1     iSCSI 00-00-c9-12-34-57   60     100
2     None  00-00-c9-12-34-58   0      0
3     None  00-00-c9-12-34-59   0      0
```

5.16.6.7.3 8 Partitions per Port – 2 Port Adapter, NParEP Mode = Enabled, Last Channel MinBandwith = 0

```
>brcmhbacmd cmsetbw 00-00-c9-12-34-56 10,100, 30,100 10,100, 10,100, 10,100,
10,100, 10,100, 10,100

>brcmhbacmd cmgetparams 00-00-c9-12-34-56
Active mode:     NPAR
Configured mode: NPAR
Available modes: NPAR

Func# Type  MAC Address       Min BW Max BW
----- ----- ----------------- ------ ------
```

```
0     NIC   00-00-c9-12-34-56   10     100
1     iSCSI 00-00-c9-12-34-57   40     100
2     NIC   00-00-c9-12-34-58   10     100
3     NIC   00-00-c9-12-34-59   10     100
4     NIC   00-00-c9-12-34-5a   10     100
5     NIC   00-00-c9-12-34-5b   10     100
6     NIC   00-00-c9-12-34-5c   10     100
7     NIC   00-00-c9-12-34-5d   0      100
```

5.16.6.8 Configuring Ports When Disabling NPar

When NPar is being disabled, each of the ports can be configured to run with a single NIC or NIC+RoCE function. Use the `SetAdapterPortConfig` command to configure these ports. See Section 5.16.4, SetAdapterPortConfig (for OCe14000-Series Adapters) for more information.

Use the following syntax when disabling NPar:

```
SetAdapterPortConfig <MAC> <p0=NIC|NIC+RoCE> <p1=NIC|NIC+RoCE> [p2=NIC|NIC+RoCE]
[p3=NIC|NIC+RoCE] [pid=20|21]
```

where:

```
MAC               : MAC address of one of the NIC functions on the adapter
  p0=NIC|NIC+RoCE : Protocol to run on 1st port (NIC or NIC+RoCE)
  p1=NIC|NIC+RoCE : Protocol to run on 2nd port (NIC or NIC+RoCE)
  p2=NIC|NIC+RoCE : Protocol to run on 3rd port, if available (NIC or NIC+RoCE)
  p3=NIC|NIC+RoCE : Protocol to run on 2nd port, if available (NIC or NIC+RoCE)
  pid=20|21       : Required when setting any port to NIC+ROCE
                    (NIC+RoCE-1=20,NIC+RoCE-2=21)
```

NOTE When specifying any of the ports to run NIC+RoCE (for example, p1=RoCE), the RoCE Profile ID must be specified. It can be either 20 or 21 depending upon the preferred protocol. See Section 5.15.1, ChangePersonality for more information.

Examples

The following table provides the locations of application examples.

Example Types	Section
Disable NPar – 4 ports, all NIC	5.16.6.8.1
2 ports – switching from all NIC to NIC+RoCE on first port, RoCE profile ID 21: RoCE-2	5.16.6.8.2

5.16.6.8.1 Disable NPar – 4 Ports, All NIC

```
>brcmhbacmd setadapterportconfig 00-00-c9-12-34-56 p0=nic p1=nic p2=nic p3=nic
mctype=none sriov=0

Successfully changed port configuration. A reboot is required to complete changes.
>brcmhbacmd getadapterportconfig 00-00-c9-12-34-56
Number of Ports             : 4
Maximum FCoE Functions      : 4
Maximum iSCSI Functions     : 4
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : Yes
ROCE Functions Allowed      : Yes
```

```
RebootRequired              : Yes

Active Configuration
  Profile ID  : 33
  MC Type     : NPAR
  SR-IOV State: Disabled
  Port Assignments:
    Port 1: NIC,iSCSI,FCoE,NIC
    Port 2: NIC,iSCSI,None,NIC
    Port 3: NIC,iSCSI,NIC,FCoE
    Port 4: NIC,iSCSI,FCoE

Reboot Configuration
  Profile ID  : 16
  MC Type     : None
  SR-IOV State: Disabled
  Port Assignments:
    Port 1 (p0): NIC
    Port 2 (p1): NIC
    Port 3 (p2): NIC
    Port 4 (p3): NIC

Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig
command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).
```

5.16.6.8.2 2 Ports – Switching from All NIC to NIC+RoCE on First Port, RoCE Profile ID 21: RoCE-2

```
>brcmhbacmd setadapterportconfig 00-00-c9-12-34-56 p0=nic+roce p1=nic pid=21
mctype=none

Successfully changed port configuration. A reboot is required to complete changes.
>brcmhbacmd getadapterportconfig 00-00-c9-12-34-56
Number of Ports             : 2
Maximum FCoE Functions      : 2
Maximum iSCSI Functions     : 2
Mixed Mode I/O Allowed      : Yes
Concurrent Mode I/O Allowed : Yes
ROCE Functions Allowed      : Yes
RebootRequired              : Yes

Active Configuration
  Profile ID  : 16
  MC Type     : None
  SR-IOV State: Disabled
  Port Assignments:
    Port 1: NIC
    Port 2: NIC

Reboot Configuration
  Profile ID  : 21
  MC Type     : None
  SR-IOV State: Disabled
  Port Assignments:
```

```
    Port 1 (p0): NIC+RoCE
    Port 2 (p1): NIC

Note: (pN) labels are the port parameter names to use in the SetAdapterPortConfig
command (e.g. p0=NIC, FCoE p1=NIC, iSCSI).
```

5.16.7 SetLinkConfig

The `SetLinkConfig` command changes the configured link configuration ID. The command can change the link configuration ID to a specific ID or to the factory default ID. When the link configuration ID is successfully changed, a reboot is required to activate the new link configuration. If the `GetAdapterPortConfig` command is run after changing the link configuration ID, the `Reboot Configuration` displayed reflects the default port configuration for the newly configured link configuration ID. You can update this configuration before rebooting to avoid an additional reboot.

NOTE The link reconfiguration feature is available on only some OCe14000-series adapters.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
SetLinkConfig <MAC|WWPN> <LinkID|default>
```

Parameters

MAC	The MAC address of any NIC+RoCE or iSCSI function on the adapter.
WWPN	The WWPN of any FCoE function on the adapter.
LinkID	The ID of the link configuration to set (use `GetLinkConfig` for list of IDs).
default	Sets the link configuration ID of the adapter to the factory default ID.

NOTE Using the defaults option for the `SetAdapterPortConfig` command also changes the link configuration ID to the factory default ID.

Example

```
>brcmhbacmd setlinkconfig 00-00-c9-12-34-56 2
Link configuration successfully changed. A reboot is required to activate the new
link configuration. However, before rebooting you can run the
GetAdapterPortConfig and SetAdapterPortConfig commands to view or change the
default link configuration.
```

5.17 UMC Commands

The adapter's physical ports provide a converged conduit for network and storage traffic. Each channel has its own unique MAC address. Each channel provides traffic management capabilities, such as enabling and disabling, minium and maximum bandwidth, and VLAN ID (for untagged packets). For additional information on UMC, refer to the *Emulex Universal Multichannel Reference Manual*.

The CLI's UMC commands allow viewing of the UMC configuration, enabling and disabling of the UMC at the adapter level, and the modification of some of the channel properties. The UMC commands cannot be used to manage other channel management types; see Section 5.5, Channel Management Commands for more information.

Considerations

- SR-IOV is not supported with UMC.
- The UMC command `UmcSetChanLink` is no longer available. Its functionality can be performed by the `UmcSetBw` command by setting the min and max bandwidths to 0.
- For Lenovo System X adapters, UMC mode might be referred to as SIMode.
- The UMC commands are not available on OCe11101-EM/EX or OCe11102-EM/EX adapters. If you attempt to use UMC commands with these adapters, the following error is returned:

```
ERROR: <251>: Hardware or firmware does not support command.
```

5.17.1 UmcEnable

This command enables or disables UMC on OCe11100-series adapters at the adapter level. A system reboot is required to make the change take effect.

The `UmcEnable` command will enable only UMC channel management, but it will disable any other channel management type that is in effect when it is used.

The `CMMode` command can also be used to enable UMC or other channel management modes.

If you are using an OCe14000-series adapter, you must set the *`mctype`* parameter in the `SetAdapterPortConfig` command to enable UMC. See Section 5.16.4, SetAdapterPortConfig (for OCe14000-Series Adapters) for more information.

NOTE If you use the `UmcEnable` command on an OCe14000-series adapter, an error message indicates that the command is not supported by the firmware or hardware.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
UmcEnable <WWPN | MAC Address> <0 | 1>
```

Parameters

WWPN	WWPN of the FCoE function on the adapter.
MAC Address	MAC address of any NIC function on the adapter.
0	Disables UMC or SIMode
1	Enables UMC or SIMode

Example

```
>brcmhbacmd UmcEnable 00-00-c9-bb-cc-aa 1
```

5.17.2 UmcGetParams

NOTE This command has been replaced by the `CMGetParams` command, and it is provided for backward compatibility only for OCe11100-series adapters.

This command shows the current UMC configuration for an adapter's physical port. The command's output is the current UMC state for the adapter followed by a table showing the port's channels and their UMC properties. The `Type` property also shows the protocol that is running on the channel. This is controlled by the `ChangePersonality` and `SetAdapterPortConfig` commands, not by the UMC commands. See Section 5.15.1, ChangePersonality and Section 5.16.4, SetAdapterPortConfig (for OCe14000-Series Adapters) for more information.

Consideration

If you specify a MAC address that belongs to a newly disabled port, an error message indicates that the operation cannot be performed on this port.

For example, if you type the following command:

```
>brcmhbacmd umcgetparams 00-90-FA-41-22-F0
```

The following error message is displayed:

```
ERROR: This port is configured to be removed after reboot. Channel
Management properties are unavailable.
```

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
UmcGetParams <WWPN | MAC Address>
```

Parameters

WWPN	WWPN of an FCoE function on the port.
MAC Address	MAC address of any NIC function on the port.

Example

All NIC

```
>brcmhbacmd UmcGetParams 00-00-c9-bb-cc-aa

Active UMC State:       Enabled
Configured UMC State:   Enabled

Func# Type  MAC Address          LPVID  Min BW  Max BW
----- ----  -----------------    -----  ------  ------
0     NIC   00-00-c9-bb-cc-aa      2       25      50
1     NIC   00-00-c9-bb-cc-ab      3       0       0
2     NIC   00-00-c9-bb-cc-ac      4       25      50
3     NIC   00-00-c9-bb-cc-ad      5       50      75
```

NIC Plus Storage

```
>brcmhbacmd UmcGetParams 00-00-c9-bb-cc-aa

Active UMC State:       Enabled
Configured UMC State:   Enabled

Func# Type  MAC Address          LPVID  Min BW  Max BW
----- ----  -----------------    -----  ------  ------
0     NIC   00-00-c9-bb-cc-aa     1001    20      50
1     FCoE  00-00-c9-bb-cc-ab     n/a     50      100
```

```
2      NIC    00-00-c9-bb-cc-ac      1002     20       50
3      NIC    00-00-c9-bb-cc-ad      1003     10       75
```

5.17.3 UmcSetBW

NOTE This command has been replaced by the `CMSetBW` command, and it is provided for backward compatibility only for OCe11100-series adapters.

This command sets the minimum and maximum bandwidths for each channel (up to four) on the physical port. This command can also be used to disable the link on a channel by setting the minimum and maximum bandwidths of that channel to 0.

The total of the minimum bandwidths must add up to 100. An exception to this rule is for UMC/SIMODE configurations, in which both the minimum and maximum bandwidth for all channels are 0, effectively bringing the logical link down on all channels. The maximum bandwidth must have a value greater than or equal to the minimum bandwidth, and it must have a maximum value of 100.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
UmcSetBW <WWPN | MAC Address> <Min0,Max0> <Min1,Max1> <Min2,Max2> [Min3,Max3]
```

Parameters

MAC Address	MAC address of any NIC function on the port.
WWPN	WWPN of the FCoE function on the port.
Min0,Max0	Minimum and maximum bandwidths for channel 0.
Min1,Max1	Minimum and maximum bandwidths for channel 1.
Min2,Max2	Minimum and maximum bandwidths for channel 2.
Min3,Max3	Minimum and maximum bandwidths for channel 3.

Example

```
>brcmhbacmd UmcSetBW 00-00-c9-bb-cc-aa 25,50 0,50 50,75 25,100
```

Considerations

- This command is not supported on 1 Gb/s ports.
- If UMC is disabled when this command is executed, a warning message indicates that UMC is currently disabled and it must be enabled for these changes to take effect.
- If you specify a MAC address that belongs to a newly disabled port, an error message indicates that the operation cannot be performed on this port.

  For example, if you type the following command:

  ```
  >brcmhbacmd UmcSetBW  00-90-FA-41-22-F0 25,100 25,100 25,100 25,100
  ```

  The following message is displayed:

  ```
  ERROR: This port is configured to be removed after reboot. Bandwidth values
  can only be set on ports that will be available after reboot.
  ```

5.17.4 UmcSetLPVID

NOTE This command has been replaced by the `CMSetLPVID` command, and it is provided for backward compatibility only for OCe11100-series adapters.

This command sets the LPVID values for the UMC and SIMode NIC channels. Use the `UmcGetParams` command to determine the number of LPVIDs required. See Section 5.17.2, UmcGetParams for more information.

A reboot is not required for these changes take effect when UMC is enabled.

NOTE If the current multichannel mode is not UMC or SIMode, the `UmcSetLPVID` command fails.

If you specify a MAC address that belongs to a newly disabled port, an error message indicates that the operation cannot be performed on this port.

For example, if you type the following command:

```
>brcmhbacmd UmcSetLPVID 00-90-FA-41-22-F0 2 3 4 5 6 7 8 9
```

The following message is displayed:

```
ERROR: This port is configured to be removed after reboot. LPVID values can
only be set on ports that will be available after reboot.
```

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
UmcSetLPVID <WWPN|MAC> <LPVID0> <LPVID1> ... [LPVIDn]
```

Parameters

WWPN	WWPN of an FCoE function on the port.
MAC	MAC address of any NIC or iSCSI function on the port.
LPVID0	LPVID for channel 0.
LPVID1	LPVID for channel 1.
LPVIDn	LPVID for channel n.

Considerations for Using UmcSetLPVID

- LPVID values are in the range of 2 to 4094.
- Every NIC channel on a physical port must have a unique LPVID.
- For FCoE and iSCSI channels, 0 must be entered because LPVIDs can be specified only for NIC channels.
- LPVIDs specified for channels with protocols set to `None` are ignored.
- This command is not supported on 1 Gb/s ports.
- If channel management is disabled when this command is executed, an error message is displayed.

Examples

All NIC Channels

```
>brcmhbacmd UmcSetLPVID 00-00-c9-12-34-56 1001 1002 1003 1004
```

Storage on second Channel

```
>brcmhbacmd UmcSetLPVID 00-00-c9-12-34-56 1001 0 1002 1003
```

5.18 vPort Commands

The vPort Commands group manages virtual ports and functions only on FCoE adapters.

These commands are only supported for FCoE adapter functions.

NOTE In Linux, vPorts do not persist across system reboots. vPorts must be re-created after a system reboot.

5.18.1 CreateVPort

This command creates a virtual port with an automatically generated WWPN or a user-specified virtual WWPN on the specified physical port. If you specify `auto`, the virtual WWPN is generated automatically. Otherwise, you must specify the virtual WWPN for this parameter. If creation is successful, the WWPN is displayed as part of the output from the command. The *`vname`* optional parameter can be specified for the virtual port's name.

Supported By

Linux, Solaris, and Windows

Syntax

```
CreateVPort <physical WWPN> auto [vname]
```

-or-

```
CreateVPort <physical WWPN> <virtual WWPN> <virtual WWNN> [vname]
```

Parameters

physical WWPN	The WWPN of an FCoE function.
auto	The virtual WWPN is automatically generated for the virtual port.
vname	The virtual port's name (optional).
virtual WWPN	The virtual WWPN to create.
virtual WWNN	The virtual WWNN to create.

5.18.2 DeleteVPort

This command deletes the virtual port specified by a physical and virtual WWPN.

Supported By

Linux, Solaris, and Windows

Syntax

```
DeleteVPort <physical WWPN> <virtual WWPN>
```

Parameters

physical WWPN	The WWPN of an FCoE function.
virtual WWPN	The WWPN of the virtual port.

5.18.3 ListVPorts

This command lists virtual ports on the specified physical FCoE function. Leaving the physical WWPN parameter blank lists all virtual ports on all manageable hosts that support virtual ports.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
ListVPorts <physical WWPN>
```

Parameters

physical WWPN	The WWPN of an FCoE function.

5.18.4 VPortTargets

This command lists targets visible to the specified virtual port.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
VPortTargets <physical WWPN> <virtual WWPN>
```

Parameters

physical WWPN	The WWPN of an FCoE function.
virtual WWPN	The WWPN of the virtual port.

5.19 WWN Management Commands

NOTE These commands are supported only for FCoE functions.

The WWN Management Commands group validates WWNs carefully to avoid WWPN duplication, but WWNN duplication is acceptable. You might see error and warning messages if a name duplication is detected. Fulfill the activation requirement after each WWN is changed or restored. If pending changes exist, some diagnostic and maintenance features are not available.

5.19.1 ChangeWWN

This command allows you to change WWPNs and WWNNs, and it allows you to change the WWN to volatile or nonvolatile. If you attempt to select volatile on an adapter that does not support volatile WWNs, a `Not Supported` error is returned.

When a volatile change is supported, a reboot is required to activate the new value. Volatile names are active until system power-down or adapter power-cycle.

For VMware ESXi:

- After changing the WWN of a function, update your zoning settings before you reboot your ESXi server. If the zoning is not updated before your reboot, the subsequent boot could take a long time.
- After changing the WWN of a function, you must reboot the ESXi system before trying to access the adapter on that system. For information on rebooting the ESXi system, refer to the VMware documentation.
- If you are using the CIM interface to access functions, after changing the WWN of a function, you must restart the CIMOM (that is, SFCB) on the ESXi system before trying to access the function on that system. For information on restarting the CIMOM, refer to the VMware documentation.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
ChangeWWN <WWPN> <New WWPN> <New WWNN> <Type>
```

Parameters

WWPN	The WWPN of an FCoE function.
New WWPN	The WWPN of the FCoE function.
New WWNN	The WWNN of an FCoE function.
Type	0 = Volatile 1 = Nonvolatile

5.19.2 GetWWNCap

This command shows if volatile change is supported for the WWPN.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
GetWWNCap <WWPN>
```

Parameters

WWPN	The WWPN of an FCoE function.

5.19.3 ReadWWN

This command reads different types of WWNs.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
ReadWWN <WWPN> <Type>
```

Parameters

WWPN	The WWPN of an FCoE function.

Type	0 = Volatile 1 = Nonvolatile 2 = Factory default 3 = Current 4 = Configured

5.19.4 RestoreWWN

This command changes the WWNs to the factory default or non-volatile values. The change is non-volatile.

NOTE A reboot is required to activate the new value.

For VMware ESXi:

- After changing the WWN of an function, you must reboot the ESXi system before trying to access the adapter on that system. For information on rebooting the ESXi system, refer to the VMware documentation available from the VMware website.
- If you are using the CIM interface to access adapters, after changing the WWN of a function, you must restart the CIMOM (that is, SFCB) on the ESXi system before trying to access the function on that system. For information on restarting the CIMOM, refer to the VMware documentation available from the VMware website.

Supported By

Linux, Solaris, Windows, and Windows + CIM Provider on a VMware host

Syntax

```
RestoreWWN <WWPN> <Type>
```

Parameters

WWPN	The WWPN of an FCoE function.
Type	0 = Restore default WWNs 1 = Restore NVRAM WWNs

Appendix A: OneCommand CNA Manager Error Messages

Table 9 contains a list of some of the error messages that may be encountered during a OneCommand CNA Manager session.

Table 9 OneCommand CNA Manager Error and Warning Messages

Error Message	Command	Description
`Error: Read-only management mode is currently set on this host. The requested command is not permitted in this mode.`	Active management commands that change a property on an adapter or host.	This message is returned when some commands are attempted when the CLI is configured for read-only mode. See Section 4.1.1, CLI in Read-Only Mode.
`ERROR: <180>: Authentication: User unknown`	All	The specified user name is not valid or could not be authenticated by the system. See Section 1.2, OneCommand CNA Manager Secure Management, for more information.
`ERROR: <181>: Authentication: Insufficient credentials`	All	The specified user name and password are valid and the user is a member of an OneCommand CNA Manager group. However, the OneCommand CNA Manager group does not have sufficient privileges to execute the specified command. See Section 1.2, OneCommand CNA Manager Secure Management, for more information.
`ERROR: <183>: Secure Mgmt: user not a member of OCM group`	All	The specified user name and password could be authenticated, but the user is not a member of an OneCommand CNA Manager group. See Section 1.2, OneCommand CNA Manager Secure Management, for more information.
`ERROR: <206>: Authentication Failed`	All	This indicates either a valid user name with an invalid password, or a general user authentication error. See Section 1.2, OneCommand CNA Manager Secure Management, for more information.
`Not supported.`	`ChangeWWN`	If a volatile change is requested on an adapter that does not support volatile WWNs, it returns a `Not Supported` error. See Section 5.19.1, ChangeWWN.
`ERROR: This port is configured to be removed after reboot. Channel Management properties are unavailable.`	`CMGetParams` `UmcGetParams`	If you specify a MAC address that belongs to a newly disabled port, an error message will be displayed indicating that the operation cannot be performed on this port. See Section 5.5.1, CMGetParams.

`ERROR: This port is configured to be removed after reboot. Bandwidth values can only be set on ports that will be available after reboot.`	`CMSetBW`
	`UmcSetBW`
	`UmcSetLPVID`
		If you specify a MAC address that belongs to a newly disabled port, this error message is displayed indicating that the operation cannot be performed on this port. See Section 5.5.3, CMSetBW, and Section 5.17, UMC Commands.

`ERROR: This port is configured to be removed after reboot. LPVID values can only be set on ports that will be available after reboot.`	`CMSetLPVID`	If you specify a MAC address that belongs to a newly disabled port, an error message will be displayed indicating that the operation cannot be performed on this port. See Section 5.5.4, CMSetLPVID.
`ERROR: Download Failed due to invalid firmware digital signature. Please contact customer support for additional help.` `ERROR: <203>: Failed validating firmware digital signature`	`Download`	If you attempt to update unauthenticated firmware for a secure OCe14000B-series adapter, this error message is displayed. See Section 5.13.4, Download.

Table 9 OneCommand CNA Manager Error and Warning Messages (Continued)

Error Message	Command	Description
ERROR: Download Failed due to missing digital signature in firmware file. Please contact customer support for additional help. ERROR: <209>: Firmware digital signature missing	Download	If you attempt to update unsecured firmware for a secure OCe14000B-series adapter, this error message is displayed. See Section 5.13.4, Download.
Warning: SR-IOV is not enabled at an adapter-wide level. Run the SetAdapterPortConfig command in order to enable SR-IOV on all NIC functions.	GetAdapterPortConfig	This message is displayed following the normal output of the GetAdapterPortConfig command indicating that not all NIC functions have SR-IOV enabled.
ERROR: <222>: DCB not available	GetDCBParams SetDCBParam GetPGInfo SetDCBPriority SetCnaPGBW	These commands are not available on OCe11101-EM/EX or OCe11102-EM/EX adapters. See Section 5.6, DCB Commands.
ERROR: HBACMD_GetDumpFile: RM_GetDumpFile call failed (2) ERROR: <2>: Not Supported	GetDumpFile	Dump files are copied from the Dump directory of the remote host to the Dump directory of the local host. Specifying a local port identifier for this command returns an error, because the source and destination directory are the same. See Section 5.9.4, GetDumpFile.
ERROR: The Adapter does not support Port Removal. F0 cannot be set to NONE. ERROR: The specified function cannot be set to NONE. ERROR: <4>: Invalid Argument	SetAdapterPortConfig	These error messages are generated if you attempt to remove a port on an unsupported adapter. See Section 5.16.4.2, Removing a Port.
ERROR: Invalid f0 parameter. Must be NIC or NIC+RoCE. ERROR: The specified function cannot be set to NONE. ERROR: <4>: Invalid Argument	SetAdapterPortConfig	These error messages are generated if you attempt to set the first function from the first port to None. See Section 5.16.4.2, Removing a Port.
ERROR: If f0 is set to None, all functions on the port must be set to NONE. ERROR: <4>: Invalid Argument	SetAdapterPortConfig	These error messages are generated if you attempt to set the first function from a port other than the first port to None. See Section 5.16.4.2, Removing a Port.
Adapter configuration successfully set to new configuration. System reboot required to activate it. The multichannel type has been changed. Be sure to use the additional multichannel commands to set the other multichannel properties in order to complete the configuration. Multichannel properties can only be set on newly added ports after system reboot.	SetAdapterPortConfig	If you add a port back while also enabling multichannel mode, a warning message is displayed to notify you that multichannel properties can only be set on this port after a system reboot has been performed. See Section 5.16.4.3, Adding a Port While Enabling Multichannel Mode.

Table 9 OneCommand CNA Manager Error and Warning Messages (Continued)

Error Message	Command	Description
`Error: <431> Cable length required for force mode and interface type`	`SetPhyPortSpeed`	This error is displayed when a length value is not included when the mode is set to 2. Example: `brcmhbacmd setphyportspeed 00-00-c9-a9-41-88 2 100Mb` See Examples.

`There are no license features for this adapter`	`ShowLicenseAdapterID`
	`InstallAdapterLicense`
	`ShowAdapterLicenseFeatures`
		Adapter License Management commands are not available on OCe11101-EM/EX or OCe11102-EM/EX adapters. See Section 5.2, Adapter License Management Commands.

`ERROR: <251>: Hardware or firmware does not support command.`	`SRIOVEnable`
	`UmcEnable`
	`UmcGetParams`
	`UmcSetBW`
	`UmcSetLPVID`
	`SetPhyPortSpeed`
		These commands are not available on OCe11101-EM/EX or OCe11102-EM/EX adapters. See Section 5.13.14, SRIOVEnable, Section 5.17, UMC Commands, and Section 5.3.5, SetPhyPortSpeed.

`Hardware Does Not Support`	`VEPAEnable`	VEPA Management is only supported for OCe1400x NIC ports that have SR-IOV enabled. For all other board types and port types, this feature is undefined. See Section 5.13.16, VEPAEnable.
`ERROR: <223>: VEPA is only settable on NIC ports when the reboot configuration is non-multichannel and SR-IOV is enabled.`	`VEPAEnable`	When the next boot configuration does not have SR-IOV enabled or multichannel is enabled, this error message is displayed.

Appendix B: License Notices

B.1 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.
*/

B.2 OpenPegasus Licensing Notice

Licensed to The Open Group (TOG) under one or more contributor license agreements. Refer to the OpenPegasusNOTICE.txt file distributed with this work for additional information regarding copyright ownership. Each contributor licenses this file to you under the OpenPegasus Open Source License; you may not use this file except in compliance with the License.
Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit

persons to whom the Software is furnished to do so, subject to the following conditions:
The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.
THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

B.3 OpenSSL Notice

This is a copy of the current LICENSE file inside the CVS repository.
LICENSE ISSUES
==========
The OpenSSL toolkit stays under a dual license, i.e. both the conditions of the OpenSSL License and the original SSLeay license apply to the toolkit. below for the actual license texts. Actually both licenses are BSD-style Open Source licenses. In case of any license issues related to OpenSSL please contact openssl-core@openssl.org.
OpenSSL License

/* ==
* Copyright (c) 1998-2011 The OpenSSL Project. All rights reserved.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
*
* 1. Redistributions of source code must retain the above copyright
* notice, this list of conditions and the following disclaimer.
*
* 2. Redistributions in binary form must reproduce the above copyright
* notice, this list of conditions and the following disclaimer in
* the documentation and/or other materials provided with the
* distribution.
*
* 3. All advertising materials mentioning features or use of this
* software must display the following acknowledgment:
* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit. (http://www.openssl.org/)"
*
* 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to
* endorse or promote products derived from this software without
* prior written permission. For written permission, please contact
* openssl-core@openssl.org.

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

Original SSLeay License

/* Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)
* All rights reserved.
*
* This package is an SSL implementation written
* by Eric Young (eay@cryptsoft.com).
* The implementation was written so as to conform with Netscapes SSL.
*
* This library is free for commercial and non-commercial use as long as
* the following conditions are aheared to. The following conditions
* apply to all code found in this distribution, be it the RC4, RSA,
* lhash, DES, etc., code; not just the SSL code. The SSL documentation
* included with this distribution is covered by the same copyright terms
* except that the holder is Tim Hudson (tjh@cryptsoft.com).

```
*
* Copyright remains Eric Young's, and as such any Copyright notices in
* the code are not to be removed.
* If this package is used in a product, Eric Young should be given attribution
* as the author of the parts of the library used.
* This can be in the form of a textual message at program startup or
* in documentation (online or textual) provided with the package.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
* 1. Redistributions of source code must retain the copyright
* notice, this list of conditions and the following disclaimer.
* 2. Redistributions in binary form must reproduce the above copyright
* notice, this list of conditions and the following disclaimer in the
* documentation and/or other materials provided with the distribution.
* 3. All advertising materials mentioning features or use of this software
* must display the following acknowledgement:
* "This product includes cryptographic software written by
* Eric Young (eay@cryptsoft.com)"
* The word 'cryptographic' can be left out if the 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.]
*/
```


BROADCOM

Emulex® Brcmflash and LpCfg Management Utilities for OneConnect® Adapters

User Guide

Version 11.2
December 30, 2016

FlashUTIL-OCA-UG112

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

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

Copyright © 2003–2017 by Broadcom. All Rights Reserved.

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

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.

Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Chapter 1: Introduction

This document explains how to use the adapter management utility kits for Linux®, FreeBSD, VMware, and Windows pre-boot environments.

1.1 Brcmflash Offline Kit

The Brcmflash Offline kit allows you to configure Emulex® adapters before you install or boot a server operating system. You can also use the Brcmflash Standalone kit, which lets you update firmware on Emulex adapters without installing any tools. See Section 4.2, Running the LpCfg Utility Included with the Brcmflash Standalone Kit, for more information.

The Brcmflash Offline kit includes:

- brcmlinlpcfg and brcmflash for FreeBSD
- brcmlinlpcfg and brcmflash for Linux
- brcmwinlpcfg and brcmflash for Windows

The brcmwinlpcfg and the brcmlinlpcfg are management utilities for adapter configuration and diagnostics for Windows and Linux, respectively. brcmlinlpcfg is also used for the FreeBSD operating system. The LpCfg utility can also download firmware, but you must manually select the adapters to be updated.

Brcmflash is a firmware download utility. When the `/auto` switch is used, Brcmflash automatically chooses the best image to download. When the `fwmatrix.txt` file is used, you choose the images to apply to the adapters.

NOTE The Offline utility must be installed on your system before you can run it. While the standalone utility does not need to be installed before it is run, you must go to the directory into which the standalone utility was extracted and run it from the command prompt.

See Chapter 4, LpCfg Utility Command Line Interface, for further information about the offline utility.

1.2 Brcmflash Utility

The Brcmflash utility uses a command line interface allowing you to build scripts for automated and unattended firmware and boot code download solutions for Emulex OneConnect® CNAs in production systems. You can also download firmware and boot code on local and remote machines simultaneously.

The Brcmflash Standalone kit, which allows you to update firmware on Broadcom Emulex adapters without needing to install any tools, is also available. See Section 4.2, Running the LpCfg Utility Included with the Brcmflash Standalone Kit, for more information.

The Brcmflash utility uses auto-discovery (`/auto`), so you do not need to maintain the `fwmatrix.txt` configuration file. When auto-discovery is invoked, the utility automatically discovers local adapters and performs selected operations, such as: `/ff` (force firmware), `/fb` (force boot), `/update`, `/downgrade`, and `/rewrite`.

When using auto-discovery, you must populate the firmware and boot subdirectories with image files for specific adapter models. Auto-discovery uses Emulex-defined firmware and boot code file names in their respective directories.

If you do not want to use auto-discovery, the Brcmflash utility can update the firmware or boot code of an adapter using the `fwmatrix.txt` configuration file. When using the `fwmatrix.txt` file, you must update each supported

adapter's type, firmware, and boot code entries and place the corresponding firmware and boot code images into their respective directories. The `fwmatrix.txt` file enables all operations that are used with auto-discovery.

1.3 Command Syntax

In all cases, the commands are given with the prefix *brcmos*. When you enter any command, you must replace *brcmos* with either:

- **./brcmlin** – for the Linux LpCfg utility commands, for use on Linux and FreeBSD machines.
- **brcmwin** – for the WinPE LpCfg utility commands, for use on WinPE machines.

The LpCfg utility commands are not case-sensitive, nor are their arguments. You can enter them in upper, lower, or mixed case. However, for Linux operating systems, use *brcmlinlpcfg* (all lowercase) for the utility commands; for example, ./brcmlinlpcfg listhba.

The LpCfg utility commands use both:

- One space between the command name and the first argument.
- One space between additional arguments.

NOTE Do not put a space before or after the equal sign within an argument.

World Wide Name (WWN) values are reported and specified with two 4-byte hexadecimal words – WWN word 0 (w0) and WWN word 1 (w1). To make the full WWN, concatenate the w0 and w1 values.

For detailed information on the command line interface, see Chapter 3, Brcmflash Utility Command Line Interface, and Chapter 4, LpCfg Utility Command Line Interface.

1.4 Supported Platforms

For supported adapters, and supported versions of operating systems and platforms, see http://www.broadcom.com.

1.5 Abbreviations

BIOS	basic input/output system
CNA	converged network adapter
DMA	direct memory access
EROM	erasable read-only memory
FCoE	Fibre Channel over Ethernet
HEX	hexadecimal
IEEE	Institute of Electrical and Electronics Engineers, Inc.
GUI	graphical user interface
iSCSI	Internet Small Computer System Interface
ISO	International Organization for Standards
LUN	logical unit number
MAC	media access control
MILI	management interface library
ms	milliseconds
NIC	network interface card
PCI	Peripheral Controller Interconnect
PCI_ID	PCI device identification number
POST	power-on self-test
ROM	read-only memory
SAN	storage area network
VPD	Vital Product Data
VWWNN	Virtual World Wide Node Name
VWWPN	Virtual World Wide Port Name
WWN	World Wide Name
WWNN	World Wide Node Name
WWPN	World Wide Port Name
XML	Extensible Markup Language

Chapter 2: Installation

This chapter details prerequisites and procedures for installing and uninstalling the Brcmflash Offline and Standalone kits for Linux, VMware, Windows, and FreeBSD systems. It also describes the update procedure for each supported operating system.

2.1 Platform Prerequisites

The utilities included in the Brcmflash Offline and Standalone kits for Linux and FreeBSD have prerequisites that must be installed prior to installing the utilities. There are no prerequisites for installing the Brcmflash Offline and Standalone kits on the VMware and Windows operating systems.

2.1.1 Linux

The Linux platform uses additional software, such as libraries, to run the different versions of the utility. This section identifies the prerequisites.

NOTE The Linux driver must be installed on the system for the applications to work properly.

2.1.1.1 Brcmflash Offline and Brcmflash Standalone Kits

The following software must be installed to run the utility.

- libnl
- ethtool
- lspci
- libsysfs (optional)

2.1.2 FreeBSD

The FreeBSD platform uses additional software, such as libraries, to run the Brcmflash utilities. These prerequisites are listed below and are provided by the operating system.

- libstdc++
- libm
- libc
- libgcc

NOTE BASH is required to run the install and uninstall scripts. If it is not installed, you must install the BASH package, which is part of the FreeBSD distribution.

2.2 Installing the Brcmflash Offline Kit for Linux

The Brcmflash Offline Kit for Linux uses an install script to install the brcmflash and linlpcfg utilty RPM packages. The install script determines the correct architecture and distribution, and it updates the existing Brcmflash Offline and

Brcmlinlpcfg RPM packages. If there are no existing Brcmflash Offline or Brcmlinlpcfg RPM packages, the install script installs the packaged Brcmflash Offline and Brcmlinlpcfg RPM packages.

To install the Brcmflash Offline Kit for Linux:

1. Untar the installation tarball.
2. Run the install script located on the root of the installation kit.

Example:

```
$ tar zxvf brcmflashOffline-linux-<version>-<rel>.tgz
$ cd brcmflashOffline-linux-<version>-<rel>
$ ./install.sh
```

2.2.1 Updating

To update an existing Brcmflash Offline Kit for a Linux installation, run the install.sh script to update the Brcmflash Offline and Brcmlinlpcfg RPM packages.

2.2.2 Uninstalling

The Brcmflash Offline Kit for Linux uses an uninstall script to uninstall the utility. The uninstall script performs specific actions depending on the switches that are used.

- `./uninstall.sh` – Uninstalls the Brcmflash Offline and Brcmlinlpcfg RPM packages.
- `./uninstall.sh -h` – Displays a summary of all available switches

To uninstall the Brcmflash Offline Kit for Linux, run the uninstall script located on the root of the installation kit.

Example:

```
$ cd brcmflashOffline-linux-<version>-<rel>
$ ./uninstall.sh
```

2.3 Installing the Brcmflash Offline Kit for Windows PE

To install the Brcmflash Offline kit for Windows PE:

1. Unzip the `brcmOffline-WinPE-<version>-<rel>.zip`.
2. Change the directory to the correct architecture subdirectory.
3. Run `setupBrcmAll-<arch>.exe`.

The following components are installed:

- Storport CNA driver
- OneConnect iSCSI and NIC drivers
- brcmwinlpcfg – Brcmflash Offline

NOTE The setupBrcmAll-<*arch*>.exe file is compatible only with WinPE. It does not function with any version of Windows Server.

2.3.1 Updating

To update an existing Brcmflash Offline kit installation, run the same steps that you used for installing the Brcmflash Offline kit for Windows PE above.

The installer uninstalls the existing version and then installs the updated version.

2.3.2 Uninstalling

To uninstall the Brcmflash Offline kit for Windows PE, run the following commands:

```
cd \ProgramData\Tarma Installer\
setup.exe /remove
```

2.4 Extracting the Driver Files from the Brcmflash Offline Kit for Windows

To extract the driver files from the Brcmflash Offline kit, you can use the installer or the command line. When extracting these files, all files, both x64 and x86, are extracted to <*Drive*>:\My Documents\Broadcom\Drivers. This location is hard-coded and cannot be changed.

NOTE Extracting all the drivers takes an additional 28 MB. By default, Windows PE allocates only 32 MB of writable memory. The size of the writable memory can be increased when building the WinPE ISO image.

You can extract the driver files using the GUI mode by running the installer, or you can use the command line in silent mode.

To extract the drivers from the Brcmflash Offline kit for Windows, perform the following steps:

1. In GUI Mode:
   a. Run the installer.
   b. From the Installation Options screen, select **Extract All Drivers**, and deselect **Install brcmApp_Drivers (64-bit)**.
   c. Click **Install** to continue.
2. In Silent Mode, from the command line, type:
   ```
   start/wait SetupBrcmAll-x64.exe/q2 extract=2
   ```

NOTE Values for the *extract* parameter are:

- 0 – Install the package; do not extract the drivers (default)
- 1 – Install the package; extract the drivers
- 2 – Do not install the package; extract the drivers

2.5 Extracting the Brcmflash Standalone Kit for Linux, Windows, and VMware

The Brcmflash Standalone kit is not installed; you must extract it from its zip file. After you extract the utility, the following directories are created. The Linux executable files are extracted to the lx\ directory, the VMware executable files are extracted to the esxi\ directory, or the Windows executable files are extracted to the win\ directory.

- boot\

- firmware\
- lx\
- esxi\
- win\

In Windows, for example, three directories are visible in the kit's root directory after you extract the standalone utility. The top-level boot directory and firmware directory are the default directories in which Brcmflash looks for boot code and firmware images. Ensure that firmware and boot code are located in these directories.

```
# dir

07/05/2012  07:02 PM    <DIR>          .
07/05/2012  07:02 PM    <DIR>          ..
07/02/2012  03:14 PM    <DIR>          boot
07/02/2012  03:14 PM    <DIR>          firmware
07/02/2012  03:15 PM    <DIR>          win
```

The Brcmflash Standalone kits for Linux have the same dependencies as the Brcmflash Offline kits. See Section 2.1, Platform Prerequisites.

NOTE Update and uninstall procedures for the Brcmflash Standalone kit are not needed. When a new version is available you extract it, replacing the older version.

NOTE To manage OneConnect adapters, you must install the operating system-specific NIC, iSCSI, and FCoE drivers.

2.6 Installing the Brcmflash Offline Kit for FreeBSD

The Brcmflash Offline kit installation for FreeBSD uses an install script. The install script determines the correct architecture and distribution and performs the following operations:

- Installs the Brcmflash utility and linlpcfg
- Temporarily mounts `/proc`

NOTE For security reasons, `/proc` is not mounted by default; however, you must mount `/proc` to run Brcmflash and linlpcfg. Verify that `/proc` is mounted before running Brcmflash or linlpcfg.

To install the Brcmflash Offline kit for FreeBSD:

1. Untar the installation tarball.
2. Run the install script located in the root directory of the installation kit.

For example:

```
# tar xvf brcmflashOffline-FreeBSD-<version>-<rel>.tgz
# cd brcmflashOffline-FreeBSD-<version>-<rel>
# ./install.sh
```

2.6.1 Updating

To update the Brcmflash Offline kit for FreeBSD, uninstall the old version and then unzip the new version.

2.6.2 Uninstalling

To uninstall the Brcmflash Offline kit for FreeBSD:

Run the uninstall script located in the root directory of the installation kit.

For example:

```
# cd brcmflashOffline-FreeBSD-<version>-<rel>
# ./uninstall.sh
```

2.7 Firmware and Boot Code

NOTE If a secure version of firmware (version 11.0 or later) is installed on an OCe14000B-series adapter and you want to update to an earlier, unsecure version of firmware, you must remove the secure firmware jumper block before performing the update. See the installation manual for the adapter for more information.

Firmware images are available on the Broadcom® support site at http://www.broadcom.com. All firmware image files are expected to be in the sub-directory named *firmware*.

NOTE OneConnect models (FCoE, iSCSI, and NIC) might share the same Emulex model names, but they can run different protocols. For example, an OCe14000B-series adapter can be an FCoE, iSCSI, or NIC CNA.

When the `fwmatrix.txt` file is used, the Brcmflash model name must include the highest protocol being used on that model. An adapter's Brcmflash model name can be seen by running the `/query` command.

NOTE On an FCoE/NIC adapter, the highest protocol is FCoE. On an iSCSI/NIC adapter, the highest protocol is iSCSI. Ensure that there are no spaces in the fwmatrix.txt model name to avoid a firmware update failure.

For example:

```
# ./brcmflash /q
HBA=OCe11100-iSCSI, Port Type=iSCSI, MAC=00:00:C9:AD:AD:21, PCI ID=712, VID=19A2,
SSID=E702, SVID=10DF, Firmware=4.0.493.0, Boot Code=2.0.21.768
HBA=OCe11100-iSCSI, Port Type=iSCSI, MAC=00:00:C9:AD:AD:25, PCI ID=712, VID=19A2,
SSID=E702, SVID=10DF, Firmware=4.0.493.0, Boot Code=2.0.21.768
HBA=OCe11100-iSCSI, Port Type=NIC, MAC=00:00:C9:AD:AD:20, PCI ID=710, VID=19A2,
SSID=E702, SVID=10DF, Firmware=4.0.493.0, Boot Code=2.0.21.768
HBA=OCe11100-iSCSI, Port Type=NIC, MAC=00:00:C9:AD:AD:24, PCI ID=710, VID=19A2,
SSID=E702, SVID=10DF, Firmware=4.0.493.0, Boot Code=2.0.21.768
brcmflash.exe: All required queries succeeded - Return Code=0
```

NOTE The OCe11100 CNA in the example above is an iSCSI adapter, which has the Brcmflash model name of OCe11100-iSCSI.

Chapter 3: Brcmflash Utility Command Line Interface

Two supported modes for each Brcmflash utility switch are available. The first mode relies on the `fwmatrix.txt` file. It is your responsibility to update the firmware directories with the appropriate `firmware/boot` code images.

NOTE Verify that there are no spaces in the `fwmatrix.txt` model name to avoid a firmware update failure.

The second mode is auto-discovery. When the `/auto` switch is used with the following switches, the Brcmflash utility automatically discovers adapters and, using the firmware and boot subdirectories, performs the specified operation on each adapter.

- /f/fb
- /ff
- /downgrade
- /rewrite
- /update

3.1 Auto-Discovery (/auto)

Usage: `/auto`

The `/auto` switch instructs the Brcmflash utility to ignore the `fwmatrix.txt` file, automatically discover local adapters, and perform specified operations employing an additional switch using the firmware directory.

The `/auto` switch must be used with an additional operational switch, such as:

- /f
- /ff
- /fb
- /downgrade
- /rewrite
- /update

Example usage:

`./brcmflash /auto /update` – Updates the firmware/boot code using the `firmware` directory.

- The `fwmatrix.txt` file is ignored. You must place the desired versions of firmware in the `firmware` directory.
- Using the `firmware` directory, the Brcmflash utility automatically discovers the best matching firmware for each installed and supported adapter.
- If multiple versions of firmware or boot code are found for an adapter, the Brcmflash utility uses the most recent version when performing the firmware download.

3.2 Downgrade (/downgrade or /g)

Usage: `/downgrade or /g`

The `/downgrade` switch downgrades the `firmware/boot` code of each adapter if the currently installed version is more recent than the downgrade version. This switch cannot be used with the `/update` or `/rewrite` commands.

NOTE If a secure version of firmware (version 11.0 or later) is installed on an OCe14000B-series adapter and you want to update to an earlier, unsecure version of firmware, you must remove the secure firmware jumper block before performing the update. Refer to the installation manual for the adapter for more information.

Example usage:

`./brcmflash /downgrade /auto` – Downgrades the firmware boot code using auto-discovery.

- The `fwmatrix.txt` file is ignored. You must place the desired downgrade version of firmware/boot code in the firmware directory.
- If the downgrade version is older than the currently installed version on the adapter, then the downgrade versions are downloaded to the adapter.
- If multiple downgrade versions of firmware/boot code are found for an adapter, the next-previous downgrade version is downloaded to the adapter.

`./brcmflash /downgrade` – Downgrades the firmware/boot code using the `fwmatrix.txt` file.

- For each installed and supported adapter, the current firmware/ boot code version is compared with the version specified in `fwmatrix.txt`.
- If the downgrade version in `fwmatrix.txt` is older than the currently installed version, the downgrade version of firmware/ boot code is downloaded to that adapter.

3.2.1 FCoE Adapter Download Summary

```
<date><time>
HBA=<model>, Port Type=<port_type>, WWN=<wwn>,
Update=Firmware, Image=<image>, New=<version>, Old=<version>,
Status=<description>
```

3.2.2 iSCSI and NIC-only Adapter Summary

```
<date><time>
HBA=<model>, Port Type=<port_type>, MAC=<mac_address>,
Update=Firmware, Image=<image>, New=<version>, Old=<version>,
Status=<description>
Return Code=<n>
```

where <*description*> is `Success` or `Error`, and <*n*> = `0` for completion with no errors or a non-zero error code for any error.

NOTE To activate new firmware, perform a system reboot.
If the preview switch is used, the `Status=`<*description*> field displays `Preview`.

3.3 Force Firmware and Boot Code (/f)

Usage: `/f`

The `/f` switch forces a firmware/ boot code download to an adapter regardless of the current version on the adapter, and it is performed regardless of any additional operational switches given on the command line. Also, see Section 3.4, Force Firmware (/ff).

NOTE If a secure version of firmware (version 11.0 or later) is installed on an OCe14000B-series adapter and you want to force an update to an earlier, unsecure version of firmware, you must remove the secure firmware jumper block before performing the update. See the installation manual for the adapter for more information.

Example usage:

`./brcmflash /f /auto` – Forces a firmware/boot code download using auto-discovery.

- The `fwmatrix.txt` file is ignored. You must place the desired version of firmware/boot code in the firmware directory.
- If multiple versions of firmware/boot code are found for an adapter, the Brcmflash utility uses the most recent version when performing the `firmware/boot` code downloads.

`./brcmflash /f` – Forces a firmware/boot code download using the `fwmatrix.txt` file.

- For each installed and supported adapter, forces a download of firmware/boot code using the version specified in the `fwmatrix.txt` file.

3.3.1 FCoE Adapter Download Summary

```
<date><time>
HBA=<model>, Port Type=<port_type>, WWN=<wwn>,
Update=Firmware, Image=<image>, New=<version>, Old=<version>,
Status=<description>
```

3.3.2 iSCSI and NIC-only Adapter Summary

```
<date><time>
HBA=<model>, Port Type=<port_type>, MAC=<mac_address>,
Update=Firmware, Image=<image>, New=<version>, Old=<version>,
Status=<description>
Return Code=<n>
```

where <*description*> is `Success` or `Error`, and <*n*> = `0` for completion with no errors or a non-zero error code for any error.

NOTE To activate new firmware, perform a system reboot.

If the preview switch is used, the `Status=`<*description*> field displays `Preview`.

3.4 Force Firmware (/ff)

Usage: `/ff`

The `/ff` switch forces a firmware download to an adapter regardless of the firmware version installed on the adapter, and it is performed regardless of any additional operational switches given on the command line.

NOTE If a secure version of firmware (version 11.0 or later) is installed on an OCe14000B-series adapter and you want to update to an earlier, unsecure version of firmware, you must remove the secure firmware jumper block before performing the update. See the installation manual for the adapter for more information.

NOTE Because boot code is included in the firmware image for OneConnect adapters, this command has the same result as the `/f` command, which forces a firmware and boot code download.

Example usage:

`./brcmflash /ff /auto` – Forces a firmware download using auto-discovery.

- The `fwmatrix.txt` file is ignored. You must place the desired version of firmware in the firmware directory.
- If multiple versions of firmware are found for an adapter, the Brcmflash utility uses the most recent version when performing the firmware download.

`./brcmflash /ff` – Forces a firmware download using the `fwmatrix.txt` file.

- For each installed and supported adapter, forces a download of firmware using the firmware version specified in the `fwmatrix.txt` file.

3.4.1 FCoE Adapter Download Summary

```
<date><time>
HBA=<model>, Port Type=<port_type>, WWN=<wwn>,
Update=Firmware, Image=<image>, New=<version>, Old=<version>,
Status=<description>
```

3.4.2 iSCSI and NIC-only Adapter Summary

```
<date><time>
HBA=<model>, Port Type=<port_type>, MAC=<mac_address>,
Update=Firmware, Image=<image>, New=<version>, Old=<version>,
Status=<description>
Return Code=<n>
```

where <*description*> is `Success` or `Error`, and <*n*> = `0` for completion with no errors or a non-zero error code for any error.

NOTE To activate new firmware, perform a system reboot.

If the preview switch is used, the `Status=`<*description*> field displays `Preview`.

3.5 Firmware Matrix Directory Change (/fmd)

Usage: `/fmd=<directory>`

The `/fmd` (firmware matrix directory change) switch changes the location of the base directory in which `fwmatrix.txt` is located.

3.6 Firmware Flash Override (/fo)

Usage: `/fo`

The `/fo` switch is used to run a firmware download using the flash executable file on supported CNAs.

NOTE The flash override switch applies to Linux Brcmflash Offline and Brcmflash Standalone utilities only.

Example usage:

`./brcmflash /f /auto /fo` – Forces a firmware/boot code download using auto-discovery, and downloads the firmware to supported CNAs using the flash executable file.

- The `fwmatrix.txt` file is ignored. The desired version of firmware/boot code must be placed in the `firmware` directory.
- If multiple versions of firmware/boot code are found for an adapter, Brcmflash uses the most recent versions when performing the firmware/boot code downloads.
- The flash executable file is used to download firmware to supported CNAs.

`./brcmflash /f /fo` – Forces a firmware/boot code download using the `fwmatrix.txt` file, and downloads the firmware to supported CNAs using the flash executable file.

- For each installed and supported adapter, a forced download of firmware/boot code occurs using the versions specified in the `fwmatrix.txt` file.
- The flash executable file is used to download firmware to supported CNAs.

3.7 Help (/h or /?)

Usage: `/h or /?`

The `/h` switch displays a help message detailing instructions on how to use the Brcmflash utility.

3.8 Image Directory Change (/id)

Usage: `/id=<image_directory>`

The `/id` switch is used to specify the location of the `firmware` directory.

Example usage:

`./brcmflash /f /auto /id=/tmp` – Forces a firmware/boot code download using auto-discovery.

- Brcmflash looks for the `firmware` directory in the `/tmp` directory.

- The `fwmatrix.txt` file is ignored. The desired version of firmware/boot code must be placed in the `firmware` directory.
- If multiple versions of firmware/boot code are found for an adapter, Brcmflash uses the most recent version when performing the firmware/boot code downloads.

`./brcmflash /f /id/tmp` – Forces a firmware/boot code download using the `fwmatrix.txt` file.

- Brcmflash looks for the `firmware` directory in the `/tmp` directory.
- For each installed and supported adapter, a forced download of firmware/boot code occurs using the version specified in the `fwmatrix.txt` file.
- In this example, firmware must be placed in the `/tmp/firmware` directory.

NOTE This switch was `/i` in previous versions of the Brcmflash utility.

3.9 Image Version Display (/iv)

Usage: `/iv=<`*`image file`*`>`

The `/iv` switch shows the firmware image file's version number. You must specify a path to the firmware image file for the command to decode the image file's version.

Example usage:

```
./brcmflash /iv=firmware/oc14-

11.2.123.45.ufi
11.2.123.45
```

3.9.1 Log(/log)

Usage: `/log=<`*`logfile.txt`*`>`

The `/log` switch appends the output of the Brcmflash utility to a text file. `Log` can be used with any switch.

3.10 Preview (/p)

Usage: `/p`

The `/p` switch provides a download preview of all adapters the Brcmflash utility can update using either auto-discovery or the `fwmatrix.txt` file. The preview switch can be used with any of the operational switches, such as:

- /ff
- /downgrade
- /rewrite
- /update

When the /p switch is used, the Brcmflash utility displays a download summary, but it does not actually perform the download.

Each adapter's download preview displays the adapter's old and new image version. The old image version represents the image version that is currently on the adapter. The new image version represents the image version the Brcmflash utility would use during a download.

Example usage:

`./brcmflash /preview /auto /update` – Previews an upgrade of firmware/boot code using auto-discovery.

- The `fwmatrix.txt` file is ignored. You must place the desired update version of firmware/boot code in the `firmware` directory.
- If the update versions are later than the currently installed versions on the adapter, the Brcmflash utility provides a download preview for each adapter that can be updated.
- If multiple update versions of firmware/boot code are found for an adapter, the Brcmflash utility provides a download preview using the most recent version.

`./brcmflash /preview` – Previews an upgrade of firmware/boot code using the `fwmatrix.txt` file.

- For each installed and supported adapter, the current firmware/boot code version is compared to the version specified in `fwmatrix.txt`.
- If the update version in `fwmatrix.txt` is more recent than the currently installed version, the Brcmflash utility provides a download preview of firmware/boot code for each adapter that can be updated.

3.10.1 FCoE Adapter Download Summary

```
<date><time>
HBA=<model>, Port Type=<port_type>, WWN=<wwn>,
Update=Firmware, Image=<image>, New=<version>, Old=<version>, Status=Preview
```

3.10.2 iSCSI and NIC-only Adapter Download Summary

```
<date><time>
HBA=<model>, Port Type=<port_type>, MAC=<mac_address>, Update=Firmware,
Image=<image>, New=<version>, Old=<version>, Status=Preview
```

3.11 Query (/q)

Usage: `/q`

The `/q` switch displays an adapter's model, WWN or MAC address, PCI_ID, and firmware/boot code version. Query can be used with any switch.

Example usage:

```
C:\brcmflashStandalone-windows-10.0.567.22-1\win>brcmflash.bat /q
```

3.11.1 FCoE Adapter Status Summary

```
HBA=<model>, Port Type=<port_type>, WWN =<wwn>, Firmware=<version>, Boot
Code=<version>
```

3.11.2 iSCSI and NIC Adapter Status Summary

```
HBA=<model>, Port Type=<port_type>, MAC =<mac_address>, Firmware=<version>, Boot
Code=<version>
```

NOTE When the query switch is used with an operational switch, the query also includes an additional field called `Supported Firmware` or `Supported Boot Code`.

3.12 Ramdrive (/ramdrive) - (Windows Offline Adapter Management Utility only)

Usage: `/ramdrive=<drive letter>[:]`

The `/ramdrive` switch specifies the drive on which to create temporary files and log files. The default is drive X; the default drive in Windows PE.

3.13 Rewrite (/rewrite or /e)

Usage: `/rewrite -or- /e`

The `/rewrite` switch updates the firmware/boot code of each adapter if the installed version is earlier than, or the same as, the rewrite version. This switch cannot be used with `/update` or `/downgrade`.

Example usage:

`./brcmflash /rewrite /auto` – Rewrites the firmware/boot code using auto-discovery.

- The `fwmatrix.txt` file is ignored. You must place the desired rewrite versions of firmware/boot code in the `/tmp/firmware` directory.
- If the rewrite version is more recent than, or the same as, the versions installed on the adapter, the rewrite version is downloaded to the adapter.
- If multiple rewrite versions of firmware/boot code are found for an adapter, the most recent version is downloaded to the adapter.

`./brcmflash /rewrite` – Rewrites the firmware/boot code using the `fwmatrix.txt` file.

- For each installed and supported adapter, the current firmware/boot code version is compared with the version in `fwmatrix.txt`.
- If the currently installed version is less than or equal to the rewrite version in `fwmatrix.txt`, the rewrite version of firmware/boot code is downloaded to the adapter.

3.13.1 FCoE Adapter Download Summary

```
<date><time>
HBA=<model>, Port Type=<port_type>, WWN=<wwn>,
Update=Firmware, Image=<image>, New=<version>, Old=<version>,
Status=<description>
```

3.13.2 iSCSI and NIC-only Adapter Download Summary

```
<date><time>
HBA=<model>, Port Type=<port_type>, MAC=<mac_address>,
Update=Firmware, Image=<image>, New=<version>, Old=<version>,
Status=<description>
Return Code=<n>
```

where <*description*> is Success or Error, and <*n*> = 0 for completion with no errors or a non-zero error code for any error.

NOTE To activate new firmware, perform a system reboot.
If the preview switch is also used, the Status =<*description*> field displays Preview.

3.14 Silent (/s)

Usage: /s

The /s switch prevents all output from being displayed.

Example usage:

./brcmflash /f /auto /s – Forces a firmware/boot code download using auto-discovery and mutes all output to stdout.

- The fwmatrix.txt file is ignored. The desired versions of firmware/boot code must be placed in the firmware directory.
- If multiple versions of firmware/boot code are found for an adapter, Brcmflash uses the most recent version when performing the firmware/boot code downloads.
- No output is printed to stdout.

./brcmflash /f /s – Forces a firmware/boot code download using the fwmatrix.txt file and mutes all output to stdout.

- For each installed and supported adapter, a download of firmware/boot code is forced using the versions specified in the fwmatrix.txt file.
- No output is printed to stdout.

3.15 Discover Adapters Using SysFS (/sysfs)

NOTE The libsysfs library must be installed. If the libsysfs library is not found, the sysfs switch is not available.

Usage: /sysfs + operational switch

/sysfs is an operational switch that discovers CNA functions using the SysFS tool and runs firmware downloads by using the ethtool utility or SysFS interface. The /sysfs switch is used with any switch that displays adapter information; for example, /q, or any combination of switches that performs a download.

The following switches can be used:

- /f - force firmware and boot code
- /ff - force firmware
- /downgrade or /g
- /query or /q
- /rewrite or /e
- /update

For example, ./brcmflash /sysfs /auto /f discovers CNA NIC functions using SysFS and forces a firmware download on CNA adapters using ethtool or SysFS.

NOTE Only CNA NIC functions are displayed. Displaying iSCSI and FCoE functions is not supported.

When `/sysfs` is used, the VPD model name is not available.

The `/sysfs` switch is available on Brcmflash Standalone on RHEL 6, RHEL 7, SLES 11 SP1+, and SLES 12.

Example usage:

```
/brcmflash /sysfs /q
```

3.16 Update (/update)

Usage: `/update`

The `/update` switch updates the firmware/ boot code of each adapter if the currently installed version is older than the update version. This switch cannot be used with `/downgrade` or `/rewrite`.

Example usage:

`./brcmflash /update /auto` – Upgrades the firmware/boot code using auto-discovery.

- The `fwmatrix.txt` file is ignored. You must place the desired update versions of firmware/ boot code in the `firmware` directory.
- If the update version is newer than the currently installed version on the adapter, the update version is downloaded to the adapter.
- If multiple update versions of firmware/ boot code are found for an adapter, the most recent version is downloaded to the adapter.

`./brcmflash /update` – Upgrades the firmware/boot code using the `fwmatrix.txt` file.

- For each installed and supported adapter, the current firmware/boot code version is compared with the versions specified in `fwmatrix.txt`.
- If the update version in `fwmatrix.txt` is more recent than the currently installed version, the update version of firmware/ boot code is downloaded to that adapter.

3.16.1 FCoE Adapter Download Summary

```
<date><time>
HBA=<model>, Port Type=<port_type>, WWN=<wwn>,
Update=Firmware, Image=<image>, New=<version>, Old=<version>,
Status=<description>
```

3.16.2 iSCSI and NIC Adapter Download Summary

```
<date><time>
HBA=<model>, Port Type=<port_type>, MAC=<mac_address>,
Update=Firmware, Image=<image>, New=<version>, Old=<version>,
Status=<description>
Return Code=<n>
```

where <*description*> is `Success` or `Error`, and <*n*> = `0` for completion with no errors or a non-zero error code for any error.

NOTE To activate new firmware, perform a system reboot.

3.17 Verbose (/v)

Usage: /v

The /v switch displays progress messages, and it can be used with any switch. When this switch is used, the following information is displayed:

- A download summary for all adapters that had successful or failed downloads
- A summary of unsupported adapters, if applicable
- A per adapter message for each adapter the Brcmflash utility did not update

3.17.1 FCoE Adapter Download Summary

```
<date><time>
HBA=<model>, Port Type=<port_type>, WWN=<wwn>,
Update=Firmware, Image=<image>, New=<version>, Old=<version>,
Status=<description>
```

3.17.2 iSCSI and NIC-only Adapter Download Summary

```
<date><time>
HBA=<model>, Port Type=<port_type>, MAC=<mac_address>,
Update=Firmware, Image=<image>, New=<version>, Old=<version>,
Status=<description>
Return Code=<n>
```

where <description> is Success or Error, and <n> = 0 for completion with no errors or a non-zero error code for any error.

3.18 VPD (/vpd)

Usage: /vpd

The /vpd command displays the VPD model name for supported adapters. You must use the /vpd command with one of the following switches:

- /f - force firmware and boot
- /ff - force firmware
- /xml - XML output
- /downgrade or /g
- /query or /q
- /rewrite or /e
- /update

NOTE The /vpd command does not display any information if used without the switches listed above. If you attempt to use the /vpd command without the switches listed above, an error message is displayed.

If the /vpd command is used in conjunction with the fwmatrix.txt file, the model name in the fwmatrix.txt file must match the VPD model name reported by Brcmflash.

3.19 XML Output (/xml)

Usage: `/xml`

The `/xml` switch displays utility output in XML format.

Chapter 4: LpCfg Utility Command Line Interface

The LpCfg utility allows you to configure Emulex adapters before you install or boot a server operating system.

The LpCfg utility includes:

- brcmwinlpcfg
- brcmlinlpcfg

NOTE Usage code and examples show `brcmoslpcfg`. Note that the *brcmos* designation changes to *brcmwin* if you are using LpCfg for Windows and to *brcmlin* if you are using LpCfg for Linux.

You can use the LpCfg utility to do the following:

- View information about an Emulex adapter
- Reset the adapter
- Return the adapter to its factory default settings
- Download firmware/boot code files
- Select a boot device
- Read and update WWNs
- Read MAC addresses
- Enable boot code
- Update configuration regions
- Set the adapter to use soft jumpers
- Run diagnostic tests
- Read and process script files
- Read, write, and reset VLAN IDs and VLAN Priorities

4.1 Running the LpCfg Utility from the Command Prompt

NOTE The LpCfg utility is a command line utility and does not include a graphical user interface.

To run the LpCfg utility from the command prompt:

1. Boot the system with a supported operating system.
2. Start the LpCfg utility with a valid command or a valid script file name.
   - To start the LpCfg utility from the command line, move to the directory where the executable file resides and type:

     `brcmoslpcfg <valid command> (replace os with ./brcmlin or brcmwin as appropriate`
     - For example, starting the Windows LpCfg utility with a reset command:

       ```
       brcmwinlpcfg reset n=2
       ```
     - For example, starting the Linux LpCfg utility with a reset command:

       ```
       ./brcmlinlpcfg reset n=2
       ```

— To start the LpCfg utility with a script file name, move to the directory where the LpCfg utility resides and type:

`brcmoslpcfg @<script file name>`

- For example, starting the Windows LpCfg utility with `script1.txt` in the `c:\test` directory:

  `brcmoslpcfg @C:\test\script1.txt`

NOTE To redirect screen output to a file, add <*filename*> at the end of each command.
For example:
`brcmoslpcfg listboot n=1 >result.out`
For more information on script files, see Section 4.3.9, Using Script Files.

4.2 Running the LpCfg Utility Included with the Brcmflash Standalone Kit

The Brcmflash Standalone kit allows you to run brcmflash and lpcfg without installing the utilities. No drivers are included in the Standalone kit. No applications are installed, because the utility runs from inside the kit.

The Brcmflash Standalone Kit has the same capability as the Brcmflash Offline Kit, except you use the following scripts to run the utility:

- **Windows** – `brcmwinlpcfg.bat` and `brcmflash.bat`
- **Linux** – `brcmlinplcfg.sh` and `brcmflash.sh`

4.2.1 Windows

1. Extract the kit contents.
2. Change directory (`cd`) to `Brcmflash Standalone-windows-<version>`.
   The following directories must be present:
   — `firmware\`
   — `win\`
3. For Brcmflash, copy the firmware/boot code images to the `firmware` directory.

   NOTE For Brcmflash, each operating system architecture directory includes a `fwmatrix.txt` file. You must use the `fwmatrix.txt` directory that matches the current operating system architecture.
4. Change directory (`cd`) to the `win` directory.
   The following files and directories must be present:
   — `win32\`
   — `x64\`
   — `brcmflash.bat`
   — `brcmwinlpcfg.bat`
   The `brcmflash.bat` script is used to configure the environment. Run Brcmflash, and revert any changes before exiting. This script installs the MILI service if it is not already installed.
   The `brcmwinlpcfg.bat` script is used to configure the environment. Run `brcmwinlpcfg`, and revert any changes before exiting. This script installs the MILI service if it is not already installed.

NOTE The MILI service is temporarily installed. The service is removed after the script runs.

If OneCommand® CNA Manager is installed, the OneCommand CNA Manager libraries and OneCommand CNA Manager MILI service are used by the Brcmflash and brcmwinlpcfg utilities.

The `brcmwinlpcfg.bat` and `brcmflash.bat` scripts call the native versions of `brcmflash.exe` and `brcmwinlpcfg.exe`. For example, on Windows x64, the 64-bit utilities are called.

The adapter being managed by the utilities included in the Brcmflash Standalone kit must not be managed simultaneously by other Emulex utilities, including OneCommand CNA Manager.

5. To update firmware/boot code on an adapter in Windows, type:

```
# brcmflash.bat /auto /up
```

To display a list of adapters in Windows, type from the `Win` directory:

```
# brcmwinlpcfg.bat  listhba
```

Each time Brcmflash or lpcfg is run, a log file is created. On Windows these files are called:

```
C:\clu\log\brcmflash.log
C:\clu\log\brcmwinlpcfg.log
```

4.2.2 Linux and VMware

NOTE Running 32-bit applications on Linux x86_64 is not supported. The `sh` scripts run only native executable files.

1. Extract the kit contents.
2. Change directory (`cd`) to `BrcmflashStandalone-linux-<`*`version`*`>`.

   The following directories must be present:

   — `boot\`
   — `firmware\`
   — `lx\`
3. For Brcmflash, copy the firmware/boot images to the `firmware` directory.

NOTE For Brcmflash, each operating system architecture directory includes a `fwmatrix.txt` file. You must use the `fwmatrix.txt` directory that matches the current operating system architecture.

4. Change directory (`cd`) to the `lx` directory.

   The following files and directories must be present:

   — `i386\`
   — `x86_64\`
   — `ppc64\`
   — `brcmflash.sh`
   — `brcmlinlpcfg.sh`

   The `brcmflash.sh` script is used to configure the environment. Run Brcmflash, and revert any changes before exiting. This script installs the MILI daemons if they are not already installed.

   The `brcmlinlpcfg.sh` script is used to configure the environment. Run `brcmlinlpcfg`, and revert any changes before exiting. This script installs the MILI daemons if they are not already installed.

NOTE

The MILI service is temporarily installed. The service is removed after the script execution is completed.

If OneCommand CNA Manager is installed, the OneCommand CNA Manager libraries and OneCommand CNA Manager `brcmmilid` daemon are used by the `brcmflash` and `brcmlinlpcfg` utilities.

The `brcmlinlpcfg.sh` and `brcmflash.sh` scripts call the native versions of `brcmflash` and `brcmlinlpcfg`. For example, on Linux x86_64, the 64-bit utilities are called.

The adapter being managed by the utilities included in the Brcmflash Standalone kit must not be managed simultaneously by other Broadcom utilities, including OneCommand CNA Manager.

5. To update firmware/boot code on an adapter, type:

```
# ./brcmflash.sh /auto up
```

To display a list of adapters, type:

```
# ./brcmlinlpcfg.sh  listhba
```

Each time `brcmflash` or `lpcfg` is run, a log file is created. On Linux these files are called:

```
/var/log/clu/brcmflash.log
/var/log/clu/brcmlinlpcfg.log
```

4.3 Supported Commands

Table 1 lists all of the LpCfg commands that are supported on Broadcom Emulex adapters on various platforms.

- ✓ indicates commands that are supported on both `brcmwinlpcfg` and `brcmlinlpcfg`.
- *FCoE only* indicates commands that are supported only by FCoE adapters.

Table 1 Supported Commands for LpCfg

COMMANDS	FCoE, NIC, and iSCSI Adapters	
	x86	x64
Operating Systems	RHEL 6.5+	RHEL 6.5+ RHEL 7.0+
	SLES 11.3+	SLES 11.3+ SLES 12
	WinPE 3.1+	WinPE 3.1+
	Debian 8.x[1]	Debian 8.x[a]
	Ubuntu 14.x[2] Ubuntu 16.x[b]	Ubuntu 14.x[b] Ubuntu 16.x[b]
config	FCoE only	FCoE only
changemac[c]	RHEL 7.3	RHEL 7.3
disablebootdevice	FCoE	FCoE
dmatest	✓	✓
download	✓	✓
enablebootdevice	FCoE	FCoE
extloopback	✓	✓

Table 1 Supported Commands for LpCfg (Continued)

COMMANDS	FCoE, NIC, and iSCSI Adapters	
	x86	x64
factorydefaults Supported on OCe14000-series adapters only.	✓	✓
hbaattr	✓	✓
h or ? (Help)	✓	✓
intloopback	✓	✓
listhba	✓	✓
listrev	FCoE only	FCoE only
listmac	NIC, iSCSI	NIC, iSCSI
listwwn	FCoE only	FCoE only
logfile	✓	✓
networkboot	✓	✓
pciloopback	FCoE only	FCoE only
personalityinfo Supported on OCe11000-series adapters only.	✓	✓
changepersonality Supported on OCe11000-series adapters only.	✓	✓
readaltboot	FCoE only	FCoE only
readbootdevice	FCoE only	FCoE only
readconfig	FCoE only	FCoE only
readmac	✓	✓
readvlanprops	✓	✓
reset	FCoE only	FCoE only
restoredefsvmac[3]	✓	✓
restoredefwwn	FCoE only	FCoE only
restorenwwn	FCoE only	FCoE only
restorevlanprops	✓	✓
restorewwn	FCoE only	FCoE only
savewwn	FCoE only	FCoE only
screendisplay	FCoE only	FCoE only
scriptvwwnn	FCoE only	FCoE only
scriptvwwpn	FCoE only	FCoE only
scriptwwnn	FCoE only	FCoE only
scriptwwpn	FCoE only	FCoE only
setaltboot	FCoE only	FCoE only
setbootdevice	FCoE only	FCoE only

Table 1 Supported Commands for LpCfg (Continued)

COMMANDS	FCoE, NIC, and iSCSI Adapters	
	x86	x64
version	✓	✓
vpd	✓	✓
writesvmac[c]	✓	✓
writevlanprops	✓	✓
writewwn	FCoE only	FCoE only

1. Debian operating systems support NIC only.
2. Ubuntu operating systems support NIC only.
3. This command is supported only on certain OEM-specific adapters.

4.3.1 Help Command and General Adapter Management Commands

Use the following commands to view the online help, reset the adapter, or reset the adapter to factory defaults.

4.3.1.1 Viewing the Syntax for Commands (help or ?)

To view the syntax for all available commands, type:

```
brcmoslpcfg help
```

or

```
brcmoslpcfg ?
```

To view the syntax for a specific command, type:

```
brcmoslpcfg help <command>
```

or

```
brcmoslpcfg ? <command>
```

For example, either:

```
brcmoslpcfg help download
```

or

```
brcmoslpcfg ? download
```

returns a response similar to the following:

- For Windows

  ```
  download <n=adapter> <i=path\image_filename>
  or
  download <a=adaptertype> <i=path\image_filename>
  ```

- For Linux, VMware, and FreeBSD:

  ```
  download <n=adapter> <i=path/image_filename>
  or
  download <a=adaptertype> <i=path/image_filename>
  ```

4.3.1.2 Resetting an Adapter (reset)

This command resets a specific adapter or all adapters in the system.

ATTENTION Executing a reset on an adapter that is being used to boot from SAN is not recommended. The reset might cause a loss of connectivity to the SAN and possible loss of data. To reset an adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter you want to reset to a non-boot from SAN host and reset it from there.
- If the host with the reset target adapter is also hosting other boot from SAN adapters, carry out a boot from SAN using one of the other boot from SAN adapters. The intended target adapter can now be reset, because it is not being actively used for boot from SAN.

To reset one adapter, type:

```
brcmoslpcfg reset n=<adapter number>
```

To reset all adapters in the system, type

```
brcmoslpcfg reset n=all
```

4.3.1.3 Resetting an Adapter to the Factory Defaults (factorydefaults)

This command returns the adapter to the factory default profile and configuration. A reboot is required for the changes to take effect.

NOTE The factory defaults command is supported on OCe14000-series adapters only.

After a successful factory reset on an OCe14000-series adapter, the adapter is offline. Perform an immediate reboot to complete the reset and return the adapter to full functionality.

To return an adapter to its default settings, type

```
brcmoslpcfg factorydefaults n=<adapter#|all>
```

The following example returns adapter number 1 to its default settings.

```
brcmoslpcfg factorydefaults 1
```

4.3.2 Viewing Adapter Information

Use the following commands to view different kinds of adapter information, such as VPD, boot device information, and adapter attributes.

4.3.2.1 Viewing Broadcom Emulex Conventional Names Instead of VPD (-/c)

Adding /c to any command that uses the a= parameter to return adapter data causes the command to return an Emulex conventional model for the adapter rather than the model name of the adapter. /c can be used for any command that uses the model name. For instance, the following commands use the a= parameter and show this behavior.

- listhba
- config
- download

NOTE The offline utility does not always display the model name of the adapter being tested.

To list adapter information using the conventional name rather than the VPD, type

```
brcmoslpcfg listhba -/c
```

The output from this command includes the Broadcom Emulex conventional model name – OCe14401-UX, rather than the VPD model name – Emulex OneConnect OCe14401-UX 40GbE 1-Port QSFP+PCIe 3.0 Universal CNA, NIC PF.

4.3.2.2 Viewing the LpCfg Utility Version Information (version)

This command shows the LpCfg utility version information.

To view this information, type:

```
brcmoslpcfg version
```

4.3.2.3 Viewing VPD

This command shows the VPD of the adapter specified by its number.

To display VPD, type:

```
brcmoslpcfg vpd n=<adapter number>
```

4.3.2.4 Viewing Boot Device Information (readbootdevice)

This command shows the WWN, the LUN (in decimal format), and the topology in use for the currently selected boot device.

To show this information, type:

```
brcmoslpcfg readbootdevice n=<adapter number>
```

The following example reads WWN and LUN for adapter number 1:

```
brcmoslpcfg readbootdevice n=1
```

4.3.2.5 Viewing Adapter Attributes (hbaattr)

This command displays adapter information.

To list the adapter attributes for all installed adapters, type:

```
brcmoslpcfg hbaattr
```

To list the adapter attributes for one adapter, type

```
brcmoslpcfg hbaattr n=<adapter_number>
```

4.3.2.6 Viewing All Adapters in the System (listhba)

This command lists all installed adapters in the system. Information includes the adapter number (base 1), the IEEE address assigned by the manufacturer, the functional firmware, the adapter type, and possible mailbox errors.

To list all adapters in the system, type:

```
brcmoslpcfg listhba
```

NOTE `listHBA` with option `-/c` displays the conventional model names instead of the model names contained in the VPD.

4.3.2.7 Viewing the WWN of All Adapters in the System (listwwn)

This command lists all adapters installed in the system and shows the factory-assigned WWN, the non-volatile WWPN, and the WWNN used to identify an adapter in the SAN.

The factory-assigned WWN is an IEEE address that cannot be changed in the field. The non-volatile WWN can be modified in the field and persists after a restart of the operating system. The full factory-assigned WWN and non-volatile WWN are a concatenation of the two 8-character values (word 0 and word 1) that are shown for each. You can modify the non-volatile WWPN and WWNN using either the `writewwn` command or the `scriptwwpn` and `scriptwwnn` commands. For more information on the `writewwn` command, see Section 4.3.4.1, Writing WWN and Updating NVPARMS (writewwn).

If the system does not have any Emulex adapters installed, it returns error code 45.

To show the WWN information, type:

brcmos`lpcfg listwwn`

4.3.2.8 Viewing the MAC Address (listmac)

This command shows the MAC address of a NIC CNA port.

To view the MAC address of a NIC port, type:

brcmos`lpcfg listmac n=<`*adapter number*`>`

4.3.2.9 Reading the MAC Address (readmac)

This command shows the current, factory and semi-volatile MAC addresses for the adapter number specified.

To view the MAC address of the adapter number specified, type:

brcmos`lpcfg.exe readmac n=<`*adapter number*`>`

The following example lists information for adapter number 1:

```
>brcmoslpcfg.exe readmac n=1
Command: readmac n=1
adapter 1:
Current MAC Address : 00-90-FA-30-43-AA
Factory MAC Address : 00-90-FA-30-43-30
Semi-Volatile MAC Address: 00-90-FA-30-43-AA
Command completed, NO Error
```

4.3.2.10 Viewing Firmware Program Revisions (listrev)

This command shows the firmware versions in the adapter's flash memory, specified by their numbers.

To show revisions, type:

brcmos`lpcfg listrev n=<`*adapter number*`>`

The following example lists information for adapter number 3:

brcmos`lpcfg listrev n=3`

4.3.2.11 Viewing Selected Configuration Regions (readconfig)

This command shows the contents of the selected configuration region up to the initialized length or the specified byte count (if the initialized length is less than the specified byte count). Valid region numbers are 0 to 32. You must initialize the configuration region first by writing data to it.

NOTE The `readconfig` command only supports reading configuration regions 0, 8, and 32 on OneConnect FCoE adapters.

To read a configuration, type:

brcmos`lpcfg readconfig n=<`*adapter number*`> r=<`*region number*`> l=<`*byte count*`>`

The following example reads the configuration for adapter number 1, region 0, byte count 20:

brcmos`lpcfg readconfig n=1 r=0 l=20`

4.3.3 Firmware and Boot Code Download Commands

The following firmware/boot code download command downloads a firmware/boot code file.

4.3.3.1 Downloading a File (download)

This command downloads a firmware/boot code file to a specific adapter.

The adapter name is the name that appears when you run the `listhba` command. For more information on the `listhba` command, see Section 4.3.2.6, Viewing All Adapters in the System (listhba).

To download a firmware image file to an adapter specified by its number, type

brcmos`lpcfg download n=<`*adapter number*`> i=<`*firmware image filename*`>`

The following example downloads the `oc14-11.2.123.45.ufi` firmware file to adapter number 6; in this example, the firmware file is for an OCe14000-series adapter:

brcmos`lpcfg download n=6 i=oc14-11.2.123.45.ufi`

NOTE Boot code on OneConnect adapters is updated by downloading firmware. LpCfg does not support updating boot code separately for these adapters.

4.3.4 World Wide Name Commands

The following commands allow you to use the WWN to update NVPARAMS, save WWN data to a file, and restore WWN data while updating NVPARAMS. You can also restore the NVPARAMS and the IEEE address.

4.3.4.1 Writing WWN and Updating NVPARMS (writewwn)

This command allows you to enter word 0 and word 1 of the WWPN or WWNN from the keyboard or from a barcode scanner to update a specified adapter's NVPARMS with a new WWPN or WWNN. The new WWPN and WWNN are used the next time the adapter is discovered. The adapter stores the original WWPN and WWNN in another region of the memory so it can be used to identify the adapter as it was manufactured. The WWN can also be read with a barcode scanner.

The `writewwn` command prompts you for the WWPN and WWNN data words, so it cannot be used in a script file. The `scriptwwnn` and `scriptwwpn` commands use values entered with the command, so they can be used in a script file.

ATTENTION Use the `writewwn` command with caution. If you use the same WWPN or WWNN on more than one adapter in a fabric, unpredictable results might occur.

NOTE Word 0 of WWNN and WWPN names must follow one of the following formats:
1 0 0 0 0 x x x
2 x x x x x x x

```
3xxxxxxx
5xxxxxxx
```

To modify the WWPN and WWNN, type:

brcmos`lpcfg writewwn n=<`*adapter number*`>`

The offline utility prompts you to enter new data:

- WWPN word 0
- WWPN word 1
- WWNN word 0
- WWNN word 1

The following example writes the WWPN and WWNN for adapter number 1.

```
brcmoslpcfg writewwn n=1
Enter or Scan value for WWPN word 0 now
```

Enter a value.

```
10000000
```

The system echoes what you entered, followed by the next prompt:

```
10000000
Enter or Scan value for WWPN word 1 now
```

4.3.4.2 Saving WWN Data to a File (savewwn)

This command reads the original words 0 and 1 of the IEEE address, installed by manufacturing, from configuration regions 16 (or 32) of the adapter (specified by its number), and it saves the configuration region information in the selected WWN file.

To save the WWN data to a file, type:

brcmos`lpcfg savewwn n=<`*adapter number*`> c=<`*wwn filename*`>`

The following example reads the configuration region information on adapter number 4 and saves it to the contents of the `ctwwn.sav` file:

brcmos`lpcfg savewwn n=4 c=ctwwn.sav`

4.3.4.3 Restoring WWN and Updating NVPARMS (restorewwn)

This command restores words 0 and 1 of the IEEE address from a specified file created with the `savewwn` command and uses them to update the NVPARMS port name with this IEEE address.

ATTENTION Running `restorewwn` command on an adapter that is being used to boot from SAN is not recommended. The `restorewwn` command might cause a loss of connectivity to the SAN and possible loss of data. To restore the WWN and Update NVPARMS, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to use the `restorewwn` command to a non-boot from SAN host, and run `restorewwn` command from there.
- If the host with the `restorewwn` command target adapter is also hosting other boot from SAN adapters, carry out a boot from SAN using one of the other boot from SAN adapters. You can now restore the WWN and update the NVPARMS on the intended

target adapter, because it is not being actively used for boot from SAN.

To restore the WWN, type:

`brcmoslpcfg restorewwn n=<adapter number> c=<wwn filename>`

The following example updates the NVPARMS on adapter number 4 with the `ctwwn.sav` file.

`brcmoslpcfg restorewwn n=4 c=ctwwn.sav`

4.3.4.4 Restoring NVPARMS (restorenvwwn)

ATTENTION Running `restorenvwwn` command on an adapter that is being used to boot from SAN is not recommended. The `restorenvwwn` command might cause a loss of connectivity to the SAN and possible loss of data. To restore the NVPARMS, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to use the `restorenvwwn` command to a non-boot from SAN host and run the command from there.
- If the host with the `restorenvwwn` target adapter is also hosting other boot from SAN adapters, carry out a boot from SAN using one of the other boot from SAN adapters. You can now restore the NVPARMS on the intended target adapter, because it is not being actively used for boot from SAN.

This command restores the non-volatile WWPN and WWNN to the adapter, replacing any volatile WWPN and WWNN data, without powering off the adapter.

If the adapter does not have firmware that supports the volatile WWN, the following error message appears:

```
Write Volatile Parms Error. Reported Error 48
```

If this occurs, install firmware that supports the volatile WWN.

To restore the non-volatile WWN, type:

`brcmoslpcfg restorenvwwn n=<adapter number>`

The following example restores the non-volatile WWN on adapter number 2.

`brcmoslpcfg restorenvwwn n=2`

4.3.4.5 Restoring the IEEE Address (restoredefwwn)

This command reads the IEEE address (assigned by the manufacturer) and writes it to the non-volatile WWPN and WWNN.

ATTENTION Running `restoredefwwn` on an adapter that is being used to boot from SAN is not recommended. The `restoredefwwn` command might cause a loss of connectivity to the SAN and possible loss of data. To restore the IEEE address, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to use the `restoredefwwn` command to a non-boot from SAN host and run restoredefwwn from there.
- If the host with the `restoredefwwn` target adapter is also hosting other boot from SAN adapters, carry out a boot from SAN using one of the other boot from SAN adapters. You can now

restore the IEEE address on the intended target adapter, because it is not being actively used for boot from SAN.

To restore the IEEE address, type:

```
brcmoslpcfg restoredefwwn n=<adapter number>
```

The following example restores the IEEE address on adapter number 2.

```
brcmoslpcfg restoredefwwn n=2
```

4.3.5 Boot Code Commands

To set the boot device with the offline utility, do the following:

1. Use the `setbootdevice` command to set the boot device. See Section 4.3.5.1, Selecting a Boot Device (setbootdevice), for more information.
2. Configure the system BIOS so the adapter boot device is the highest in the boot order.

4.3.5.1 Selecting a Boot Device (setbootdevice)

This command sets the boot device specified by its WWN, LUN, and desired topology.

The selected device boots when the system reboots.

ATTENTION Selecting a boot device on an adapter that is being used to boot from SAN is not recommended. The `setbootdevice` command might cause a loss of connectivity to the SAN and possible loss of data. To run the `setbootdevice` command, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to run the `setbootdevice` command to a non-boot from SAN host and run the command from there.
- If the host with the `setbootdevice` target adapter is also hosting other boot from SAN adapters, carry out a boot from SAN using one of the other boot from SAN adapters. The `setbootdevice` command can now be run on the intended target adapter, because it is not being actively used for boot from SAN.

To set the boot device, type (all on one line):

```
brcmoslpcfg setbootdevice n=<adapter number> w0=<wwpn word 0> w1=<wwpn word 1>
l=<Decimal ID of LUN> t=<topology>
```

NOTE Enter the LUN in decimal format.

The following example sets the boot device on adapter number 1, LUN number 46, with a desired topology of arbitrated loop:

```
brcmoslpcfg setbootdevice n=1 w0=a1b2c3d4 w1=b946a4e8 l=46 t=0
```

NOTE If port login fails after 50 ms, the command is retried once.

4.3.5.2 Enabling or Disabling Boot Devices (enablebootdevice or disablebootdevice)

After using the `setbootdevice` command, you can either enable or disable the boot device by using the `enablebootdevice` or `disablebootdevice` command.

ATTENTION Enabling or disabling the boot device on an adapter that is being used to boot from SAN is not recommended. The `enablebootdevice` or `disablebootdevice` command might cause a loss of connectivity to the SAN and possible loss of data. To run either of the commands, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to enable or disable the boot device to a non-boot from SAN host and run the command from there.
- If the host with the target adapter is also hosting other boot from SAN adapters, carry out a boot from SAN using one of the other boot from SAN adapters. Either command can now be run on the intended target adapter, because it is not being actively used for boot from SAN.

To enable the boot device, type:

`brcmoslpcfg enablebootdevice n=<adapter number>`

To disable the boot device, type:

`brcmoslpcfg disablebootdevice n=<adapter number>`

NOTE For the change to take effect, perform a system reboot.

4.3.5.3 Read All Alternative Boot Devices (readaltboot)

This command shows the WWN and LUN numbers in decimal format of all possible alternate boot devices. You can have up to seven alternate boot devices.

To read all alternate boot devices, type (all on one line):

`brcmoslpcfg readaltboot n=1`

4.3.5.4 Selecting One or More Alternate Boot Devices (setaltboot)

This command sets alternate boot devices. You can set up to seven alternate boot devices; that is, *index* `i` can be from one to seven.

To set up one or more alternate boot devices, type (all on one line):

`brcmoslpcfg setaltboot n=<adapter number> i=<index> w0=<wwpn word 0> w1=<wwpn word 1> l=<Decimal ID of LUN>`

The following example sets the alternate boot device on adapter number 1, LUN number 3:

`brcmoslpcfg setaltboot n=1 i=1 w0=12345678 w1=a842b6ed l=3.`

4.3.5.5 Enable or Disable PXE Boot on NIC Devices (networkboot)

This command allows you to enable or disable PXE boot on NIC devices.

To enable or disable PXE boot on a NIC device, type:

`brcmoslpcfg networkboot n=<adapter_number> <t=type> <s=status>`

The following example enables PXE boot on adapter 1:

`brcmoslpcfg networkboot n=1 t=pxe s=1`

The following example disables PXE boot on adapter 1

`brcmoslpcfg networkboot n=1 t=pxe s=0`

4.3.6 Configuration Commands

Configuration commands allow you to update the configuration region of an adapter by name or by number.

4.3.6.1 Updating Configuration Regions (config)

Two forms of configuration are available:

- Configure all adapters of a given adapter name at once
- Configure a single adapter by its number

Valid region numbers range from 0 to 32.

4.3.6.1.1 Update by Name

To update a specified configuration region on all adapters of the same selected name, type (all on one line)

`brcmoslpcfg config a=<adapter name> r=<region number> c=<configuration filename>`

NOTE The adapter name is the name that appears when you run the `listHBA` command. For more information on the `listHBA` command, see Section 4.3.2.6, Viewing All Adapters in the System (listhba).

The following example updates region 6 of all OCe14000 adapters with `ctplus1.cfl`:

`brcmoslpcfg config a=oce14000 r=6 c=ctplus1.cfl`

NOTE The size of the `.cfl` file for configuration region update can be up to 2028 bytes.

4.3.6.1.2 Update by Number

To update a specified configuration region for one adapter, type (all on one line):

`brcmoslpcfg config n=<adapter number> r=<region number> c=<configuration filename>`

The following example updates region 17 of adapter number 4 with `heplus1.cf1`:

`brcmoslpcfg config n=4 r=17 c=heplus1.cfl`

The following example updates region 6 of adapter number 2 with `d:\dfplus1.cf1`:

`brcmoslpcfg config n=2 r=6 c=d:\dfplus1.cfl`

4.3.6.2 Writing a Semi-Volatile MAC Address on a NIC Device (writesvmac)

NOTE This command is supported only on OEM-specific adapters.

To write a semi-volatile MAC address to an adapter, type:

`brcmoslpcfg writesvmac n=<adapter_number> <m=mac_address>`

The following example writes the MAC address `0090FA112233` to adapter 1:

`brcmoslpcfg writesvmac n=1 m=0090FA112233`

4.3.6.3 Writing a Non-Volatile MAC Address on a NIC Device (changemac)

NOTE This command is supported only on OEM-specific adapters.

To write a non-volatile MAC address to an adapter, type:

`brcmoslpcfg changemac n=<adapter_number> <m=mac_address>`

The following example writes the non-volatile MAC address 0090FA112233 to adapter 1:

`brcmoslpcfg changemac n=1 m=0090FA112233`

4.3.6.4 Restoring the Default Semi-Volatile MAC Address on a NIC Device (restoredefsvmac)

NOTE This command is supported only on OEM-specific adapters.

To restore the default semi-volatile MAC address to an adapter, type:

`brcmoslpcfg restoredefsvmac n=<adapter_number>`

The following example restores the default MAC address to adapter 1:

`brcmoslpcfg restoredefsvmac n=1`

4.3.7 Personality Commands

The following commands allow you to manage the personality of the adapter.

NOTE Personality commands are supported on OCe11100-series adapters only.

4.3.7.1 Viewing Personality Information (personalityInfo)

This command displays the current personality running on the adapter and other available personalities.

To view personality information, type:

`brcmoslpcfg personalityInfo n=<adapter index>`

where <*adapter index*> is obtained from the `listhba` command.

The following example lists the `personalityInfo` on adapter number 3:

`brcmoslpcfg personalityInfo n=3`

4.3.7.2 Changing Personality Information (changePersonality)

This command changes the personality of the adapter to the selected personality. You must reboot the system to activate the personality change.

To change personality, type:

`brcmoslpcfg changePersonality n=<adapter index> p=<personality string>`

The following example changes the personality information on adapter number 3:

`brcmoslpcfg changePersonality n=3 p=iSCSI`

4.3.8 Diagnostic Tests

The following commands let you run diagnostic tests on your adapter.

NOTE Because the `extloopback`, `intloopback`, and `pciloopback` commands do not support testing on all installed adapters, the `o=3` option applies to running `brcmlinlpcfg` with the script command.

The o=3 option instructs LpCfg to ignore errors and continue script execution.

When the o=2 option is used, LpCfg ignores up to 3 errors during an adapter test before continuing to the next adapter.

4.3.8.1 Running the External Loopback Test (extloopback)

ATTENTION Performing an `extloopback` test on an adapter that is being used to boot from SAN is not recommended. The `extloopback` command might cause a loss of connectivity to the SAN and possible loss of data. To perform an `extloopback` command, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to perform the external loopback test onto a non-boot from SAN host and perform the test from there.
- If the host with the test target adapter is also hosting other boot from SAN adapters, carry out a boot from SAN using one of the other boot from SAN adapters. The `extloopback` command can now be run on the intended target adapter, because it is not being actively used for boot from SAN.

This command runs the external loopback test. You must put a loopback plug in each adapter port to be tested. You can test a specific adapter in the system. Specify the number of times you want the test to repeat, and direct the test response if an error is found.

The option on error choices are:

- o=1 stops the test on the first error
- o=2 ignores three errors and stops the test on the fourth error
- o=3 ignores errors and continues the test

NOTE The `extloopback` command does not support testing all installed adapters using `n=all`.

To run the external loopback test, type (all on one line):

brcmos`lpcfg extLoopback n=<`*adapter index*`> p=<`*pattern*`> c=<`*byte count*`> r=<`*repeat count*`>`

where:

- *p* = 3 to 8-HEX byte pattern
- *c* = 1500 to 8192
- *r* = 1 to 4096

The following example runs the external loopback test 50 times on adapter number 1 and stops the test if an error occurs:

brcmos`lpcfg extloopback n=1 r=50 o=1`

4.3.8.2 Running the Internal Loopback Test (intloopback)

ATTENTION Performing an internal loopback test on an adapter that is being used to boot from SAN is not recommended. The `intloopback` command might cause a loss of connectivity to the SAN and possible loss of data. To perform an `intloopback` command, make sure the

adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to perform the internal loopback test onto a non-boot from SAN host and perform the test from there.
- If the host with the test target adapter is also hosting other boot from SAN adapters, carry out a boot from SAN using one of the other boot from SAN adapters. The `intloopback` command can now be run on the intended target adapter, because it is not being actively used for boot from SAN.

This command runs the internal loopback test. You can run the test on a specific adapter in the system. Specify the number of times you want the test to repeat, and direct the test response if an error occurs.

The option on error choices are:

- o=1 stops the test on the first error
- o=2 ignores three errors and stops the test on the fourth error
- o=3 ignores errors and continues the test

NOTE The `intloopback` command does not support testing all installed adapters using `n=all`.

Do not use the internal loopback test on Emulex blade adapters.

To run the internal loopback test, type (all on one line):

```
brcmoslpcfg intLoopback n=<adapter index> p=<pattern> c=<byte count> r=<repeat count> t=<type>
```

where:

- *p* = 3 to 8-HEX byte pattern
- *c* = 1500 to 8192
- *r* = 1 to 4096
- *t* = 2

The following example runs the internal loopback test 100 times on adapter number 1 and stops the test if an error occurs:

```
brcmoslpcfg intloopback n=1 r=100 o=1
```

4.3.8.3 Running the PCI Loopback Test (pciloopback)

ATTENTION Performing a PCI loopback test on an adapter that is being used to boot from SAN is not recommended. The `pciloopback` command might cause a loss of connectivity to the SAN and possible loss of data. To perform a `pciloopback` command, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to perform the PCI loopback test onto a non-boot from SAN host and perform the test from there.
- If the host with the test target adapter is also hosting other boot from SAN adapters, carry out a boot from SAN using one of the other boot from SAN adapters. The `pciloopback` command can now be run on the intended target adapter, because it is not being actively used for boot from SAN.

This command runs the PCI loopback test. You can run the test on a specific adapter or on all adapters in the system. Specify the number of times you want the test to repeat, and direct the test response if an error occurs.

The option on error choices are:

- o=1 stops the test on the first error
- o=2 ignores three errors and stops the test on the fourth error
- o=3 ignores errors and continues the test

NOTE The `pciloopback` command does not support testing all installed adapters using `n=all`.

To run the PCI loopback test, type (all on one line):

brcmos`lpcfg pciloopback n=<`*all|adapter number*`> r=<`*repeat count*`> o=<`*option on error*`>`

The following example runs the PCI loopback test 100 times on all adapters in the system and stops the testing if any errors occur.

brcmos`lpcfg pciloopback n=all r=100 o=1`

4.3.8.4 Running the DMATest

This command runs the DMA test, which is a test of the memory on the adapter. You can run the test on a specific adapter in the system and specify the number of times you want the test to repeat.

NOTE `DMATest` does not support testing all installed adapters using `n=all`.
Do not use the DMA test on Emulex blade adapters.

To run the DMA test, type (all on one line):

brcmos`lpcfg dmatest n=<`*adapter*`> p=<`*pattern*`> c=<`*byte-count*`> r=<`*repeat-count*`>`

The following example runs the DMA test two times on adapter number 1:

brcmos`lpcfg dmatest n=1 p=abcde c=1500 r=2`

4.3.9 Using Script Files

Use script files to efficiently perform tasks. Script files are common sequences of commands you use when performing tasks.

You can group commands together and run them using a script file. You can also enter comment lines, which begin with a semicolon. Each line follows the same command syntax as those documented in this manual. Using the offline utility, you can:

- Run commands entered in a script file. Use the @ command to run the script file.
- Run commands multiple times. Add the `repeat` command as the last line of the script file.
- Create a log of test results. Add the `logfile` command as the first line of the script file.

To run a script file type:

brcmos`lpcfg @<`*scriptname.txt*`>`

The following example runs the script `script1.txt`, which resides in the current directory and runs all the commands in that script file.

brcmos`lpcfg @script1.txt`

NOTE To interrupt and stop any script, press <S> on the keyboard.

The following is a sample script file. Each command follows the syntax covered above. The comment lines begin with a semicolon (;):

```
version
screendisplay o=0
;download a=oce14000 i=c:\temp\oc14-11.2.123.45.ufi
;reset n=1 s=0
;reset n=2 s=0
reset n=all s=0
; pciloopback n=1 r=10 o=1
; pciloopback n=2 r=10 o=2
; pciloopback n=all r=50 o=3
intloopback n=1 r=10 o=1
intloopback n=2 r=10 o=1
extloopback n=1 r=40 o=3
extloopback n=2 r=40 o=3
;repeat r=10
```

4.3.9.1 Repeating a Series of Commands (repeat)

Enter this command at the end of a script file to repeat a series of commands from the beginning of the script file a specific number of times.

NOTE To interrupt and stop the repeat command, press <S> on the keyboard.

To repeat the series of commands in the script file, add the following as the last line of the file:

`repeat r=<repeat count>`

The following example repeats the series of commands in the script file ten times.

```
repeat r=10
```

4.3.9.2 Enabling or Disabling Test Messages on the Screen (screendisplay)

This command enables or disables test message displays on the screen.

- `o=0` – Prevents messages from appearing
- `o=1` – Enables messages

NOTE This command is supported only in script files.

To enable or disable test message displays on the screen, add the following line to the script:

`screendisplay o=<display option>`

In the following script file example, messages from the `version` and `listhba` commands are output to the log file and to the screen. After the `screendisplay` command is set to `0`, the result messages for all successive commands (`download` and `enablebootdevice`) are output only to the log file, not to the screen.

```
version
listhba
screendisplay o=0
download n=2 i=oc14-11.2.123.45.ufi
enablebootdevice n=2
```

4.3.9.3 Updating Non-volatile WWNN (scriptwwnn)

ATTENTION Running `scriptwwnn` on an adapter that is being used to boot from SAN is not recommended. The `scriptwwnn` command might cause a loss of connectivity to the SAN and possible loss of data. To update the non-volatile WWNN on an FCoE adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to use the `scriptwwnn` command to a non-boot from SAN host and run the script from there.
- If the host with the target adapter is also hosting other boot from SAN adapters, carry out a boot from SAN using one of the other boot from SAN adapters. You can now run the script to update the non-volatile WWNN on the intended target adapter, because it is not being actively used for boot from SAN.

This command reads the WWNN words 0 and 1 from the command line to update the non-volatile WWNN. You can also include this command in a script file. When the adapter is discovered, the new WWNN value is used. The adapter retains the original WWNN in another region of the firmware.

ATTENTION Use the `scriptwwnn` command with caution. If you use the same WWNN on more than one adapter in a fabric, unpredictable results might occur.

To change WWNN words 0 and 1 from the command line, type (all in one line):

brcmos`lpcfg scriptwwnn n=<`*adapter number*`> w0=<`*wwnn word 0*`> w1=<`*wwnn word 1*`>`

The following example updates non-volatile WWNN word 0 and word 1 for adapter number 1:

brcmos`lpcfg scriptwwnn n=1 w0=10000345 w1=B620A1B2`

NOTE Word 0 of WWNN and WWPN names must follow one of the following formats:
1 0 0 0 0 x x x
2 x x x x x x x
3 x x x x x x x
5 x x x x x x x
If the `scriptvwwnn` command has been used previously, the adapter continues to use that WWNN until you change the WWNN with the `restorenvwwn` command.

4.3.9.4 Updating Non-volatile WWPN (scriptwwpn)

ATTENTION Running `scriptwwpn` on an adapter that is being used to boot from SAN is not recommended. The `scriptwwpn` command might cause a loss of connectivity to the SAN and possible loss of data. To update the non-volatile WWPN on an FCoE adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to use the `scriptwwpn` command to a non-boot from SAN host and run the script from there.
- If the host with the target adapter is also hosting other boot from SAN adapters, carry out a boot from SAN using one of the other

boot from SAN adapters. You can now run the script to update the non-volatile WWPN on the intended target adapter, because it is not being actively used for boot from SAN.

This command reads WWPN words 0 and 1 from the command line to update the non-volatile WWPN. You can also include this command in a script file. When it is discovered, the adapter uses the new WWNN value. It does not use the original IEEE address assigned by manufacturing (located in Configuration Region 16 or 32).

ATTENTION Use the `scriptwwpn` command with caution. If you use the same WWPN on more than one adapter in a fabric, unpredictable results might occur.

NOTE If the `scriptvwwnn` command has been used previously, the adapter continues to use that WWPN until you change the WWPN with the `restorenvwwn` command.

To change WWPN words 0 and 1 from the command line, type (all in one line):

brcmos`lpcfg scriptwwpn n=<`*adapter number*`> w0=<`*wwpn word 0*`> w1=<`*wwpn word 1*`>`

The following example updates the non-volatile WWPN word 0 word 1 for adapter number 1:

brcmos`lpcfg scriptwwpn n=1 w0=20A2D6B8 w1=C920A1B2`

4.3.9.5 Updating Volatile WWNN (scriptvwwnn)

ATTENTION Running `scriptvwwnn` on an adapter that is being used to boot from SAN is not recommended. The `scriptvwwnn` command might cause a loss of connectivity to the SAN and possible loss of data. To update the volatile WWNN on an FCoE adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to use the `scriptvwwnn` command to a non-boot from SAN host and run the script from there.
- If the host with the target adapter is also hosting other boot from SAN adapters, carry out a boot from SAN using one of the other boot from SAN adapters. You can now run the script to update the volatile WWNN on the intended target adapter, because it is not being actively used for boot from SAN.

This command reads the WWNN words 0 and 1 from the command line to update the volatile WWNN. The next time the adapter is discovered, it uses this new WWNN. It does not use the original IEEE address assigned by manufacturing (located in Configuration Region 16 or 32), nor does it use the value entered by the `writewwn` or `scriptwwnn` commands.

If the adapter does not have firmware that supports the volatile WWN, the following error message appears:

```
Write Volatile Parms Error. Reported Error 48
```

If this occurs, install firmware that supports the volatile WWN.

NOTE Word 0 of WWNN and WWPN names must follow one of the following formats:
1 0 0 0 0 x x x
2 x x x x x x x
3 x x x x x x x
5 x x x x x x x

After you issue this command, the volatile WWNN is used by the adapter until the `restorenvwwn` command is issued or the system is restarted.

ATTENTION Use the `scriptvwwnn` command with caution. If you use the same volatile WWNN on more than one adapter in a fabric, unpredictable results might occur.

To change volatile WWNN words 0 and 1 from the command line, type (all in one line):

brcmos`lpcfg scriptvwwnn n=<`*adapter number*`> w0=<`*wwnn word 0*`> w1=<`*wwnn word 1*`>`

The following example updates the volatile WWNN word 0 word 1 for adapter number 1:

brcmos`lpcfg scriptvwwnn n=1 w0=20A2D6B8 w1=C920A1B2`

4.3.9.6 Updating Volatile WWPN (scriptvwwpn)

ATTENTION Running `scriptvwwpn` on an adapter that is being used to boot from SAN is not recommended. The `scriptvwwpn` command might cause a loss of connectivity to the SAN and possible loss of data. To update the volatile WWPN on an FCoE adapter, make sure the adapter is not currently being used to boot from SAN. There are two ways to do this:

- Move the adapter on which you want to use the `scriptvwwpn` command to a non-boot from SAN host and run the script from there.
- If the host with the target adapter is also hosting other boot from SAN adapters, carry out a boot from SAN using one of the other boot from SAN adapters. You can now run the script to update the volatile WWPN on the intended target adapter, because it is not being actively used for boot from SAN.

This command reads the WWPN words 0 and 1 from the command line to update the volatile WWPN. The next time the adapter is discovered, it uses this new WWPN. It does not use the original IEEE address assigned by manufacturing (located in Configuration Region 16 or 32), nor does it use the value entered by the `writewwn` or `scriptwwpn` commands.

If the adapter does not have firmware that supports the volatile WWN, the following error message appears:

```
Write Volatile Parms Error. Reported Error 48
```

If this occurs, install firmware that supports the volatile WWN.

NOTE Word 0 of WWNN and WWPN names must follow one of the following formats:
1 0 0 0 0 x x x
2 x x x x x x x
3 x x x x x x x
5 x x x x x x x
After you issue this command, the volatile WWNN is used by the adapter until the `restorenvwwn` command is issued or the system is restarted.

ATTENTION Use the `scriptvwwnn` command with caution. If you use the same VWWNN on more than one adapter in a fabric, unpredictable results might occur.

To change VWWPN words 0 and 1 from the command line, type (all in one line):

`brcmoslpcfg scriptvwwpn n=<adapter number> w0=<wwpn word 0> w1=<wwpn word 1>`

The following example updates the volatile WWPN word 0 word 1 for adapter number 1:

```
brcmoslpcfg scriptvwwpn n=1 w0=20A2D6B8 w1=C920A1B2
```

4.3.9.7 Creating a Log (logfile)

This command creates a log file with a specified directory and file name. The default log file is `lpcfglog.txt` and is created in the system's current directory.

NOTE This command is supported only in script files. Make it the first command in the script.

To create a log file, type the following command in the script file:

`logfile l=<filename>`

The following example creates a file called `lplog.txt` in the `d:\` directory:

```
logfile l=d:\log\lplog.txt
```

Results of all commands are recorded in a log file. Unless otherwise specified by the `logfile` command, the default log file is `lpcfglog.txt` in the current directory.

4.3.10 VLAN Commands

Using the following commands, you can manage the VLAN ID and VLAN priorities on specified adapters.

4.3.10.1 Reading VLAN Priorities (readvlanprops)

This command shows the NIC function on an adapter's VLAN ID and VLAN Priority for the adapter number specified.

To read the VLAN priorities for the adapter specified, type:

`>brcmoslpcfg.exe readmac n=<adapter number>`

The following example lists information for adapter number 1:

```
>brcmoslpcfg.exe readvlanprops n=1
Command: readvlanprops n=1
adapter 1:
VLANId          : 1024
VLANPriority: 1
Command completed, NO Error
```

4.3.10.2 Restoring the Default Values for the VLAN ID and VLAN Priorities (restorevlanprops)

This command restores the adapter's default VLAN ID and VLAN Priority values, which are both set to zero, on the adapter number specified.

To restore the VLAN ID and VLAN priorities for the adapter specified, type:

`>brcmoslpcfg.exe restorevlanprops n=<adapter number>`

4.3.10.3 Changing the VLAN ID and VLAN Priority on a Specified Adapter's NIC Function (writevlanprops)

This command lets you change the VLAN ID and VLAN Priority on the specified adapter's NIC function. You can specify the VLAN ID and the VLAN Priority with the following values:

- `vlan_id` – 0 to 4095. A value of 0 disables the VLAN ID.
- `vlan_priority` – 0 to 7.

To change the VLAN ID and the VLAN Priority on a specified adapter's NIC function, type:

`brcmwinlpcfg.exe writevlanprops n=<adapter number> i=<vlan_id> p=<vlan_priority>`

The following example changes the VLAN ID and VLAN Priority for adapter number 1:

```
>brcmoslpcfg.exe writevlanprops n=1 i=1024 p=1
Command: writevlanprops n=1 i=1024 p=1
adapter 1:
Command completed, NO Error
```

Chapter 5: Status Messages and Error Codes

5.1 Brcmflash Utility Status Messages

The Brcmflash utility takes a less stringent approach to reporting status messages when used with auto-discovery. An error message is returned only when the Brcmflash utility reports a failure during a download or if the firmware and boot directories cannot be located.

Table 2 lists the Brcmflash utility status messages that are supported.

Table 2 Brcmflash Status Messages

Message Number	Message Title	Message Details
0	GOOD_ALL_UPGRADES_OK	Returned if any of the discovered adapters had a successful download performed.
1	GOOD_NO_UPDATES_NEEDED	Returned if any of the discovered adapters did not need an update, downgrade, or rewrite operation. This message is never returned during a force firmware or force boot code download operation.
2	ERROR_ALL_UPGRADES_FAILED	Returned if all attempted downloads failed.
3	ERROR_SOME_UPGRADES_FAILED	Returned if some of the attempted downloads failed.
4	ERROR_EMULEX_APPS_COMMAND	Returned if Brcmflash cannot run an external executable such as: ■ *brcmos*lpcfg (brcmflashOffline)
5	ERROR_NO_SUPORTED_HBA_FOUND	Returned if no supported adapters are found.
6	ERROR_DIRECTORY_NOT_FOUND	Returned if the `firmware` directory is missing. Depending on the operation, auto-discovery expects the `firmware` directory to exist as a subdirectory in the root of the package directory.
7	GOOD_NO_UPGRADES_AVAILABLE	Returned if none of the discovered adapters had a matching image in the firmware directory. In this case only, a matching image is an image that the adapter accepts regardless of the download operation being performed.
8	ERROR_NOT_ADMIN_USER	Returned when you attempt to run the executable file and do not have administrator (Windows) or root (Linux) privileges.
9	ERROR_UNSUPPORTED_OS	Used by the Exlflash Standalone execution scripts. Returned when the scripts are run on an unsupported operating system.
10	ERROR_UNSUPPORTED_ARCH	Used by the Exlflash Standalone execution scripts. Returned when the scripts are run on an unsupported architecture.
19	ERROR_IMAGE_VERSION	An error occurred when decoding the image version. The image file has an invalid extension, or Brcmflash is unable to decode the image file's version. **NOTE** This is returned only when using the /iv=<image_version> switch.
20	ERROR_MISSING_DEP	Returned when a missing dependency is detected.

5.2 LpCfg Error Codes

Table 3 lists the LpCfg error codes that are supported.

Table 3 LpCfg Error Codes

Error Code	Description
0	No error
1	Invalid adapter number
2	Mailbox command error
4	Open file error
5	Invalid configuration region
6	Invalid adapter name
7	Download error
9	Link not up for external loopback test
10	Link not up for internal loopback test
11	Invalid jumper selection (in jumper command)
12	Invalid alternate configuration region (in jumper command)
13	PCI loopback test fails
14	Adapter reset error
15	Read configuration region error
16	No VPD information available
17	No command in command line
18	Open log file error
19	Read wakeup parameters error
20	Update wakeup parameters error
21	Incorrect test parameters
22	Stopped by user
23	Internal loopback test fails
24	External loopback test fails
25	Error exists after four retries
26	Invalid command
27	Incorrect syntax
28	Command supported only in script files
29	Read_rev error
30	Dump configuration region error
31	Read file error
32	Short file error
33	Read NVPARMS error
34	Write NVPARMS error
35	Command does not support all adapters
36	Invalid LUN
38	Update configuration region error
40	Dump memory error

Table 3 LpCfg Error Codes (Continued)

Error Code	Description
41	Update EROM error
42	Delete load entry error
43	Write WWN error
44	Not supported in script files
45	No Broadcom Emulex adapter found
46	Invalid Alternate Boot Device Index
47	Cannot restart adapter
48	Write volatile parameters error
50	Incorrect symbols
51	Invalid length
52	Invalid topology
53	No event log
54	Read event log
55	Invalid input value
56	No libdfc library
57	Non-numeric input
58	No valid WWN
59	Region cleanup
60	Region initialization
62	Unable to allocate memory
63	DFC_InitDiagEnv error
64	DFC_ReadPciCfg error
65	No driver installed
66	No valid driver
67	Not valid adapter type
68	Not valid image
69	Long file error
70	Incompatible image
71	Not supported
72	MILI service not started
73	Script not supported
74	MILI not started
75	No NIC adapter
76	Personality information
77	Personality change
78	Administrator rights error
79	SLI-4 management error
80	Reboot required
83	Application is missing one or more dependencies.
87	The system must be rebooted to activate the new firmware image and the image must be downloaded again.

Table 3 LpCfg Error Codes (Continued)

Error Code	Description
88	Duplicate MAC address
89	Download failed due to missing digital signature. Please contact customer support for additional help. Download failed on adapter <adapter number> Stat 89.
90	Download failed due to invalid firmware digital signature. Please contact customer support for additional help. Download failed on adapter <adapter number> Stat 90.
141	General error from MILI
200	General error

Chapter 6: Troubleshooting

When using Brcmflash, some commands take precedence over others. For instance, the force firmware (`/ff`) and force firmware/boot code (`/f`) switches take precedence over the `/downgrade`, `/update`, and `/rewrite` switches. If `/ff` or `/f` are used, the `/downgrade`, `/update`, and `/rewrite` switches are ignored.

The following examples illustrate this behavior:

1. `./brcmflash /ff /update`
   a. `/update` is ignored.
   b. Performs a force firmware operation on all installed and supported adapters.
2. `./brcmflash /ff /downgrade`
   a. `/downgrade` is ignored
   b. Performs a force firmware operation on all installed and supported adapters.
3. `./brcmflash /ff /rewrite`
   a. `/rewrite` is ignored
   b. Performs a force firmware operation on all installed and supported adapters.

6.1 Unsupported Driver

The Brcmflash utility reports an error similar to the following if an unsupported driver is installed on the system:

```
brcmflash: no supported Emulex HBA's found - Return Code=1
```

If an error similar to the above occurs, verify that the correct version of the driver is installed. For supported adapters and supported versions of operating systems and platforms, go to http://www.broadcom.com.


BROADCOM

Emulex® Drivers for Windows for OneConnect® Adapters

Quick Installation Guide

Version 11.2
December 30, 2016

DRVWINOCA-IG112

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

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

Copyright © 2003–2017 Broadcom. All Rights Reserved.

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

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.

Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Chapter 1: Introduction

This document explains how to install the Emulex® drivers for Windows on your system.

This product supports the following Emulex OneConnect® converged network adapters (CNAs).

- OCe11000-series adapters
- OCe14000-series adapters

Chapter 2: Driver Installation Options

Install the Windows drivers in one of two ways:

- OneInstall Installer – Contains all the applicable Emulex drivers for Windows (Emulex Storport Miniport and NDIS Miniport drivers) and the OneCommand® CNA Manager Application for OneConnect® Adapters in a single download package. The OneInstall Installer can be run in Interactive Mode or Silent Mode.
- Driver kits and AutoPilot Installer – Provides installation options ranging from simple installations with a few mouse clicks to unattended installations that use predefined script files and text-only installations.

NOTE Microsoft patch KB2846340 must be installed on your system to successfully install the NIC driver. If the patch is not installed on your system, the installation stops and prompts you to install it. This patch, from Microsoft's Knowledge Base, is required for Windows Server 2008 SP2, Windows Server 2008 R2, Windows Server 2012, and for Windows 7 and Windows 8 client operating systems.

NOTE If you are installing the NIC driver kit as an update to the Windows Server 2012 driver, Windows Server 2012 driver R2, or Windows Server 2016 driver, some parameter defaults are different from the inbox driver. After you install the Emulex out-of-box driver, select **reset to default** on the Advanced tab of the Device Manager property page. This action returns all adapter and driver settings to the default values listed in this guide.

NOTE Low performance might occur if the Emulex NIC driver is installed on a system meeting the following conditions before installing Microsoft KB2846837:

- A Windows 8, Windows 8.1, or Windows Server 2012 computer with multi-core processors is in use
- Three or more Ethernet ports are installed on the computer
- RSS is enabled and sets the RSS profile to use the `Closest` parameter for the Ethernet adapters

If any or all of these conditions exist, install KB2846837 before installing the Emulex NIC driver.

NOTE Windows 8 x64, Windows 8.1 x64, and Windows 10 drivers are Emulex signed. You must accept the Emulex certificate to install these kits. Support is provided by Broadcom, but not by Microsoft.

NOTE Check the Broadcom website for required updates to the Windows operating system or the Emulex drivers.

2.1 OneInstall Installer

The OneInstall Installer can be run in Interactive mode or Silent mode.

NOTE The OneInstall Installer does not allow you to perform preinstallation tasks or text-only installations. For these tasks, use the driver kits.

NOTE The OneInstall Installer is not supported on Window Server 2016 Nano Server.

The OneInstall package is a self-extracting executable file that installs the following software on your system:

- All compatible protocol drivers:
  - Fibre Channel over Ethernet (FCoE)
  - Internet Small Computer System Interface (iSCSI)
  - Network Interface Card or Controller plus RDMA over Converged Ethernet (NIC + RoCE)

    NOTE RoCE is available as Technical Preview only. Do not use RoCE in a production environment.
- ElxPlus driver (supports the OneCommand CNA Manager Application, persistent binding, LUN mapping and masking, and virtual port creation)
- OneCommand CNA Manager Application (management utility for Emulex adapters)

  NOTE The Enterprise kit for the OneCommand CNA Manager Application does not operate locally on Windows Server Core. You must install the OneCommand CNA Manager Application Core Kit (command-line interface only) to the Windows Server Core.

  NOTE The One Command CNA Manager Enterprise and Core kit do not work on Windows Server 2016 Nano Server. You must install the OneCommand CNA Manager application for Windows Server 2016 Nano Server. See the *Emulex OneCommand CNA Manager application for OneConnect Adapters User Guide* for installation instructions.

2.1.1 Loading the OneInstall Package Using Interactive Mode

To install the drivers using Interactive mode, perform these steps:

1. Download the OneInstall package from the Broadcom website.
2. Navigate to the OneInstall package in Windows Explorer.
3. Double-click the OneInstall package.

   The **Welcome** screen displays.
4. Click **Next**.

   The **Installation Options** screen displays.
5. Select the drivers and applications that you want to install and click **Next**.

   A progress screen displays while the OneInstall installer loads the selected drivers and applications. After the drivers and application software are loaded, an **Installation completed** screen displays.
6. Click **Finish**.

2.1.2 Loading the OneInstall Package Using Silent Mode

Silent mode installation must be run from a batch file or from the command line.

If you run OneInstall from a batch file or from a command line prompt, the default Windows behavior starts OneInstall, then immediately continues with the next command. It does not wait until OneInstall has finished.

As a result, the value of %ERRORLEVEL% will always be 0, because Windows successfully started OneInstall. It does *not* reflect an accurate OneInstall exit code.

To remedy this, run setup as follows:

```
START /wait brcmOneInstall-Setup-<version>.exe
echo %ERRORLEVEL%
```

The "start /wait" ensures that the command does not return until setup has exited. The value of %ERRORLEVEL% now accurately reflects the OneInstall exit code.

2.1.2.1 Command Format

The format of the command is:

`brcmOneInstall-Setup-<version>.exe <install-mode> <options>`

Where:

<*version*> is the version number of the OneInstall package

<*install-mode*> is one of the following:

- `/q0` – (Interactive, non-silent install) This is the default.
- `/q1` – (non-Interactive install) This option displays status pages.
- `/q2` – (Silent install) This option is completely silent, no pages are displayed.
- `/q` – This is the same as `/q1`.

<*options*> specifies the kit, or kits, to install:

`ALL=1` – Install all drivers and the OneCommand CNA Manager Application (Default).

NOTE On a CORE system, this installs all drivers and the OneCommand CNA Manager Application Core Kit.

`ALLCORE=1` – Install all drivers and the OneCommand CNA Manager Application Core Kit.

`DRIVERS=1` – Install all drivers.

`FCOE=1` – Install the FCoE driver only.

`NIC=1` – Install the NIC driver only.

`ISCSI=1` – Install the iSCSI driver only.

`OCM=1` – Install the OneCommand CNA Manager Application for OneConnect Adapters Enterprise Kit only.

NOTE On a CORE system, this installs the OneCommand CNA Manager Application Core Kit.

`OMCORE=1` – Install the OneCommand CNA Manager Application Core Kit only.

To install the drivers using Silent mode, perform these steps:

1. Download the OneInstall package from the Broadcom website.
2. Open a DOS window.
3. Change directory to the folder containing your OneInstall package.

The following are examples of Silent mode commands:

```
Start /wait brcmOneInstall-Setup-<version>.exe  /q2  ALL=1
Start /wait brcmOneInstall-Setup-<version>.exe  /q2  DRIVERS=1
Start /wait brcmOneInstall-Setup-<version>.exe  /q2  FCOE=1  NIC=1  OCM=1
Start /wait brcmOneInstall-Setup-<version>.exe  /q2
Start /wait brcmOneInstall-Setup-<version>.exe  /q2  ALLCORE=1
Start /wait brcmOneInstall-Setup-<version>.exe  /q2  OCMCORE=1
```

2.2 Driver Kit Installer

Each driver kit contains and loads all the Windows drivers for a specific protocol, and includes ElxPlus.

- FCoE driver package (`brcmdrv-fcoe-<version>.exe`)
- iSCSI driver package (`brcmdrv-iscsi-<version>.exe`)
- NIC + RoCE driver package (`brcmdrv-nic-<version>.exe`)

NOTE Updating the NIC protocol driver can temporarily disrupt operation of any NIC teams configured on the system.

2.2.1 Loading the Driver Kit

The driver kit copies the selected Emulex drivers and applications onto your computer.

NOTE This procedure does not install drivers, and no driver changes are made until you run the AutoPilot Installer.

To load the driver kit, perform these steps:

1. Download the driver kit from the Broadcom website to your system.
2. Double-click the driver kit to run it.

   The Emulex Kit Welcome page displays.
3. Click **Next**.

   The **Installation Options** window displays.
4. Select one or both of the following options:
   - **Perform Installation of Software** – Copies the driver kit for your operating system to your computer.
   - **Unpack All Drivers** – Extracts all drivers to the current user's Documents folder. Select this option to perform boot from SAN installations.

   The **Operation in progress** window shows the kit file loading progress. After the kit files are loaded, the **Installation completed** window displays.
5. To continue with the installation, ensure that Start AutoPilot Installer is checked. Click **Next**.

2.2.1.1 Using the Driver Kit on Windows Server 2016 Nano Server

Refer to the *Emulex Drivers for Windows for OneConnect Adapters User Guide* for information about using the driver kit on Windows Server 2016 Nano Server.

2.3 AutoPilot Installer

AutoPilot Installer runs after the driver kit is loaded and the OneCommand CNA Manager Application is installed. AutoPilot Installer can be installed at these times:

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

AutoPilot Installer provides the following functions:

- Command line functionality – Initiates an installation from a command prompt or script. Configuration settings can be specified in the command line.
- Compatibility verification – Verifies that the driver to be installed is compatible with the operating system and platform.
- Driver installation and update – Installs and updates drivers.
- Multiple adapter installation capability – Installs drivers on multiple adapters, alleviating the requirement to manually install the same driver on all adapters in the system.

  NOTE Refer to the *Emulex Software Kit Migration User Guide* for information about installing drivers on a system containing a mix of OneConnect and LightPulse® adapters.
- Driver diagnostics – Determines whether the driver is operating properly.
- Silent installation mode – Suppresses all screen output (necessary for unattended installation).

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

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

NOTE These installation methods are not mutually exclusive.

- **Hardware-first installation**. At least one Emulex adapter must be installed before you can install the Emulex drivers and utilities.
- **Software-first installation**. You can install drivers and utilities using AutoPilot Installer prior to the installation of any adapters. You do not need to specify the adapter models to be installed later. The appropriate drivers and utilities automatically load when you install the adapters.
- **Utility-only installation**. If the drivers in the driver kit share the same version with those already installed on the system, you can reinstall or update the previously installed utility without reinstalling the drivers.
- **Text-only installation**. Text-based installation mode is used automatically when AutoPilot Installer is run on a Server Core system.
- **Network installation.** You can place the driver kit installers on a shared network drive and install them across your local area network (LAN). Network-based installation is often used with unattended installation and scripting. This allows you to configure and install the same driver and utility versions on all the hosts in a storage area network (SAN).
- **Unattended installation**. You can run the driver kit installers and AutoPilot Installer with no user interaction from a command line or script. Unattended installation works for both hardware-first and software-first installations and all driver kits. An unattended installation operates in Silent mode (also referred to as Quiet mode) and creates an extensive report file with installation status.

  NOTE Complete driver and utilities documentation can be downloaded from the Broadcom website.

2.3.1 Starting Installers from a Command Prompt or Script

If a driver kit or an AutoPilot Installer is run from a command prompt or command script (batch file), the Windows command processor does not wait for the installer to run to completion. As a result, you cannot check the exit code of the installer before the next command is executed. Use the "start" command with the "/wait" option. This causes the command processor to wait for the installer to finish before it continues.

For additional information on command line installation and configuration parameters, refer to Section D, AutoPilot Installer Command Line and Configuration File Parameters in the *Emulex Drivers for Windows for OneConnect Adapters User Guide*.

2.3.2 Running a Software Installation Interactively

Two options are available when performing an installation interactively. These options assume you have already downloaded the driver kit installer from the Broadcom website.

- Option 1 allows you to automatically run the AutoPilot Installer, which completes the driver kit loading and installation with a few mouse clicks.
- Option 2 allows you to run the AutoPilot Installer separately. This option is recommended when:
  - Changing installation settings for a limited number of systems
  - Familiarizing yourself with AutoPilot Installer configuration options

2.3.2.1 Option 1: Automatically Run AutoPilot Installer

Use this option unless you have specific configuration requirements.

1. Double-click the driver kit or run it from a command line. The command line parameter APargs allows you to specify arguments that are automatically passed to the AutoPilot Installer command.

   A **Welcome** window displays with driver kit version information and Broadcom contact information (refer to the AutoPilot Installer Command Line and Configuration File Parameters topic in the *Emulex Driver for Windows for OneConnect Adapters User Guide* for additional information on command line installations).
2. Click **Next** to proceed to the **Installation Options** window.

   For each installation option, the default installation location for that option is displayed. Browse to a different location, if desired.
3. Click **Install** to continue the installation.

   The **Progress** dialog displays.

   After all tasks complete, a **Finish** window displays. The Start AutoPilot Installer box is automatically selected.
4. Click **Finish**.

   AutoPilot Installer runs automatically. AutoPilot Installer completes one of the installations:
   - Section 2.3.3, Hardware-First Installation or Driver Update
   - Section 2.3.4, Software-First Installation

2.3.2.2 Option 2: Run AutoPilot Installer Separately

To access these options, run AutoPilot Installer after the driver kit loading has been completed. This allows you to change the configuration options supplied to the AutoPilot Installer.

1. Perform steps 1 through 3 for Option 1: Automatically Run AutoPilot Installer.
2. Clear the **Run AutoPilot Installer** check box on the **Finish** dialog.
3. Click **Finish**.

   The driver kit installer exits.

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

- Change the configuration file.
- Supply parameters on the command line.

NOTE See Appendix D, AutoPilot Installer Command Line and Configuration File Parameters, in the *Emulex Drivers for Windows for OneConnect Adapters User Guide* for additional information on either of these configuration methods.

After you have finished this step, you can run AutoPilot Installer at a later time, using either of the following methods:

NOTE If you are supplying options using the command line, you must run AutoPilot Installer from the command line.

- Run AutoPilot Installer from a command line by running the following command:

```
C:\Program Files\Broadcom\AutoPilot Installer\BrcmAPInstall.exe
```

NOTE The location of `BrcmAPInstall.exe` might differ on your system, depending on your system's Program Files location. You can also specify a different location when you install the driver package.

2.3.3 Hardware-First Installation or Driver Update

The driver kit installer must be downloaded from the Broadcom website and installed before performing this installation.

NOTES
- Updating the NIC protocol driver can temporarily disrupt operation of any NIC teams configured on the system.
- To update the Emulex protocol drivers, begin this procedure at step 2.

To perform a hardware-first installation, perform these steps:

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

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

2. Run AutoPilot Installer using any of the options listed in Running a Software Installation Interactively.
3. When the **AutoPilot Installer Welcome** window displays, select an adapter in the list and click **Next**.

   The installation continues.

Consider the following:

- If you are updating the driver, the existing port settings are used unless otherwise specified in the configuration file. These settings are pre-selected but can be changed. Set or change the settings, then click **Next**.
- If you are initially installing a vendor-specific version of the Emulex driver installation program, a **Driver Configuration** window may be displayed. This window includes one or more windows with questions that you must answer before continuing the installation process. In this case, answer each question and click **Next** on each window to continue.

4. Click **Next**. The installation is completed automatically.

   A dialog displays if Windows requires a reboot. After the installation is successful, a **Finish** window displays.
5. View or print a report, if desired.

— View Installation Report – The installation report is a text file with current Emulex adapter inventory, configuration information, and task results.
— Print Installation Report – The Windows print dialog is displayed to select options for printing the installation report.

6. Click **Finish** to exit AutoPilot Installer.

NOTE If the system must be rebooted, you are prompted to do so as indicated in step 4; you must reboot before using the drivers or utilities.

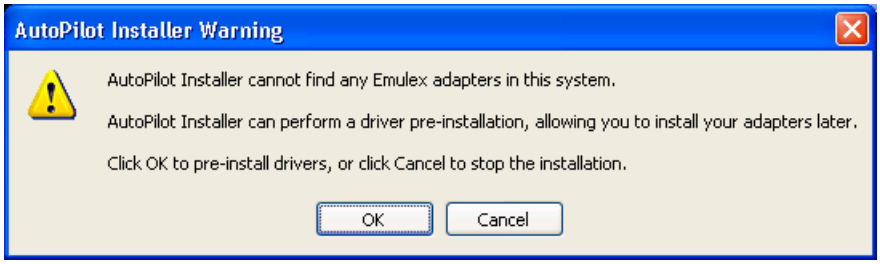
2.3.4 Software-First Installation

The driver kit must be downloaded from the Broadcom website and loaded before performing this installation.

To perform a software-first installation, perform these steps:

1. Run AutoPilot Installer using any of the options listed in Running a Software Installation Interactively.
   Figure 1 displays.

Figure 1 AutoPilot Installer Warning (Software-First Installation)



2. Click **OK**.
   A **Welcome** window displays.
3. Click **Next**. The installation automatically progresses.
   After the installation is successful, the **Finish** window displays.
4. View or print a report, if desired.
   — **View Installation Report** – The installation report is a text file with current Emulex adapter inventory, configuration information, and task results.
   — **Print Installation Report** – The Windows print dialog box allows you to select options for printing the installation report.
5. Click **Finish** to exit AutoPilot Installer.

2.3.5 Text-Only Driver Installation

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

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

2.4 Installing the Emulex iSCSI Crash Dump Driver (OCe14000-Series Adapters Only)

When you install the Emulex NIC driver, you can choose to also install the Emulex iSCSI crash dump driver (also referred to as the iSCSI Boot Firmware Table (iBFT) crash dump driver). The Emulex iSCSI crash dump driver is intended only for crash dump/hibernation and is not loaded during a normal boot for regular network operations. The Emulex iSCSI crash dump driver is named `ocibftcd`.

Because the Emulex iSCSI crash dump driver is a special driver used only for dump, the driver cannot be installed using Windows Device Manager. The driver is provided as an installation option in all the Emulex driver installer utilities.

When running the NIC Driver Kit in silent mode, add the `crashdriver=1` parameter to enable the Emulex iSCSI Crash Dump Driver, (The default is 0.) For example:

```
start /wait elxdrvr-nic-<version>.exe /q2 crashdriver=1
```

When running OneInstall in silent mode, add the `ibft=1` parameter to enable the Emulex iSCSI Crash Dump Driver. (The default is 0.) For example:

```
start /wait OneInstall-Setup-<version>.exe /q2 ibft=1
```

NOTES

- The iSCSI crash dump driver is not supported on Windows Server 2008 systems. It is supported only on OCe14000-series adapters on systems that support booting into Windows using iBFT configuration.
- The Emulex iSCSI crash dump driver is not supported on 32-bit, IA64, or ARM systems.

2.4.1 Supported Configurations

The dump miniport driver does not support network configurations on its own. The driver works with any iBFT configuration setting used to successfully boot into Windows. The supported network configurations are listed in Table 1.

Table 1 Supported Network Configurations

Network Configurations	Supported
Crash dump modes	Small Memory Dump Kernel Memory Dump Complete Memory Dump
Internet Protocol options	IPv4 IPv6 Autoconfigure
IP configuration methods	Static IP DHCP
Port configuration	Supported on any physical port used to boot into the system.
MTU settings	1514 bytes. However, the actual iSCSI boot might have occurred over a connection that used jumbo MTUs.
MPIO configuration	Yes

Table 1 Supported Network Configurations (Continued)

Network Configurations	Supported
Target discovery	Retrieved either through explicit configuration in iBFT menu or using DHCP
iSCSI Login redirection	Supports targets that enable iSCSI login redirection
Boot LUN Id	Supports crash dump on any LUN ID

Chapter 3: Installation Failure

The two possible failure scenarios are described in this chapter.

3.1 AutoPilot Installer Failure

If the AutoPilot Installer fails, the **Diagnostics** window shows that the adapter failed. If the adapter fails, perform these steps:

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

NOTE You can run `BrcmAPInstall.exe` from a command prompt.

3.2 OneInstall Installer Failure

The OneInstall Installer might fail for any of the following reasons:

- The operating system prerequisites have not been met.
- The individual kit installation failed. To check, run the installation interactively. If you encounter error messages when you run the installation interactively, those issues would also apply to an unattended installation.
- If an individual package failed to install properly, run that package's installer directly. This method displays status and error messages that can be used to diagnose the issue. (The OneInstall Installer does not provide this information because each package is installed silently.)


BROADCOM

Emulex® Drivers for Windows for OneConnect® Adapters

User Guide

Version 11.2
December 30, 2016

DRVWINOCA-UG112

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

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

Copyright © 2003–2017 Broadcom. All Rights Reserved.

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

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.

Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Chapter 1: Introduction

This product supports the Emulex® OneConnect® converged network adapters (CNAs).

- OCe11000-series adapters
- OCe14000-series adapters

1.1 Driver Information

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

- FCoE
  - Configuring the FCoE driver parameters
  - Improving server performance with FC/FCoE drivers
- Ethernet and TOE
  - Configuring NIC driver options
  - Configuring SR-IOV
  - Configuring Network Virtualization using Generic Routing Encapsulation (NVGRE)
  - Configuring RoCE supporting Server Message Block (SMB) Direct

    NOTE RoCE is available as Technical Preview only. Do not use RoCE in a production environment.
  - Configuring Universal Multichannel (UMC)
  - Configuring NIC partitioning (NPar) for Dell adapters only
  - Tuning network driver performance
- iSCSI
  - Configuring iSCSI driver options
  - Creating non-bootable targets
  - Configuring Multipath Input/Output (I/O)

A NIC teaming package driver and manager are also available as a separate download. The *Emulex OneCommand® NIC Teaming and VLAN Manager User Manual* is available for download as well. Refer to the Broadcom website for more information.

1.2 Abbreviations

AL_PA	arbitrated loop physical address
API	application programming interface
ARI	alternative routing-ID interpretation
BIOS	basic input-output system
CHAP	Challenge Handshake Authentication Protocol
CLI	command line interface
CNA	Converged Network Adapter
CNP	Congestion Notification Packet

CPU	central processing unit
CRC	cyclic redundancy check
CTCP	Compound TCP
DCB	Data Center Bridging
DCBX	Data Center Bridging Capabilities Exchange
DPC	deferred procedure call
DHCP	Dynamic Host Control Protocol
DID	device ID
DIMM	dual in-line memory module
DMA	direct memory access
DNS	Domain Name Server
DOS	disk operating system
DSM	Device Specific Module
ECN	Explicit Congestion Notification
ETO	extended timeout
ETS	Enhanced Transmission Selection
FC	Fibre Channel
FC-AL	Fibre Channel Arbitrated Loop
FCoE	Fibre Channel over Ethernet
FCP	Fibre Channel Protocol
FDMI	Fabric-Device Management Interface
FLOGI	fabric login
FW	firmware
FSB	front-side bus
GB	gigabyte
GbE	gigabit Ethernet
Gb/s	gigabits per second
GUI	graphical user interface
HBA	host bus adapter
iBFT	iSCSI Boot Firmware Table
ICMP	Internet Control Message Protocol
IEEE	Institution of Electrical and Electronics Engineers
IET	iSCSI Enterprise Target
I/O	input/output
IOCB	input/output control block
IOCTL	input/output control
IOMMU	input/output memory management unit
IOPs	I/O operations per second
IP	Internet Protocol
IPsec	Internet Protocol Security

IPL	initial program load
IPs	IP Security
iSCSI	internet Small Computer System Interface
IQN	iSCSI Qualified Name
KB	Knowledge Base
kb	kilobyte
LACP	Link Aggregation Control Protocol
LAN	local area network
LDTO	link down timeout
LRO	large receive offload
LSO	large send offload
LUN	logical unit number
MAC	Media Access Control
MPIO	multipath input/output
MSI	message signaled interrupts
MSS	maximum segment size
MTU	maximum transmission unit
N/A	not applicable
NAT	network address translation
NDIS	Network Driver Interface Specification
NIC	network interface card
NPar	NIC partitioning
NPIV	N_Port ID virtualization
NTFS	New Technology File System
NUMA	non-uniform memory access
NVGRE	network virtualization using generic routing encapsulation
OS	operating system
PCI®	PCI Express
PCIe	Peripheral Component Interconnect express
PDU	protocol data unit
PF	PCI function
PFC	process flow control or priority flow control
PLOGI	port login
POST	power-on self-test
PT-PT	point-to-point
PXE	Preboot Execution Environment
QCN	Quantized Congestion Notification
QFE	Quick Fix Engineering
QoS	quality of service

RAID	redundant array of independent disks
RCMD	Remote Command Service
RDMA	remote direct memory access
RFC	Request for Comments
RoCE	RDMA over Converged Ethernet
ROM	read-only memory
RSC	receive segment coalescing
RSCN	registered state change notification
RSS	receive-side scaling
RX	receive
SACK	selective acknowledgement
SAN	storage area network
SAS	serial attached SCSI
SCSI	Small Computer System Interface
SFP	small form factor pluggable
SLI	Service Level Interface
SMB	Server Message Block
SR-IOV	Single Root I/O Virtualization
SSH	Secure Shell network
TCP	Transmission Control Protocol
TCP/IP	TCP over Internet Protocol
TOE	TCP Offload Engine
TSO	TCP segmentation offload
TX	transmit
UDP	User Datagram Protocol
UEFI	Unified Extensible Firmware Interface
UMC	Universal Multichannel
VF	virtual function
VLAN	virtual local area network
VLAN ID	VLAN identifier
VM	virtual machine
VMQ	virtual machine queue
VPN	virtual private network
vPort	virtual port
WMI	Window Management Instrumentation
WWN	World Wide Name
WWNN	World Wide Node Name
WWPN	World Wide Port Name
XRI	extensible resource indicator

Chapter 2: Installation

Install the Windows drivers in one of two ways:

- OneInstall Installer – Contains all the applicable Emulex drivers for Windows (Emulex Storport Miniport and NDIS Miniport drivers) and the OneCommand® CNA Manager Application for OneConnect® Adapters in a single download package. The OneInstall Installer can be run in Interactive Mode or Silent Mode.
- Driver kits and AutoPilot Installer – Provides installation options ranging from simple installations with a few mouse clicks to unattended installations that use predefined script files and text-only installations.

NOTE Microsoft patch KB2846340 must be installed on your system to successfully install the NIC driver. If the patch is not installed on your system, the installation stops and prompts you to install it. This patch, from Microsoft's Knowledge Base, is required for Windows Server 2008 SP2, Windows Server 2008 R2, Windows Server 2012, and for Windows 7 and Windows 8 client operating systems.

NOTE If you are installing the NIC driver kit as an update to the Windows Server 2012 driver, Windows Server 2012 driver R2, or Windows Server 2016 driver, some parameter defaults are different from the inbox driver. After you install the Emulex out-of-box driver, select **reset to default** on the Advanced tab of the Device Manager property page. This action returns all adapter and driver settings to the default values listed in this guide.

NOTE Low performance might occur if the Emulex NIC driver is installed on a system meeting the following conditions before installing Microsoft KB2846837:

- A Windows 8, Windows 8.1, or Windows Server 2012 computer with multi-core processors is in use
- Three or more Ethernet ports are installed on the computer
- RSS is enabled and sets the RSS profile to use the `Closest` parameter for the Ethernet adapters

If any or all of these conditions exist, install KB2846837 before installing the Emulex NIC driver.

NOTE Windows 8 x64, Windows 8.1 x64, and Windows 10 drivers are Emulex signed. You must accept the Emulex certificate to install these kits. Support is provided by Broadcom, but not by Microsoft.

NOTE Check the Broadcom website for required updates to the Windows operating system or the Emulex drivers.

2.1 OneInstall Installer

The OneInstall Installer can be run in Interactive mode or Silent mode.

NOTE The OneInstall Installer does not allow you to perform preinstallation tasks or text-only installations. For these tasks, use the driver kits.

NOTE The OneInstall Installer is not supported on Window Server 2016 Nano Server.

The OneInstall package is a self-extracting executable file that installs the following software on your system:

- All compatible protocol drivers:
  - Fibre Channel over Ethernet (FCoE)
  - Internet Small Computer System Interface (iSCSI)
  - Network Interface Card or Controller plus RDMA over Converged Ethernet (NIC + RoCE)
- ElxPlus driver (supports the OneCommand CNA Manager Application, persistent binding, LUN mapping and masking, and virtual port creation)
- OneCommand CNA Manager Application (management utility for Emulex adapters)

NOTE The Enterprise kit for the OneCommand CNA Manager Application does not operate locally on Windows Server Core. You must install the OneCommand CNA Manager Application Core Kit (command-line interface only) to the Windows Server Core.

NOTE The One Command CNA Manager Enterprise and Core kit do not work on Windows Server 2016 Nano Server. You must install the OneCommand CNA Manager application for Windows Server 2016 Nano Server. See the *Emulex OneCommand CNA Manager application for OneConnect Adapters User Guide* for installation instructions.

2.1.1 Loading the OneInstall Package Using Interactive Mode

To install the drivers using Interactive mode, perform these steps:

1. Download the OneInstall package from the Broadcom website.
2. Navigate to the OneInstall package in Windows Explorer.
3. Double-click the OneInstall package.

   The **Welcome** screen displays.
4. Click **Next**.

   The **Installation Options** screen displays.
5. Select the drivers and applications that you want to install and click **Next**.

   A progress screen displays while the OneInstall installer loads the selected drivers and applications. After the drivers and application software are loaded, an **Installation completed** screen displays.
6. Click **Finish**.

2.1.2 Loading the OneInstall Package Using Silent Mode

Silent mode installation must be run from a batch file or from the command line.

If you run OneInstall from a batch file or from a command line prompt, the default Windows behavior starts OneInstall, then immediately continues with the next command. It does not wait until OneInstall has finished.

As a result, the value of %ERRORLEVEL% will always be 0, because Windows successfully started OneInstall. It does *not* reflect an accurate OneInstall exit code.

To remedy this, run setup as follows:

```
START /wait brcmOneInstall-Setup-<version>.exe
echo %ERRORLEVEL%
```

The "start /wait" ensures that the command does not return until setup has exited. The value of %ERRORLEVEL% now accurately reflects the OneInstall exit code.

2.1.2.1 Command Format

The format of the command is:

```
brcmOneInstall-Setup-<version>.exe  <install-mode>  <options>
```

Where:

<*version*> is the version number of the OneInstall package

<*install-mode*> is one of the following:

- `/q0` – (Interactive, non-silent install) This is the default.
- `/q1` – (non-Interactive install) This option displays status pages.
- `/q2` – (Silent install) This option is completely silent, no pages are displayed.
- `/q` – This is the same as `/q1`.

<*options*> specifies the kit, or kits, to install:

`ALL=1` – Install all drivers and the OneCommand CNA Manager Application for OneConnect Adapters (Default).

NOTE On a CORE system, this installs all drivers and the OneCommand CNA Manager Application for OneConnect Adapters Core Kit.

`ALLCORE=1` – Install all drivers and the OneCommand CNA Manager Application for OneConnect Adapters Core Kit.

`DRIVERS=1` – Install all drivers.

`FCOE=1` – Install the FCoE driver only.

`NIC=1` – Install the NIC driver only.

`ISCSI=1` – Install the iSCSI driver only.

`OCM=1` – Install the OneCommand CNA Manager Application for OneConnect Adapters Enterprise Kit only.

NOTE On a CORE system, this installs the OneCommand CNA Manager Application for OneConnect Adapters Core Kit.

`OMCORE=1` – Install the OneCommand CNA Manager Application for OneConnect Adapters Core Kit only.

To install the drivers using Silent mode, perform these steps:

1. Download the OneInstall package from the Broadcom website.
2. Open a DOS window.
3. Change directory to the folder containing your OneInstall package.

The following are examples of Silent mode commands:

```
Start /wait brcmOneInstall-Setup-<version>.exe  /q2  ALL=1
Start /wait brcmOneInstall-Setup-<version>.exe  /q2  DRIVERS=1
```

```
Start /wait brcmOneInstall-Setup-<version>.exe  /q2  FCOE=1  NIC=1  OCM=1
Start /wait brcmOneInstall-Setup-<version>.exe  /q2
Start /wait brcmOneInstall-Setup-<version>.exe  /q2  ALLCORE=1
Start /wait brcmOneInstall-Setup-<version>.exe  /q2  OCMCORE=1
```

2.2 Driver Kit Installer

Each driver kit contains and loads all the Windows drivers for a specific protocol, and includes ElxPlus.

- FCoE driver package (`brcmdrv-fcoe-<version>.exe`)
- iSCSI driver package (`brcmdrv-iscsi-<version>.exe`)
- NIC + RoCE driver package (`brcmdrv-nic-<version>.exe`)

NOTE Updating the NIC protocol driver can temporarily disrupt operation of any NIC teams configured on the system.

2.2.1 Loading the Driver Kit

The driver kit copies the selected Emulex drivers and applications onto your computer.

NOTE This procedure does not install drivers, and no driver changes are made until you run the AutoPilot Installer.

To load the driver kit, perform these steps:

1. Download the driver kit from the Broadcom website to your system.
2. Double-click the driver kit to run it.

   The Emulex Kit Welcome page displays.
3. Click **Next**.

   The **Installation Options** window displays.
4. Select one or both of the following options:
   - **Perform Installation of Software** – Copies the driver kit for your operating system to your computer.
   - **Unpack All Drivers** – Extracts all drivers to the current user's Documents folder. Select this option to perform boot from SAN installations.

   The **Operation in progress** window shows the kit file loading progress. After the kit files are loaded, the **Installation completed** window displays.
5. To continue with the installation, ensure that Start AutoPilot Installer is checked. Click **Next**.

2.3 AutoPilot Installer

AutoPilot Installer runs after the driver kit is loaded and the OneCommand CNA Manager Application for OneConnect Adapters is installed. AutoPilot Installer can be installed at these times:

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

AutoPilot Installer provides the following functions:

- Command line functionality – Initiates an installation from a command prompt or script. Configuration settings can be specified in the command line.
- Compatibility verification – Verifies that the driver to be installed is compatible with the operating system and platform.
- Driver installation and update – Installs and updates drivers.
- Multiple adapter installation capability – Installs drivers on multiple adapters, alleviating the requirement to manually install the same driver on all adapters in the system.

NOTE Refer to the *Emulex Software Kit Migration User Guide* for information about installing drivers on a system containing a mix of OneConnect and LightPulse® adapters.

- Driver diagnostics – Determines whether the driver is operating properly.
- Silent installation mode – Suppresses all screen output (necessary for unattended installation).

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

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

NOTE These installation methods are not mutually exclusive.

- **Hardware-first installation**. At least one Emulex adapter must be installed before you can install the Emulex drivers and utilities.
- **Software-first installation**. You can install drivers and utilities using AutoPilot Installer prior to the installation of any adapters. You do not need to specify the adapter models to be installed later. The appropriate drivers and utilities automatically load when you install the adapters.
- **Utility-Only installation**. If the drivers in the driver kit share the same version with those already installed on the system, you can reinstall or update the previously installed utility without reinstalling the drivers.
- **Text-Only installation**. Text-based installation mode is used automatically when AutoPilot Installer is run on a Server Core system.
- **Network installation.** You can place the driver kit installers on a shared network drive and install them across your local area network (LAN). Network-based installation is often used with unattended installation and scripting. This allows you to configure and install the same driver and utility versions on all the hosts in a storage area network (SAN).
- **Unattended installation**. You can run the driver kit installers and AutoPilot Installer with no user interaction from a command line or script. Unattended installation works for both hardware-first and software-first installations and all driver kits. An unattended installation operates in Silent mode (also referred to as Quiet mode) and creates an extensive report file with installation status.

NOTE Complete driver and utilities documentation can be downloaded from the Broadcom website.

2.3.1 Starting Installers from a Command Prompt or Script

If a driver kit or an AutoPilot Installer is run from a command prompt or command script (batch file), the Windows command processor does not wait for the installer to run to completion. As a result, you cannot check the exit code of the installer before the next command is executed. Use the "start" command with the "/wait" option. This causes the command processor to wait for the installer to finish before it continues.

For additional information on command line installation and configuration parameters, see Appendix D, AutoPilot Installer Command Line and Configuration File Parameters.

2.3.2 Running a Software Installation Interactively

Two options are available when performing an installation interactively. These options assume you have already downloaded the driver kit installer from the Broadcom website.

- Option 1 allows you to automatically run the AutoPilot Installer, which completes the driver kit loading and installation with a few mouse clicks.
- Option 2 allows you to run the AutoPilot Installer separately. This option is recommended when:
  - Changing installation settings for a limited number of systems
  - Familiarizing yourself with AutoPilot Installer configuration options

2.3.2.1 Option 1: Automatically Run AutoPilot Installer

Use this option unless you have specific configuration requirements.

1. Double-click the driver kit or run it from a command line. The command line parameter APargs allows you to specify arguments that are automatically passed to the AutoPilot Installer command.

   A **Welcome** window displays with driver kit version information and Broadcom contact information (refer to *Appendix D, AutoPilot Installer Command Line and Configuration File Parameters* for additional information on command line installations).
2. Click **Next** to proceed to the **Installation Options** window.

   For each installation option, the default installation location for that option displays. Browse to a different location, if desired.
3. Click **Install** to continue the installation.

   The **Progress** dialog is displayed.

   After all tasks complete, a **Finish** window displays. The Start AutoPilot Installer box is automatically selected.
4. Click **Finish**.

   AutoPilot Installer runs automatically. AutoPilot Installer completes one of the installations:
   - Section 2.3.3, Hardware-First Installation or Driver Update
   - Section 2.3.4, Software-First Installation

2.3.2.2 Option 2: Run AutoPilot Installer Separately

To access these options, run AutoPilot Installer after the driver kit loading has been completed. This allows you to change the configuration options supplied to the AutoPilot Installer.

1. Perform steps 1 through 3 for Option 1: Automatically Run AutoPilot Installer.
2. Clear the **Run AutoPilot Installer** check box on the **Finish** dialog.
3. Click **Finish**.

   The driver kit installer exits.

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

- Change the configuration file.
- Supply parameters on the command line.

NOTE Refer to Appendix D, AutoPilot Installer Command Line and Configuration File Parameters, AutoPilot Installer Command Line and Configuration File Parameters, for additional information on either of these configuration methods.

After you have finished this step, you can run AutoPilot Installer at a later time, using either of the following methods:

NOTE If you are supplying options using the command line, you must run AutoPilot Installer from the command line.

- Run AutoPilot Installer from a command line by running the following command:

```
C:\Program Files\Broadcom\AutoPilot Installer\BrcmAPInstall.exe
```

NOTE The location of `BrcmAPInstall.exe` might differ on your system, depending on your system's Program Files location. You can also specify a different location when you install the driver package.

2.3.3 Hardware-First Installation or Driver Update

The driver kit installer must be downloaded from the Broadcom website and installed before performing this installation.

NOTES
- Updating the NIC protocol driver can temporarily disrupt operation of any NIC teams configured on the system.
- To update the Emulex protocol drivers, begin this procedure at step 2.

To perform a hardware-first installation, perform these steps:

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

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

2. Run AutoPilot Installer using any of the options listed in Section 2.3.2, Running a Software Installation Interactively.
3. When the **AutoPilot Installer Welcome** window displays, select an adapter in the list and click **Next**.

   The installation continues.

Consider the following:

- If you are updating the driver, the existing port settings are used unless otherwise specified in the configuration file. These settings are pre-selected but can be changed. Set or change the settings, then click **Next**.
- If you are initially installing a vendor-specific version of the Emulex driver installation program, a **Driver Configuration** window may be displayed. This window includes one or more windows with questions that you must answer before continuing the installation process. In this case, answer each question and click **Next** on each window to continue.

4. Click **Next**. The installation is completed automatically.

   A dialog displays if Windows requires a reboot. After the installation is successful, a **Finish** window displays.
5. View or print a report, if desired.
   - View Installation Report – The installation report is a text file with current Emulex adapter inventory, configuration information, and task results.
   - Print Installation Report – The Windows print dialog is displayed to select options for printing the installation report.
6. Click **Finish** to exit AutoPilot Installer.

   NOTE If the system must be rebooted, you are prompted to do so as indicated in step 4; you must reboot before using the drivers or utilities.

2.3.4 Software-First Installation

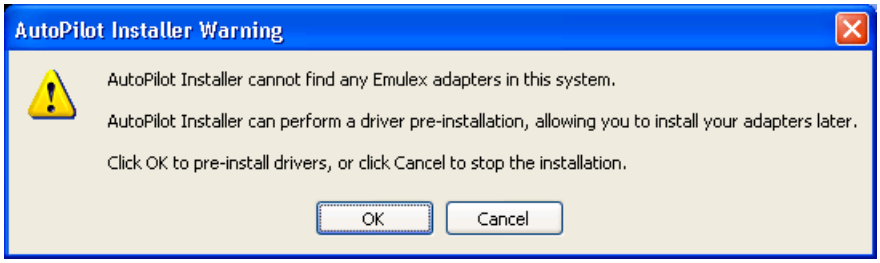
The driver kit must be downloaded from the Broadcom website and loaded before performing this installation.

To perform a software-first installation, perform these steps:

1. Run AutoPilot Installer using any of the options listed in Section 2.3.2, Running a Software Installation Interactively.

   Figure 1 displays.

Figure 1 AutoPilot Installer Warning (Software-First Installation)



2. Click **OK**.

   A **Welcome** window displays.
3. Click **Next**. The installation automatically progresses.

   After the installation is successful, the **Finish** window displays.
4. View or print a report, if desired.
   - **View Installation Report** – The installation report is a text file with current Emulex adapter inventory, configuration information, and task results.
   - **Print Installation Report** – The Windows print dialog box allows you to select options for printing the installation report.
5. Click **Finish** to exit AutoPilot Installer.

2.3.5 Text-Only Driver Installation

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

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

2.3.6 Unattended Driver Installation

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

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

- Reads the configuration file

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

An unattended installation can be performed in two ways:

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

2.3.6.1 Option 1: Install the Driver Silently

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

```
brcmdrvr-fcoe<version>.exe /q
```

NOTE The name of the driver kit depends on the current version identifier. For other command line options, see Appendix D, AutoPilot Installer Command Line and Configuration File Parameters.

2.3.6.2 Option 2: Run the Driver Kit Installer Separately

1. Follow steps 1 to 3 from Section 2.3.2, Running a Software Installation Interactively.
2. Clear the **Run AutoPilot Installer** check box on the Finish dialog.
3. Choose one of the following options:
   — Run the AutoPilot Installer from a command prompt or script with the silent option:
   `APInstall.exe /silent`
   — Edit the AutoPilot Installer configuration file before running AutoPilot Installer. The configuration file is typically located in:
   `C:\Program Files\Emulex\AutoPilot Installer\<driver type>\APInstall.cfg`

   Uncomment the line that sets `SilentInstallEnable` to `True`. You may also want to edit other settings in the same section of the configuration file related to unattended installations. See Appendix D.2.3, Software Configuration Parameters for additional information. After editing the file, run the AutoPilot Installer from the Start menu, a command prompt, or a script.

2.3.7 Installation Failure

The two possible failure scenarios are described in this section.

2.3.7.1 AutoPilot Installer Failure

If the AutoPilot Installer fails, the **Diagnostics** window shows that the adapter failed. If the adapter fails, perform these steps:

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

NOTE You can run `BrcmAPInstall.exe` from a command prompt.

2.3.7.2 OneInstall Installer Failure

The OneInstall Installer might fail for any of the following reasons:

- The operating system prerequisites have not been met.
- The individual kit installation failed. To check, run the installation interactively. If you encounter error messages when you run the installation interactively, those issues would also apply to an unattended installation.

NOTE If an individual package failed to install properly, run that package's installer directly. This method displays status and error messages that can be used to diagnose the issue. (The OneInstall Installer does not provide this information because each package is installed silently.)

2.4 Manually Installing or Updating the Emulex Protocol Drivers

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

The ElxPlus driver supports the OneCommand CNA Manager application, persistent binding, and LUN mapping and masking.

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

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

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

To install the ElxPlus driver from the desktop, perform these steps:

1. Run the driver kit installer, but do not run AutoPilot Installer. See Section 2.3.2, Running a Software Installation Interactively for instructions.
2. Select **Start>Settings>Control Panel>Add Hardware**. The Add Hardware Wizard window displays. Click **Next**.
3. Select **Yes, I have already connected the hardware** and click **Next**.
4. Select **Add a new hardware device** and click **Next**.
5. Select **Install the hardware that I manually select from a list (Advanced)** and click **Next**.
6. Select **Show All Devices** and click **Next**.
7. Click **Have Disk** and direct the Device Wizard to the location of `elxplus.inf`. If you have installed the driver installer kit in the default folder and `C:\` is your Windows system drive, the path is:

```
C:\Program Files\Broadcom\AutoPilot
Installer\FCoE\Drivers\Storport\x64\<OS>
```

8. Click **OK**.
9. Select **elxplus**. Click **Next** and click **Next** again to install the driver.
10. Click **Finish**.

    The initial ElxPlus driver installation has completed. Continue with manual installation of the Storport Miniport Driver. Refer to Section 2.5.2, Installing or Updating the FCoE Storport Miniport Driver for this procedure.

2.5.1 Updating the Emulex PLUS (ElxPlus) Driver

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

To update an existing ElxPlus driver from the desktop, perform these steps:

1. Run the driver kit installer, but do not run AutoPilot Installer. See Section 2.3.2, Running a Software Installation Interactively for instructions.
2. Select **Start>Settings>Control Panel>Administrative Tools>Computer Management**.
3. Click **Device Manager** (left pane).
4. Click **Emulex PLUS class** (right pane) to show the ElxPlus driver entry.
5. Right-click the ElxPlus driver entry and select **Update Driver** from the menu.
6. Select **No, not this time**. Click **Next** on the Welcome to the Hardware Update Wizard window. Click **Next**.
7. Select **Install from a list or specific location (Advanced)** and click **Next**.
8. Select **Don't Search. I will choose the driver to install**.
9. Click **Have Disk** and direct the Device Wizard to the location of the driver's distribution kit. If you have installed the driver installer kit in the default folder, the path is:

   ```
   C:\Program Files\Broadcom\AutoPilot Installer\FCoE\Drivers\Storport\x64<OS>
   ```

10. Click **OK**. Select Emulex PLUS.
11. Click **Next** to install the driver.
12. Click **Finish**.

    The ElxPlus driver update is finished. Continue with the manual installation of the Storport Miniport Driver.

2.5.2 Installing or Updating the FCoE Storport Miniport Driver

To update or install the FCoE Storport Miniport driver from the desktop, perform these steps:

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

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

6. Select the Driver tab.
7. Click **Update Driver**. The Update Driver wizard starts.

8. Select **No, not this time**. Click **Next** on the Welcome to the Hardware Update Wizard window.
9. Select **Install from a list or specific location (Advanced)** and click **Next**.
10. Select **Don't search. I will choose the driver to install** and click **Next**.
11. Click **Have Disk** and direct the Device Wizard to the location of `oemsetup.inf`. If you have installed the driver installer kit in the default folder, the path is:

    ```
    C:\Program Files\Broadcom\AutoPilot
    Installer\FCoE\Drivers\Storport\x64\<OS>
    ```

12. Click **OK**. Select Emulex OneConnect OCeXX000, PCI Slot X, Storport Miniport Driver (your adapter model is displayed here).
13. Click **Next**.
14. Click **Finish**.

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

2.5.3 Installing or Updating the iSCSI Driver

To update or install the iSCSI driver from the desktop, perform these steps:

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

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

10. Select I**nstall from a list or specific location (Advanced)** and click **Next**.
11. Select **Don't search. I will choose the driver to install** and click **Next**.
12. Click **Have Disk** and direct the Device Wizard to the location of `be2iscsi.inf`. If you have installed the driver installer kit in the default folder, the path is:

    ```
    C:\Program Files\Broadcom\AutoPilot
    Installer\iSCSI\Drivers\Storport\x64\[Windows Version]
    ```

13. Click **OK**. Select Emulex OneConnect OCe<*your adapter model*>, iSCSI Initiator.
14. Click **Next**.
15. Click **Finish**.

    The driver installation has completed. The driver will start automatically.

2.5.4 Installing or Updating the NIC Driver

NOTE The Microsoft patch KB2846340 must be installed on your system. This patch, from Microsoft's KB, is available for Windows Server 2012,

Windows Server 2008 R2, and Windows Server 2008 SP2 on the Microsoft website.

2.5.4.1 Installing or Updating the NIC Driver on Windows Server 2008 and Windows Server 2008 R2

1. Select **Start>Settings>Control Panel>Device Manager**.
2. Open the Network Adapters item.
3. Double-click the desired Emulex adapter.
4. Select the Driver tab.
5. Click **Update Driver**.

   The Update Driver wizard starts.
6. Click **Browse my computer** for driver software.

   NOTE The driver affects only the selected adapter. If there are other adapters in the system, you must repeat this process for each adapter. All dual-channel adapter models are displayed in the Device Manager as two adapters, therefore, you must update each adapter.
7. Click **Let me pick from a list of device drivers on my computer** and click **Next**.
8. Select the network adapter that matches your hardware and click **Have Disk**.
9. Direct the Device Wizard to the location of `be2nd6x.inf` for Windows Server 2008 or `be2nd62.sys` for Windows Server 2008 R2.

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

   `C:\Program Files\Broadcom\AutoPilot Installer\NIC\Drivers\NDIS\x64\<os>`
10. Click **OK**.

    The Windows Security dialog box opens.
11. Click **Install**.
12. After the device driver finishes installing, click **Close**.

    The driver installation is completed. The driver will start automatically.

2.5.4.2 Installing or Updating the NIC Driver on Windows Server 2012, Windows Server 2012 R2, and Windows Server 2016

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

   NOTE Server Manager is set to open by default when booting Windows Server 2012. If it does not open automatically, you can open it with the Server Manager icon at the bottom left of the screen.
2. Open the Network Adapters item.
3. Double-click the desired Emulex adapter.
4. Select the Driver tab.
5. Click **Update Driver**.

   The Update Driver wizard starts.
6. Click **Browse my computer for driver software**.

   The driver affects only the selected adapter. If there are other adapters in the system, you must repeat this process for each adapter. All dual-channel adapter models are displayed in the Device Manager as two adapters, therefore, you must update each adapter.
7. Click **Let me pick from a list of device drivers on my computer**.
8. Select the network adapter that matches your hardware and click **Have Disk**.

9. Direct the Device Wizard to the location of `ocnd63.inf` for Windows Server 2012, `ocnd64.inf` for Windows Server 2012 R2 and `ocnd65.inf` for Windows Server 2016.

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

   `C:\Program Files\Emulex\AutoPilot Installer\NIC\Drivers\NDIS\x64\<os>`

10. Click **Next**.
11. After the device driver finishes installing, click **Close**.

    The driver installation has completed. The driver will start automatically.

2.5.4.3 Installing the Emulex iSCSI Crash Dump Driver (OCe14000-Series Adapters Only)

When you install the Emulex NIC driver, you can choose to also install the Emulex iSCSI crash dump driver (also referred to as the iSCSI Boot Firmware Table (iBFT) crash dump driver). The Emulex iSCSI crash dump driver is intended only for crash dump/hibernation and is not loaded during a normal boot for regular network operations. The Emulex iSCSI crash dump driver is named `ocibftcd`.

Because the Emulex iSCSI crash dump driver is a special driver used only for dump, the driver cannot be installed using Windows Device Manager. The driver is provided as an installation option in all the Emulex driver installer utilities.

When running the NIC Driver Kit in silent mode, add the `crashdriver=1` parameter to enable the Emulex iSCSI Crash Dump Driver, (The default is 0). For example:

```
start /wait elxdrvr-nic-<version>.exe /q2 crashdriver=1
```

When running OneInstall in silent mode, add the `ibft=1` parameter to enable the Emulex iSCSI Crash Dump Driver. (The default is 0.) For example:

```
start /wait OneInstall-Setup-<version>.exe /q2 ibft=1
```

NOTES

- The iSCSI crash dump driver is not supported on Windows Server 2008 systems. It is supported only on OCe14000-series adapters on systems that support booting into Windows using iBFT configuration.
- The Emulex iSCSI crash dump driver is not supported on 32-bit, IA64, or ARM systems.

2.5.5 Supported Configurations

The dump miniport driver does not support network configurations on its own. The driver works with any iBFT configuration setting used to successfully boot into Windows. The supported network configurations are listed in Table 1.

Table 1 Supported Network Configurations

Network Configurations	Supported
Crash dump modes	Small Memory Dump Kernel Memory Dump Complete Memory Dump
Internet Protocol options	IPv4 IPv6 Autoconfigure
IP configuration methods	Static IP DHCP
Port configuration	Supported on any physical port used to boot into the system.

Table 1 Supported Network Configurations (Continued)

Network Configurations	Supported
MTU settings	1514 bytes. However, the actual iSCSI boot might have occurred over a connection that used jumbo MTUs.
MPIO configuration	Yes
Target discovery	Retrieved either through explicit configuration in iBFT menu or using DHCP
iSCSI Login redirection	Supports targets that enable iSCSI login redirection
Boot LUN Id	Supports crash dump on any LUN ID

2.5.6 Adding Emulex Out-of-Box Drivers to a Nano Server Virtual Hard Disk

Update the out-of-box (OOB) drivers prior to booting to the virtual hard disk (VHD) for the first time. The OOB drivers can be added to the VHD right after it is created. This procedure helps prevent the inbox driver from being loaded and linked to an Emulex device that is already present in the system.

To install the drivers, perform these steps:

1. Obtain the Driver Installer Kits from the following link:

   http://www.avagotech.com/support/download-search
2. Select the appropriate Product Family, Product and Asset type from the Support Documents and Downloads page.
3. In an elevated command prompt, navigate to the directory where the downloaded driver kit is located and run the following commands to unpack the drivers:

   ```
   brcmdrvr-nic-<version>.exe /q2 extract=2
   brcmdrvr-fcoe-<version>.exe /q2 extract=2
   brcmdrvr-fc-<version>.exe /q2 extract=2
   ```

   The drivers are extracted to the \Documents folder. For example:

   ```
   C:\Users\Administrator\Documents\Emulex\Drivers
   ```
4. In an elevated PowerShell prompt, navigate to the Nano Server VHD directory and run the following commands:

   ```
   md mountdir
   dism /Mount-Image /ImageFile:.\NanoServer.vhd /Index:1 /Mountdir:.\mountdir
   dism /Add-Driver /image:.\mountdir
   /driver:C:\Users\Administrator\Documents\Emulex\Drivers\NIC-<VERSION>\x64\
   win2016\ocnd65.inf
   dism /Unmount-Image /Mountdir:.\mountdir /commit
   ```

2.5.7 Adding Emulex OOB Drivers to running Nano Server with Pnputil.exe

NOTE This installation process assumes that the Nano Server system, and the system with a remote PowerShell connection to manage it, are both booted up and running.

To install the drivers, perform these steps:

1. Obtain the Driver Installer Kits from the following link:

   http://www.avagotech.com/support/download-search
2. Select the appropriate Product Family, Product and Asset type from the Support Documents and Downloads page.

3. In an elevated command prompt, navigate to the directory where the downloaded driver kit is located and run the following commands to unpack the drivers:

   ```
   brcmdrvr-nic-<version>.exe /q2 extract=2
   brcmdrvr-fcoe-<version>.exe /q2 extract=2
   brcmdrvr-fc-<version>.exe /q2 extract=2
   ```

   The drivers are extracted to the \Documents folder. For example:

   ```
   C:\Users\Administrator\Documents\Emulex\Drivers
   ```

4. Using a remote PowerShell connection, create the directory on the Nano Server machine to which you want to copy the drivers. For example, you can create a directory under the C: drive of the Nano Server machine as follows:

   ```
   md C:\<Directory Name>
   ```

5. In an elevated PowerShell prompt, navigate to the directory to which the Ethernet driver kit was extracted, and copy its contents to the Nano Server system using the following commands:

   ```
   $ip = "<Nano server IP Address>"
   $s = New-PSSession -ComputerName $ip -Credential ~\Administrator
   copy-ToSession $s -Path <full path to drivers from step 2>
   -Destination <Full path to Directory created in step 3> -Force - Recurse
   ```

   For example (typed all on one line):

   ```
   copy-ToSession $s -Path C:\EmulexDrivers -Destination C:\EmulexDrivers
   -Force -Recurse
   ```

6. Enter the remote Powershell session on the Nano server and navigate to the directory you created in step 3.
7. On the command line type

   ```
   Pnputil /add-driver elxplus.inf /install
   Pnputil /add-driver ocnd65.inf /install
   ```

   Additional help for pnputil.exe can be found on the command line by typing

   ```
   Pnputil /?
   ```

   Or on the Microsoft website.

2.6 Removing Emulex Driver Kits and Drivers

This section details procedures to uninstall the driver kits.

2.6.1 Uninstalling Emulex Driver Kits

NOTE If you uninstall the Emulex driver kit, AutoPilot Installer is automatically uninstalled.

2.6.1.1 Uninstalling an Emulex Driver Kit on Windows Server 2008

To uninstall a driver kit on a Windows Server 2008 system, perform these steps:

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

2.6.1.2 Uninstalling an Emulex Driver Kit on Windows Server 2012

To uninstall a driver kit on a Windows Server 2012 system, perform these steps:

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

2.6.1.3 Uninstalling an Emulex Driver Kit on a Server Core System

To uninstall a driver kit on a Server Core system, perform these steps:

1. From the system prompt, navigate to the Program Files folder.
2. Navigate to Emulex\AutoPilot Installer.
3. Run the following batch files:
   - `Brcmuninstall_cna_kit.bat`
   - `Brcmuninstall_nic_kit.bat`
   - `Brcmuninstall_iscsi_kit.bat`

The driver files are removed from the system.

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

2.6.2 Uninstalling the Emulex Drivers

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

2.6.2.1 Uninstalling an Emulex Drivers on Windows Server 2008

NOTE On Windows 2008, after the message `Warning – you are about to uninstall this device from your system` is displayed, you must select **Delete the software for this device** to uninstall the driver.

2.6.2.1.1 Uninstalling an Emulex Storport Miniport Driver

To uninstall the Emulex Storport Miniport driver, perform these steps:

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

2.6.2.1.2 Uninstalling an ElxPlus Driver

NOTE Uninstall the ElxPlus driver only if all adapters and installations of Emulex miniport drivers are uninstalled.

To uninstall the ElxPlus driver, perform these steps:

1. Select **Start>All Programs>Administrative Tools>Computer Management**.
2. Click **Device Manager**.
3. Click the **Emulex PLUS driver** class.
4. Right-click the Emulex driver and click **Uninstall**.
5. Click **OK** in the Confirm Device Removal window.

2.6.2.2 Uninstalling the Emulex Driver on Windows Server 2012

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

NOTE On Windows 2012 and Windows 2012 R2, after the message `Warning – you are about to uninstall this device from your system` is displayed, you must select the checkbox **Delete the software for this device** to uninstall the driver.

2.6.2.2.1 Uninstalling the Emulex Storport Miniport Driver

To uninstall the Emulex Storport Miniport driver in Windows Server 2012, perform these steps:

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

2.6.2.2.2 Uninstalling the ElxPlus Driver

NOTE Uninstall the ElxPlus driver only if all adapters and installations of Emulex miniport drivers are uninstalled.

To uninstall the ElxPlus driver, perform these steps.

1. Select **Server Manager>Dashboard>Tools>Computer Management>Device Manager**.
2. Click the **Emulex PLUS driver** class.
3. Right-click the Emulex driver and click **Uninstall**.
4. Click **OK** in the Confirm Device Removal window.

Chapter 3: Configuration

NOTE For information on configuring profile management, refer to the *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide* or the *Emulex OneCommand CNA Manager Command Line Interface for OneConnect Adapters User Guide*.

3.1 FCoE Driver Configuration

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

3.1.1 Configuring FCoE Driver Parameters

Table 2, Storport Miniport Driver Parameters, provides information such as the allowable range of values and factory defaults. Parameters can be entered in decimal or hexadecimal format.

A parameter has one of the following activation requirements:

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

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

Most parameters default to a setting that optimizes adapter performance.

Table 2 Storport Miniport Driver Parameters

Parameter	Definitions	Activation Requirement
AutoMap	AutoMap controls the way targets are assigned SCSI IDs. Discovered targets are assigned persistent SCSI IDs according to the selected binding method. Persistent bindings do not take effect with the driver in stand-alone mode. ■ 0 = automap is disabled. The OneCommand CNA Manager application persistently sets the SCSI address of a discovered FCP-capable FC node (target). ■ 1 = automap by WWNN. ■ 2 = automap by WWPN. ■ 3 = automap by DID. Value: 0–3 Default = 2	Reboot
DriverTraceMask	The DriverTraceMask parameter is only available on operating systems that support extended system event logging. ■ If set to 0 = the parameter is disabled. ■ If set to 1 = error events logging is enabled. ■ If set to 4 = warning events logging is enabled. ■ If set to 8 = informational events logging is enabled. The values can be masked to generate multi-levels of events logging. Values: 0, 1, 4, and 8. Default = 0.	Dynamic
EnableFDMI	If set to 1, enables management server login on fabric discovery. This allows FDMI to operate on switches that have FDMI-capable firmware. FDMI operates as FDMI-1. If set to 2, FDMI operates as FDMI-2. Value: 0–2 (decimal) Default = 2	Reset
EnableNPIV	If set to 1, enables NPIV. Requires NPIV supported firmware for the adapter. Value: 0–1 Default = 1 (enabled)	Reboot
FrameSizeMSB	FrameSizeMSB controls the upper byte of receive FrameSize if issued in PLOGI. This allows the FrameSize to be constrained on 256-byte increments from 256 (1) to 2048 (8). Value: 0–8 Default = 0	Reset

Table 2 Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement
InitTimeout	Determines the number of timeout seconds during driver initialization for the link to come up. If the link fails to come up by InitTimeout, driver initialization exits but is still successful. If the link comes up before InitTimeout, the driver sets double the amount for discovery to complete. Value: 5–30 seconds or 0x5–0x1E (hexadecimal) Default = 15 seconds (0xF)	Reboot
LimTransferSize	Limits the maximum transfer size to selectable values if this parameter is nonzero. Values: ■ 0 = Port (default) ■ 1 = 64 KB ■ 2 = 128 KB ■ 3 = 256 KB	Reboot
LinkTimeOut	LinkTimeOut applies to a private loop only. A timer is started on all mapped targets using the link timeout value. If the timer expires before discovery is resolved, commands issued to timed-out devices returns a SELECTION_TIMEOUT. The Storport driver is notified of a bus change event, which leads to the removal of all LUNs on the timed-out devices. Value: 1–500 seconds or 0x0–0xFE (hexadecimal) Default = 30 (0x1E)	Dynamic
LogErrors	LogErrors determine the minimum severity level required to enable entry of a logged error into the system event log. Errors are classified as severe, malfunction, or command level. A severe error requires user intervention to correct a firmware or adapter issue. An invalid link speed selection is an example of a severe error. A malfunction error indicates a system problem, but user intervention is not required. An invalid fabric command type is an example of a malfunction error. An object allocation failure is an example of a command error. ■ If set to 0 = All errors are logged. ■ If set to 1 = Command level errors are logged. ■ If set to 2 = Malfunction errors are logged. ■ If set to 3 = Severe errors are logged. Value: 0–3 Default = 3	Dynamic

Table 2 Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement
NodeTimeout	The node timer starts when a node (that is, a discovered target or adapter) becomes unavailable. If the node fails to become available before the NodeTimeout interval expires, the operating system is notified so that any associated devices (if the node is a target) can be removed. If the node becomes available before NodeTimeout expires, the timer is canceled and no notification is made. Value: 1–255 seconds or 0x0–0xFF (hexadecimal) Default = 30 (0x1E)	Dynamic
QueueDepth	QueueDepth requests per LUN/target (see the QueueTarget parameter). If you expect the number of outstanding I/Os per device to exceed 32, then you must increase to a value greater than the number of expected I/Os per device (up to a value of 254). If the QueueDepth value is set too low, a performance degradation can occur due to driver throttling of its device queue. QueueDepth supports more than 1000 outstanding commands per port. Value: 1–254 or 0x1–0xFE (hexadecimal) Default = 32 (0x20)	Dynamic
QueueTarget	QueueTarget controls I/O depth limiting on a per-target or per LUN-basis. ■ If set to 0 = Depth limitation is applied to individual LUNs. ■ If set to 1 = Depth limitation is applied across the entire target. Value: 0–1 or 0x0–0x1 (hexadecimal) Default = 0 (0x0)	Dynamic
RmaDepth	RmaDepth sets the remote management buffer queue depth. The greater the depth, the more concurrent management controls can be handled by the local node. Value: 8–64, or 0x8–0x40 (hexadecimal) Default = 16 (0x10) **NOTE** The RmaDepth driver parameter pertains to the functionality of the OneCommand CNA Manager application.	Reboot

Table 2 Storport Miniport Driver Parameters (Continued)

Parameter	Definitions	Activation Requirement
SrbTimeout	SrbTimeout limits the SCSI timeout value to 60 seconds if set to 1 or enabled. This parameter is a non-displayed parameter that must be set manually in the registry. This option alters the I/O timeout behavior, where an I/O will be returned in a maximum timeout of 60 seconds on long I/O timeouts. ■ If set to 1 = Enabled ■ If set to 0 = Disabled Values: 0, 1 Default = 0	
TraceBufSiz	TraceBufSiz sets the size in bytes for the internal driver trace buffer. The internal driver trace buffer acts as an internal log of the driver's activity. Value: 250,000–2,000,000 or 0x3D090–0x1E8480 (hexadecimal). Default = 250,000 (0x3D090)	Reboot
XLanePriority	Express Lane CS_CTL priority value. Refer to the switch vendor administration guide to set the value. Value: 0–7F (hexadecimal) Default = 0	Dynamic

3.2 NIC Driver Configuration

TOE is enabled by default on OCe11000-series adapters. It is not supported on OCe14000-series adapters.

3.2.1 Configuring NIC Driver Options

The Windows Server NIC driver supports configurable driver options through the Advanced Property page in the Windows Device Manager. For information on how to configure the options through the Advanced Property page, refer to Section 3.2.2.1, Modifying Advanced Properties.

For additional information on NIC driver options, see Section 27, Resource Counts on a 2-Port 10GbE OCe14000-Series Adapter.

You can also set configurable driver options using Microsoft PowerShell on Windows Server 2012 and Windows Server 2016.

Use the `Get-NetAdapter` PowerShell command to list all available adapters in the system. The `Get-Help` <cmdl> -full command returns descriptions and help for the cmdlets.

Use the following PowerShell commands to get and set driver parameter values.

To get the driver parameter value, perform this step:

```
Get-NetAdapterAdvancedProperty –Name <adapter name> -AllProperties
-RegistryKeyword <registry keyword>
```

Example:

```
Get-NetAdapterAdvancedProperty –Name ""SLOT 6 Port 1"" -AllProperties
-RegistryKeyword *RSS
```

To set the driver parameter value, perform this step:

```
Set-NetAdapterAdvancedProperty –Name <adapter name> -AllProperties
-RegistryKeyword <registry keyword> -RegistryValue <valid registry value>
```

Example:

```
Set-NetAdapterAdvancedProperty –Name ""SLOT 6 Port 1"" -AllProperties
-RegistryKeyword *RSS -RegistryValue 1
```

NOTE Refer to the documentation that accompanies the Windows Server 2012 and Windows Server 2016 operating systems for additional information on using PowerShell.

Refer to Table 3 for a list of NIC driver options.

Table 3 Windows Server NIC Driver Options

Option Name	Acceptable Values	Supported Operating Systems	Definition
`Class of Service (802.1p)`	Automatic Priority (default) Filtered Priority User Priority Disable Priority	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	The following modes are supported for selecting 802.1p priority tags: ■ Automatic Priority – The DCBX standard allows the network adapter to negotiate priority class usage with DCBX aware endpoints such as switches or network cards. If the peer indicates that priority pause is supported for a nonzero priority, the NIC automatically inserts the default priority in all transmitted packets. This mode is the default mode, which allows priority pause to operate for both storage and network traffic. ■ Filtered Priority – This mode coerces the user priorities in each packet to avoid sending packets on the network function that might disrupt the adapter's storage traffic. The network device uses the next lower priority if a conflict exists. This mode is useful if multiple network priorities are necessary. Only a limited number of classes are supported for priority pause, so typically, it does not function optimally in this mode. ■ User Priority – This mode allows any user-specified priority value and must be limited to cases where storage functions are not used. ■ Disable Priority – The adapter always transmits either untagged packets, or VLAN ID (802.1q) tagged packets with a priority value (802.1p) of zero.
`Enhanced Transmission Selection (ETS)`	Disabled (default) Enabled	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016 **NOTE** For OCe11102 CNAs only. **NOTE** ETS is not supported with VMQ technology. ETS is not available if SR-IOV is enabled.	If ETS is enabled, the driver filters transmit packets based on the 802.1p priority tag into multiple separate transmit rings. The network switch must be configured for ETS to group priorities into a priority group (or traffic class). Each priority group can be assigned a QoS bandwidth limit. For example, one network priority can support priority flow control to achieve loss-less network traffic. Using separate hardware interfaces in the driver allows each priority to progress at a different rate, or pause temporarily without affecting the other priorities. If ETS is enabled, all configurations regarding bandwidth and priority flow control must be performed on the network switch. The adapter will learn the configuration using the DCBX protocol.
`Flow Control`	Disabled RX and TX Enabled (default) RX Enable/TX Disable TX Enable/RX Disable	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Flow control is almost always advantageous to avoid packet drops on the network. The switch or network peer must also have flow control enabled. The IEEE 802.3x Ethernet specification defines a control frame between peers that can request a pause in packet transmissions. Flow control allows one system to request a temporary halt of all incoming traffic if receive buffer space is exhausted. The network device can be configured to respond to pause frames (RX Enable) and also to send pause frames (TX Enable). **NOTE** When NPar is enabled, flow control is set using UEFI.

Table 3 Windows Server NIC Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
Interrupt Moderation	0 (None) 8 (Static 90k Int/sec) 9 (Static 70k Int/sec) 10 (Static 50k Int/sec) 11 (Static 40k Int/sec) 12 (Static 30k Int/sec) 2 (Static 25k Int/sec) 13 (Static 20k Int/sec) 14 (Static 15k Int/sec) 15 (Static 10k Int/sec) 16 (Static 5k Int/sec) 4 (Adaptive) (default)	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	The network device uses interrupt moderation algorithms to reduce the total amount of CPU cycles spent processing interrupts, which increases efficiency for the system. However, interrupt moderation increases the latency of each send and receive. It should only be disabled when short latencies are more important than efficient CPU utilization. The "No Moderation" setting disables all delays to minimize latency. The "Static Moderation" setting uses a constant interrupt delay to avoid any spikes in the interrupt rate. The Adaptive (default) setting causes the driver to dynamically maintain a target interrupt rate. The Adaptive setting value is controlled by a dynamic algorithm that scales well for various adapter link speeds.
IP Checksum Offload (IPv4)	Disabled RX and TX Enabled (default) RX Enabled TX Enabled	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	This option offloads the transmit and the receive IPv4 checksum computation. Offloading checksums increases system efficiency.
Large Send Offload v1 (IPv4)	Disabled Enabled (default)	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Large Send Offload allows the NIC hardware to segment large TCP packets (up to 64 KB) into smaller packets (≤ "Packet Size") that can be transmitted. This segmentation increases transmit efficiency for TCP applications that send large buffers. During segmentation, the hardware computes the IPv4 and TCP checksums for each individual packet. The Windows Version 1 LSO supports only IPv4.
Large Send Offload v2 (IPv4)	Disabled Enabled (default)	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Large Send Offload allows the NIC hardware to segment large TCP packets (up to 64 KB) into smaller packets (≤ "Packet Size") that can be transmitted. This segmentation increases transmit efficiency for TCP applications that send large buffers. During segmentation, the hardware computes the IPv4 and TCP checksums for each individual packet. The Windows Version 2 LSO supports larger offload sizes.
Large Send Offload v2 (IPv6)	Disabled Enabled (default)	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Large Send Offload allows the NIC hardware to segment large TCP packets (up to 64 KB) into smaller packets (less than the MTU) that can be transmitted. This segmentation increases transmit efficiency for TCP applications that send large buffers. During segmentation, the hardware computes the TCP checksums for each individual packet. IPv6 support requires LSO Version 2.
Maximum Number of RSS Processors	Min: 0 Max: The number of CPU cores installed on your system	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	This property sets the maximum number of processors that can be used for RSS.

Table 3 Windows Server NIC Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
`Maximum Number of RSS Queues`	Windows Server 2012, Windows Server 2012 R2; ■ OCe11102, legacy: Min 1, Max 4, default 4 ■ OCe11102, advanced mode: Min 1, Max 8, default 8 ■ OCe14000: Min 1, Max 16, default 8	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	If RSS is enabled, this parameter controls the number of receive queues. Typically, this option is left at the maximum value. Windows reduces the number of queues as necessary based on the number of installed CPU cores. This value can be reduced during performance tuning for a particular application. It is possible that system performance might improve by limiting the number of RSS queues. Windows Server 2016For OCe11102 adapters, greater than four RSS queues requires that Advanced Mode Support be enabled in the BIOS controller configuration.
`Maximum RSS Processor Number`	Min: 1 Max: The number of CPU cores installed on your system	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	This parameter sets the maximum processor number for the RSS CPUs. This value is the highest processor number of any processors from the RSSMaxProcGroup parameter.
`Network Address`	Valid MAC address The default setting is None	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	This option overrides the permanent MAC address for the interface. The MAC address must follow this format *XX:XX:XX:XX:XX:XX*, where *X* is a hexadecimal digit (0–9 or A–F). ■ The address cannot be a multicast address, which has the lowest bit in the first byte set. ■ The address cannot be all zeros. For example, 01:00:00:00:00:00 is not valid, while 02:00:00:00:00:00 is valid.
`NetworkDirect`	0: Disabled 1: Enabled (default)	Windows Server 2012 R2 Windows Server 2016	The Network Direct feature enables an offloaded RDMA interface for SMB 3.0 network attached storage traffic using Microsoft's SMB Direct protocol. For best performance, PFC must be configured on the network switch. The Emulex driver defaults to priority (PFC) 5 for RoCE traffic, although it will still work without PFC enabled.
`Network Direct MTU`	256 512 1024 (default) 2048 4096	Windows Server 2012 R2 Windows Server 2016	The MTU or frame size for RoCE traffic can be configured with this parameter.
`NVGRE Task Offload` (also known as Encapsulated Task Offload)	Disabled Enabled (default)	Windows Server 2012 Windows Server 2012 R2 Windows Server 2016 **NOTE** For OCe14000-series adapters only.	NVGRE Task Offload enables the offloading of network virtualization using GRE tunneling on the NIC adapter. NVGRE offload works in conjunction with VMQ.
`Packet Size`	1514 (default) 9014 8222 4088	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Configures packet size for OneConnect NIC adapters only. This parameter determines the maximum packet size transmitted and received on the interface. A 1514 byte frame size is standard, while larger packets are called jumbo frames. Using a higher frame size is generally more efficient, but it uses more system memory. A larger frame size also requires support on the network switch. Jumbo frames are IPv4-only frames; IPv6 packets will be fragmented by LSO. Switches and the peer must be configured to accept the specified packet size, or the size will be negotiated to the common smallest size.

Table 3 Windows Server NIC Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
Performance Tuning	■ Maximum performance (default) ■ Dynamically balanced ■ Statically balanced	Windows Server 2012 Windows Server 2016	This parameter selects the driver algorithm for performance tuning, which allows you to balance raw networking throughput with overall system fairness among multiple devices and applications. ■ Maximum Performance – This mode maximizes the network performance for this adapter. This mode is the recommended mode. However, in systems with a large number of network or storage adapters, this mode can limit the performance of other devices. ■ Statically Balanced – This mode configures the network adapter to throttles CPU usage in all cases, allowing more balance among hardware devices and applications. If system responsiveness is poor, this mode can improve the overall system behavior. ■ Dynamically Balanced – Dynamic balancing adjusts the network adapter's performance based on system metrics, such as CPU usage. This mode can aggressively limit performance for the most stressful networking applications to ensure that all network adapters can share limited computer resources, yet it can maintain maximum performance if the system has resources available.
Preferred NUMA Node	Not present or a value from 0–65535. Optional. No default setting is set.	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2	Most modern multi-socket servers have separate memory controllers for each CPU socket. These systems have NUMA latencies for a given CPU core to access the local versus remote memory node. By setting this property, the driver attempts to use both memory and CPU cores from the given NUMA node. If the Preferred NUMA node is not set, the driver uses the preferred NUMA node as specified by the computer's BIOS. For best performance, the network applications must use memory and CPU affinity from the same NUMA node. This level of tuning is primarily noticeable when multiple adapters are running.
Physical Link Tracking	Disabled Enabled (default)	**NOTE** Supported for OCe11102 and OCe14000-series adapters only. Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	This parameter enables or disables physical link tracking when SR-IOV is used. When SR-IOV is enabled, the virtual Ethernet bridging (VEB) switch is used and the driver link status does not reflect the physical link status. This parameter forces the driver link status to reflect the physical link status. By default, physical link status tracking is enabled. Disable physical link tracking to allow physical function (PF) and VFs to communicate via the VEB switch regardless of the physical link status. When physical link status tracking is disabled, the driver always reports the link as UP. When physical link tracking is disabled, teaming failover does not work. When SR-IOV is disabled, the driver always reports physical link status.
Receive Buffers	64–32768, inclusive The default value is 896.	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	This option determines the number of Ethernet receive buffers allocated per receive queue. This number can be adjusted by the driver as needed.

Table 3 Windows Server NIC Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
Receive CPU	"Not Present" or a value from 0 through (number of CPUs on the system – 1). Optional. A default setting is not available.	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Sets the logical CPU used for processing the non-RSS receive packets. By default, the driver intelligently chooses a CPU in the system, so this parameter must only be used for advanced performance tuning. RSS packets are processed by the set of RSS CPUs provided by the Windows operating system.
Receive Side Scaling	Disabled Enabled (default)	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Support for multiple RSS queues if enabled. RSS scales receive processing over multiple CPUs in parallel. This scaling typically improves application performance; however, it tends to increase CPU usage on low end machines. For the OCe11102 adapter, RSS is only supported on two primary adapters per device (PF0 and PF1). For additional PCI functions, RSS does not appear in the Properties List.
Recv Segment Coalescing (IPv4)	Disabled Enabled (default)	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	RSC merges multiple TCP segments and identifies them as a single coalesced unit to the operating system's TCP/IP stack. This option reduces the per-packet receive processing overhead and CPU usage if standard 1514 byte sized frames are in use. **NOTE** If checksum offloads are disabled, RSC must also be disabled. RSC depends on checksum offloads for better performance. **NOTE** Both RSC (IPv4) and RSC (IPv6) are coerced to zero if TCP Connection Offload (IPv4) is enabled.
Recv Segment Coalescing (IPv6)	Disabled Enabled (default)	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	RSC merges multiple TCP segments and identifies them as a single coalesced unit to the operating system's TCP/IP stack. This reduces the per-packet receive processing overhead and CPU usage if standard 1514 byte sized frames are in use. **NOTE** If checksum offloads are disabled, RSC must also be disabled. RSC depends on checksum offloads for better performance. **NOTE** Both RSC (IPv4) and RSC (IPv6) are coerced to zero if TCP Connection Offload (IPv4) is enabled.
RoCE Mode	Acceptable values: 1, 2	Windows Server 2012 R2 Windows Server 2016	Configures routeable RoCE. 2 (Default) is routable RoCE and 1 is native RoCE.
RSS Base Processor Group	Min: 1 Max: 63	Windows Server 2012 Windows Server 2016	This option defines the base processor group for the RSS queues on the network adapter. A processor group contains 64 logical processors. This value can be modified with the "RSS Base Processor Number" to explicitly select the desired RSS processors for the adapter.
RSS Base Processor Number	Min: 1 Max: 63	Windows Server 2012 Windows Server 2016	This defines the base processor number for the RSS queues on the network adapter within the given processor group. A processor group contains 64 logical processors, so this value ranges from 0 to 63. This value can be modified with the "RSS Base Processor Group" to explicitly select the desired RSS processors for the adapter.
RSS Max Processor Group	Min: 0 Max: The number of processor groups present on your system	Windows Server 2012 Windows Server 2016	RSS Max Processor Group allows you to set the maximum number of processor groups for the RSS CPUs.

Table 3 Windows Server NIC Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
RSS Profile	Closest processor (default) Closest processor static NUMA scaling NUMA scaling static Conservative scaling	Windows Server 2012 Windows Server 2016	The RSS Profile setting determines the RSS load balancing profile implemented by Microsoft for this network adapter. The "Closest Processor" settings tend to localize the RSS CPUs to one NUMA node, allowing the device driver to allocate memory from the local node. The "NUMA Scaling" settings use all NUMA nodes on the system, and the memory allocation is not specific to a particular node. The driver ignores the Preferred NUMA node setting.
SpeedDuplex	AutoNeg (default) 10 Gb/sFullDuplex 1 Gb/sFullDuplex	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	SpeedDuplex is used for selecting link speed, mainly for 10GBASE-T adapters. If it is set to the default, it auto negotiates 100 Mbps/1 Gb/s/10 Gb/s with the switch/peer. Link speed can be forced to 1 Gb/s, if option 1Gb/sFullDuplex is selected. Link speed can be forced to 10 Gb/s, if option 10Gb/sFullDuplex is selected. 10 Gb/s is the maximum supported link speed.
SR-IOV	Disabled (default) Enabled	**NOTE** Supported for OCe11102, and OCe14000-series adapters only. Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	SR-IOV enables the adapter to allocate virtual PCI functions for each virtual machine in Hyper-V. **NOTE** The virtual switch and virtual network adapter must have SR-IOV enabled in the Hyper-V Manager. SR-IOV requires a platform with IOMMU virtualization (VT-d, AMD-Vi). If using SR-IOV, the Emulex NIC driver must be installed on each virtual function within the virtual machine. SR-IOV provides a direct hardware interface from the virtual machine to the networking adapter, which reduces latency and improves performance. The Windows Server 2012 and Windows Server 2012 R2 SR-IOV architecture establishes each Emulex virtual NIC with a corresponding emulated NIC. This allows the virtual machine to seamlessly failover to the emulated NIC if SR-IOV is disabled. It also allows Live Migration to another system, regardless of the installed NIC hardware. **NOTE** The driver currently supports the following virtual functions for the following adapter families: ■ OCe11100-series adapters support a maximum of 24 virtual functions per port. ■ OCe14000-series adapters support a maximum ofthe following functions: — 2-port 10 Gb/s: 31 virtual functions/physical function. — 4-port 10 Gb/s: 31 virtual functions/physical function — 1-port 40 Gb/s: 63 virtual functions/physical function
TCP Checksum Offload (IPv4)	Disabled RX and TX Enabled (default) RX Enabled TX Enabled	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	TCP Checksum Offload (IPv4) offloads the transmit or receive IPv4 TCP checksum computation. Offloading checksums increases system efficiency.
TCP Checksum Offload (IPv6)	Disabled RX and TX Enabled (default) RX Enabled TX Enabled	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	TCP Checksum Offload (IPv6) offloads the transmit or receive IPv6 TCP checksum computation. Offloading checksums increases system efficiency.

Table 3 Windows Server NIC Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
`TCP Connection Offload (IPv4)`	Enabled Disabled (default)	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	**NOTE** `TCP Connection Offload` is not supported on 16GFC adapters. If TCP Connection Offload is enabled, the device offloads the entire TCP protocol, including acknowledgement processing, retransmits, and timers. Applications that prepost receive buffers (before the data arrives) might avoid data copies in the receive path, which substantially increases the system efficiency and data rates. Windows does not offload TCP connections if any of the following are enabled: ■ Network Load Balancing ■ IPsec ■ Network Address Translation ■ NDIS 5.1 Intermediate Drivers TCP offload must be enabled in the Windows operating system with the shell command: `netsh int tcp set global chimney=enabled` This parameter appears disabled if the firmware installed on your device does not support TCP connection offload. Upgrading the firmware might resolve this issue. View the "Statistics" property page to ensure that TCP connection offload is working. **NOTE** Both RSC (IPv4) and RSC (IPv6) are coerced to zero if TCP Connection Offload (IPv4) is enabled.
`TCP Offload Optimization`	Optimize Latency Optimize Throughput (default)	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	This parameter only applies to TCP connection offload, which must be enabled in the "Protocol Offloads" section. Most applications perform better with TCP Offload Optimization set to `Optimize Throughput`, which handles large data transfers with minimal CPU impact. Setting this parameter to "Optimize Latency" causes receive data to be delivered to the application without waiting for a TCP push flag. This causes additional receive indications that typically decrease total throughput.
`Transmit Buffers`	64–256, inclusive The default setting is 256.	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	`Transmit Buffers` sets the number of Ethernet transmits that might be posted to the hardware at any given time. The default value is sufficient to achieve maximum performance. Reducing this value conserves system memory.
`Transmit CPU`	"Not Present" or a value from 0 through (number of CPUs – 1). Optional. A default setting is not available.	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Sets the CPU to be used to process transmit completions. By default, the driver intelligently chooses a CPU in the system, so this parameter must only be set for advanced performance tuning.
`Transmit Side Scaling (TSS)`	Enabled Disabled	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	`Transmit Side Scaling` distributes transmit completions to be processed on multiple CPUs in parallel. It uses the RSS CPU table for distribution, and therefore, requires RSS to be enabled.

Table 3 Windows Server NIC Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
`UDP Checksum Offload (IPv4)`	Disabled RX and TX Enabled (default) RX Enabled TX Enabled	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	`UDP Checksum Offload` settings offload the transmit or receive IPv4 UDP checksum computation. Offloading checksums increases system efficiency.
`UDP Checksum Offload (IPv6)`	Disabled RX and TX Enabled (default) RX Enabled TX Enabled	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	`UDP Checksum Offload` settings offload the transmit or receive IPv6 UDP checksum computation. Offloading checksums increases system efficiency.
`Virtual Machine Queues`	Enabled (default) Disabled	VMQs require Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016 with Hyper-V	VMQs are dedicated hardware receive queues for virtual machines that filter receive packets based on the destination MAC address or VLAN. Receive buffers can be allocated for each queue from VM memory. This improves network throughput by distributing processing of network traffic for multiple VMs among multiple processors. It reduces CPU utilization by offloading receive packet filtering to NIC hardware. VMQs prove beneficial when four or more VMs are in use.
`Virtual Machine Queues Lookahead Split`	Enabled (default) Disabled	**NOTE** Supported for OCe11102 CNAs only. Not applicable for OCe14000-series adapters. Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	VMQ enables direct DMA to VM memory. Lookahead improves packet steering performance by PCI prefetching an adjacent header buffer into a cache when examining a packet. Header buffers are continuous in physical memory because they belong to one pool. For OCe11102, Lookahead split requires Advanced Mode Support and is enabled in the BIOS controller configuration. **NOTE** Lookahead split is not supported for jumbo frames.
`Virtual Machine Queues Transmit`	Enabled (default) Disabled	**NOTE** Supported for OCe11102 and OCe14000-series adapters only. Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	If this option is enabled with VMQs, separate transmit queues are created for each VM network interface. Send and receive interrupts for a VM network interface are processed on the same CPUs. Separate transmit queues increase system overall CPU utilization, but offer greater system scalability.
`VLAN Identifier (802.1q)`	Not Present (default) 1 to 4094	Windows Server 2008 Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	If selected, the adapter adds a VLAN tag to all transmitted packets, and only receives packets with the matching VLAN tag. **NOTE** Do not use this property if the Emulex Teaming Driver is enabled. In that case, perform VLAN configuration in the Teaming Driver application. **NOTE** Do not use this property with Hyper-V. Use the Microsoft Hyper-V Manager to configure VLANs on each virtual machine.

Table 3 Windows Server NIC Driver Options (Continued)

Option Name	Acceptable Values	Supported Operating Systems	Definition
`Wake on LAN`	Enabled (default) Disabled	Windows Server 2008	Enabling `Wake on LAN` allows the network device to wake up the computer if a magic packet is received during standby. In Blade server configurations, `Wake on LAN` is only supported on two primary adapters per device. Additional PCI functions appear disabled.
`Wake on Magic Packet`	0 – Disabled 1 (Default) – Enabled	Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Defines if a network adapter is enabled to wake a computer on the magic packet.
`Wake on Pattern Match`	0 – Disabled 1 (Default) – Enabled	Windows Server 2008 R2 Windows Server 2012 Windows Server 2012 R2 Windows Server 2016	Defines if a network adapter is enabled to wake the computer on pattern matches.

3.2.1.1 Advisory: PowerShell Behavior

3.2.1.1.1 Issues with Capabilities Reported by Standard PowerShell Commands (Get-NetAdapter)

Driver parameter default registry values are initially populated from the driver installation INF file. Thereafter, the registry is written to only if the default settings are explicitly overridden. PowerShell uses these registry values to report capabilities with the result that the registry values might not always reflect what is supported in the current configuration.

The default settings can be modified through the Driver Properties page, standard PowerShell commands, and utilities like occfg (for additional information on occfg, see Section 3.2.3, Using OCCFG for Windows NIC Driver Options).

Standard PowerShell (Get-NetAdapter) commands function in the following manner:

- If the feature is currently enabled, the driver reports its current capabilities. PowerShell reports all of the feature capabilities based on what the driver indicates. The feature capabilities are guaranteed to be what the NIC driver supports in the current configuration.
- If the feature is not enabled, the driver does not report any current capabilities. At that point, PowerShell searches the registry for keys related to the feature and reports their values. These values are either the default values (INF) or the last configured user values (if overwritten by the user). Default values are only intended as maximum upper bounds; they are not guaranteed resources supported in every configuration.

As a result, the driver can only report a feature's current capabilities (accurate for the present configuration) if the feature is currently enabled. However, standard PowerShell commands report whatever is present in the registry if the feature is not enabled. This information can conflict with what the driver actually supports in the current configuration.

3.2.1.1.2 Determining What PowerShell is Reporting (Registry and Driver-Reported Capabilities)

You can usually tell whether PowerShell is using capabilities reported by the driver or is picking up registry values.

- SR-IOV

  Check the output of (Get-NetAdapterSRIOV). CurrentCapabilities for CurrentCapabilities.

  If CurrentCapabilities is empty, the driver is not currently enabled for SR-IOV. Any reported fields in Get-NetAdapterSriov | fl * are based on registry values. If CurrentCapabilities is not null, the driver is enabled for SR-IOV. Get-NetAdapterSriov fields are based on what the driver reports.

NetAdapter*) commands function in this manner.

- RDMA

  Check the output of (Get-NetAdapterRdma). RdmaAdapterInfo for RdmaAdapterInfo.

  If RdmaAdapterInfo is empty, any reported fields in Get-NetAdapterRdma | fl * are based on registry values. If RdmaAdapterInfo is not null, the driver is reporting RDMA capabilities. Get-NetAdapterRdma fields are based on what the driver reports.

3.2.1.2 Considerations for Using UMC and NIC

NOTE 64 VLAN IDs can be used with each UMC virtual channel.
SR-IOV must be disabled if UMC is enabled.

For additional information on UMC, refer to the *Emulex Universal Multichannel Reference Manual,* which is available for download from the Broadcom website.

3.2.1.3 ARI Considerations

NOTE RoCE is not supported with ARI.

The PCIe standard limits an adapter to a maximum of eight physical functions. This means that a 2-port adapter can only have four functions per port and a 4-port adapter can only have two functions per port. The following requirements must be met to fully support ARI and expose more than eight functions on an adapter:

- ARI must be available on the system to support up to 16 functions on an adapter. If these conditions are not met, although you can configure all 16 functions, only eight functions will be present and discovered by the OneCommand CNA Manager application after a reboot.
- Only OCe14000-series adapters support ARI.
- The system hardware, such as the motherboard and BIOS, must support ARI.
- ARI must be enabled in the system BIOS.
- The operating system must support ARI, such as the Windows Server 2012 and later.
- Any management tools that you use must support ARI, such as the OneCommand CNA Manager application.

For Dell NPar support, refer to Section 3.2.8.1, NPar Configuration (Dell Only).

3.2.2 Configuring Windows Server NIC Driver Parameters

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

NOTE Ensure that the OneCommand CNA Manager application GUI is closed before opening the Windows Device Manager.

3.2.2.1 Modifying Advanced Properties

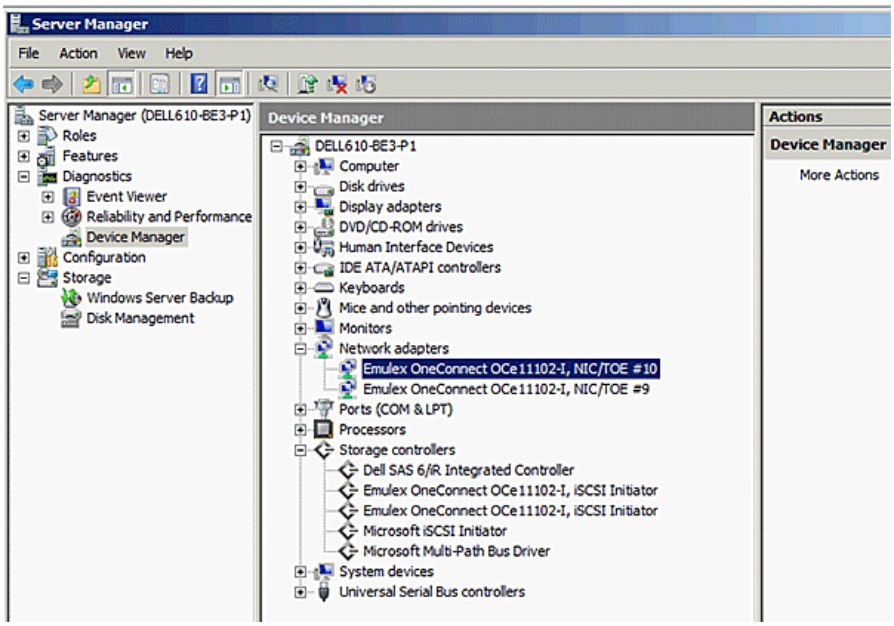
Modify the advanced properties for the driver for Windows with the Windows Device Manager. For additional information on advanced properties, refer to Section 3.2.9, Network Driver Performance Tuning.

To modify the advanced properties, perform these steps:

1. Enter the Windows Device Manager using one of the following options:
   - Click **Start> Control Panel>System** and click the Device Manager hyperlink.
   - Click **Start>Run**, and type
     `devmgmt.msc`
2. Click **OK**.

The Windows Device Manager is displayed (refer to Figure 2).

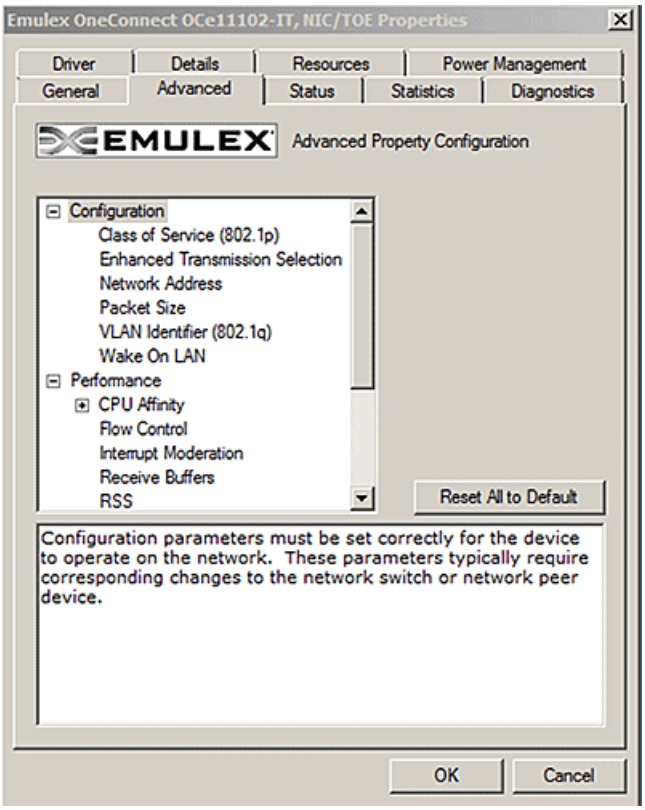
Figure 2 Partial View of Windows Device Manager



3. Right-click the network adapter for which you want to modify advanced properties.
4. Click **Properties**, and click the **Advanced** tab (refer to Figure 3).
5. From the list of properties, click the property (parameter) you want to modify, then select the new value of the property by selecting it from the Value list.
6. Click **OK**.

NOTE Modifying properties causes the network driver to reload, and some TCP connections might be temporarily dropped.

Figure 3 NIC Advanced Properties in Windows Server 2008



3.2.2.2 Statistics Property Page

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

To view the statistics properties, perform these steps:

1. Enter the Windows Device Manager using one of the following options:

- Click **Start>Control Panel>System** and click the Device Manager hyperlink.
- Click **Start>Run**, then type:

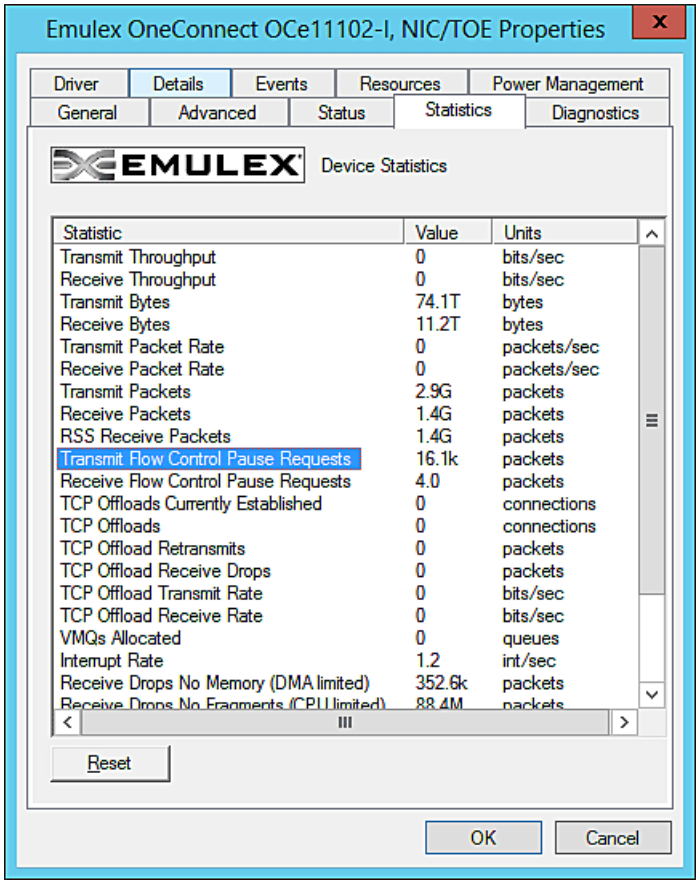
  `devmgmt.msc`

2. Click **OK**.

   The Windows Device Manager is displayed (refer to Figure 2).

3. Right-click the network adapter for which you want to view the statistics properties.
4. Click **Properties,** then click the **Statistics** tab (refer to Figure 4).

Figure 4 NIC Statistics Properties in Windows Server 2008



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

Table 4 lists the NIC driver properties statistics.

Table 4 NIC Driver Properties Statistics

Statistic Name	Description
`Transmit Throughput`	The data rate for this adapter on the network, including all packet headers. It is expressed in terms of bits per second, where 1 byte = 8 bits. This range is computed as the average over approximately three seconds.
`Receive Throughput`	The receive rate for this adapter.
`Transmit Bytes`	The total number of bytes transmitted by this adapter, since the last statistics reset or the last driver reload.
`Receive Byte`	The total number of bytes received by this adapter.
`Transmit Packet Rate`	The rate of transmit packets for the adapter.
`Receive Packet Rate`	The rate of receive packets for the adapter.
`Transmit Packets`	The total number of packets transmitted by the adapter since the last statistics reset, or the driver was reloaded.
`Receive Packets`	The total number of packets received. This number includes both RSS and non-RSS packets.
`RSS Receive Packets`	The number of receive packets that were suitable for RSS.
`Transmit Flow Control Pause Requests`	The number of times the network adapter sent a PAUSE frame to request that the peer stop sending data temporarily. This number indicates a potential bottleneck in the system. Typically, this bottleneck is the result of the DMA of packets from the adapter to host memory.

Table 4 NIC Driver Properties Statistics (Continued)

Statistic Name	Description
Transmit Flow Control PFC Pause Requests	IEEE 802.1Qbb PFC extends Ethernet PAUSE frame to each of the 8 traffic classes. The PFC PAUSE frame is a link-level mechanism used by overwhelmed receiver nodes to halt the transmission from the peer node for a specified period of time.
Receive Flow Control Pause Requests	The number of times the network adapter received a PAUSE frame from the peer. This number indicates a potential bottleneck in the attached switch or network peer device. This statistic increments only if the switch is correctly configured for flow control.
Receive Flow Control PFC Pause Requests	IEEE 802.1Qbb PFC extends Ethernet PAUSE frame to each of the 8 traffic classes. The PFC PAUSE frame is a link-level mechanism used by overwhelmed receiver nodes to halt the transmission from the peer node for a specified period of time.
TCP Offloads Current Established	The current number of TCP connections offloaded to the adapter's TOE.
TCP Offloads	The total number of TCP connections that have been offloaded since the last statistics reset or the driver was reloaded.
TCP Offload Retransmits	The number of packets retransmitted for TCP offloaded connections.
TCP Offload Receive Drops	The number of packets dropped by in the offloaded TCP stack. These drops might be the result of TCP protocol errors or bottlenecks in the system for consuming receive data.
TCP Offload Transmit Rate	The transmit data rate of the offloaded TCP connections. This value is the portion of the total "Transmit Throughput" contributed by offloaded TCP connections.
TCP Offload Receive Rate	The receive data rate of the offloaded TCP connections.
VMQs Allocated	The current number of Virtual Machine Queues allocated.
Interrupt Rate	The number of interrupts per second generated by the adapter. The interrupt rate can be tuned by modifying the Interrupt Moderation parameter.
Receive Drops No Memory (DMA Limited)	The number of packets dropped as a result of insufficient buffers posted by the driver. This value is generally the result of the CPU core used for any receive queue reaching 100%. The system might lack sufficient CPU cycles to post receive buffers at the necessary rate. Many small packets lead to this behavior on almost any CPU, because the processing time for small packets is very high in the networking stack. Using a teaming driver might also lead to this behavior, because it increases the CPU load during receive. Increasing the number of "Receive Buffers" in the advanced property page might alleviate some of these drops, in particular if the drops are the result of bursts of small receive packets on the network. However, if the CPU is the limit, increasing the buffer resources does not help because the driver cannot post them fast enough. Enabling RSS is another strategy to reduce drops because it allows the NIC driver to use additional CPU cores. The number of RSS queues can be increased to increase the total number of posted buffers available to the adapter. Enabling RSC can also reduce CPU consumption in the networking stack by combining multiple TCP packets into one larger packet. For best performance, the system BIOS must be set to "Maximum Performance" or manually disable C-states. The transitions to low power C-states might cause a steady trickle of drops due to increased latencies from packet reception until the driver's interrupt processing code is invoked.

Table 4 NIC Driver Properties Statistics (Continued)

Statistic Name	Description
`Receive Drops No Fragments (CPU Limited)`	The number of receive packets dropped because of a DMA bottleneck from the network adapter to host memory. This situation might be caused by bottlenecks in either the PCIe bus or main memory. In the Status tab of the Custom property page, the Emulex NIC reports the PCIe link parameters and the maximum supported parameters. For example, installing a 8x device in a 4x PCIe slot cuts the available PCIe bandwidth in half. The PCIe MTU and Read Request size are also reported, and these can be configurable in the system BIOS. The performance of the main memory is the other major concern for networking throughput. The ideal situation uses high speed memory with all memory channels populated per CPU; typically, three or four DIMMs per CPU socket. For the ideal performance, use the same DIMM size in each memory channel to allow perfect memory channel interleaving. Features, such as memory sparing or memory mirroring, dramatically decrease the memory bandwidth of the system and cause drops. TCP connection offload might lead to increased drops as a result of "no memory." If TCP connection offload is used, enabling flow control might reduce the drops. Alternatively, disabling TCP connection offload might improve performance.
`Receive CRC Errors`	The number of packets dropped as the result of CRC errors on the layer 2 Ethernet packet. In products that expose multiple PCIe functions per Ethernet port, this statistic is incremented only for the lowest PCI function per port since the packet cannot be further classified because of the error.
`Receive IP Checksum Errors`	The number of receive packets with an incorrect IPv4 checksum. These packets are provided to the TCP/IP stack for disposal in the operating system.
`Receive UDP Checksum Errors`	The number of receive packets with an incorrect UDP checksum. These packets are provided to the TCP/IP stack for disposal in the operating system.
`Receive TCP Errors`	The number of receive packets with an incorrect TCP checksum. These packets are provided to the TCP/IP stack for disposal in the operating system.
`Tunnels allocated`	The number of interfaces converted to tunnel interfaces. Used with NVGRE offload enabled and on.
`Tenants allocated`	The number of interfaces converted into tenant interfaces. Used with NVGRE offload enabled and on and VMQ.
`Virtual Functions allocated`	The number of PCIe virtual functions created by the SR-IOV supporting adapter.

3.2.3 Using OCCFG for Windows NIC Driver Options

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

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

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

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

3.2.3.1 Displaying OCCFG Help

To display help, use the -? option by typing

```
occfg -?
```

The following text displays:

```
OneConnect Network Config (0.0.9999.0)
Copyright 2011 Emulex
Usage:   occfg.exe [-options]
```

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

Options:

`-a str[,str]`	Selects all adapters with any of the given strings in the connection or device name. If omitted, occfg prompts for an adapter from a list.
`-s name=v,[name=v]`	Sets the parameter's value and reloads the devices.
`-g name[,name]`	Gets parameter value.
`-r`	Skips reloading the driver when setting a parameter.
`-f`	Forces reloading the driver.
`--`	Forces disabling the driver.
`-+`	Forces enabling the driver.
`-l`	Lists available adapters and exits.
`-T filename`	Saves tinylog to a binary file.
`-L filename`	Loads a binary file and replays tinylog.
`-x`	Resets all parameters to the default values.
`-p`	Shows all registry parameter values.
`-q`	Shows all driver parameter values.
`-h`	Shows help text for all parameters.
`-?`	Shows this help.
`-M module=trace level [,module=trace level]`	Continously downloads ARM log into a file. Arguments set a specific trace level on listed modules. Default argument is `all=error`. Refer to ARM firmware for list of modules and debug trace levels. This is a special command argument.

3.2.3.1.1 Examples:

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

3.2.3.2 Selecting an Adapter

In batch mode, the `-a` parameter must be followed by a substring that is contained within the adapter name. The name is a combination of the device manager name (for example, Emulex OneConnect OCe11102) and the network connection name (for example, Local Area Connection). The latter can be modified by using the Windows Network Connections applet (`ncpa.cpl`).

The most typical scenario involves setting parameters to be the same for all ports of a network adapter by specifying `-a emulex`.

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

3.2.3.3 Configuring Device Parameters

The `occfg` program is used to query and modify registry parameters for Emulex network devices. The registry keys are stored at:

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

where "####" is the device instance number.

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

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

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

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

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

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

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

NOTE Quotation marks are required if the parameter name contains a space character.

To modify a parameter without a driver reload, use `-r`. This setting is used to modify several parameters in sequence, without forcing a driver reload. To force a driver reload, use the `-f` parameter.

The following is an example of such a sequence:

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

Registry keys can be set to two special values:

- The "delete" value causes the key to be entirely deleted and the driver uses the default value. This value is appropriate for keys that are optional, such as the "Network Address".
- The "default" value sets the key to the driver's default value. If the key is optional, the default value might be equivalent to deleting the key.

For example:

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

3.2.3.4 Viewing Device Parameters

The `occfg.exe` program can query device parameters from either the registry or the device driver (if running driver version greater than or equal to 2.103.*x.x*).

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

NOTE If the driver has been disabled or if the driver failed to load due to any error, the driver query returns the error, "Failed to query driver for the parameter."

The following are batch mode examples:

```
occfg -a emulex -g "Interrupt Moderation"
occfg -a "(Local Area Connection)" -g interrupt,rss
Emulex OneConnect OCe11102-I, NIC/TOE (Local Area Connection):
   [Registry] Interrupt Moderation = 4 (Adaptive [Default])
   [Driver] Interrupt Moderation = 4 (Adaptive [Default])
Emulex OneConnect OCe11102-I, NIC/TOE (Local Area Connection):
   [Registry] RSS = 0 (Disable)
   [Driver] RSS = 0 (Disable)
```

3.2.3.5 Resetting All Parameters

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

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

3.2.3.6 Displaying All Parameters

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

For example:

```
occfg.exe -a "SLOT 5 Port 1" -pq
OneConnect Network Config (10.4.164.0)
Copyright 2011 Emulex

Emulex OneConnect OCe14102-UX-D 2-port PCIe 10GbE CNA (SLOT 5 Port 1)
   Display all properties.

   [Registry] Class of Service (802.1p) = 1 (Auto Priority Pause)
   [Driver]   Class of Service (802.1p) = 1 (Auto Priority Pause)

   [Registry] Encapsulated Task Offload = 1 (Enabled)
   [Driver]   Encapsulated Task Offload = 1 (Enabled)

   [Registry] Enhanced Transmission Selection = 0 (Disabled)
   [Driver]   Enhanced Transmission Selection = 0 (Disabled)

   [Registry] Flow Control = 3 (Rx & Tx Enabled)
   [Driver]   Flow Control = 0 (Disabled)

   [Registry] IPv4 Checksum Offload = 3 (Rx & Tx Enabled)
   [Driver]   IPv4 Checksum Offload = 3 (Rx & Tx Enabled)

   [Registry] Interrupt Moderation = 4 (Adaptive (default))
   [Driver]   Interrupt Moderation = 4 (Adaptive (default))
```

```
[Registry] Large Send Offload V1 (IPv4) = 1 (Enabled)
[Driver]   Large Send Offload V1 (IPv4) = 1 (Enabled)

[Registry] Large Send Offload V2 (IPv4) = 1 (Enabled)
[Driver]   Large Send Offload V2 (IPv4) = 1 (Enabled)

[Registry] Large Send Offload V2 (IPv6) = 1 (Enabled)
[Driver]   Large Send Offload V2 (IPv6) = 1 (Enabled)

[Registry] Maximum Number of RSS Processors = <not set>
[Driver]   Maximum Number of RSS Processors = <not set>

[Registry] Maximum Number of RSS Queues = 6
[Driver]   Maximum Number of RSS Queues = 6

[Registry] Maximum RSS Processor Number = <not set>
[Driver]   Maximum RSS Processor Number = <not set>

[Registry] Network Address = <not set>
[Driver]   Network Address = <not set>

[Registry] NetworkDirect = 1 (Enabled)
[Driver]   NetworkDirect = 1 (Enabled)

[Registry] NetworkDirect MTU = 1024 (1024)
[Driver]   NetworkDirect MTU = 1024 (0x400) (1024)

[Registry] Packet Size = 9014 (9014)
[Driver]   Packet Size = 9014 (0x2336) (9014)

[Registry] Performance Tuning = 0 (Maximum Performance)
[Driver]   Performance Tuning = 0 (Maximum Performance)

[Registry] Preferred NUMA Node = <not set>
[Driver]   Preferred NUMA Node = <not set>

[Registry] RSS Base Processor Group = <not set>
[Driver]   RSS Base Processor Group = <not set>

[Registry] RSS Base Processor Number = <not set>
[Driver]   RSS Base Processor Number = <not set>

[Registry] RSS Max Processor Group = <not set>
[Driver]   RSS Max Processor Group = <not set>

[Registry] RSS Profile = 1 (Closest Processor)
[Driver]   RSS Profile = 1 (Closest Processor)

[Registry] Receive Buffers = 896
[Driver]   Receive Buffers = 1280 (0x500)

[Registry] Receive CPU = <not set>
[Driver]   Receive CPU = <not set>
```

```
[Registry] Receive Side Scaling = 1 (Enabled)
[Driver]   Receive Side Scaling = 1 (Enabled)

[Registry] Recv Segment Coalescing (IPv4) = 1 (Enabled)
[Driver]   Recv Segment Coalescing (IPv4) = 1 (Enabled)

[Registry] Recv Segment Coalescing (IPv6) = 1 (Enabled)
[Driver]   Recv Segment Coalescing (IPv6) = 1 (Enabled)

[Registry] SR-IOV = 0 (Disabled)
[Driver]   SR-IOV = 0 (Disabled)

[Registry] TCP Checksum Offload (IPv4) = 3 (Rx & Tx Enabled)
[Driver]   TCP Checksum Offload (IPv4) = 3 (Rx & Tx Enabled)

[Registry] TCP Checksum Offload (IPv6) = 3 (Rx & Tx Enabled)
[Driver]   TCP Checksum Offload (IPv6) = 3 (Rx & Tx Enabled)

[Registry] Transmit = 1 (Enabled)
[Driver]   Transmit = 1 (Enabled)

[Registry] Transmit Buffers = 256 (256)
[Driver]   Transmit Buffers = 256 (0x100) (256)

[Registry] Transmit CPU = <not set>
[Driver]   Transmit CPU = <not set>

[Registry] Transmit Side Scaling = 1 (Enabled)
[Driver]   Transmit Side Scaling = 0 (Disabled)

[Registry] UDP Checksum Offload (IPv4) = 3 (Rx & Tx Enabled)
[Driver]   UDP Checksum Offload (IPv4) = 3 (Rx & Tx Enabled)

[Registry] UDP Checksum Offload (IPv6) = 3 (Rx & Tx Enabled)
[Driver]   UDP Checksum Offload (IPv6) = 3 (Rx & Tx Enabled)

[Registry] VLAN Identifier (802.1q) = 102
[Driver]   VLAN Identifier (802.1q) = 102 (0x66)

[Registry] Virtual Machine Queues = 1 (Enabled)
[Driver]   Virtual Machine Queues = 1 (Enabled)

[Registry] Wake On LAN = 1 (Enabled)
[Driver]   Wake On LAN = 0 (Disabled)
```

3.2.3.7 Using Interactive Mode

The `occfg.exe` program also supports interactive mode with a set of menus. To start this utility in interactive mode, perform these steps:

1. Run `occfg.exe` from a command console.

   A list of adapters displays on which to operate.

2. Type either a number of the list or a substring from any part of the name (for additional information, see Section 3.2.3.2, Selecting an Adapter).

The program prompts for an operation, such as modifying or querying a parameter value.

3. Follow the prompt.

   The program provides a list of available registry parameters to modify or query.

4. Type either the number of the corresponding option or a substring in the parameter name. The substring must uniquely identify the parameter or occfg displays all potential options.
5. To apply the parameters, select the menu item to exit and reload the drivers. Pressing [Ctrl+ C] at any point might leave modifications in the registry, but the driver does not use the new parameters until it is reloaded.

3.2.3.8 Parameter Help

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

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

```
RSS:
Receive Side Scaling (RSS) scales receive processing over multiple CPUs in
parallel. This scaling typically improves application performance; however, it
tends to increase CPU usage on low end machines.
RSS is only supported on two primary adapters per device. It will appear disabled
for additional PCI functions in blade server configurations.
RSS requires Windows 2008 and later.
Registry Key: *RSS
Default Value : 1 (Enable)
Valid Values :
   0 = Disable
   1 = Enable
```

3.2.4 Extracting the occfg.exe Utility on Nano Server

Use the `occfg.exe` utility to display, or configure, parameters for the network functions on Emulex OneConnect adapters.

NOTE This installation process assumes that you have a Nano Server system booted up and running, as well as a system from which a remote PowerShell connection can be established to manage it.

To extract and use the `occfg.exe` utility on Nano Server, perform these steps:

1. Download the Ethernet driver installer kit from the following link:
   http://www.avagotech.com/support/emulex/windows-server-2016
2. In an elevated PowerShell prompt, navigate to the directory where the downloaded ethernet driver kit is located, and run the following commands to extract the files:

   ```
   brcmdrvr-nic-<VERSION>.exe /q2 extract=2
   ```

   The `occfg.exe` for utility Nano Server is extracted to the \Documents folder. For example:

   ```
   C:\Users\Administrator\Documents\Emulex\Drivers\NIC-<VERSION>\x64\
   xx\occfg.exe
   ```

3. Using a remote PowerShell connection, create the directory on the Nano Server machine to which you want the `occfg.exe` utility copied. For example, you can create a directory under the C: drive of the Nano Server machine as follows:

   ```
   md C:\<Directory Name>
   ```

4. In an elevated PowerShell prompt, navigate to the directory to which the Ethernet driver kit was extracted, and copy its contents to the Nano Server system using the following commands:

```
$ip = "<Nano server IP Address>"
$s = New-PSSession -ComputerName $ip -Credential ~\Administrator
copy -ToSession $s -Path <full path to occfg.exe from step 2> -Destination
<Full path to Directory created in step 3> -Force
```

For example (typed all on one line):

```
copy -ToSession $s -Path C:\Users\Administrator\Documents\Emulex\
Drivers\NIC-<VERSION>\x64\xx\occfg.exe -Destination C:\temp -Force
```

5. Using a remote PowerShell connection, navigate to the directory you created in step 3.

3.2.4.1 Using the occfg.exe Utility

The `occfg.exe` utility for Nano Server is a non-interactive mode program. You must specify a full command. For example:

To get all registry parameters of all Emulex adapters, perform this step:

```
-p
```

To set a parameter on all Emulex adapters, perform this step:

```
occfg.exe -a Emulex -s "Receive Side Scaling=1"
```

To set multiple parameters on one adapter, perform this step:

```
occfg.exe -a "Local Area Connection 23" -s "Flow=3,Receive Side Scaling=0"
```

The `occfg.exe` utility supports the following commands:

Table 5 occfg.exe Commands

Command	Description
`-a str,[str]`	Selects all adapters with any of the given strings in the connection or device name. Use "@" to select all Emulex adapters.
`-s name=v,[name=v]`	Sets the parameter's value and reloads the devices.
`-g name,[name]`	Gets the parameter value.
`-r`	Skips reloading the driver when setting a parameter.
`-f`	Forces a driver reload.
`-l`	Lists the available adapters and exits the utility.
`-x`	Resets all parameters to their default values.
`-p`	Displays all registry parameter values.
`-q`	Displays all driver parameter values.
`-h`	Displays help text for all parameters.

3.2.5 Using SR-IOV with Emulex Devices

This section describes how to use SR-IOV with Emulex devices.

3.2.5.1 Advisory

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

This advisory highlights a use case where a "rogue" (non-Emulex) digitally signed driver is installed by the system administrator on a virtual machine. It is then possible for that rogue driver to cause problems with an OCe11100-series networking adapter. While there are many benefits to using SR-IOV with virtualized workloads, these benefits must be weighed against the potential risks in doing so.

Notes

- The operating system comes with an Emulex inbox driver. Use the Emulex out-of-box driver.
- For a list of supported drivers and adapters, refer to the latest Windows Drivers release notes, which are available for download from the documents and downloads area of the Broadcom website.
- SR-IOV is not supported with RoCE configurations.
- SR-IOV is not supported with UMC.
- SR-IOV is not supported on 1Gb/s adapters.
- SR-IOV is supported only on OCe11000-series and OCe14000-series adapters in NIC mode installed on Windows Server 2012 and Windows Server 2012 R2 with an installed Emulex NIC driver.
- The driver supports the following virtual functions for the following adapter families:
  - OCe11100-series adapters support a maximum of 24 virtual functions per port.
  - OCe14000-series adapters support a maximum of:
    - 2-port 10Gb Ethernet: 31 virtual functions per physical function
    - 4-port 10GbE: 31 virtual functions per physical function
    - 1-port 40GbE: 63 virtual functions per physical function
    - 2-port 40GbE: 63 virtual functions per physical function

3.2.5.2 Server BIOS Configuration

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

- Enable "Virtualization", such as Intel VT-x or AMD-V. This is required for any virtual machine.
- Explicitly enable SR-IOV in the system BIOS. The specific name for this option varies between vendors. For instance, it might be called Intel VT-d (Virtualization Technology for Direct I/O), AMD-Vi (AMD I/O Virtualization Technology), or IOMMU.

3.2.5.3 Emulex PXESelect Configuration for SR-IOV

The Emulex OCe11000-series adapters require enabling firmware support for SR-IOV within the Emulex PXESelect BIOS. Refer to the *Emulex Boot for NIC, iSCSI, and FCoE Protocols for OneConnect Adapters User Guide* for more information.

To enable firmware support in the PXESelect BIOS, perform these steps:

1. Press [Ctrl+ P] during the Emulex PXE Select splash screen as the server boots.

   A screen appears showing global options.
2. Set the following options to use SR-IOV:
   - Advanced mode = Enable
   - Universal Multichannel (UMC) = Disable
3. Save the settings and enable SR-IOV for each PCI function.

   The server reboots after this modification.

3.2.5.4 SR-IOV Server Validation

Use the following Microsoft PowerShell commands to determine if your server is capable of SR-IOV.

- `Get-NetAdapterSriov`

- `Get-VmHost`
- `Get-VmNetworkAdapter`
- `Get-VmSwitch`

Refer to the Microsoft documentation for additional information.

NOTE Early SR-IOV-capable chipsets had errors that might prevent SR-IOV from operating in Windows Server 2012 and Windows Server 2012 R2. The PowerShell command `Get-VmHost | fl *` includes `IovSupportReasons` that indicates if the chipset suffers from this issue.

3.2.5.4.1 Enabling SR-IOV on Unqualified Servers

If Windows Server 2012 or Windows Server 2012 R2 detects a problem with the system I/O remapping hardware, you might still be able to use SR-IOV by explicitly enabling SR-IOV in the registry using `IovEnableOverride`.

NOTE Only use this procedure for trusted virtual machines.

CAUTION Use the registry editor at your own risk. Using the registry editor can cause serious issues that might require you to reinstall the computer's operating system. Broadcom cannot guarantee that issues resulting from changes you make to the registry can be repaired. Make a backup of your registry before making any changes.

3.2.5.4.2 Backing Up and Editing the Registry

To back up and edit the registry, perform these steps:

1. Create a system restore point.
2. Open the registry editor by running `regedit.exe` at the command prompt.
3. Select the hive (the top-level key) and export it to a `.reg` file.
4. Save the `.reg` file to a location off of the server as a precaution.
5. Navigate to:

   `HKLM\Software\Microsoft\Windows NT\CurrentVersion\Virtualization`
6. Create a DWORD type entry named `IovEnableOverride`.
7. Set the value of `IovEnableOverride` to 1.
8. Reboot the system.

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

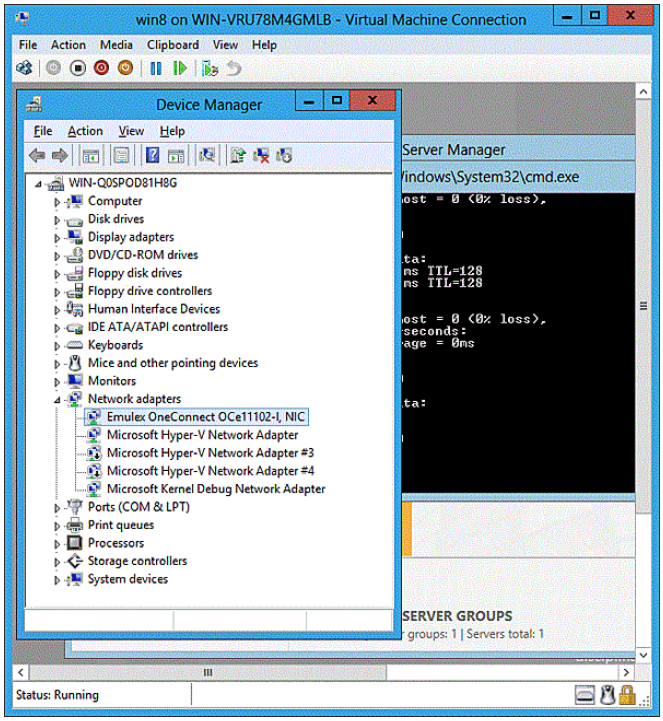
3.2.5.5 Verifying the Driver Version

To verify that the device driver meets the minimum requirements, perform these steps:

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

   The Device Manager opens (see Figure 5.)

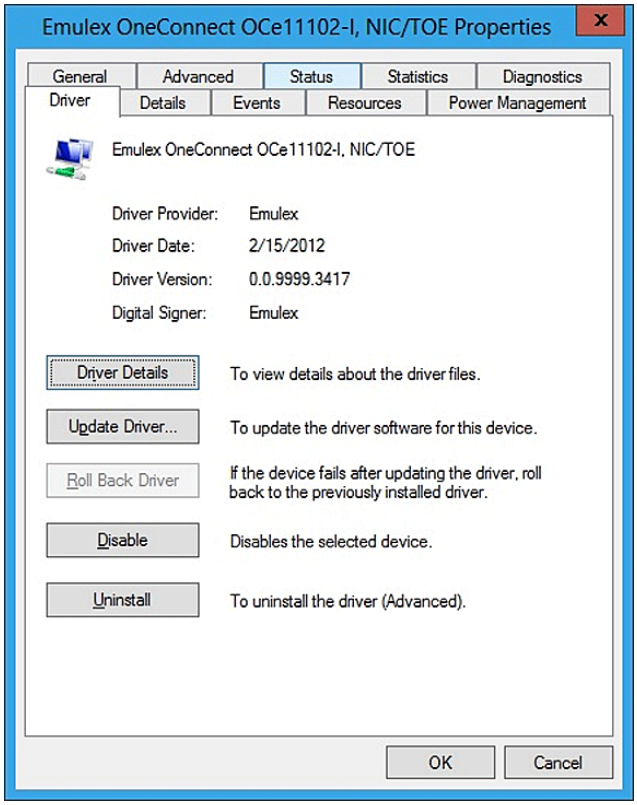
Figure 5 Device Manager for Windows Server 2012



3. Open the Network Adapters item, find the Emulex device and right-click it.
4. Select **Properties** from the context menu.

   The Properties dialog box opens showing the Driver page (see Figure 6). The Driver page contains the driver version number.

Figure 6 Emulex NIC Driver Properties Page



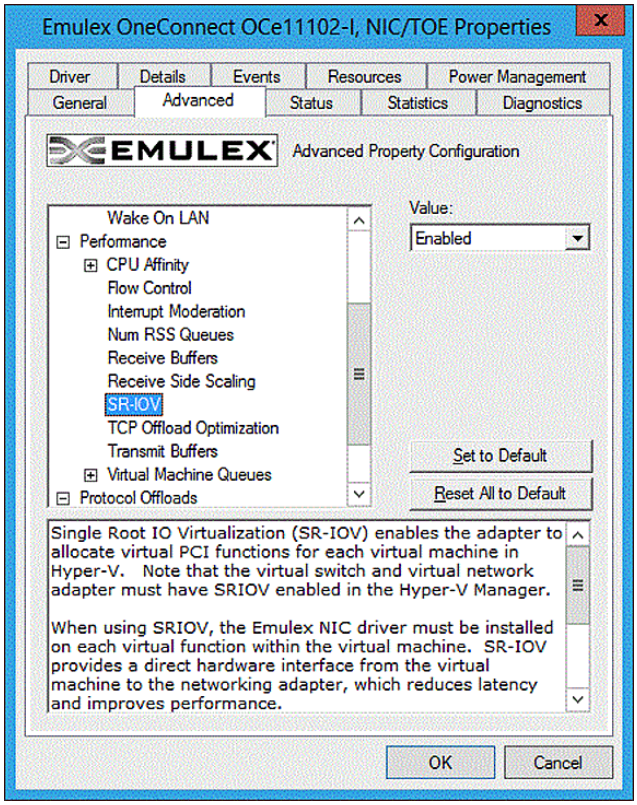
5. Click **Driver Details**.
   A window opens that displays the driver name.

3.2.5.6 Enabling SR-IOV in the Emulex Device

To enable SR-IOV in the Emulex device, perform these steps:

1. Select **Server Manager>Dashboard>Tools>Computer Management**.
2. Click **Device Manager**.
   The Device Manager opens (see Figure 5).
3. Open the Network Adapters item, find the Emulex device and right-click it.
4. Select **Properties** from the context menu.
   The Properties dialog opens (see Figure 6).
5. Click the **Advanced** tab.
   The Advanced Property Configuration page opens (see Figure 7).

Figure 7 Emulex NIC Advanced Properties Page



6. Select **SR-IOV** from the list and select **Enabled** from the **Value** drop-down list.

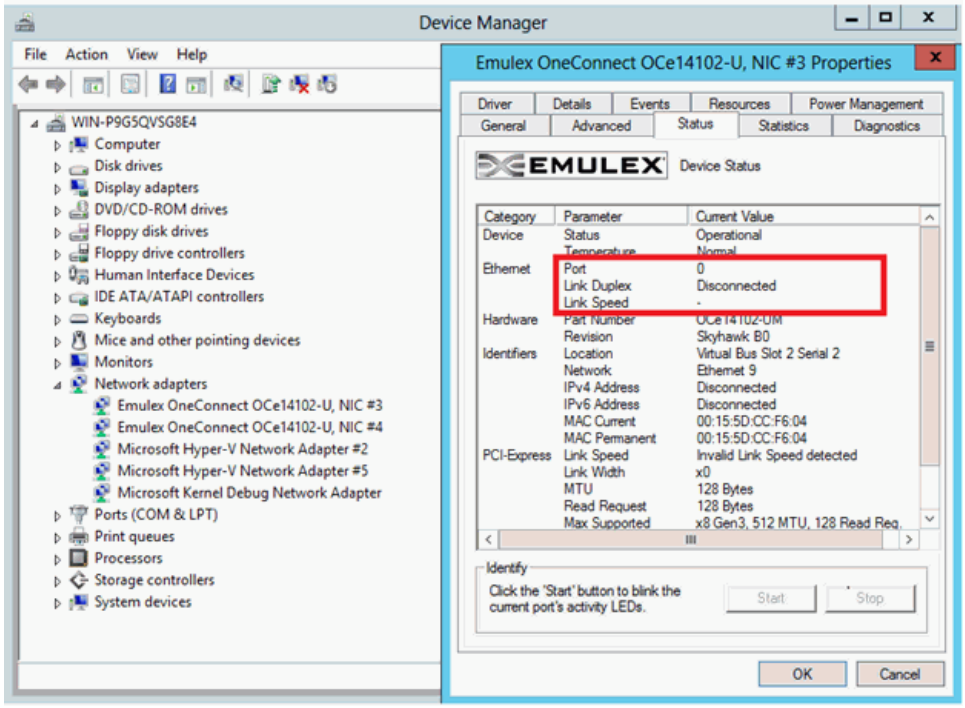
NOTE You must configure Hyper-V to create an SR-IOV-enabled virtual machine. See the Microsoft Hyper-V documentation for additional information.

3.2.5.7 SR-IOV Link Behavior

Link Status information for the Virtual Function (VF) is not available from the Network Connection Manager of the virtual machine. The Link Status of the VF adapter is displayed in the Status tab of the Device Manager Advanced Property page for the corresponding adapter.

If the link status of the VF is up, then it is shown as Connected along with the Link Speed. When PF link is disconnected, the VF adapter will be shown as Disconnected.

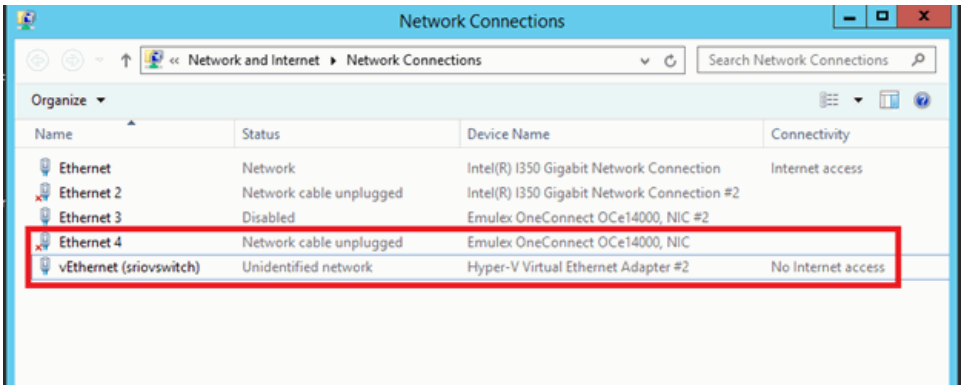
Figure 8 PF Link Disconnected, VF Adapter Disconnected



When a VF adapter link is down, the communication between the VF and its corresponding Physical function (PF) adapter is still possible. The default synthetic data path provided by hypervisor will be used instead of the VF data path of the adapter. In this case, statistics counters of the VF adapter are not incremented.

The virtual switch connected to the PF is always shown as link up regardless of the PF link status. The link status of the switch is controlled by the operating system and not the PF driver. Figure 9 depicts the Network Connections Panel when the PF link is disconnected, but the SR-IOV enabled virtual switch link is shown as connected.

Figure 9 Network Connection Panel PF Link Disconnected



3.2.5.8 Hyper-V

The Hyper-V role must be added using the Server Manager. After the Hyper-V role is added, enable SR-IOV in the Hyper-V Manager by doing one of the following:

- Creating the virtual switch
- Creating each virtual NIC

Refer to the Microsoft documentation for additional information.

NOTE Ensure that SR-IOV is enabled on the server and on the Emulex adapter prior to configuring the Hyper-V virtual switch.

The Windows Server 2012 and Windows Server 2012 R2 servers treat SR-IOV as an offload; an active-active team with virtual function and an emulated adapter, which means each Emulex SR-IOV adapter is accompanied by a fully functional, emulated NIC. The emulated NIC is named "Microsoft Virtual Network Adapter," and the TCP/IP stack is bound only to this device.

After the Emulex driver is loaded, the Emulex SR-IOV virtual function is used for all unicast receive and transmit traffic. The emulated NIC handles multicast and broadcast traffic. If SR-IOV is disabled, the Emulex adapter is removed from the virtual machine, and all traffic automatically uses the emulated NIC. This technology allows live migration of virtual machines if using SR-IOV.

NOTE If multiple adapters are added to the virtual machine, use MAC addresses to map the Emulex Network adapter to the corresponding Microsoft Virtual Network adapters.

3.2.5.9 Verifying SR-IOV

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

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

3.2.5.9.1 Verifying SR-IOV from the Virtual Machine

To verify SR-IOV from within the virtual machine perform these steps:

1. From within the virtual machine, select **Server Manager>Dashboard>Tools>Computer Management**.
2. Click **Device Manager**.

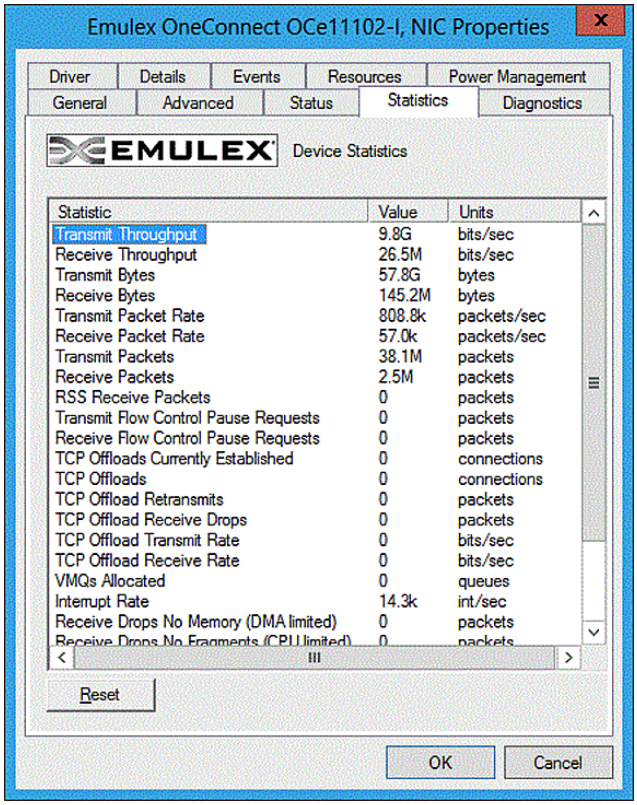
   The Device Manager opens (see Figure 5).
3. Open the Network Adapters item, select the Emulex device and right-click.
4. Select **Properties** from the context-menu.

   The Properties dialog opens showing the Driver page (see Figure 6).

   NOTE The Emulex adapter might initially appear as a "Network Adapter" before the driver is loaded.
5. Select the **Statistics** tab (see Figure 10).

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

Figure 10 Emulex NIC Statistics Properties Page



3.2.5.9.2 Verifying SR-IOV from the Host Hyper-V Server

To verify SR-IOV from the Host Hyper-V Server, perform these steps:

1. From the Device Manager, open the Network Adapters item, select the Microsoft Hyper-V Network adapter and right-click.
2. Select **Properties** from the context-menu.

   The Hyper-V Network adapter Properties dialog box opens showing the Driver page.
3. Select the **Statistics** tab (see Figure 10).
4. From the Statistics tab, locate the "Virtual Functions Allocated" item.

   "Virtual Functions Allocated" shows the count of currently enabled virtual functions.

NOTE The Microsoft Powershell command "Get-NetAdapterSriovVf" lists each SR-IOV virtual function. Refer to the Microsoft documentation for additional information.

3.2.6 Configuring NVGRE for the OCe14000-series Adapters

Network virtualization using NVGRE is a network virtualization method that uses encapsulation and tunneling to create large numbers of VLANs for subnets that can extend across dispersed data centers and layer 2 (the data link layer) and layer 3 (the network layer). The purpose is to enable multi-tenant and load-balanced networks that can be shared across on-premises and cloud environments.

NVGRE was designed to solve issues caused by the limited number of VLANs that the IEEE 802.1Q specification enables, which are inadequate for complex virtualized environments, and make it difficult to stretch network segments over the long distances required for dispersed data centers.

3.2.6.1 Setup

Hardware resources:

- Two host servers
- Virtual machines (two per Hyper-V host)
- One 10GbE or 40GbE Ethernet switch
- Two OCe14000-series adapters (one per host server)

Software resources:

- Windows Server 2012 with Hyper-V
- Windows Server 2012 on the Virtual Machines
- Add and Remove PowerShell Policy Scripts for each host server

To set up NVGRE, perform these steps:

1. On the Hyper-V hosts and peer, change the execution policy to allow PowerShell scripts to run:
   — Set-Execution Policy unrestricted –Force.
   — Run HostRegedit (run this once on the Hyper-V host only).

   This sets the registry key to use VMQs and allows remote PowerShell scripts to be run on the host.
2. Set up non-blank administrator passwords on the peer to run remote PowerShell scripts.
3. Copy the NIC driver to C : \ driver on the Hyper-V Hosts.

3.2.6.2 Configuration

To create a VM, perform these steps:

1. Use a 10 GB disk image size and 1 GB RAM.
2. Install Windows Serer 2012 RTM.
3. Turn off automatic administrator login by using `control userpasswords2`.
4. Turn off the Windows Firewall.
5. Create a vswitch for NVGRE (for example, vport0).
6. Create a vswitch for non-NVGRE (normal traffic).
7. Expose a NIC interface into the VM for each of the vSwitches:
   — Make sure "Enable virtual machine queue" is selected under Network Adapter > Hardware Acceleration of the VMs.
   — Record the MAC addresses located under the Network Adapter > Advanced Features. These will be used in the add/remove policy scripts.
8. Rename the Network Connection name being used for NVGRE to WNVNIC (for example: **Control Panel>Network and Internet>Network Connections**, rename **Ethernet 3** to **WNVNIC**).
9. Set up an NVGRE script.

3.2.6.2.1 Setting up an NVGRE Script

The following sample script is required for network virtualization:

Example of a Script Adding the NVGRE Tunnel Between Two Hosts

```
# Add the locator records for Blue subnet
Get-NetVirtualizationLookupRecord | where {$_.VirtualSubnetID -eq "5001"} |
Remove-NetVirtualizationLookupRecord;
```

```
New-NetVirtualizationLookupRecord -VirtualSubnetID "5001" -CustomerAddress
"10.0.0.5" -ProviderAddress "192.x.x.x" -MACAddress "060600000005" -Rule
"TranslationMethodEncap"
New-NetVirtualizationLookupRecord -VirtualSubnetID "5001" -CustomerAddress
"10.0.0.7" -ProviderAddress "192.x.x.x" -MACAddress "060600000007" -Rule
"TranslationMethodEncap"

# Add the customer route records for Blue subnet
Get-NetVirtualizationCustomerRoute | where {$_.VirtualSubnetID -eq "5001"} |
Remove-NetVirtualizationCustomerRoute;
New-NetVirtualizationCustomerRoute -RoutingDomainID
"{11111111-2222-3333-4444-000000005001}" -VirtualSubnetID "5001"
-DestinationPrefix "10.0.0.0/24" -NextHop "0.0.0.0"  -Metric 255

##########################################################################################
################################
# Red Virtual Network Information
#
# RoutingDomainID="{11111111-2222-3333-4444-000000006001}"
# VirtualSubnetID=6001
# (Both RDID and VSID are defined by administrators, MUST be unique in the
datacenter)
#
# [Customer Addresses]
# VM Name      Host    VSID  CA        PA             MAC                DefaultGW
#
-------------------------------------------------------------------------------
-----
# Red1         Host1   6001  10.0.0.5  192.x.x.x   08-08-00-00-00-05  10.0.0.1
# Red2         Host2   6001  10.0.0.7  192.x.x.x   08-08-00-00-00-07  10.0.0.1
#
# [Customer Routes]
# DestPrefix   NextHopGW  Note
#
-------------------------------------------------------------------------------
----
# 10.0.0.0/24  0.0.0.0    Onlink route for Red subnet

# Add the locator records for Red subnet
Get-NetVirtualizationLookupRecord | where {$_.VirtualSubnetID -eq "6001"} |
Remove-NetVirtualizationLookupRecord;

New-NetVirtualizationLookupRecord -VirtualSubnetID "6001" -CustomerAddress
"10.0.0.5" -ProviderAddress "192.x.x.x" -MACAddress "080800000005" -Rule
"TranslationMethodEncap"
New-NetVirtualizationLookupRecord -VirtualSubnetID "6001" -CustomerAddress
"10.0.0.7" -ProviderAddress "192.x.x.x" -MACAddress "080800000007" -Rule
"TranslationMethodEncap"

# Add the customer route records for Red subnet
Get-NetVirtualizationCustomerRoute | where {$_.VirtualSubnetID -eq "6001"} |
Remove-NetVirtualizationCustomerRoute;
```

```
New-NetVirtualizationCustomerRoute -RoutingDomainID
"{11111111-2222-3333-4444-000000006001}" -VirtualSubnetID "6001"
-DestinationPrefix "10.0.0.0/24" -NextHop "0.0.0.0"  -Metric 255

#
# [2] Configure the Host Provider Addresses and Routes required for this setup
#
# [Host PA Address & Route information required by the VM policy]
#
# Host    Hostname        {PA's}          {VM:VirtualSubnetID} ==> Set on the
VMNetworkAdapter on eash host
#
--------------------------------------------------------------------------------
-----
# Host1   example-host1  192.x.x.x    {Blue1:5001, Red1:6001}
# Host2   example-host2  192.x.x.x    {Blue2:5001, Red2:6001}

# [2-1] Host1
#
# (a) Configure Provider Address and Route:
#     Get the interface, assign the PA and the default route
Get-NetVirtualizationProviderAddress | where {$_.ProviderAddress -eq "192.x.x.x"}
| Remove-NetVirtualizationProviderAddress;

$iface = Get-NetAdapter $WNVNIC
New-NetVirtualizationProviderAddress -InterfaceIndex $iface.InterfaceIndex
-ProviderAddress "192.x.x.x" -PrefixLength 24

# (b) Set VirtualSubnetID on the VM network port
Get-VMNetworkAdapter "Blue1" | where {$_.MacAddress -eq "060600000005"} |
Set-VMNetworkAdapter -VirtualSubnetID 5001;
Get-VMNetworkAdapter "Red1"  | where {$_.MacAddress -eq "080800000005"} |
Set-VMNetworkAdapter -VirtualSubnetID 6001;
```

Example of a Script Removing the NVGRE Tunnel Between Two Hosts

```
################################################################################
################################
# Blue Virtual Network Information
#
# RoutingDomainID="{11111111-2222-3333-4444-000000005001}"
# VirtualSubnetID]=5001
# (Both RDID and VSID are defined by administrators, MUST be unique in the
datacenter)
#
# [Customer Addresses]
# VM Name      Host   VSID  CA        PA            MAC               DefaultGW
#
--------------------------------------------------------------------------------
# Blue1        Host1  5001  10.0.0.5  192.x.x.x   06-06-00-00-00-05  10.0.0.1
# Blue2        Host2  5001  10.0.0.7  192.x.x.x   06-06-00-00-00-07  10.0.0.1
#
# [Customer Routes]
# DestPrefix   NextHopGW  Note
```

```
#
-------------------------------------------------------------------------------
--
# 10.0.1.0/24  0.0.0.0    Onlink route for Blue subnet

# Remove the locator records for Blue subnet
Get-NetVirtualizationLookupRecord | where {$_.VirtualSubnetID -eq "5001"} |
Remove-NetVirtualizationLookupRecord;

# Remove the customer route records for Blue subnet
Get-NetVirtualizationCustomerRoute | where {$_.VirtualSubnetID -eq "5001"} |
Remove-NetVirtualizationCustomerRoute;

###############################################################################
################################
# Red Virtual Network Information
#
# RoutingDomainID="{11111111-2222-3333-4444-000000006001}"
# VirtualSubnetID=6001
# (Both RDID and VSID are defined by administrators, MUST be unique in the
datacenter)
#
# [Customer Addresses]
# VM Name      Host    VSID  CA        PA             MAC                DefaultGW
#
-------------------------------------------------------------------------------
-----
# Red1         Host1   6001  10.0.0.5  192.x.x.x   08-08-00-00-00-05  10.0.0.1
# Red2         Host2   6001  10.0.0.7  192.x.x.x   08-08-00-00-00-07  10.0.0.1
#
# [Customer Routes]
# DestPrefix   NextHopGW  Note
#
-------------------------------------------------------------------------------
-----
# 10.0.0.0/24  0.0.0.0    Onlink route for Red subnet

# Remove the locator records for Red subnet
Get-NetVirtualizationLookupRecord | where {$_.VirtualSubnetID -eq "6001"} |
Remove-NetVirtualizationLookupRecord;

# Remove the customer route records for Red subnet
Get-NetVirtualizationCustomerRoute | where {$_.VirtualSubnetID -eq "6001"} |
Remove-NetVirtualizationCustomerRoute;

#
# [2] Configure the Host Provider Addresses and Routes required for this setup
#
# [Host PA Address & Route information required by the VM policy]
#
# Host    Hostname        {PA's}          {VM:VirtualSubnetID} ==> Set on the
VMNetworkAdapter on eash host
```

```
#
--------------------------------------------------------------------------------
-----
# Host1   example-host1  192.x.x.x    {Blue1:5001, Red1:6001}
# Host2   example-host2  192.x.x.x    {Blue2:5001, Red2:6001}

# [2-1] Host1
#
# (a) Configure Provider Address and Route:
#     Get the interface, assign the PA and the default route
Get-NetVirtualizationProviderAddress | where {$_.ProviderAddress -eq "192.x.x.x"}
| Remove-NetVirtualizationProviderAddress;

# (b) Set VirtualSubnetID on the VM network port
Get-VMNetworkAdapter "Blue1" | where {$_.MacAddress -eq "060600000005"} |
Set-VMNetworkAdapter -VirtualSubnetID 0;
Get-VMNetworkAdapter "Red1"  | where {$_.MacAddress -eq "080800000005"} |
Set-VMNetworkAdapter -VirtualSubnetID 0;
```

3.2.6.2.2 Policy Script Information

Verify or modify the following list of components for your specific conditions:

- Network connection name of the NIC attached to the virtual switch
- Binding network connection to the Microsoft filter driver (ms_netwnv)

CAUTION Do not change the filter driver name.

- VM Names
- VM MAC Addresses
- VM SID (Virtual Machine Subnet ID)
- HOST IP Address (Provider Address = Physical NIC IP Address for NVGRE)
- VM IP Address (Customer Address = IP Address on each vNIC per VM)
- Subnet Masks

3.2.6.2.3 Verification and Troubleshooting

Verify that the host provider addresses can access each other. If a ping fails, perform these steps:

1. Shutdown all VMs on all of the hosts.
2. Remove the policies.
3. Remove the vport0 virtual switch.
4. Reboot the hosts.
5. Create a vport0 switch.
6. On each VM, add a new adapter, and use the vswitch name that you just created.

NOTE Do not share the NIC with the host operating system.

7. Apply the policies.
8. Power up the VMs.
9. Send a ping to the provider addresses.

`MS_NETWNV.sys` must only be bound to WNVNIC (physical NIC). If `ms_netwnv` is bound to a vswitch or Hyper-V adapter, unbind it from the host server. For example:

```
Disable-netadapterbinding vEthernet* -ComponentID ms_netwnv
```

Verify that all of the IP addresses and MAC addresses used in the add and remove policy scripts match the VM IP's/MACs using `ipconfig /all`.

Create firewall rules to allow ICMP (ping) packets:

1. New-NetFirewallRule –DisplayName "Allow ICMPv4-In"
   `–Protocol ICMPv4`
2. New-NetFirewallRule –DisplayName "Allow ICMPv4-Out"
   `–Protocol ICMPv4 –Direction Outbound`

PowerShell commands shown w/"WNVNIC" used as the network connection name:

```
Get-VM
Get-vmswitch
Get-vmnetworkadapter –VMName * | fl
vmname,switch*,macaddress,ipaddress*,virtualsub*
Get-netvirtualizationprovideraddress
Get-netvirtualizationlookuprecord
Get-netvirtualizationcustomerroute
Get-netadapter wnvnic
Get-netadapterbinding –componentID ms_netwnv
Get-netadapterencapsulatedpackettaskoffload wnvnic
Get-netadapteradvancedproperty wnvnic
Disable-netadatperSriov wnvnic
Disable-netadapterEncapsulationPacketTaskOffload wnvnic
Get-help *-NetVirtualization*
Get-netadapterstatistics wnvnic
```

Under the **Host Device Manage> – Network Adapters > Emulex Statistics** tab, check the following:

- VMQs Allocated
- Tunnels Allocated
- Tenants Allocated

To verify that all of the IP addresses and MAC addresses used in the add and remove policy scripts match the VM IP's/MACs, perform these steps:

1. Send a ping.
   a. Launch Policy Scripts on each host.
   b. Send a ping using the `–t` option.
      Pings will respond.
   c. Run the Remove Policy script on one host server.
      Pings will stop responding.
   d. Add policies.
      Pings will respond.
2. Change the VM IP Address.
   For example: change 10.0.0.5 on one host and 10.0.0.7 on the other.

Without NVGRE, you would not be able to use the same IP address as the other VM.

Ensure that you see the WNVNIC interface using the PowerShell command get-netadapter.

If you are unable to see the NVGRE Ethernet connection, make sure that the Hyper-V's control panel/network connections property and advanced property windows are all closed. The Control Panel renames the NDI/Params registry keys and causes the policy scripts to be inoperative.

3.2.6.2.4 Cleaning Up Outdated Network Adapter Data

To clean up outdated network adapter data, perform these steps:

1. Start a win32 console window (command prompt).
2. From the command prompt, type

   ```
   Set devmgr_show_nonpresent_devices=1
   ```
3. Start `devmgmt.msc`.
4. In the device manager console, go to the view menu and select **Show hidden devices**.
5. Open the network devices tree view.
6. Uninstall all Emulex entries.
7. Rescan for hardware changes.
8. Uninstall Emulex devices until they are not recognized.
9. Install the new driver.

3.2.7 Configuring RoCE for the OCe14000-Series Adapters

Both Windows SMB Direct and Windows NetworkDirect, which are included as part of the Windows operating system, are required for RoCE. Windows Server 2012 or Windows Server 2012 R2 is required to use the RoCE features on RoCE-capable adapters.

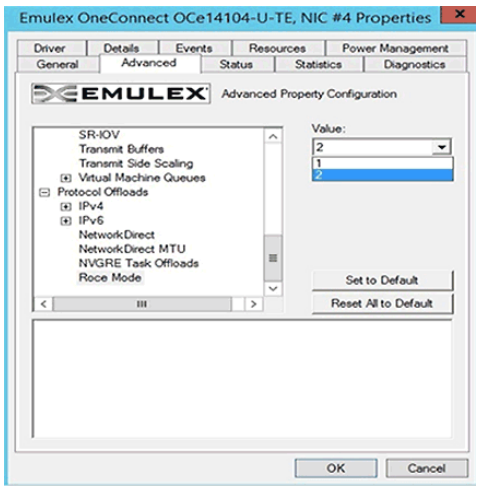
NOTE RoCE is not supported if UMC is enabled and RoCE configurations are not supported with SR-IOV.

3.2.7.1 Configuring Routable RoCE

Routable RoCE is enabled by default. A drop-down menu is added in property page. RoCE Mode 2 is Routable (default setting) and RoCE Mode 1 is Native RoCE.

NOTE Only OCe14000B-series adapters support Routable RoCE.

Figure 11 Advanced Tab RoCE Mode Selected



To enable or disable routeable RoCE, perform these steps:

1. From the **Advanced** tab of the **Network Property** page, select **RoCE Mode** (Figure 11).
2. From the Value menu, choose **2** to enable routable RoCE, or choose **1** to select Native RoCE.

3.2.7.2 Enabling the RoCE Profile on the Client-Side

The RoCE profile can be enabled by using one of the following:

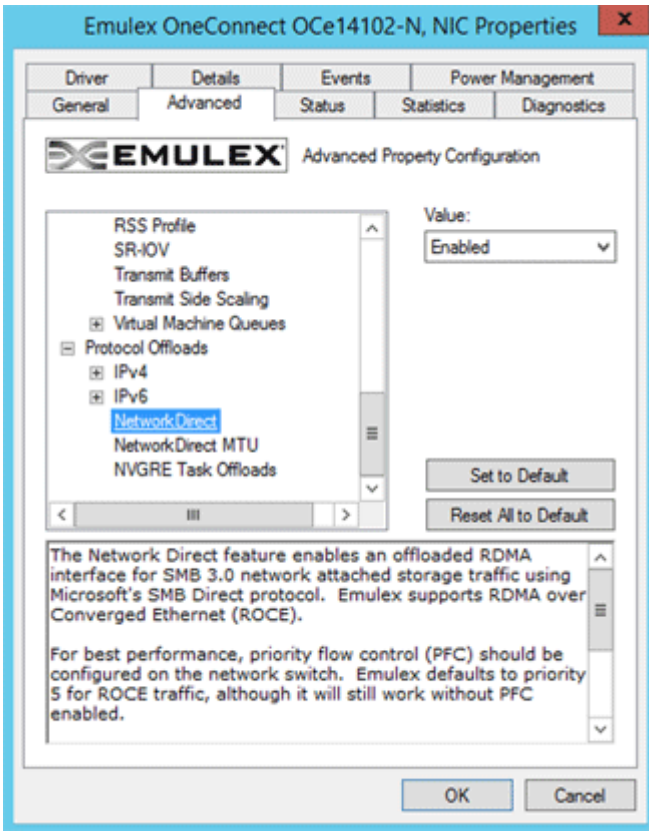
- PXESelect BIOS. Refer to the *Emulex Boot for NIC, iSCSI, and FCoE Protocols for OneConnect Adapters User Guide* for details on the PXESelect BIOS utility.
- OneCommand CNA Manager GUI. Refer to the *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide* for information about enabling the RoCE profile using the OneCommand CNA Manager GUI.
- OneCommand CNA Manager CLI. Refer to the *Emulex OneCommand CNA Manager Command Line Interface for OneConnect Adapters User Guide* for information about enabling the RoCE profile using the OneCommand CNA Manager CLI.

3.2.7.3 Confirming That the RoCE Profile Is Enabled

To confirm that the RoCE profile is enabled, use one of the following methods:

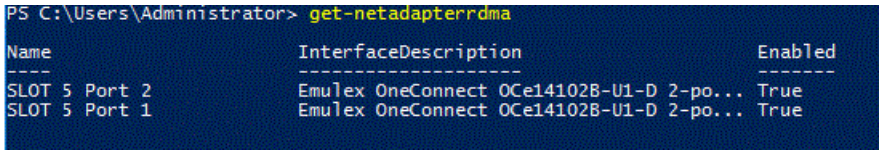
- In the **Advanced** tab of the **Network Property** page, ensure that **NetworkDirect** is enabled (see Figure 12).

Figure 12 Advanced Property Configuration - RoCE - Enabled



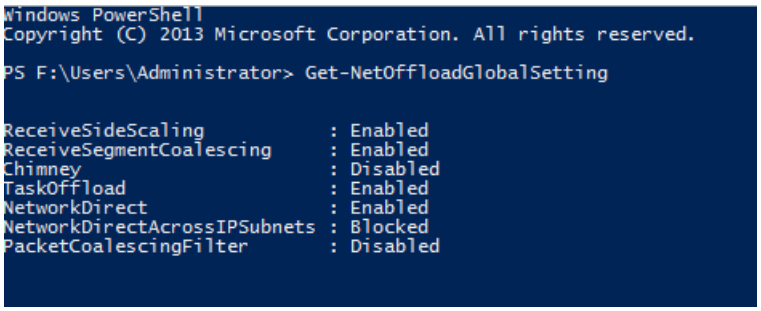
- By using a PowerShell script:
  - Get-NetAdapterRDMA (Figure 13)

Figure 13 Get-NetAdapterRDMA - RoCE-Enabled)



— Get-NetOffloadGlobal (Figure 14)

Figure 14 Get-NetOffloadGlobal - RoCE-Enabled



If the profile is correct and NetworkDirect is enabled, you will see active NetworkDirect listeners on IP addresses (port 445) assigned to the NICs using
`netstat -xan`.

3.2.7.4 Configuring Host Mode RDMA

To configure Host Mode RDMA, perform these steps:

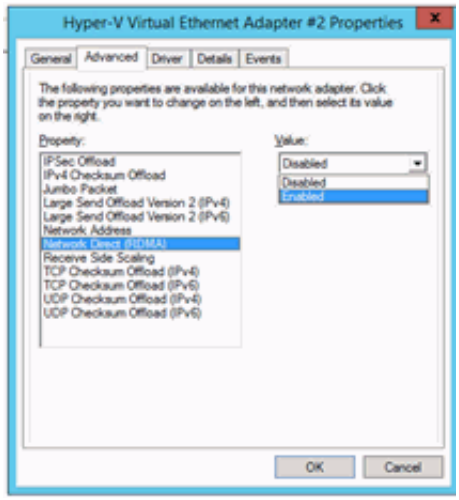
1. Load the driver.

   NOTE Do not add a virtual local area network (VLAN) to the network adapter advanced settings prior to creating the vSwitch.

2. Using the Hyper-V Manager Virtual Switch Manager, create a new external virtual switch and attach it to the Emulex adapter.
3. Using either PowerShell or Virtual Switch Manager, assign any required VLAN IDs to the management operating system.
4. From the Device Manager of the host operating system, select the **Advanced** page of the Hyper-V Virtual Ethernet Adapter, and enable Network Direct (RDMA).

   NOTE Network Direct (RDMA) is disabled by default.

Figure 15 Network Direct (RDMA) Enabled



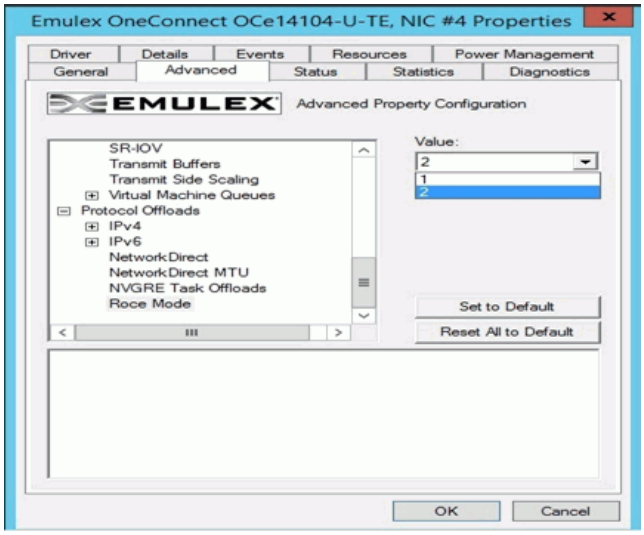
3.2.7.5 Configuring Routable RoCE

Routable RoCE is enabled by default and is only supported on OCe14000B and OCm14000B adapters.

To configure routable RoCE, perform these steps:

1. From the OneConnect **Advanced tab**, choose **RoCE Mode**.

Figure 16 Routable RoCE Enabled (default)



2. From the **Value** menu, choose **2** for Routable RoCE (default setting) or **1** for Native RoCE.
3. Click **OK**.

3.2.7.6 VxLAN

3.2.7.6.1 Enabling Encapsulated Task Offload and VxLAN UDP

You can enable Encapsulated Task Offload and modify the VxLAN user datagram protocol (UDP) destination port number using the OneConnect Advanced tab or by using Powershell commands.

NOTE VxLAN is enabled and NVGRE is enabled by default. To enable VxLAN, you do not need to disable NVGRE Encapsulated Task Offload in the adapter property page before creating a vswitch with encapsulation.

To enable Encapsulated Task Offload (Enabled is the default), perform these steps:

1. From the OneConnect Advanced tab, select **Encapsulated Task Offload,** or **VxLAN Encapsulated Task Offload**.
2. Set the value to **Enabled**.

Figure 17 Encapsulated Task Offload (The default is enabled)

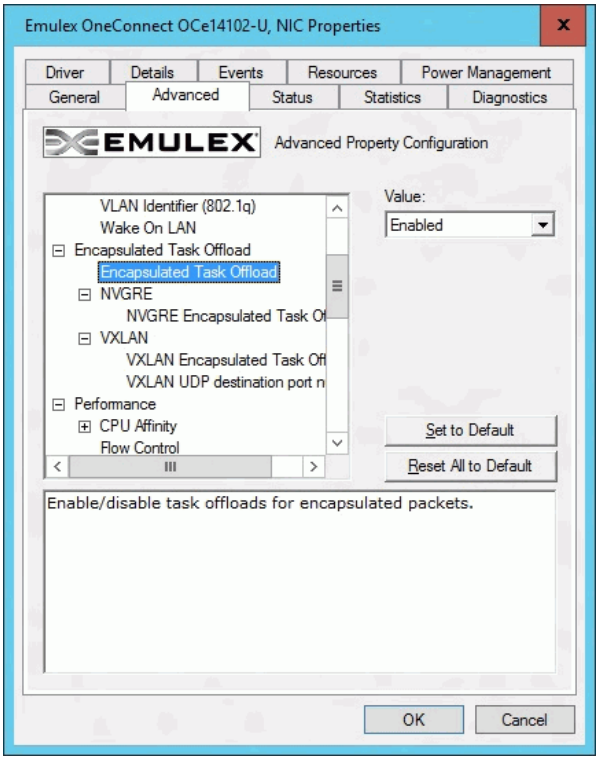
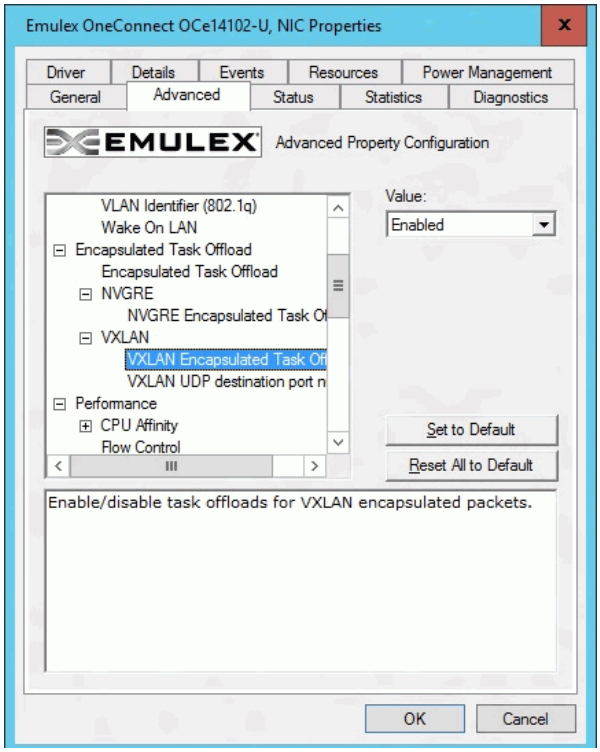


Figure 18 VXLAN Encapsulated Task Offload (The default is enabled)

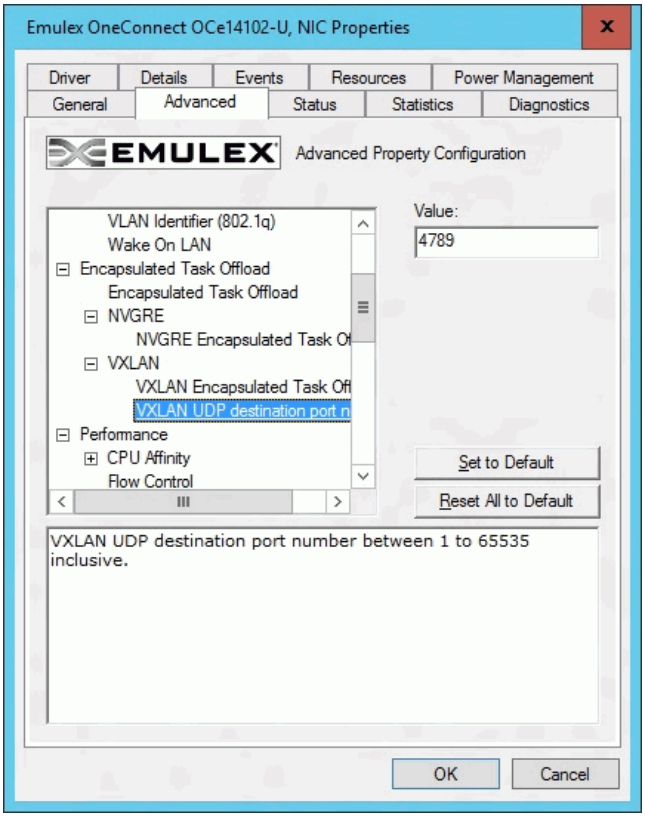


3. Click **OK**.

To modify the VxLAN UDP destination port number parameter, perform these steps:

1. From the OneConnect **Advanced** tab, select **VxLAN UDP destination port number**.

Figure 19 VXLAN UDP Destination Port Number



2. Enter the port number you want to use.
3. Click **OK**.

3.2.7.6.2 Using Powershell Commands

You can use PowerShell commands to set the values for the following subkeys:

- `EncapsulatedPacketTaskOffload`
- `EncapsulatedPacketTaskOffloadVxlan`
- `VxlanUDPPortNumber`

For `EncapsulatedPacketTaskOffload` and `EncapsulatedPacketTaskOffloadVxlan` subkeys, a value of 0 disables the feature and a value of 1 enables the feature.

The `VxlanUDPPortNumber` subkey has a default value of 4789 and a value range of 1–65535.

The following commands are also available:

- `Disable-NetAdapterEncapsulatedPacketTaskOffload`
- `Enable-NetAdapterEncapsulatedPacketTaskOffload`
- `Get-NetAdapterEncapsulatedPacketTaskOffload`
- `Set-NetAdapterEncapsulatedPacketTaskOffload`

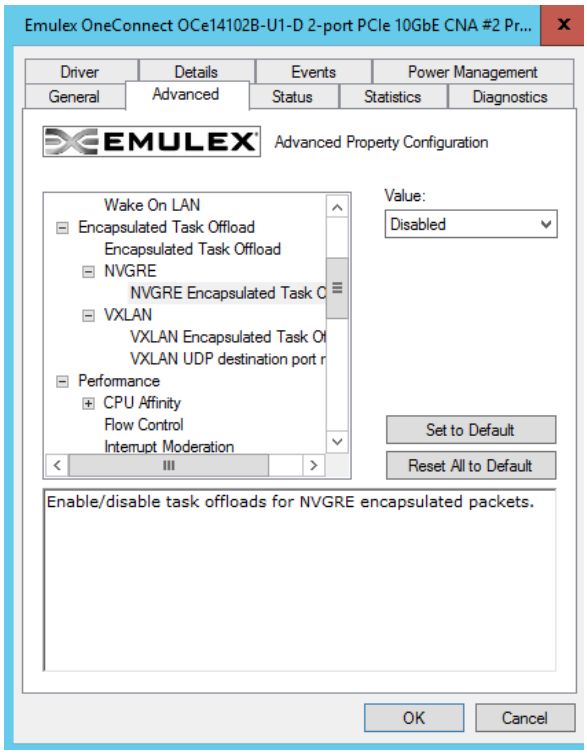
3.2.7.7 Enabling NVGRE Encapsulated Task Offload

NOTE VxLAN is enabled and NVGRE is enabled by default. To enable NVGRE, you must disable VxLAN Encapsulated Task Offload in the adapter property page before creating a vswitch with encapsulation.

To enable NVGRE Encapsulated Task Offload (Enabled is the default) using the OneConnect Advanced tab, perform these steps:

1. From the OneConnect **Advanced** tab, choose **NVGRE Encapsulated Task Offload**.

Figure 20 NVGRE Encapsulated Task Offload (The default is disabled)



2. From the **Value** menu, choose **Enabled**.
3. Choose **VxLAN Encapsulation Task offload**.
4. From the **Value** menu, choose **Disabled**.

NOTE VxLAN must be disabled when NVGRE is enabled.

5. Click **OK**.

3.2.7.8 Configuring Packet Direct

By default, the advanced driver parameter Packet Direct is enabled. When you attach a virtual machine (VM) to a Packet Direct enabled switch, the virtual machine NIC (vmNIC) automatically takes the Packet Direct path.

To use Packet Direct, perform these steps:

1. On the host powershell prompt, limit the number of Event Contexts ECs based on the logical processors present in the system.

```
Set-NetAdapterRss –Name "SLOT 6 Port 2" -Enabled $True –MaxProcessors 2
```

2. Create a PacketDirect switch using the powershell command.

```
New-VMSwitch -Name PDSwitch -NetAdapterName "SLOT 6 Port
2"-EnablePacketDirect $True
```

3. Enable the Azure Virtual Filtering Platform (VFP) switch extension on the switch.

```
Enable-VMSwitchExtension -VMSwitchName PDSwitch -Name "Windows Azure VFP
Switch Extension"
```

4. Attach the vmNIC of the VM to packet direct switch (PDSwitch) in the VM settings.

5. Start the VM.
6. Run the `disableVFP.ps1` script on the host using powershell prompt.
7. Open perfmon, add packetdirect counters, and confirm they are updated.

3.2.7.8.1 The disableVFP.ps1 script

Ensure the file extension is `.ps1` before you run attempt to run the script.

```
param(
    [string]$switchName = $(throw "please specify a switch name")
    )

$switches = Get-WmiObject -Namespace root\virtualization\v2 -Class
Msvm_VirtualEthernetSwitch
foreach ($switch in $switches) {
              if ( $switch.ElementName -eq $switchName) {
                            $ExternalSwitch = $switch
                            break
              }
}

$vfpCtrlExe = "vfpctrl.exe"
$ports = $ExternalSwitch.GetRelated("Msvm_EthernetSwitchPort",
"Msvm_SystemDevice", $null, $null, $null, $null, $false, $null)
foreach ($port in $ports) {
#if ($port.ElementName -eq "Dynamic Ethernet Switch Port")
#{
    $portGuid = $port.Name
    echo "Disabling VFP on port: " $portGuid
    & $vfpCtrlExe  /disable-port /port $portGuid
#}
}
```

3.2.7.8.2 disableVFP.ps1 script Errors

The `disbleVFP.ps1 script` generates the following expected errors.

```
PS C:\Users\Administrator> .\disableVFP.ps1 vsw_sh2
Disabling VFP on port:
1EDA9CAD-89F3-4019-98EC-96C8E077DD2F
ERROR: failed to execute disable-port
Error (1): Incorrect function.
Disabling VFP on port:
2632C1DA-97A3-4091-A694-60A2F57079F6
ERROR: failed to execute disable-port
Error (1): Incorrect function.
Disabling VFP on port:
B1CDC522-4EA0-4F3C-8323-90E8C70F3FC0
Command disable-port succeeded!
```

3.2.7.9 Using SMB Direct with NetworkDirect

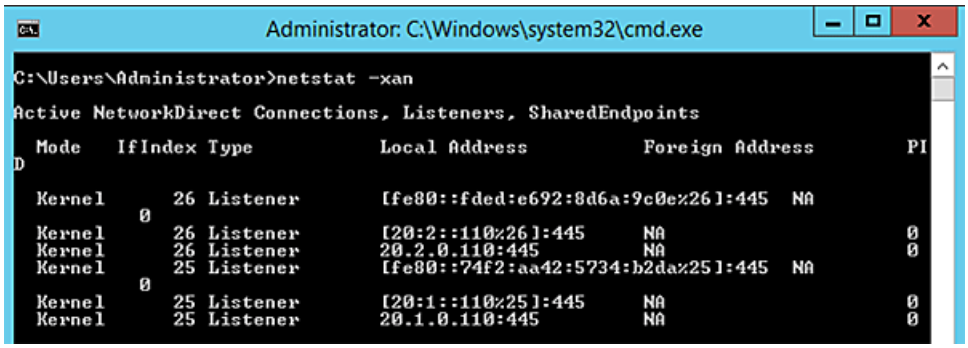
Because RoCE is supported in Windows using SMB Direct with NetworkDirect, SMB Direct and NetworkDirect must be configured correctly.

To use SMB Direct with NetworkDirect, perform these steps:

1. From the **Advanced** tab of the **Network Interface Properties** page, enable the `NetworkDirect` parameter.
2. Set the NetworkDirect MTU. Use a NetworkDirect MTU of 4096 bytes for OCe14400-series adapters.

   NOTE The NetworkDirect MTU affects only RoCE traffic, but the NIC traffic still uses the "Packet Size" MTU. An SMB Server accepts an incoming connection request from an SMB Client when the NetworkDirect MTU on the server is at least as large as the NetworkDirect MTU on the initiating client.
3. Use the `netstat -xan` command to enumerate the active NetworkDirect connections and listeners (Figure 21). A NetworkDirect enabled driver creates listeners on any configured IPv4 or IPv6 addresses, and the link-local IPv6 address. SMB Direct listeners listen on port 445.

Figure 21 Active NetworkDirect Connections and Listeners



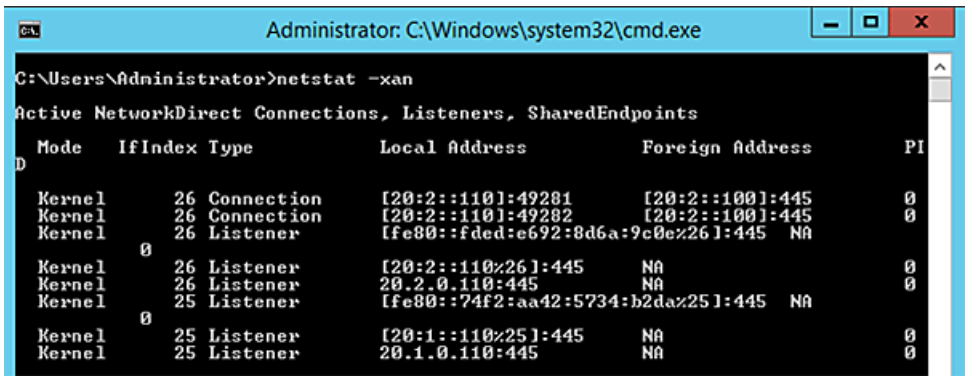
3.2.7.10 Mapping the RoCE-Enabled Client to the Server-Side Storage

Using an available network share with the proper permissions configured, open an SMB share from the Windows Run command or from the command prompt, by typing:

```
net use [devicename:*] [\\computername\sharename]
```

By default, this creates two RDMA connections per SMB Direct-enabled network interface on a particular server (see Figure 22). Each SMB Direct connection maps to an RDMA queue pair. Both the client and server must negotiate support for SMB Direct. If available, each TCP connection is offloaded to an RDMA queue pair.

Figure 22 SMB Share - Two RDMA Connections Per RDMA-Enabled Network Interface



The PowerShell command, `Get-NetAdapterStatistics` (see Figure 23), shows the RDMA Statistics which indicate the number of failed connection attempts.

Figure 23 Get-NetAdapterStatistics

```
PS C:\Users\Administrator> get-netadapterrdma

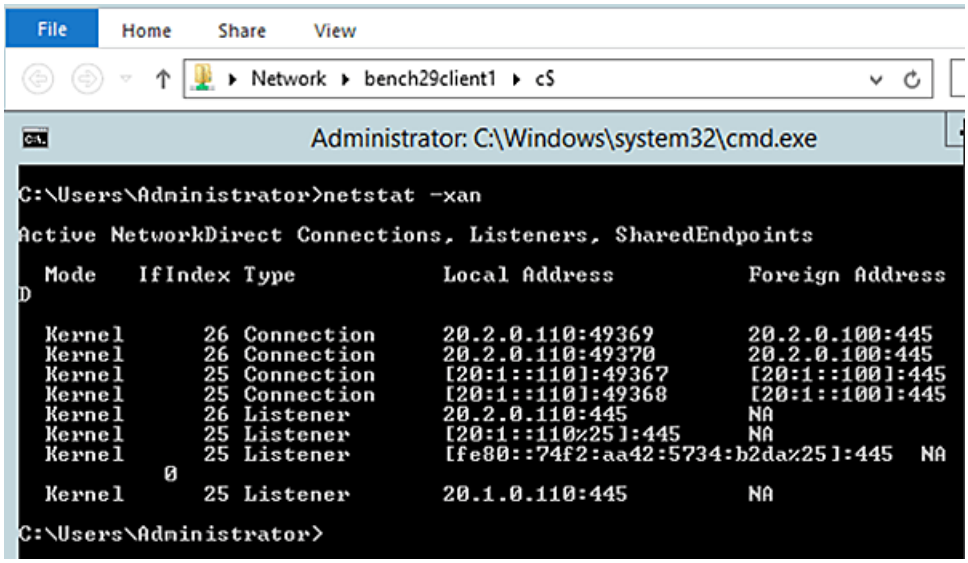
Name                    InterfaceDescription                    Enabled
----                    --------------------                    -------
SLOT 5 Port 2           Emulex OneConnect OCe14102B-U1-D 2-po... True
SLOT 5 Port 1           Emulex OneConnect OCe14102B-U1-D 2-po... True
```

3.2.7.11 SMB Multichannel

For each SMB session, SMB multichannel establishes two SMB Direct connections to a particular server by default. It also makes use of multiple RDMA-capable NIC interfaces, if available.

Opening a file share from a 2-port OCe14000-series adapter (both ports are RDMA-enabled) connected back-to-back to another 2-port OCe14000-series adapter (both ports are also RDMA-enabled) creates two SMB Direct connections per interface (see Figure 24).

Figure 24 Two SMB Direct Connections Per Interface



Use SMB multichannel constraints to limit an SMB connection to specific network interfaces. For example, if you have a 1GbE interface meant for a management path and one or more 10GbE interfaces meant for the RDMA traffic, you can restrict the RDMA traffic to the faster 10GbE interfaces by setting SMB multichannel constraints.

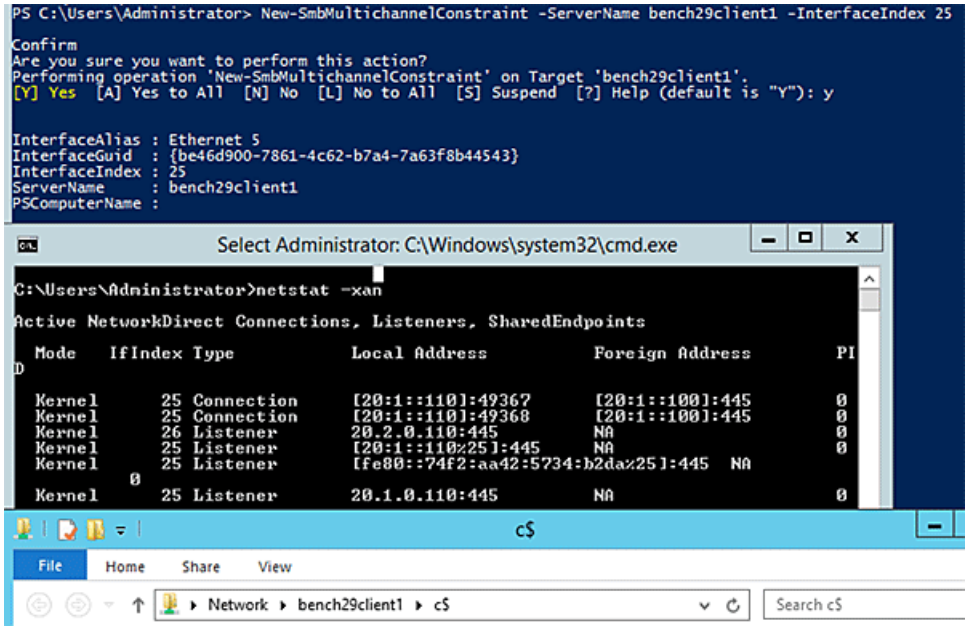
You can specify SMB multichannel constraints on the SMB client using certain RDMA network interfaces to access a particular server.

Figure 25 shows that an SMB multichannel constraint has been added to specify that only InterfaceIndex 25 must be used to connect to the server `bench29client1`.

Two SMB Direct connections are established on the RDMA Network Interface with IfIndex: 25. None are established on the other OCe14000-series adapter ports; for example, the RDMA Network Interface with IfIndex: 26 (20.2.0.110).

NOTE On Windows Server 2012, depending on the path used to open the previous SMB sessions to the SMB server, you might notice multiple instances of connections being created while establishing a single SMB session to the server. When adding or removing adapters from your system, you must readjust the multichannel constraints for optimal performance.

Figure 25 SMB Multichannel Constraint



3.2.7.12 SMB Direct Resource Usage

This section describes how to use SMB Direct resources.

3.2.7.12.1 Active Connections

Active connections describe the connections that a client makes to a server. Passive connections describe the connections that a server allows the client to complete.

The maximum number of active connections per port for an adapter are required if setting the `ConnectionCountPerRdmaNetworkInterface` parameter. Use Table 6 on page 79 to determine the correct number of maximum active connections based on your OCe14000-series adapter and operating system.

NOTE Passive and active connection limits for OCe14000-series adapters followe following equation:
511 - (2 *:

- Active mode – See Table 6 on page 79.
- Passive mode – Maximum passive connections are calculated using the following equation:
  511 - (2 x number of ports)
- Passive and Active mode – Active and passive connection counts together must not exceed 511 - (2 x number of ports)

NOTE Passive and active connection limits for OCe14000B-series adapters follow:

- Active mode – See Table 6 on page 79.
- Passive mode
- 40Gb 1 port is 254 passive connections.
- 10Gb 2 port is 253 passive connections per port.
- 10Gb 4 port is 126 passive connections per port.

Table 6 shows the maximum number of SMB Direct active (client mode) connections that can be initiated on Windows Server 2012 and Windows Server 2012 R2 using the OCe14000-series adapters.

Table 6 SMB Direct Active Connections (Client Mode) Per Port for OCe14000-Series Adapters

Adapter Type	Windows Server 2012	Windows Server 2012 R2
1-port 40GbE adapter	31	15
2-port 10GbE adapter	15	7
4-port 10GbE adapter	7	3
1-port 40Gb adapter (OCe14000B-series adapters)	47	23
2-port 10Gb adapter (OCe14000B-series adapters)	31	11
4-port 10Gb adapter (OCe14000B-series adapters)	11	6

The Maximum Queue Pair counts on a 1-port, 2-port, and 4-port OCe14000-series adapter are shown in Figure 26, Figure 27, and Figure 28.

Figure 26 Resource Counts on a 1-Port 10GbE or 40GbE OCe14000-Series Adapter

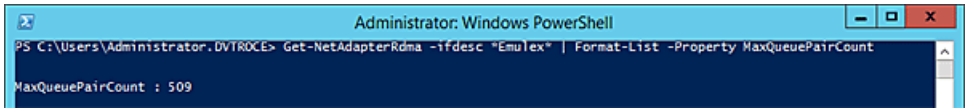


Figure 27 Resource Counts on a 2-Port 10GbE OCe14000-Series Adapter

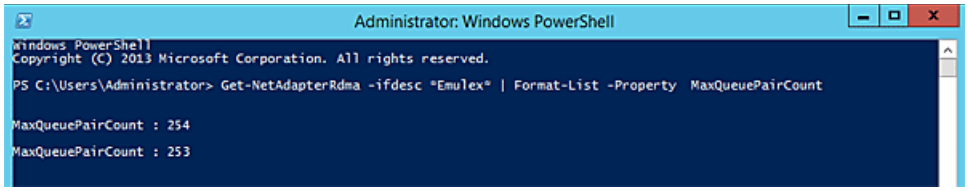
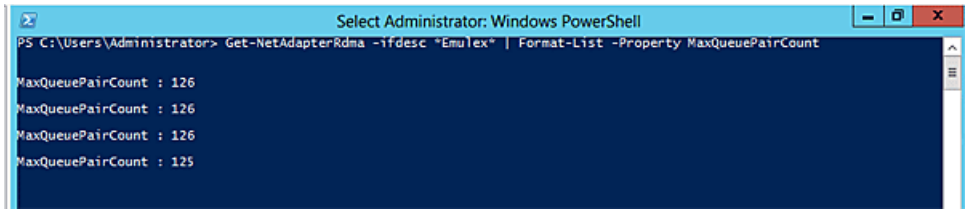


Figure 28 Resource Counts on a 4-Port 10GbE OCe14000-Series Adapter



SMB Direct does not take these adapter-reported limits into consideration when creating SMB Direct connections because the driver might still get a request to create more queue pairs or memory regions than are supported.

NOTE OCe14000-series adapters will fail to create a memory region or queue pair if they exceed the limits of what is supported. The per-port maximum active connections and the maximum passive connections cannot exceed the MaxQueuePairCount. See Table 6 on page 79, for more information.

For Windows Server 2012 and Windows Server 2012 R2, all existing RDMA connections between the particular client-server pair, on which the failure occurred, are torn down and re-created. After a certain number of unsuccessful retries, SMB traffic falls back to TCP/IP.

The following event warning message is placed in the Windows System Log under the source `be2net` which indicates the adapter is running out of resources:

```
"The Adapter ran out of resources while creating the requested number of SMB
Direct connections. Please reduce the connection count to a supported value."
```

3.2.7.12.1.1 Setting RoCE Parameters

You can set the RoCE adapter parameters using OneCommand CNA Manager, PowerShell scripts, or by using the Network Interface Property page.

NOTE Refer to the *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide* for more information on using the OneCommand CNA Manager GUI application to configure RoCE, or refer to the *Emulex OneCommand CNA Manager Command Line Interface for OneConnect Adapters User Guide* for information on using the OneCommand CNA Manager CLI to configure RoCE.

The parameters in Table 7 can be modified from the Network Interface Property page (see Figure 7).

Table 7 Parameters for RoCE

Parameter	Description
`NetworkDirect`	This parameter enables an offloaded RDMA interface for SMB 3.0 network-attached storage traffic using Microsoft's SMB Direct protocol.
`NetworkDirect MTU`	This parameter configures the maximum transmission unit (frame size) for RoCE traffic. **NOTE** For optimal performance, Emulex recommends setting the MTU size to 4096.
`RoCE Mode`	This parameter controls the RoCE mode as native or routable. ■ 1= Native RoCE ■ 2 = Routable RoCE

The parameters in Table 8 on page 80 can be viewed using the Statistics Property Page (see Figure 10).

Table 8 RoCE Statistics for Viewing

Parameter	Description
`RoCE QP Allocated`	Indicates the number of established queue pairs for RoCE.
`RoCE Transmit Throughput`	The transmit data rate of RoCE traffic.
`RoCE Receive Throughput`	The receive data rate of RoCE traffic.

3.2.7.13 QoS Concepts Related to RoCE

This section describes QoS concepts and how to configure QoS for RoCE.

3.2.7.13.1 Priority Groups

It is advisable to split traffic into two or more priority groups; one priority group for RoCE and other groups for non-RoCE traffic. Many of the cluster applications use TCP and RoCE traffic simultaneously. Some of them use TCP for establishing connections and share connection-specific information. As a result, it is important to allocate enough bandwidth (greater than 1%) to non-RoCE (NIC traffic) to avoid a slow connection establishment rate and starvation of

NIC traffic. Work conserving behavior ensures that each priority group gets enough bandwidth. Based on this behavior, non-RoCE traffic must be given sufficient bandwidth; ideally 30 percent to 70 percent.

3.2.7.13.2 L2 Flow Control

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

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

3.2.7.14 Configuring QoS for RoCE

If configuring QoS for RoCE, consider the following points:

- A limited QoS configuration is available through the OneCommand CNA Manager application.
- A single traffic class group for RoCE exists per port.
- A single RoCE priority exists in PFC mode.
- Bandwidth allocation for priority groups is supported.

NOTE

- The Windows NIC driver does not support the Microsoft DCB/QoS API.
- PowerShell commands cannot be used to configure QoS-related parameters for the RoCE profile.

3.2.7.14.1 RoCE Adapter Side

With DCBX enabled, the switch settings are used for PFC. Ensure that the switch settings match the adapter default Priority 5 used for RoCE and PFC.

NOTE PFC is enabled by default in OCe14000-series adapters.

3.2.7.14.2 Switch Side

For information on switch-side configurations, see Section E, RoCE Switch Support.

3.2.7.14.3 OCe14000-Series Adapter Defaults

If using the OCe14000-series adapters for RoCE functionality, the following defaults apply:

- Adapter boot time
  - PFC is disabled on all the ports in the NIC, and the RoCE profile is enabled.
  - Generic Pause is enabled on all the ports in the NIC and RoCE profiles.
- Back-to-back connection (OCe14000 to OCe14000)
  - PFC is disabled by default.
  - Generic Pause is enabled on the connected port.
- DCBX-enabled switch connection
  - If an OCe14000-series adapter is connected to a DCBX-enabled switch, it shifts the mode from Generic Pause to PFC.
  - An OCe14000-series adapter configures RoCE traffic for Priority 5.
- Manually enable Priority 5 on the switch under a different priority group other than the FCoE/iSCSI/NIC priority group.

NOTE If you do not enable Priority 5 on the switch side, the OCe14000-series adapter maintains its configuration for PFC mode Priority 5. This configuration can result in packet losses, unrecoverable errors, or infinite retries for RoCE traffic.

- DCBX-disabled switch connection
  - If an OCe14000-series adapter is connected to a DCBX-disabled switch, generic pause mode is enabled.

3.2.7.15 Congestion Management Options for RoCE

In addition to QoS settings, the OCe14000-series adapter supports congestion management protocols. The supported modes are Quantized Congestion Notification (QCN) for RoCE ports or Explicit Congestion Notification (ECN) for RoCE v2 ports.

NOTE For each RoCE port, the RoCE traffic must be limited to either RoCE or Routable RoCE (default).

If the port is configured for RoCE, it can support QCN. This feature must be enabled on both the RoCE port and the switch port. The port can be enabled using the OneCommand Management application. Refer to the *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide* for details. The switch must also be enabled to generate QCN packets. QCN is active only if both the RoCE port and the switch port are enabled. The switch port must toggled after enabling QCN from the OneCommand Management application.

If the port is configured for Routable RoCE, it can support ECN. If this feature is enabled on the network switches, the Routable RoCE port will recognize ECN marked packets, send CNPs and react to CNPs. ECN is automatically detected by OCe14000B-series adapters.

3.2.7.16 Performance Considerations

The following recommended settings can improve SMB performance over TCP, including RDMA. However, the configuration must be tuned to provide line rate with TCP network traffic.

1. Disable TCP Autotuning.
2. Set the RDMA MTU size to 4096.
3. Set the NIC MTU size to 9014. The NIC MTU must be greater than the RDMA MTU.
4. On Windows Server 2012 R2, disable the bandwidth throttling option on the SMB client side to improve throughput (this parameter has no effect on Windows Server 2012):

```
Set-SmbClientConfiguration -EnableBandwidthThrottling 0
```

3.2.8 Configuring UMC

NOTE RoCE is not supported if UMC is enabled.

UMC allows you to divide a 10GbE port into multiple physical functions with flexible bandwidth capacity allocation. These functions appear to the operating system and network as separate physical devices.

UMC can be configured on OCe14000-series adapters through the adapter BIOS or the OneCommand CNA Manager application.

- To configure UMC using the adapter BIOS, refer to the *Emulex Boot for NIC, iSCSI, and FCoE Protocols for OneConnect Adapters User Guide.*
- To configure UMC using the OneCommand CNA Manager application, refer to the *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide* or the *Emulex OneCommand CNA Manager Command Line Interface for OneConnect Adapters User Guide*.

Refer to the *Emulex Universal Multichannel Reference Manual* for additional information on UMC.

3.2.8.1 NPar Configuration (Dell Only)

NOTE

- NPar is available only on Dell OCe14000-series adapters.
- If NPar is enabled, RoCE cannot be configured on any function.
- Each partition must have standard NIC properties for stateless offload.

NPar enables you to divide a 10GbE NIC port into multiple PCI functions with flexible bandwidth capacity allocation that appears to the operating system and network as separate NIC ports. For example, a single 10GbE port appears as multiple physical devices showing in PCI configuration space as multiple functions.

3.2.8.1.1 Adapter Configuration

NPar can be configured on OCe14000-series adapters in several ways, including the Emulex adapter driver properties, the adapter BIOS, or using the OneCommand CNA Manager application.

Refer to the *Emulex Boot for NIC, iSCSI, and FCoE Protocols for OneConnect Adapters User Guide* for information on configuring the adapter BIOS. If you are using the OneCommand CNA Manager application, refer to the *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide*.

On the host operating system side, NPar provides up to 16 PCI functions per device using standard PCI configuration space if NParEP is enabled. Four PCI functions can be mapped to a physical port on a 4-port adapter. Eight PCI functions can be mapped to a physical port on an 2-port adapter. Each function or partition is assigned a unique MAC address.

Partitions are available for virtual function assignment and for application segmentation using VLAN or IP subnets. The partitions can be on separate subnets or VLANs.

3.2.8.1.2 NPar Partition Support

- Flexible bandwidth allocation with no changes required for the operating system or BIOS.
- The switch is independent; changes to the external switch are not required.
- NIC teaming is supported.
- The following items are supported on a per-partition basis:
  - Statistics
  - LSO, LRO, RSS, TSO, and MTU
  - Support for NetQueues

3.2.8.2 NPar Considerations

- NPar can use virtual adapters using VLAN tagging per partition.
- NPar can use RSS queuing support per partition.
- DCBX is supported while in NPar mode.
- If iSCSI or FCoE functions are not enabled, they are available as NIC functions.
- Only one iSCSI function is allowed per physical port.
- Only one FCoE function is allowed per physical port.
- The second to fourth functions on a particular port are available for storage protocols if desired, which allows you to configure up to two storage functions.
- If NParEP is disabled, a total of eight functions are available evenly distributed across the ports on the adapter. For example, a 2-port adapter can have four functions per port and a 4-port adapter can have two functions per port.

NOTE A system reboot is required if any of the NPar function mode settings are modified. A reboot is not required for bandwidth and MTU settings modifications.

3.2.9 Network Driver Performance Tuning

This section describes the tuning and configuration of the network drivers.

3.2.9.1 Optimizing Server Hardware and BIOS Configuration

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

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

Some servers might allow memory mirroring or memory sparing, where the total memory is divided in half and each location is stored twice. Memory mirroring and memory sparing provide fault recovery if one memory location detects an error, but they greatly reduce the perceived memory bandwidth of the system.

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

3.2.9.2 Windows Server Network Driver

Table 9 describes ways to use various NIC driver properties and Microsoft Windows properties to performance-tune a system.

Table 9 Windows Server Performance Tuning Situations

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

3.2.9.2.1 Analyzing Performance Issues

Use the Windows Performance Monitor (perfmon) to view statistics for each network device.

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

Table 10 is a partial list of the statistics to use to troubleshoot performance issues. For network performance, all of the counters from the table are useful: Network Interface, TCPv4, IPv4, and Processor.

Table 10 Statistics and Fine Tuning

Situation	Answer/Solution
Network Interface > Packets Received Errors.	If this value is incrementing even a small amount, a physical problem might exist on the network, such as a loose connection or bad cable, which causes CRC errors in Ethernet packets. Find and eliminate the physical problem.
Network Interface > Packets Received Discarded.	If this value is incrementing dramatically, the computer system might be receiving a lot of unsolicited traffic using network resources.
IPv4 > Fragmented Datagrams / sec.	If this value is greater than 0, the computer system is sending or receiving IP fragments. This problem impacts performance. See Section 3.2.9.2.2, Jumbo Packet.
TCPv4 > Segments Retransmitted / sec.	TCP retransmits indicate that packets are being dropped by the receiving system (or in a network switch). Ideally, reduce retransmits to 0.
Processor > % Processor Time.	If CPU usage is high, try to enable all available offloads, such as TCP offload or checksum offloads, and use jumbo packets.

3.2.9.2.2 Jumbo Packet

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

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

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

To set a jumbo packet value, go to the Advanced Properties page in Windows Device Manager. For information on how to configure the options through the Advanced Property page, Section 3.2.2.1, Modifying Advanced Properties.

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

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

Typical cases for using the MTU include the following:

- Server interconnects are typically deployed using jumbo frames. This configuration is the most efficient for high bandwidth server-to-server communication, such as Network Attached Storage, iSCSI, and database transactions.
- Servers connected to client systems that run desktop operating systems typically use standard 1500-byte frames. Most desktop systems do not support jumbo packets.

- Servers that require both high performance server-to-server communication and client access can be configured with jumbo frames with Path MTU Discovery enabled. Path MTU Discovery is enabled by default in the Windows Server, and it allows TCP connections to negotiate the optimal packet size that avoids IP fragmentation.

3.2.9.2.3 Flow Control

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

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

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

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

3.2.9.2.3.1 Examples

Flow control greatly improves the following situations:

- The adapter is installed in a 4x PCIe slot or an underpowered server system.

  If the PCIe bus does not provide 10 Gb/s of throughput due to chipset limitations or the bus width, the adapter cannot maintain 10 Gb/s of incoming receive data. It starts dropping packets quickly. In this situation, it might be beneficial to enable receive flow control in the adapter, and enable flow control in the attached switch for all devices, which helps to slow down the transmitters.
- The adapter transmits to 1GbE devices, especially non-TCP protocol.

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

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

For information on modifying the Flow Control parameter, see "Configuring NIC Driver Options" on page 29.

3.2.9.2.4 NUMA Considerations for Windows Server 2012 R2 and Windows Server 2016

NUMA assignments can affect network performance and CPU efficiency. If your application is not NUMA aware and network traffic is moderate to heavy, the CPU and memory access are managed by the operating system. As a result, the operating system can cross NUMA nodes or your application might be on the same NUMA node as other applications, decreasing your network efficiency. Regardless of whether your application is multi-threaded, and if data is not in parallel, consider the NUMA CPU defaults.

To improve network and CPU performance for heavy network loads under these conditions, you might have to make an appropriate NUMA CPU selection. For example, in Windows Server 2012 R2 and Windows Server 2016, you can use the Task Manager to adjust the `Set Affinity` property to bind the application to a specific NUMA node for maximum network performance and CPU efficiency.

3.2.9.2.5 Checksum Offloading and Large Send Offloading (LSO)

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

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

NOTE On Windows Server 2012 and later, Recv Segment Coalescing is enabled by default. You must disable Recv Segment Coalescing if you want to set the Checksum Offload setting to anything other than enabled.

For information on modifying the CheckSum Offload or Large Send Offload parameter, Section 3.2.1, Configuring NIC Driver Options.

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

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

3.2.9.2.7 Understanding RSS

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

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

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

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

For information on modifying the RSS Queues parameter, see “Configuring NIC Driver Options” on page 29.

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

3.2.9.2.8 Enabling Windows to Use Up to Eight Processors

Windows Server 2008 uses only four processors by default. For the driver to use up to eight processors, the registry must be changed and the system restarted.

CAUTION Use the registry editor at your own risk. Using the registry editor can cause serious issues that might require you to reinstall the computer's operating system. Emulex cannot guarantee that issues resulting from changes you make to the registry can be repaired. Back up your registry before making any changes.

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

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

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

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

3.2.9.3 TCP Offloading (TOE)

NOTE TCP Offloading (TOE) is not supported by OCe14000-series adapters.

Some Emulex adapters and drivers support TCP offload, which provides significant performance improvements. The performance improvements follow:

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

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

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

3.2.9.3.1 TCP Offload Exclusions

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

If an application sends less data than the MSS, the driver, like most TCP stacks, uses a Nagling algorithm. Nagling reduces the number of TCP packets on the network by combining small application sends into one larger TCP packet. Nagling typically reduces the performance of a single connection to allow greater overall performance for a large group of connections.

During Nagling, a single connection might have long pauses (200 ms) between sending subsequent packets, because the driver waits for more data from the application to append to the packet. An application can disable Nagling using the TCP_NO_DELAY parameter. TCP offload does not improve the performance for connections that Nagle, because the performance is intentionally limited by the Nagling algorithm. Telnet and SSH consoles are examples of connections that typically use Nagling.

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

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

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

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

3.2.9.3.2 TCP Offload Optimization Settings

The adapter supports an option for optimizing TCP connection offload characteristics for throughput or latency. This option is available through the Advanced Property Page. See Section 3.2.1, Configuring NIC Driver Options for the TCP Offload Optimization settings.

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

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

Use the Optimize Throughput (default) parameter setting.

3.2.9.3.3 Windows Networking and TOE

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

Installing or activating firewall applications causes no connections to be offloaded by the Windows Server 2008 and Windows Server 2008 R2 network stack. By default, Windows Firewall Services are enabled at operating system

installation time, and they must be explicitly disabled to use TOE. Firewall services can be disabled through the Service Control panel, or the following commands at the command line prompt:

To set firewall services to load on demand:	`sc config MpsSvc start= demand`
To stop firewall services:	`Net stop MpsSvc`
To temporarily disable firewall services:	`netsh advfirewall set all state off`

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

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

3.2.9.3.4 Windows TCP Parameters

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

3.2.9.4 Receive Window Auto Tuning and Compound TCP

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

Performance of some 10GbE stress applications can suffer with these features enabled. In particular, when receive window auto-tuning is enabled, some bi-directional data stream test performance degradation might occur. This situation is due to increased receive performance that adversely affects the same TCP connection's transmit performance.

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

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

3.2.9.5 Interrupt Coalescing

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

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

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

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

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

3.2.9.6 CPU Binding Considerations

Windows applications can set a processor affinity, which binds a program to a particular CPU in a multiple processor computer. Do not manually configure CPU affinity due to the recent additions to the Windows networking stack.

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

The driver uses multiple parallel DPCs that are explicitly assigned to particular CPUs for processing both RSS and TCP offloading tasks. Each TCP connection is assigned to a particular CPU for processing. This situation provides the advantage of assigning CPU affinities by reducing CPU cache misses, without any user configuration.

Explicit processor affinity assignments are not necessary for the driver because the advantages of assigning processor affinities are realized by using RSS. The only reason to experiment with application and interrupt CPU affinity is when performing isolated networking benchmarks.

3.2.9.7 Single TCP Connection Performance Settings

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

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

3.3 iSCSI Driver Configuration

Table 11 lists the user-configurable iSCSI driver options available on Windows Server. It includes a description of the parameters, their default values, and their configuration limits.

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

3.3.1 Configuring iSCSI Driver Options

The OneConnect Windows iSCSI driver parameters can be configured using the Advanced tab of the Device Manager Property Page.

To modify a configuration parameter, perform these steps:

1. Select the Emulex OneConnect iSCSI adapter in the Windows Device Manager under Storage Controllers.
2. Right-click and select **Properties**.

The Device Manager Property page opens.

3. Select the **Advanced** tab and make appropriate changes to the parameter as required.
4. Reboot the system for the changes to take effect.

NOTE

- The modifications to the driver parameters made using the Device Manager Property page are not immediately applied. They take effect during the next driver load sequence; either during the next reboot or a driver unload or load operation.
- Although the Advanced Tab of the Device Manager Property page can be accessed from any OneConnect iSCSI adapter, the parameter changes are uniformly applied to all the OneConnect iSCSI adapter PCIe functions on the system. Individual iSCSI PCIe functions cannot have their own set of parameters.

Table 11 provides descriptions of the different iSCSI driver options and their possible values.

Table 11 iSCSI Driver Options

Parameter	Default Value	Minimum Value	Maximum Value	Description
ETO	90 seconds	0 seconds	3600 seconds	ETO in seconds. This parameter determines the amount of time the driver waits for the target to be available after it has lost connection.
im_policy	2	0	4	The Interrupt Moderation policy parameter controls the rate of interrupts for the adapter. For more information, see Section 3.3.2, Interrupt Moderation Policy Settings.

Table 11 iSCSI Driver Options (Continued)

Parameter	Default Value	Minimum Value	Maximum Value	Description
large_io	64	64	512	Maximum transfer size in a single I/O request, in KB. By default, the iSCSI driver supports a maximum of 64 KB of data and 16 scatter/gather entries in a single I/O request. This option enables support for 512 KB of data in a single I/O request. If an application issues an I/O request that is larger than 64 KB or that needs more than 16 scatter/gather entries, the request is split into multiple requests by the Storport driver. **NOTE** If the large_io parameter is set to 512, the amount of physical memory consumed by the driver increases. Also, although intermediate values between 64 and 512 are accepted, the memory used by the driver is the same as is used if large_io is set to 512.
LDTO	20 seconds	0 seconds	3600 seconds	LDTO, in seconds. This parameter determines the amount of time the driver waits for the controller's physical link to be available before reporting that the LUNs are unavailable to the operating system.
lqd	128	1	255	The LUN queue depth parameter configures the number of concurrent commands to a logical unit via Storport API StorPortSetDeviceQueueDepth. The lqd parameter also sets the maximum number of concurrent commands allowed per LUN.

3.3.2 Interrupt Moderation Policy Settings

The Interrupt Moderation policy settings control the rate of interrupts for adapter hardware. By default, the driver implements an interrupt moderation scheme that is based on the I/O load and the interrupt rate. The default setting for im_policy tries to vary the interrupt rate between 3500 to 10000 interrupts per second. In addition, the iSCSI driver allows other configuration settings, as shown in Table 12.

Table 12 im_policy Settings

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

The default setting works for most configurations; however, there are instances when the setting might need to be altered. The im_policy parameter setting must be based on the adapter system configuration, the number of iSCSI targets to be connected, the I/O load, and the throughput and latency offered by these iSCSI targets.

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

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

3.3.3 Creating Non-Bootable Targets

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

3.3.3.1 Using the Microsoft iSCSI Initiator Service

Use the Microsoft iSCSI Initiator Service to configure and manage the adapter. The Microsoft Initiator Service is available as a free download from www.microsoft.com. Refer to the documentation that accompanies it for detailed information.

NOTE If you install the Microsoft iSCSI Initiator Service, you need to select only the **Initiator Service** check box and not the **Software Initiator** check box.

The Microsoft iSCSI Initiator Service sets its own initiator name. After you have installed it, you must replace this with a custom initiator name.

To assign a name, perform these steps:

1. In the Microsoft iSCSI Initiator Service, under the **General** tab, or the **Configuration** tab (Windows Server 2008 R2 systems and later), and click **Change**.
2. Type your initiator name and click **OK**.

3.3.3.2 Logging into a Target Using the Microsoft Software Initiator

If you install the Software Initiator, you must select the adapter initiator when logging into the target.

To select the initiator, perform these steps:

1. From the **Targets** tab, select the target and click **Logon**, or **Connect** (Windows Server 2008 R2 systems and later).
2. Click **Advanced**.

   Under the **General tab**, everything appears as default.
3. Select the adapter initiator as the local adapter, select your source IP, and click **OK**

3.3.3.3 Windows Multipath I/O Support

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

3.3.3.4 Multipath Support

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

The following steps describe the installation process for setting up Microsoft iSCSI DSM and enabling multipath I/O for all iSCSI devices irrespective of their vendor and device IDs. Use the MPIO GUI to configure DSMs other than Microsoft

iSCSI DSM. Also, use the GUI to enable multipath support for a specific vendor ID and device ID. For more information, refer to the Microsoft TechNet Library on the Microsoft website.

In a multipath configuration, the driver parameters LDTO and ETO can be configured to control the amount of time it takes for the failover operation to complete. The default value of LDTO is 20 seconds and the default value of ETO is 90 seconds.

For information on modifying the timeout parameters in a failover configuration, see Section 3.3.4.2.1, Error Handling Under MultiPath (MPIO) and Cluster Configurations.

If the ETO or LDTO value must be modified, perform these steps:

1. Select the Emulex OneConnect iSCSI adapter in the **Windows Device Manager** under **Storage Controllers**.
2. Right-click and select **Properties**.

   The Device Manager Property page opens.
3. Select the **Advanced** tab and set the desired value of ETO and LDTO; for example, ETO=120 and LDTO=60.
4. Reboot the system for the changes to take effect.
5. Log into the iSCSI target using WMI.

   For more information, see Section 3.3.3.4.1, Logging into Targets for Multipath Support.
6. Enable MPIO.
   a. Select **Start>Administrative Tools>Server Manager**.
   b. In the Server Manager tree, click **Features**.
   c. In the Features area, click **Add Features**.
   d. In the Add Features wizard on the Select Features page, select the **Multipath I/O** check box and click **Next**.
   e. On the Confirm Installation Selections page, click **Install**.
   f. After the installation is completed, click **Close** on the Installation Results page.
   g. When prompted to restart the computer, click **Yes**.
   h. Click **Close**.
7. Discover all possible paths to all devices on the system.
   a. Open the MPIO control pane and select **Start>Administrative Tools>MPIO**.
   b. On the User Account Control page, click **Continue**.

      The Properties dialog box is displayed.
   c. Select the **Discover Multi-Paths** tab.
   d. Select **Add support for iSCSI Devices** and click **Add**.
8. Reboot the system when prompted to do so.

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

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

3.3.3.4.1 Logging into Targets for Multipath Support

After you have successfully installed and enabled MPIO support on a Windows Server, you must log into the target. This section describes the steps to log into iSCSI targets through the WMI GUI. For information on using the iSCSISelect utility to log into an iSCSI target, refer to the *Emulex Boot for NIC, iSCSI, and FCoE Protocols for OneConnect Adapters User Guide*.

To log into a target using WMI, perform these steps:

1. Select the Enable multi-path in the **Log On to Target** window.

This setting must be selected for every target to be logged in for MPIO. Use the Advanced tab to select the specific adapter port to use for login.

2. After the target login is complete, select the target and click the **Details** or **Properties** button (depending on the Windows operating system) to see the multiple sessions connected.

3.3.4 Maximum Transmission Unit (MTU) for iSCSI Connections

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

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

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

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

3.3.4.1 iSCSI Error Handling

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

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

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

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

3.3.4.2 Configuring LDTO and ETO on the Windows Server

LDTO and ETO values are configured using the Advanced tab of the Device Manager Property page. Table 13 on page 97 lists the default values of LDTO and ETO on the Windows Server and the configuration limits.

NOTE If the ETO is set to a number between 0 and 19, the driver assumes the value of 20 seconds internally. The registry is not modified.

Table 13 LDTO and ETO Information on the Windows Server

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

To modify LDTO and ETO values, edit the driver parameters for the iscsi service, perform these steps:

1. Select the Emulex OneConnect iSCSI adapter in the Windows Device Manager under Storage Controllers.
2. Right-click and select **Properties**.

   The Device Manager Property page opens.
3. Select the **Advanced** tab and make the following changes:
   - LDTO = 25
   - ETO = 50
4. Reboot the system for the changes to take effect.

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

3.3.4.2.1 Error Handling Under MultiPath (MPIO) and Cluster Configurations

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

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

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

Chapter 4: Troubleshooting

Your system may operate in an unexpected manner in certain circumstances. This section contains reference tables on event codes and error messages and provides information regarding unusual situations.

4.1 General Troubleshooting

The following table describes issues you may encounter and their solutions.

Table 14 General Troubleshooting

Issue	Answer/Solution
The operating system fails to install or does not successfully install the driver.	Verify that the operating system is supported by the driver.
The AutoPilot Installer fails.	If the AutoPilot Installer fails, the Diagnostics window shows that the adapter failed. If the adapter fails: 1. Select the adapter to view the reason why the adapter failed. The reason and suggested corrective action are displayed. 2. Perform the suggested corrective action and run AutoPilot Installer again. **NOTE** You can run AutoPilot Installer again from the Start menu (**Programs>Broadcom>AutoPilot Installer**), or you can run APInstall.exe from a command prompt.
The OneInstall Installer fails.	If the OneInstall Installer fails, it may be because: ■ The operating system prerequisites have not been met. ■ The individual kit installation failed. To check, run the installation interactively. If you encounter error messages when you run the installation interactively, those issues would also apply to an unattended installation. ■ If an individual package failed to install properly, run that package's installer directly. This method displays status and error messages that can be used to diagnose the issue. (The OneInstall Installer does not provide these displays because each package is installed silently.)
Windows Device Manager shows a code 10 or code 39 with a yellow or red exclamation point on the device.	The firmware image does not match the installed device drivers, or the firmware is corrupt. Using the OneCommand CNA Manager application or one of the Windows PE offline or online utilities, install a version of firmware that is compatible with the driver.
The firmware is corrupt or non-responsive.	Using the OneCommand CNA Manager application or one of the Windows PE offline or online utilities, install a version of firmware that is compatible with the driver.
The Emulex iSCSI BIOS banner is not displayed during system POST.	Configure the motherboard BIOS to enable the Option ROM for the PCIe slot in which the adapter is installed.

4.2 Troubleshooting the FCoE Driver

4.2.1 Event Trace Messages

4.2.1.1 ELS Log Messages (0100–0130)

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

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

4.2.1.1.2 lpfc_mes0101: FLOGI completes successfully – NPortId: x%x, RaTov: x%x, EdTov: x%x

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

4.2.1.1.3 lpfc_mes0102: PLOGI completes to NPortId: x%x

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

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

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

4.2.1.1.5 lpfc_mes0104: ADISC completes to NPortId x%x

Description	The adapter performed an ADISC into the remote NPort.
Severity	Information
Log	LOG_ELS verbose
Action	No action needed, informational.

4.2.1.1.6 lpfc_mes0105: LOGO completes to NPortId: x%x

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

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

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

4.2.1.1.8 lpfc_mes0114: PLOGI chkparm OK

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

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

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

4.2.1.1.10 lpfc_mes0128: Accepted ELS command: OpCode: x%x

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

4.2.1.1.11 lpfc_mes0129: Rejected ELS command: OpCode: x%x

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

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

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

4.2.1.2 Discovery Log Messages (0202–0262)

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

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

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

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

4.2.1.2.3 lpfc_mes0214: RSCN received: Word count:%d

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

4.2.1.2.4 lpfc_mes0215: RSCN processed: DID: x%x

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

4.2.1.2.5 lpfc_mes0225: Device Discovery completes

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

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

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

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

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

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

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

4.2.1.2.9 lpfc_mes0235: New RSCN being deferred due to RSCN in process

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

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

Description	The driver is issuing a nameserver request to the fabric. Also recorded if a GID_FT is sent.
Severity	Information
Log	LOG_DISCOVERY verbose

Action	No action needed, informational.
See Also	lpfc_mes0239 or lpfc_mes0240

4.2.1.2.11 lpfc_mes0238: NameServer response DID count:%d

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

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

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

4.2.1.2.13 lpfc_mes0240: NameServer Response Error – CmdRsp:x%x, ReasonCode: x%x, Explanation x%x

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

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

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

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

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

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

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

4.2.1.3 Mailbox Log Messages (0310–0326)

4.2.1.3.1 lpfc_mes0310: Mailbox command timeout – HBA unresponsive

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

4.2.1.3.2 lpfc_mes0326: Reset HBA – HostStatus: x%x

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

4.2.1.4 INIT Log Messages (0400–0463)

4.2.1.4.1 lpfc_mes0400: Initializing discovery module: OptionFlags: x%x

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

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

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

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

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

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

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

4.2.1.4.5 lpfc_mes0436: Adapter not ready: hostStatus: x%x

Description	The adapter failed during powerup diagnostics after it was reset.
Severity	Error
Log	LOG_INIT verbose
Action	This error could indicate a hardware or firmware issue. If issues persist report these errors to Broadcom technical support.

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

Description	Adapter initialization failed when issuing CONFIG_PORT mailbox command.
Severity	Error
Log	LOG_INIT verbose
Action	This error could indicate a hardware or firmware issue. If issues persist report these errors to Broadcom technical support.

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

Description	Adapter initialization failed when issuing CFG_RING mailbox command.
Severity	Error
Log	LOG_INIT verbose
Action	This error could indicate a hardware or firmware issue. If issues persist report these errors to Broadcom technical support.

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

Description	Adapter initialization failed when issuing INIT_LINK mailbox command.
Severity	Error
Log	LOG_INIT verbose
Action	This error could indicate a hardware or firmware issue. If issues persist report these errors to Broadcom technical support.

4.2.1.4.9 lpfc_mes0458: Bring Adapter online

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

4.2.1.4.10 lpfc_mes0460: Bring Adapter offline

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

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

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

4.2.1.5 FCP Log Messages (0701–0749)

4.2.1.5.1 lpfc_mes0701: Issue Abort Task Set to PathId: x%x, TargetId: x%x, Lun: x%x

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

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

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

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

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

4.2.1.5.4 lpfc_mes0714: Issued Bus Reset for PathId:%d

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

4.2.1.5.12 lpfc_mes0749: Abort task set complete: Did: x%x, SCSIId: x%x

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

4.2.1.6 Link Log Messages (1302–1306)

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

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

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

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

4.2.1.6.3 lpfc_mes1305: Link down even: tag x%x

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

4.2.1.6.4 lpfc_mes1306: Link Down timeout

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

4.2.1.7 Tag Messages (1400–1401)

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

Description	Firmware has generated an invalid response.
Severity	Error
Log	`LOG_LINK_EVENT verbose`
Action	Review hardware configuration. Contact Broadcom technical support.

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

Description	Firmware has generated an invalid response.
Severity	Error
Log	`LOG_LINK_EVENT verbose`
Action	Review hardware configuration. Contact Broadcom technical support.

4.2.1.8 NPIV Messages (1800–1899)

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

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

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

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

4.2.1.8.3 lpfc_mes1802: Exceeded the MAX NPIV port: x%x

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

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

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

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

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

4.2.1.9 ELS Messages (1900–1999)

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

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

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

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

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

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

4.2.1.9.4 lpfc_mes1903: Receives ELS_AUTH_RJT x%x x%x

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

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

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

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

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

4.3 Troubleshooting the NIC Drivers

Table 15 provides troubleshooting information for the NIC drivers.

Table 15 Troubleshooting the NIC Drivers

Issue	Answer/Solution
Performance is not as expected.	The adapter may be installed in the wrong type of PCIe slot. Verify that the adapter has been properly installed. If TOE is enabled and performance is not as high as expected, the operating system may not offload TOE connections. For more information, see Section 3.2.9.3, TCP Offloading (TOE).
Frequent event log entries for link changes, or statistics that show more than expected CRC errors, occur.	Unload and reload the driver to reset available target IDs. Ensure that the SAN configuration is correct prior to reloading the driver. This clears the driver's consistent binding table and frees target IDs for new target nodes.
The driver fails to load, and an event log entry states that the driver failed to load due to memory constraints.	There may not be enough memory installed in the system to provide sufficient memory for all devices installed in the system. Try installing more memory if possible.
Unpredictable results occur if the flow control setting differs among PCI functions.	If multiple PCI functions are exposed for a single 10GbE port, such as in blade configurations, the flow control parameter must be set the same on all adapters for the port. Results are unpredictable if the setting differs among PCI functions because this is a shared property of the 10GbE port.
On servers that support PCIe hot unplug, the system may hang or produce a bugcheck if a PCIe hot unplug or replace is attempted.	Hot unplug is not supported in this release.
If Wake on LAN is set through the advanced properties page, the system does not wake when receiving a magic packet or a packet that would normally wake the system.	The system may not support Wake on LAN on the PCIe slot in which the adapter is installed. Check the system documentation to determine whether the system is capable of Wake on LAN operation. A system BIOS setting may not be correct for Wake on LAN to work as expected. Check the system documentation to determine whether Wake on LAN must be enabled in the system BIOS. Wake on LAN may not be supported by the chipset as reported by the firmware. The driver reports the value that is reported by the firmware. The system may not go to a lower power state because another software component, device, or driver is preventing it. Microsoft provides several useful references for troubleshooting Wake on LAN configuration issues in the Microsoft TechNet Library on the Microsoft website.

Table 15 Troubleshooting the NIC Drivers (Continued)

Issue	Answer/Solution
When running Windows Server 2008, the computer restarts and shows various Stop error codes when performing one of the following operations: ■ Enabling or disabling TCP Chimney Offload ■ Changing the network adapter settings ■ Upgrading the NIC drivers	Apply the 979614 hotfix as described on the Microsoft website.
If an NDIS driver is being installed manually on a Windows Server 2008 system, the installer installs the first driver it finds, even if it is not the latest version of the driver.	Windows Server 2008 picks up the first available driver it finds if an NDIS driver is being installed manually. Therefore, an NDIS5 driver will be installed even if a Windows NDIS6 driver is available. An event log message advises you to update to the latest driver for best performance.
The system crashes or appears to hang. In the case of a hang, there could be a message indicating that the driver experienced a hardware malfunction.	Several possible causes for this issue exist. ■ Certain systems require an updated BIOS to properly manage the power states of newer Intel and AMD processors. Check with your vendor for information regarding BIOS and firmware updates that may be required to run well with the latest releases of the Windows operating systems. Also, certain BIOS settings may be required. For example, disable any low power processor states and low power settings for PCIe. ■ On certain AMD systems, it is possible the intelppm.sys driver is enabled, and should not be. To query this system driver's run state, log in as administrator and at the command line, type: sc query intelppm If the results indicate that the intelppm driver is running, you must disable it. At the command line, type: sc config intelppm start= disabled ■ On all systems, it may be necessary to set the power options to High Performance. See the operating system documentation for details.

4.3.1 Monitoring TCP Offloads

To monitor TCP offloads, in a command window type

```
netstat -t
```

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

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

4.3.2 TCP Offload Failure

Table 16 lists common reasons why TCP offloads do not occur and their potential solutions.

Table 16 Troubleshooting TCP Offload Failures

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

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

4.4 Troubleshooting the iSCSI Driver

This section provides troubleshooting information for the iSCSI driver.

4.4.1 Troubleshooting the Cisco Nexus Switch Configuration

NOTE LACP cannot be used on an iSCSI port.

Table 17 provides troubleshooting information for the Cisco switch.

Table 17 Cisco Nexus Switch Situations for iSCSI

Issue	Solution
The system is showing an excessive number of I/O timeouts as a result of the switch routing frames to the incorrect port.	Ensure that LACP is not used on the iSCSI port.

4.4.2 iSCSI Driver Troubleshooting

Table 18 provides troubleshooting information for the iSCSI driver.

Table 18 Troubleshooting the iSCSI Driver

Issue	Answer/Solution
Overall failure.	Use the iSCSISelect utility to clear the Adapter Configuration. See the *Emulex Boot for NIC, iSCSI, and FCoE Protocols for OneConnect Adapters User Guide* for instructions.
The driver loads successfully, but there are event 11 entries in the event log for the iSCSI driver.	The most common cause is that the adapter link is down. See Section A.3.1, Viewing the iSCSI Error and Event Log on Windows Server 2008 and look for specific event codes to confirm.
Under Windows Server 2008 R2, you cannot inject the version 11.x iSCSI driver during boot from SAN .	Perform BFS using the existing (earlier than version 11.x) iSCSI driver. After the operating system is booted, update to the version 11.x iSCSI driver.
Unable to create a memory dump file on a system booted over iSCSI.	Make sure the disk has enough free disk space to create the dump file. If a full memory dump is selected, the disk must have free space at least equivalent to the amount of physical memory in the system.
Unable to log in to target from WMI.	Ensure that the IP address on the adapter is valid and the network connection has been set up to reach the target. If login is attempted after discovering the target, ensure that the correct adapter port has been selected for the login.
The iSCSI WMI GUI shows the target state as connected, but LUNs are not seen from the disk manager.	Verify that the adapter name used to connect to the target matches the adapter name configured on the iSCSI target.
Multipath configuration shows duplicate LUNs on the disk manager.	Ensure that the MPIO software is installed and the login options have selected the MPIO flag. On Windows Server 2008 and Windows Server 2008 R2 Operating Systems, the server role must be set up for Multipath. See the *Emulex Boot for NIC, iSCSI, and FCoE Protocols for OneConnect Adapters User Guide* for additional information on MPIO.
Multipath configuration takes a long time to fail over or failover does not occur.	Ensure that LDTO settings and ETO settings have been configured for MPIO. These values must be set to 0. For more information, see Section 3.3.4.2, Configuring LDTO and ETO on the Windows Server.
Sendtargets to a IET fails because it violates the iSCSI specification.	If you still want to add a IET, you must add the target manually. This issue affects Sendtargets only.
The following POST error message appears: Initiator iSCSI Name mismatch, Please use iSCSISelect to set a single name for all controllers. Press <Ctrl><S> to enter iSCSISelect. (Reboot required)	In the iSCSI BIOS, the Emulex iSCSI initiator name may be different if more than one OneConnect adapters are in the system. This message appears if the iSCSI initiator name is different on multiple controllers. Enter the iSCSISelect utility and save a new initiator name on the first iSCSISelect utility menu window so that the iSCSI initiator name on all controllers match. All logins from the multiple controllers will use the new name. See the *Emulex Boot for NIC, iSCSI, and FCoE Protocols for OneConnect Adapters User Guide* for more information.

Table 18 Troubleshooting the iSCSI Driver (Continued)

Issue	Answer/Solution
If an iscsicli logouttarget command is issued back-to-back in a script, event 12 errors from the PlugPlayManager are seen in the Windows Event Viewer. The error message is similar to this string: The device 'SE ISCSI 00 SCSI Disk Device'(SCSI\Disk&Ven_SE_ISCSI&Prod_00&Rev_3.64\5&17659873&2&020000) disappeared from the system without first being prepared for removal.	This behavior is not specific to the OneConnect adapter.
On a system running Windows Server 2008, or Windows Server 2008 R2, the iSCSI driver fails to load after many iterations of enable/disable from Device Manager. Because the system failed to allocate contiguous uncached extension memory, the iSCSI driver failed to load, and an attention icon is displayed next to the OneConnect iSCSI device. The Device Status shows "This device cannot start. (Code 10)", and an Event 11 error is logged in the Windows system event log for the iSCSI driver with 0x31840006 in the 5th DWORD.	No workaround exists for this issue.
If an iSCSI adapter is used to login to an iSCSI target and the LUN configuration on the target is changed, neither the adapter nor the WMI GUI see the updated LUN configuration.	If an iSCSI target provides an asynchronous event notification to the adapter when its logical unit inventory has changed, the iSCSI driver initiates a bus rescan and the LUNs are updated dynamically. However, if an iSCSI target does not provide an asynchronous event notification, the LUN list is not updated dynamically. Perform a manual rescan in Disk Management.
A login to new target fails after Microsoft iSCSI Initiator Service is installed.	If Microsoft iSCSI software is installed, the service chooses a default IQN name for the adapter. The Microsoft iSCSI service issues the request to the iSCSI driver using the WMI interface to set this new IQN name. Therefore, any IQN name that was configured earlier (such as by using iSCSISelect) is overridden and the new IQN name will be in effect. Although this does not affect existing boot sessions and persistent sessions, new target logins could fail because the new IQN name does not match the incoming IQN name configured on the target. After the Microsoft iSCSI Initiator Service is installed, the initiator name must be renamed to the previous name configured from the WMI GUI.
If software-based iSCSI targets are logged into the adapter, Event ID 56 (Driver SCSI (000000)) appears in the Windows event viewer. This issue has been observed on Windows Server 2008 R2 under the following conditions: The iSCSI target is a software-based target (MSiSCSI, IET, StarWind) that uses a local hard drive or a RAM disk for its backend LUN. Different adapter ports are involved in the login. A SAS controller is present on the system.	This occurs caused because of an issue with the data reported by the iSCSI target in the Product Identification field in response to the standard inquiry from the adapter. This field should be unique among different targets' LUNs, but software-based targets report the same pre-formatted data for all the LUNs across all targets. lf Windows encounters the same Product Identification field for different LUNs with the same Bus Target LUN field, it records error in the event log. No other effect has been found as a result of this behavior. The workaround for this error is to use non-overlapping LUN numbers for the various LUNs across the various iSCSI targets. On the iSCSI target system, LUNs can be numbered sequentially; they do not have to start at zero.

Appendix A: Error and Event Log Information

A.1 FC/FCoE Error and Event Logs

A.1.1 Viewing the FC/FCoE Error Log

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

To view the error log:

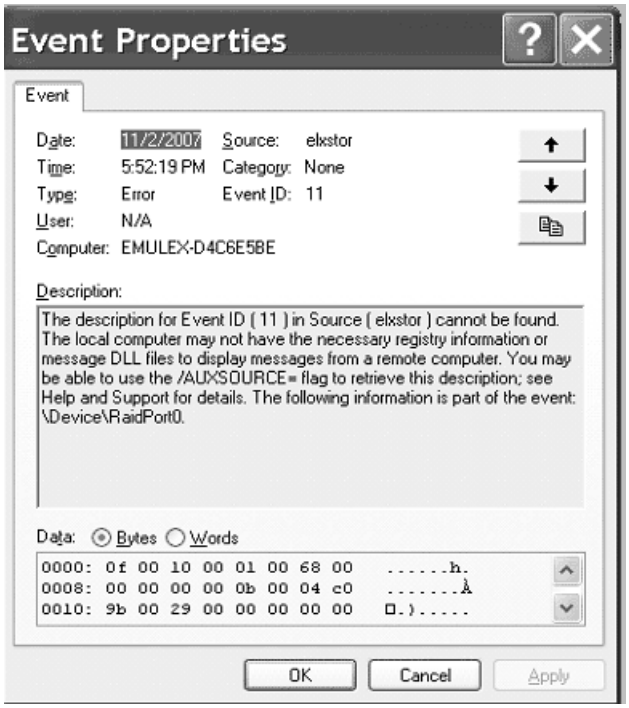
1. Open the Event Viewer window by doing one of the following:
   - Click Start>Programs>Administrative Tools>Event Viewer.
   - Right-click My Computer, Manage and Event Viewer in Computer Management.

   The Event Viewer window is displayed.
2. Double-click any event with the source name ELXCNA.
3. Examine the entry at offset 0x10 and Event ID 11. The Emulex event code is found in byte 0010 and supplementary data is in the byte offsets 0011 through 0013.

For example, in Figure 29:

byte 0010 = 9b, byte 0011 = 00, byte 0012 = 29 and byte 0013 = 00

Figure 29 Event Properties



A.1.1.1 Severity Scheme

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

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

A.1.1.2 Related Driver Parameter: LogError

The LogError driver parameter determines the minimum severity level to enable entry of a logged error into the system. See Chapter 3 Configuration for instructions on how to set driver parameters.

- If set to 0 = All errors regardless of severity are logged.
- If set to 1 = Severe, malfunction, and command level errors are logged.
- If set to 2 = Both severe and malfunction errors are logged.
- If set to 3 = Only severe errors are logged.

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

A.1.1.3 Format of an Error Log Entry

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

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

A.1.1.4 Error Codes Tables

This section provides tables listing error codes and their descriptions.

A.1.1.4.1 Severe Errors

Table 19 lists severe errors and their codes.

Table 19 Severe Errors

Bits 0–7	Interpretation
0x04	CONFIG_PORT failed
0x06	READ_CONFIG_failed
0x07	CONFIG_RING 0 failed
0x08	CONFIG_RING 2 failed
0x09	CONFIG_RING 1 failed

Table 19 Severe Errors (Continued)

Bits 0–7	Interpretation
0x0A	CONFIG_RING 3 failed
0x10	CONFIG_PORT failed (reinitialization)
0x12	READ_CONFIG command failed (reinitialization)
0x13	CONFIG_RING 0 failed (reinitialization)
0x14	CONFIG_RING 1 failed (reinitialization)
0x15	CONFIG_RING 2 failed (reinitialization)
0x16	CONFIG_RING 3 failed (reinitialization)
0x28	Failure on REG_LOGIN mailbox command
0x29	Unable to initiate fabric binding operation
0x42	Re-simulate FCF after exhausted retries on FLOGI
0x51	ABTS timeout on path and target (Bits 8-15: path id; Bits 16-23: target id)
0x2A	Attempted ADISC to non-existent node
0x2B	Failure on iocb context allocation
0x2C	Unable to initiate nport unbinding operation
0x2D	Unable to initiate nport binding operation
0x2E	Failed to allocate resources for Express Lane
0x30	Failure on mailbox context allocation
0x7C	Menlo initialization error
0x7D	Menlo initialization error
0x7E	Menlo initialization error
0xA0	Failed to initialize adapter port
0xCA	Invalid scatter gather list size
0xCB	Unsupported IFType (SLI4 mode)
0xC1	Failed to allocate miniport un-cached extension
0xC2	Insufficient un-cached extension space
0xC3	Port initialization failure
0xC5	Utility mailbox command error
0xC6	SLI4 Pre-initialization failure
0xC7	UNREG_VPI failure requiring reset
0xC8	Invalid FLOGI response failure requiring reset
0xC9	REG_FCFI failure requiring resolicitation (SLI4 mode)
0xD3	NPIV memory allocation failure
0xE0	Unable to allocate exchange for unsolicited ELS command
0xE1	Misconfigured port event on indicated port, link effect and link state (SLI4 mode) Bits 31–24: Port Name; Bits 23–16: Link effect; Bits 15–8: Link state. Missing or unqualified SFP installed.
0xF0	Unresponsive adapter port (SLI4 mode)
0xF4	ULP Unrecoverable Error: low part (SLI4 mode)
0xF5	ULP Unrecoverable Error: high part (SLI4 mode)
0xF6	ARM Unrecoverable Error (SLI4 mode)
0xF7	READ_NV failed (SLI4 mode)

Table 19 Severe Errors (Continued)

Bits 0–7	Interpretation
0xF8	READ_NV failed (SLI4 mode)
0xF9	READ_REV failed (SLI4 mode)
0xFA	READ_CONFIG failed (SLI4 mode)
0xFB	Failed to post header templates (SLI4 mode)
0xFC	Invalid Completion Queue Entry (SLI4 mode)
0xFD	Invalid Completion Queue Entry (SLI4 mode)
0xFE	Invalid Completion Queue Entry (SLI4 mode)

A.1.1.4.2 Malfunction Errors

Table 20 lists malfunction errors and their codes.

Table 20 Malfunction Errors

Bits 0–7	Interpretation
0x05	SET_VAR command failed
0x11	SET_VAR command failed (reinitialization)
0x21	Spurious mailbox command interrupt
0x31	Unrecognized mailbox command completion
0x32	Duplicate link attention: event tag unchanged
0x33	Invalid link attention: no link state indicated
0x34	Duplicate link attention: link state unchanged
0x35	Error reading common service parameters for port
0x36	Error reading common service parameters for fabric
0x37	Error reading common service parameters for nport
0xB1	Write check error
0x3B	Failed to create node object
0x3C	PRLI initiation failure
0x3D	Recoverable UNREG base VPI error (Bits 8–15: mailbox status)
0x3E	Recoverable UNREG VPI error (Bits 8–15: mailbox status)
0x42	Exhausted retries on FLOGI
0x45	ELS command rejected
0x49	Exhausted retries on PLOGI
0x4E	World Wide Port Name mismatch on ADISC
0x4F	World Wide Node Name mismatch on ADISC
0x50	ADISC response failure
0x55	LOGO response failure
0x57	PRLI to non-existent node
0x5A	PRLI response error
0x5F	CT command error
0x62	Name server response error
0x66	State Change Notification registration failure
0x6A	Unrecognized ELS command received

Table 20 Malfunction Errors (Continued)

Bits 0–7	Interpretation
0x6F	Received PRLI from un-typed source
0x73	Failed to pend PRLI for authentication
0x77	Failed to allocate Node object
0x7A	REG_VPI failed
0xA3	Command context allocation failure
0xAB	SCSI command error
0xAC	Read check error
0xB0	Node timeout: device removal signaled to Storport
0xB2	FCP_RSP short frame received
0xE1	Misconfigured port event on indicated port, link effect and link state (SLI4 mode) Bits 31–24: Port Name; Bits 23–16: Link effect; Bits 15–8: Link state. Missing or unqualified SFP installed.

A.1.1.4.3 Command Errors

Table 21 lists command errors and their codes.

Table 21 Command Errors

Bits 0–7	Interpretation
0x43	Fabric login succeeded
0x46	ELS command failed
0x47	Exhausted retries on ELS command
0x4A	PLOGI accepted
0x56	LOGO accepted
0x59	PRLI accepted
0x63	Fabric name server response
0x6B	ELS RSCN processed
0x71	LOGO received from fabric
0x79	FDISC accepted
0xA2	SCSI address assigned to discovered target
0xA4	Report LUNs error (initial I/O to discovered target)
0xA5	Local error indication on FCP command
0xA6	FCP Command error
0xA8	Data overrun
0xA9	FCP command error
0xAA	SCSI check condition
0xAD	Local reject indication on FCP command
0xAE	Error on SCSI pass-through command
0xAF	Error on Menlo CT command
0xE1	Misconfigured ports event on indicated port, link effect and link state (SLI4 mode) Bits 31–24: Port Name; Bits 23–16: Link effect; Bits 15–8: Link state. Missing or unqualified SFP installed.

A.1.1.4.4 Event Indicators

Table 22 lists event indications and their codes.

Table 22 Event Indications

Bits 0–7	Interpretation
0xA7	Data underrun
0xD0	NPIV Virtual Port creation success (Virtual Port Did in bits 8–31)
0xD1	NPIV Virtual Port creation failed (Virtual Port index in bits 8–31)
0xD2	NPIV Virtual Port FDISC failed (Virtual Port index in bits 8–31)
0xD4	Exceeded max Virtual Port supported (Virtual Port index in bits 8–31)
0xD5	NPIV Virtual Port removal (Virtual Port Did in bits 8–31)
0xEE	Authenticated successfully (remote Did in bits 8–31)
0xEF	Failed to authenticate (remote Did in bits 8–31)
0xE2	Authentication not support (remote Did in bits 8–31)
0xE3	Authentication ELS command timeout (remote Did in bits 8–31)
0xE4	Authentication transaction timeout (remote Did in bits 8–31)
0xE5	LS_RJT other than Logical Busy received for Authentication transaction (remote Did in bits 8–31)
0xE6	LS_RJT Logical Busy received for Authentication Transaction (remote Did in bits 8–31)
0xE7	Received Authentication Reject other than Restart (remote Did in bits 8–31)
0xE8	Received Authentication Reject Restart (remote Did in bits 8–31)
0xE9	Received Authentication Negotiate (remote Did in bits 8–31)
0xEA	Authentication spurious traffic (remote Did in bits 8–31)
0xEB	Authentication policy has been changed (remote Did in bits 8–31)
0xED	Same passed were set for both local and remote entities (remote Did in bits 8–31)
0xF1	Port shutdown event (SLI4 mode)
0xF2	Port in off-line state (SLI4 mode)
0xF3	Port in on-line state (SLI4 mode)

A.1.2 Viewing the FCoE Event Log

This section provides information on the FCoE event logs.

A.1.2.1 Event Log Interpretation

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

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

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

A.1.2.2 Additional Event Log Information

The following tables are not comprehensive but do include the codes that are most likely to show up in SAN environments where issues occur.

A.1.2.2.1 ELS/FCP Command Error Status Codes

Table 23 lists the internal firmware codes posted by the adapter firmware that explain why a particular ELS or FCP command failed at the FC level.

Table 23 ELS/FCP Command Error Status Codes

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

A.1.2.2.2 CT Command Response Codes

Table 24 lists the codes that indicate the response to a FC Common Transport protocol command.

Table 24 CT Command Response Codes

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

A.1.2.2.3 FC-CT Reject Reason Codes

Table 25 lists the codes that indicate the reason a CT command was rejected.

Table 25 FC-CT Reject Reason Codes

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

A.1.2.2.4 ELS Command Codes

Table 26 lists the FC protocol codes that describe the Extended Link Services commands that were sent.

Table 26 ELS Command Codes

Explanation	Code
Link Service Reject (LS_RJT)	0x01
Accept (ACC)	0x02
N_Port Login (PLOGI)	0x03
Fabric Login (FLOGI	0x04
N_Port Logout (LOGO)	0x05
Process Login (PRLI)	0x20

Table 26 ELS Command Codes (Continued)

Explanation	Code
Process Logout (PRLO)	0x21
Discover F_Port Service Params (FDISC)	0x51
Discover Address (ADISC)	0x52
Register State Change Notify (RSCN)	0x61

A.1.2.2.5 SCSI Status Codes

Table 27 lists the SCSI status codes returned from a SCSI device which receives a SCSI command.

Table 27 SCSI Status Codes

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

A.1.2.2.6 Local Reject Status Codes

Table 28 list the codes supplied by the Emulex adapter firmware that indicate why a command was failed by the adapter.

Table 28 Local Reject Status Codes

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

A.1.2.2.7 SRB Status Codes

Table 29 lists the SCSI Request Block status codes provided by the driver to the operating system based upon the response from a SCSI device in the SAN.

Table 29 SRB Status Codes

Explanation	Code
ERROR	0x04
BUSY	0x05
TIMEOUT	0x09

Table 29 SRB Status Codes (Continued)

Explanation	Code
SELECTION TIMEOUT	0x0A
COMMAND TIMEOUT	0x0B
BUS RESET	0x0E
DATA OVERUN	0x12

A.1.2.3 ASC/ASCQ

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

A.1.2.4 Additional Notes on Selected Error Codes

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

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

A.1.2.4.1 Node Timeout (Code 0xAA)

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

A.1.2.4.2 SCSI Command Error (0x9A) and SCSI Check Condition (code 0x9B)

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

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

A.1.2.4.3 Nameserver Response (Code 0x98)

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

A.1.2.4.4 Context Allocation Failures

A number of event codes for which the interpretation contains the phrase "context allocation failure" exist. These types of events are referring to the internal memory constructs of the Emulex Storport Miniport driver and as such are intended for Broadcom design engineers' information. If you encounter this type of code, contact Broadcom technical support for assistance.

NOTE Context allocation failures are rare.

A.2 NIC Error and Event Logs

This section provides information on NIC error and event logs.

A.2.1 Viewing the NIC Error Log

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

To view the error codes, perform these steps:

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

A.2.2 RoCE Event Log

The Windows Device Manager generates error log codes if any errors occur during the installation of the NIC or RoCE driver. Each log contains a Message ID, Severity, and Symbolic Link. The Message ID is unique and tracks the error message if it is not displayed.

Table 30 shows the list of error codes, the severity of the error, the message displayed, the meaning of the error, and the recommended resolutions. When reporting an issue with the adapter to Broadcom, check the event log and report any of these entries that are displayed.

Table 30 RoCE Event Log Entries

Message ID Hexadecimal/ Decimal	Severity	Message	Recommended Resolution
0x3F/63	Warning	<adapter>: Incorrect RoCE profile selected, select the RoCE-2 profile.	Select the RoCE-2 profile.
0x3C/60	Warning	The adapter ran out of resources while creating the requested number of SMB Direct connections. Reduce the connection count to a supported value.	Reduce the connection count to a supported value.
0x3B/59	Warning	RoCE is not enabled. Update the firmware and ensure that the RoCE-2 profile is selected.	Update the firmware and ensure that the RoCE-2 profile is selected in the OneCommand CNA Manager application, the OneCommand CNA Manager CLI, or the PXESelect utility.
0x004B	Warning	Driver received an unexpected RoCE packet.	Ensure the same RoCE mode is enabled on both sides of the system.

A.2.3 NIC Event Log

The Windows Device Manager generates error log codes if any errors occur during the installation of the NIC driver. Each log contains a Message ID, Severity, and Symbolic Link. The Message ID is unique and tracks the error message (if not displayed). Table 31 shows the list of error codes, the severity of the error, the message displayed, the meaning of the error, and recommended resolutions. When reporting an issue with the adapter to Broadcom, check the event log and report any of these entries that are displayed.

Table 31 NIC Event Log Entries

Message ID Hexadecimal/ Decimal	Severity	Message	Recommended Resolution
0x80000039	Warning	Firmware does not support GRE encapsulation. Encapsulation disabled.	Issue configuring NVRE hardware offloads. Update firmware on the OCe14000 adapter.
0x4000003AL	Informational	N/A	N/A
0x0000038L/56	Warning	The device firmware does not support ETS functionality in SR-IOV or multichannel mode.	Revert to default mode for ETS support.
0x0049/73	Informational	%2 : Correct optics installed. (%2 is a place holder for the NIC controller name.)	This message is informational.
0x0046/70	Warning	Unqualified SFP+ module detected on %2, Port %3 from %4 part number %5. (%2 is a holder for the NIC controller name. %3, %4, %5 are placeholders for other arguments in the log message.)	Replace the SFP+ module.
0x0045/69	Informational	SFP+ module detected on %2, Port %3 from %4 part number %5. (%2 is a place holder for the NIC controller name. %3, %4, %5 are placeholders for other arguments in the log message.)	This message is informational.
0x00037/55	Warning	This adapter may have an issue recovering from corrupted use of SR-IOV. Assigning an SR-IOV device to a Virtual Machine could leave the system vulnerable, and lead to instability. Assign SR-IOV devices only to Virtual Machines that run trusted workloads, or consider disabling the use of SR-IOV.	This adapter exposes a vulnerability to the VM that may allow the VM to crash the entire physical computer. This is no different than running a physical adapter. SR-IOV should only be used if the VM has a trusted server administrator.
0x00036/54	Warning	Incompatible optics. Replace with compatible optics for card to function.	Replace the incompatible SFP transceivers with compatible ones for the adapter to function correctly.
0x00035/53	Warning	Optics of two types installed-Remove one optic or install matching pair of optics.	Remove one SFP transceiver or install a matching pair of SFP transceivers.
0x00034/52	Warning	Optics faulted/incorrectly installed/not installed. Reseat optics, if issue not resolved, replace.	Reseat the SFP transceiver. If the issue is not resolved, replace the transceiver.
0x00033/51	Warning	SR-IOV virtualization failed initialization. Check system BIOS settings, or disable SR-IOV for the adapter.	Check system BIOS settings, or disable SR-IOV for the adapter.

Table 31 NIC Event Log Entries (Continued)

Message ID Hexadecimal/ Decimal	Severity	Message	Recommended Resolution
0x00032/50	Warning	The Ethernet link is down due to PHY over-temperature condition. Improve cooling for the device.	Improve the cooling conditions for the device.
0x00031/49	Warning	RSS is limited to 4 queues. Enable Advanced Mode in the PXE BIOS to use up to 16 queues. This may require a firmware update.	Enable Advanced Mode in the PXESelect BIOS utility during boot to use up to 16 queues. This may require a firmware update. See the Broadcom website for compatible firmware.
0x00030/48	Warning	SR-IOV is not enabled. Update the firmware, enable SR-IOV in the server BIOS, and enable SR-IOV and Advanced Mode in the PXE BIOS.	Update the firmware, enable SR-IOV in the server BIOS, and enable SR-IOV and Advanced Mode in the PXESelect BIOS utility. See the Broadcom website for compatible firmware.
0x0002f/47	Warning	VMQ offload is disabled. Disable SR/IOV support in PXE BIOS to use VMQ.	Disable SR-IOV support in the PXESelectBIOS utility to use VMQ.
0x0002e/46	Error	Device is not supported on Windows 7 Operating System.	
0x0002d/45	Error	Error recovery failed. The device is no longer operational. Update all drivers and firmware.	See the Broadcom website for compatible firmware and drivers.
0x0002c/44	Warning	Error recovery is disabled on the system. The device is no longer operational.	This message is informational.
0x0002b/43	Informational	The driver successfully recovered from an error.	This message is informational.
0x0002a/42	Warning	Legacy driver loaded. Move to the NDIS 6.20 driver for Windows Server 2008 R2 for best performance.	
0x0029/41	Warning	Legacy driver loaded. Move to the NDIS 6.x driver for Windows Server 2008 for best performance.	
0x0028/40	Warning	The firmware is outdated and does not support TOE offloads for this driver. Update the firmware.	The firmware and the driver are not compatible versions. See the Broadcom website for compatible firmware and drivers.
0x0026/38	Warning	The device firmware does not support RSS functionality for this network adapter.	The firmware and the driver are not compatible versions. See the Broadcom website for compatible firmware and drivers.
0x0025/37	Warning	The device firmware does not support TCP offload functionality.	The firmware and the driver are not compatible versions. See the Broadcom website for compatible firmware and drivers.
0x0024/36	Error	The device firmware does not support network functionality.	The firmware and the driver are not compatible versions. See the Broadcom website for compatible firmware and drivers.
0x0023/35	Warning	The Ethernet link is down due to a remote fault.	The Ethernet link is down due to the remote partner signaling a fault. Check the peer device for errors.

Table 31 NIC Event Log Entries (Continued)

Message ID Hexadecimal/ Decimal	Severity	Message	Recommended Resolution
0x0022/34	Warning	The Ethernet link is down due to a local fault.	The Ethernet link is down due to a link-down event detected at the driver.
0x0021/33	Informational	Network device is operating in Gen2 mode and installed in a 4x PCIe slot.	For best performance, install the adapter in an 8x Gen2 PCIe slot. Note: A 16x slot does not provide any additional performance.
0x0020/32	Informational	The network device is operating in Gen2 mode and installed in a 1x PCIe slot.	For best performance, install the adapter in an 8x Gen2 PCIe slot. Note: A 16x slot does not provide any additional performance.
0x001f/31	Informational	The network device is operating in Gen1 mode and installed in a 8x PCIe slot.	For best performance, install the adapter in an 8x Gen2 PCIe slot. Note: A 16x slot does not provide any additional performance.
0x001e/30	Informational	The network device is operating in Gen1 mode and installed in a 4x PCIe slot.	For best performance, install the adapter in an 8x Gen1 PCIe slot. Note: A 16x slot does not provide any additional performance.
0x001d/29	Informational	The network device is operating in Gen1 mode and installed in a 1x PCIe slot.	For best performance, install the adapter in an 8x Gen1 PCIe slot. Note: A 16x slot will not provide any additional performance.
0x001c/28	Error	Vital product data is not initialized correctly.	Use the offline flash utility to reconfigure the device.
0x0015/21	Warning	Firmware version does not match driver version.	The firmware version and driver must match. This is a warning message, but Broadcom recommends that you reinstall matching versions of the firmware and driver.
0x0014/20	Error	Failed to read registry configuration.	The registry is corrupted. Reinstall the driver or the operating system.
0x0013/19	Error	Resource conflict.	The operating system failed to allocate resources for the device. Check low memory conditions and operating system hardware resource conflicts.
0x0012/18	Error	Failed to enable bus mastering.	Verify that the BIOS allows bus mastering and that no resource conflicts exist.
0x0011/17	Error	The driver is incompatible with the device.	The driver is loaded on the incorrect hardware device. Verify that the correct driver is installed.
0x0010/16	Warning	The network driver was reset.	This may indicate a system hang or hardware issue. Verify that other system devices are working properly.
0x000c/12	Informational	The Ethernet link is down.	This message is informational.
0x000b/11	Informational	The Ethernet link is up.	This message is informational.
0x000a/10	Error	The network device detected an error.	A hardware error occurred. Verify that the firmware flash image is not corrupted. Contact Broadcom technical support.

Table 31 NIC Event Log Entries (Continued)

Message ID Hexadecimal/ Decimal	Severity	Message	Recommended Resolution
0x0009/9	Error	Failed to register interrupt service routine.	This is an NDIS error. Verify that hardware resource conflicts do not exist.
0x0008/8	Error	Failed to get TCP offload handlers.	This is an NDIS error. Verify that the NDIS version is valid for the driver.
0x0007/7	Warning	A memory allocation failure occurred during driver load. Performance may be reduced.	This warning occurred due to a failed memory allocation. Check low memory conditions. Use a smaller MTU or disable TCP offload to reduce driver memory requirements.
0x0006/6	Error	Driver load failed due to memory allocation failure	This failure occurred due to a failed memory allocation in the driver. Check low memory conditions.
0x0005/5	Error	Failed to register scatter gather DMA.	This failure occurred due to a failed memory allocation in the operating system. Check low memory conditions.
0x0004/4	Error	Failed to map device registers.	This failure occurred due to a failed memory allocation in the operating system. Check low memory conditions.
0x0003/3	Error	Unsupported medium.	This is an internal NDIS error. Check the operating system installation.
0x0002/2	Error	The network driver initialization failed.	This may be a firmware driver mismatch or corrupt installation. Check the firmware version, reinstall the firmware and try again. This may also indicate a hardware issue.
0x0001/1	Informational	The driver successfully loaded.	This message is informational and indicates successful loading of the device driver.

A.3 iSCSI Error and Event Log

This section provides information on iSCSi error and event logs.

A.3.1 Viewing the iSCSI Error and Event Log on Windows Server 2008

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

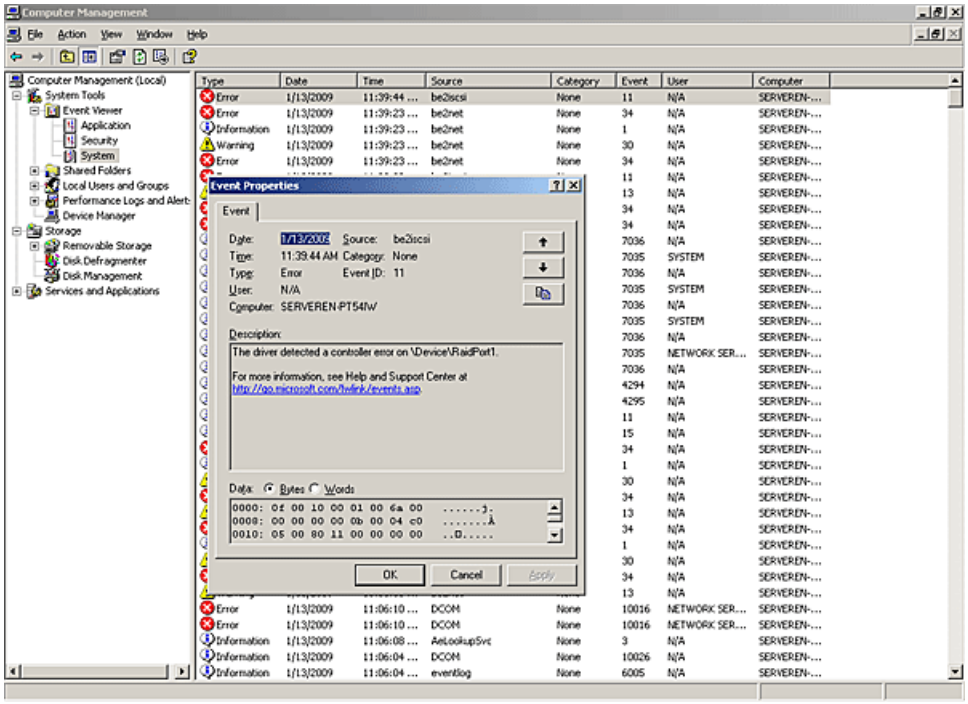
1. To view the error codes:
2. Click the Start tab on the bottom of the screen.
3. Click Run.
4. Type `eventvwr` and click OK.
5. Click Windows Log.
6. Click System.
7. Click the be2iscsi error under System Events to show the details of the event.

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

The following is an example of the iSCSI driver error code 0x11800003 viewed with the Event Viewer application (see Figure on page 132). The window shows the driver-generated error code in the fifth DWORD (offset 0x10) of the word dump.

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

Figure 30 iSCSI Error



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

A.3.2 iSCSI Error Log on Windows Server 2008

Table 32 lists brief descriptions of the error log codes generated by the iSCSI driver for Windows Server 2008. It includes the error code, the message displayed, and the meaning of the message with the recommended resolution.

Table 32 iSCSI Error Log Entries on Windows Server 2008

Message ID	Message	Description/Recommended Resolution
0x348d0008	The iSCSI driver failed a WMI IOCTL request from the port driver because the request was failed by the ARM firmware. This error is immediately followed by another error code entry indicating the WMI request code in error.	This failure indicates that an operation attempted from the Microsoft WMI application resulted in an error. Check the operation being attempted for user or configuration errors.
0x348d0007	The iSCSI driver failed a WMI IOCTL request from the port driver. This error is immediately followed by another error code entry indicating the WMI request code in error.	This failure indicates that an operation attempted from the Microsoft WMI application resulted in an error. Check the operation being attempted for user or configuration errors.
0x33900002	The iSCSI driver failed an I/O request because it failed to retrieve a scatter gather list from the Storport driver.	This failure occurred due to a failed memory allocation in the operating system. Check low memory conditions.
0x31880001	The iSCSI driver failed to load because initialization failed during a power management bootup.	This failure may be due to the firmware not being present or currently running. This failure may also indicate a hardware issue.
0x3184000c	The iSCSI driver was unable to map one or more PCI Base Address Registers and failed to load.	This failure may indicate a low memory condition or a hardware error.
0x3184000b	The iSCSI driver ignored a configuration entry because the entry was invalid.	The invalid entry must be removed or corrected. Check the registry configuration for any new valid values added to the driver parameters. For more information on valid driver values, see Table 11, iSCSI Driver Options.
0x31840009	The iSCSI driver failed to load a configuration value specified in the registry because the value was out of range. The driver will use the default value for this configuration parameter instead.	The range specified for a configuration parameter is too large or too small and must be corrected. Check the registry configuration for any new valid values added to the driver parameters. For more information on valid driver values, see Table 11, iSCSI Driver Options.
0x31840006	The iSCSI driver failed to load due to memory allocation failure.	This failure occurred due to a failed memory allocation in the driver. Check low memory conditions.
0x31840001	The iSCSI driver failed to load because initialization failed during normal bootup.	This failure may be due to the firmware not being present or currently running. This failure may also indicate a hardware issue.
0x31640004	An internal API failed in the iSCSI driver during initialization.	This failure may indicate a low memory condition.
0x3164000D	The driver failed to allocate its complete memory requirement and will attempt to load with reduced capabilities. Total number of targets available will be reduced.	This message indicates a low memory condition.
0x14831000	There was an Unrecoverable Error detected by the iSCSI driver. Following this error log entry, the next 3 entries indicate the error codes.	This may be due to hardware errors or due to unhandled exceptions in the hardware or firmware.

Table 32 iSCSI Error Log Entries on Windows Server 2008 (Continued)

Message ID	Message	Description/Recommended Resolution
0x138e0103	The iSCSI driver failed an IOCTL request because the number of scatter gather elements required for the IOCTL buffer exceeded the firmware limit. Following this error log entry, the next entry will indicate the IOCTL opcode and the payload length requested.	This error may indicate an incorrect configuration option for the iSCSI driver. It may also indicate a low memory condition.
0x138d0101	The iSCSI driver detected an error offloading the iSCSI connection. The operation will be retried again. Following this error log entry, the next entry will indicate the session handle and the firmware error code.	This may indicate a target is in error or may point to transient network connectivity issues. It may also indicate a firmware error.
0x12990013	The iscsi driver did not receive an iSCSI command window update within 25 seconds during I/O operations. Following this error log entry, the next entry will indicate the session handle where this error occurred. The iSCSI driver will trigger a session recovery on the session and continue.	■ Verify that the iSCSI target is certified by Microsoft. Check for errors reported at the target. The Emulex iSCSI adapter is only supported with certified targets. ■ Check for software updates at the target vendor's website. If applicable, update the software. ■ Check for software updates at the Broadcom website. If applicable, update the software.
0x127b0012	The iSCSI driver received an invalid iSCSI Command Sequence Number update from the target. Following this error log entry, the next three entries will indicate the session handle and the iSCSI parameters – MaxCmdSN and ExpCmdSN respectively.	■ Verify that the iSCSI target is certified by Microsoft. Check for errors reported at the target. The Emulex iSCSI adapter is only supported with certified targets. ■ Check for software updates at the target vendor's website. If applicable, update the software. ■ Check for software updates at the Broadcom website. If applicable, update the software.
0x12790006	A connection to the target was lost for a period exceeding the ETO. The error log entry immediately following this entry will indicate the session ID of the target that lost the connection. There will be event log entries from the disk subsystem indicating that the drives were lost. If any I/Os were in progress, the system may see I/O errors or failures.	Check the connection to the target or the state of the target device. If the target is made available, any sessions that existed previously are re-established and the devices are available for I/O.
0x11990007	The iSCSI driver received a TMF that is not supported and rejected this request. The error log entry immediately following this entry will indicate the TMF function code that was rejected.	The operating system version is not supported.
0x11940008	The iSCSI driver received a TMF Abort request for an I/O request that is not present with the driver.	This may indicate a slow connection to the target. Check network connectivity to the target for any errors.
0x1184000B	Firmware returned invalid data in its configuration. iSCSI login and offload are disabled.	Reload the firmware.

Table 32 iSCSI Error Log Entries on Windows Server 2008 (Continued)

Message ID	Message	Description/Recommended Resolution
0x11840002	The iSCSI driver encountered a mismatched version of the firmware running on the board. This error may be followed by error codes 0x31840001 or 0x31880001 indicating that the iSCSI driver failed to load.	This failure indicates that the driver version running on the system does not match the firmware version on the board. Correct this by running the installer from the desired version.
0x11840001	The iSCSI driver detected a failure in the hardware during initialization. This error may be followed by error codes 0x31840001 or 0x31880001 indicating that the iSCSI driver failed to load.	This failure indicates that the hardware has not been initialized or is malfunctioning. This may also indicate that the firmware is not running correctly.
0x11800005	Both Port 0 and Port 1 links were down for a period exceeding the LDTO. If the adapter has a connection to the target, there will be event log entries from the disk subsystem indicating that the drives were lost. If any I/Os were in progress, the system may see I/O errors or failures.	Check the links to the adapter. If the link is re-established, any sessions that previously existed are re-established and the devices are available for I/O.
0x11800003	Both Port 0 and Port 1 links are down.	Check the links to the adapter.
0x31840005	Driver load failed because the PCI Vendor ID and Device ID are not supported.	Check the configuration on the adapter.
0x1180000A	The logical link on the OneConnect Port is down, traffic is disallowed on this function.	The iSCSI function may have been disabled in the PXESelect utility. If you disabled it intentionally, you can ignore this message.
0x13612000	A PHY event has been detected. Following this log entry, the next log entry indicates an additional dword with four fields: bits field 0-7 link state 8-15 link effect 16-23 port name 24-31 port number	The first byte of the additional dword (link state) indicates the type of the PHY event received. link state Event 00 Physical Link is functional 01 Optics faulted /incorrectly installed/not Installed - Reseat optics, if issue not resolved, replace. 02 Optics of two types installed - Remove one optic or install a matching pair of optics. 03 Incompatible optics - Replace with compatible optics. 04 Unqualified optics - Replace with Avago optics for warranty and technical support. 05 Uncertified optics - Replace with Avago certified optics to enable link operation. other values Unrecognized optics state

A.3.3 Viewing the iSCSI Error Log on Windows 7, Windows 8, Windows 8.1, Windows Server 2008 R2, Windows Server 2012, and Window Server 2012 R2

The iSCSI driver on the Windows 7, Windows 8, Windows 8.1, Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012 R2 operating systems support the new event logging mechanism introduced by Storport. Custom event messages are logged for a variety of events with different severity, such as informational, warning, or error. The source of the events indicates the service name and every event includes a unique ID and a symbolic name. See Table 33 for a description of the error log codes.

Table 33 iSCSI Error Log Entries on Windows 7, Windows 8, Windows 8.1, Windows Server 2008 R2,Windows Server 2012, and Windows Server 2012 R2

Message ID	Severity	Message	Recommended Resolution
0x02	Info	Driver loaded successfully.	N/A
0x04	Error	Firmware version does not match with driver version.	The driver version that is running on the system does not match the version of the firmware on the adapter. Install a driver that is compatible with the firmware.
0x05	Warning	Port link is down, check connection to the adapter.	Check the links to the adapter.
0x06	Info	Port link is up.	N/A
0x07	Error	Link down timeout expired on the port, all targets are lost.	The link on the adapter is down for a period exceeding the LDTO value. If the adapter has a connection to the target, event log entries from the disk subsystem indicate that the drives were lost. If any I/O was in progress, the system may see I/O errors or failures. Check the links to the adapter. If the link is re-established, any sessions that previously existed are reestablished and the devices are available for I/O.
0x08	Error	Target with session id N failed to connect within the configured timeout.	A connection to the target was lost for a period exceeding the ETO. The error log entry includes the session ID of the target that lost the connection. Event log entries from the disk subsystem indicate that the drives were lost. If any I/O was in progress, the system may see I/O errors or failures. Check the connection to the target or the state of the target device. If the target is made available, any sessions that previously existed are re-established and the devices are available for I/O.
0x09	Error	Task Management request N was unhandled.	The iSCSI driver received a Task Management Function that is not supported, and it rejected this request. An application or service that is installed on the system may not be compatible with the driver.
0x0a	Warning	Task Management Function abort was received on a task that is not present.	The iSCSI driver received a Task Management Function Abort request for an I/O request that is not present with the driver. This may indicate a slow connection to the target. Check network connectivity to the target for any errors.
0x0b	Error	Error in determining firmware configuration.	An error in determining the firmware configuration occurred. The firmware on the adapter may not be functioning properly. Check the adapter and reinstall the firmware if required.

Table 33 iSCSI Error Log Entries on Windows 7, Windows 8, Windows 8.1, Windows Server 2008 R2,Windows Server 2012, and Windows Server 2012 R2 (Continued)

Message ID	Severity	Message	Recommended Resolution
0x0e	Warning	iSCSI error was detected on session A, ExpCmdSn B, MaxCmdSn C.	The iSCSI driver received an invalid iSCSI Command Sequence Number update from the target. The event log entry indicates the session handle, MaxCmdSN, and ExpCmdSN. ■ Verify that the iSCSI target is certified by Microsoft. Check for errors reported at the target. The Emulex iSCSI adapter is only supported with certified targets. ■ Check for software updates at the target vendor's website. If applicable, update the software. ■ Check for driver and firmware updates at the Broadcom website. If applicable, update the driver and firmware.
0x0f	Warning	The iSCSI target on session id N failed to open the command window within configured timeout.	The iSCSI driver did not receive an iSCSI command window update for up to 25 seconds during I/O operations. The event log entry indicates the session handle on which the error occurred. The iSCSI driver triggers a session recovery on the session and continues. ■ Verify that the iSCSI target is certified by Microsoft. Check for errors reported at the target. The Emulex iSCSI adapter is only supported with certified targets. ■ Check for software updates at the target vendor's website. If applicable, update the software. ■ Check for driver and firmware updates at the Broadcom website. If applicable, update the driver and firmware.
0x10	Warning	Encountered an error offloading an iSCSI connection, error code N.	The iSCSI driver detected an error while offloading the iSCSI connection. The operation is retried up to five times. The session handle and the adapter firmware error code are included in the event log message. This may indicate a target is in error or it may point to transient network connectivity issues. It may also indicate an adapter firmware error.
0x11	Warning	The IOCTL opcode A requires more scatter gather elements than allowed. Transfer length is B.	The iSCSI driver failed an IOCTL request because the number of scatter/gather elements required for the IOCTL buffer exceeded the adapter firmware limit. The IOCTL opcode and the payload length requested are included in the event log entry. This error may indicate an incorrect configuration option for the iSCSI driver. It may also indicate a low memory condition.
0x12	Error	Unrecoverable error detected. UE Low: A UE High: B FW Line: C.	An unrecoverable error was detected by the iSCSI driver. This may be caused by hardware errors or by unhandled exceptions in the hardware or firmware.

Table 33 iSCSI Error Log Entries on Windows 7, Windows 8, Windows 8.1, Windows Server 2008 R2,Windows Server 2012, and Windows Server 2012 R2 (Continued)

Message ID	Severity	Message	Recommended Resolution
0x13	Error	Hardware initialization failed, failing driver load	The iSCSI driver detected a failure in the hardware during initialization. This failure indicates that the hardware has not been initialized or is malfunctioning. This may also indicate that the firmware is not running correctly.
0x14	Warning	Failed to retrieve scatter gather list for a SCSI Request Block, an IO has failed.	The iSCSI driver failed an I/O request because it failed to retrieve a scatter/gather list from the Storport driver. This failure occurred because of a failed memory allocation in the operating system. Check low memory conditions.
0x15	Error	ACIT library table initialization failed.	An internal API failed in the iSCSI driver during initialization. This failure may indicate a low memory condition.
0x16	Error	An ACIT API failed.	An internal API failed in the iSCSI driver during initialization. This failure may indicate a low memory condition.
0x17	Error	Unsupported hardware, failing driver load.	Driver loading failed because the PCI Vendor ID and Device ID are not supported. Check the adapter configuration.
0x18	Error	Memory could not be allocated, failing driver load.	This failure occurred because of a failed memory allocation in the driver. This failure may indicate a low memory condition
0x19	Warning	WMI driver error, code A.	The iSCSI driver failed a WMI IOCTL request from the port driver. The event log entry includes the WMI request code in error. An operation attempted from the Microsoft WMI application resulted in an error. Check the operation being attempted for errors.
0x1a	Warning	WMI IOCTL error, code A.	The iSCSI driver failed a WMI IOCTL request from the port driver because the request was failed by the ARM firmware. The event log entry includes the request code in error. An operation attempted from the Microsoft WMI application resulted in an error. Check the operation being attempted for errors.
0x1b	Warning	A configuration parameter is out of range.	The iSCSI driver failed to load a configuration value specified in the registry because the value was out of range. The driver uses the default value for this configuration parameter. The range specified for a configuration parameter is either too large or too small, and it must be corrected. Check the registry configuration for new driver parameter entries. See Table 11, iSCSI Driver Options for the correct range of values

Table 33 iSCSI Error Log Entries on Windows 7, Windows 8, Windows 8.1, Windows Server 2008 R2,Windows Server 2012, and Windows Server 2012 R2 (Continued)

Message ID	Severity	Message	Recommended Resolution
0x1d	Warning	A configuration parameter is invalid.	The iSCSI driver ignored a configuration entry because the entry was invalid. Check the registry configuration for new driver parameter entries. The invalid entry must be removed or corrected. See Table 11, iSCSI Driver Options for the correct range of values.
0x1e	Error	Failed to map Base Address Register, failing driver load.	The iSCSI driver was unable to load because it was unable to map one or more PCI Base Address registers. This failure may indicate a low memory condition or a hardware error.
0x1f	Error	Hardware initialization has failed – error code A.	The hardware initialization has failed. This error causes the driver load to fail. The error code included in the event log entry identifies the specific point of failure. This failure indicates that the hardware has not been initialized or is malfunctioning. This may also indicate that the firmware is not running correctly.
0x20	Warning	Initial memory allocation failed, driver is running with reduced capabilities.	The driver failed to allocate its complete memory requirement and attempts to load with reduced capabilities. The total number of targets available is reduced. This message indicates a low memory condition.
0x21	Info	Target Reconnected for Session id N.	N/A
0x22	Info	Interrupt Redirection capability is enabled.	N/A
0x23	Warning	Interrupt Redirection capability is not supported by this firmware. Update your firmware.	Update the firmware to the latest version.
0x24	Error	Device is not supported on Windows 7 Operating System, failing driver load.	The iSCSI adapter family is not supported on the Windows 7 client operating systems.
0x25	Info	Interrupt Redirection capability is not supported by this hardware.	N/A
0x26	Warning	Logical link on the OneConnect Port is down, traffic is disallowed on this function.	The iSCSI function may have been disabled in the PXESelect utility. If you disabled it intentionally, you can ignore this message.
0x27	Error	Firmware returned invalid data in its configuration. iSCSI login and offload are disabled.	Reload the firmware.
0x28	Warning	Error Recovery is not being attempted. Adapter is no longer functional.	An unrecoverable error has occurred, but error recovery is not enabled. A system reboot is required to make the adapter operational again.
0x2b	Informational	The storage device is operating in Gen<xx> mode and installed in a <yy>x PCIe slot.	Informational message that provides the slot capabilities where the iSCSI adapter is installed.
0x2c	Error	The firmware appears unresponsive; Unrecoverable Error.	The adapter is no longer functional. A system reboot is required to make the adapter operational again.

Table 33 iSCSI Error Log Entries on Windows 7, Windows 8, Windows 8.1, Windows Server 2008 R2,Windows Server 2012, and Windows Server 2012 R2 (Continued)

Message ID	Severity	Message	Recommended Resolution
0x2d	Informational	Reporting X of the total Y sessions logged in by firmware.	Informational event to indicate that not all targets logged in by the firmware were reported as available to the operating system. If the total number of targets logged by the firmware is over the specified limits, this error can be ignored.
0x2f	Warning	Solicited command was invalidated internally due to a Data Digest error.	This condition is detected by the OneConnect firmware. If this message is unexpected, check the iSCSI configuration.
0x30	Warning	Connection was invalidated internally; the received PDU size was greater than the DSL.	This condition is detected by the OneConnect firmware. If this message is unexpected, check the iSCSI configuration.
0x31	Warning	Connection was invalidated internally; the received PDU sequence size was greater than the FBL/MBL.	This condition is detected by the OneConnect firmware. If this message is unexpected, check the iSCSI configuration.
0x32	Warning	Connection was invalidated internally; a received PDU HDR had AHS.	This condition is detected by the OneConnect firmware. If this message is unexpected, check the iSCSI configuration.
0x33	Warning	Connection was invalidated internally due to a Header Digest warning.	This condition is detected by the OneConnect firmware. If this message is unexpected, check the iSCSI configuration.
0x34	Warning	Connection was invalidated internally due to a bad opcode in the PDU header.	This condition is detected by the OneConnect firmware. If this message is unexpected, check the iSCSI configuration.
0x35	Warning	Connection was invalidated internally due to a received ITT/TTT that did not belong to this connection.	This condition is detected by the OneConnect firmware. If this message is unexpected, check the iSCSI configuration.
0x36	Warning	Connection was invalidated internally; the received ITT/TTT value was greater than the maximum supported ITTs/TTTs.	This condition is detected by the OneConnect firmware. If this message is unexpected, check the iSCSI configuration.
0x37	Warning	Connection was invalidated internally due to an incoming TCP RST.	This condition is detected by the OneConnect firmware. If this message is unexpected, check the iSCSI configuration.
0x38	Warning	Connection was invalidated internally due to TCP protocol warning (SYN received, maximum retransmits exceeded, urgent received, etc.).	This condition is detected by the OneConnect firmware. If this message is unexpected, check the iSCSI configuration.
0x39	Warning	Connection was invalidated internally due to TCP RST sent by the transmit side.	This condition is detected by the OneConnect firmware. If this message is unexpected, check the iSCSI configuration.
0x3a	Warning	Connection was invalidated internally due to an incoming TCP FIN.	This condition is detected by the OneConnect firmware. If this message is unexpected, check the iSCSI configuration.
0x3b	Warning	Connection was invalidated internally due to a bad unsolicited PDU (unsolicited PDUs are PDUs with ITT=0xffffffff).	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.

Table 33 iSCSI Error Log Entries on Windows 7, Windows 8, Windows 8.1, Windows Server 2008 R2,Windows Server 2012, and Windows Server 2012 R2 (Continued)

Message ID	Severity	Message	Recommended Resolution
0x3c	Warning	Connection was invalidated internally due to a bad WRB index.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x3d	Warning	Command was invalidated internally; the received command had residual overrun bytes.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x3e	Warning	Command was invalidated internally; the received command had residual underrun bytes.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x3f	Warning	Command was invalidated internally; a received PDU had an invalid StatusSN.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x40	Warning	Command was invalidated internally; a received R2T had invalid field(s).	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x41	Warning	Command was invalidated internally; the received PDU had an invalid LUN.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x42	Warning	Command was invalidated internally; the corresponding ICD was not in a valid state.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x43	Warning	Command was invalidated; the received PDU had an invalid ITT.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x44	Warning	Command was invalidated; the received sequence buffer offset was out of order.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x45	Warning	Command was invalidated internally; a received PDU had an invalid DataSN.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x46	Warning	Connection invalidation completion notification.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x47	Warning	Connection invalidation completion with data PDU index.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x48	Warning	Command invalidation completion notification.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x49	Warning	Unsolicited header notification.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x4a	Warning	Unsolicited data notification.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x4b	Warning	Unsolicited data digest warning notification.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.

Table 33 iSCSI Error Log Entries on Windows 7, Windows 8, Windows 8.1, Windows Server 2008 R2,Windows Server 2012, and Windows Server 2012 R2 (Continued)

Message ID	Severity	Message	Recommended Resolution
0x4c	Warning	TCP acknowledge based notification.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x4d	Warning	Connection was invalidated internally; the command and data were not on the same connection.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x4e	Warning	Solicited command was invalidated internally due to DIF warning.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x4f	Warning	Connection was invalidated internally due to an incoming unsolicited PDU that had immediate data on the connection that does not support it.	This condition is detected by the OneConnect firmware. If this message is unexpected, check your iSCSI configuration.
0x50	Informational	Physical Link is functional.	N/A
0x51	Error	Optics faulted/incorrectly installed/not installed - Reseat optics, if issue not resolved, replace.	The optics are not present.
0x52	Error	Optics of two types installed - Remove one optic or install matching pair of optics.	A mismatched pair of optics (for example:1 FC and 1 Ethernet) has been installed.
0x53	Error	Incompatible optics - Replace with compatible optics for adapter to function.	Incompatible optics (for example: FC optics in Ethernet adapter) have been used.
0x54	Warning	Unqualified optics - Replace with Avago optics for Warranty and Technical Support.	Replace with Avago optics for Warranty and Technical Support.
0x55	Error	Uncertified optics - Replace with Avago Certified optics to enable link operation.	Replace with Avago Certified optics to enable link operation.
0x56	Warning	Unrecognized optics state	This condition is reported if an unknown state for the optics is detected.

Appendix B: Configuring iSCSI through DHCP

B.1 Dynamic Host Configuration Protocol (DHCP) Recommendations

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

B.2 Vendor-Specific Option 43

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

B.2.1 Format of Vendor-Specific Option 43

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

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

- Strings shown in quotes are part of the syntax and is therefore mandatory.
- Fields enclosed in angular brackets (including the angular brackets) should be replaced with their corresponding values. Some of these fields are optional and can be skipped.
- If an optional field is skipped, a colon must be used as a placeholder to indicate the default value for that field.
- If specified, the value of each parameter should be enclosed in double quotes. See "Examples" on page 144.
- All options are case-sensitive.

B.2.2 Description of Mandatory and Optional Parameters

Table 34 describes the parameters used in the data string for option 43.

Table 34 Data String Parameters for Option 43

Parameter	Description	Field Type
<TargetIP>	A valid IPv4 address in dotted decimal notation	Mandatory
<TargetTCPPort>	A decimal number ranging from 1 to 65535 (inclusive). The default TCP port is 3260.	Optional
<LUN>	A hexadecimal representation of the LUN of the boot device. By default, LUN 0 is assumed to be the boot LUN. It is an eight-byte number that should be specified as a hexadecimal number consisting of 16 digits, with an appropriate number of zeroes padded to the left, if required.	Optional
<TargetName>	A valid iSCSI name of up to 223 characters.	Mandatory

Table 34 Data String Parameters for Option 43 (Continued)

Parameter	Description	Field Type
<InitiatorName>	A valid iSCSI iqn name of up to 223 characters. If it is not provided, the default initiator name (generated by the adapter based on its MAC address) is used.	Optional
<HeaderDigest>	Either E (the header digest is enabled) or D (the header digest is disabled).	Optional
<DataDigest>	Either E (the data digest is enabled) or D (the data digest is disabled).	Optional
<AuthenticationType>	D (authentication is disabled), E (one-way CHAP is enabled; the user name and secret must be specified by non-DHCP means), or M (mutual-CHAP is enabled; the user name and passwords must be specified by non-DHCP means). D is the default setting.	Optional

B.2.2.1 Examples

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

```
iscsi:"192.168.0.2":"3261":"000000000000000E":"iqn.2009-4.com:1234567890"::"E"::"
E"
```

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

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

```
iscsi:"192.168.0.2"::"000000000000000E":"iqn.2009-4.com:1234567890"::"E":"D":"M"
```

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

Appendix C: Port Speed Specifications

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

Table 35 lists negotiated speed specifications per an adapter's port connection:.

Table 35 Negotiated Speed Specification per Adapter Port Connection

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

C.1 Negotiating Speed on a Mezzanine Card

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

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

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

Appendix D: AutoPilot Installer Command Line and Configuration File Parameters

The AutoPilot Installer can initiate an installation from a command prompt or script. You can run the AutoPilot Installer manually from the command line or a script, or you can run it automatically through the driver kit. When run manually from the command line or script, the command line parameters can be passed.

If you specify the "/q" switch with the driver kit installer command, the driver kit installer runs in unattended mode and automatically invokes the BrcmAPInstall.exe with its "/silent" switch. See Section 2.3.6, Unattended Driver Installation for additional information.

D.1 AParg Driver Kit Parameter and Appending to the BrcmAPInstall.exe File

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

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

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

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

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

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

This results in the AutoPilot Installer being run as:

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

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

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

This results in AutoPilot Installer being run as:

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

To pass multiple parameters to the AutoPilot Installer and minimize errors, you can run the utility kit installer interactively, delay AutoPilot Installer execution, and then run the AutoPilot Installer command. The procedure is described in Section 2.3.2.2, Option 2: Run AutoPilot Installer Separately and Section 2.3.6, Unattended Driver Installation.

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

```
brcmdrvr-fcoe<version>.exe install:"C:\Broadcom"
```

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

D.2 AutoPilot Installer Syntax

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

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

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

The syntax of a parameter setting is:

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

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

The APInstall command can contain the settings listed below. Each setting, except ConfigFileLocation, can also be specified in the AutoPilot Configuration file. For descriptions of each parameter, see "Software Configuration Parameters" on page 148.

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

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

D.2.1 Path Specifiers

Paths can be specified as:

- An explicit path:

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

- A relative path:

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

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

- With the %ProgramFiles% environment variable:

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

D.2.2 Configuration File Location

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

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

Example:

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

D.2.3 Software Configuration Parameters

D.2.3.1 DiagEnable (Running Diagnostics)

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

Default: True

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

D.2.3.2 ForceDriverTypeChange (Forcing a Driver Type Change)

Default: False

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

D.2.3.3 ForceDriverUpdate (Forcing a Driver Version Update)

Default: False

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

NOTE ForceDriverUpdate applies only to unattended installations; in interactive installations, this parameter is ignored and you are asked if the driver should be updated.

D.2.3.4 ForceRegUpdate (Forcing an Update of an Existing Driver Parameter Value)

Default: False

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

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

D.2.3.5 LocalDriverLocation (Specifying Location to Search for Drivers)

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

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

Example:

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

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

D.2.3.6 NoSoftwareFirstInstalls (Prohibiting Software First Installations)

Default: False

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

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

D.2.3.7 ReportLocation (Setting Up an Installation Report Title and Location)

The automatically generated file name for this report is

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

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

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

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

D.2.3.8 SilentInstallEnable (Enabling Unattended Installation)

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

Default: False

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

D.2.3.9 SilentRebootEnable (Enabling Silent Reboot)

Default: False

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

D.2.3.10 InstallWithoutQFE (Enabling Installation if a QFE Check Fails)

Default: False

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

D.3 AutoPilot Configuration File

The AutoPilot configuration file is organized into sections, grouped according to related commands. Six main sections exist:

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

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

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

D.3.1 Using the Windows Environment Variable (%ProgramFiles%)

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

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

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

D.3.2 Configuration Identification [AUTOPILOT.ID]

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

D.3.3 Software Configuration [AUTOPILOT.CONFIG]

This section can contain settings that control and configure AutoPilot Installer and the OneCommand CNA Manager application operation. This section can appear only once in the AutoPilot configuration file. See "Software Configuration Parameters" on page 148 for information about settings that can be specified in this section.

D.3.4 Configuration Prompts/Vendor-Specific Questions [STORPORT.CONFIGURATION]

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

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

```
question= “question?”, “explanation”, “answer0”, “answer1”, “answer2”,.... ,
“answern”
```

Where:

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

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

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

D.3.4.1 Example of [STORPORT.CONFIGURATION] section:

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

D.3.5 QFE Checks [STORPORT.QFES]

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

```
qfe = "qfe name", "path and file name", "file version", "applicable OS"
```

qfe name	The name of the item being checked. For example, QFE 2846340. The name should facilitate searching Microsoft's website for any required code updates.
path and file name	This string identifies the file to be checked and its location relative to the Windows home folder. In most cases, the file to check is the Microsoft Storport driver. For example: "\system32\drivers\storport.sys". This string is also used in dialogs and log file messages.
file version	This is the minimum version that the file to be checked must have for the QFE to be considered installed. It is specified as a text string using the same format as is used when displaying the files property sheet. For example: "5.2.1390.176".
applicable OS	This is used to determine if the QFE applies to the operating system platform present. The acceptable value is "Win2008".

For example:

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

D.3.6 Setting Up FC Driver Parameters [STORPORT.PARAMS]

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

For example:

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

See Table 2, Storport Miniport Driver Parameters for a listing of driver parameters, defaults, and valid values.

D.3.7 Setting Up System Parameters [SYSTEM.PARAMS]

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

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

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

D.4 AutoPilot Installer Exit Codes

AutoPilot Installer sets an exit code to indicate whether an installation was successful or an error occurred. These exit codes allow AutoPilot Installer to be used in scripts with error handling. In unattended installations, AutoPilot Installer sets the following exit codes listed in Table 36.

Table 36 Unattended Installation Error Codes

Error Code	Hex	Description
0	0x00000000	No errors.
2399141889	0x8F000001	Unsupported operating system detected.
2399141890	0x8F000002	The AutoPilot Configuration file is not found.
2399141891	0x8F000003	Disabled adapters detected in the system.
2399141892	0x8F000004	The selected driver is 64-bit and this system is 32-bit.
2399141893	0x8F000005	The selected driver is 32-bit and this system is 64-bit.
2399141894	0x8F000006	Installation activity is pending. AutoPilot Installer cannot run until it is resolved.
2399141895	0x8F000007	(GUI Mode only) You cancelled execution because you did not wish to perform a software-first install.
2399141896	0x8F000008	No drivers found.
2399141897	0x8F000009	One or more adapters failed diagnostics.
2399141904	0x8F000010	(GUI Mode only) You chose to install drivers even though a recommended QFE or Service Pack was not installed.
2399141920	0x8F000020	(GUI Mode only) You chose to stop installation because a recommended QFE or Service Pack was not installed.
2399141899	0x8F00000B	Unattended installation did not find any drivers of the type specified in the config file.
2399141900	0x8F00000C	A silent reboot was attempted, but according to the operating system a reboot is not possible.
2399141901	0x8F00000D	(GUI Mode only) A driver package download was cancelled.
2399141902	0x8F00000E	(Non-Enterprise) No adapters were found in the system.
2399141903	0x8F00000F	A required QFE or Service Pack was not detected on the system.
2399141836	0x8F000030	AutoPilot Installer was not invoked from an account with Administrator-level privileges.
2391419952	0x8F000040	AutoPilot Installer has detected unsupported adapters on the system.
2399141968	0x8F000050	Unattended software-first installations are disallowed.
2399141984	0x8F000060	You cancelled APInstall before any driver/utility installation occurred.
2399142000	0x8F000070	You cancelled APInstall after driver/utility installation occurred.
2399142032	0x8F000090	APInstaller encountered an error while parsing the command line (Report file contains details).

D.5 AutoPilot Installer Installation Reports

During each installation, the AutoPilot Installer produces a report describing events that occurred during the installation. This report contains the following sections:

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

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

D.6 Command Script Example

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

```
rcmd \\server1 g:\Broadcom\autopilot installer\fcoe\brcmapinstall.exe
if errorlevel 1 goto server1ok
echo AutoPilot reported an error upgrading Server 1.
if not errorlevel 2147483650 goto unsupported
  echo Configuration file missing.
goto server1ok
 :unsupported
if not errorlevel 2147483649 goto older
echo Unsupported operating system detected.
 :older
if not errorlevel 2001 goto none
   echo The driver found is the same or older than the existing driver.
   goto server1ok
:none
if not errorlevel 1248 goto noreport
   echo No Emulex adapter found.
goto server1ok
:noreport
 if not errorlevel 110 goto nocfg
   echo Could not open installation report file.
 goto server1ok
:nocfg
 if not errorlevel 87 goto badcfg
   echo Invalid configuration file parameters.
   goto server1ok
:badcfg
   if not errorlevel 2 goto server1ok
   echo No appropriate driver found.
server1ok
rcmd \\server2 g:\autopilot\ApInstall
ConfigFileLocation=g:\autopilot\mysetup\brcmapinstall.cfg
```

```
if errorlevel 1 goto server2ok
echo AutoPilot reported an error upgrading Server 2.
if not errorlevel 2147483650 goto unsupported
   echo Configuration file missing.
goto server2ok
 :unsupported
 if not errorlevel 2147483649 goto older
   echo Unsupported operating system detected.
 :older2
 if not errorlevel 2001 goto none2
   echo The driver found is the same or older than the existing driver.
   goto server2ok
:none2
 if not errorlevel 1248 goto noreport2
   echo No adapter found.
goto server2ok
:noreport
 if not errorlevel 110 goto nocfg2
   echo Could not open installation report file.
 goto server2ok
:nocfg2
 if not errorlevel 87 goto badcfg2
   echo Invalid configuration file parameters.
   goto server2ok
:badcfg2
   if not errorlevel 2 goto server2ok
   echo No appropriate driver found.
server2ok
```

Appendix E: RoCE Switch Support

Some switches do not support DCBX, and most DCBX-enabled switches do not fully support RoCE as a protocol. At this time, none of the known switch vendors (Arista, Brocade, Cisco, and Juniper) allow configuring priority for RoCE-specific traffic. Additionally, most of the known switch vendors do not support APP TLV of 0x8915 for RoCE ETS bandwidth and PFC configuration.

In addition to QOS settings, OCe14000-series adapters support congestion management protocols. The supported modes are QCN for RoCE ports or ECN for Routable RoCE ports.

NOTE For each RoCE port, the RoCE traffic must be limited to either RoCE or Routable RoCE (Default).

E.1 DCBX-Enabled Switch Connection PFC Mode

Manually enable Priority 5 on the switch under a different priority group other than FCOE/ISCSI/NIC priority group.

NOTE If an OCe14000-series adapter is connected to a DCBX-enabled switch, the mode shifts from generic pause to PFC mode.

NOTE If an OCe14000-series adapter is connected to a DCBX-disabled switch, generic pause mode is enabled.

NOTE In absence of Priority 5 on the switch side, the OCe14000-series adapter maintains configuration for RoCE and PFC Priority 5. This can result in packet losses, unrecoverable errors, or infinite retries for RoCE traffic.

E.1.1 Switch Configuration for PFC Priority 5

Using the documentation provided by your switch vendor, configure your switch for the following:

- Priority pause frames using Priority 5
- MTU size of 4200 or higher
- No-drop policy

To configure the switch:

1. Create PG 1 as Priority 5 with a no-drop policy for RoCE traffic.
2. Assign the appropriate bandwidth to PG 1; for example, 90%.
3. Create PG 2 (or something different from above, which is priority group 1) and assign NIC traffic to it.
4. Assign remaining bandwidth to PG 2.
5. Enable priority flow control on all ports participating in the cluster and at a global level in the switch.

   NOTE Some switches have global and port level settings for flow control and bandwidth allocation. Make sure the PFC flow control setting is performed on all the ports that participate in the cluster.
6. Configure a valid VLAN with an ID other than 0 or 1.
7. Ensure that Jumbo Frames is enabled, or at minimum set the MTU >= 4200.
8. Specify each switch port service policy rather than using the system QoS.

NOTE Some switches have jumbo frame size support disabled on the port and global level by default.

NOTE Some switches show the priority for FCoE on the switch itself. Use a policy with zero bandwidth for the FCoE priority.

E.1.2 Host—Client Configuration

For all host and clients participating in the network:

1. Create a VLAN using the VLAN ID you configured in step 6 above.
2. Assign an appropriate IP address to the VLAN interface.

NOTE Ensure that all the traffic is flowing through the VLAN interface.

E.1.2.1 DCBX-Disabled Switch Connection (Generic Pause Mode)

1. Host Configuration:
   a. On the host and peer systems, ensure that TX pause flow control and RX pause flow control are enabled on all the ports and interfaces that are RoCE enabled using operating system standard tools.
2. Switch Configuration:
   a. Enable TX generic pause flow control and RX generic pause flow control on each port participating in the cluster.
   b. Enable Jumbo Frames, or set the MTU to at least 4200 or greater.

NOTE Some switches have jumbo frame size support disabled on the port and global level by default.

E.1.2.2 Examples for Cisco Switch

This section provides information for configuring a Cisco switch.

E.1.2.3 Sample Class-maps for RoCE on a Cisco Switch

NOTE Not all switch settings are shown.

Cisco5548UP2(config)# show class-map

```
Type qos class-maps
===================
class-map type qos match-any class-fcoe
match cos 3
class-map type qos match-all class-roce
match cos 5
class-map type qos match-any class-default
match any

Type queuing class-maps
=======================
class-map type queuing class-fcoe
match qos-group 1
class-map type queuing class-roce
match qos-group 5
class-map type queuing class-default
match qos-group 0
```

```
Type network-qos class-maps
===========================
class-map type network-qos class-fcoe
match qos-group 1
class-map type network-qos class-roce
match qos-group 5
class-map type network-qos class-default
match qos-group 0
```

E.1.2.3.1 Sample Policy-maps for RoCE on a Cisco Switch

NOTE Not all switch settings are shown.

```
Type qos policy-maps
====================
policy-map type qos class-roce
class type qos class-roce
set qos-group 5
class type qos class-fcoe
set qos-group 1
class type qos class-default
set qos-group 0
policy-map type qos class-rocenofcoe
class type qos class-roce
set qos-group 5
class type qos class-default
set qos-group 0

Type queuing policy-map
=======================
policy-map type queuing class-roce90
class type queuing class-roce
bandwidth percent 90
class type queuing class-default
bandwidth percent 10

Type network-qos policy-maps
============================
policy-map type network-qos class-rocenofcoe
class type network-qos class-roce
pause no-drop
mtu 4200
class type network-qos class-default
mtu 9216
multicast-optimize
```

E.1.2.3.2 Sample Port Configuration for RoCE on a Cisco Switch

NOTE The Port flow control should be off and the Priority Flow Control should be on Auto. PFC flow is not explicitly displayed on this switch.

```
interface Ethernet1/15
description RoCE configuration
switchport mode trunk
switchport trunk allowed vlan 102
spanning-tree port type edge trunk
```

```
service-policy type qos input class-rocenofcoe
service-policy type queuing input class-roce90
service-policy type queuing output class-roce90
```

E.1.2.3.3 Sample Switch PFC Verification on a Cisco Switch

Ensure that the "Mode" is set to "Auto." and "Operational (Oper)" is "On".

```
Cisco5548UP2(config)# show int eth 1/11 priority-flow-control
================================================================
Port               Mode Oper(VL bmap)  RxPPP      TXTXPPP
================================================================

Ethernet1/11       Auto On  (28)      873        694950
```

E.1.2.4 Verifying Switch Configuration in OneCommand Manager

NOTE You do not need to configure the OCe14000-series adapter in the OneCommand Manager application to enable RoCE with PFC.

You can use the OneCommand Manager GUI application or the OneCommand Manager CLI application to verify the switch configuration.

See the *Emulex OneCommand Manager Application User Manual for OneConnect Adapters* for more information on using the OneCommand Manager GUI application to verify the switch configuration.

See the *Emulex OneCommand Manager Command Line Interface User Manual for OneConnect Adapters* for more information on using the OneCommand Manager CLI application to verify the switch configuration.

Appendix F: License Notices

F.1 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.
*/


BROADCOM

Emulex® Drivers for Linux for OneConnect® Adapters

User Guide

Version 11.2
December 30, 2016

DRVLin-OCA-UG112

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

Broadcom, the pulse logo, Connecting everything, Avago Technologies, Avago, the A logo, Emulex, OneCommand, and OneConnect are the trademarks of Broadcom in the United States, certain other countries and/or the EU.

Copyright © 2003–2017 Broadcom. All Rights Reserved.

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

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.

Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Chapter 1: Introduction

1.1 Overview

This guide provides installation, updating, uninstalling, configuring, and troubleshooting procedures for several types of Emulex®-supported drivers for Linux.

This guide is applicable to several versions of drivers, operating systems, firmware, and adapters.

- For supported firmware versions and their latest release, refer to the Documents and Downloads area of http://www.broadcom.com for the specific adapter.
- This product supports the following Emulex OneConnect® converged network adapters (CNAs):
  - OCe11000-series adapters
  - OCe14000-series adapters

1.2 Abbreviations

AIC	adaptive interrupt coalescing
ANSI	American National Standards Institute
API	application programming interface
ARI	alternative routing-ID interpretation
ASIC	application-specific integrated circuit
BIOS	basic input/output system
CLI	command line interface
CNA	converged network adapter
CPU	central processing unit
CQ	completion queue
CQE	completion queue entry
DAPL	Direct Access Programming Library
DCBX	Data Center Bridging Capabilities Exchange
DHCHAP	Diffie-Hellman Challenge Handshake Authentication Protocol
DIF	Data Integrity Field
DIMM	dual in-line memory module
DMA	direct memory access
EQ	event queue
ETO	extended timeout
FC	Fibre Channel
FCF	Fibre Channel over Ethernet Forwarder
FCFI	Fibre Channel Forwarder Indicator
FCoE	Fibre Channel over Ethernet
FCP	Fibre Channel Protocol

FSB	front side bus
FW	firmware
GbE	gigabit Ethernet
Gb/s	gigabits per second
GID	group identifier
GnuPG	GNU Privacy Guard
GRO	Generic Receive Offload
GUI	Graphical User Interface
HBA	host bus adapter
HII	Human Interface Infrastructure
IB	Infiniband
IEEE	Institute of Electrical and Electronics Engineers
I/O	input/output
IOCB	input/output control block
IOPs	I/O operations per second
IOV	I/O virtualization
IP	Internet Protocol
IPL	initial program load
IQN	iSCSI qualified name
IRQ	interrupt request
iSCSI	internet Small Computer System Interface
iSER	iSCSI Extensions for RDMA
iSNS	Internet Storage Name Service
KVM	kernel-based virtual machine
LACP	Link Aggregation Control Protocol
LAN	local area network
LBA	logical block address
LDTO	link down timeout
LLC	logical link control
LOM	LAN on motherboard
LPVID	logical port VLAN ID
LRO	large receive offload
LUN	logical unit number
MAC	Media Access Control
Mb/s	megabits per second
MPIO	multipath I/O
MR	memory region
MSI	message signaled interrupts
MSI-X	message signaled interrupts - extended
MTU	maximum transmission unit

NAA	Network Address Authority
NCSI	Network Communication Services Interface
NFS RDMA	network file system over RDMA
NIC	network interface card (or controller)
NPar	NIC partitioning
NPIV	N_Port ID Virtualization
NUMA	non-uniform memory access
OFA	Open Fabric Alliance
OFED	OpenFabrics Enterprise Distribution
PCIe	Peripheral Component Interconnect Express
PD	protection domain
PF	physical function
PFC	priority flow control
POST	power-on self-test
QCN	quantized congestion notification
QoS	quality of service
QP	queue pair
RDMA	remote direct memory access
RHEL	Red Hat Enterprise Linux
RoCE	RDMA over converged Ethernet
RPI	remote port indicator
RPM	resource package manager
RQ	receive queue
RSCN	registered state change notification
RSS	receive-side scaling
Rx	receive
SAN	storage area network
SCSI	Small Computer System Interface
SGE	Oracle (formerly Sun) grid engine
SLES	SUSE Linux Enterprise Server
SLI	Service Level Interface
SNAP	Subnetwork Access Protocol
SNMP	Simple Network Management Protocol
SQ	send queue
SR-IOV	single-root I/O virtualization
SRQ	shared receive queue
tar	tape archive
TCP	Transmission Control Protocol
TSO	TCP Segmentation Offload
Tx	transmit

UDP	User Datagram Protocol
UEFI	Unified Extensible Firmware Interface
UMC	Universal multichannel
VEB	virtual Ethernet bridging
VEPA	virtual Ethernet port aggregator
VF	virtual function
VGT	virtual guest tagging
VLAN	virtual local area network
VM	virtual machine
VPD	vital product data
VPort	virtual port
VXLAN	Virtual eXtensible LAN
WWPN	World Wide Port Name
XRI	extensible resource indicator

Chapter 2: Installing and Uninstalling

Emulex releases Linux binary RPMs that are digitally signed using the GNU Privacy Guard (GnuPG) standard. This signature will allow certification of the contents of the RPMs and verification that the contents have not been modified since they were created by Emulex. The RPMs have been digitally signed by Emulex with a GnuPG private key that is only held by Emulex. Instructions for creating the GnuPG public key file are located at:

http://www.broadcom.com/docs/elx-rpm-public-key

2.1 General Installation Requirements

NOTE You must install the driver before updating the firmware.

Prior to driver installation, follow these general requirements:

- Install a supported Emulex adapter in the system. Refer to the adapter installation guide for specific hardware installation instructions.
- Use a supported operating system. Refer to the Documents and Downloads area of http://www.broadcom.com for supported operating systems.

2.2 Binary RPM FCoE Driver Kit

The binary resource package manager (RPM) Fibre Channel over Ethernet (FCoE) driver kit contains the following:

- A zipped tape archive (tar) file that includes the driver binary RPMs for a specific driver version and Linux distribution

  NOTE Use only officially released Linux distribution kernels. The binary RPM packages only support officially released Linux distribution kernels and do not support pre-release distribution kernels.
- An installation script, `brcmfcoe_install.sh`, that installs by default the FCoE driver binary RPM that corresponds to the target system's architecture and kernel memory variant
- A README file that provides a description of the kit structure, its contents, and distribution support scope

2.2.1 Installing the Binary RPM FCoE Driver Kit

NOTE You must uninstall any driver kits that are not part of this distribution. For example, you must uninstall any previous FCoE driver kits before installing this driver kit. This installation fails if a previous version of the FCoE driver kit is detected. For more information, see Section 2.2.2, Uninstalling the Binary RPM FCoE Driver Kit.

To install the binary RPM FCoE driver:

1. Download the appropriate driver kit from the Documents and Downloads area of http://www.broadcom.com.
2. Log in as *root* to a terminal, and unpack the tarball:

   `tar xzf brcm-lpfc-dd-<Linux distribution version>-<driver version>.tar.gz`
3. Change to the directory into which the tarball was extracted:

   `cd brcm-lpfc-dd-<Linux distribution version>-<driver version>/`

4. Run the `brcmfcoe_install.sh` script without options to install the driver kit:

   ```
   ./brcmfcoe_install.sh
   ```

   After the `brcmfcoe_install.sh` script has finished running successfully, the Emulex Fibre Channel (FC) and FCoE driver is loaded, and devices that are connected to the system are accessible.
5. Reboot the system to enable the newly added driver options in the ramdisk. You can also reboot the system later.

2.2.2 Uninstalling the Binary RPM FCoE Driver Kit

NOTE You must run the uninstall script that shipped with the version of the driver kit you want to remove.

To uninstall the binary RPM FCoE driver:

1. Log in as *root*.
2. If possible, exit all applications that use FCoE-attached drives, then unmount the drives. If you cannot exit all applications that use FCoE-attached drives, the uninstall script works properly, but you must reboot after the uninstall is complete.
3. Run the `brcmfcoe_install.sh` script with the `--uninstall` option:

   ```
   ./brcmfcoe_install.sh --uninstall
   ```

2.3 Ethernet Driver Kit

The Ethernet driver kit includes the driver that supports the network interface card (NIC) protocol. The Ethernet driver kit contains the following:

- A zipped tar file that includes the driver binary RPMs for a specific driver version, and for all of the supported Linux distribution kernels

  NOTE Use only officially released Linux distribution kernels. The binary RPM packages only support officially released Linux distribution kernels and do not support pre-release distribution kernels.
- An installation script, `elx_net_install.sh`, which installs (by default) the Ethernet driver binary RPM that corresponds to the target system's architecture and kernel memory variant
- A README file that provides a description of the kit structure, its contents, and distribution support scope

2.3.1 Installing the Ethernet Driver Kit

NOTE Remove any previously installed Ethernet driver kits (that is, those that were not part of a distribution's kernel), before proceeding. For more information on uninstalling the driver, see Section 2.3.2, Uninstalling the Ethernet Driver Kit.

To install the Ethernet driver:

1. Download the appropriate driver kit from the Documents and Downloads area of http://www.broadcom.com.
2. Log in as *root* to a terminal, and unpack the tarball:

   ```
   tar xzf elx-be2net-dd-<driver version>.tar.gz
   ```
3. Change to the directory that is extracted:

   ```
   cd elx-be2net-dd-<driver version>/
   ```

4. Run the `elx_net_install.sh` script without options to install the driver kit:

   ```
   ./elx_net_install.sh
   ```

   After the `elx_net_install.sh` script has finished running successfully, the Emulex Ethernet driver is loaded, and devices that are connected to the system are accessible.
5. Reboot the system to enable the newly added driver options in the ramdisk. You can also reboot the system later.

2.3.2 Uninstalling the Ethernet Driver Kit

NOTE You must run the uninstall script that shipped with the version of the driver kit you want to remove.

To uninstall the Ethernet driver:

1. Log in as *root*.
2. If possible, exit all applications that use Ethernet-attached drives, then unmount the drives. If you cannot exit all applications that use Ethernet-attached drives, the uninstallation works properly, but you must reboot after the uninstallation is complete.
3. Run the `elx_net_install.sh` script with the `--uninstall` option:

   ```
   ./elx_net_install.sh --uninstall
   ```

2.4 RoCE Driver for the OCe14000-Series Adapters

NOTE RDMA over converged Ethernet (RoCE) is available as Technical Preview only. Do not use RoCE in a production environment.

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

Network-intensive applications, such as networked storage or cluster computing require, a network infrastructure with high bandwidth and low latency. The advantages of remote direct memory access (RDMA) over other network application programming interfaces are lower latency, lower central processing unit (CPU) load, and higher bandwidth.

For a list of operating systems supported for RoCE, refer to the latest *Emulex Driver for Linux for OneConnect Adapters Release Notes*.

2.4.1 Installing the Infiniband Stack

Before installing the OpenFabrics Enterprise Distribution (OFED) package/Infiniband stack (IB stack) and RoCE driver/libraries, ensure that the OCe14000 adapter uses the NIC+RoCE profile. You can verify this by pressing **Ctrl+P** during system booting, or by checking the profile using the Emulex OneCommand® CNA Manager application, or by inspecting the Ethernet driver load messages in `/var/log/messages` as detailed in Section 3.3.1.2, Confirming that the RoCE Profile Is Enabled.

Before you proceed with installation, ensure that the installed operating system supports the driver. Check the latest *Emulex for Linux Driver for OneConnect Adapters Release Notes* for the list of supported operating systems.

2.4.1.1 Installing the IB Stack from the OFA Distribution

The Open Fabric Alliance (OFA) distributes the IB stack in the form of a OFED tarball. For a list of available OFED versions that are compatible with the supported operating systems, refer to the latest Emulex *Driver for Linux for OneConnect Adapters Release Notes*.

Supported OFED packages can be downloaded from the following website:

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

To install the OFED package:

1. Download the appropriate tarball file to the `/tmp` directory.
2. Extract the downloaded `OFED-x.tgz` tarball.

```
# cd /tmp
# tar -xvzf OFED-x.tgz
```

3. Install the OFED package by typing:

```
#cd /tmp/<OFED-version>
```

— To install the OFED package in SUSE Linux Enterprise Server (SLES) 12 platforms, type:

```
# ./install.pl --without-ocrdma --without-be2net
--without-infiniband-diags
```

NOTE If the `infiniband-diags` package is removed, the ibstat and ibxxx tools are not available.

— To install the OFED package in any other platform, type:

```
# ./install.pl --without-ocrdma --without-be2net
```

4. When installation of the OFED stack is complete, reboot the system.

NOTE The OFED installation process might list a few missing dependency packages that you can install from the operating system CD.

2.4.1.2 Installing the IB Stack from the Operating System Distribution

To install the IB stack during the operating system installation process:

1. Select **Infiniband Support** in the customization menu. The Infiniband Support option will install all IB stack kernel components, user space libraries, and utilities.

   The following list of RPM packages are typically present in a fully configured IB system:

   Binary packages

   — `librdmacm-utils`
   — `libibverbs-utils`
   — `libibverbs`
   — `libibmad`
   — `librdmacm`
   — `libibumad`
   — `perftest`
   — `qperf`

   Development packages

   — `libibumad-devel`
   — `libibverbs-devel`
   — `librdmacm-devel`

   To verify whether the RPM packages are present in your system, type:

```
# rpm -qa | grep <rpm name>
```

2.4.1.3 Installing the IB stack using the yum groupinstall Command

If the operating system has already been installed or you did not choose the **Infiniband Support** option while installing the operating system, you can use the `yum groupinstall` command to install the IB stack. Before you

run the `yum groupinstall` command, a yum repository must be set up. For information on setting up the yum repository, refer to the Red Hat Enterprise Linux (RHEL) documentation.

To install the IB stack using the `yum groupinstall` command, type:

```
# yum groupinstall "Infiniband Support"
```

NOTE This command resolves any RPM package dependencies that might exist.

2.4.1.4 Installing the RoCE Driver Using the elx_roce_install.sh Script

NOTE Always use the `elx_roce_install.sh` script to install the RoCE driver.

To install the ROCE driver using the `elx_roce_install.sh` script:

1. Download the Emulex RoCE driver package from the Documents and Downloads area of http://www.broadcom.com.
2. Use the following package to install the driver:
   `elx-ocrdma-dd-<release>-<version>.tar.gz`
3. Copy the package to `/tmp` and run:
   `tar -xvzf elx-ocrdma-dd-<release>-<version>.tar.gz`
4. Change directory to the RoCE packages directory:
   `cd /tmp/elx-<release>-ocrdma-dd-<version>`
5. Select the required package based on the host on which the driver is being installed. The current Linux distribution on the host might be determined as follows:
   On an RHEL host, type:
   `# cat /etc/redhat-release`
   On an SLES host, type:
   `# cat /etc/SuSE-release`
6. Select the appropriate package based on the output of step 5.
7. Change the directory to the resulting directory from the tar extraction command to find the installer script.
   `./elx_roce_install.sh`
8. Run the installer script to install the required RoCE driver and user library RPMs.
   `#./elx_roce_install.sh`

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

- Set `allow_unsupported_modules` to `1` in `/etc/modprobe.d/unsupported-modules`.
- Specify `--allow-unsupported` on the command line.

2.4.1.5 Manually Installing the Out-of-Box Driver rpms (OFED-3.12 or Later)

In situations where the `elx_roce_install.sh` script cannot be used, use the manual procedure to install the RoCE driver.

Use the `elx_roce_install.sh` script to install the out-of-box driver rpms. If the `elx_roce_install.sh` script is not used to update the drivers and the inbox OFED drivers (be2net and ocrdma) are also present, use the

following script to update and load the drivers. The following script prevents the situation where fresh OOB rpms are installed, but the `modinfo be2net/ocrdma` command displays the older driver versions.

```
Goto /lib/modules/<kernel
version>/updates/drivers/net/ethernet/emulex/benet
If be2net.ko exists /bin/rm be2net.ko
Goto /lib/modules/<kernel version>/updates/drivers/infiniband/hw/ocrdma/
If ocrdma.ko exists /bin/rm ocrdma.ko
rpm –ivh kmod-be2net-11.0.xx.x.rpm
rpm -ivh --nodeps kmod-ocrdma-11.0.xxx.ofed-xxx.x.rpm
rpm -ivh libocrdma-11.0.xx.ofed-xx.rpm
rmmod be2net
modprobe be2net
modprobe ocrdma
```

2.4.2 Uninstalling the RoCE Driver

Before you attempt to uninstall the RoCE driver, you must perform these steps:

1. Stop all active network file system over RDMA (NFS RDMA) mounts and internet Small Computer System Interface (iSCSI) Extensions for RDMA (iSCSI Extensions for RDMA [iSER]) sessions.
2. Stop the ibacm service on systems where OFED is installed using the command:

   ```
   # service ibacm stop
   ```
3. Stop all user space RoCE applications.

To uninstall the RoCE driver, type:

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

2.5 iSCSI Driver Kit

The iSCSI driver kit includes the driver that supports the iSCSI protocol. The iSCSI driver kit contains the following:

- A zipped tar file that includes the binary RPMs for a specific driver version, and for all of the supported Linux distribution kernels

  NOTE Use only officially released Linux distribution kernels. The binary RPM packages only support officially released Linux distribution kernels and do not support pre-release distribution kernels.
- An installation script, `elx_iscsi_install.sh`, that installs by default the iSCSI driver binary RPM that corresponds to the target system's architecture and kernel memory variant
- A README file that provides a description of the kit structure, its contents, and distribution support

2.5.1 Compiling the iSCSI Driver

To compile the iSCSI driver from source, you must run the following `make` command from the driver `src` directory:

```
make -C /usr/src/<kernel dir> M=`pwd` CONFIG_BE2ISCSI=m
```

Where the <*kernel dir*> is:

For RHEL 6.5 SS#:

```
/usr/src/kernels/2.6.32-431.el6/
```

For SLES 11 SP# default variant:

```
/usr/src/linux-obj/x86_64/default
```

To compile and install the iSCSI driver, use following instructions:

```
make -C /lib/modules/`uname –r` M=`pwd` CONFIG_BE2ISCSI=m modules
make -C /lib/modules/`uname –r` M=`pwd` CONFIG_BE2ISCSI=m modules_install
```

2.5.2 Installing the iSCSI Driver Kit

Remove any previously installed iSCSI driver kits, Application Helper Modules (that is, those that were not part of a distribution's kernel), or both before proceeding. For more information on uninstalling the driver, see Section 2.5.3, Uninstalling the iSCSI Driver Kit.

When using this driver for storage area network (SAN)-boot applications, the following dependencies must be installed:

- iscsi-initiator-utils
- dracut-network
- For SAN boot with an SLES operating system, kernel parameter `withiscsi=1` should be passed during installation.

Because this driver is based on Open-iSCSI, attempting to perform a SAN-boot without these dependencies will fail.

To install the iSCSI driver:

1. Download the appropriate driver kit from the Documents and Downloads area of http://www.broadcom.com.
2. Log in as *root* to a terminal, and unpack the tarball:

   ```
   tar xzf elx-be2iscsi-dd-<driver version>.tar.gz
   ```
3. Change to the directory that is extracted:

   ```
   cd elx-be2iscsi-dd-<driver version>/
   ```
4. Run the `elx_iscsi_install.sh` script with no options to install the driver kit:

   ```
   ./elx_iscsi_install.sh
   ```
5. After the `elx_iscsi_install.sh` script has finished running successfully:
   - For an iSCSI boot case, you must reboot the system now to load the driver.
   - For all other iSCSI cases, the Emulex iSCSI driver is loaded, and devices that are connected to the system are accessible. Reboot the system now to enable the newly added driver options in the ramdisk. You can also reboot the system later.

2.5.3 Uninstalling the iSCSI Driver Kit

NOTE You must run the uninstall script that shipped with the version of the driver kit you want to remove.

To uninstall the iSCSI driver:

1. Log in as *root*.
2. If possible, exit all applications that use iSCSI-attached drives, then unmount the drives. If you cannot exit all applications that use iSCSI-attached drives, the uninstallation works properly, but you must reboot after the uninstallation is complete.
3. Run the `elx_iscsi_install.sh` script with the `--uninstall` option:

   ```
   ./elx_iscsi_install.sh --uninstall
   ```

2.6 Booting from a Nonzero LUN Attached to a Emulex FCoE Adapter

To configure SLES 11 SPx to boot from an FCoE-attached disk device other than `/dev/sda`, refer to the *Boot for NIC, iSCSI, and FCoE Protocols User Guide* which is available at the Documents and Downloads area of http://www.broadcom.com.

2.7 Emulex OneCommand CNA Manager Application

The Emulex OneCommand CNA Manager application is a powerful and centralized adapter management suite. It provides discovery, reporting, and management of local and remote adapters from a single console anywhere in the SAN and across platforms. Both a Graphical User Interface (GUI) and command line interface (CLI) are provided. For instructions on installing and using the Emulex OneCommand CNA Manager application, refer to the *OneCommand CNA Manager Application for OneConnect Adapters User Guide*, which is available at the Documents and Downloads area of http://www.broadcom.com.

2.8 Updating the Adapter Firmware

NOTE This section can be ignored if the adapter already has the required firmware version and RoCE profile.

2.8.1 Determining the Firmware Version

The adapter firmware should be updated to version 11.2. You can update the firmware manually or by using the Emulex OneCommand CNA Manager application.

1. Determine if the firmware must be updated. List the ocrdma interfaces by typing:

   ```
   #ibdev2netdev
   ```

   NOTE The `ibdev2netdev` command is supported only in an RoCE set up. In a NIC, an FCoE, or an iSCSI set up, you must use the `ifconfig` command to find the Emulex NIC interfaces.

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

   If the reported firmware version does not match the version listed previously, the firmware for that adapter must be updated.

   The latest firmware can be downloaded from the Documents and Downloads area of http://www.broadcom.com. The name of the firmware file has a format of:

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

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

2.8.2 Updating the Firmware Manually

NOTE These steps must be performed only once for each adapter.

To update the firmware:

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

   ```
   #ibdev2netdev
   ```

3. Update the firmware for ethX:

   ```
   cp oc14-x.x.x.x.ufi /lib/firmware
   ```

   to

   ```
   /lib/firmware
   # ethtool –f ethX  oc14***.ufi
   ```

   NOTE Multiple ocrdma interfaces can be present on an adapter, so by using the mapping of ocrdma interfaces to Ethernet interfaces (through `ibdev2netdev`) and `ethtool -i ethX`, only one Ethernet interface per adapter can be selected.

Chapter 3: Configuration

3.1 FCoE Driver Configuration

The following section describes how to configure parameters for the FCoE driver.

3.1.1 FCoE Driver Parameters

The FCoE driver parameters determine some aspects of the driver's behavior. There are two main types, static and dynamic. Changes to the static parameters require a driver reload for the change to take effect. Changes to most dynamic parameters take effect immediately; some do not take effect until there is a link-down/link-up sequence.

For more information on driver parameters, see Section 3.1.1.1, Static FCoE Driver Parameters, and Section 3.1.1.2, Dynamic FCoE Driver Parameters.

3.1.1.1 Static FCoE Driver Parameters

Changes to static parameters require a driver reload for the change to take effect. The following table lists the static FCoE driver parameters.

Table 1 Static FCoE Driver Parameters

Parameter	Description	sysfs Visible
`lpfc_ack0`	When enabled, ACK0 is used for Class 2. The enabled value is 1. The disabled value is 0 (default).	Yes
`lpfc_discovery_threads`	Specifies the maximum number of ELS commands that can be outstanding for a discovery. **NOTE** The `lpfc_discovery_threads` parameter defaults to a value of 64 for private loop topologies regardless of the configured value. If multiple ports are configured on the host, the value of 64 is used only for those ports that are connected in a private loop topology. The configured value is used for all other ports. The minimum value is 1. The maximum value is 64. The default value is 32.	No
`lpfc_enable_bg`	Enables the BlockGuard (T10-Data Integrity Field [DIF]) feature. The minimum value is 0 (default). The maximum value is 1.	Yes
`lpfc_enable_da_id`	When enabled, the FCoE driver issues a `DA_ID CT` command to the fabric when VPorts log out of the fabric. The enabled value is 1. The disabled value is 0 (default).	No
`lpfc_enable_hba_heartbeat`	When enabled, the heartbeat logic in the FCoE driver detects whether the adapter is functional. If the heartbeat logic detects the adapter is not functional, the driver will shut down the adapter. The disabled value is 0 (default). The enabled value is 1.	Yes
`lpfc_enable_hba_reset`	When enabled, the FCoE driver passes resets to the adapter. This parameter is typically used for debugging purposes. The enabled value is 1 (default). The disabled value is 0.	Yes
`lpfc_enable_npiv`	When enabled, the FCoE driver uses N_Port ID Virtualization (NPIV) to create VPorts (if supported by the fabric). The enabled value is 1 (default). The disabled value is 0.	Yes
`lpfc_fcp_class`	Specifies either FC Class 2 or 3 for Fibre Channel Protocol (FCP) data transmission. For Class 2, the value is 2. For Class 3, the value is 3 (default).	Yes

Table 1 Static FCoE Driver Parameters (Continued)

Parameter	Description	sysfs Visible
lpfc_fdmi_on	Specifies if FDMI support is enabled or disabled. ■ 0 = Disabled (default) ■ 5 = Enable FDMI without a 60-second delay, use all adapter and port attributes. ■ 7 = Enable FDMI with a 60-second delay, use all adapter and port attributes. Use this value to introduce a delay if FDMI is not operational. The delays begins after a linkup.	Yes
lpfc_lun_queue_depth	Specifies the default maximum number of commands sent to a single logical unit (disk drive). The minimum value is 1. The maximum value is 128. The default value is 30.	Yes
lpfc_max_luns	Specifies the highest available logical unit number (LUN) ID that is valid, per target. For example, a value of 19 means that LUN IDs from 0 to 19 are valid for the target. The Small Computer System Interface (SCSI) layer scans each target until it reaches this specified LUN ID. The minimum value is 0. The maximum value is 65535. The default value is 255.	Yes
lpfc_max_scsicmpl_time	Uses command completion time to control queue depth. The units are in milliseconds. The minimum value is 0 (default). The maximum value is 6000.	Yes
lpfc_multi_ring_rctl	When lpfc_multi_ring_support is enabled, identifies the routing control (R_CTL) for the additional ring configuration. The minimum value is 1. The maximum value is 255. The default value is 4.	Yes
lpfc_multi_ring_support	Determines the number of primary Service Level Interface (SLI) rings over which to spread input/output control block (IOCB) entries. The minimum value is 1 (default). The maximum value is 2.	Yes
lpfc_multi_ring_type	When lpfc_multi_ring_support is enabled, identifies the TYPE of the additional ring configuration. The minimum value is 1. The maximum value is 255. The default value is 5 (logical link control [LLC]/ Subnetwork Access Protocol [SNAP]).	Yes
lpfc_restrict_login	When enabled, restricts VPorts login to remote initiators. The enabled value is 1 (default). The disabled value is 0.	No
lpfc_scan_down	When enabled, selects the *scan down* method (scanning the AL_PA from high to low) to assign a SCSI ID. The enabled value is 1 (default). The disabled value is 0.	Yes
lpfc_sg_seg_cnt	Controls the scatter/gather maximum segment count passed to the FCoE driver. This variable is applicable per SCSI command. The minimum value is 64 (default), and the maximum value is 510.	Yes (sg_tablesize)
lpfc_use_msi	When enabled, determines whether the driver uses message signaled interrupts (MSI) or message signaled interrupts - extended (MSI-X). ■ 0 = MSI disabled; INTx mode is used ■ 1 = MSI; allows a maximum of 32 interrupts. ■ 2 = MSI-X; allows a maximum of 2048 interrupts (default). **NOTE** The default is 2. This value reverts to 1 if the system does not support MSI-X. This value reverts to 0 if the system does not support MSI.	Yes

3.1.1.2 Dynamic FCoE Driver Parameters

Changes to the dynamic parameters take effect immediately. All lpfc dynamic parameters are read/write using sysfs. The following table lists the dynamic FCoE driver parameters.

Table 2 Dynamic FCoE Driver Parameters

Parameter	Description
`lpfc_devloss_tmo`	Specifies the number of seconds to hold an input/output (I/O) error when a device disappears. The minimum value is 0. The maximum value is 255. The default value is 30.
`lpfc_fcp_imax`	Specifies the maximum number of fast-path FCP interrupts per second. The minimum value is 636. The maximum value is 651042. The default value is 5000.
`lpfc_log_verbose`	Specifies the log verbosity level of the messages posted by the driver. Extra activity logging (bit mask). The minimum value is 0x0 (default). The maximum value is 0xFFFFFFFF.
`lpfc_nodev_tmo (deprecated)`	**NOTE** This is a deprecated parameter and the `lpfc_devloss_tmo` parameter should be used instead. This parameter will not work if you altered `lpfc_devloss_tmo`. Specifies the number of seconds to hold an I/O error when a device disappears. The minimum value is 1. The maximum value is 255. The default value is 30.
`lpfc_pci_max_read`	Specifies the maximum direct memory access (DMA) read byte count. The possible values are 512, 1024, 2048 (default), and 4096.
`lpfc_poll_tmo`	Specifies the number of milliseconds that the driver waits between polling FCP ring interrupts. The minimum value is 1. The maximum value is 255. The default value is 10.
`lpfc_throttle_log_cnt`	Specifies the number of messages logged within `throttle_log_time`. The minimum value is 1. The maximum value is 1000. The default value is 10.
`lpfc_throttle_log_time`	Specifies the time limit for `throttle_log_cnt`. The minimum value is 1. The maximum value is 60. The default value is 1.
`lpfc_use_adisc`	When enabled, an ADISC is sent instead of a PLOGI for device discovery or registered state change notification (RSCN). The enabled value is 1. The disabled value is 0.

3.1.2 Configuring FCoE Driver Parameters

You can configure the FCoE driver parameters using any of the following methods:

- The `modprobe` Linux program for temporary configuration
- The `lpfc.conf` file for persistent configuration
- The sysfs interface (to view and modify parameters after loading the FCoE driver)

  NOTE Not all parameters visible in the sysfs directory can be modified; some are read-only.
- The Emulex OneCommand CNA Manager application (refer to the *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide* for more information)

  NOTE The FCoE driver parameter changes made using the Emulex OneCommand CNA Manager application remain after the FCoE driver is uninstalled. To return to the default settings, you must reset them using the OneCommand CNA Manager application.

3.1.2.1 Temporary Configuration with modprobe

When you manually load the FCoE driver as a module using the `modprobe` command, and you change one or more driver parameter values in the command line, the configuration is temporary. These changes are considered temporary because they are valid only for the current session or until the FCoE driver is unloaded.

The `modprobe` program uses the `lpfc.conf` file, but parameters passed to it using the command line override the parameters in the `lpfc.conf` file. Values can be expressed in hexadecimal or decimal notation.

If you want to temporarily set `lun_queue_depth` to 20 (the default is 30) for all HBAs in your system, load the FCoE driver with the following command:

```
modprobe lpfc lpfc_lun_queue_depth=20
```

3.1.2.2 Persistent Configuration with the lpfc.conf File

To make the FCoE driver parameters persist across module loads and reboots, perform these steps:

1. In the `/etc/modprobe.d` directory, create a file with the driver name `lpfc.conf`.
2. In `/etc/modprobe.d/lpfc.conf`, use the `options` command to add the appropriate FCoE driver parameters and their desired values. For example, adding the following command to the `lpfc.conf` file sets the verbose flag:

   ```
   options lpfc lpfc_log_verbose=0x3ffff
   ```

If driver parameters are added to the `lpfc.conf` file, the FCoE driver must be reloaded for the parameters to take effect. Also, a new ramdisk image is required if you want the changes to take effect in the next boot. For information on creating a new ramdisk, See Section 3.1.3, Creating a New Ramdisk Image.

If the same parameter is specified on the `modprobe` command line and in the `lpfc.conf` file, then the value specified in the modprobe command line takes precedence.

3.1.2.3 Configure Parameters with a Read/Write to sysfs

Sysfs is a virtual file system that exposes the structure of the system. It also includes interfaces to driver parameters through which the FCoE driver parameters can be viewed and modified. Because these interfaces are available only after driver load, only dynamic FCoE driver parameters can be changed. However, both static and dynamic FCoE driver parameters can be read through sysfs.

NOTES

- Sysfs changes exist only during driver load and are lost when the FCoE driver is unloaded or the system is rebooted.
- Driver parameters that are set through module parameters are global; setting them through sysfs is on a SCSI host (adapter port) basis.

  For example:

  `echo 0x7f >> /sys/class/scsi_host/host7/lpfc_log_verbose` only affects host 7.

  `modprobe lpfc lpfc_log_verbose=0x7f` applies to all SCSI host (ports) managed by the brcmlpfc driver.

Viewing Parameters with sysfs

The sysfs file system is mounted and available as `/sys`. You must first identify the scsi_host that represents the adapter for which you want to modify the FCoE driver parameters. All scsi_hosts bound to the FCoE driver can be viewed with the following command:

```
ls -d /sys/bus/pci/drivers/lpfc/*/host*
```

Assuming you are interested in adapter scsi_host 7, you can list the FCoE driver parameters for this particular adapter with the following command:

```
ls -l /sys/class/scsi_host/host7/lpfc*
```

An example output follows:

```
-r--r--r--  1 root root 4096 Feb 28 17:03
/sys/class/scsi_host/host7/lpfc_ack0
-r--r--r--  1 root root 4096 Feb 28 17:03
/sys/class/scsi_host/host7/lpfc_fcp_class
-rw-r--r--  1 root root 4096 Feb 28 17:03 /sys/class/scsi_host/host7/
lpfc_fdmi_on
-r--r--r--  1 root root 4096 Feb 28 17:03
/sys/class/scsi_host/host7/lpfc_link_speed
-rw-r--r--  1 root root 4096 Feb 28 15:34
/sys/class/scsi_host/host7/lpfc_log_verbose
-r--r--r--  1 root root 4096 Feb 28 17:03
/sys/class/scsi_host/host7/lpfc_lun_queue_depth
-rw-r--r--  1 root root 4096 Feb 28 17:03
/sys/class/scsi_host/host7/lpfc_max_luns
-rw-r--r--  1 root root 4096 Feb 28 17:03
/sys/class/scsi_host/host7/lpfc_nodev_tmo
-rw-r--r--  1 root root 4096 Feb 28 17:03
/sys/class/scsi_host/host7/lpfc_scan_down
-r--r--r--  1 root root 4096 Feb 28 17:03
/sys/class/scsi_host/host7/lpfc_topology
-rw-r--r--  1 root root 4096 Feb 28 17:03
/sys/class/scsi_host/host7/lpfc_use_adisc
```

Temporary Configuration Parameters with sysfs

In the previous example, notice that the FCoE driver parameters are available as files. Reading a file displays the current value of a driver parameter. If the permissions allow it, you can write a value to the file, and it will take effect immediately.

Reading the `lpfc_log_verbose` file may show that its value is `0`:

```
cat /sys/class/scsi_host/host7/lpfc_log_verbose
0
```

To modify the `lpfc_log_verbose` value to 0xffffffff:

```
echo 0xffffffff > /sys/class/scsi_host/host7/lpfc_log_verbose
```

Reading the lpfc_log_verbose file now shows a value of 0xffffffff:

```
cat /sys/class/scsi_host/host7/lpfc_log_verbose
0xffffffff
```

NOTE Setting the `lpfc_log_verbose` value to 0xffffffff might cause degraded system performance.

3.1.3 Creating a New Ramdisk Image

The `lpfc-install` script creates a ramdisk image containing the FCoE driver for the currently running kernel.

NOTE You must create a new ramdisk image whenever the lpfc options in `/etc/modprobe.conf` are changed and you want the change to take effect on the next reboot.

To create a new initial ramdisk image for inbox FCoE drivers and installed binary RPM FCoE driver kits, type

```
dracut -f /boot/initramfs-<kernel-version>.img <kernel-version>
```

3.1.4 Dynamically Recognizing LUNs and Targets (Using scan)

The FCoE driver enables you to dynamically recognize LUNs and targets without unloading or reloading the FCoE module and without resetting the adapter.

To rescan an adapter's targets with sysfs, given the adapter's host number (in this example, 3), type:

```
echo "- - -" > /sys/class/scsi_host/host3/scan
```

To limit the rescan to a particular target, given the adapter's host number (in this example, 3) and the target number (in this example, 2), type:

```
echo "- 2 -" > /sys/class/scsi_host/host3/scan
```

You can also use the `lun_scan` script in the `/usr/sbin/brcmlpfc` directory.

3.1.5 Persistent Naming

The generic device manager for the Linux kernel is *udev*, which primarily manages device nodes in the `/dev` directory.

3.1.5.1 Using udev to Discover Logical to Physical Mappings for sd Devices

In Linux, the driver for SCSI disk drives is *sd*. A disk device name has an sd prefix. Persistent names for sd devices are provided in the `/dev/disk/by-id` directory. To find the persistent udev name for the disk, which is currently *sdc*, type:

```
cd /dev/disk/by-id
ls -l | grep sdc
```

The sample output is:

```
lrwxrwxrwx 1 root root 9 2006-08-01 19:08 scsi-32000000c5005d6e6 ->
../../sdc
```

In the previous example, the disk has no partitions. If the disk had two partitions, the output would look like the following:

```
lrwxrwxrwx 1 root root 9 2006-08-01 19:08 scsi-32000000c5005d6e6 ->
../../sdc
lrwxrwxrwx 1 root root 10 2006-08-01 19:08 scsi-32000000c5005d6e6-part1 ->
../../sdc1
lrwxrwxrwx 1 root root 10 2006-08-01 19:08 scsi-32000000c5005d6e6-part2 ->
../../sdc2
```

3.1.5.2 Configuring the System to Boot Using Persistent Names

For SLES 11 SPx and SLES 12 SPx:

NOTE SLES 11 SPx and SLES 12 SPx are configured by default with udev to provide persistent names for hard disks, including FCoE-attached disks.

To use a persistent name for a boot device:

1. In `/boot/grub/menu.lst`, find the kernel line for the default boot. For example:
   ```
   kernel /boot/vmlinuz root=/dev/sda2 vga=0x314
   ```
2. Find the persistent name for the root partition (following `root=` on the kernel line) by using the instructions in Section 3.1.5.1, Using udev to Discover Logical to Physical Mappings for sd Devices.
3. In the same file, `/boot/grub/menu.lst`, replace the text after `root=` with the partition's persistent name. For example:
   ```
   kernel /boot/vmlinuz root=/dev/disk/by-id/scsi-32000000c5005d6e6-part2
   vga=0x314
   ```
4. Change any mounts listed in `/etc/fstab` that refer to this root partition by either its `/dev/sd` name or a file system label to use the persistent name as well.

For RHEL 6.x and RHEL 7.x:

To use a persistent name for a boot device:

1. In `/boot/grub/grub.conf`, find the kernel line for the default boot. For example:
   ```
   kernel /boot/vmlinuz -<kernel version> ro root=/dev/sda2
   ```
2. Find the persistent name for the root partition (following `root=` on the kernel line) by using the instructions in Section 3.1.5.1, Using udev to Discover Logical to Physical Mappings for sd Devices.
3. In the same file, `/boot/grub/menu.lst`, replace the text after `root=` with the partition's persistent name. For example:
   ```
   kernel /boot/vmlinuz -<kernel version> ro
   root=/dev/disk/by-id/scsi-32000000c5005d6e6-part2
   ```
4. Change any mounts listed in `/etc/fstab` that refer to this root partition by either its `/dev/sd` name or a file system label to use the persistent name as well.

3.1.5.3 Using udev with st Devices

In Linux, the driver for SCSI tape drives is *st*. A tape device name has an *st* prefix. The udev rules for tape devices are the same as for disk devices. There must be a unique ID that persists across initiator reboots and persists regardless of discovery order.

You must consider whether the tape device is an FC tape device or an FC-SCSI tape device (in which there are multiple SCSI tape devices that reside behind an FC controller). If it is an FC tape device, then the World Wide Port Name (WWPN) is unique and can be used to create the persistent name. In this case, the `scsi_id` command should return this as the unique identifier with a single digit prefix. If the FC controller has multiple SCSI tape devices behind it, the WWPN is not unique, and the persistent name must use multiple information elements to build the unique ID. FC Tape Device Examples and FC-SCSI Tape Device Example are examples of each scenario.

FC Tape Device Examples

The following examples use the `scsi_id` command to retrieve and generate a unique SCSI identifier:

```
scsi_id [options]
```

For these examples, the following [options] are used:

- -g Treats the device as white listed. It is required on the command line or in the scsi_id.config file for the scsi_id command to generate any output. In the examples, the -g option is needed on the command line because the vendor and model for this tape device were not in the /etc/scsi_id.config file.
- -s Generates an id for the sysfs-device. Note that -s is an invalid option for scsi_id version 147.

NOTE Because the [options] can vary depending on the version of the scsi_id command, refer to the *scsi_id* man page on your system for the correct and complete list of the [options].

The following example is an FC tape device using the SCSI generic driver (sg) rather than the SCSI tape driver. The value returned has a leading prefix of 3, which is the Network Address Authority (NAA) type. The remaining digits represent the FC controller's WWPN.

```
scsi_id -g -s /sys/class/scsi_generic/sg0
350060b000029b592
```

The following example is an FC tape device using the SCSI tape driver. The value returned is the same as the previous example.

```
scsi_id -g -s /sys/class/scsi_tape/nst0
350060b000029b592
```

The following example uses a different FC tape vendor. Notice that the value returned is similar to the previous examples, with respect to the leading digit and the WWPN.

```
/sbin/scsi_id -g -s sys/class/scsi_tape/nst0
35005076300015101
```

FC-SCSI Tape Device Example

The following is an example of a FC controller with multiple SCSI tape devices behind it (FC-SCSI tape device). When the Emulex driver is loaded, the SCSI mid-level discovers the SCSI tape devices as follows:

```
scsi scan: INQUIRY to host 14 channel 0 id 0 lun 0
scsi: unknown device type 12
Vendor: ADIC      Model: SNC 4000          Rev: 42d4
Type:   RAID                               ANSI SCSI revision: 03
Attached scsi generic sg5 at scsi14, channel 0, id 0, lun 0,  type 12
scsi scan: INQUIRY to host 14 channel 0 id 0 lun 1
Vendor: ADIC      Model: Scalar 24         Rev: 227A
Type:   Medium Changer                     ANSI SCSI revision: 02
Attached scsi generic sg6 at scsi14, channel 0, id 0, lun 1,type 8
scsi scan: INQUIRY to host 14 channel 0 id 0 lun 2
Vendor: IBM       Model: ULTRIUM-TD2       Rev: 38D0
Type:   Sequential-Access                  ANSI SCSI revision: 03
Attached scsi tape st0 at scsi14, channel 0, id 0, lun 2
st0: try direct i/o: yes (alignment 512 B), max page reachable by HBA
4503599627370495
Attached scsi generic sg7 at scsi14, channel 0, id 0, lun 2,  type 1
scsi scan: INQUIRY to host 14 channel 0 id 0 lun 3
Vendor: IBM       Model: ULTRIUM-TD2       Rev: 38D0
Type:   Sequential-Access                  ANSI SCSI revision: 03
Attached scsi tape st1 at scsi14, channel 0, id 0, lun 3
st1: try direct i/o: yes (alignment 512 B), max page reachable by HBA
4503599627370495
Attached scsi generic sg8 at scsi14, channel 0, id 0, lun 3,  type 1
```

This log output shows a controller at LUN 0, the medium changer at LUN 1, and two SCSI tape devices at LUNs 2 and 3.

The following example is the result of a `scsi_id` call:

```
scsi_id -g -s /sys/class/scsi_tape/nst0
1IBM     ULTRIUM-TD2     1110133831
scsi_id -g -s /sys/class/scsi_tape/nst1
1IBM     ULTRIUM-TD2     1110133994
```

Notice that the unique ID contains three values with space delimiters. A udev rule must have a unique ID for the device, so that all three parts of this returned string are required. To do this, use the following command:

```
scsi_id -u -g -s /sys/class/scsi_tape/nst0
1IBM_____ULTRIUM-TD2_____1110133831
scsi_id -u -g -s /sys/class/scsi_tape/nst1
1IBM_____ULTRIUM-TD2_____1110133994
```

Creating the udev Persistent Name for SCSI Tape Device

After you know the SCSI ID call needed to extract a unique ID, use the same process to create a udev persistent name for a SCSI tape device as on a SCSI disk device.

The rule for the FC tape device follows:

```
BUS="scsi", SYSFS{vendor}="HP", SYSFS{model}="ULTRIUM 3-SCSI",
PROGRAM="/sbin/scsi_id -p 0x83 -u -g -s
/sys/class/scsi_tape/nst%n",RESULT="350060b000029b592",
SYMLINK="fc_lun_st%n"
```

The rule for the FC-SCSI tape device follows:

```
BUS="scsi", SYSFS{vendor}="IBM", SYSFS{model}="ULTRIUM-TD2",
PROGRAM="/sbin/scsi_id -p 0x83 -u -g -s
/sys/class/scsi_tape/nst%n",RESULT="1IBM_____ULTRIUM-TD2_____1110133831",
SYMLINK="fc_lun_st%n"
BUS="scsi", RESULT="1IBM_____ULTRIUM-TD2_____1110133994",
SYMLINK="fc_lun_st%n"
```

Create a new file named `/etc/udev/rules.d/45-local.rules`, and put the appropriate rule in it. Then run `udevtrigger` to reload the udev rules, and the output of the rules will follow:

```
udevtrigger
ls -al /dev/fc*
lrwxrwxrwx  1 root root 3 Apr  7 15:03 fc_lun_st0 -> st0
lrwxrwxrwx  1 root root 3 Apr  7 15:03 fc_lun_st1 -> st1
```

3.1.5.4 Persistent Naming References

Refer to the following references for more information on persistent naming:

- http://www.reactivated.net/writing_udev_rules.html by Daniel Drake (dsd)
- https://groups.google.com/forum/#!topic/fa.linux.kernel/20FrKJkUKAA

3.2 Ethernet Driver Configuration

The following section describes how to configure parameters for the Ethernet driver.

3.2.1 Ethernet Driver Configuration Parameters

The Ethernet driver configuration parameters are described in the following table.

Table 3 Ethernet Driver Configuration Parameters

Parameter	Description
num_vfs	In systems supporting single-root I/O virtualization (SR-IOV), when I/O virtualization (IOV) is enabled, this parameter indicates the number of VFs to be enabled per physical function (PF). For configuring SR-IOV, see Section 3.2.7, SR-IOV Configuration. The default value is 0 (SR-IOV is not enabled). The possible values are 0 to 63. **NOTE** This parameter is obsolete. Use the sysfs method to enable or disable VFs.
rss_on_mc	Enables receive-side scaling (RSS) on multichannel functions that have the capability. The default value is 0 (disabled). The enabled value is 1. **NOTE** Currently, this parameter is ignored by the driver. The driver enables RSS on multichannel functions that have the capability by default.
rx_frag_size	The size of fragments used to DMA received data. The possible values are 2048 (default), 4096 and 8192.

You can configure Linux to automatically load the driver with any of these options after each reboot. To do so, add a line to `/etc/modprobe.conf` with the required options. For example, to load the driver with the fragment size of 4096, add the following line:

```
options be2net rx_frag_size=4096
```

3.2.2 VLAN Support on UMC

In the Universal multichannel (UMC) mode, a logical port VLAN ID (LPVID) for each channel configured in the basic input/output system (BIOS) is transparently added in the transmit path and removed in the receive path by the adapter. When a virtual local area network (VLAN) is configured in the host using `vconfig` or `ip link` on any of the functions, the host VLAN ID overrides the corresponding LPVID channel configured in the BIOS. Up to 15 VLAN IDs can be configured in the host for each PF using `vconfig` or `ip link`. The VLAN IDs configured in the host should be different from the channel VLAN IDs configured in the BIOS. For additional information on UMC, refer to the *Emulex Universal Multichannel Reference Manual*.

NOTE

- Link Aggregation Control Protocol (LACP) is not supported when UMC is enabled.
- SR-IOV is not supported when UMC is enabled.

3.2.3 Support for Ethtool set-channels and get-channels Commands

The Ethtool `get-channels` command – `ethtool -l <ethx>` displays the following information:

- The current number of receive (Rx)/ transmit (Tx)/interrupt queue pairs (a *combined channel*) created by the NIC function
- The maximum number of channels supported by the NIC function

The NIC driver supports a combination of Rx/Tx-only channels along with combined channels.

Use the following command to create the combined channels:

```
ethtool -L <ethx> combined N
```

Use the following commands to create N Tx-queues and M Rx-queues:

When N < M use:

```
ethtool -L ethX combined N rx M-N
```

When M < N use:

```
ethtool -L ethX combined M tx N-M
```

NOTE

- You can increase (up to the maximum supported limit) or decrease the number of channels used by the NIC function.
- Setting both Rx-only and Tx-only channels is not supported; it is mandatory to create at-least one combined channel.
- The maximum number of channels supported by an interface is also limited to the number of cpu-cores in the server.

NOTE In earlier Linux distributions, use the sysfs interface as follows:

- The maximum number of channels supported by the NIC function:
  `cat /sys/class/net/<ethx>/max_qs`
- The current number of channels configured:
  `cat /sys/class/net/<ethx>/num_qs`
- To configure the 'N' number of channels requested for the NIC function:
  `echo N > /sys/class/net<ethx>/num_qs`

3.2.4 Transmit/Receive Queue Counts

The number of queue counts created by the NIC driver is provided in the following table.

Table 4 serves as an example in some combinations of profile configurations and adapters. This table is not intended to be a comprehensive list of all profile configurations and adapters.

Table 4 Tx and Rx Queue Counts

Profile/Configuration			**Combined IRQs (for example, RSS and Tx queues) (Initial/Max) Initial: number of queues available on load Max: number of queues that can be set by the `set-channels` command**		
	OCe14002 (ASIC REV 0x11)	OCe14401	OCe14002 (ASIC REV 0x10)	OCe14004	OCe11102
NIC only	8/31	8/31	8/31	8/31	8/16
NIC only SR-IOV PF (sriov_numvfs>0)	[1-8]/[1-31]	[1-8]/[1-31]	[1-8]/[1-31]	[1-8]/[1-31]	1/1

Table 4 Tx and Rx Queue Counts (Continued)

Profile/Configuration			Combined IRQs (for example, RSS and Tx queues) (Initial/Max) Initial: number of queues available on load Max: number of queues that can be set by the `set-channels` command		
NIC only SR-IOV virtual function (VF)	[1-3]/[1-3]	[1-3]/[1-3]	[1-3]/[1-3]	[1-3]/[1-3]	1/1
NIC only SR-IOV PF (sriov_numvfs=0)	8/31	8/31	8/31	8/31	8/16
NIC (only with ARI) (UMC andNIC partitioning [NPar])	8/8	8/8	4/4	8/8	8/16 (first interface, 1/1 otherwise)

3.2.5 Support for Ethtool set-rxfh/get-rxfh Commands

The Ethtool `get-rxfh` and `set-rxfh` commands are available in kernel versions of RHEL 7.1, 3.16.0 and later.

The Ethool `get-rxfh` command – `ethtool – x <`*`ethx`*`>` shows the following information:

- The current indirection table entries
- The programmed RSS hash key

The Ethtool `set-rxfh` command – `ethtool –X <`*`ethx`*`> [hkey aa:bb:cc...] [equal N | weight W0 W1...]` sets the following parameters:

- Indirection table entries
- RSS hash key

NOTE In earlier Linux distributions, use the sysfs interface as follows:

- To show the currently configured RSS hash key:
  `cat /sys/class/net/<`*`ethx`*`>/rss_hashkey`
- To configure the new RSS hash key
  `echo -e "aa:bb:cc...." > /sys/class/net/<`*`ethx`*`>/rss_hashkey`

3.2.6 Support for Showing Onboard Die Temperature

Onboard die temperature is exported through the sysfs interface. On kernels that support the hardware monitoring framework, the temperature is exported using the hwmon sysfs interface. Temperatures exported using the hwmon interface can also be read using the lm-sensors utility with the `sensors` command.

To show the current board temperature using the sysfs node:

```
cat /sysclass/hwmon/<hwmonx>/device/temp1_input
```

Otherwise, the current board temperature is exported with the regular sysfs interface:

```
cat /sys/class/net/<ethx>/device/temp1_input
```

3.2.7 SR-IOV Configuration

3.2.7.1 Introduction

This section contains requirements and instructions to use SR-IOV with the following host operating systems:

- Red Hat Enterprise Linux 6.6 – 64-bit kernel-based virtual machine (KVM)
- Red Hat Enterprise Linux 6.7– 64-bit KVM
- Red Hat Enterprise Linux 6.8 – 64-bit KVM
- Red Hat Enterprise Linux 7.1 – 64-bit KVM
- Red Hat Enterprise Linux 7.2 – 64-bit KVM
- Red Hat Enterprise Linux 7.3 – 64-bit KVM
- –
- SuSE Linux Enterprise Server 11 SP4 – 64-bit Xen and KVM
- SuSE Linux Enterprise Server 12 SPI – 64-bit Xen and KVM
- SuSE Linux Enterprise Server 12 SP2 – 64-bit Xen and KVM
- SLES and RHEL supported with Xen

These environments support capabilities of the Emulex OneConnect adapter to enable multiple Peripheral Component Interconnect Express (PCIe) virtual functions (VFs) for a PCIe physical function (PF). Each of these VFs can be assigned to virtual machines (VMs). A VF enables the guest operating system direct access to the Emulex OneConnect adapter, such that guest performance is not limited by the overheads of the hypervisor.

With SR-IOV, VMs directly drive I/Os on the NIC. Therefore, SR-IOV has the following advantages over traditional virtualized I/O:

- Improved device performance for virtual guests
- Increased scalability
- Reduced CPU utilization
- Reduced latency

Known issues include:

- The Kdump (kernel dump) feature is not supported when SR-IOV is enabled.

3.2.7.2 Setting Up SR-IOV

NOTE SR-IOV is not supported in multichannel setups and Dell NPAR.

Prerequisites

To set up SR-IOV on your system, you need the following:

- A server or blade with an SR-IOV-capable motherboard BIOS

  NOTE Configuration mechanisms for parameters, such as Media Access Control (MAC) address, VLAN, and quality of service (QoS) for VF, are supported in RHEL 6 and later, and SLES 11 and later distributions.
- The Emulex OneConnect OCx11102-xT or later adapter versions.
- If a KVM hypervisor is installed, it must contain the qemu-kvm packages.

Depending on your system, perform one or more of the following tasks to set up your BIOS. For more information, refer to the manufacturer's instructions for your system.

- Enable SR-IOV in the system BIOS, which supports SR-IOV functionality.
- Enable Intel Virtualization Technology support for Direct I/O VT-d.

You can use the PXESelect utility, the Emulex OneCommand CNA Manager application, or the Unified Extensible Firmware Interface (UEFI) Human Interface Infrastructure (HII) utility to set up SR-IOV.

To enable and configure SR-IOV using either the UEFIBoot utility or the PXESelect utility, refer to the *Boot for NIC, iSCSI, and FCoE Protocols User Guide*.

1. Install the required Linux operating system that serves as a hypervisor.
2. Update the `/boot/grub/menu.lst` file to include the following command line load parameter for the Linux kernel:

   ```
   intel_iommu=on
   ```
3. Install the appropriate version of the Emulex Ethernet driver that supports SR-IOV for the operating system version that you are using.
4. Reboot the server for the new changes to become operational.
5. Use the `lspci -vvv` output command to check if SR-IOV is properly enabled.

   The `lspci-vvv` command returns an SR-IOV capability section for each Ethernet PF. The Initial VFs and Total VFs should be nonzero. Make a note of the lspci command output in the hypervisor. This output is needed in step 10. Specifically, make note of the pci-id of the VFs that have been created.

   The following is an example of the SR-IOV capabilities section output of the PF with SR-IOV enabled:

   ```
   Capabilities: [180] Single Root I/O Virtualization (SR-IOV)
   IOVCap: Migration-, Interrupt Message Number: 000
   IOVCtl: Enable- Migration- Interrupt- MSE- ARIHierarchy-
   IOVSta: Migration-
   Initial VFs: 32, Total VFs: 32, Number of VFs: 0, Function Dependency
   Link: 00
   VF offset: 0, stride: 1, Device ID: 0710
   Supported Page Size: 00000557, System Page Size: 00000001
   Region 0: Memory at 0000000000000000 (64-bit, non-prefetchable)
   VF Migration: offset: 00000000, BIR: 0
   ```

   The following is an example of the SR-IOV capabilities section output of the PF with SR-IOV disabled:

   ```
   Capabilities: [180] Single Root I/O Virtualization (SR-IOV)
   IOVCap: Migration-, Interrupt Message Number: 000
   IOVCtl: Enable- Migration- Interrupt- MSE- ARIHierarchy-
   IOVSta: Migration-
   Initial VFs: 0, Total VFs: 0, Number of VFs: 0, Function Dependency
   Link: 00
   VF offset: 0, stride: 1, Device ID: 0710
   Supported Page Size: 00000557, System Page Size: 00000001
   Region 0: Memory at 0000000000000000 (64-bit, non-prefetchable)
   VF Migration: offset: 00000000, BIR: 0
   ```
6. Create a virtual machine (VM) using the Virtual Machine Manager utility, and install any supported operating system on the VM.
7. Shut down the VM.
8. Echo the number of VFs/PF to the `sriov_numvfs` file, where *X* is the number of VFs per PF:

   ```
   echo X > /sys/bus/pci/devices/..../sriov_numvfs
   ```

   The possible values are 0 to 64 per physical port. The default value is 0 (SR-IOV is not enabled).

   With the latest kernels, the sysfs path to enable VFs is:

   ```
   /sys/bus/pci/devices/<device-id>/sriov_numvfs
   ```

   For earlier kernels, the sysfs path to enable VFs is

   ```
   /sys/bus/pci/devices/<device-id>/net/<ifaceName>/sriov_numvfs
   ```

The total number of VFs can be distributed among available ports as required, but each port has a maximum of 64 VFs. Table 5 lists the total number of VF counts that are supported for various adapter configurations.

The maximum number of VFs supported per PF can be read from:

```
cat/sys/bus/pci/devices/..../sriov_totalvfs
```

NOTE

- On 4-port OCx11104 adapters, VFs can be configured only for ports 0 and 1, not for ports 2 or 3.
- VFs are supported only for network functions; they are not supported for storage functions.

Table 5 serves as an example in some adapter configurations. This table is not intended to be a comprehensive list of all adapter configurations.

Table 5 Total VF Counts Supported for Some Adapter Configurations

Adapter Configuration	Number of Ports	PF Count	VF Count per ASIC	Comments
OCe11102 2-port 10Gb NIC, stand-alone adapter	2	2	60	Per-port VF count is 30. P0/P1: 30/30.
OCe11102 2-port 10Gb NIC, stand-alone, Network Communication Services Interface (NCSI)-enabled initial program load (IPL)	2	2	59	Per-port VF count can be P0/P1: 30/29 or P0/P1: 29/30.
OCm11102 2-port 10Gb, HP/Lenovo System X LAN on motherboard (LOM)/Mezzanine	2	2	60	Per-port VF count is 30. P0/P1: 30/30.
OCm11102 2-port 10Gb, HP/Lenovo System X LOM/Mezzanine, NCSI-enabled IPL	2	2	59	Per-port VF count can be P0/P1: 30/29 or P0/P1: 29/30.
OCm11104 2-port 10Gb + 2-port 1Gb (4-port), Lenovo System X LOM/Mezzanine	4	4	55	Per-port VF count can be P0/P1: 28/27 or P0/P1: 27/28. VFs are not supported on 1Gb/s ports.
OCe11102 2-port 10Gb HP (all IPLs)	2	4/6/8	48	Per-port VF count is 24. VFs will be distributed across the PFs in multiples of 8.
OCe14401 1-port 40Gb NIC	1	1	63	63 VFs for NIC port.
OCe14102 (ASIC REV 0X10) 2-port 10Gb NIC	2	2	62	Per-port VF count is 31.
OCe14102 (ASIC REV 0X11) 2-port 10Gb NIC	2	2	126	Per-port VF count is 63.
OCe14104 4-port 10Gb NIC	4	4	124	Per-port VF count is 31.
OCe14401 1-port 40Gb NIC plus storage (iSCSI/FCoE)	1	2 (one NIC, one storage)	63	63 VFs for NIC port.
OCe14102 (ASIC REV 0x10) 2-port 10Gb NIC plus storage (iSCSI/FCoE)	2	4 (two NIC, two storage)	62	Per-NIC-port VF count is 31.
OCe14102 (ASIC REV 0X11) 2-port 10Gb NIC plus storage (iSCSI/FCoE)	2	4 (two NIC, two storage)	126	Per-NIC-port VF count is 63.

9. Detach VFs from the host before adding them to the guest.

   Example 1 (using only one VF per physical port):

   RHEL 6.x KVM: View Emulex PCI devices with the `lspci` command:

   ```
   # lspci | grep Emulex     16:00.0 Ethernet controller: Emulex Corporation OneConnect 10Gb NIC (be3) (rev 03) <-- LOM Port 0 (Function 0)     16:00.1 Ethernet controller: Emulex Corporation OneConnect 10Gb NIC (be3) (rev 03) <-- LOM Port 1 (Function 1)     16:04.0 Ethernet controller: Emulex Corporation OneConnect 10Gb NIC (be3) (rev 03) <-- VF LOM Port 0 (Function 0)     16:08.0 Ethernet controller: Emulex Corporation OneConnect 10Gb NIC (be3) (rev 03) <-- VF LOM Port 1 (Function 1)
   ```

   Detach desired VFs:

   ```
   # virsh nodedev-dettach pci_0000_16_04_0     Device pci_0000_16_04_0 detached <-- VF LOM Port 0 (Function 0)     # virsh nodedev-dettach pci_0000_16_08_0     Device pci_0000_16_08_0 detached <-- VF LOM Port 1 (Function 1)
   ```

   Per-port VF count SLES 11.*x* Xen: View Emulex PCI devices with the `lspci` command:

   ```
   # lspci | grep Emulex     16:00.0 Ethernet controller: Emulex Corporation OneConnect 10Gb NIC (be3) (rev 03) <-- LOM Port 0 (Function 0)     16:00.1 Ethernet controller: Emulex Corporation OneConnect 10Gb NIC (be3) (rev 03) <-- LOM Port 1 (Function 1)     16:04.0 Ethernet controller: Emulex Corporation OneConnect 10Gb NIC (be3) (rev 03) <-- VF LOM Port 0 (Function 0)     16:08.0 Ethernet controller: Emulex Corporation OneConnect 10Gb NIC (be3) (rev 03) <-- VF LOM Port 1 (Function 1)     # lspci -nn | grep Emulex     16:00.0  Ethernet controller: [0200]: Emulex Corporation OneConnect 10Gb NIC (be3) [19a2:0710] (rev 03) <-- LOM Port 0 (Function 0)     16:00.1  Ethernet controller: [0200]: Emulex Corporation OneConnect 10Gb NIC (be3) [19a2:0710] (rev 03) <-- LOM Port 1 (Function 1)     16:04.0  Ethernet controller: [0200]: Emulex Corporation OneConnect 10Gb NIC (be3) [19a2:0710] (rev 03) <-- VF LOM Port 0 (Function 0)     16:08.0  Ethernet controller: [0200]: Emulex Corporation OneConnect 10Gb NIC (be3) [19a2:0710] (rev 03) <-- VF LOM Port 1 (Function 1)
   ```

   List Emulex PCI device by device code:

   ```
   # virsh nodedev-list | grep 19a2     pci_19a2_710  <-- LOM Port 0 (Function 0)     pci_19a2_710_0 <-- LOM Port 1 (Function 1)     pci_19a2_710_1 <-- VF LOM Port 0 (Function 0)     pci_19a2_710_2 <-- VF LOM Port 1 (Function 1)     # virsh nodedev-dettach pci_19a2_710_1     Device pci_19a2_710_1 detached <-- VF LOM Port 0 (Function 0)     # virsh nodedev-dettach pci_19a2_710_2     Device pci_19a2_710_2 detached <-- VF LOM Port 1 (Function 1)
   ```

10. Use the Virtual Machine Manager GUI to attach the VF (step 5.) to the guest operating system by using the add physical device option.

    NOTE To reconfigure a system that is already set up, remove the attached VF from the guest operating system by selecting the VF and using the remove option. Refer to the documentation for the host operating system for information on using the Virtual Machine Manager to attach and remove virtual interfaces.

11. Start the RHEL 6.*x* guest operating system. After the guest operating system is booted, use the `lspci` command to confirm the visibility of the NIC function. The output shows a NIC function, for example:

    ```
    03:00.0 Ethernet controller: Emulex Corp. Emulex OneConnect 10Gb NIC (be3)
    ```

12. The Ethernet driver automatically loads with the out-of-box driver and creates the network interfaces. Use the `ifconfig` command to verify that the interface is created.

13. After configuring the network interfaces with proper Internet Protocol (IP) addresses, you can send network traffic to and receive network traffic from the VM. Refer to the documentation for the host and guest operating systems for information on network configuration.

3.2.7.3 Edge Virtual Bridging

The OCe14000-series chip supports both virtual Ethernet bridging (VEB) and tagless virtual Ethernet port aggregator (VEPA) modes. VEB mode is enabled by default in the chip. Switching to VEB/VEPA mode can be done using the bridge utility or the Emulex OneCommand CNA Manager application. The bridge utility is available in RHEL 7 distributions.

3.2.7.4 Assigning VFs to a VM on the SLES Operating System

To assign VFs to the VM in the SLES Xen kernel, the VF must be unbound from the NIC module and then bound to the pciback module.

NOTE In the following steps, *0000:07:0b.5* is used as an example. To match those instances to the port that you want to use, you need to select the entry that matches the PCI bus, device, or function that corresponds to the port that you want to assign. Use the `ethtool` utility to determine this information, such as `ethtool -i eth0` (where `eth0` is the interface you want to assign).

1. Load the pciback driver `modprobe pciback`.
2. Navigate to the `/sys/bus/pci/drivers/pciback` directory, and ensure that the following is displayed:

```
Panama-Sles11sp2:/sys/bus/pci/drivers/pciback # ls -lrt
total 0
--w------- 1 root root 4096 Sep 5 15:29 unbind
--w------- 1 root root 4096 Sep 5 15:29 uevent
-r-------- 1 root root 4096 Sep 5 15:29 slots
--w------- 1 root root 4096 Sep 5 15:29 remove_slot
--w------- 1 root root 4096 Sep 5 15:29 remove_id
-rw------- 1 root root 4096 Sep 5 15:29 quirks
-rw------- 1 root root 4096 Sep 5 15:29 permissive
--w------- 1 root root 4096 Sep 5 15:29 new_id
lrwxrwxrwx 1 root root 0 Sep 5 15:29 module ->
../../../../module/pciback
--w------- 1 root root 4096 Sep 5 15:34 new_slot
--w------- 1 root root 4096 Sep 5 15:34 bind
```

3. Navigate to the `/sys/bus/pci/drivers/be2net` directory, and ensure that the following is displayed:

```
--w------- 1 root root 4096 Sep 5 15:32 uevent
--w------- 1 root root 4096 Sep 5 15:32 remove_id
--w------- 1 root root 4096 Sep 5 15:32 new_id
lrwxrwxrwx 1 root root 0 Sep 5 15:32 module ->
../../../../module/be2net
--w------- 1 root root 4096 Sep 5 15:32 bind
--w------- 1 root root 4096 Sep 5 15:33 unbind
lrwxrwxrwx 1 root root 0 Sep 5 15:32 0000:07:0b.5 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.5
lrwxrwxrwx 1 root root 0 Sep 5 15:32 0000:07:0b.4 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.4
lrwxrwxrwx 1 root root 0 Sep 5 15:32 0000:07:0b.3 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.3
lrwxrwxrwx 1 root root 0 Sep 5 15:32 0000:07:0b.2 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.2
lrwxrwxrwx 1 root root 0 Sep 5 15:32 0000:07:0b.1 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.1
lrwxrwxrwx 1 root root 0 Sep 5 15:32 0000:07:0b.0 ->
../../../../devices/pci0000:00/0000:00:09.0/0000:07:0b.0
```

4. Unbind the VF from the be2net driver.

```
echo –n “0000:07:0b.5” > /sys/bus/pci/drivers/be2net/unbind
```

5. Bind the driver to the pciback module.

```
echo –n “0000:07:0b.5” > /sys/bus/pci/drivers/pciback/new_slot
echo –n “0000:07:0b.5” > /sys/bus/pci/drivers/pciback/bind
```

6. Navigate to the `/sys/bus/pci/drivers/pciback` directory and ensure that the device *0000:07:0b.5* is listed under it. In addition, check that dmesg logs report the same device.

   You can now launch `qemu-kvm` and attach VF *0000:07:0b.5* to any desired VM.

3.2.7.5 Link State Reporting with SR-IOV

When VEB is used for switching traffic between the functions of an SR-IOV-enabled port, the link status reported to the operating system stack when the physical port is not connected varies with the adapter type. In OCe11000-series adapters, the VFs continue to indicate the link to be up even when the physical port is not linked up. In OCe14000-series and OCe11100 series adapters, by default, the VFs indicate the link as down when the physical port is not linked up. This behavior enables two or more VFs to be configured in a bond.

The following table summarizes the default link status of the PFs and VFs when SR-IOV is enabled in the BIOS and in the driver.

Table 6 Default Link Status of the PFs and VFs

Function	Physical Link is Not Connected	
	OCe14000	OCe11100
Physical function (PF)	Down	Down
Virtual function (VF)	Down	Down

The default behavior of VFs in OCe14000-series and OCe11100 series adapters can be changed using the `ip link` command in distributions of Linux that support iproute version 3.11.0 or later. For details on configuring the link state for VFs, see Section 3.2.7.7, Link State Configuration.

When VEPA is configured to switch traffic between the functions instead of VEB, the VF link state always reflects the physical state of the associated port, and this behavior cannot be changed.

3.2.7.6 Configuring VFs

In operating system distributions with newer IP commands that support VF configuration options, the host administrator can perform the following by using the `ip link set` command:

- Change the default MAC address
- Configure VLAN
- Configure the transmission rate
- Configure the link state
- Set the QoS parameter on VFs

MAC Address Configuration

The Emulex OneConnect adapter is shipped with factory-configured MAC addresses for the network interfaces corresponding to the PFs. The driver generates random MAC addresses for the network interfaces corresponding to the VFs based on the factory-configured MAC address. Other MAC addresses can be assigned for the interfaces corresponding to the VF using IP utility commands in the hypervisor.

To configure the MAC address for the virtual function, run the following command in the hypervisor:

```
# ip link set eth<X> vf <VFN> mac <MAC-ADDR>
```

where:

- eth<*X*> is the interface corresponding to the physical function.
- <*VFN*> is the VF number (0-based) corresponding to the interface for which you are configuring the MAC address.
- <*MAC-ADDR*> is the MAC address you are configuring.

For example, to configure the MAC address for VF number 0 on eth0 (PF), run the following command in the hypervisor:

```
# ip link set eth0 vf 0 mac 00:16:88:AA:BB:AA
```

If the VM is already running and the VF driver is loaded, the VF MAC address can be updated using one of the following two methods in the VM:

- If the VF interface is configured using `ifcfg` script:
  a. Update the `ifcfg` script for the VF, with the MAC address configured in the hypervisor in the earlier step. For example, update the HWADDR line in the VF `ifcfg-eth`<X> script.
     `HWADDR=00:16:88:AA:BB:AA`
  b. Restart the network service.
     `# service network restart`

Or

- If the VF interface is not configured using `ifcfg` script:
  Update the VF MAC address using the `ifconfig` or `ip link` command.
  `# ifconfig eth0 hw ether 00:16:88:AA:BB:AA`

VLAN Configuration

This section includes examples for transparent tagging and virtual guest tagging (VGT).

When transparent tagging is configured for a VF, the NIC transparently tags all non-VLAN traffic from the VF with the configured transparent VLAN ID. The VM is not aware of the VLAN tag.

To assign a transparent VLAN ID to the VF, run the following IP command in the hypervisor:

```
# ip link set eth<X> vf <VFN> vlan <VLAN>
```

where:

- eth<*X*> is the interface corresponding to the physical function
- <*VFN*> is the VF number corresponding to the interface for which you are configuring the VLAN
- <*VLAN*> is the VLAN ID you are configuring

Example:

```
# ip link set eth0 vf 0 vlan 5
```

NOTE Guest VLAN tagging is not allowed on the VF in the guest operating system when transparent VLAN tagging is enabled on the VF.

Transmission Rate Configuration

Configure the transmission rate limit (TX-RATE) on a VF interface from the hypervisor using the following IP command syntax:

```
# ip link set eth<X> vf <VFN> rate <TX-RATE>
```

where:

- eth<*X*> is the interface corresponding to the physical function
- <*VFN*> is the VF number corresponding to the interface that you are configuring the TX-RATE
- <*TX-RATE*> is the transmission rate limit, in Mb/s

For example, to set a TX-RATE of 5000 Mb/s for the VF 0, run the following command:

```
# ip link set eth0 vf 0 rate 5000
```

NOTE For OCe14000-series adapters, the TX-RATE must be multiples of 400 Mb/s on a 40Gb port, and multiples of 100 Mb/s on a 10Gb port.

3.2.7.7 Link State Configuration

NOTE The link state configuration feature is supported on RHEL 7.*x* and SLES 12 onwards and is not supported in earlier kernels.

This section provides link state configuration for a VF in OCe14000-series adapters.

Configure the link state on a VF interface from the hypervisor using the following IP command syntax:

```
#ip link set eth<X> vf <VFN> state < auto | enable | disable >
```

where:

- `eth<X>` is the interface corresponding to the physical function
- <*VFN*> is the VF number corresponding to the interface for which the link state is being configured
- `auto` – VF link state will reflect the PF link state
- `enable` – VF link state will be always up
- `disable` – VF link state will be always down

3.2.7.8 Spoof Check Configuration

Configure the HW packet source MAC spoof check on a VF interface from the hypervisor using the following IP command syntax:

```
# ip link set eth<X> vf <VFN> spoofchk [on|off]
```

where:

- `eth<X>` is the interface corresponding to the physical function
- <*VFN*> is the VF number corresponding to the interface for which you are configuring the spoof check

For example, to turn on spoof checking for the VF 0, run the following command:

```
# ip link set eth0 vf 0 spoofchk on
```

3.2.7.9 Viewing VF Properties

To view the properties configured to VFs attached to a PF, use the following IP command syntax:

```
# ip link show eth<X>
```

where `eth<X>` is the interface corresponding to the physical function.

For example, to view the properties of a PF at eth0 (which has VFs 0, 1 associated with it), run the following command in the hypervisor:

```
# ip link show eth0
```

Expected example output:

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen
1000
link/ether 00:00:c9:bb:16:ee brd ff:ff:ff:ff:ff:ff
vf 0 MAC 00:00:c9:9d:90:80, tx rate 10000 (Mbps)
vf 1 MAC 00:00:c9:9d:90:81, tx rate 10000 (Mbps)
```

3.2.8 Bonding Considerations

VFs can be added to a bonding configuration. But because VFs are typically assigned to a VM, it might not be possible to identify two VFs from the same PF or physical port. This situation could lead to incorrect configurations. To avoid this issue, the driver exports a physical port identifier corresponding to a given network interface. VFs from the same physical port will have the same identifier.

For example, to obtain the physical port identifier of a VF `eth0` in a VM, read the corresponding sysfs file `phys_port_id` in the VM as follows. The number in the output shown is the physical port identifier of eth0.

```
# cat /sys/class/net/eth0/phys_port_id
0100000000000000000000383838314643
```

The NetworkManager reports incorrect bonding configuration with a message like this:

```
"NetworkManager: <warn> (bond0): slave eth1 shares a physical port with
existing slave eth0"
```

3.2.9 ARI Considerations for OCe14000-Series Adapters

Alternative routing-ID interpretation (ARI) is supported on the OCe14000-series adapters. The maximum number of functions allowed on an adapter is controlled by the adapter's IPL file and the system's support for ARI.

When UMC is enabled on an OCe14000-series network adapter, each port can be partitioned into isolated PFs (channels). You can configure the following number of functions:

- Up to 16 functions on a one-port OCe14400-series adapter
- Up to eight functions per port on a one or two-port OCe14100-series adapter
- Up to four functions per port on a four-port OCe14100-series adapter

NOTE Refer to the *Emulex Universal Multichannel Reference Manual* for additional information on UMC.

ARI must be available to support up to 16 functions on an adapter. OCe14000-series adapters automatically support ARI. However, the following requirements must be met to support more than eight functions on an adapter:

- The system hardware (the motherboard and BIOS) must support ARI.
- ARI must be enabled in the system BIOS.
- The host or guest operating system must support ARI:
  - RHEL 6.4 and later versions
  - RHEL 7
  - SLES 11 SP2 and later versions
- The application management tools, including the Emulex OneCommand CNA Manager application, must support ARI.

If these conditions are not met, you may be able to configure more than eight functions, but only up to eight functions will be running and discovered after a reboot.

3.2.10 NIC Partitioning (NPar) Configuration (Dell Only)

NOTE

- NPar support is only available on OCe14000-series adapters running in 10Gb mode.
- NPar is tested and supported on Dell 12G servers.
- NPar + ARI = NParEP is tested and supported on Dell 13G servers.

NPar mode enables the capability to divide a 10Gb NIC port into multiple PCI functions, with flexible bandwidth capacity allocation that appear to the operating system and network as separate NIC ports. A single 10-gigabit Ethernet (GbE) port appears as multiple physical devices showing in PCI configuration space as multiple functions.

3.2.10.1 Adapter Configuration

NPar can be configured on OCe14000-series adapters by using the BIOS or the Emulex OneCommand CNA Manager application.

- To configure NPar using the BIOS, refer to the *Boot for NIC, iSCSI, and FCoE Protocols User Guide*.
- To configure NPar using the Emulex OneCommand CNA Manager, refer to the latest *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide* or the *Emulex OneCommand CNA Manager CLI for OneConnect Adapters User Guide*.

On the host operating system side, NPar provides up to eight PCI functions per device using standard PCI configuration space. Four PCI functions can be mapped to a physical port. Each function or partition will be assigned a unique MAC address.

Partitions will be available for virtual function assignment and for application segmentation through VLAN or IP subnets.

3.2.10.2 NPar Features

- The partitions can be on separate subnets or VLANs.
- Bandwidth allocation is flexible.
- No operating system or BIOS changes are required.
- No external switch changes are required.
- Each partition should have standard NIC properties for stateless offload.
- NIC teaming on the same port should be avoided.

The following items are supported on a per-partition basis:

- Per-partition statistics
- LSO, large receive offload (LRO), RSS, and TCP Segmentation Offload (TSO) needed per partition
- Maximum transmission unit (MTU) per partition
- Support for NetQueues

3.2.10.3 Using NParEP

NOTE

- NParEP is available only on OCe14000-series adapters.
- On a four-port adapter, ARI functionality must be enabled in the PCIe subsystem on a particular system to support NParEP on all four ports.
- NParEP support is available only on Dell 13G or newer systems.

- SR-IOV must be enabled in the system BIOS for NParEP to work properly on Linux systems.

NParEP can be configured on the OCe14000 family of adapters by using the adapter BIOS utility or the Emulex OneCommand CNA Manager application.

- To configure NParEP using the adapter BIOS utility, refer to the *Boot for NIC, iSCSI, and FCoE Protocols User Guide*.
- To configure NParEP using the Emulex OneCommand CNA Manager application, refer to the *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide* or the *Emulex OneCommand CNA Manager CLI for OneConnect Adapters User Guide*.

3.3 RoCE Configuration for OCe14000-Series Adapters

NOTE

- RoCE is not supported with multichannel.
- RoCE configurations are not supported with SR-IOV.
- Enable priority flow control (PFC) as the default mode whenever possible while using RoCE. See Section 3.4, iSCSI Driver Configuration, to enable QoS for RoCE. In addition, VLAN interfaces must be configured and used for RoCE trafficr for PFC to work correctly (see Section 3.3.2, Pinning Interrupts to a CPU Core). Use VLANs greater than 1 for best interoperability.

3.3.1 Basic Configuration

3.3.1.1 Setting the Profile

Specify and enable RoCE profiles can be specified and enabled using the OneCommand CNA Manager GUI application, the OneCommand CNA Manager CLI application, and the PXESelect BIOS utility.

For specific information on selecting RoCE profiles, refer to the applicable guide:

- *Boot for NIC, iSCSI, and FCoE Protocols User Guide* for details on the PXESelect BIOS utility.
- *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide* for information about setting the RoCE profiles using the OneCommand CNA Manager application GUI.
- *Emulex OneCommand CNA Manager CLI for OneConnect Adapters User Guide* for information on setting the RoCE profiles using the OneCommand CNA Manager CLI.

3.3.1.2 Confirming that the RoCE Profile Is Enabled

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

For example: Using profile 0x15.

NOTE If the correct profile is not reported, update the profile ID using the Emulex OneCommand CNA Manager application or the PXESelect BIOS utility.

CAUTION You must have a homogeneous set up in terms of firmware version, active profile, and driver version. If not, the applications could fail.

3.3.1.3 Interface Configuration

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

To list the ocrdma interfaces:

```
ibv_devinfo –l
```

To list the corresponding NIC interfaces:

```
ibdev2netdev
```

Follow the standard procedure to assign a valid IP address to the desired Ethernet interface that corresponds to the RoCE port. You can assign an IP address to eth0 to use `ocrdma0` for RoCE.

Example output:

```
# ibv_devinfo –l
2 HBAs found:
ocrdma1
ocrdma2
# ibdev2netdev
ocrdma0 port 1 ==> eth0 (Up)
ocrdma1 port 1 ==> eth1 (Up)
ocrdma2 port 1 ==> eth2 (Up)
ocrdma3 port 1 ==> eth3 (Up)
```

On some systems, the `ibdev2netdev` command might not be available. On these systems, perform the following steps to find the Ethernet interface names from the RoCE interface name.

```
cat /sys/class/infiniband/ocrdma<X>/nic_interface
```

where `ocrdma<X>` is the output of the `ibv_devinfo -l` command.

For example:

```
# cat /sys/class/infiniband/ocrdma0/nic_interface
```

3.3.1.4 VLAN Configuration

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

To configure a VLAN interface, perform the following steps:

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

   ```
   modprobe 8021q
   ```

   NOTE `modprobe 8021q` must be loaded even if VLANs are not configured.

2. Create a VLAN interface by typing:

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

   For example:

   ```
   # vconfig add eth4 100
   ```

   On systems where `vconfig` is not available, the `ip` command must be used.

   ```
   ip link add link <base interface name> name <vlan interface name> type vlan
   id <vlan-id>
   ```

   For example:

   ```
   # ip link add link eth4 name eth4.100 type vlan id 100
   ```

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

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

For example:

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

Testing the RoCE Interface Connectivity

You can test the RoCE interface connectivity using the `rping` command.

On a server, type:

```
# rping  -sva <server IP>
```

On a client, type:

```
# rping -C <iteration> -cva <server IP>
```

For example:

On a server, when you type:

```
# rping -sva  3.3.3.2
```

The following sample output appears:

```
server ping data: rdma-ping-0: ABCDEFGHIJKLMNOPQRSTUVWXYZ[\]'_`abcdefghijklmnopqr
server ping data: rdma-ping-1: BCDEFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrs
server ping data: rdma-ping-2: CDEFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrst
server ping data: rdma-ping-3: DEFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstu
server ping data: rdma-ping-4: EFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuv
server DISCONNECT EVENT...
wait for RDMA_READ_ADV state 10
```

On a client, when you type:

```
# rping -C 5 -cva 3.3.3.2
```

The following sample output appears:

```
ping data: rdma-ping-0: ABCDEFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqr
ping data: rdma-ping-1: BCDEFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrs
ping data: rdma-ping-2: CDEFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrst
ping data: rdma-ping-3: DEFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstu
ping data: rdma-ping-4: EFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuv
client DISCONNECT EVENT...
```

3.3.1.5 Sample Applications

The following additional in-built tools are available for RoCE testing.

- ib_send_bw
- ib_read_bw
- ib_write_bw

NOTE The use of these commands assumes that the client (ocrdma0) and server (ocrdma0) interfaces are configured with IPs 11.192.168.x and 11.192.168.y, respectively.

For example:

- ib_send_bw

  Server:

  ```
  # ib_send_bw -Rd ocrdma0
  ```

  Client:

  ```
  # ib_send_bw -Rd ocrdma0 11.192.168.x
  ```

- ib_read_bw

  Server:

  ```
  # ib_read_bw –Rd ocrdma0
  ```

  Client:

  ```
  # ib_read_bw –Rd ocrdma0 11.192.168.x
  ```

- ib_write_bw

  Server:

  ```
  # ib_write_bw –Rd ocrdma0
  ```

  Client:

  ```
  # ib_write_bw –Rd ocrdma0 11.192.168.x
  ```

3.3.2 Pinning Interrupts to a CPU Core

On a multicore system, distribute RoCE interrupt request (IRQ) vectors among CPU cores to achieve better system utilization.

RoCE interface IRQs can be pinned to different CPU cores by disabling `irqbalance` and running:

```
ocrdma_irq_affinity.sh script
```

(included with Linux RoCE driver package)

```
# service irqbalance stop
# ./ocrdma_irq_affinity.sh  <RoCE interface name>
```

3.3.3 MTU Configuration

The minimum supported MTU for RoCE is 512. However, if you use the `ifconfig` command to set the MTU value on a corresponding Ethernet interface, 80 bytes must be added to the MTU of the RoCE interface. For example, if 512 is the desired MTU on RoCE interface, 592 is the minimum MTU size to be configured on the corresponding Ethernet interface.

To utilize the maximum RoCE MTU of 4096, the Ethernet interface must be configured with an MTU greater than 4176.

Perform the following procedure when changing the MTU value on a network interface or a switch port:

1. Stop all running RoCE applications.
2. Change the MTU value (on the host interface and the switch).
3. Restart all the RoCE applications.

3.3.4 QoS Configuration

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

In addition to QOS settings, the OCe14000-series adapter can support quantized congestion notification (QCN) for RoCE ports. To use QCN, the feature must be enabled on both the RoCE port and the switch port. The RoCE port can be enabled using the Emulex OneCommand CNA Manager application in the Physical Port Info tab. For more

information, refer to the latest *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide*. The switch must also be enabled to generate QCN packets.

3.3.4.1 QoS Behavior

Supported QoS behavior follows:

- PFC configuration for back to back
- Limited QoS configuration using the Emulex OneCommand CNA Manager application
- A single traffic class group for RoCE per port
- A single RoCE priority in PFC mode
- Bandwidth allocation for priority groups

OCe14000-Series Adapter Defaults

NOTE If generic pause mode is used, ensure that the switches support this feature. Use PFC with priority 5 if the switch does not support the generic pause mode.

- Adapter boot time
  - PFC is disabled on all the ports at adapter boot time in the NIC + RoCE profile.
  - The generic pause mode is enabled on all the ports at adapter boot time in the NIC + RoCE profile.
- Back-to-back connection (OCe14000-to-OCe14000)
  - PFC is enabled by default.
  - The generic pause mode is disabled on that port.
- Data Center Bridging Capabilities Exchange (DCBX)-enabled switch connection
  - If the OCe14000-series adapter is connected to a DCBX-enabled switch, the mode changes from generic pause to PFC mode.
  - The OCe14000-series adapter configures RoCE traffic for priority 5.
  - Priority 5 is manually enabled on a switch under a priority group other than an FCoE, ISCSI, or NIC priority group.

NOTE In the absence of priority 5 at the switch side, the OCe14000-series adapter maintains its configuration for PFC mode priority 5. This configuration can result in packet losses, unrecoverable errors, or infinite retries for RoCE traffic.

- DCBX-disabled switch connection
  - If the OCe14000-series adapter is connected to a DCBX-disabled switch, the generic pause mode is enabled.

3.3.4.2 QoS Configuration Guidelines

This section details QoS configuration.

Priority Groups

Split the traffic into two or more priority groups:

- One priority group for RoCE
- Other groups for non-RoCE traffic

Many RoCE applications use Transmission Control Protocol (TCP)/IP for out-of-band connection establishment. Therefore, you should allocate sufficient bandwidth to non-RoCE priority groups.

L2 Flow Control

When a port is operating in generic pause mode, RoCE latencies can increase. In this situation, configure RoCE to use PFC for better results.

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

DCBX-Enabled Switch Configuration

Currently, switches do not support configuring priority for RoCE specific traffic. Priority 5 must be manually enabled on the switch under a priority group other than the FCoE, iSCSI, or NIC priority group.

NOTE In the absence of priority 5 at the switch side, the OCe14000-series adapter maintains its configuration for PFC mode for priority 5. This can result in packet losses, unrecoverable errors, or infinite retries for RoCE traffic.

To configure the switch:

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

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

To configure the host:

1. Enable PFC using the Emulex OneCommand CNA Manager application (refer to the *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide)*.
2. Create a VLAN.
3. Assign an appropriate IP address to the VLAN interface.

Example Switch PFC Configuration

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

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

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

To configure the switch PFC:

1. Use default-group as the non-RoCE priority group (see DCBX-Enabled Switch Configuration, step 4).
2. Create another group for priority 5, such as PG 5 (see DCBX-Enabled Switch Configuration, step 1).
3. Assign 90 percent bandwidth to the PG 5 group and 10 percent to the default-group. No other changes are required to the default-group or fcoe-group.
4. Configure the individual ports in trunk mode, enable PFC-allowed VLANs and disable generic pause mode.

The following are examples of switch PFC configurations:

Cisco Global QoS Configuration

Global QoS configuration on the Cisco switch:

```
class-map type qos roce
      match qos 5
class-map type queuing roce
      match qos-group 5
class-map type network-qos roce
      match qos-group 5
policy-map type qos roce
      class roce
            set qos-group 5
      class class-fcoe
            set qos-group 1
      class class-default
policy-map type queuing roce
      class type queuing roce
            bandwidth percent 90
      class type queuing class-fcoe
            bandwidth percent 0
      class type queuing class-default
            bandwidth percent 10
policy-map type network-qos roce
      class type network-qos roce
      pause no-drop
      mtu 4200
      class type network-qos class-default
      mtu 9216
class type network-qos class-fcoe
      pause no-drop
      mtu 2158
system qos
      service-policy type qos input roce
      service-policy type queuing input roce
      service-policy type queuing output roce
      service-policy type network-qos roce

      #Individual port configuration
      interface ethernet 1/13
            switchport mode trunk
            switchport trunk allowed vlan 4
            priority-flow-control mode auto
            flowcontrol receive off
                  flowcontrol send off
```

DCBX-Disabled Switch Connection (Generic Pause Mode)

1. Host configuration:

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

   a. To verify status:

   ```
   ethtool -a ethX
   ```

b. To configure:

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

2. Switch configuration:
   a. Enable Tx and Rx generic pause flow control on each port.
   b. Some switches have jumbo-frame size support disabled by default on the port or global level. Enable jumbo-frame support, or set MTU to at least 4200.

3.3.5 Advanced Applications

3.3.5.1 NFS over RDMA

On NFS servers where *fsid* must be specified while exporting file systems, ensure that they are unique for each exported file system.

When trying multiple mounts using NFS RDMA, the following error might be seen when the hardware resources are exhausted.

```
mount.nfs4: Cannot allocate memory
```

In this situation, use smaller `rdma_slot_table_entries` counts until the mount succeeds.

NOTE For the new `rdma_slot_table_entries` counts to take effect across all mounts, you must unmount the existing mounts before retrying all the mounts.

To get the current slot table entries, type:

```
# cat /proc/sys/sunrpc/rdma_slot_table_entries
32
```

To modify the slot table entries, type:

```
echo <X> >/proc/sys/sunrpc/rdma_slot_table_entries
```

where <*X*> is the new count for slot table entries.

For example:

```
# echo 16 >/proc/sys/sunrpc/rdma_slot_table_entries
```

3.3.5.2 Server Configuration

To configure the server:

1. Load and configure the ocrdma driver/library.
2. Load NFS-RDMA server module:

   ```
   #modprobe svcrdma
   ```

3. Start the NFS server:

   `# service nfs start` (for RHEL)

   Or

   `#service nfsserver start` (for SLES)

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

   ```
   #echo “rdma 20050”>/proc/fs/nfsd/portlist
   ```

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

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

NOTE Each `fsid` value must be unique.

For example:

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

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

```
#exportfs -a
```

3.3.5.3 Client Configuration

To configure the client:

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

```
#modprobe xprtrdma
```

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

```
#showmount –e <server_roce_ip>
```

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

For example:

```
#showmount -e 11.192.168.1
```

Export list for `11.192.168.1`:

```
/export (everyone)
```

4. Mount the file system:

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

For example:

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

5. Verify the NFS mount using RDMA:

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

3.3.5.4 Using Multiple Queue Pairs with NFS-RDMA

In OCe14400-series 40GbE adapters, it is possible to use multiple queue pairs per NFS-RDMA mount for improved performance in terms of I/O operations per second (IOPs) and throughput. If this feature is enabled on the adapter port, the provider creates multiple queue pairs for every individual mount point. This feature is disabled by default and must be enabled manually for every port on both the NFS client and NFS server host. In case multiple queue pairs are enabled only on one of the hosts, the feature is disabled by the Emulex RoCE driver.

3.3.5.5 Enabling Multiple Queue Pairs

To enable multiple queue pairs for a port, type:

```
# echo 1 > /sys/class/infiniband/ocrdma<x>/qp_expansion
```

where `ocrdma`<*x*> is the interface exported for OCe14400-series; for example, `ocrdma0`.

3.3.5.6 Verifying Multiple Queue Pairs

To verify if multiple queue pairs are enabled, type:

```
# tail /var/log/messages
```

A sample output indicating four queue pairs would look like the following:

```
ocrdma(0) 40GBps  Adapter. qp_expansion : Setting the expansion factor to 4
```

3.3.5.7 Disabling Multiple Queue Pairs

To disable multiple queue pairs, type:

```
# echo 0 > /sys/class/infiniband/ocrdma<x>/qp_expansion
```

where `ocrdma<x>` is the interface exported for OCe14400-series; for example, `ocrdma0`.

NOTE You cannot disable mount points that are multiple queue pair enabled after they are mounted.

3.3.6 Installing and Using DAPL

The native RDMA stack may not install the Direct Access Programming Library (DAPL)-related RPMs by default. In this case, you must use the following steps to manually download and install the DAPL and DAPL utilities:

1. Install the following RPMs:

   ```
   libibverbs-devel
   librdmacm-devel
   ```

2. Download the DAPL tar ball. A supported DAPL tar ball can be downloaded from the following website:

   http://downloads.openfabrics.org/downloads/dapl/

3. Type the following for the build steps:

   ```
   # tar -zxvf dapl-x.x.x.tar.gz
   # cd dapl-x.x.x
   # ./configure
   # make
   # make install
   ```

4. Configure DAPL. Edit the `/etc/dat.conf` file to append following lines:

   ```
   ofa-v2-scm-roe-ocrdma0-1 u2.0 nonthreadsafe default libdaploscm.so.2
   dapl.2.0 "ocrdma0 1" ""
   ofa-v2-scm-roe-ocrdma1-1 u2.0 nonthreadsafe default libdaploscm.so.2
   dapl.2.0 "ocrdma1 1" ""
   ofa-v2-scm-roe-ocrdma2-1 u2.0 nonthreadsafe default libdaploscm.so.2
   dapl.2.0 "ocrdma2 1" ""
   ofa-v2-scm-roe-ocrdma3-1 u2.0 nonthreadsafe default libdaploscm.so.2
   dapl.2.0 "ocrdma3 1" ""
   ```

3.4 iSCSI Driver Configuration

3.4.1 Open-iSCSI Support

Open-iSCSI is a high-performance, transport independent, multi-platform implementation of Request for Comments (RFC) 3720.

For the inbox Open-iSCSI driver, use the Open-iSCSI tools to configure and manage Open-iSCSI. For the Emulex out-of-box iSCSI driver, use the Emulex OneCommand CNA Manager application to configure and manage Open-iSCSI. The applications are not interchangeable between the drivers.

The `iscsiadm` utility is a command-line tool allowing discovery of iSCSI targets; logging into iSCSI targets; as well as, access and management of the open-iSCSI database on all Linux installations. This utility presents a set of operations that you can perform on iSCSI nodes, sessions, connections, and discovery records.

3.4.2 Discovering and Adding iSCSI Targets

You can add iSCSI targets using a discovery phase first or add them directly (see Adding iSCSI Targets Directly (without Discovering)).

Follow these steps to discover and add a target. In this example, based on the IPL file flashed, the adapter has two NIC functions and two iSCSI functions.

1. To configure iface, type:

   ```
   iscsiadm -m iface
   ```

   Or, to set the network properties of iface, type:

   ```
   iscsiadm -m iface -I <iface name> --op=update -n iface.bootproto -v
   <static|dhcp>
   ```

2. If you selected static in step 1, you must configure the following manually:

   ```
   iscsiadm -m iface -I <iface name> --op=update -n iface.ipaddress -v <IPv4
   address>
   iscsiadm -m iface -I <iface name> --op=update -n iface.subnet_mask -v
   <netmask>
   iscsiadm -m iface -I <iface name> --op=update -n iface.gateway -v <IPv4
   address>
   To configure VLAN following commands need to be executed:
   iscsiadm -m iface -I <iface name> --op=update -n iface.vlan_id -v <1…4095>
   iscsiadm -m iface -I <iface name> --op=update -n iface.vlan_state -v
   <enable/disable>
   iscsiadm -m iface -I <iface name> --op=update -n iface.vlan_priority -v
   <0…7>
   iscsiadm -m iface -I <iface name> -o apply
   ```

   NOTE The apply operation pushes the configuration to the hardware using the iSCSI driver.

3. Log in to the discovered targets.

   ```
   iscsiadm -m node -p 192.168.65.196:3260 -l
   ```

   Example output:

   ```
   Logging in to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt7, portal:
   192.168.65.196,3260]
   Logging in to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt3, portal:
   192.168.65.196,3260]
   ```

```
Logging in to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt1, portal:
192.168.65.196,3260]
Logging in to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt0, portal:
192.168.65.196,3260]
Logging in to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt6, portal:
192.168.65.196,3260]
Logging in to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt4, portal:
192.168.65.196,3260]
Logging in to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt2, portal:
192.168.65.196,3260]
Logging in to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt5, portal:
192.168.65.196,3260]
Login to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt7, portal:
192.168.65.196,3260] successful.
Login to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt3, portal:
192.168.65.196,3260] successful.
      Login to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt1,
portal: 192.168.65.196,3260] successful.
Login to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt0, portal:
192.168.65.196,3260] successful.
Login to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt6, portal:
192.168.65.196,3260] successful.
Login to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt4, portal:
192.168.65.196,3260] successful.
Login to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt2, portal:
192.168.65.196,3260] successful.
Login to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt5, portal:
192.168.65.196,3260] successful.
```

4. Display the session information after a successful login.

   After the login is successful, the session information can be dumped with the following command:

   ```
   iscsiadm -m session -P 3
   ```

5. Log out from the targets.

   To log out from a single target, use the following command:

   ```
   iscsiadm -m node -T <Target_Name> -u
   ```

   To log out from all targets, use the following command:

   ```
   iscsiadm -m node -u
   ```

Adding iSCSI Targets Directly (without Discovering)

If the target details are known, you can skip the discovery phase and add the targets to the iscsiadm node database directly. Then you can log in to those targets.

1. Add individual target details to the nodes database.

   ```
   iscsiadm -m node -o new -T iqn.tgt1 -p 192.168.65.196:3260 -I
   be2iscsi.00:00:c9:f2:73:91
   ```

   Example output:

   ```
   New iSCSI node
   [be2iscsi:[hw=00:00:c9:f2:73:91,ip=,net_if=,iscsi_if=be2iscsi.00:00:c9:f2:7
   3:91] 192.168.65.196,3260,-1 iqn.tgt1] added
   ```

   The `-I be2iscsi.00:00:c9:f2:73:91` option is passed so that the target is bound with the interface name, and the login occurs through those interfaces.

2. Log in to the specific target.

```
iscsiadm -m node -T iqn.tgt1 -l
```

Example output:

```
Logging in to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt1, portal:
192.168.65.196,3260]
Login to [iface: be2iscsi.00:00:c9:f2:73:91, target: iqn.tgt1, portal:
192.168.65.196,3260] successful.
```

3. To display the session after a successful login or to log out, see step 4 and step 5 in Section 3.4.2, Discovering and Adding iSCSI Targets.

3.4.3 iscsiadm Commands for Configuring the Target

After setting up the target and initiator machines, use the following procedure to configure the iSCSI target through Open-iSCSI.

NOTE The following instructions assume that the initiator machine is a Linux machine with an Emulex OneConnect adapter installed.

1. Create a new interface (iface):

```
iscsiadm -m iface -o new -I <ifacename>
```

In this command, <*ifacename*> is the name you provide for the iface.

2. Add a NIC media access control (MAC) address:

```
iscsiadm -m iface -I <ifacename> --op=update -n iface.hwaddress -v <NIC
mac_address>
```

In this command, <*ifacename*> is the name of the iface created in step 1, and <*NIC mac_address*> is the NIC MAC address you are adding for discovery.

3. Add the transport name:

```
iscsiadm -m iface -I <ifacename> --op=update -n iface.transport_name -v
be2iscsi
```

In this command, <*ifacename*> is the name of the iface created in step 1.

4. Restart the service:

```
service open-iscsi restart
```

5. Perform target discovery using `SendTargets`:

```
iscsiadm -m discovery -t st -p <ip:port> -I <iface> -P 1
```

In this command, <*ip:port*> is the IP address and port number and <*iface*> is the name of the iface created in step 1.

6. Add the iSCSI driver MAC address:

```
iscsiadm -m iface -I <ifacename> --op=update -n iface.hwaddress -v <iSCSI
MAC Address>
```

In this command, <*ifacename*> is the name of the iface created in step 1.

7. Set the initiator's IP address:

```
iscsiadm -m iface -I <ifacename> -o update -n iface.ipaddress -v <IP>
```

In this command, <*ifacename*> is the name of the iface created in step 1 and <*IP*> is the IP address.

8. Log into the target:

```
iscsiadm -m node -T <targetname> -p <ip:port> -I <iface> -l
```

In this command, <*targetname*> is the target name you want to log into; <*ip:port*> is the IP address and port number; and <*iface*> is the name of the iface created in step 1.

9. Verify that the disk shows:

```
fdisk -l
```

10. Log out:

```
iscsiadm -m node -T <targetname> -p <ip:port> -I <iface> -u
```

In this command, <*targetname*> is the target name, <*ip:port*> is the IP address and port number, and <*iface*> is the name of the iface created in step 1.

Example

```
iscsiadm -m iface -o new -I iface0
iscsiadm -m iface -I iface0 --op=update -n iface.hwaddress -v
00:0F:1F:62:2B:BF
iscsiadm -m iface -I iface0 --op=update -n iface.transport_name -v be2iscsi
service open-iscsi restart
iscsiadm -m discovery -t st -p 20.0.0.107:3260 -I iface0 -P 1
iscsiadm -m iface -I iface0 --op=update -n iface.hwaddress -v
00:0F:1F:92:6B:BF
iscsiadm -m iface -I iface0 -o update -n iface.ipaddress -v 20.0.0.107
iscsiadm -m node -T iqn.tgt0 -p 20.0.0.107:3260 -I iface0 -l fdisk -l
iscsiadm -m node -T iqn.tgt0 -p 20.0.0.107:3260 -I iface0 –u
```

3.4.4 iSNS Discovery Using iscsiadm

Internet Storage Name Service (iSNS discovery uses a standard NIC interface. The default interface is the TCP_interface.

```
localhost:~ # iscsiadm -m iface
default tcp,<empty>,<empty>,<empty>,<empty>
iser iser,<empty>,<empty>,<empty>,<empty>
```

Perform the following steps to discover targets using iSNS through the NIC interface:

1. Ensure that a ping to the iSNS server IP from the initiator machine is successful through the NIC interface.
2. Run the discovery command:

```
iscsiadm -m discoverydb -t isns -p <iSNS_SERVER_IP> -I default --discover -P
1
```

For example:

```
iscsiadm -m discoverydb -t isns -p 10.192.204.116 -I default --discover -P1
```

Example output:

```
Target: iqn.tgt:TGT1
      Portal: 192.168.65.197:3260,1
            Iface Name: default
Target: iqn.tgt:TGT2
      Portal: 192.168.65.197:3260,1
            Iface Name: default
Target: iqn.tgt:TGT3
      Portal: 192.168.65.197:3260,1
            Iface Name: default
```

3. Display the interfaces that are present:

```
iscsiadm -m iface
```

Example output:

```
default tcp,<empty>,<empty>,<empty>,<empty>
```

```
iser iser,<empty>,<empty>,<empty>,<empty>
bnx2i.00:25:b3:df:8d:ad bnx2i,00:25:b3:df:8d:ad,<empty>,<empty>,<empty>
be2iscsi.00:00:c9:f2:73:91
```

4. Add the targets listed in the discovery command to the iSCSI driver interface through which the connection will be offloaded:

   ```
   iscsiadm -m node -T <Target_Name> -p <Target_Portal:Port_Number,Portal Tag>
   -I
   <be2iscsi Interface> -o new
   ```

   For example:

   ```
   iscsiadm -m node -T <Target_Name> -p <Target_Portal:Port_Number,Portal Tag>
   -I
   ```

   New iSCSI node

   ```
   iscsiadm -m node -T <Target_Name> -p <Target_Portal:Port_Number,Portal Tag>
   -I
   ```

3.5 VPort Configuration

This section describes how to create, delete, and display virtual ports VPorts.

3.5.1 VPort Configuration Prerequisites

Before configuring VPorts, consider the following points:

- Ensure that you are using the latest recommended firmware for VPort functionality. For the latest firmware, check the Documents and Downloads area of http://www.broadcom.com.

  NOTE Before performing a firmware update, driver installation is required. For more information on installing the driver, see Chapter 2, Installing and Uninstalling.

- Loop devices and NPIV are not supported on the same port simultaneously. If you are running a loop topology and you create a VPort, the VPorts link state is offline.
- VPorts do not persist across system reboots.

3.5.2 Creating, Deleting, and Displaying VPorts

VPorts are created through sysfs entries that are presented in the physical port's sysfs directory. The `vport_create` and `vport_delete` sysfs entries are described in Section 3.5.3.2, VPort sysfs Entries. There are three scripts create, delete, and display VPorts. The scripts reside in the `/usr/sbin/brcmlpfc` directory and are part of the OneCommand CNA Manager application kit.

When NPIV is enabled and VPorts are configured, it can take longer for the adapter to finish discovery in some cases because each VPort must perform discovery independently. As more VPorts are configured, the amount of time that the driver and adapter take to finish discovery of remote ports on the SAN increases. To compensate for this extended amount of time taken in discovery, set the `lpfc_devloss_tmo` parameter to 60 seconds when NPIV is enabled.

3.5.2.1 Creating VPorts Using the mkvport.sh Script

You can use the mkvport script to create VPorts. To see the usage information, run the script with no parameters specified. The `mkvport.sh` script uses the following syntax:

```
./mkvport.sh <Physical Port's Host number> <Port Name> <Node Name>
```

You must supply the physical port's host number, WWPN, and WWNN when using the `mkvport.sh` script. For example, to create a VPort with port name of 10000000c94ac63a and a node name of 20010000c94ac63a on the physical port with scsi_host name *host7*, type:

```
./mkvport.sh host7 10000000c94ac63a 20010000c94ac63a
```

This script fails if the VPort is not created.

NOTE It is possible for a VPort to be created successfully but to be in a failed state. For example, loop devices and NPIV are not supported on the same port simultaneously. If you are running a loop topology and you create a VPort, the VPorts link state will be offline.

3.5.2.2 Deleting VPorts Using the rmvport.sh Script

NOTE You must unmap, unmount, and flush I/Os to VPort-connected devices before deleting the VPort.

You can use the `rmvport.sh` script to delete VPorts. To see the usage information, run the script with no parameters specified. The `rmvport.sh` script uses the following syntax:

```
./rmvport.sh <VPort's Host number>
```

Or

```
./rmvport.sh <Port Name> <Node Name>
```

To delete the VPort with a port name of 10000000c94ac63a and a node name of 20010000c94ac63a, type:

```
./rmvport.sh 10000000c94ac63a 20010000c94ac63a
```

This script may take up to 30 seconds to finish. The script fails if the VPort is not deleted.

3.5.2.3 Displaying VPorts Using the lsvport.sh Script

You can use the `lsvport.sh` script to display the VPorts and physical ports that are present on the system. Run the script with no parameters to display port information. For example:

```
./lsvport.sh
lpfc0: host6 10000000c93a5b5e:20000000c93a5b5e LP10000 NPIV Not Supported
lpfc1: host7 10000000c93a5b5d:20000000c93a5b5d LP10000 NPIV Not Supported
lpfc2: host8 10000000c93cc8dd:20000000c93cc8dd LPe12000 NPIV Physical
lpfc3: host9 10000000c93cc8dc:20000000c93cc8dc LPe12000 NPIV Physical
lpfc4: host10 10000000c94ac63a:20010000c94ac63a  NPIV Virtual (VPI 1)
```

In reference to the previous example:

- For LPFC0 and LPFC1, *NPIV Not Supported* indicates that this adapter/firmware combination does not support the creation of VPorts.
- For LPFC2 and LPFC3, *NPIV Physical* refers to a physical port of this adapter.
- For LPFC4, *NPIV Virtual* refers to a VPort of this adapter.

3.5.3 VPort sysfs

This section describes the VPort sysfs structure and VPort sysfs entries.

3.5.3.1 VPort sysfs Tree

When a VPort is created, three new directories are created in the class tree:

```
/sys/class/scsi_host/hostY/
```

```
/sys/class/fc_host/hostY/
/sys/class/fc_vports/vport-X:0-Z/-
```

Creating a new VPort also creates a new sysfs directory in the bus and devices tree:

```
/sys/bus/pci/drivers/lpfc/0000:A:B:C/hostX/vport-X:0-Z/hostY
/sys/devices/pci0000:A/0000:A:B:C/hostX/vport-X:0-Z/hostY
```

In both directories, there is a hostY directory that contains the remote ports that this new host can access:

```
/sys/bus/pci/drivers/lpfc/0000:A:B:C/hostX/vport-X:0-Z/hostY
/sys/bus/pci/drivers/lpfc/0000:A:B:C/hostX/vport-X:0-Z/hostY/rport-Y:0-0
/sys/bus/pci/drivers/lpfc/0000:A:B:C/hostX/vport-X:0-Z/hostY/rport-Y:0-1
/sys/bus/pci/drivers/lpfc/0000:A:B:C/hostX/vport-X:0-Z/hostY/rport-Y:0-2
```

In this example:

- X indicates the host value for the parent fc_host from which this VPort was created.
- Y indicates the new host value for the VPort that was created.
- Z indicates the instance of VPort created from the parent fc_host. A, B, and C indicate the PCI hierarchy for each physical FCoE port.

hostY is the new host created for the new VPort. *vport-X:0-Z* uniquely identifies the VPort and indicates the parent host structure (X) that created this VPort.

For example, if a VPort is created from host5, a new scsi_host, a new fc_host, a new fc_vport, and a new entry under the bus tree are created as well.

```
ls /sys/class/scsi_host/
host0  host1  host4  host5  host6
ls /sys/class/fc_host/
host4  host5  host6
ls /sys/class/fc_vports/
vport-5:0-0
```

3.5.3.2 VPort sysfs Entries

NOTE VPort sysfs entries shown in the following table are only present if the driver was loaded with `lpfc_enable_npiv` enabled.

Table 7 VPort sysfs Entries

VPort sysfs Entries	Type	Range/Input	Location and Description
`lpfc_peer_port_login`	Read/Write	0=Off (default) 1=On	`/sys/class/scsi_host/hostX/lpfc_peer_port_login` This entry sets the port's behavior when discovering targets in the SAN. The default behavior (value=0) will log in only to N_Ports that are physically located on a different port. The port will still attempt to log in to targets on all other ports (including the other port in a dual-port adapter). If this parameter is turned on (value=1), the port attempts to log in to all N_Ports, even if they are physically located on the same port. **NOTE** This parameter was created to reduce the amount of hardware resources (for example, RPIs) that the driver requires. In a configuration where many VPorts are on one physical port, this capability greatly reduces the number of RPIs that the driver uses.
`lpfc_restrict_login`	Read/Write	0=Off 1=On (default)	`/sys/class/scsi_host/hostX/lpfc_restrict_login (VPorts only)` This entry sets the VPort behavior when discovering targets in the SAN. The default behavior (value=1) prevents the VPort from logging into other initiator ports in the SAN. It also rejects logins from other ports in the SAN, because it assumes that all ports that send a PLOGI are initiators. If this sysfs entry is turned off, the driver attempts to log in to every port that it can access in the SAN, and it accepts logins from all ports. **NOTE** This parameter was created to reduce the amount of hardware resources (for example, RPIs) that the driver requires. In a SAN where there are other initiators, this capability greatly reduces the number of RPIs that the driver uses.
`max_npiv_vports`	Read-only	integers	`/sys/class/fc_host/hostX/max_npiv_vports` This entry displays the maximum number of VPorts that are supported by the underlying fc_host hardware. This sysfs entry exists only if the `vport_create` and `vport_delete` sysfs entries exist. If an fc_host does not support NPIV, this sysfs entry may not exist. Use this sysfs entry with the `npiv_vports_inuse` entry to determine whether the maximum number of VPorts have been created on this fc_host.

`node_name`	Read-only	16-byte hexadecimal value	For the FC RHEL 5.*x* driver:
			`/sys/class/fc_host/hostX/node_name`
			For FC RHEL 6.*x*/SLES 11 SPx drivers:
			`/sys/class/fc_host/hostX/node_name/sys/class/fc_vports/vport-X:0-Z/node_name`
			These entries display the physical or VPort node name. You assign this value when the VPort is created, and it is transmitted to the fabric upon fabric login.

`npiv_vports_inuse`	Read-only	integers	`/sys/class/fc_host/hostX/npiv_vports_inuse` This entry displays the number of VPorts that were created on this fc_host. This sysfs entry exists only if the `vport_create` and `vport_delete` sysfs entries exist. If an fc_host does not support NPIV, this sysfs entry might not exist. Use this sysfs entry with `max_npiv_vports` to determine whether the maximum number of VPorts have been created on this fc_host.

Table 7 VPort sysfs Entries (Continued)

VPort sysfs Entries	Type	Range/Input	Location and Description
`port_name`	Read-only	16-byte hexadecimal value	`/sys/class/fc_host/hostX/port_name/sys/class/fc_vports/vport-X:0-Z/port_name` This entry displays the physical or VPort's port name. You assign this value when the VPort is created, and it is transmitted to the fabric upon fabric login.
`vport_create`	Write-only	WWPN; WWNN	`/sys/class/fc_host/hostX/vport_create` This entry creates a VPort on the physical port on which hostX is located. The new VPort will have a WWPN and WWNN present on the fabric based on the WWPN and WWNN that are entered with this sysfs entry. This entry returns a 0 if the VPort creation was successful. A nonzero value indicates that the VPort was not created. If an fc_host does not support NPIV, this sysfs entry might not exist. **NOTE** It is possible for the VPort creation to succeed, but the VPort will be in a failed or an inoperative state. Use the new sysfs tree created by the new VPort to check the state of the new VPort.
`vport_delete`	Write-only	WWPN; WWNN	`/sys/class/fc_host/hostX/vport_delete` This entry deletes a VPort on the physical port on which hostX is located. The VPort matching the WWPN and WWNN is immediately deleted. This entry returns a 0 if the VPort deletion was successful. A nonzero value indicates that the VPort was not deleted. If an fc_host does not support NPIV, this sysfs entry may not exist. **NOTE** This entry deletes the VPort even if there are mounted file systems being accessed through this VPort, or if there are open files on it.

3.5.4 Monitoring VPorts with fc_vport (FCoE Drivers)

The transport creates an `fc_vports` directory that you can use to monitor VPorts. This directory is populated entirely of VPorts and has links from each to the fc_host associated with that VPort.

```
ls /sys/class/fc_vports/
vport-5:0-0
ls -d /sys/bus/pci/drivers/lpfc/*/host*/*/host*
/sys/bus/pci/drivers/lpfc/0000:03:06.1/host5/vport-5:0-0/host6
ls /sys/devices/pci*/*/host5/vport-5*/host6
power rport-6:0-0 rport-6:0-1 rport-6:0-2 uevent
ls /sys/devices/pci*/*/host5/vport-5*/host6/rport-*
/sys/devices/pci00:03/00:03:06.1/host5/vport-5:0-0/host6/rport-6:0-0:
power uevent

/sys/devices/pci00:03/00:03:06.1/host5/vport-5:0-0/host6/rport-6:0-1:
power uevent

/sys/devices/pci00:03/00:03:06.1/host5/vport-5:0-0/host6/rport-6:0-2:
power target6:0:0 uevent
```

In this example:

- A new entry exists in the `fc_vports` directory for the VPort (vport-5:0-0). The vport-5:0-0 entry indicates that the VPort was created from host5 and it is the first (0) VPort to be created on that fc_host.
- The new host for the VPort is host6, and it will appear in the usual directories.

- A new directory also exists in the bus tree. This new directory indicates that host6 was created under vport-5:0-0 (which was created from host5).

3.5.5 VPort Configuration Limits

VPort configuration limits are designated as enforced or unenforced. Enforced limits are limits that the driver enforces and prevents you from exceeding. Unenforced limits are limits that the driver cannot enforce, but configurations that exceed them are unsupported.

The following VPort configuration limits have been tested with and are supported by the Emulex driver. Configurations that exceed one or more of these limits are unsupported.

- Before the VPort is deleted or the driver is unloaded, I/O devices accessed through a VPort must be stopped and file systems must be unmounted.
- For enterprise-class adapters, the maximum number of VPorts configurable on a physical port is 64. The hardware allows more than 64 VPorts to be created, but the driver has only been qualified at 64. For mid-range adapters, the maximum number of VPorts configurable on a physical port is 16.
- The maximum number of LUNs supported on each driver port is 256.
- The maximum number of targets supported for each driver port is 255.
- The maximum number of driver ports in one zone is 64. This limit is based on the system's ability to recover from link events within the time constraints of the default timers.

The NPIV use-cases that involve a virtual server environment, include associating a VPort with a virtual machine, and placing the virtual machine in its own zone. This situation results in one VPort per zone. In the case of load-balanced environments, this situation can increase typically to two VPorts per virtual machine, to a practical limit of something far less than 50.

In the NPIV cases not related to virtual server environments, zoning is typically initiator-zoning, again resulting in one VPort, or a low number of VPorts in the case of load-balancing, within a given zone. If too many VPorts are within a single zone, expected behavior includes devices being lost after link events.

- The minimum lifetime of a VPort is 60 seconds. An unenforced limit of 60 seconds exists between the creation of a VPort and the deletion of the same VPort. VPorts exist for a long time in the system, and the creation of VPorts is asynchronous, which means that a VPort might not be finished with FC or SCSI discovery when the command to create a VPort is finished.

3.6 FCoE Driver Performance Tuning

This section describes how to tune the FC/FCoE driver for best performance.

3.6.1 Overview

The following configurable parameters can enhance performance:

- `lpfc_fcp_io_channel`
- `lpfc_fcp_io_sched`
- `lpfc_fcp_imax`

These features are available through module parameters that are defined in the FCoE driver, as well as sysfs entries defined by the Linux kernel.

In addition, you can use the `lpfc_vector_map.sh` script to map a specific I/O channel to a specific CPU. Determination of the mapping of the I/O channel to a specific CPU is also handled by the driver.

This section provides more information about how the tuning parameters and script can improve Emulex adapter performance.

3.6.1.1 lpfc_fcp_io_channel

The `lpfc_fcp_io_channel` module parameter can be configured at driver load time. It defines the number of I/O channels supported by the driver. The driver is capable of supporting parallel I/O paths, and each I/O path is capable of posting and completing FCP commands independent of the other.

Emulex adapters that are running in MSI-X interrupt mode can use more than one I/O path. Each I/O channel is composed of a unique MSI-X vector- event queue (EQ)/ completion queue (CQ)/WQ tuple. This parameter will override the value of the `lpfc_fcp_eq_channel` and the `lpfc_fcp_wq_count` parameters.

By default, the driver is configured for four I/O channels per port. The driver will also limit the number of I/O channels to not exceed the number of online *logical* CPUs (as reported by `/proc/cpuinfo`).

3.6.1.2 lpfc_fcp_io_sched

The `lpfc_fcp_io_sched` module parameter can be configured at driver load time. It can also be set dynamically as an sysfs entry. The driver uses the parameter to determine which algorithm to use when scheduling an FCP I/O to an I/O channel.

When multiple I/O channels are in use, I/Os can be scheduled to an I/O channel in a round-robin fashion, or by determining which CPU is running when the I/O is submitted.

The default value (0) configures the driver for round-robin scheduling. A value of 1 configures scheduling by CPU.

3.6.1.3 lpfc_fcp_imax

The `lpfc_fcp_imax` can be configured at driver load time. It can also be set dynamically as an sysfs entry. This parameter defines the maximum number of interrupts per second that each adapter port will support.

Considerations

- The lower the value set, the more completions are coalesced by the adapter, which causes the driver to handle multiple completions under the context of one interrupt. The higher the value, the faster an interrupt is generated for a completed command. Therefore, a balanced or *tuned* system must be found.
- A lower value equals higher interrupt latency; a higher value equals lower interrupt latency.
- Faster completions consume more system resources and CPU cycles, as the overhead of one interrupt completes fewer commands. The value is divided by the number of I/O channels, and each I/O channel is separately configured for its own interrupt latency.

By default, the module parameter is configured for 50,000 interrupts per second per adapter port. Older driver versions have a default value of 10,000.

3.6.1.4 lpfc_vector_map.sh

The `lpfc_vector_map.sh` script uses kernel sysfs entry points to map a specific I/O channel (MSI-X vector-EQ/CQ/WQ tuple) to a specific CPU. Run the script immediately after the driver is loaded.

This script resides in the `/usr/sbin/lpfc` directory. Run this script by adding the following line to the `/etc/modprobe.d/lpfc.conf` file and through `/etc/rc.d` for the initial boot:

```
install lpfc /sbin/modprobe --ignore-install lpfc;
/usr/sbin/lpfc/lpfc_vector_map.sh
```

The script maps each interrupt vector allocated by the driver to a CPU, thereby spreading the interrupt load of the ports across multiple CPUs. Each vector, with its associated I/O channel, is sequentially mapped to a CPU in a round-robin fashion. The number of vectors assigned to each adapter port is defined by the `lpfc_fcp_io_channel` module parameter.

I/O channels, which correspond to vectors, are typically mapped to unique CPUs to enhance the ability of the driver to run multiple FCP commands in parallel. In addition, running this script forces I/O scheduling to be by CPU (`lpfc_fcp_io_sched = 1`), which increases performance when a specific I/O channel is mapped to a specific CPU.

The script has two modes of operation: Driver mode and host bus adapter (HBA) mode. By default, the script runs in Driver mode.

Driver Mode

Driver mode maps all vectors for all driver ports, starting with CPU0, sequentially assigning a new CPU for each vector belonging to the entire driver. If there are more interrupt vectors than CPUs, the vector assignment wraps back to CPU0 as needed.

HBA Mode

HBA mode maps all vectors for each specific adapter port, starting with CPU0, sequentially assigning a new CPU for each vector belonging to that adapter port.

3.7 Network Performance Tuning

This section describes optimizing network performance.

3.7.1 Memory Bandwidth Considerations

The availability of higher memory bandwidth leads to better network performance. The following sections describe how memory bandwidth can be increased.

3.7.1.1 Enabling Optimal Bandwidth Options

Most computers offer multiple distinct memory channels, or memory interleaves, that might not be enabled by default. Check the manufacturer's documentation and BIOS parameters for details on enabling optimal memory bandwidth options.

3.7.1.2 Populating DIMM Slots

Typically, all the dual in-line memory module (DIMM) slots must be populated to make use of all the memory channels. As a general rule, using more DIMMs provides better performance by allowing a higher degree of memory-access interleaving to occur.

3.7.1.3 Disabling Memory Mirroring

Some servers may allow memory mirroring, where the total memory is divided in half and each location is stored twice. Memory mirroring allows fault recovery if one memory location detects an error, but it greatly reduces the perceived memory bandwidth of the system. Consider disabling memory mirroring if it is not needed.

3.7.1.4 Using a Fast Clock Speed for the FSB

Nearly any desktop or low-end server has enough memory bandwidth for Emulex adapters to support DMA at 20Gb/s of data (10Gb/s read, 10Gb/s write). However, most of the memory demands come from the processor accessing the data for either packet copies in the nonoffloaded networking stack or application accesses. All processor memory accesses use the front side bus (FSB). The clock speed of this bus is critical for allowing efficient memory bandwidth. A system with a faster processor FSB clock speed performs better than a system with a slower FSB clock speed.

3.7.2 Network Memory Limits

The default values of tunable parameters in the Linux network stack are optimal for most network applications involving several TCP/User Datagram Protocol (UDP) streams. The optimal size for the network queues and buffers depends on several factors, such as protocol, number of streams (connections), request size, and application behavior.

The following network configuration settings are a good combination to get the best unidirectional transmit and receive performance with six or more TCP connections/UDP streams:

```
echo 4096 87380 4194304 > /proc/sys/net/ipv4/tcp_rmem
echo 4096 16384 4194304 > /proc/sys/net/ipv4/tcp_wmem
echo 64000000 > /proc/sys/net/core/rmem_default
echo 64000000 > /proc/sys/net/core/rmem_max
echo 32000000 > /proc/sys/net/core/wmem_default
echo 32000000 > /proc/sys/net/core/wmem_max
```

These settings assume ideal conditions, such as low latency, zero (or close-to-zero) packet loss in the network, enough free memory, and 10Gb/s path-to-peer system.

These `tcp_rmem` and `tcp_wmem` values are also the default values in recent RHEL 5 distributions. If your application requires best throughput with very small number of connections (less than four), it may help to increase the `tcp_rmem` and `tcp_wmem` to much larger values:

```
echo 4096 87380 16777216 > /proc/sys/net/ipv4/tcp_rmem
echo 4096 65536 16777216 > /proc/sys/net/ipv4/tcp_wmem
```

3.7.3 TCP Segmentation Offload (TSO)

In low-loss networks, TSO considerably improves performance and, therefore, must be enabled. TSO is enabled by default in the OneConnect network driver.

The `/proc/sys/net/ipv4/tcp_tso_win_divisor` process variable controls how aggressive the network stack can be in making TSO requests. For low-loss networks, the TSO divisor values must be in the range of 2 to 16. In most distributions, the default value of 3 may be the optimal choice for a no-loss network.

Smaller divisor values result in larger TSO chunks and better throughput, as well as better CPU utilization. However, if the receiver or the network is dropping frames (too many retransmits on the transmit side as indicated by `netstat -st`), it might help to make TSO chunks smaller (by increasing the divisor value) or to turn TSO off. For example, to set the divisor level to a value of 8, run:

```
echo 8 > /proc/sys/net/ipv4/tcp_tso_win_divisor
```

To turn TSO on or off, run one of the following `ethtool` commands:

```
ethtool -K eth<N> tso on
ethtool -K eth<N> tso off
```

where `eth`<*N*> is the name of the Ethernet device on which you are working (for example, eth0).

3.7.4 Flow Control

Emulex adapters support Institute of Electrical and Electronics Engineers (IEEE) 802.3x standard flow control, which uses control packets to temporarily pause the transmission of packets between two endpoints. These control messages are point-to-point and are not forwarded by switches or routers. The adapter can respond to flow control packets by temporarily pausing transmits. The adapter can send flow control pause packets when the transmitter is overwhelming the system's receive bandwidth.

Flow control can greatly improve performance, as described in the following examples:

- The adapter is installed in a 4x PCIe slot or an underpowered server system.

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

- The adapter transmits to 1Gb/s devices, especially when using a non-TCP protocol.

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

You can configure the adapter to respond to flow control pause frames from the other side (switch or router) using the following `ethtool` commands:

```
ethtool -A eth<N> pause rx on
ethtool -A eth<N> pause rx off
```

where `eth<N>` is the name of the Ethernet device on which you are working (for example, eth0).

You can configure the adapter to send flow control pause frames using the following ethtool commands:

```
ethtool -A eth<N> pause tx on
ethtool -A eth<N> pause tx off
```

where `eth<N>` is the name of the Ethernet device on which you are working (for example, eth0).

RX and TX flow control are enabled by default. When priority flow control is enabled in the adapter, normal flow control cannot be enabled.

Refer to the switch or router documentation to determine how link level flow control can be configured on the switch or router to which the port is connected.

NOTE In multichannel configurations where multiple PCI functions are exposed for a single 10GbE port, the flow control parameter for a port can be configured through any interface associated with the physical port, and the configured property will apply to all interfaces associated with the port.

3.7.5 RX Frame Coalescing

The Ethernet driver coalesces regular-sized TCP segments to a large frame before passing it to the network stack, which may improve TCP receive performance. RX frame coalescing is implemented using the Generic Receive Offload (GRO) mechanism (in Linux driver versions that support GRO) or the LRO mechanism (in older Linux driver versions).

RX frame coalescing is enabled by default. In some configurations where the end point for the TCP connection to which the packets belong is not in the current server (for example, the end point is a router), RX coalescing should not be enabled.

GRO can be disabled using the –K option with the ethtool command:

```
ethtool –K eth<N> gro off
```

LRO can be disabled using the –C option with the ethtool command:

```
ethtool -C eth<N> rx-frames 1
```

where `eth<N>` is the name of the Ethernet device you are working on (for example, eth0).

3.7.6 Maximum Transmission Unit (MTU)

The Ethernet driver supports MTUs between 256 bytes and 9000 bytes. The default MTU is set to 1500. If other elements in the network path support a larger MTU, you can increase the MTU up to 9000 using the `ifconfig` command. To do this, run:

```
ifconfig eth<N> mtu 9000
```

where `eth<N>` is the name of the Ethernet device on which you are working (for example, eth0).

3.7.7 Interrupt Coalescing

The Ethernet driver tries to reduce the number of interrupts by delaying the interrupts from the adapter or CNA. This action reduces CPU utilization during a high traffic rate. The interrupt delay duration can be set to change dynamically within a range of values, depending on the receive rate (known as Adaptive Interrupt Coalescing [AIC]), or can be set to a constant value.

3.7.7.1 Setting the Interrupt Delay Duration to a Range of Values (AIC)

For receive interrupts, AIC is enabled by default. When AIC is enabled, the default low limit is 0 microseconds, and the default high limit is 96 microseconds. In low traffic, the interrupt delay is set to 0 for best latency. As the number of interrupts per second increases, the delay is increased to higher values proportional to the receive rate, up to the default high limit of 96 microseconds. You can change the low and high limits using ethtool. For example, to set a low limit of 8 and a high limit of 40, run

```
ethtool -C eth<N> rx-usecs-low 8
ethtool -C eth<N> rx-usecs-high 40
```

where `eth<N>` is the name of the Ethernet device on which you are working (for example, eth0).

For transmit interrupts, AIC is not supported.

To disable AIC and set the interrupt delay duration to a constant value, see the following section.

3.7.7.2 Setting the Interrupt Delay Duration to a Constant Value

The interrupt delay duration can be set to a constant value for both receive and transmit interrupts. The possible interrupt delay duration values are 0 to 96 microseconds, in 8-microsecond increments.

For receive interrupts, disable AIC (because it is enabled by default) and set the interrupt delay duration using `ethtool`. For example, to disable AIC and set the constant RX interrupt delay to 8-microseconds, run

```
ethtool -C eth<N> adaptive-rx off rx-usec 8
```

where `eth<N>` is the number of the Ethernet interface on which you are working.

If your application requires low or predictive latency, turn off AIC and set `rx-usecs` to 0.

For transmit interrupts, the default interrupt delay duration is 96 microseconds. You can change this value using ethtool. For example, to set the transmit interrupt delay to 64 microseconds run

```
ethtool -C eth<N> tx-usec 64
```

where `eth<N>` is the number of the Ethernet interface on which you are working.

3.7.8 Receive-Side Scaling (RSS)

Distributing the incoming traffic across several receive rings with separate interrupt vectors helps to distribute the receive processing across several CPU cores. This could reduce the packet drop and improve the packet rate in certain

applications. RSS is enabled in non-SR-IOV and non-multichannel configurations. In multichannel configurations, RSS is enabled in the first section of each port.

3.7.9 Analyzing Performance Issues

MSI-x interrupts are required for RSS to work. If your motherboard and operating system version support MSI-X, the Ethernet driver automatically uses MSI-X interrupts. If not enough MSI-X vectors are available, the Ethernet driver uses INTx interrupts, which might decrease performance. The `/proc/interrupts` proc node shows the interrupts and their types.

The Linux performance *top* utility can monitor the CPU utilization while troubleshooting performance issues. A low idle CPU percentage in any CPU core is an indication of excessive processing load for that CPU. The `/proc/interrupts` proc node shows the distribution of the interrupts across the CPU cores. If you see too many interrupts per second directed to one CPU, check to see if the `irqbalance` program is running. The `irqbalance` program is normally started at system boot. In some cases, you can get better performance by disabling irqbalance and manually distributing interrupts. You can manually distribute the interrupt load across the available CPU cores by setting the CPU affinity for any interrupt vector by setting the mask in the `/proc/irq/<int-vector>/smp_affinity` proc node.

Use the `netstat` command to look for excessive TCP retransmits or packet drops in the network stack.

In systems having more than one non-uniform memory access (NUMA) node, you can get better performance by pinning interrupts to the NUMA node local to the PCIe device.

Use the –S option of `ethtool` to see all statistics counters maintained by the Ethernet and driver. Excessive drop or error counters are an indication of a bad link or defective hardware. See Table 19, Ethtool -S Option Statistics, and Table 20, Transmit/Receive Queue Statistics.

Turning off auditing and SELinux can improve CPU utilization and, in some cases increase throughput. You can disable auditing by appending `audit=0` in the boot command line. You can turn off SELinux by specifying `selinux=0` in the boot command line. For example, the following command boots the Linux kernel with the SELinux and auditing options disabled:

```
kernel /boot/vmlinux-2.6.18 ro root=/dev/md0 selinux=0 audit=0
```

You can get better CPU utilization, and in some cases better throughput, by disabling kernel debug options, such as `CONFIG_DEBUG_SLAB`. This situation requires you to build the kernel image and modules. Turning off the firewall and disabling Hyper-Threading can also improve performance.

Chapter 4: Troubleshooting

This section explains some of the situations in which your system can operate in an unexpected manner and some possible resolutions.

4.1 Hardware Situations and Resolutions

Table 8 lists the hardware situations and their resolutions.

Table 8 Hardware Situations and their Resolutions

Situation	Resolution
An unapproved optical transceiver is used to connect the adapter.	Unapproved optical transceivers include: ■ 10GBASE-SR ■ 10GBASE-LR ■ 40GBASE-SR The system log can generate one or more of these events if an unapproved optical transceiver is detected: ■ Unsupported module ■ Optics faulted, incorrectly installed, or not installed ■ Incompatible optics ■ Unknown port SFB status To resolve this issue, do the following: 1. Power the system off. 2. Replace the unapproved optical transceiver with an approved one. 3. Power the system on. **NOTE** For more information on replacing an optical transceiver, refer to the adapter hardware guide available at the Documents and Downloads area of http://www.broadcom.com.

4.2 FCoE Driver Situations and their Resolutions

Table 9 lists the FCoE driver situations and their resolutions.

Table 9 FCoE Driver Situations and their Resolutions

Situation	Resolution

"Authentication is enabled but authentication service is not running." error message	If you see this message in `/var/log/messages` file and the adapter is in an error state, the `fcauthd` daemon probably is not running. To determine whether `fcauthd` is running, run:
	`/etc/init.d/fcauthd status`
	To start `fcauthd`, run:
	`/etc/init.d/fcauthd start`

If a SAN configuration has 256 targets mapped by the 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 re-initializing the link does not solve this issue. Unload and reload the driver to reset available target IDs. Ensure that the SAN configuration is correct prior to rebooting the driver. This clears the driver's consistent binding table and frees target IDs for new target nodes.
rmmod fails to unload FCoE driver module due to "ERROR: Module lpfc is in use."	This message can appear when you attempt to remove the driver and a Logical Volume Group is dependent on the driver. If you have configured boot from a SAN, you must reboot the system. Otherwise, use these steps to resolved this situation: 1. Make the Logical Volume Group unavailable. Type: `lvchange -a n xxxxxxx` The *xxxxxxx* parameter is the Volume Group Name. 2. Stop the Emulex OneCommand CNA Manager application. 3. Stop Device Mapper.
rmmod hangs and module reference count is 0.	Due to a small race condition in the kernel, it is possible for an `rmmod` command to hang. Issue the `rmmod -w` command. If this command does not help, reboot the computer.

rmmod fails to unload driver due to device or resource busy.	This message occurs when you attempt to remove the driver without first stopping the Emulex OneCommand CNA Manager application or the `fcauthd` daemon when the Emulex OneCommand CNA Manager application is installed and running. To resolve this situation:
	1. Stop the Emulex OneCommand CNA Manager application before attempting to unload the driver. The script is located in the `/usr/sbin/ocmanager` directory. Type:
	`./stop_ocmanager`
	2. Unmount any disks connected to the adapter.
	3. Unload the driver. Type:
	`rmmod lpfc`

An lspci shows recent Emulex adapters as unknown.	This situation occurs because of the delay of getting new product IDs into the Linux development cycle. No resolution exists at this time.
Slow targets or extended link faults on the storage side might result in storage being marked offline by the mid-level and remaining offline (not recovered) when the link faults are corrected.	This version of the driver should eliminate this issue. However, if you experience offline device issues, increase the SCSI command timeout to a value greater than or equal to 60 seconds. Emulex also provides a script that addresses this issue. To access the `lun_change_state.sh` script, go to the Documents and Downloads area of http://www.broadcom.com.

Table 9 FCoE Driver Situations and their Resolutions (Continued)

Situation	Resolution
Under certain conditions of an I/O load, some targets cannot complete an I/O issued by a Linux initiator within the default timeout of 30 seconds given by the SCSI mid-level.	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 I/O error of hex 6000000. Emulex provides a script that addresses this issue. To access the `set_target_timeout.sh` script, go to the Documents and Downloads area of http://www.broadcom.com.
The 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. Install the latest firmware on the adapter. **NOTE** Before performing a firmware update, driver installation is required. For more information on installing the driver, see Section 2.2.1, Installing the Binary RPM FCoE Driver Kit.
Loading the FCoE driver on SLES 11 SPx reports "unsupported module, tainting kernel" in the system log.	This message is logged by the kernel whenever a module that is not shipped with the kernel is loaded. Ignore this message.
The system panics when it is booted with a failed adapter installed.	Remove the failed adapter, and reboot the system.
Unloading the FCoE driver on SLES 11 SPx might cause a message to be logged in the system log, such as the following: `umount: /dev/disk/bypath/pci-0000:02:04.0-scsi-0:0:1:0: not mounted`	These messages are normal output from the SLES 11 SPx hot plug scripts and can be ignored.
Driver installation fails.	The `brcmfcoe_install` script fails to install the driver. The installation script might fail for the following reasons: ■ A previous version of the driver is installed. Run the `brcmfcoe_install` --uninstallation script and then try to install the driver. ■ The current driver is already installed. Ensure that you are running a supported RHEL or SLES kernel.
"No module lpfc found for kernel KERNELVERSION" RPM error message appears when upgrading the kernel. A recently upgraded kernel cannot find the ramdisk. After upgrading the kernel, the kernel cannot find the ramdisk, which halts or panics the system. The driver is not loaded after a system reboot after upgrading the kernel.	These three situations can be resolved by upgrading the kernel. There are two ways to install the driver into an upgraded kernel. The method you use depends on whether you are updating the driver. ■ Upgrade the kernel using the same version of the driver. ■ Upgrade the kernel using a new version of the driver. For more information on upgrading the kernel, see Chapter 2, Installing and Uninstalling.
Driver uninstall fails.	The `brcmfcoe_install` --uninstallation script fails with an error. Try the following solutions: ■ Uninstall the Emulex OneCommand CNA Manager application. For instructions, refer to the *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide*. ■ Unmount all FC disk drives. ■ Unload the lpfcdfc and FCoE driver. ■ Use rpm `-e brcmfcoedriver` and `-e ocmanager` and uninstall the new kits.

Table 9 FCoE Driver Situations and their Resolutions (Continued)

Situation	Resolution
`brcmfcoe_install` script exit code.	The `brcmfcoe_install` script contains exit codes that are useful in diagnosing installation issues. Refer to the `brcmfcoe_install` script for a complete listing of codes and definitions.

The Emulex driver for Linux does not load in ramdisk for a custom-built kernel.	Broadcom® does not support custom built kernels. However, the Emulex installation script attempts to install the driver into a ramdisk that follows the naming scheme used by RHEL or SLES kernels.
	■ The RHEL Hat naming scheme for ramdisk images is:
	`/boot/initrd-KERNELVERSION.img.`
	■ The SLES naming scheme for ramdisk images is:
	`/boot/initrd.`
	If a custom-built kernel has a ramdisk image that does not follow the appropriate naming scheme, change the name of the image using the following procedure:
	1. Change the name of the ramdisk image to match the SLES naming scheme.
	2. Update any file links to the ramdisk image.
	3. Edit the boot loader configuration file (for example, `/etc/lilo.conf`, `/etc/yaboot.conf`, `/boot/grub/grub.conf`, `/boot/grub/menu.lst`), find any references to the old ramdisk image name, and replace them with the new name.
	4. Reboot the system to verify the changes.
	5. Install the brcmfcoe driver kit.

The Linux SCSI subsystem sees only eight LUNs when more are present.	Some SCSI drivers do not scan past eight LUNs when the target reports itself as a SCSI-2 device. To resolve this situation, force a SCSI bus scan with the following command: `/usr/sbin/brcmlpfc/lun_scan.` SUSE supplies a `/bin/rescan-scsi-bus.sh` script, which can be changed to scan everything.

4.3 Ethernet Driver Situations and their Resolutions

The following table lists the Ethernet driver situations and their resolutions.

Table 10 Ethernet Driver Situations and their Resolutions

Situation	Resolution
The ethtool configuration settings are not restored after system reboot.	The ethtool settings do not persist across reboot. For persistence, invoke configuration commands from a boot script that is executed at system start, such as `/etc/rc.local`.
The Ethernet driver works, but the transmit and receive data rates are not near a 10Gb/s line rate.	Several reasons exist for poor performance. For best performance practices, see Section 3.7, Network Performance Tuning.

When MILI and Simple Network Management Protocol (SNMP) daemons start, they trigger warning messages within SELinux for certain operations.	This issue is known, and no solution is available. However, to avoid SElinux warning messages, you can disable SELinux. To disable SELinux, open a terminal, and enter the following command at the prompt:
	`echo 0 > /selinux/enforce`
	To enable SELinux, use the following command:
	`echo 1 > /selinux/enforce`
	In addition, see Section 3.7.9, Analyzing Performance Issues.

4.4 iSCSI Driver Situations and their Resolutions

The following table lists the iSCSI driver situations and their resolutions for OneConnect CNAs.

Table 11 iSCSI Driver Situations and their Resolutions

Situation	Resolution
When you log out of a target while an I/O is running and you log into the target again, you will get an error trace in `/var/log/messages` beginning with the error message, "trying to free buffer."	No solution is available.

4.5 Log Messages

4.5.1 FCoE Driver Log Messages

The following section describes retrieving and interpreting FC and FCoE driver log messages.

4.5.1.1 Retrieving FC and FCoE Driver Log Messages

FCoE error log messages are logged in the `/var/log/messages` file.

An example of an FCoE message:

```
Jul 2 04:23:34 daffy kernel: lpfc 0000:03:06.0: 0:1305 Link Down Event x2f2
received Data: x2f2 x20 x110
```

In this example:

- `lpfc 0000:03:06.0` – Identifies the PCI location of the particular LPFC hardware port.
- `0:` – Indicates Emulex adapter 0
- `1305` – Indicates a log message number of 1305.

NOTE

- If `Data:` is present in a log message, any information following `Data:` is intended only for Broadcom Technical Support and Engineering use.
- If an error message instructs you to perform a firmware update, ensure that the driver is installed first. For more information on installing the driver, see Chapter 2, Installing and Uninstalling.

4.5.1.2 FCoE Error Log Messages and their Descriptions

The following table lists the FCoE error log messages and their descriptions.

Table 12 FCoE Error Log Messages and their Descriptions

`0111: Dropping received ELS cmd` The driver decided to drop an ELS Response ring entry. Data: (1) ulpStatus (2) ulpWord[4] (3) ulpTimeout Severity: Error Log: Always Action: This error could indicate a software driver or firmware issue. If this issue persists, report these errors to Broadcom Technical Support.
`0113: An FLOGI ELS command <elsCmd> was received from DID <did> in Loop Mode` While in Loop Mode, an unknown or unsupported ELS command was received. Data: None Severity: Error Log: Always Action: Check the device DID.
`0115: Unknown ELS command <elsCmd> received from NPORT <did>` Received an unsupported ELS command from a remote N_Port. Data: None Severity: Error Log: Always Action: Check the remote N_Port for a potential issue.
`0125: FDISC Failed (x%x). Fabric out of resources` The fabric rejected an FDISC because the switch cannot support additional virtual ports. Data: lsRjtError Severity: Error Log: Always Action: Reconfigure the switch to support more NPIV logins. If this issue persists, contact Broadcom Technical Support.
`0126: FDISC failed ulpStatus ulpWord4` Data: lsRjtError Severity: Error Log: Always Action: Reconfigure the switch to support more NPIV logins. If this issue persists, contact Broadcom Technical Support.
`0127: ELS timeout` An ELS IOCB command was posted to a ring and did not complete within ULP timeout seconds. Data: (1) elscmd (2) remote_id (3) ulpcommand (4) ulpIoTag Severity: Error Log: Always Action: If no ELS command is going through the adapter, reboot the system. If the issue persists, contact Broadcom Technical Support.
`0133: PLOGI: no memory for reg_login` A memory allocation error occurred. Data: (1) nlp_DID (2) nlp_state (3) nlp_flag (4) nlp_rpi Severity: Error Log: LOG_ELS Action: Memory allocation error. Check system resources. Unload unused modules.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

`0134: PLOGI cannot issue reg_login` The ELS PLOGI mailbox command has failed. Data: (1) nlp_DID (2) nlp_state (3) nlp_flag (4) nlp_rpi Severity: Error Log: LOG_ELS Action: Check the port and switch configuration.
`0135: cannot format reg_login` Could not allocate an remote port indicator (RPI) or DMA buffer for the mailbox command. Data: (1) nlp_DID (2) nlp_state (3) nlp_flag (4) nlp_rpi Severity: Error Log: LOG_ELS Action: None required.
`0136: PLOGI completes to NPort <DID> completion` A PLOGI has completed for which there is no NDLP. Data: (1) ulpStatus (2) ulpWord[4] Severity: Error Log: LOG_ELS Action: None required.
`0137: No retry ELS command <ELS_CMD> to remote` Data: (1) ulpStatus (2) ulpWord[4] Severity: Error Log: LOG_ELS Action: None required.
`0138: ELS rsp: Cannot issue reg_login for <DID>` REG_LOGIN mailbox command failed. Data: (1) nlp_DID (2) nlp_state (3) nlp_flag (4) nlp_rpi Severity: Error Log: LOG_ELS Action: None required.
`0139: Ignoring ELS cmd tag <ioTag> completion Data` This ELS command was aborted. Data: (1) ulpStatus (2) ulpWord[4] (3) ulpTimeout Severity: Error Log: LOG_ELS Action: None required.
0140: PLOGI Reject: invalid name Invalid node WWN was provided. Data: None Severity: Error Log: LOG_ELS Action: None required.
0141: PLOGI Reject: invalid pname Invalid port WWN provided. Data: None Severity: Error Log: LOG_ELS Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0142: PLOGI RSP: Invalid WWN The PLOGI sent to the port by a remote port had an invalid WWN. Data: None Severity: Error Log: LOG_ELS Action: None required.
0143: SLI4 Adapter Hardware Error Data: <status0>/<status1> The HBA has encountered an unrecoverable error. Data: None Severity: Error Log: LOG_INIT Action: Use hbacmd to retrieve a dump file.
0144: Not a valid WCQE code: <Completion Code> The completion queue handler detected an invalid type. Data: None Severity: Error Log: LOG_SLI Action: None required.
0147: Failed to allocate memory for RSCN event Memory could not be allocated to send the RSCN event to the management application. Data: None Severity: Error Log: LOG_ELS Action: None required.
0148: Failed to allocate memory for LOGO event Memory could not be allocated to send the LOGO event to the FC transport. Data: None Severity: Error Log: LOG_ELS Action: None required.
0149: Failed to allocate memory for ELS event Memory could not be allocated to send the ELS event to the FC transport. Data: None Severity: Error Log: LOG_ELS Action: None required.
0154: Authentication not complete Authentication was restarted because the previous authentication did not complete. Data: None Severity: Error Log: LOG_DISCOVERY Action: Check the switch configuration.
0200: CONFIG_LINK bad hba state <hba_state> A CONFIG_LINK mailbox command completed, and the driver was not in the right state. Data: None Severity: Error Log: Always Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0203: Devloss timeout on WWPN <address> NPort <nlp_DID> A remote N_Port that was discovered by the driver disappeared for more than `lpfc_devloss_tmo` seconds. Data: (1) nlp_flag (2) nlp_state (3) nlp_rpi Severity: Error Log: Always Action: If the device generating this message is not a target to which the HBA is connected, this error will not affect the data integrity of the I/O between the HBA and the attached storage and can be ignored.
0206: Device discovery completion error This message indicates that an uncorrectable error was encountered during device discovery after a link up. FC devices will not be accessible if this message is displayed. Data: None Severity: Error Log: Always Action: Reboot the system. If this issue persists, report the error to Broadcom Technical Support. Run with verbose mode enabled for more information.
0207: Device <DID> (<WWN>) sent invalid service parameters. Ignoring device. Invalid service parameters were received from DID. Ignoring this remote port. Data: DID, WWN Severity: Error Log: Always Action: Verify the remote port's configuration. If this issue persists, report the error to Broadcom Technical Support. Run with verbose mode on for more details.
0217: Block sgl registration required DMAsize <reqlen> great than a page The request to post SGL pages does not fit on a page. Data: None Severity: Warning Log: LOG_INIT Action: None required.
0221: FAN timeout A linkup event was received without the login bit set, so the driver waits E_D_TOV for the fabric to send a FAN. If no FAN is received, a FLOGI will be sent after the timeout. Data: None Severity: Warning Log: LOG_DISCOVERY verbose Action: None required. The driver recovers from this condition by issuing a FLOGI to the fabric.
0222: Initial FLOG/FDISKI timeout The driver sent the initial FLOGI or FDISK to the fabric and never received a response. Data: None Severity: Error Log: Always Action: Check the fabric configuration. The driver recovers from this situation and continues with device discovery.
0223: Timeout while waiting for NameServer login A login request to the NameServer was not acknowledged within R_A_TOV. Data: None Severity: Error Log: Always Action: Check the fabric configuration. The driver recovers from this and continues with device discovery.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0224: NameServer Query timeout Node authentication timeout or node Discovery timeout occurred. A NameServer Query to the fabric or discovery of reported remote N_Ports is not acknowledged within R_A_TOV. Data: (1) fc_ns_retry (2) fc_max_ns_retry Severity: Error Log: Always Action: Check fabric configuration. The driver recovers from this and continues with device discovery.
0226: Device discovery completion error This indicates that an uncorrectable error was encountered during device discovery after a link up. FC devices will not be accessible if this message is displayed. Data: None Severity: Error Log: Always Action: Reboot the system. If this issue persists, report the error to Broadcom Technical Support. Run with verbose mode on for more details.
0227: Node Authentication timeout The driver has lost track of what N_Ports are being authenticated. Data: None Severity: Error Log: Always Action: None required. The driver should recover from this event.
0228: CLEAR LA timeout The driver issued a CLEAR_LA that never completed. Data: None Severity: Error Log: Always Action: None required. The driver should recover from this event.
0230: Unexpected timeout, hba linkstate <link_state> Discovery has timed out, and the HBA state is not ready. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.
0231: RSCN timeout The driver has lost track of which N_Ports have RSCNs pending. Data: (1) fc_ns_retry (2) lpfc_max_ns_retry Severity: Error Log: Always Action: None required. The driver should recover from this event.
0233: Nodelist not empty Driver unloaded or hot plug detected a node still in use. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0237: Pending Link Event during Discovery: State <hba_state> Received link event during discovery. Causes discovery restart. Data: None Severity: Warning Log: LOG_DISCOVERY verbose Action: None required, unless this issue persists. If persistent, check cabling.
0241: NameServer rsp error The driver received a NameServer response containing a status error. Data: (1) CommandResponse.bits.CmdRsp (2) ReasonCode (3) Explanation (4) fc_flag Severity: Error Log: LOG_DISCOVERY verbose Action: Check the fabric configuration. The driver recovers from this and continues with device discovery.
0246: RegLogin failed The firmware returned a failure for the specified RegLogin. Data: (1) Did (2) mbxStatus (3) hbaState Severity: Error Log: Always Action: This message indicates that the firmware could not perform a RegLogin for the specified DID. A limitation may exist on how many nodes an HBA can see.
0249: Cannot issue Register Fabric login: Err %d\ Could not issue the fabric reg login, the error value is unique for each possible failure. Data: None Severity: Error Log: LOG_ELS Action: None required.
0251: NameServer login: no memory Could not allocate memory for the NDLP structure. Data: None Severity: Error Log: LOG_ELS Action: None required.
0252: Cannot issue NameServer login Could not issue an ELS PLOGI to the NameServer DID. Data: None Severity: Error Log: LOG_ELS Action: Check the port connection and the switch configuration.
0253: Register VPI: Can't send mbox\ Could not issue the REG_VPI mailbox command for this VPort. Data: None Severity: Error Log: LOG_MBOX Action: None required.
0254: Register VPI: no memory" goto mbox_err_exit Could not allocate memory for the REG_VPI mailbox command. Data: None Severity: Error Log: LOG_MBOX Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0255: Issue FDISC: no IOCB All of the pre-allocated IOCBs are in use. Data: None Severity: Error Log: LOG_ELS Action: None required.
0256: Issue FDISC: Cannot send IOCB\ Unable to send the fabric IOCB. Data: None Severity: Error Log: LOG_ELS Action: None required.
0257: GID_FT Query error The GID_FT CT request for the NameServer has failed. Data: None Severity: Error Log: LOG_ELS Action: Check the switch configuration.
0258: Register Fabric login error: The REG_LOGIN for the fabric has failed. Data: None Severity: Error Log: LOG_MBOX Action: Check the port connection and the switch configuration.
0259: No NPIV Fabric support The switch to which the port is connected does not support NPIV. Data: None Severity: Error Log: LOG_ELS Action: Check the switch configuration.
0260: Register NameServer error: The REG_LOGIN mailbox command has failed for the NameServer. Data: None Severity: Error Log: LOG_ELS Action: Check the switch configuration
0261: Cannot register NameServer login: Either a memory allocation issue occurred or an invalid parameter was sent to the REG_LOGIN. Data: None Severity: Error Log: LOG_ELS Action: At least one message (0142 0121 0133 0134 0135) should precede this message.
0262: No NPIV Fabric support The switch to which the port is connected does not support NPIV. Data: None Severity: Error Log: LOG_ELS Action: Check the switch configuration.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0263: Discovery Mailbox error: state: Either the driver could not allocate resources or it could not send sparam_mbox or cfglink_mbox. Data: (1) address of sparam_mbox command (2) address of cfglink_mbox command Severity: Error Log: LOG_MBOX Action: Attempt to unload and reload the driver when it is convenient.
0264: No NPIV Fabric support The switch to which the port is connected does not support NPIV. Data: None Severity: Error Log: LOG_ELS Action: Check the switch configuration.
0266: Issue NameServer Req <cmdcode> err <rc> Data: <fc_flag> <fc_rscn_id_cnt> The driver was unable to send the NameServer CT command. Data: (1) vports fc_flag (2) vports fc_rscn_id_cnt Severity: Error Log: LOG_DISCOVERY Action: Check the port and switch configurations.
0267: NameServer GFF Rsp <did> Error (<ulpStatus> <un.ulpWord[4]>) Data: <fc_flag> <fc_rscn_id_cnt> The NameServer GFF CT request failed. Data: (1) vports fc_flag (2) vports fc_rscn_id_cnt Severity: Error Log: LOG_DISCOVERY Action: Check the port and switch configurations.
0268: NS cmd <cmdcode> Error (<ulpStatus> <un.ulpWord[4]>) The NameServer CT request failed. Data: None Severity: Error Log: LOG_DISCOVERY Action: Check the port and switch configurations.
0271: Illegal State Transition: node <nlp_DID> event <evt>, state <nlp_state> Data: <nlp_rpi> <nlp_flag> The current node state does not have a handler for this event. Data: (1) nlp_rpi (2) nlp_flag Severity: Error Log: LOG_DISCOVERY Action: Verify that all targets are still visible to the SCSI mid-layer.
0272: Illegal State Transition: node <nlp_DID> event <evt>, state <nlp_state> Data: <nlp_rpi> <nlp_flag> The driver is completing a PLOGI but does not have the rcv_plogi flag set. Data: (1) nlp_rpi (2) nlp_flag Severity: Error Log: LOG_DISCOVERY Action: Verify that all targets are still visible to the SCSI mid-layer.
0273: Unexpected discovery timeout, vport State x%x The discovery process has timed out. Data: None Severity: Error Log: LOG_DISCOVERY Action: Verify that all targets are still visible.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0274: lpfc_nlp_put: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount) Data: None Severity: Warning Log: LOG_NODE Action: None required.
0275: lpfc_nlp_put: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount) A kref_put was called again after the node was already inactive. Data: None Severity: Warning Log: LOG_NODE Action: None required.
0276: lpfc_nlp_get: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount) A kref_get was attempted on a node that was being released. Data: None Severity: Warning Log: LOG_NODE Action: None required.
0277: lpfc_enable_node: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount) An enable node was attempted on an inactive node. Data: None Severity: Warning Log: LOG_NODE Action: None required.
0278: lpfc_enable_node: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount) An enable node was attempted on an inactive node. Data: None Severity: Warning Log: LOG_NODE Action: None required.
0280: lpfc_cleanup_node: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount) Node cleanup was attempted on a node that has already been marked for memory free. Data: None Severity: Warning Log: LOG_NODE Action: None required.
0281: lpfc_cleanup_node: ndlp:x%pusgmap:x%x refcnt:%d, void *)ndlp, ndlp->nlp_usg_map, atomic_read(&ndlp->kref.refcount) Node cleanup was called to prepare the node for release. Data: None Severity: Warning Log: LOG_NODE Action: None required.
0282: ldid:x%x ndlp:x%pusgmap:x%x refcnt:%d, ndlp->nlp_DID, (void *)ndlp, lpfc_init.c-ndlp->nlp_usg_map, Driver clean-up has found a node that is still on the node list during driver unload or PCI hotplug removal. Data: None Severity: Error Log: LOG_NODE Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0283: Failed to allocate mbox cmd memory A mailbox allocation error occurred. Data: None Severity: Error Log: LOG_INIT Action: None required.
0285: Allocated DMA memory size <alloclen> is less than the requested DMA memorysize<reqlen> Memory allocation was truncated. Data: None Severity: Error Log: LOG_INIT Action: None required.
0286: lpfc_nlp_state_cleanup failed to allocate statistical data buffer <nlp_DID> Memory allocation failed for the node statistical data. Data: None Severity: Error Log: LOG_INIT Action: None required.
0287: lpfc_alloc_bucket failed to allocate statistical data buffer DID <nlp_DID> Memory allocation failed for the node statistical data. Data: None Severity: Error Log: LOG_NODE Action: None required.
0288: Unknown FCoE event type <event_type> event tag <event_tag> The firmware has detected an unknown FCoE event. Data: None Severity: Error Log: LOG_SLI Action: Check the FCoE switch configuration and the HBA DCBX mode.
0289: Issue Register VFI failed: Err <rc> The driver could not register the Virtual Fabric Index for the FCFI. Data: None Severity: Error Log: LOG_ELS Action: Check the switch and port configurations.
0290: The SLI4 DCBX asynchronous event is not handled yet The SLI-4 DCBX asynchronous event is not handled yet. Data: None Severity: Error Log: LOG_SLI Action: None required.
0291: Allocated DMA memory size (x%x) is less than the requested DMA memory size (x%x) The asynchronous DCBX events are not handled in the driver. Data: None Severity: Error Log: LOG_INIT Action: Check the switch configuration.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0293: PM resume failed to start worker thread: error=<error> The PCI resume (hotplug) could not start the worker thread for the driver. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.
0294: PM resume Failed to enable interrupt The PCI resume (hotplug) could not get an interrupt vector. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.
0297: Invalid device group<pci_dev_grp> While unloading the driver, the driver detected a PCI device that it should not have claimed. Data: None Severity: Error Log: LOG_INIT Action: None required.
0299: Invalid SLI revision <sli_rev> While processing a host attention or unrecoverable error, the driver detected an invalid SLI revision. Data: None Severity: Error Log: LOG_INIT Action: None required.
0300: LATT: Cannot issue READ_LA: Data:<rc> The link attention handler could not issue a READ_LA mailbox command. Data: None Severity: Error Log: LOG_MBOX Action: None required.
0301: READ_SPARAM: no buffers The driver attempted to issue a READ_SPARAM mailbox command to the adapter, but there no buffers were available. Data: None Severity: Warning Log: LOG_MBOX verbose Action: This message indicates: (1) Kernel virtual memory is depleted. Check that the system meets minimum RAM requirements for the Emulex FC adapter. Try closing other applications to free some memory. (2) A possible driver buffer management issue occurred. If this issue persists, report the error to Broadcom Technical Support.
0302: REG_LOGIN: no buffers The driver attempted to issue a REG_LOGIN mailbox command to the adapter, but no buffers were available. Data: (1) Did, (2) flag Severity: Warning Log: LOG_MBOX verbose Action: This message indicates: (1) Kernel virtual memory is depleted. Check that the system meets minimum RAM requirements for the Emulex FC adapter. Try closing other applications to free some memory. (2) A possible driver buffer management issue occurred. If this issue persists, report the error to Broadcom Technical Support.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0313: Ring <ringno> handler: unexpected Rctl <Rctl> Type <Type> received The RCTL/Type of a received frame did not match any for the configured masks for the specified ring. Data: None Severity: Warning Log: LOG_SLI verbose Action: This error could indicate a software driver, firmware, or hardware issue. Report these errors to Broadcom Technical Support.
0303: Ring <ringno> handler: portRspPut <portRspPut> is bigger then rsp ring <portRspMax> The port rsp ring put index is larger than the size of the rsp ring. Data: None Severity: Error Log: Always Action: This error could indicate a software driver, firmware or hardware issue. Report these errors to Broadcom Technical Support.
0304: Stray mailbox interrupt, mbxCommand <mbxcommand> mbxStatus <mbxstatus> Received a mailbox completion interrupt, and no outstanding mailbox commands exist. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
0306: CONFIG_LINK mbxStatus error <mbxStatus> HBA state <hba_state> The driver issued a CONFIG_LINK mailbox command to the HBA that failed. Data: None Severity: Error Log: Always Action: This error could indicate a firmware or hardware issue. Report these errors to Broadcom Technical Support.
0310: Mailbox command <mbxcommand> timeout A mailbox command was posted to the adapter and did not complete within 30 seconds. Data: (1) hba_state (2) sli_flag (3) mbox_active Severity: Error Log: Always Action: This error could indicate a software driver or firmware issue. If no I/O is going through the adapter, reboot the system. If this issue persists, report the error to Broadcom Technical Support.
0312: Ring <ringno> handler: portRspPut <rspPutInx> is bigger then rsp ring <numRiocb> The IOCB command rings put pointer is ahead of the get pointer. Data: None Severity: Error Log: LOG_SLI Action: None required.
0313: Ring <ringno> handler: unexpected Rctl <Rctl> Type <Type> received The RCTL/Type of a received frame did not match any of the configured masks for the specified ring. Data: None Severity: Warning Log: LOG_SLI verbose Action: This error could indicate a software driver, firmware, or hardware issue. Report these errors to Broadcom Technical Support.
0315: Ring <ringno> issue: portCmdGet <local_getidx> is bigger then cmd ring <max_cmd_idx> The port command ring get index is greater than the size of the command ring. Data: None Severity: Error Log: Always Action: This error could indicate a software driver, firmware, or hardware issue. Report these errors to Broadcom Technical Support.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0317: iotag <ulp_IoTag> is out of range: max iotag <max_iotag> wd0 <wd0> The IoTag in the completed IOCB is out of range. Data: None Severity: Error Log: Always Action: This error could indicate a software driver, firmware, or hardware issue. Report these errors to Broadcom Technical Support.
0318: Failed to allocate IOTAG. last IOTAG is <last_allocated_iotag> The driver cannot allocate an IoTag. Display the last value used. Data: None Severity: Error Log: Always Action: This message indicates the adapter HBA I/O queue is full. Typically, this occurs when heavy I/O is running on a low-end (3 digit) adapter. Upgrade to a higher-end adapter.
0319: READ_SPARAM mbxStatus error <mbxStatus> hba state <hba_state> The driver issued a READ_SPARAM mailbox command to the HBA that failed. Data: None Severity: Error Log: Always Action: This error could indicate a firmware or hardware issue. Report these errors to Broadcom Technical Support.
0320: CLEAR_LA mbxStatus error <mbxStatus> hba state <hba_state> The driver issued a CLEAR_LA mailbox command to the HBA that failed. Data: None Severity: Error Log: Always Action: This error could indicate a firmware or hardware issue. Report these errors to Broadcom Technical Support.
0322: Ring <ringno> handler: unexpected completion IoTag <IoTag> The driver could not find a matching command for the completion received on the specified ring. Data: (1) ulpStatus, (2) ulpWord[4], (3) ulpCommand, (4) ulpContext Severity: Warning Log: LOG_SLI verbose Action: This error could indicate a software driver or firmware issue. If this issue persists, report these errors to Broadcom Technical Support.
0323: Unknown Mailbox command <mbxCommand> Cmpl An unknown mailbox command completed. Data: None Severity: Error Log: Always Action: This error could indicate a software driver, firmware, or hardware issue. Report these errors to Broadcom Technical Support.
0324: Config port initialization error, mbxCmd <mbxCommand> READ_NVPARM, mbxStatus <mbxStatus> The READ_NVPARMS mailbox command failed during port configuration. Data: None Severity: Error Log: Always Action: This error could indicate a software driver, firmware, or hardware issue. Report these errors to Broadcom Technical Support.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0328: Rsp Ring <ring number> error: IOCB Data: The firmware has returned an error for this IOCB. Data: (1) <iocb word[0]:iocb word[7]>, (2) <rsp word[0]:rsp[word[7]> Severity: Warning Log: LOG_SLI Action: None required.
0330: IOCB wake NOT set The completion handler associated with the IOCB was never called. Data: (1) timeout (2) timeleft/jiffies Severity: Error Log: Always Action: This error could indicate a software driver, firmware or hardware issue. If this issue persists, report the error to Broadcom Technical Support.
0334: Unknown IOCB command Received an unknown IOCB command completion. Data: (1) type (2) ulpCommand (3) ulpStatus (4) ulpIoTag (5) ulpContext) Severity: Error Log: Always Action: This error could indicate a software driver or firmware issue. If this issue persists, report these errors to Broadcom Technical Support.
0335: Unknown IOCB command Received an unknown IOCB command completion. Data: (1) ulpCommand (2) ulpStatus (3) ulpIoTag (4) ulpContext) Severity: Error Log: Always Action: This error could indicate a software driver or firmware issue. If this issue persists, report these errors to Broadcom Technical Support.
0336: Rsp Ring <ringno> error: IOCB An IOCB error has occurred on the specified ring. Data: (1) ulpWord[0], (2) ulpWord[1], (3) ulpWord[2], (4) ulpWord[3], (5) ulpWord[4], (6) ulpWord[5], (7) irsp+6, (8) irsp+7 Severity: Warning Log: LOG_SLI verbose Action: If this issue persists, check the targets. If the targets are correct, report the error to Broadcom Technical Support.
0340: Adapter temperature is OK now The adapter temperature has reverted to normal range. Data: Temperature in Celsius Severity: Error Log: LOG_TEMP verbose Action: No action needed, informational.
0341: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <un.ulpWord[3]> No more pre-allocated buffers are available to handle unsolicited buffers. Data: None Severity: Error Log: LOG_SLI Action: Ensure this port is not being managed by multiple ports.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0342: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <unsli3.sli3Words> This is a multiple IOCB unsolicited command and sufficient buffer space cannot be allocated for it. Data: None Severity: Error Log: LOG_SLI Action: None required.
0343: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <un.ulpWord[3]> No more pre-allocated buffers are available to handle unsolicited buffers. Data: None Severity: Error Log: LOG_SLI Action: None required.
0344: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <unsli3.sli3Words[7]> No more pre-allocated buffers are available to handle unsolicited buffers. Data: None Severity: Error Log: LOG_SLI Action: None required.
0345: Resetting board due to mailbox timeout iocb. tag 0x%x A mailbox command failed to complete. The driver is resetting the port. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: If the mailbox command fails again, set the lpfc_log_verbose to LOG_MBOX and retry.
0346: Ring <ring number> handler: unexpected ASYNC_STATUS evt_code <evt code> W0 <hex w0> W1 <hex w1> W2 <hex W2> W3 <hex W3> W4 <hex W4> W5 <hex W5> W6 <hex W6> W7 <hex W7> W8 <hex W8> W9 <hex W9> W10 <hex W10> W11<hex W11> The HBA received an asynchronous event that was not a temperature event. Data: None Severity: Error Log: LOG_SLI Action: None required.
0347: Adapter is very hot, please take corrective action Adapter temperature is above normal range. Data: Temperature in Celsius Severity: Error Log: LOG_TEMP verbose Action: Shutdown and remove the HBA. Contact Broadcom Technical Support.
0348: NameServer login: node freed The enable mode failed to free up the NameServer login. Data: None Severity: Error Log: LOG_ELSI Action: None required.
0349: rc should be MBX_SUCCESS The next mailbox command on the mailbox queue has failed. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0350: rc should have been MBX_BUSY Attempting to unregister a default RPI from an interrupt context and the mailbox state is not busy. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
0351: Config MSI mailbox command failed, mbxCmd <u.mb.mbxComm>, mbxStatus <u.mb.mbxStatus> The mailbox command sent to the firmware to configure the adapter to use MSI-X has failed. Data: None Severity: Warning Log: LOG_MBOX Action: Ensure that the hardware platform supports MSI-X.
0352: Config MSI mailbox command failed, mbxCmd <u.mb.mbxCommand>, mbxStatus <u.mb.mbxStatus> The mailbox command sent to the firmware to configure the HBA to use MSI-X has failed. Data: None Severity: Error Log: LOG_MBOX Action: Ensure that the hardware platform supports MSI-X.
0353: Active Mailbox cleared - mailbox timeout exiting The mailbox timeout handler has determined that the driver is in the process of completing this mailbox command. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
0357: MSI-X interrupt with no EQE An SLI-4 adapter interrupt on the slow path occurred, but there is no associated EQE. Data: None Severity: Warning Log: LOG_SLI Action: None required.
0358: MSI-X interrupt with no EQE An SLI-4 adapter interrupt on the fast path occurred, but there is no associated EQE. Data: None Severity: Warning Log: LOG_SLI Action: None required.
0359:Not a valid slow-path completion `event: majorcode=x%x, minorcode=x%x\n`, bf_get(lpfc_eqe_major_code, eqe), bf_get(lpfc_eqe_minor_code, eqe)); SLI-4: The EQE is invalid. Data: None Severity: Error Log: LOG_SLI Action: None required.
0360:Unsupported EQ count. <entry_count> Cannot create an event queue of this size. Data: None Severity: Error Log: LOG_SLI Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0361:Unsupported CQ count. <entry_count> Cannot create an completion queue of this size. Data: None Severity: Error Log: LOG_SLI Action: None required.
0362:Unsupported MQ count. <entry_count> Cannot create MQ of this size. Data: None Severity: Error Log: LOG_SLI Action: None required.
0364:Invalid param: SLI-4: The post-SGL function was passed an invalid extensible resource indicator (XRI). Data: None Severity: Error Log: LOG_SLI Action: None required.
0365:Slow-path CQ identifier <CQID> does not exist: The Completion Queue ID passed in the event queue entry does not reference a valid completion queue. Data: None Severity: Error Log: LOG_SLI Action: None required.
0366:Not a valid fast-path completion event: majorcode=<major code hex>, minor-code=<minor code hex> The major or minor code in the Event Queue field is invalid. Data: None Severity: Error Log: LOG_SLI Action: None required.
0367: Fast-path completion queue does not exist The fast path completion queue referenced by the CQID does not exist. Data: None Severity: Error Log: LOG_SLI Action: None required.
0368: Mis-matched fast-path completion queue identifier: eqcqid=%d, fcpcqid=%d The CQID in the event queue entry does not match the fcp_cqid that was passed into the routine. Data: None Severity: Error Log: LOG_SLI Action: None required.
0369: No entry from fast-path completion queue fcpcqid=<queue_id> No completions in the completion queue were referenced by fcp_cqid. Data: None Severity: Error Log: LOG_SLI Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0370: Invalid completion queue type <type> The event queue entry is not for a mailbox or a work queue entry. Data: None Severity: Error Log: LOG_SLI Action: None required.
0371: No entry from the CQ: identifier <queue_id>, type <type> No completion queue event existed for this event queue entry. Data: None Severity: Error Log: LOG_SLI Action: None required.
0372: iotag <iotag> is out of range: max iotag (<sli.last_iotag>) The IOCB lookup cannot be performed because the iocb_tag is out of range. Data: None Severity: Error Log: LOG_SLI Action: None required.
0373: FCP complete error: status=<status> hw_status=<hw status>, total_data_specified=<total data transferred>, parameter=<rsp word[4]>, word3=<wcqe word 3> Logs the FCP failure. Status and parameter are equivalent to ulpStatus and ulpWord[4]. Data: None Severity: Warning Log: LOG_SLI Action: None required.
0374: FCP complete with no corresponding cmdiocb: iotag <iocb iotag> No IOCB was on the in-progress list that matched this iotag. Data: None Severity: Warning Log: LOG_SLI Action: None required.
0375: FCP cmdiocb not callback function iotag: <iocb iotag> The IOCB found for this iotag does not have a completion handler set in it. Data: None Severity: Warning Log: LOG_SLI Action: None required.
0377: Error <rc> parsing vpd. Using defaults. Could not parse the vital product data (VPD) data, so the driver is using the default values. Data: None Severity: Error Log: Always Action: None required.
0378: No support for fcpi mode. Could not configure the port to run in FCP initiator mode. Data: None Severity: Warning Log: LOG_MBOX, LOG_SLI Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0379: Feature Mismatch Data: <req ftr word2 hex> <req_ftr word3 hex> <cfg_enable_npiv> <max vpi hex> The features passed in to the driver as module parameters do not match what the firmware can do. Set to default values. Data: None Severity: Warning Log: LOG_MBOX, LOG_SLI Action: None required.
0381: Error %d during queue setup. Could not set up all the queues that the driver requires to exchange I/Os with the HBA. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: Reload the driver.
0382: READ_SPARAM command failed status <issue status>, mbxStatus <mailbox status> The READ_SPARAM mailbox command has failed during initialization. The HBA has been set to error state. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI: Action: Perform a dump with `hbacmd` and then try reloading the driver.
0383: Error <rc> during scsi sgl post operation The SGL entries could not be registered with the adapter. Data: None Severity: Warning Log: LOG_MBOX, LOG_SLI Action: Reset the adapter using `hbacmd`.
0384: There is pending active mailbox cmd The mailbox commands have overlapped. This command should have been added to the mailbox queue. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
0385: rc should have been MBX_BUSY The completion handler for REG_LOGIN detected the IMMED_UNREG flag and tried to issue the UNREG_LOGIN command from an interrupt level. The mailbox status should still be busy. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
0386: ELS complete with no corresponding cmdiocb: iotag <iotag> The completion that the ISR is handling cannot find a tag associated with the IOTAG. Data: None Severity: Warning Log: LOG_SLI Action: None required.
0387:Failed to allocate an iocbq Failed to get an IOCBQ from the list of available IOCBQs. Data: None Severity: Error Log: LOG_SLI Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0388:Not a valid WCQE code: x<hex cqe_code> The event code is invalid. This event will be dropped. Data: None Severity: Error Log: LOG_SLI Action: Ensure the adapter firmware is current.
0391:Error during rpi post operation The driver was trying to post pages to the firmware to be used to keep target login information and encountered a failure. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: Unload and reload the driver.
0393:Error <rc> during rpi post operation The driver was trying to post pages to the firmware to keep target login information and encountered a failure. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: Unload and reload the driver.
0394: Failed to allocate CQ_EVENT entry The asynchronous event handler was unable to allocate an event queue entry to which to transfer the asynchronous event. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: This could be a V-LINK clear from the switch or a fatal error from the firmware. Perform a dump from the Emulex OneCommand CNA Manager application.
0395: The mboxq allocation failed The asynchronous link event handler could not allocate a mailbox command to issue the READ_LA (read link attention) mailbox command. Data: None Severity: Error Log: LOG_SLI Action: None required.
0396:The lpfc_dmabuf allocation failed The asynchronous link event handler could not allocate a mailbox command to issue the READ_LA mailbox command. Data: None Severity: Error Log: LOG_SLI Action: None required.
0397:The mbuf allocation failed The asynchronous link event handler could not allocate DMA-able memory for the READ_LA mailbox command. Data: None Severity: Error Log: LOG_SLI Action: None required.
0398 Invalid link fault code: <hex link_fault> The attempt to read the link attention register has returned an unknown value. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0399 Invalid link attention type: <hex link_type> The READ_LA mailbox command has returned an invalid link type. Data: None Severity: Error Log: LOG_INIT Action: None required
0400: lpfc_nodev_tmo attribute cannot be set to <val>, allowed range is [<LPFC_MIN_DEVLOSS_TMO>, <LPFC_MAX_DEVLOSS_TMO>] The attempt to set the devloss timeout value failed because the value is out of the allowable range. Data: None Severity: Error Log: LOG_INIT Action: Use a value between the minimum and maximum values.
0401: Ignoring change to nodev_tmo because devloss_tmo is set Attempted to change the nodev timeout when the devloss has already been set. Data: None Severity: Error Log: LOG_INIT Action: None required.
0402:Cannot find virtual addr for buffer tag on ring <ringno> A DMA buffer is unavailable for this unsolicited command. Data: (1) tag (2) next (3) prev (4) postbufq_cnt Severity: Error Log: LOG_INIT Action: None required.
0403: lpfc_nodev_tmo attribute cannot be set to <val>, allowed range is [<LPFC_MIN_DEVLOSS_TMO>, <LPFC_MAX_DEVLOSS_TMO>] Attempted to set the nodev timeout value is outside the range of the devloss timeout range. Data: None Severity: Error Log: LOG_INIT Action: Set the nodev timeout between the minimum and maximum timeout range.
0404: lpfc_devloss_tmo attribute cannot be set to <val>, allowed range is [<LPFC_MIN_DEVLOSS_TMO>, <LPFC_MAX_DEVLOSS_TMO>] Attempted to set the devloss timeout value is outside the allowed range. Data: None Severity: Error Log: LOG_INIT Action: Set the devloss timeout between the minimum and maximum devloss range.
0405: lpfc_link_speed attribute cannot be set to %d, allowed values are ["LPFC_LINK_SPEED_STRING"] Attempted to set the link speed value is outside the allowed range. Data: None Severity: Error Log: LOG_INIT Action: Set the link speed between 0 and the maximum.
0406: Adapter maximum temperature exceeded <temperature>, taking this port offline The driver has received an error for the HBA indicating that the maximum allowable temperature has been exceeded. Data: (1) work_hs (2) work_status[0] (3) work_status[1] Severity: Error Log: LOG_INIT Action: Ensure that the server fans are not blocked. Shut down the server if the airflow is restricted.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0407: Ignoring nodev_tmo module parameter because devloss_tmo is set. Both module parameters (`nodev` and `devloss`) were set so the driver is ignoring the `nodev` parameter. Data: None Severity: Error Log: LOG_INIT Action: Only one of these parameters must be set.
0410: Cannot find virtual addr for mapped buf on ring <ringno> The driver cannot find the specified buffer in its mapping table. Therefore, it cannot find the virtual address needed to access the data. Data: (1) phys (2) next (3) prev (4) postbufq_cnt Severity: Error Log: Always Action: This error could indicate a software driver or firmware issue. If this issue persists, report these errors to Broadcom Technical Support.
0421: MSI-X slow-path request_irq failed <rc> The kernel application programming interface (API) to request an IRQ has failed. Data: None Severity: Warning Log: LOG_INIT Action: Use module parameter `lpfc_use_msi=0 (INTx)`.
0422: lpfc_restrict_login attribute cannot be set to <val>, allowed range is [0, 1] Attempted to set the restrict login parameter to something other than on or off. Data: None Severity: Error Log: LOG_INIT Action: Use 0 (Off) or 1 (On)
0423: lpfc_"#attr" attribute cannot be set to %d, allowed range is ["#minval", "#maxval"] This compile time macro is used by several module parameters during initialization. Each module parameter has its own minimum and maximum values that are displayed. Data: None Severity: Error Log: LOG_INIT Action: Set the module parameter between the minimum and maximum values.
0424:lpfc_"#attr" attribute cannot be set to %d, allowed range is ["#minval", "#maxval"] This compile time macro is used by several module parameters to set the value. Data: None Severity: Error Log: LOG_INIT Action: Set the module parameter between the minimum and maximum values.
0425:lpfc_restrict_login attribute cannot be set to %d, allowed range is [0, 1] The module parameter `lpfc_restrict_login` can only be set to 0 (off) or 1 (on). Data: None Severity: Error Log: LOG_INIT Action: Set `lpfc_restrict_login=[0,1]`.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0426: Failed to enable interrupt The driver failed to start the interrupt. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.
0427: Cannot re-enable interrupt after slot reset The driver was unable to enable the interrupt after an HBA reset. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.
0429: MSI-X fast-path request_irq failed (<rc>) The driver received an error for the request_irq_call. Data: None Severity: Warning Log: LOG_INIT Action: Unload and reload the driver.
0430: PM resume Failed to enable interrupt The driver's power management resume function could not enable the interrupt. Data: None Severity: Error Log: LOG_INIT Action: Perform another PM suspend and resume or HBA reset.
0431: Failed to enable interrupt. The driver failed to start the interrupt. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.
0433: Wakeup on signal: rc=<rc> A signal other than the LPFC_DATA_READY was received on the worker thread. Data: None Severity: Error Log: LOG_ELS Action: Unload and reload the driver.
0434: PM resume failed to start worker thread: error=<error> The driver's power management resume function could not start the worker thread. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.
0435: Adapter failed to get Option ROM version status <rc> The driver could not read the HBA's option ROM. Data: None Severity: Error Log: LOG_INIT Action: Reset the HBA. Ensure that the adapter's firmware is current.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0436: Adapter failed to init, timeout, status reg <status> The adapter failed during power-up diagnostics after it was reset. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
0437: Adapter failed to init, chipset, status reg <status> The adapter failed during power-up diagnostics after it was reset. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
0438: Adapter failed to init, chipset, status reg <status> The adapter failed during power-up diagnostics after it was reset. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
0439: Adapter failed to init, mbxCmd <mbxCommand> READ_REV, mbxStatus <mbxStatus> Adapter initialization failed when issuing a READ_REV mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
0440: Adapter failed to init, READ_REV has missing revision information A firmware revision initialization error was detected. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. Install the latest firmware revision. If this issue persists, report the error to Broadcom Technical Support.
0442: Adapter failed to init, mbxCmd <mbxCommand> CONFIG_PORT, mbxStatus <mbxStatus> Adapter initialization failed when issuing a CONFIG_PORT mailbox command. Data: (1) hbainit Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
0443: Adapter failed to set maximum DMA length mbxStatus <u.mb.mbxStatus> Cannot set the maximum DMA length to reflect cfg_pci_max_read. Data: None Severity: Error Log: LOG_INIT Action: Set module parameter lpfc_pci_max_read to 512, 1024, 2048, or 4096.
0446: Adapter failed to init, mbxCmd <mbxCommand> CFG_RING, mbxStatus <mbxStatus>, ring <num> Adapter initialization failed when issuing a CFG_RING mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0447: Adapter failed init, mbxCmd <mbxCommand> CONFIG_LINK mbxStatus <mbxStatus> Adapter initialization failed when issuing a CONFIG_LINK mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
0448: Adapter failed to init, mbxCmd <mbxCommand> READ_SPARM, mbxStatus <mbxStatus> Adapter initialization failed when issuing a READ_SPARM mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
0449: lpfc_%attr attribute cannot be initialized to %d, allowed range is [%min, %max] The sysfs attribute value written exceeds the attribute range. Data: (1) attribute name (2) value written (3) minimum value (3) maximum value Severity: Error Log: Always Action: Write a value within the supported range.
0450: lpfc_%attr attribute cannot be set to %d, allowed range is [%min, %max] The sysfs attribute value written exceeds the attribute range. Data: (1) attribute name (2) value written (3) minimum value (3) maximum value Severity: Error Log: Always Action: Write a value within the supported range.
0451: Enable interrupt handler failed The driver attempted to register the HBA interrupt service routine with the host operating system but failed. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or driver issue. If this issue persists, report the error to Broadcom Technical Support.
0453: Adapter failed to init, mbxCmd <mbxCommand> READ_CONFIG, mbxStatus <mbxStatus> Adapter initialization failed when issuing a READ_CONFIG mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
0454: Adapter failed to init, mbxCmd <mbxCommand> INIT_LINK, mbxStatus <mbxStatus> Adapter initialization failed when issuing an INIT_LINK mailbox command. Data: None Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
0456: Adapter failed to issue ASYNCEVT_ENABLE mbox status x%x The mailbox command to enable an asynchronous event notification failed. Data: None Severity: Error Log: LOG_INIT Action: Ensure that the adapter firmware is current. Reload the driver.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0457: Adapter Hardware Error The driver received an interrupt indicating a possible hardware issue. Data: (1) status (2) status1 (3) status2 Severity: Error Log: Always Action: This error could indicate a hardware or firmware issue. If this issue persists, report the error to Broadcom Technical Support.
0458: Bring adapter online The FC driver has received a request to bring the adapter online. This error might occur when running `lputil`. Data: None Severity: Warning Log: LOG_INIT verbose Action: None required.
0459: Adapter heartbeat failure, taking this port offline. The Heartbeat mailbox command failed. Data: None Severity: Error Log: LOG_INIT Action: Ensure the adapter firmware is current. Reload the driver.
0460: Bring adapter offline The FC driver has received a request to bring the adapter offline. This error might occur when running `lputil`. Data: None Severity: Warning Log: LOG_INIT verbose Action: None required.
0466: Outstanding I/O when bringing Adapter offline The I/O is still pending while attempting to stop the driver. Data: None Severity: Warning Log: LOG_INIT Action: None required.
0467: lpfc_topology attribute cannot be set to %d, allowed range is [0, 6], phba->brd_no, val. The `lpfc_topology` module parameter is invalid. Data: None Severity: Error Log: LOG_INIT Action: Use a value in the valid range.
0468: lpfc_restrict_login must be 0 for Physical ports. “vport->cfg_restrict_login = 0; Cannot restrict the login for the physical port. Data: None Severity: Error Log: LOG_INIT Action: None required.
0469: lpfc_link_speed attribute cannot be set to %d, allowed range is [0, 8] The link speed module parameter is invalid. Data: None Severity: Error Log: LOG_INIT Action: Use a link speed parameter in the valid range.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0472: PCI channel I/O permanent failure The PCI bus has detected an error. Data: None Severity: Error Log: LOG_INIT Action: Issue an HBA reset.
0474: Unable to allocate memory for issuing MBOX_CONFIG_MSI command A mailbox memory pool allocation error occurred. Data: None Severity: Error Log: LOG_INIT Action: None required.
0475: Not configured for supporting MSI-X cfg_use_msi: 0x%x The lpfc_use_msi module parameter should have been set to 2. Data: None Severity: Error Log: LOG_INIT Action: Set the module parameter `lpfc_use_msi = 2`.
0476: HBA not supporting SLI-3 or later SLI Revision: <sli_rev> The HBA does not support SLI-3 or SLI-4. Data: None Severity: Error Log: LOG_INIT Action: This HBA does not support MSI. Set `lpfc_use_msi=0`.
0478: MSI request_irq failed (<rc>). The request_irq kernel API has failed. Data: None Severity: Warning Log: LOG_INIT Action: Set lpfc_use_msi=0.
0479: Deferred Adapter Hardware Error An adapter hardware error was sent to the driver. Data: (1) work_hs, (2) work_status[0], (3) work_status[1] Severity: Error Log: LOG_INIT Action: Perform a dump using `hbacmd`.
0483:Invalid link-attention link speed: x%x", bf_get(lpfc_acqe_link_speed, acqe_link). The link speed reported in the link attention interrupt is invalid. Data: None Severity: Error Log: LOG_INIT Action: Check the switch configuration.
0485: MSI-X slow-path request_irq failed (<rc>). The request_irq kernel API has failed. Data: None Severity: Warning Log: LOG_INIT Action: Set the module parameter `lpfc_use_msi=0`.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0486: MSI-X fast-path (<index>) request_irq failed (<rc>). The request_irq kernel API has failed. Data: None Severity: Warning Log: LOG_INIT Action: Set the module parameter `lpfc_use_msi=0`.
0490: MSI request_irq failed (<rc>). The request_irq kernel API has failed. Data: None Severity: Warning Log: LOG_INIT Action: Set module parameter lpfc_use_msi=0.
0492: Unable to allocate memory for issuing SLI_CONFIG_SPECIAL mailbox command A mailbox memory pool allocation error occurred. Data: None Severity: Error Log: LOG_INIT Action: None required.
0493: SLI_CONFIG_SPECIAL mailbox failed with status<rc> A mailbox command failed. Data: None Severity: Error Log: LOG_INIT Action: Ensure that the adapter firmware is current. Unload and reload the driver.
0494: Unable to allocate memory for issuing "SLI_FUNCTION_RESET mailbox command" A mailbox memory pool allocation error occurred. Data: None Severity: Error Log: LOG_INIT Action: None required.
0495: SLI_FUNCTION_RESET mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>. A mailbox command failed. Data: None Severity: Error Log: LOG_INIT Action: Reset the HBA.
0496: Failed allocate slow-path EQ The event queue for the slow path was not allocated. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.
0497: Failed allocate fast-path EQ The event queue for the fast path was not allocated. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0499: Failed allocate fast-path FCP CQ (<fcp_cqidx>). The completion queue event for the fast path could not be allocated. Data: None Severity: Error Log: LOG_INIT Action: Unload and reload the driver.
0500: Failed allocate slow-path mailbox CQ Failed to allocate a slow-path mailbox CQ. Data: None Severity: Error Log: LOG_INIT Action: None required.
0501: Failed allocate slow-path ELS CQ Failed to allocate a slow-path ELS CQ. Data: None Severity: Error Log: LOG_INIT Action: None required.
0503: Failed allocate fast-path FCP Failed to allocate a fast-path FCP. Data: None Severity: Error Log: LOG_INIT Action: None required.
0504: Failed allocate slow-path ELS WQ Failed to allocate a slow-path ELS WQ. Data: None Severity: Error Log: LOG_INIT Action: None required.
0505: Failed allocate a slow-path ELS MQ Data: None Severity: Error Log: LOG_INIT Action: None required.
0506: Failed allocate receive HRQ\n Data: None Severity: Error Log: LOG_INIT Action: None required.
0507: Failed allocate receive DRQ Failed to allocate a receive DRQ. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0520: Slow-path EQ not allocated The slow-path EQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.
0522: Fast-path EQ <fcp_eqidx> not allocated The fast-path EQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.
0523: Failed setup of fast-path EQ <fcp_eqidx>, rc = <rc> The fast-path EQ set up failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
0526: Fast-path FCP CQ <fcp_cqidx> not allocated The fast-path FCP is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.
0527: Failed setup of fast-path FCP CQ <fcp_cqidx>, rc = <rc> The fast-path FCP CQ set up failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
0528: Mailbox CQ not allocated The mailbox CQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.
0530: ELS CQ not allocated The ELS CQ is not allocated Data: None Severity: Error Log: LOG_INIT Action: None required.
0534: Fast-path FCP WQ <fcp_wqidx> not allocated The fast-path FCP WQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0535: Failed setup of fast-path FCP WQ <fcp_wqidx>, rc = <rc> The fast-path FCP WQ set up failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
0536: Slow-path ELS WQ not allocated The slow-path ELS WQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.
0538: Slow-path MQ not allocated The slow-path MQ is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.
0540: Receive Queue not allocated The Receive Queue is not allocated. Data: None Severity: Error Log: LOG_INIT Action: None required.
0542: lpfc_create_static_vport failed to allocate mailbox memory Failed to allocate mailbox memory for VPort creation. Data: None Severity: Error Log: LOG_INIT Action: None required.
0543: lpfc_create_static_vport failed to allocate vport_info\n")) Failed to allocate vport_info. Data: None Severity: Error Log: LOG_INIT Action: None required.
0544: lpfc_create_static_vport failed to issue dump mailbox command ret <rc> status <mbxStatus> Failed to issue a dump mailbox command for static VPort creation. Data: None Severity: Warning Log: LOG_INIT Action: None required.
0545: lpfc_create_static_vport bad information header 0x%x 0x%x\n"\, le32_to_cpu(vport_info->signature), le32_to_cpu(vport_info->rev) & VPORT_INFO_REV_MASK); Invalid information header; the signature or revision is invalid. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0546: lpfc_create_static_vport failed to create vport Failed to create a VPort. Data: None Severity: Warning Log: LOG_INIT Action: None required.
0560: lpfc_enable_auth attribute cannot be set to <val>, allowed range is [0, 1] The lpfc_enable_auth attribute can only be 0 or 1. Data: None Severity: Error Log: LOG_INIT Action: None required.
0582: Error <rc> during sgl post operation The SGL post operation failed. Data: None Severity: Error Log: LOG_MBOX, LOG_IP verbose Action: None required.
0602: Failed to allocate CQ_EVENT entry Failed to allocate a CQ_EVENT entry. Data: None Severity: Error Log: LOG_SLI Action: None required.
0603: Invalid work queue CQE subtype (x%x)\n", cq-<subtype> Invalid work queue completion queue entry (CQE) occurred. Data: None Severity: Error Log: LOG_SLI Action: None required.
0700: Bus Reset on target <i> failed The bus reset for the specified target failed. Data: None Severity: Error Log: LOG_FCP Action: None required.
0704: At limitation of <total> preallocated command buffers The maximum number of command buffers have already been allocated. Data: None Severity: Warning Log: LOG_FCP verbose Action: None required.
0705: Allocation request of <num> command buffers will exceed max of <hba_queue_depth>. Reducing allocation request to <size> The number of command buffers requested will exceed the maximum so a smaller quantity will be allocated. Data: None Severity: Warning Log: LOG_FCP verbose Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0708: Allocation request of <num_to_alloc> command buffers did not succeed. Allocated <num_allocated> buffers. The allocation request for the specified command buffers did not succeed. However, the specified number of buffers has been allocated. Data: None Severity: Warning Log: LOG_FCP Action: None required.
0711: detected queue full - lun queue depth adjusted to%d The driver detected a queue full status on a SCSI command response. A new LUN queue depth is reported. Data: (1) New LUN queue depth Severity: Warning Log: LOG_FCP verbose Action: This might indicate an oversubscribed target array. Check your SAN configuration and I/O workload.
0713: SCSI layer issued Device Reset (%d, %d) A device reset was issued. Data: None Severity: Error Log: LOG_FCP Action: None required.
0714: SCSI layer issued bus reset The SCSI layer is requesting the driver to abort all I/Os to all targets on this HBA. Data: (1) ret Severity: Error Log: Always Action: Check the state of the targets in question.
0720: FCP command <cmnd[0]> residual overrun error A residual overrun error has occurred while processing the specified FCP command. Data: (1) request_bufflen (2) resid Severity: Warning Log: LOG_FCP verbose Action: If this issue persists, check the targets for errors.
0721: Device Reset rport failure: rdata <rdata> The reset of the R_Port failed. Data: None Severity: Error Log: LOG_FCP Action: None required.
0722: Target Reset rport failure: rdata <rdata> The reset of the target failed. Data: None Severity: Error Log: LOG_FCP Action: None required.
0723: SCSI layer issued Target Reset (%d, %d) The SCSI layer issued a target reset. Data: None Severity: Error Log: LOG_FCP Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0724: I/O flush failure for context <"LUN","TGT","HOST","Unknown">: cnt <cnt> The I/O flush to the LUN, target, or host has failed. Data: None Severity: Error Log: LOG_FCP Action: None required.
0727: TMF <cmd> to TGT <TGT#> LUN <LUN#> failed (<ulpStatus>, <ulpWord[4]>) The task management command failed. Data: None Severity: Error Log: LOG_FCP Action: None required
0729: FCP cmd <cmnd> failed <target>/<lun> status: <status> result: <result> The specified device failed an FCP command. Data: (1) ulpContext (2) iotag Severity: Warning Log: LOG_FCP verbose Action: Check the state of the target in question.
0730: FCP command failed: RSP The FCP command failed with a response error. Data: (1) resp_info (2) scsi_status (3) ResId (4) SnsLen (5) RspLen (6)rspInfo3 Severity: Warning Log: LOG_FCP verbose Action: Check the state of the target in question.
0734: FCP read check error The issued FCP command returned a read check error. Data: (1) fcpDl (2) rspResId (3) fcpi_parm (4) cmd[0] Severity: Warning Log: LOG_FCP verbose Action: Check the state of the target in question.
0735: FCP Read Check Error and Underrun Data HBA reported under run from storage array. Data: (1) vpi (2) fcpDI (3) res_id (4) fcpi_parm Severity: Warning Log: LOG_FCP_ERROR verbose Action: No action needed, informational.
0748: Abort handler timed out waiting for abort to complete:ret <status> D <target id> LUN <lun id> The abort handler timed out waiting for abort to complete. Data: None Severity: Error Log: Always Action: None required.
0749: SCSI layer issued abort device The SCSI layer aborted a device. Data: (1) ret, (2) id, (3) lun, (4) snum Severity: Warning Log: LOG_FCP verbose Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

0915 Register VPI failed:<mbxStatus> Could not register the VPI. Data: None Severity: Error Log: LOG_MBOX Action: None required.
1005: AUTHENTICATION_FAILURE Nport:<port> The system detected a Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) authentication failure on a port. Data: (1) nlp_DID Severity: Error Log: LOG_SECURITY Action: Verify the authentication settings and keys on the local and remote ports.
1006: Bad Name tag in auth message < message > The DHCHAP Authentication process failed when an invalid tag was detected. Data: (1) message Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1007: Bad Name length in auth message < message > The DHCHAP Authentication process failed when an invalid name was detected. Data: (1) message Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1008: Bad Number of Protocols <message> The DHCHAP Authentication process failed due to an unexpected protocol number. Data: (1) message Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1009: Bad param type <message> The DHCHAP Authentication process failed when an invalid protocol was detected. Data: (1) message Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1010: Bad Tag 1 <message> The DHCHAP Authentication process failed when a bad Tag was detected. Data: (1) message Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1011: Auth_neg no hash function chosen The DHCHAP Authentication process failed when an incorrect hash function was specified. Data: (1) message Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1012: Auth_negotiate Bad Tag <message> DHCHAP Authentication process failed due to bad Tag for auto negotiation. Data: (1) message Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1013: Auth_negotiate no DH_group found DHCHAP Authentication process failed when incorrect or missing DH Group was detected. Data: (1) message Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to the Broadcom Technical Support.
1014: dhchap challenge bad name tag <message> DHCHAP Authentication process failed when incorrect Challenge name tag was detected. Data: (1) message Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to the Broadcom Technical Support.
1015: dhchap challenge bad name length <message> DHCHAP Authentication process failed due to unexpected Challenge name length. Data: (1) message Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1016: dhchap challenge Hash ID not Supported <message> The DHCHAP Authentication process failed due to an uncorroborated Challenge Hash ID. Data: (1) message Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1017: dhchap challenge could not find DH Group The DHCHAP Authentication process failed due to an uncorroborated Challenge Group. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1018: dhchap challenge No Public key for non-NULL DH Group No Public key exists for the non-NULL DH Group. Data: None Severity: Error Log: LOG_SECURITY Action: None required.
1019: Request tranid <tran_id> timed out A transaction with the storage array could not complete due to timeout. Data: (1) tran_id Severity: Warning Log: LOG_SECURITY verbose Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1021: ERROR: attempted to queue security work, when no workqueue created The driver encountered a missing queue required for processing security information. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.
1022: Security request does not exist A security request operation failed because no match was found for such request. Data: None Severity: Warning Log: LOG_SECURITY Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.
1023: Warning - data may have been truncated. Data: <data> reqdl: <data_len> mesdl:<data_len> A security message exchange operation failed because the response was missing or unreliable. Data: None Severity: Warning Log: LOG_SECURITY Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.
1028: Start Authentication: No buffers The authentication failed because some memory resources were not allocated. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1029: Reauthentication Failure The driver encountered errors and a failure to re-authenticate occurred. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1031: Start Authentication: Get config failed The authentication failed due to some error during port configuration. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1032: Start Authentication: get config timed out The node authentication was aborted because a timeout occurred while waiting for port configuration to complete. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1033: Received auth_negotiate from Nport: < nlp_DID> An unsolicited authentication negotiation message was received from a port. Data: (1) nlp_DID Severity: Warning Log: LOG_SECURITY Action: No action, this message is informational.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1034: Not Expecting Challenge - Rejecting Challenge An unsolicited authentication challenge received from a port was rejected. Data: None Severity: Warning Log: LOG_SECURITY Action: Software driver warning. If this issue persists, report errors to Broadcom Technical Support.
elx_mag1036: Authentication transaction reject - re-auth request reason <reason> exp <explanation> An authentication was rejected and requested again due to the reason as displayed with the explanation. Data: (1) reason (2) explanation. Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1037: Authentication transaction reject - restarting authentication, reason <reason> exp <explanation> An authentication process was rejected and then restarted and the authentication requested again due to the reason as displayed with the explanation. Data: (1) reason (2) explanation. Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1039: Not Expecting Reply - rejecting. State <state> An unanticipated reply was received during authentication and was subsequently rejected. Data: (1) auth_state. Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1040: Bad Reply trans_id- rejecting. Trans_id < trans_id > Expecting: < trans_id> An unexpected transaction id was received during authentication and was subsequently rejected. Data: (1) auth_state Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1043: Authentication LS_RJT The authentication request was rejected. Data: None Severity: Error Log: LOG_ELS Action: None required.
1045: Issue AUTH_NEG failed Status:%x The authentication negotiation failed. Data: None Severity: Error Log: LOG_ELS Action: None required.
1048: Issue AUTH_REJECT failed Could not issue the reject for the authentication request. Data: None Severity: Error Log: LOG_ELS Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1049: Authentication is enabled but authentication service is not running Discovery failed because DHCHAP authentication was enabled while no authentication service was established. Data: None Severity: Error Log: LOG_SECURITY Action: Start the authentication daemon (`fcauthd`).
elx_ msg1050: Authentication mode is disabled, but is required by the fabric Discovery failed because the switch fabric required authentication, but authentication was not configured or the authentication mode for this port pair is disabled. Data: None Severity: Error Log: LOG_SECURITY Action: Configure the driver to authenticate with the switch or disable authentication on the switch to this port.
1053: Start Authentication: Security service offline The authentication failed because security service was unavailable. Data: None Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1055: Authentication parameter is disabled, but is required by the fabric FLOGI failed because the fabric has indicated that authentication is required, but authentication has not yet been configured or enabled on the HBA. Data: None Severity: Error Log: LOG_SECURITY Action: Configure authentication on this HBA.
1057: Authentication transaction reject. reason <reason> exp <explanation> An authentication was rejected and requested again due to the reason as displayed with the explanation. Data: (1) reason (2) explanation. Severity: Error Log: LOG_SECURITY Action: Software driver error. If this issue persists, report errors to Broadcom Technical Support.
1058: Waiting for authentication service A delay occurred when the authentication service was not initially available as expected. Data: None Severity: Warning Log: LOG_SECURITY Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.
1059: Authentication became available The authentication service came online, but it was not initially available as expected. Data: None Severity: Warning Log: LOG_SECURITY Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1201: Failed to allocate dfc_host Failed to allocate memory for the dfc_host_struct. Data: None Severity: Error Log: LOG_ELS Action: None required.
1210: Invalid cmd size: cmd <cmd> cmdsz <cmdsize> rspsz <rspsize> The management command for LPFC 2100 has failed. Data: None Severity: Error Log: LOG_LIBDFC Action: None required.
1211: genreq alloc failed\n"); A resource allocation failure occurred. Data: (1) return code Log: LOG_LIBDFC Severity: Error Action: Kernel memory resources too low.
1213: FCoE cmd overflow: off <#> + cnt <#> > cmdsz <#> The application has tried to read more data than originally requested. Data: (1) response offset (2) size (3) cmd size Log: LOG_LIBDFC Severity: Error Action: The application might have sent an invalid command.
1214: Cannot issue FCoE cmd SLI not active: <#> rc= -EACCESS The SLI layer has not been initialized. Data: (1) offset Log: LOG_LIBDFC Severity: Error Action: Restart the HBA.
1215: Cannot issue FCoE cmd: not ready or not in maint mode The external link is unplugged, the link is down, or the FCoE is not in maintenance mode. Data: (1) current offset (2) return code. Log: LOG_LIBDFC Severity: Error Action: Plug in the external cable or set FCoE in maintenance mode.
1216: FCoE IOCB failed: off <#> rc <#> The FCoE command generated by the application has failed. Data: (1) offset (2) return code. Log: LOG_LIBDFC Severity: Error Action: The application should retry the command.
1223: menlo_write: couldn't alloc genreq Resource allocation failure. Data: None Log: LOG_LIBDFC Severity: Error Action: Kernel memory resources too low.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1224: FCoE iocb failed off <#> rc=<#>", FCoE command failed in SLI. Data: (1) offset (2) return code Log: LOG_LIBDFC Severity: Informational. Action: Retry the command, if it fails again, reset HBA when convenient.
1227: FCoE IOCB TMO: handler set for <context3> The management command for the LPFC 2100 has timed out. Data: None Severity: Warning Log: LOG_LIBDFC Action: None required.
1228: FCoE IOCB TMO: handler set for <context3> A management IOCB for the LPFC 2100 has timed out Data: None. Severity: Warning Log: LOG_LIBDFC Action: None required.
1229: Waiting for menlo mnt Waiting for the LPFC 2100 to enter maintenance method. Data: None. Severity: Warning Log: LOG_LIBDFC Action: None required.
1230: Could not find buffer for FCoE cmd:off <#> indmp <addr> off <#> Could not find resources associated with this FCoE command. Data: (1) current offset (2) buffer desc pointer (3) size. Severity: Error Log: LOG_LIBDFC Action: Try reloading the driver when convenient.
1231: bad bpl: A invalids buffer list was detected upon completion. Data: None. Severity: Error Log: LOG_LIBDFC Action: None required.
1235: Could not find buffer for FCoE cmd: off:<#> poff:<#> cnt:<#> mlastcnt:<#> addl:<x> addh:<x> mdsz:<#> FCoE command failed because it could not find the resource. Data: (1) current offset (2)previous offset (3) count (4) last count (5) address low (6) address high Severity: Error Log: LOG_LIBDFC Action: No action needed, informational.
1238: FCoE IOCB failed: off <#> rc=<#> The command generated by the driver to check the FCoE has failed. Data: (1) offset (2) return code Log: LOG_LIBDFC Severity: Error Action: Make sure link is up or the adapter has set menlo in maintenance mode.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1240: Unable to allocate command buffer memory Could not allocate memory for the command buffer. Data: None. Severity: Error Log: LOG_LINK_EVENT Action: None required.
1243: Menlo command error. code=%d.\n", mlorsp->code The Menlo maintenance command failed. Data: None. Severity: Error Log: LOG_LINK_EVENT Action: None required.
1244: Unable to allocate response buffer memory. Could not allocate memory for the management command response. Data: None. Severity: Error Log: LOG_LINK_EVENT Action: None required.
1246: FCoE chip is running golden firmware. Update FCoE chip firmware immediately <fw_type> The FCoE is running the golden firmware. Data: (1) firmware-type Severity: Error Log: LOG_LINK_EVENT Action: Try resetting the FCoE to operational mode and disable maintenance mode.
1247: FCoE chip is running diagnostic firmware. Operational use suspended. <fw_type> The FCoE is running a diagnostic. DATA:(1) firmware-type Severity: Error Log: LOG_LINK_EVENT Action: Try resetting the FCoE to operational mode.
1248: FCoE chip is running unknown firmware. <fw_type> The FCoE is running an unknown firmware version. Data: (1) firmware-type Severity: Error Log: LOG_LINK_EVENT Action: Try resetting the FCoE to operational mode. Try loading latest FCoE firmware.
1249: Invalid FRU data found on adapter. Return adapter to Emulex for repair. The FRU data on the FCoE chip is invalid. Data: (1) firmware-type Severity: Error Log: LOG_LINK_EVENT Action: Try resetting the FCoE to operational mode. Try loading latest FCoE firmware or send the HBA back to Broadcom for repair.
1250: Menlo command error. code=<#> The IOCB driver sent to check FCoE state has bad header size. Data: (1) return code Log: LOG_LINK_EVENT Severity: Error Action: Try resetting the FCoE to operational mode.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1251: Menlo command error. code=<#> The IOCB driver sent to check FCoE state has failed, no resources. Data: (1) return code Log: LOG_LINK_EVENT Severity: Error Action: Try resetting the FCoE to operational mode.
1252: Menlo command error. code=<#> The IOCB driver sent to check FCoE state has failed. Data: (1) return code Log: LOG_LINK_EVENT Severity: Error Action: Try resetting the FCoE to operational mode.
1257: lpfc_menlo_issue_iocb: handler set for <context3>. Data: None Log: LOG_LIBDFC Severity: Warning Action: None required.
1259: mbox: Issued mailbox cmd <u.mb.mbxCommand> while in stopped state. Only the dump mailbox command and reset adapter mailbox command are allowed when in the stopped state. Data: None Severity: Warning Log: LOG_MBOX Action: None required.
1262: Failed to allocate dfc_host Could not allocate memory the dfc_host_struct. Data: None Log: LOG_LIBDFC Severity: Error Action: None required.
1268: Find ndlp returned NULL for oxid:x%x SID:x%x, oxid, sid.(int)off, rc. Could not find the node for this DID. Data: None Severity: Warning Log: LOG_ELS Action: None required.
1302: Invalid speed for this board: Reset link speed to auto: <cfg_link_speed> The driver is re-initializing the link speed to auto-detect. Data: None Severity: Warning Log: LOG_LINK_EVENT verbose Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1303: Link Up Event <eventTag> received A link up event was received. It is possible for multiple link events to be received together. Data: (1) fc_eventTag (2) granted_AL_PA (3) ULnkSpeed (4) alpa_map[0] Detail: If link events received, log (1) last event number received, (2) ALPA granted, (3) Link speed (4) number of entries in the loop init LILP ALPA map. An ALPA map message is also recorded if LINK_EVENT verbose mode is set. Each ALPA map message contains 16 AL_PAs. Severity: Error Log: Always Action: If numerous link events are occurring, check the physical connections to the FC network.
1304: Link Up Event ALPA map A link up event was received. Data: (1) wd1, (2) wd2, (3) wd3, (4) wd4 Severity: Warning Log: LOG_LINK_EVENT verbose Action: If numerous link events are occurring, check the physical connections to the FC network.
1305: Link Down Event <eventTag> received A link down event was received. Data: (1) fc_eventTag (2) hba_state (3) fc_flag Severity: Error Log: Always Action: If numerous link events are occurring, check the physical connections to the FC network.
1306: Link Up Event in loop back mode x%x received Data: x%x x%x x%x x%x Link up notification; configured for loopback. Data: (1) fc_eventTag (2) granted_AL_PA (3) ULnkSpeed (4) alpa_map[0] Severity: Error Log: LOG_LINK_EVENT Action: None required.
1308: Menlo Maint Mode Link up Event x%x rcvd Data: x%x x%x x%x Link down notification; configured for loopback. Data: (1) fc_eventTag (2) port_state (3) vport fc_flag Severity: Error Log: LOG_LINK_EVENT Action: None required.
1309: Link Up Event npiv not supported in loop topology NPIV is not supported in loop topology. Data: None Severity: Error Log: LOG_LINK_EVENT Action: None required.
1310: Menlo Maint Mode Link up Event <eventTag> rcvd The link is up in maintenance mode; only management commands are allowed. Data: (1) fc_eventTag (2) port_state (3) vport fc_flag Severity: Error Log: LOG_LINK_EVENT Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1312: Link Down Event <eventTag> received Maintenance mode link up notification received without entering link down. Data: (1) fc_eventTag (2) port_state (3) vport fc_flag Severity: Error Log: LOG_LINK_EVENT Action: None required.
1400: Failed to initialize sgl list. Failed to initialize SGL list during initialization. Data: None Severity: Error Log: LOG_INIT Action: None required.
1401: Failed to enable pci device. Failed to enable PCI device during initialization. Data: None Severity: Error Log: LOG_INIT Action: None required.
1402: Failed to set up pci memory space. PCI initialization failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
1403: Failed to set up driver resource. Driver resource initialization failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
1404: Failed to set up driver resource. Driver resource initialization failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
1405: Failed to initialize iocb list. Driver resource initialization failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
1406: Failed to set up driver resource. Initialization failed to set up driver resource. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1407: Failed to create scsi host. Initialization failed to create SCSI host. Data: None Severity: Error Log: LOG_INIT Action: None required.
1408: Failure HBA POST Status: sta_reg=0x%x, ""perr=x%x, sfi=x%x, nip=x%x, ipc=x%x, xrom=x%x, ""dl=x%x, pstatus=x%x\n", sta_reg.word0, bf_get(lpfc_hst_state_perr, &sta_reg), The HBA power-on self-test (POST) has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
1409: Failed to enable pci device. Failed to enable PCI device during initialization. Data: None Severity: Error Log: LOG_INIT Action: None required.
1410: Failed to set up pci memory space. Initialization failed to set up PCI memory space. Data: None Severity: Error Log: LOG_INIT Action: None required.
1411: Failed to set up driver resource. Data: None Severity: Error Log: LOG_INIT Action: None required.
1412: Failed to set up driver resource. Initialization failed to set up driver resource. Data: None Severity: Error Log: LOG_INIT Action: None required.
1413: Failed to initialize iocb list. Initialization failed to initialize the IOCB list. Data: None Severity: Error Log: LOG_INIT Action: None required.
1414: Failed to set up driver resource. Initialization failed to set up driver resource. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1415: Failed to create scsi host. Initialization failed to create SCSI host. Data: None Severity: Error Log: LOG_INIT Action: None required.
1416: Failed to allocate sysfs attr Initialization failed to sysfs attribute. Data: None Severity: Error Log: LOG_INIT Action: None required.
1418: Invalid HBA PCI-device group: <dev_grp> Invalid HBA PCI-device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1419: Invalid HBA PCI-device group: <dev_grp> Invalid HBA PCI-device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1420: Invalid HBA PCI-device group:<dev_grp> Invalid HBA PCI-device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1421: Failed to set up hba Initialization failed to set up the HBA. Data: None Severity: Error Log: LOG_INIT Action: None required.
1422: HBA Unrecoverable error: uerr_lo_reg=<ue lo>, uerr_hi_reg=<ue hi>, online0_reg=<Online0>, online1_reg=<Online1> The HBA has notified the driver that it has encountered an unrecoverable error. Data: None Severity: Error Log: LOG_INIT Action: A dump from the Emulex OneCommand CNA Manager application should be taken. Then, the driver should be unloaded and reloaded.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1423: HBA Unrecoverable error: uerr_lo_reg=<ue lo>, uerr_hi_reg=<ue hi>, online0_reg=<Online0>, online1_reg=<Online1> The HBA has notified the driver that it has encountered an unrecoverable error. Data: None Severity: Error Log: LOG_INIT Action: A dump from the Emulex OneCommand CNA Manager application should be taken. Then, the driver should be unloaded and reloaded.
1424: Invalid PCI device group:<pci_dev_grp> Invalid HBA PCI-device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1425: Invalid PCI device group: <pci_dev_grp> Invalid HBA PCI-device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1426: Invalid PCI device group: <pci_dev_grp> Invalid HBA PCI-device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1427: Invalid PCI device group: <pci_dev_grp> Invalid HBA PCI-device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1428: Invalid PCI device group: <pci_dev_grp> Invalid HBA PCI-device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1429: Invalid PCI device group: <pci_dev_grp> Invalid HBA PCI-device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1430: Failed to initialize sql list Failed to initialize SQL list. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1431: Invalid HBA PCI-device group: <pci_dev_grp> Invalid HBA PCI-device group detected. Data: None Severity: Error Log: LOG_INIT Action: None required.
1432: Failed to initialize rpi headers. Failed to initialize RPI headers. Data: None Severity: Error Log: LOG_INIT Action: None required.
1476: Failed to allocate sysfs attr Failed to allocate sysfs attrIbutes. Data: None Severity: Error Log: LOG_INIT Action: None required.
1477: Failed to set up hba Failed to set up the HBA. Data: None Severity: Error Log: LOG_INIT Action: None required.
1603: Loopback test did not receive expected data length. actual length <len>expected length <full_size>. The loopback test did not receive the same amount of data that it transmitted. Data: None Severity: Error Log: LOG_LIBDFC Action: None required.
1800: Could not issue unreg_vpi Driver attempt to unregister VPI failed. Data: None Severity: Error Log: LOG_VPORT verbose Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.
1801: Create vport work array FAILED: cannot do scsi_host_get The driver was unable to get a reference to a SCSI host. Data: None Severity: Warning Log: LOG_VPORT verbose Action: Software driver warning. If this issue persists, report these errors to Broadcom Technical Support.
1816: FLOGI NPIV supported, response data <port> The fabric reports support for NPIV upon FLOGI. Data: (1) response_multiple_NPort Severity: Warning Log: LOG_VPORT verbose Action: No action needed, informational.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1817: Fabric does not support NPIV - configuring single port mode The fabric reports no support for NPIV upon FLOGI. Data: None Severity: Warning Log: LOG_VPORT verbose Action: No action needed, informational.
1818: VPort failed init, mbxCmd <mailbox command> READ_SPARM mbxStatus <mailbox status> , rc = <status> A pending mailbox command that was issued to initialize the port failed. Data: (1) mbxCommand (2) mbxStatus (3) rc Severity: Error Log: LOG_VPORT verbose Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.
1819: Unrecognized lpfc_sli_mode parameter: <mode> An attempt was made to set the SLI mode to an invalid value. The valid values for the SLI mode are 0, 2, and 3. Data: (1) lpfc_sli_mode Severity: Error Log: LOG_VPORT verbose Action: The lpfc_sli_mode driver parameter setting must be corrected. Valid values are 0, 2, and 3.
1820: Unable to select SLI-3. Not supported by adapter. The HBA is incapable of operating in a given mode. Data: None Severity: Error Log: LOG_VPORT verbose Action: SLI-3 mode is only available on some HBAs. Do not attempt to force the SLI mode to 3 on HBAs that do not support SLI-3 mode. This is an informational message. HBAs that do not support SLI-3 will be configured to run in SLI-2 mode, but it is recommended to use the auto setting (0).
1821: Create VPORT failed. Invalid WWN format The port could not be created due to an invalid WWNN or WWPN format. Data: None Severity: Error Log: LOG_VPORT verbose Action: Provide a valid WWN when creating Vports.
1822: Invalid <name>: <xx: xx: xx: xx: xx: xx: xx: xx> An invalid WWN was used when creating a VPort. Data: (1) type_name (2) wwn[1] (3) wwn[3] (3) wwn[5] (4) wwn[7] Severity: Error Log: LOG_VPORT verbose Action: When creating a VPort you must furnish a valid WWN.
1823: Create VPORT failed. Duplicate WWN on HBA. The port could not be created because it would duplicate an existing WWNN HBA address. The resources for the port had to be discarded. Data: None Severity: Error Log: LOG_VPORT verbose Action: Provide a WWN that is unique.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1824: NPIV enabled: Override lpfc_sli_mode parameter (<mode>) to auto(0) The lpfc_enable_npiv and lpfc_sli_mode driver parameter settings conflict. The HBA must be configured for SLI-3 mode to support NPIV. Data: (1) lpfc_sli_mode Severity: Error Log: LOG_VPORT verbose Action: This is an informational message that indicates that the lpfc_enable_npiv and lpfc_sli_mode parameter settings are not compatible. Resolve the parameter conflict by setting the SLI mode to 0 or 3 or, if SLI-2 mode is required then disable NPIV.
1825: Vport Created. This message is displayed to indicate that a port was created in the system. It is displayed at this level to ensure it is always appears at all log levels. Data: None Severity: Error Log: LOG_VPORT verbose Action: No action, informational.
1826: Vport Disabled. The port had to be disabled in the system. Data: None Severity: Error Log: LOG_VPORT verbose Action: No action, informational.
1827: Vport Enabled The port had to be enabled after possible recovery from some errors. Data: None Severity: Error Log: LOG_VPORT verbose Action: No action, informational.
1828: Vport Deleted A Vport was deleted. Data: None Severity: Error Log: LOG_VPORT verbose Action: No action, informational.
1829: CT command failed to delete objects on fabric. A command issued to the fabric to delete an associated resource for an object, such as for a port, failed. Data: None Severity: Error Log: LOG_VPORT verbose Action: Software driver error. If this issue persists, report these errors to Broadcom Technical Support.
1830: Signal aborted mbxCmd <command> A pending mailbox command was aborted because the thread received a signal. Data: None Severity: Error Log: LOG_VPORT verbose Action: You should retry the attempted command.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

1831: Create VPORT Interrupted The port creation process was unexpectedly interrupted at a critical time and the operation was unsuccessful. Data: None Severity: Error Log: LOG_VPORT verbose Action: The process was interrupted while creating a VPort. Retry the command.
1832: No pending MBOX command to handle Data: None Severity: Error Log: LOG_MBOX Action: None required.
1835: Vport discovery quiesce failed: state <port_state> fc_flags <fc_flag> wait msecs <jiffies_to_msecs(jiffies - start_time)> Could not pause discovery on this VPort. Data: None Severity: Error Log: LOG_VPORT Action: None required.
1836: Could not issue unreg_login(all_rpis) status <rc> The unreg_login cannot be issued. Data: None Severity: Error Log: LOG_MBOX, LOG_VPORT Action: None required.
1837: Vport_delete failed: Cannot delete static vport Static VPorts cannot be deleted. Data: None Severity: Error Log: LOG_VPORT Action: None required.
1838: Failed to INIT_VPI on vpi <vpi> status <rc> Failed to INIT_VPI. Data: None Severity: Error Log: LOG_VPORT Action: None required.
2000: Failed to allocate mbox for read_FCF cmd Failed to allocate mailbox for READ_FCF command. Data: None Severity: Error Log: LOG_INIT Action: None required.
2001: Unable to allocate memory for issuing SLI_CONFIG_SPECIAL mailbox command Unable to allocate memory for issuing the SLI_CONFIG_SPECIAL mailbox command. Data: None Severity: Error Log: LOG_SLI Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

2002: Error Could not grow rpi count An error occurred because the RPI count could not be increased. Data: None Severity: Error Log: LOG_SLI Action: None required.
2004: Failed to allocate XRI.last XRITAG is <XRI> Max XRI is <MAX_XRI>, Used XRI is <USED_XRI>. All XRIs are in use. Data: None Severity: Warning Log: LOG_SLI Action: None required.
2005: Unable to deregister pages from HBA: <rc> The SGL pages could not be unregistered from the firmware. Data: None Severity: Error Log: LOG_SLI Action: None required.
2007: Only Limited Edition cmd Format supported <iocb.ulpCommand> The SGL pages could not be unregistered from the firmware. Data: None Severity: Error Log: LOG_SLI Action: None required.
2008: Error <rc> posting all rpi headers The RPI headers could not be posted to the firmware. Data: None Severity: Error Log: LOG_SLI Action: None required.
2009: Failed to allocate mbox for ADD_FCF cmd Failed to allocate mailbox for ADD_FCF command. Data: None Severity: Error Log: LOG_INIT Action: None required.
2010: Resume RPI Mailbox failed status <status>, mbxStatus <mbx status> Data: None Severity: Error Log: LOG_SLI Action: None required.
2011: Unable to allocate memory for issuing SLI_CONFIG_SPECIAL mailbox command Data: None Severity: Error Log: LOG_SLI Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

2012: Mailbox failed , mbxCmd <mbx_cmd> READ_CONFIG, mbxStatus <mbx status> The READ_CONFIG mailbox command failed. Data: None Severity: Error Log: LOG_SLI Action: None required.
2013: Could not manually add FCF record 0, status <rc> Could not add Fibre Channel over Ethernet Forwarder (FCF) record to the FCF list. Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
2014: Invalid command <iocb.ulpCommand> The IOCB command is invalid. Data: None Severity: Error Log: LOG_SLI Action: None required.
2015: Invalid CT %x command <iocb.ulpCommand> Invalid Command-Type in the IOCB is not supported. Data: None Severity: Error Log: LOG_SLI Action: None required.
2017: REG_FCFI mbxStatus error <mbx status> HBA state <port_state> The REG_FCFI mailbox command has failed. Data: None Severity: Error Log: LOG_MBOX Action: None required.
2018: REG_VFI mbxStatus error <mbx status> HBA state <port_state> The REG_VFI mailbox command has failed. Data: None Severity: Error Log: LOG_MBOX Action: None required.
2020: Failed to allocate mbox for ADD_FCF cmd Failed to allocate mailbox for ADD_FCF command. Data: None Severity: Error Log: LOG_INIT Action: None required.
2022: VPI Mailbox failed status <status>, mbxStatus <mbxStatus> The INIT VPI mailbox command has failed. Data: None Severity: Error Log: LOG_SLI Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

2401: Failed to allocate memory for ELS XRI management array of size <els_xri_cnt>. Initialization failed to allocate memory for the ELS XRI management array. Data: None Severity: Error Log: LOG_SLI Action: None required.
2500: EQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc> The mailbox command sent to create the event queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2501: CQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc> The mailbox command sent to create the completion queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2502: MQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc> The mailbox command sent to create the mailbox queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2503: WWQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc> The mailbox command sent to create the work queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2504: RQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc> The mailbox command sent to create the receive queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2505: EQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc> The mailbox command sent to delete the event queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2506: CQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc> The mailbox command sent to delete the completion queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

2507: MQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc> The mailbox command sent to delete the mailbox queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2508: WQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc> The mailbox command sent to delete the work queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2509: RQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc> The mailbox command sent to delete the work queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2510: RQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc> The mailbox command sent to delete the work queue has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2511: POST_SGL mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>n The mailbox command sent to post the SGL pages to the firmware has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2512: REMOVE_ALL_SGL_PAGES mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc> The mailbox command sent to delete the SGL pages from the firmware has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2513: POST_SGL_BLOCK mailbox command failed status <shdr_status> add_status <shdr_add_status> mbx status <rc> The mailbox command sent to post the SGL pages to the firmware has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2514: POST_RPI_HDR mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc> The mailbox command sent to post the RPUI header pages to the firmware has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

2515: ADD_FCF_RECORD mailbox failed with status <rc> The mailbox command to add the FCF record has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2516: DEL FCF of default FCF Index failed mbx status <rc>, status <shdr_status> add_status<shdr_add_status> The mailbox command to delete the FCF record has failed. Data: None Severity: Error Log: LOG_SLI Action: None required.
2517: Unregister FCFI command failed status %d, mbxStatus x%x", rc, bf_get(lpfc_mqe_status, &mbox->u.mqe) The driver was unable to unregister the Fibre Channel Forwarder Indicator (FCFI) from the firmware. Data: None Severity: Error Log: LOG_SLI Action: None required.
2518: Requested to send 0 NOP mailbox cmd Data: None Severity: Warning Log: LOG_INIT Action: None required.
2519: Unable to allocate memory for NOP mailbox command Memory allocation for this mailbox command has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2520: NOP mailbox command failed status x%x add_status x%x mbx status x%x, shdr_status, shdr_add_status, rc. The NOP mailbox command has failed. Data: None Severity: Warning Log: LOG_INIT Action: None required.
2521: READ_FCF_RECORD mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx The READ_FCF_RECORD mailbox command failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2523: Allocated DMA memory size (<alloc_len>) is less than the requested DMA memory size (<req_len>) The ADD_FCF_RECORD mailbox command failed to retrieve the length required from the firmware. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

2524: Failed to get the non-embedded SGE virtual address The READ_FCF_RECORD mailbox command could not retrieve the Oracle (formerly Sun) grid engine (SGE) that was requested. Data: None Severity: Error Log: LOG_MBOX Action: None required.
2527: Failed to allocate non-embedded SGE array. Failed to allocate the non-embedded SGE array. Data: None Severity: Error Log: LOG_MBOX Action: None required.
2528: Mailbox command <vpi> cannot issue The mailbox command could not be issued because the mailbox interrupt is disabled. Data: (1) mbxCommand (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
2529: Mailbox command <vpi> cannot issue Data: (1) mbxCommand (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
2530: Mailbox command <vpi> cannot issue The SLI layer in the driver is inactive. Data: (1) mb.mbxCommand (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
2531: Mailbox command <cpi> cannot issue Data: (1) mb.mbxCommand (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
2532: Mailbox command <vpi> (<mbxCommand>) cannot issue The mailbox bootstrap code detected that the SLI layer is active. Data: (1) sli4_mbox_opcode (2) sli_flag,(3) MBX_POLL Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
2533: Mailbox command <vpi> (<mbxCommand>) cannot issue Data: (1) sli4_mbox_opcode (2) sli_flag (3) MBX_NOWAIT Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

2535: Unsupported RQ count. (<entry_count>) The receive queue ring can only be 512, 1024, 2048, or 4096. Data: None Severity: Error Log: LOG_SLI Action: None required.
2536: Unsupported RQ count. (<entry_count>) The receive queue ring can only be 512, 1024, 2048, or 4096. Data: None Severity: Error Log: LOG_SLI Action: None required.
2537: Receive Frame Truncated! The receive unsolicited handler detected a truncated frame. Data: None Severity: Error Log: LOG_SLI Action: None required.
2539: Dropped frame rctl:%s type:%s\n An unsupported frame was received by the port and dropped. Data: (1) rctl_names[fc_hdr->fh_r_ctl], (2) type_names[fc_hdr->fh_type] Severity: Error Log: Always Action: No action needed, informational.
2540: Ring <ring #> handler: unexpected Rctl <fh_rctl> Type <fh_type> The received frame has an unsupported RCTL or FH_TYPE. Data: None Severity: Warning Log: LOG_SLI Action: None required.
2541: Mailbox command <vpi> (<mbxCommand>) cannot issue Data: (1) sli_mbox_opcode (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
2542: Try to issue mailbox command <vpi> (<mbxCommand>) synchronously ahead of async mailbox command queue Attempting to send a synchronous mailbox command ahead of the asynchronous mailbox commands. Data: (1) sli4_mbx_opcode or sli_mbox_opcode, (2) sli_flag, (3) flag Severity: Warning Log: LOG_MBOX, LOG_SLI Action: None required.
2543: Mailbox command <vpi> (<mbxCommand>) cannot issue The mailbox command does not have all of the fields set correctly. Data: (1) sli_mbox_opcode (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

2544: Mailbox command <vpi> (<mbxCommand>) cannot issue The HBA cannot be accessed on the PCI bus. Data: (1) sli_mbox_opcode (2) sli_flag (3) flag Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
2546: New FCF found index <index> tag <event_tag> A new FCF has been found. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.
2547: Read FCF record failed Could not read the FCF record from the firmware. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.
2548: FCF Table full count <count> tag <event_tag> The FCF table is full. Data: None Severity: Error Log: LOG_SLI Action: None required.
2549: FCF disconnected from network index <index> tag <event_tag> The FCF has disconnected from the network. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.
2550: UNREG_FCFI mbxStatus error <u.mb.mbxStatus> HBA state <port_state>. The UNREG_FCFI mailbox command has failed. Data: None Severity: Error Log: LOG_DISCOVERY. LOG_MBOX Action: None required.
2551: UNREG_FCFI mbox allocation failed HBA state <port_state> The allocation for the UNREG_FCFI mailbox command has failed. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
2552: UNREG_FCFI issue mbox failed rc <rc> HBA state <port_state>. The UNREG_FCFI mailbox command has failed. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

2553: lpfc_unregister_unused_fcf failed to read FCF record HBA state. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
2554: Could not allocate memory for fcf record Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
2555: UNREG_VFI mbxStatus error <u.mb.mbxStatus> HBA state <port_state> The UNREG_VFI mailbox command has failed. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
2556: UNREG_VFI mbox allocation failed HBA state <port_state> Could not allocate memory for UNREG_VFI mailbox command. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
2557: UNREG_VFI issue mbox failed rc <rc> HBA state <port_state> Could not issue the UNREG_VFI mailbox command. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
2558: ADD_FCF_RECORD mailbox failed with status<shdr_status> add_status <shdr_add_status> The ADD_FCF_RECORD mailbox command has failed. Data: None Severity: Error Log: LOG_INIT Action: None required.
2559: Block sgl registration required DMA size <reqlen> great than a page. Attempting to register more SGEs with the firmware than can fit in a page. Data: None Severity: Warning Log: LOG_INIT Action: None required.
2560: Failed to allocate mbox cmd memory\n Failed to allocate mailbox command memory. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

2561: Allocated DMA memory size (<alloclen>) is less than the requested DMA memory size (<reqlen>) Could not get the memory required for the number of XRIs that are attempting to be posted. Data: None Severity: Error Log: LOG_INIT Action: None required.
2562: No room left for SCSI XRI allocation max_xri=<sli4_hba.max_cfg_param.max_xri>, els_xri=<els_xri_cnt>n The number of allocated XRIs has reached the max_xri value. Data: None Severity: Error Log: LOG_SLI Action: None required.
2563: Failed to allocate memory for SCSI XRI management array of size <sli4_hba.scsi_xri_max>. Initialization could not allocate memory to hold the XRIs. Data: None Severity: Error Log: LOG_SLI Action: None required.
2564: POST_SGL_BLOCK mailbox command failed status <shdr_status> add_status <shdr_add_status> mbx status <rc> The list of XRI SGEs failed to be registered with the firmware. Data: None Severity: Error Log: LOG_SLI Action: None required.
2566: Failed to allocate table entry Failed to allocate connection table entry. Data: None Severity: Error Log: LOG_INIT Action: None required.
2567: Config region 23 has bad signature The driver was unable to read Config Region 23 because it has an invalid signature. Data: None Severity: Error Log: LOG_INIT Action: None required.
2568: Config region 23 has bad version The driver was unable to read Config Region 23 because it is an invalid version. Data: None Severity: Error Log: LOG_INIT Action: None required.
2569: lpfc_dump_fcoe_param: memory allocation failed Memory allocation has failed. Data: None Severity: Warning Log: LOG_MBOX Action: None required

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

2570: Failed to read FCoE parameters The driver failed to read FCoE parameters. Data: None Severity: Error Log: LOG_MBOX, LOG_INIT Action: None required.
2572: Failed allocate memory for fast-path per-EQ handle array Failed to allocate memory for the fast-path per-EQ handle array. Data: None Severity: Error Log: LOG_INIT Action: None required.
2573: Failed allocate memory for msi-x interrupt vector entries The driver was unable to allocate memory during initialization of the MSI-X interrupt array. Data: None Severity: Error Log: LOG_INIT Action: None required.
2574: Not enough EQs (<sli4_hba.max_cfg_param.max_eq>) from the pci function for supporting FCP EQs (<cfg_fcp_eq_count>) Failed to create the minimum fast-path event queues. Data: None Severity: Error Log: LOG_INIT Action: None required.
2575: Not enough EQs (<max_eq>) from the pci function for supporting the requested FCP EQs (<cfg_fcp_eq_count>), the actual FCP EQs can be supported: <eq_count> The driver was not configured with enough fast-path event queues. Data: None Severity: Warning Log: LOG_INIT Action: None required.
2576: Failed allocate memory for fast-path EQ record array Failed to allocate memory for the fast-path EQ record array. Data: None Severity: Error Log: LOG_INIT Action: None required.
2577: Failed allocate memory for fast-path CQ record array Failed to allocate memory for the fast-path EQ record array. Data: None Severity: Error Log: LOG_INIT Action: None required.
2578: Failed allocate memory for fast-path WQ record array Failed to allocate memory for the fast-path EQ record array. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

2579: Slow-path wqe consume event carries miss-matched qid: wcqe-qid=<wcqe_quid>, sp-qid=<sp_quid> The consumed entry does not have the slow path's queueID. Data: None Severity: Warning Log: LOG_SLI Action: None required.
2580: Fast-path wqe consume event carries miss-matched qid: wcqe-qid=<fcp_wqid>. The consumed entry does not have the fast path's queueID. Data: None Severity: Warning Log: LOG_SLI Action: None required.
2581: Not enough WQs (<sli4_hba.max_cfg_param.max_wq>) from the pci function for supporting FCP WQs (<cfg_fcp_wq_count>) The driver was not configured with the minimum number of fast-path work queues. Data: None Severity: Error Log: LOG_INIT Action: None required.
2582: Not enough WQs (<max_wq>) from the pci function for supporting the requested FCP WQs (<cfg_wq_count>), the actual FCP WQs can be supported: <wq_count> The driver was not configured with enough fast-path work queues. Data: None Severity: Warning Log: LOG_INIT Action: None required.
2593: The FCP EQ count(<cfg_fcp_eq_count>) cannot be greater than the FCP WQ count(<cfg_fcp_wq_count>), limiting the FCP EQ count to <cfg_fcp_wq_count> The fast-path event queue cannot be greater than the fast-path work queue count. Data: None Severity: Warning Log: LOG_INIT Action: None required.
2597: Mailbox command <vpi> (<mbxCommand>) cannot issue Synchronou(2) sli_flag (3) flag Data: None Severity: Error Log: LOG_MBOX, LOG_SLI Action: None required.
2598: Adapter Link is disabled. The adapter link has been disabled. Data: None Severity: Error Log: LOG_INIT Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

2599: Adapter failed to issue DOWN_LINK mbox command rc <rc> The driver was unable to issue the DOWN_LINK mailbox command. Data: None Severity: Error Log: LOG_INIT Action: None required.
2600: lpfc_sli_read_serdes_param failed to allocate mailbox memory Failed to allocate mailbox memory. Data: None Severity: Error Log: LOG_INIT Action: None required.
2605: lpfc_dump_static_vport: memory allocation failed Failed to allocate mailbox memory. Data: None Severity: Error Log: LOG_MBOX Action: None required.
2606: No NPIV Fabric support No NPIV fabric support. Data: None Severity: Error Log: LOG_ELS Action: None required.
2607: Failed to allocate init_vpi mailbox Failed to allocate the INIT_VPI mailbox command. Data: None Severity: Error Log: LOG_MBOX Action: None required.
2608: Failed to issue Init VPI mailbox The driver was unable to send an INIT_VPI mailbox command. Data: None Severity: Error Log: LOG_MBOX Action: None required.
2609: Init VPI mailbox failed <u.mb.mbxStatus> The INIT_VPI mailbox command failed. Data: None Severity: Error Log: LOG_MBOX Action: None required.
2610: UNREG_FCFI mbox allocation failed Failed to allocate mailbox memory. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

2611: UNREG_FCFI issue mbox failed Could not issue the UNREG_FCFI mailbox command. Data: None Severity: Error Log: LOG_DISCOVERY, LOG_MBOX Action: None required.
2619: Config region 23 has bad signature Configuration region 23 has an invalid signature. Data: None Severity: Error Log: LOG_INIT Action: None required.
2620: Config region 23 has bad version Configuration region 23 has an invalid version. Data: None Severity: Error Log: LOG_INIT Action: None required.
2621: Failed to allocate mbox for query firmware config cmd Failed to allocate mailbox memory. Data: None Severity: Error Log: LOG_INIT Action: None required.
2622: Query Firmware Config failed mbx status <rc>, status <shdr_status> add_status <shdr_add_status> Could not read the firmware configuration. Data: None Severity: Error Log: LOG_SLI Action: None required.
2623: FCoE Function not supported by firmware. Function mode = <function_mode>> FCoE is not supported by this firmware. Data: None Severity: Error Log: LOG_SLI Action: Use the Emulex OneCommand CNA Manager application to update to the latest firmware.
2707: Ring <Ring#> handler: Failed to allocate iocb Rctl <fh_rctl> Type <fh_type> received The driver was unable to allocate memory to send a query config mailbox command. Data: None Severity: Error Log: LOG_SLI Action: None required.
2717: CT context array entry [<index>] over-run: oxid:<fh_ox_id>, sid:<fh_SID> All of the array slots to hold buffers that are passed to the application are in use. Data: None Severity: Warning Log: LOG_ELS Action: None required.

Table 12 FCoE Error Log Messages and their Descriptions (Continued)

2718: Clear Virtual Link Received for VPI <index> tag <event_tag> A Clear virtual link was received from the fabric for this VPI. Data: None Severity: Error Log: LOG_DISCOVERY Action: None required.
2719: Invalid response length: tgt <TGT_ID> lun <LUN> cmnd <CMD> rsplen <RSPLEN> The response length for this FCP command is not supported. Data: None Severity: Error Log: LOG_FCP Action: None required.
2721: ndlp null for oxid %x SID %x\n, icmd->ulpContext, dfchba->ct_ctx[tag].SID); The Node value for this SID is not in the node list. Data: None Severity: Warning Log: LOG_ELS Action: None required.
2726: READ_FCF_RECORD Indicates empty FCF table The driver requested the firmware provide a list of FCF entries to connect to and the firmware responded that the FCF table is empty. Data: None Severity: Error Log: LOG_INIT Action: None required.

4.5.2 Ethernet Driver Log Messages

The following section describes how to retrieve and interpret Ethernet log messages.

4.5.2.1 Retrieving Ethernet Driver Log Messages

Ethernet driver messages are logged in the `/var/log/messages` file. This log file is an ASCII text file and can be viewed and searched with a text editor, such as vim. A log file automatically rotates as it gets larger. Rotated log files are named `messages.`*x*, where *x* is an integer.

To search the log file for error messages, at the command prompt, type:

```
# cd /var/log
# vim messages
```

A message is displayed similar to the following:

```
Aug 15 09:57:48 S74 kernel: Invalid MTU requested. Must be between 256 and
8974 bytes
```

4.5.2.2 Ethernet Driver Log Messages and their Descriptions

When reporting an issue with the adapter, check the kernel message log using the `dmesg` command or the `/var/log/messages` file, and report any of these entries that may be present.

There are three types of Ethernet log messages: error, information, and warning.

Ethernet driver warning messages logged by an adapter start with `be2net <BID>`, where <*BID*> is the PCI bus identifier string. For example:

```
be2net 0000:0d:00.1: MTU must be between 256 and 9000 bytes.
```

NOTE In the following table, <*D*>, <*DD*>, or <*DDD*> refers to decimal values that appear in the log messages, and <*S*> refers to strings.

Table 13 Ethernet Driver Log Messages of Type

Ethernet Log Message	Description	Type
`<S> initialization failed`	Initialization of the adapter or allocation of a resource for initializing the driver failed. In most cases, this message is accompanied by a more specific error message. Try rebooting the system after power cycling. If the issue persists, this could be a symptom of a hardware issue or corrupted firmware.	Error
`<S>: Physical link <S>`	This informational message is about a change in physical link status.	Information
`Adapter busy, could not reset FW. Reboot server to activate new FW`	After flashing firmware on the adapter, the adapter is busy activating the new firmware. Reboot the machine for the new firmware to be active.	Error
`Cannot set rx filter flags <XX> Interface is capable of <XX> flags only`	Failed to change RX filter settings on an interface as requested, as interface does not have the capability.	Warning
`Could not create sysfs group`	The creation of the `flash_fw` entry under the `/sys/class/net/eth`<*x*> failed. The driver is fully functional, but you cannot install later firmware versions on the adapter.	Error
`Failed to enable VLAN promiscuous mode`	The firmware command failed to enable VLAN promiscuous mode.	Error
`Could not enable Wake-on-lan`	Enabling Wake-on-LAN on the adapter failed.	Error
`Could not get CRC for <DD> region`	The driver could not get enough information from the adapter to decide whether a region from a firmware image should be flashed. The driver skips updating this section. This is a very unlikely error.	Error
`Could not get <DD> MSI-x vector(s) Using <DD> vector(s)`	Enabled only a subset of MSI-x vectors was requested.	Information
`Could not set PCI DMA Mask`	The operating system call to set the DMA mask failed.	Error
`created <DD> RX queue(s)`	An informational message logging the number of receive rings was created.	Information
`created <DD> TX queue(s)`	An informational message logging number of transmit rings created.	Information
`<DD> VFs are already enabled`	When unloading the driver while a VF is still assigned to the VM, the VFs remain enabled. When the driver is reloaded, this message is logged and indicates the number of VFs that are enabled and active.	Information
`Debug data event - <DDDD>, <DDDD>, <DDDD>, <DDDD>`	Firmware generates these log messages to identify certain infrequent hardware and firmware events and aid further analysis if required. These messages by themselves are not indicative of any problem. Data from these logs might help Broadcom troubleshoot other problems, if any, that are observed along with these messages.	Information
`enabled <D> (out of <DD>) MSI-x vector(s) for NIC`	The number of MSI-x vectors enabled for the NIC function.	Information
`Enable VLAN promiscuous mode Disabling VLAN promiscuous mode`	This informational log message occurs when the driver enables or disables VLAN promiscuous mode.	Information

Table 13 Ethernet Driver Log Messages of Type (Continued)

Ethernet Log Message	Description	Type

`Error detected in the card`
`ERR: sliport status <DDDDDD>`
`ERR: sliport error1 <DDDDDD>`
`ERR: sliport error2 <DDDDDD>`
	An SLIPORT error is detected in the adapter, followed by status and error messages.	Error

`Error in mapping PCI BARs`	Initialization of the adapter failed due to an error while mapping PCI BARs.	Error
`Failed to optimize SR-IOV resources`	When the adapter is in SRIOV capable configuration, the driver failed to redistribute the PF-pool resources among the PF and requested number of VFs.	Error
`FW config: function_mode=<DD>, function_caps=<DD>`	This message indicates the function mode and function capabilities set on the adapter during initialization.	Information
`FW dump deleted successfully`	The previously generated firmware (FW) dump is successfully deleted.	Information
`FW dump initiated successfully`	FW dump is successfully initiated on the adapter.	Information
`FW dump not generated`	Failed to initiate FW dump on the adapter.	Error
`Firmware flashed successfully`	This informational message states that the firmware on the adapter has been updated.	Information
`Firmware load error`	Updating the adapter with new firmware failed. Usually this message is accompanied by a detailed message on the failure.	Error
`Firmware load not allowed (interface is down)`	Flashing firmware failed because the status of the interface is down.	Error
`Firmware update in progress`	A firmware update is in progress on the adapter.	Information
`Flash incomplete. Reset the server` `Download FW image again after reset`	A firmware is older than version 10.0.*xx.xxx* is upgraded to a newer version on an OCe14000-series adapter, reboot the server and flash the firmware a second time.	Error
`Flashing firmware file <filename>`	This informational message states that the firmware in the adapter is being updated with the firmware image in the indicated file.	Information
`Flashing section type <DD> failed`	The driver could not get enough information from the FW UFI image to flash a section in the UFI image, after which the FW download will be aborted.	Error
`FW image size should be multiple of 4`	The FW UFI image size is invalid. Download the correct FW image and try again.	Error
`FW version is <DD.DD.DDD.DDD>`	This is an informational log to log the FW version loaded on the adapter.	Information
`INTx request IRQ failed – err <DDD>`	The request for INTx interrupt registration failed. The driver is nonfunctional if the INTx interrupt cannot be registered.	Error
`Invalid Cookie. FW image may be corrupted`	The firmware image under `/lib/firmware/<`*`filename`*`>` does not have the expected cookie. The firmware in this file will not be flashed. Copy the proper file and try flashing again.	Error
`Invalid digital signature`	Updating the adapter with new firmware failed. The firmware image being updated does not have a valid digital signature.	Error
`invalid mailbox completion`	The driver received an unexpected completion status for a firmware command.	Error

Table 13 Ethernet Driver Log Messages of Type (Continued)

Ethernet Log Message	Description	Type
`Link down event type: <DD>`	The reason that the application-specific integrated circuit (ASIC) signaled the link status as down. Possible values are: ■ 0 = Link down due to reasons other than those listed here. ■ 1 = Link down caused by Dynamic Control channel protocol. ■ 3 = Link down triggered by Virtual NIC configuration (for example, a zero bandwidth is assigned to a VNIC). ■ 4 = Link down caused by Ethernet Pause frame flooding. ■ 5 = Link down due to physical thermal temperature going up.	Information
`Link down/Link up`	This is an informational message about a change in link status.	Information
`LPVID <DD>`	VLAN filter is configured in BIOS in multichannel configuration.	Information
`Mac address assignment failed for VF <DD>` `MAC <S>M set on VF <DD> Failed`	Firmware command fails to add a MAC address to a virtual function during initialization, or when requested.	Error
`MAC address change to <S> failed`	Failed to change to new MAC address as requested.	Warning
`Mac hash table alloc failed`	Hash table to remember learned MAC addresses could not be allocated.	Error
`Max: txqs <DD>, rxqs <DD>, rss <DD>, eqs <DD>, vfs <DD>` `Max: uc-macs <DD> , mc-macs <DD>, vlans <DD>`	An information message logging maximum resources available to the function. TX rings, RX rings, RX rings with RSS capability, number of VFs supported, unicast MACs, multicast MACs and VLAN filters, respectively.	Information
`Memory allocation failure`	The driver could not allocate the memory required for the requested operation.	Error
`Memory allocation failure during GET_MAC_LIST`	Failed to allocate memory to issue GET_MAC_LIST FW command.	Error
`Missing digital signature`	Updating the adapter with new firmware failed. The firmware image being updated does not have a digital signature.	Error
`Module param rx_frag_size must be 2048/4096/8192.` `Using 2048`	An unsupported receive buffer size was passed for the rx_frag_size module parameter. The driver ignores the specified value and uses the default RX buffer size of 2048.	Warning
`MSIx enable failed`	Request for enabling MSIx interrupts registration failed. Driver will then use INTx interrupts.	Warning
`MAC address changed to <S>`	MAC address is changed successfully as requested.	Information
`MSIX request IRQ failed –err <DDD>`	The request for MSI-X interrupt registration failed. The driver will use INTx interrupts.	Warning
`MTU changed from <DDD> to <DDD> bytes`	This is an informational message that the MTU value changed as requested.	Information
`MTU must be between 256 and 9000 bytes`	Request to change the MTU was issued with an invalid MTU value. The request failed and MTU will not be changed.	Warning
`opcode <DDD>-<D> failed: status <DD>-<DD>`	A firmware command with opcode failed with the indicated status code and extended status code.	Error
`PCIe error reporting enabled`	PCIe error reporting is successfully enabled on this function.	Information
`POST timeout; stage=<DD>`	The power-on self test of the adapter failed. This is an indication of a hardware or firmware issue. Try rebooting the system after a reset.	Error
`Previous dump not cleared, not forcing dump`	Initiating dump on the adapter failed, as previous dump is present. Clear the previous dump (Using `ethtool -W eth<x> 2` command) before initiating new dump.	Error
`queue_setup failed`	Firmware command failed to create requested number of queues.	Error
`Setting HW VLAN filtering failed`	Adding a VLAN filter to HW failed.	Error

Table 13 Ethernet Driver Log Messages of Type (Continued)

Ethernet Log Message	Description	Type
SRIOV enable failed	Could not enable SRIOV since the call to enable SRIOV failed.	Error
TX-rate must be between 100 and <DDDD> Mbps	Request to change transmission rate was issued with an invalid TX rate value. Request failed and speed will not be changed.	Error
TX-rate must be a multiple of <DDD> Mbps	Transmission rate on a virtual function should be given as multiples of %1Mb/s of link speed.	Error
TX-rate setting not allowed when link is down	Transmission rate on a virtual function cannot be modified when physical link is down.	Error
Tx-rate setting of <DDD>Mbps on VF<DD> failed : <DDD>	Firmware command failed to change transmission rate as requested.	Error
txq<D>: cleaning <D> pending tx-wrbs	The driver did not get completions for some transmit requests from the adapter while unloading the driver. This usually indicates an issue with the adapter.	Error
Flash image is not compatible with adapter	Requested UFI image is not compatible with the chip on which the flash was requested.	Error
Unknown debug event <DD>	Other than QNQ type debug event, driver logs as unknown for other asynchronous debug events received.	Warning
Unqualified SFP+ detected on <D> from <S> part no: <S>	The SFP module indicated in the message is not qualified or supported by Broadcom.	Information
Unrecoverable Error detected in the adapter Please reboot server to recover UE LOW: <S> bit set UE HIGH: <S> bit set	There is an unrecoverable error detected in the adapter that requires a reboot to recover. Low and high bits set in the data path in which error occurred.	Error
User has aborted FW download	User requested abort when FW download is in progress.	Error
Using profile <DD>	An informational message of profile type currently enabled in the adapter.	Information
VF<DD> has FILTMGMT privilege	VF is provided with FILTMGMT privilege to program MAC/VLAN filters.	Information
VF is not privileged to issue opcode <DD>-<DD>	VF does not have enough privileges to issue opcode mentioned in the log message.	Warning
VFs are assigned to VMs: not disabling VFs	Do not disable virtual functions on the port during driver unload, on which VFs are assigned to guests which are powered ON.	Warning
VF setup failed	Failed to create VFs as FW commands failed to provide required resources.	Error
VLAN <DDDD> config on VF <DD> failed	Firmware command failed to set VLAN filter as requested.	Error
Waiting for FW to be ready after EEH reset	After a PCI EEH reset, wait until firmware becomes ready.	Information
Waiting for POST aborted	Waiting for power-on self-test of the adapter is aborted.	Error
Waiting for POST, <D>s elapsed	This is an informational log which logs the seconds elapsed while waiting for power-on self-test of the adapter.	Information
Disable/re-enable i/f in VM to clear Transparent VLAN tag	After clearing transparent VLAN tagging for a VF, disable and re-enable the VF interface in guest operating system to clear VLAN-tagging for the traffic from guest.	Warning
Cannot disable VFs while they are assigned	Cannot disable VFs on the PF, when any of the VF is assigned to guest. Detach any VFs from guest to disable VFs on that PF.	Error
Invalid FW UFI file	Firmware UFI file is corrupted. Try flashing after copying correct UFI file.	Error

Table 13 Ethernet Driver Log Messages of Type (Continued)

Ethernet Log Message	Description	Type
RSS hash key is longer than <DD> bytes	Requested RSS hash key is longer than 40 bytes. Request for a 40-byte RSS hash key.	Error
Invalid RSS hash key format	Hash key format must be in xx:yy:zz:aa:bb:cc format meaning both the nibbles of a byte should be mentioned even if a nibble is zero.	Error
RSS hash key is too short (<DD> < <DD>)	Requested RSS hash key is less than 40 bytes. Request for only a 40 byte RSS hash key.	Error
<be2net> version is <DD.DD.DDD.DDD>	This is an informational logging of be2net driver version loaded.	Information
MGMT_STATUS_FLASHROM_SAVE_FAILED (0x17) / MGMT_ADDI_STATUS_DIGITAL_SIGNATURE_MISSING(0x57)	The firmware being used for flash is not digitally signed.	Error
MGMT_STATUS_FLASHROM_SAVE_FAILED (0x17) / MGMT_ADDI_STATUS_DIGITAL_SIGNATURE_INVALID (0x56)	The digital signature on the firmware being used for flash cannot be verified.	Error

4.5.3 RoCE Error Log Messages

The following table lists the RoCE error log messages and their descriptions.

Table 14 RoCE Error Log Messages

RoCE Log Message	Description	Type
ocrdma_check_qp_params(<D>) unsupported inline data size=0x<DD> requested ocrdma_check_qp_params(<D>) supported inline data size=0x<DD>	Validation checks during queue pair (QP) create.	Error
ocrdma_check_qp_params(<D>) unsupported send_sge=0x<DD> requested ocrdma_check_qp_params(<D>) supported send_sge=0x<DD>	Validation checks during QP create.	Error
<pci bus info> <hca_name>: <speed> "<model_number> port <port_num>	Prints the information about the RoCE PCI function. For example: 0000:04:00.1 Emulex OneConnect RoCE HCA: 10 Gbps “OneConnect OCe14000” port 1	Information
<pci bus info> ocrdma<D> driver loaded successfully	Driver loaded successfully on the device.	Information
crdma_mbx_cmd() cqe_status=0x<D>, ext_status=0x<DD>	Completion and extended status in case of mailbox errors.	Error
ocrdma is using default service level	Using the default service level.	Information
ocrdma_add_stat: No space in stats buff	Response data for debugfs request has exceeded the stats buffer size.	Error
ocrdma_add() leaving. ret=<D>	Adding device failed with Error = D.	Error
ocrdma_alloc_resources(<D>) error	Failed to allocate driver resources.	Error

Table 14 RoCE Error Log Messages (Continued)

RoCE Log Message	Description	Type
`ocrdma_alloc_stats_mem: stats debugfs mem allocation failed`	Failed to allocate memory for debugfs.	Error
`ocrdma_alloc_stats_mem: stats mbox allocation failed`	Failed to allocate memory for statistics command.	Error
`ocrdma_build_inline_sges() supported_len=0x<L>,unspported len req=0x<LL>`	Failed to build inline SGEs.	Error
`ocrdma_check_qp_params(<D>) Consumer QP cannot use GSI CQs`	Consumer QPs should not use the CQ of GSI QP.	Error
`ocrdma_check_qp_params(<D>) GSI special QPs already created`	Validation checks during QP create.	Error
`ocrdma_check_qp_params(<D>) unsupported recv_sge=0x<DD> requested` `ocrdma_check_qp_params(<D>) supported recv_sge=0x<DD>`	Validation checks during QP create.	Error
`ocrdma_check_qp_params(<D>) unsupported recv_wr=0x<DD> requested` `ocrdma_check_qp_params(<D>) supported recv_wr=0x<DD>`	Validation checks during QP create.	Error
`ocrdma_check_qp_params(<D>) unsupported send_wr=0x<DD> requested` `ocrdma_check_qp_params(<D>) supported send_wr=0x<DD>`	Validation checks during QP create.	Error
`ocrdma_check_qp_params(<D>) unsupported qp type=0x<DD> requested`	Validation checks during QP create.	Error
`ocrdma_check_qp_params(<D>) Userspace can't create special QPs of type=0x<DD>`	Validation checks during QP create.	Error
`ocrdma_copy_cq_uresp(<D>) copy error cqid=0x<DD>`	Failed to copy CQ create response.	Error
`ocrdma_copy_qp_uresp(<D>) user copy error`	Failed to copy the QP create response back to the user.	Error
`ocrdma_create_qp(<D>) error=<DD>`	Failed QP create command.	Error
`ocrdma_dealloc_ucontext_pd(D) Freeing in use pdid=0x<DD>`	Dealloc ucontext requested on a protection domain (PD) that is in use.	Information
`ocrdma_dereg_mr(<D>) fw not responding`	De-registration of (memory region) MR failed because the firmware is not responding.	Information
`ocrdma_dispatch_ibevent () unknown type=0x<D>`	Received unknown event from the hardware.	Error
`ocrdma_dispatch_ibevent: Fatal event received`	Device reported a fatal event.	Error
`ocrdma_get_dma_mr err, invalid access rights`	Invalid access rights while allocating lkey.	Error
`ocrdma_init_hw() status=<D>`	Initialization of the hardware failed.	Error
`ocrdma_init_service_level(): status=<D>`	Failed to get the DCBX configuration from the adapter.	Error

Table 14 RoCE Error Log Messages (Continued)

RoCE Log Message	Description	Type
ocrdma_irq_handler(): Fatal Error, EQ full eq_id = 0x<D>,eqe = 0x<DD>	EQ full detected.	Error
ocrdma_mbx_create_cq (<D>) max_cqe=0x<DD>, requester_cqe=0x<DDD>	Requesting more CQ entries than what the device supports.	Error
ocrdma_mbx_create_qp(<D>) rq_err	Failed to create QP.	Error
ocrdma_mbx_create_qp(<D>) sq_err	Failed to create QP.	Error
ocrdma_mbx_create_srq() req. max_wr=0x<D>	The total number of shared receive queue (SRQ) entries requested is greater than what the device supports.	Error
ocrdma_modify_port(D) invalid_port=0x<DD>	Invalid port specified in Modify port.	Error
ocrdma_modify_qp(<D>) invalid attribute mask=0x<M> specified for qpn=0x<QP> of type=0x<T> old_qps=0x<OS>, new_qps=0x<NS>	Parameter error while trying to modify the QP.	Error
ocrdma_parse_dcbxcfg_rsp(): DCBX state is disabled.	DCBX state is disabled in the adapter.	Information
ocrdma_parse_dcbxcfg_rsp(): pfc is disabled.	PFC is disabled in the adapter.	Information
ocrdma_process_acqe(<D>) invalid evt code=0x<DD>	Invalid event code <*DD*> reported on the device <*D*>.	Error
ocrdma_process_mcqe() cqe for invalid tag0x<D> expected=0x<DD>	Invalid completion tag reported.	Error
ocrdma_query_port(<D>) invalid_port=0x<DD>	Invalid port specified in Query port.	Error
ocrdma_reg_mr() status=<D>	Failed to register MR.	Error
ocrdma_resolve_dmac () fail to resolve mac_addr	Failed to resolve MAC address.	Error
ocrdma_set_create_qp_rq_cmd() req. max_recv_wr=0x<D>	The total number of receive queue (RQ) entries requested is greater than what the device supports.	Error
ocrdma_set_create_qp_sq_cmd:() req. max_send_wr=0x<D>	The total number of send queue (SQ) entries requested is greater than what the device supports.	Error
ocrdma_update_stats: stats mbox failed with status = <D>	The Statistics command failed from the hardware.	Error
ocrdma_update_wc() invalid opcode received = 0x%x	Invalid opcode received from the hardware completion.	Information
ocrdma_wait_mqe_cmpl(<D>) mailbox timeout: fw not responding	Mailbox failed because of timeout.	Error
opcode=0x<D>, subsystem=0x<DD>	Opcode and subsystem IDs of the failed mailbox commands.	Error
Unable to allocate ib device	ib_alloc_device failed.	Error

4.5.4 iSCSI Driver Log Messages

The following section describes how to retrieve and interpret iSCSI log messages.

4.5.4.1 Retrieving iSCSI Driver Error Log Messages

NOTE These error log messages are specific to the open-iscsi be2iscsi driver.

For Linux systems, the iSCSI driver generates log messages to the `/var/log/messages` file. The log file is an ASCII text file and can be viewed and searched with your preferred text editor.

To search the log file for error messages, at the command prompt, type:

```
# cd /var/log
# vim messages
```

Archived files for older logs also exist in the same directory.

You can view the latest recently generated messages for the current system boot by running:

```
#dmesg
```

4.5.4.2 iSCSI Driver Error Log Messages and their Descriptions

Table 15 iSCSI Log Messages and Descriptions

Log Message	Description
`beiscsi_module_init - Unable to register beiscsi transport.`	Driver registration failure.
`beiscsi_module_init - Unable to register beiscsi pci driver.`	Driver registration failure.
`BM_# : mgmt_invalidate_icds could not be submitted`	Driver error messages for error handling. Cannot submit an abort request due to no memory available.
`BM_# : Unsupported fw version`	Driver init error message. This failure indicates that the driver version that is running on the system does not match the version of the firmware flashed on the board. This issue can be addressed by running the installer from the desired version.
`BM_# : hwi_init_controller failed`	Driver init error message. This failure may be due to the firmware not being present or running currently. This failure may also indicate a hardware issue.
`BM_# : beiscsi_dev_probe -Failed in beiscsi_alloc_memory`	Driver init error message. Cannot allocate required memory for driver initialization.
`BM_# : No boot session`	Driver init error message. Informational message indicating this port does not have a boot-able session configured.
`BM_# : EEH error detected`	Driver pci error message. PCI error detected by system.
`BM_# : EEH : State PERM Failure`	Driver pci error message. PCI error cannot be recovered.
`BM_# : AER EEH Resume Failed`	Driver pci error message. Driver failed to resume after PCI error recovery via a chip reset.

Table 15 iSCSI Log Messages and Descriptions (Continued)

Log Message	Description
`beiscsi_ep_connect shost is NULL`	Driver session management error message. This host no longer exists. Indication of system trying to connect to previously configured sessions through a port that no longer exists, or its MAC address has been changed.
`BS_# : PCI_ERROR Recovery`	Driver session management error message. Cannot create session, controller is busy recovering from PCI error.
`BS_# : Failed in beiscsi_open_conn`	Driver session management error message. Some possible reasons for this include: non-existing targets, wrong destination address, and target rejected login.
`BS_# : The Adapter Port state is Down!!!`	Driver session management error message. An attempt was made to login to a target through a port with link down.
`BS_# : upload failed for cid #`	Driver session management error message. Failed to properly disconnect.

Appendix A: iSCSI Driver

The following section describes iSCSI error handling and log messages.

A.1 iSCSI Error Handling

This section describes iSCSI error handling.

A.1.1 Error Handling Using ETO and LDTO Parameters

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

The error handling is triggered under the following conditions:

- Loss of connection to the target due to target or network disconnection at the target.

  If the driver has I/O requests pending with the target and the target becomes unavailable (due to the target going down or failing over, or network issues at the target), the driver queues up the I/O requests internally for a configured period of time. The threshold value of this period is the extended timeout (ETO) value.
- Loss of immediate link to the initiator (such as cable disconnect or port failure).

  The adapter firmware detects and notifies the driver of a loss of the link. When this happens, the driver queues up the I/O requests internally to a configured period of time so that the operating system does not see I/O errors. The threshold value of this period is the link down timeout (LDTO) value.

When the configured ETO or LDTO value is reached, and the initiator is still unable to connect to the target, the driver fails all I/O requests. At this point, I/O errors are seen by the application and operating system.

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

A.1.2 Error Handling Under MPIO and Cluster Configurations

In an multipath I/O (MPIO) or cluster configuration, fault tolerant software is present on the system that makes the iSCSI driver error handling redundant. These configurations also require that I/O errors be reported as soon as they are detected so that the software can fail over to an alternate path or an alternative node as quickly as possible.

When the iSCSI driver runs under these configurations, the error handling implemented in the driver must be turned off by setting the default values of LDTO and ETO to zero. The changes take effect during the next driver load.

A.2 iSCSI Driver Log Messages

This section describes retrieving and interpreting iSCSI log messages.

A.2.1 Retrieving iSCSI Driver Error Log Messages

NOTE These error log messages are specific to the proprietary `be2iscsi` driver.

For Linux systems, the iSCSI driver generates log messages to the `/var/log/messages` file. The log file is an ASCII text file and can be viewed and searched with your preferred text editor.

To search the log file for error messages, at the command prompt, type:

```
# cd /var/log
# vim messages
```

A.2.2 iSCSI Driver Error Log Messages and their Descriptions

All iSCSI driver error log messages are preceded by a `scsi`*X*: prefix (if the SCSI host controller can be determined, with *X* being a number representing the Linux SCSI host controller) and OneConnect iSCSI Driver.

For example:

```
scsi2: OneConnect iSCSI Driver: Rejected IOCTL since buffer size limit
exceeded
scsi2: OneConnect iSCSI Driver: Subsystem / Opcode = 0x12345678 / 0x11223344
scsi2: OneConnect iSCSI Driver: Payload Length = 0x1000
```

NOTE

- Some error log entries may be followed by additional entries that provide further information.
- In the following table, `0x`X refers to a hexadecimal value that appears in the log messages.

The following table lists the iSCSI driver error log messages and their descriptions.

Table 16 iSCSI Driver Error Log Messages

iSCSI Error Log Message	Description
`Did not receive an iSCSI Command window update from Target for at least 25 Secs. Session Handle`	Check for any errors reported at the target. The Emulex iSCSI initiator is only supported with certified targets. Check for software updates at the target vendor's website and the Documents and Downloads area of http://www.broadcom.com. If this fails, contact Broadcom Technical Support.
`Driver version does not match Firmware. Please run Installer.`	This failure indicates that the driver version that is running on the system does not match the version of the firmware flashed on the board. This issue can be addressed by running the installer from the desired version.
`Error during iSCSI offload Session Handle / Firmware Error code`	This may indicate a target is in error or may point to transient network connectivity issues. It may also indicate a firmware error.
`Initialization Failure`	This failure may be due to the firmware not being present or running currently. This failure may also indicate a hardware issue.
`Initialization failure during Power Management Bootup`	This failure may be due to the firmware not being present or running currently. This failure may also indicate a hardware issue.
`Internal API failed during Initialization`	This failure may indicate a low memory condition.

Table 16 iSCSI Driver Error Log Messages (Continued)

iSCSI Error Log Message	Description
`Hardware Initialization Failed. Either Hardware/Firmware is not initialized or is malfunctioning.`	This failure indicates that the hardware has not been initialized or is malfunctioning. This may also indicate that the firmware is not running correctly.
`OSM Hardware Initialization Failure`	This failure indicates that the hardware has not been initialized or is malfunctioning. This may also indicate that the firmware is not running correctly.
`OSM Resource Allocation Failure`	The operating system failed to allocate resources for the device. Check low memory conditions and operating system hardware resource conflicts.
`Received a TMF Abort for an I/O that is not present with the driver.`	This may indicate a slow connection to the target. Check network connectivity to the target for any errors.

`Received invalid iSCSI Command Sequence Number update from Target.`
`Session Handle = 0xX`
`MaxCmdSN = 0xX`
`ExpCmdSN = 0xX`
	Check for any errors reported at the target. The Emulex iSCSI initiator is only supported with certified targets. Check for software updates at the target vendor's website and the Documents and Downloads area of http://www.broadcom.com. If this fails, contact Broadcom Technical Support.

`Received unsupported Task Management Function.` `Task Management Function code = 0xX`	The operating system version is not supported.

`Rejected IOCTL since buffer size limit exceeded.`
`Subsystem / Opcode = 0xX / 0xX`
`Payload Length = 0xX`
	This error may indicate an incorrect configuration option for the iSCSI driver. It may also indicate a low memory condition.

`Unrecoverable Error`
`UE_LOW = 0xX`
`UE_HIGH = 0xX`
`Firmware Line Number = 0xX`
	This may be due to hardware errors or due to unhandled exceptions in the hardware or firmware.

Appendix B: Configuring iSCSI Through DHCP

This section describes configuring iSCSI using DHCP.

B.1 IP Address Reservation

If you are using the DHCP server to obtain an IP address for your iSCSI initiator, set up a reservation. A reservation assigns a specific IP address based on the MAC address of your iSCSI function.

If you do not reserve an IP address through DHCP, you must set the lease length for the iSCSI initiator's IP address to unlimited, which allows the IP address lease not to expire.

B.2 DHCP Option 43 (Vendor-Specific Information)

This section describes the format for the data returned in DHCP option 43. The method and format for specifying the Vendor ID is outside the scope of this document and is not included here. The initiator offers this Vendor ID to the DHCP server to retrieve data in the format as described in Section B.2.2, DHCP Option 43 Format.

B.2.1 DHCP Option 43 Parameter Descriptions

The following table describes the parameters used in the format (data string) for option 43.

Table 17 DHCP Option 43 (Vendor-Specific Information)

Parameter	Description	Field Type
<AuthenticationType>	If applicable, replace with D, E, or M. ■ D denotes that authentication is disabled. ■ E denotes that one-way CHAP is enabled (the user name and secret to be used for one-way CHAP must be specified by non-DHCP means). ■ M denotes that mutual CHAP is enabled (the user name and passwords required for mutual CHAP authentication must be specified by non-DHCP means). If a value is not specified, this field defaults to authentication disabled.	Optional
<DataDigest>	Replace with either E or D. ■ E denotes that the data digest is enabled. ■ D denotes that the data digest is disabled. If a value is not provided, it is assumed that the Data Digest is disabled by default.	Optional
<HeaderDigest>	Replace with either E or D. ■ E denotes that the header digest is enabled. ■ D denotes that the header digest is disabled.	Optional
<InitiatorName>	Replace with a valid initiator iSCSI qualified name (IQN) of up to 223 characters. If a value is not provided, the default initiator name (generated by the OneConnect adapter based on the board's MAC address) is used.	Optional
<LUN>	A hexadecimal representation of the logical unit number (LUN) of the boot device. Replace with an 8-byte number that should be specified as a hexadecimal number consisting of 16 digits, with an appropriate number of zeros padded to the left, if required. If a value is not provided, LUN 0 is assumed to be the boot LUN.	Optional

Table 17 DHCP Option 43 (Vendor-Specific Information) (Continued)

Parameter	Description	Field Type
<TargetIP>	Replace with a valid IPv4 address in dotted decimal notation.	Mandatory
<TargetName>	Replace with a valid target iSCSI Qualified Name (IQN) name of up to 223 characters.	Mandatory
<TargetTCPPort>	Replace with a decimal number ranging from 1 to 65535 (inclusive). The default TCP port (3260) is assumed if a value is not specified.	Optional

B.2.2 DHCP Option 43 Format

The following is the format of DHCP option 43 and its guidelines for creating the data string:

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

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

B.2.2.1 Default Initiator Name and Data Digest Settings Example

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

```
iscsi:"192.168.0.2":"3261":"000000000000000E":"iqn.2009-4.com:1234567890"::
"E"::"E"
```

In this example, the field values are:

- TargetIP: 192.168.0.2
- TargetTCPPort: 3261
- LUN: 0x0E
- TargetName: iqn.2009-04.com:1234567890
- InitiatorName: Not specified. Use the Initiator name already configured. Use the default name if none was configured.
- HeaderDigest: Enabled
- DataDigest: Not specified. Assume disabled.
- AuthenticationType: One-way CHAP is enabled.

B.2.2.2 Default TCP and Mutual CHAP Settings Example

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

```
iscsi:"192.168.0.2"::"000000000000000E":"iqn.2009-4.com:1234567890"::"E":"D
":"M"
```

In this example, the field values are:

- TargetIP: 192.168.0.2
- TargetTCPPort: Use default from RFC 3720 (3260).
- LUN: 0x0E
- TargetName: iqn.2009-04.com:1234567890
- InitiatorName: Not specified. Use the Initiator name already configured. Use the default name if none was configured.

- HeaderDigest: Enabled
- DataDigest: Disabled
- AuthenticationType: Mutual CHAP is enabled.

Appendix C: OneConnect 10GbE Adapter Port Speed Specifications

This section describes how to negotiate port speed on non-mezzanine and mezzanine cards.

C.1 Port Speed Negotiation on Non-Mezzanine Cards

OneConnect 10GbE adapters can support only one Ethernet port speed at a time. On non-mezzanine cards, its preference is 10Gb/s. The type of module used (copper/optical) does not make a difference. If a 10Gb/s module is plugged into one of the ports, the adapter runs at a 10Gb/s speed regardless of its other port's speed, even if I/Os are running on that port. This behavior is an adapter constraint; another adapter can be running on a different speed.

The following table lists negotiated speed specifications per OneConnect 10GbE adapter port connection.

Table 18 OneConnect 10GbE Adapter Negotiated Speed Specifications

Port 0 Speed (Gb/s)	Port 1 Speed (Gb/s)	Port Link Status	OneConnect 10GbE Speed (Gb/s)
10	10	Both ports are link up	10
10	1	Only Port 0 is link up	10
1	10	Only Port 1 is link up	10
1	1	Both ports are link up	1
1	-	Only Port 0 is link up	1
-	1	Only Port 1 is link up	1
10	-	Only Port 0 is link up	10
-	10	Only Port 1 is link up	10

C.2 Port Speed on Mezzanine Cards

For mezzanine cards, only one Ethernet port speed is supported at a time. Its speed is the first negotiated speed (either 1GbE or 10GbE), depending on the switch that is connected.

To change the speed on these cards:

1. Remove the switch from both of the ports.
2. Insert the switch for one port, and wait for the link to come up.

The mezzanine card retains this speed until both links are down.

Appendix D: Updating Ethernet Firmware

NOTE Driver installation is required before performing a firmware update. For more information on installing the driver, see Section 2.3.1, Installing the Ethernet Driver Kit.

NOTE Starting with firmware version 11.0, firmware supports secure firmware update by digitally signing firmware images and verifying authenticity before updating the image to flash memory.

The Emulex Ethernet driver supports updating the firmware image in the adapter flash through the `request_firmware` interface in Linux. You can perform this update when the adapter is online and passing network and storage traffic.

To update the Ethernet firmware image:

1. Copy the latest firmware image under the `/lib/firmware` directory:
   ```
   # cp  be3flash.ufi /lib/firmware
   ```
2. Start the update process.
   - In Linux distributions that support the flash option in `ethtool` (for example, SLES 11 SP*x* and RHEL 6.*x*), use the following command:
     ```
     # ethtool -f eth<X>  be3flash.ufi
     ```
   - In earlier Linux distributions (for example, RHEL 5.*x*), write the name of the flash image file in the sysfs node:
     ```
     # echo 60 > /sys/class/firmware/timeout
     # echo be3flash.ufi > /sys/class/net/eth<X>/flash_fw
     ```
3. Reboot the system to enable the new firmware image to take effect.

D.1 Ethtool Support for Secure Firmware Update

If a secure version of firmware (version 11.0 or later) is installed on an Emulex OCe14000B-series adapter and you want to update to an earlier unsecured version of firmware (version 10.6 or earlier), you must remove the secure firmware jumper block before performing the update.

Appendix E: Ethtool -S Option Statistics

The following table contains a list of `ethtool -S` option statistics and their descriptions.

Table 19 Ethtool -S Option Statistics

Name	Description
`dma_map_errors`	The number of packets dropped due to DMA mapping errors.
`eth_red_drops`	Received packets dropped due to the ASIC's Random Early Drop policy.
`forwarded_packets`	The number of packets generated by the ASIC internally. These packets are not handed to the host. This counter is shared across ports and all functions (NIC/FCoE/iSCSI).
`jabber_events`	The number of jabber packets received. Jabber packets are packets that are longer than the maximum size Ethernet frames and that have bad CRC.
`link_down_reason`	The reason that the ASIC signaled the link status as down. The various values are: ■ 0 – Link down due to reasons other than those listed here. ■ 1 – Link down caused by Dynamic Control channel protocol. ■ 3 - Link down triggered by Virtual NIC configuration (for example: zero bandwidth assigned to a VNIC). ■ 4 – Link down caused by Ethernet Pause frame flooding. ■ 5 – Link down due to physical thermal temperature going up.
`pmem_fifo_overflow_drop`	Received packets dropped when an internal FIFO going into main packet buffer tank (PMEM) overflows.
`rx_address_filtered`	Received packets dropped when they do not pass the unicast or multicast address filtering.
`rx_alignment_symbol_errors`	The number of packets dropped due to L1 alignment errors. This counter is on a per-port basis.
`rx_compl_err`	The number of RX completion errors received.
`rx_control_frames`	The number of control frames received.
`rx_crc_errors`	The number of packets dropped due to CRC errors.
`rx_dropped_header_too_small`	Received packets dropped when the IP header length field is less than 5.
`rx_dropped_runt`	Dropped receive packets due to runt packets (for example, packets shorter than the Ethernet standard).
`rx_dropped_tcp_length`	Received packets dropped when the TCP header length field is less than 5 or the TCP header length + IP header length is more than IP packet length.
`rx_dropped_too_short`	Received packets dropped when IP length field is greater than the actual packet length.
`rx_dropped_too_small`	Received packets dropped when IP packet length field is less than the IP header length field.
`rx_drops_mtu`	Received packets dropped when the frame length is more than 9018 bytes.
`rx_drops_no_erx_descr`	Received packets dropped due to the input receive buffer descriptor FIFO overflowing.
`rx_drops_no_pbuf`	Packets dropped due to lack of available HW packet buffers used to temporarily hold the received packets.
`rx_drops_no_tpre_descr`	Packets dropped because the internal FIFO to the offloaded TCP receive processing block is full. This could happen only for offloaded iSCSI or FCoE traffic.
`rx_drops_too_many_frags`	Received packets dropped when they need more than 8 receive buffers. This counter will always be 0.
`rx_frame_too_long`	Received packets dropped when they are longer than 9216 bytes.
`rx_in_range_errors`	Received packets dropped when the Ethernet length field is not equal to the actual Ethernet data length.
`rx_ip_checksum_errs`, `rx_tcp_checksum_errs`, `rx_udp_checksum_errs`	Packets dropped due to TCP/IP/UDP checksum errors.

Table 19 Ethtool -S Option Statistics (Continued)

Name	Description
rx_out_range_errors	Received packets dropped when their length field is >= 1501 bytes and <= 1535 bytes.
rx_pause_frames	The number of Ethernet pause frames (flow control) received.
rx_priority_pause_frames	The number of Ethernet priority pause frames (priority flow control) received per port.
rx_switched_unicast_packets, rx_switched_multicast_packets, rx_switched_broadcast_packets	The number of unicast, multicast, and broadcast packets switched internally.
rxpp_fifo_overflow_drop, rx_input_fifo_overflow_drop	Number of received packets dropped when a FIFO for descriptors going into the packet demux block overflows. In normal operation, this FIFO must never overflow.
tx_controlframes	The number of Ethernet control frames transmitted per port.
tx_dma_err	The number of errors occurred in the DMA operation associated with the transmit request from the host to the device.
tx_hdr_parse_err	The number of errors while parsing the packet header of a transmit request.
tx_internal_parity_err	The number of parity errors in the transmit request.
tx_pauseframes	The number of Ethernet pause frames (flow control) transmitted per port.
tx_priority_pauseframes	The number of Ethernet priority pause frames transmitted per port.
tx_qinq_err	The number of transmit requests with Q-in-Q style VLAN tagging, when such tagging is not expected on the outgoing interface.
tx_spoof_check_err	The number of spoof TX request failures, when MAC or VLAN spoof checking is enabled on the interface.
tx_tso_err	The number of transmit request errors, while performing TSO offload.

The following table contains a list of transmit/receive statistics per receive queue basis.

Table 20 Transmit/Receive Queue Statistics

Statistic	Description
rxq<x>:rx_bytes	The number bytes received by the driver.
rxq<x>:rx_pkts	The number of packets received by the driver.
rxq<x>:rx_compl	The number of receive completions signaled to the driver by the ASIC.
rxq<x>:rx_mcast_pkts	The number of multicast packets received by the driver.
rxq<x>:rx_post_fail	The number of times the driver could not post received buffers to the ASIC.
rxq<x>:rx_drops_no_skbs	The number of times the driver could not allocate socket buffers.
rxq<x>:rx_drops_no_frags	Packets dropped due to insufficient buffers posted by the driver.
txq<x>:tx_compl	The number of transmit completions signaled by the ASIC.
txq<x>:tx_bytes	The number of bytes transmitted by the driver.
txq<x>:tx_pkts	The number of packets transmitted by the driver.
txq<x>:tx_reqs	The number of transmit request generated by the driver.
txq<x>:tx_stops	The number of times the driver requests the host to stop giving further transmit requests since the hardware transmit queue is filled up.
txq<x>:tx_drv_drops	The number of transmit packets dropped by the driver.

Appendix F: brcmfcoe Driver BlockGuard Functionality

This section describes how to enable BlockGuard and set `brcmfcoe` driver module parameters.

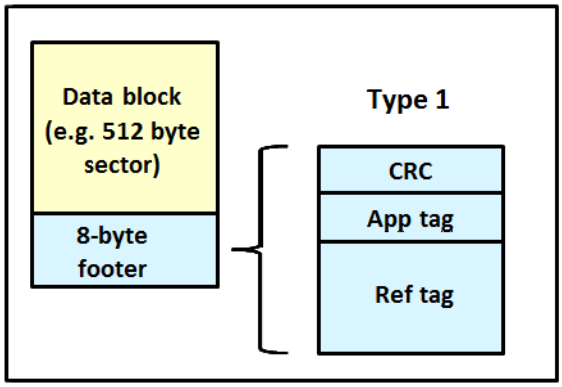
F.1 Overview

The BlockGuard feature checks the integrity of data read from and data written to the host to the disk and back through the SAN. This check is implemented through the Data Integrity Field (DIF) defined in the American National Standards Institute (ANSI) T10 standard.

The Emulex `brcmfcoe` driver supports T10 DIF Type 1. In the Type 1 implementation, the 8-byte DIF consists of a Ref Tag (or logical block address [LBA]), an App Tag, and a Guard Tag (or CRC). A Type 1 DIF is defined as a having a 2-byte Guard Tag, a 2-byte App tag, and a 4-byte Ref tag, which consists of the lower 32 bits of the logical block address.

The following figure shows a data block (with a 512-byte sector) with the 8-byte footer attached to the end. The contents of the 8-byte footer are shown with the fields that make up the Type 1 DIF: the Guard Tag, the App Tag, and the Ref Tag. The App Tag is not used by the lpfc driver.

Figure 1 Data Block Showing Type 1 DIF



When data is written, the DIF is generated by the host, or by the adapter, based on the block data and the logical block address. The DIF field is added to the end of each data block, and the data is sent through the SAN to the storage target. The storage target validates the CRC and Ref tag and, if correct, stores both the data block and DIF on the physical media. If the CRC does not match the data, then the data was corrupted during the write. A Check Condition is returned back to the host with the appropriate error code. The host records the error and retransmits the data to the target. In this way, data corruption is detected immediately on a write and never committed to the physical media. On a read, the DIF is returned along with the data block to the host, which validates the CRC and Ref tags. Because the hardware performs this validation, it adds a very small amount of latency to the I/O.

The format of the Guard Tag can optionally be an IP Checksum instead of the CRC mandated by T10 DIF, which can be beneficial because the Initiator Host uses less CPU overhead to generate an IP Checksum than it does with a CRC. The IP Checksum is typically passed as the Guard Tag between the Initiator Host and the adapter. The adapter hardware will translate the IP Checksum into a CRC, or visa versa, on data being sent to or received from on the wire. The CRC is called a DIF protection type, and the IP Checksum is referred to as DIX protection type.

F.2 Enabling BlockGuard

BlockGuard is disabled by default. To enable it, pass the parameter `lpfc_enable_bg` to the driver as follows:

```
insmod brcmfcoe.ko lpfc_enable_bg=1
```

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

```
options brcmfcoe lpfc_enable_bg=1
```

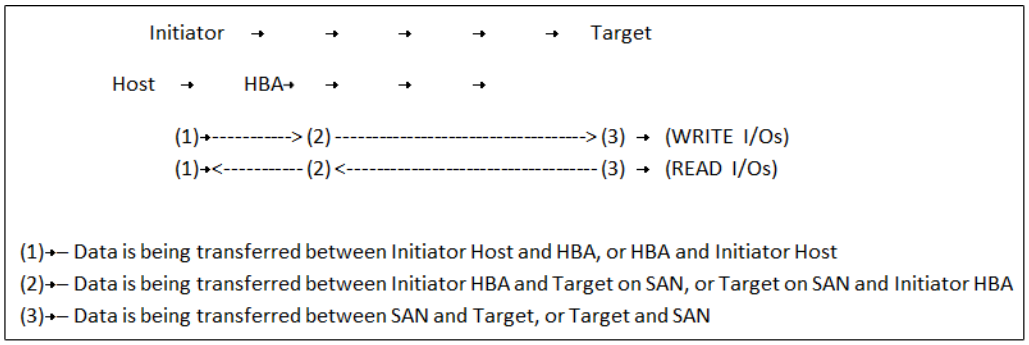
Additional module parameters may be added to this line, separated by spaces.

F.3 SCSI Command Support – SCSI Operation Codes

When both READ and WRITE requests exist, the command descriptor block (CDB) passed to the adapter from the Initiator Host has a read protect/write protect (RDPROTECT/WRPROTECT) field that indicates to the target whether to perform data integrity verification. It also indicates whether to transfer protection data between initiator and target. The adapter does not know if a target supports protection information or with which type of protection it is formatted. The Initiator Host, which has this knowledge, will always prepare a CDB with the appropriate RDPROTECT/WRPROTECT information, depending on target format and capabilities. The request will also include information about with which protection type the target has been formatted.

In addition, the Initiator Host will also provide the adapter with an operation code that tells the controller how to place the protection data for the type of I/O to perform. Each I/O is logically a two-step process. The data is transferred between the Initiator Host and the adapter (over the PCI bus) and between the adapter and the target (over the SAN) as shown in the following figure. The type of operation defines whether the data transfer has protection data.

Figure 2 Data Transfer between Initiator Host and the Adapter



The initiator operations are listed in the following table.

Table 21 Initiator Operations

Initiator Operation	Initiator Host <-> Adapter	Adapter <-> Target	Comment
NORMAL	Unprotected	Unprotected	Used for unprotected I/O
READ_INSERT	Protected	Unprotected	Reads the data from the target. The adapter then generates the protection data and transfers both data and protection data to the Initiator Host. No protection data is sent on the SAN. The adapter can insert the protection data guard tag as CRC or IP CSUM.
READ_PASS	Protected	Protected	Reads the data and protection data from the target on the SAN. The adapter will verify data integrity and transfer both data and protection data to the Initiator Host. The adapter can convert the protection data guard tag from CRC to IP CSUM.
READ_STRIP	Unprotected	Protected	Reads data and protection data from the target. The adapter will verify data integrity, discard protection data, and only transfers the data to the Initiator Host. It does not send the protection data to the Initiator Host. Protection data is only sent on the SAN.
WRITE_INSERT	Unprotected	Protected	Transfers the data from the Initiator Host. The adapter will then generate protection data and write both the data and protection data to the target. Protection data is only sent on the SAN.
WRITE_PASS	Protected	Protected	Transfers the data and protection data from the Initiator Host to the adapter. The adapter will verify protection data and write both data and protection data to the target on the SAN. The adapter can convert the protection data guard tag from IP CSUM to CRC.
WRITE_STRIP	Protected	Unprotected	Transfers data and protection data from the Initiator Host. The adapter will verify data integrity, discard protection data, and writes only the data to the target. No protection data is sent on the SAN.

F.4 brcmfcoe Driver Module Parameters

The brcmfcoe driver has two module parameters: lpfc_prot_mask and lpfc_prot_guard. Using these parameters, you can control which DIF capabilities the lpfc driver registers with the Linux SCSI subsystem. This, in turn, controls which initiator operations (BlockGuard profiles) are used during I/O operations. These parameters are set up when the driver loads and cannot be changed while the driver is running.

F.4.1 lpfc_prot_mask

This parameter controls the DIF operations that the driver registers with the operating system. The operating system selects an operation to use for each I/O command that matches the adapter DIF capability. The driver indicates its capabilities by the operations it registers with the operating system.

If the parameter is not passed to the driver, the default DIX_TYPE0 is used.

The SCSI layer will typically use the bit masks listed in the following table to determine how to place the protection data associated with I/Os to the SCSI Host. The default value for `lpfc_prot_mask` is to allow all of the options.

Table 22 `lpfc_prot_mask` **Protection Types**

Flag	Value	Indicates	Description
SHOST_DIF_TYPE1_PROTECTION	1	Adapter supports T10 DIF Type 1	Adapter to Target Type 1 protection
SHOST_DIX_TYPE0_PROTECTION	8	Adapter supports DIX Type 0	Host to adapter protection only
SHOST_DIX_TYPE1_PROTECTION	16	Adapter supports DIX Type 1	Host to adapter Type 1 protection

The following table shows how protection data gets placed for each supported profile.

Table 23 Protection Data Placement for Supported Profiles

Flag	Value	BlockGuard Profile	Operation
SHOST_DIF_TYPE1_PROTECTION	1	A1	READ_STRIP/ WRITE_INSERT
SHOST_DIX_TYPE0_PROTECTION	8	AST2	READ_INSERT/ WRITE_STRIP
SHOST_DIX_TYPE1_PROTECTION SHOST_DIF_TYPE1_PROTECTION	17	AST1 / C1	READ_PASS/ WRITE_PASS

F.4.2 lpfc_prot_guard

This parameter correlates to the `SCSI_host_guard_type` of the Linux kernel, which specifies the type of CRC the Linux operating system will pass to the lpfc driver. The following table shows the two guard types: CRC, and IP-CSUM, with values of 0x1 and 0x2, respectively.

Table 24 lpfc_prot_guard Guard Types

Flag	Value	Indicates
SHOST_DIX_GUARD_CRC	1	Adapter supports T10 DIF CRC
SHOST_DIX_GUARD_IP	2	Adapter supports both T10 DIF CRC and IP-CSUM

The default value for `lpfc_prot_guard` is `SHOST_DIX_GUARD_IP`. This defines the format for the guard tag when the data is transferred between the Host and the adapter. When data is transferred on the wire, the protection data guard tag is always translated into a T10 DIF CRC. To override the default, you can pass a module parameter value with either `insmod` or `modprobe`.

The SCSI layer will typically use an IP-CSUM as the method for computing the protection data guard tag because it uses less CPU overhead.

Appendix G: Setting Up and Configuring VXLAN

Virtual eXtensible LAN (VXLAN) technology allows a physical layer 3 network to host multiple logical, or virtual, layer 2 networks, which allows for these key advantages over nonvirtualized networks:

- Expands the number of isolated network segments that can exist on a single wire beyond what VLANs alone can provide (significantly beyond 4096).
- Physical switches no longer have to track virtual machine MAC addresses that reside on VXLAN segments.
- VXLANs are seamless to the virtual machines that reside on them (for example, the virtual machines are unaware of the VXLANs presence, reducing complexity in deployment).

The Emulex OCe14000-series converged network adapters support this technology and also provide offloading capabilities of some of the VXLAN functions, decreasing CPU utilization, and potentially increasing I/O throughput, depending on the configuration of the network. This VXLAN offloading technology is only supported in RHEL 7 and SLES 12.

G.1 Setting up VXLAN Networks Across Two Systems

To set up VXLAN networks across two systems:

1. Install all required Emulex software on both systems, such as the NIC driver and the Emulex OneCommand CNA Manager application. Verify that your NIC is recognized.
2. Configure the desired NIC ports with IP addresses, and verify that they can communicate. Set the Maximum Transmission Unit (MTU) option to the desired size.
3. After the physical network has been set up, set up the VXLAN interface on top of it.
   a. From a console, type the following command to create the VXLAN function:

NOTE The default UDP port is 8472.

- For SLES 12:

```
ip link add <vxlan_name> type vxlan id <vxlan_id> group
<multicast_group> dev <ethernet_interface>
```

- For RHEL 7.0:

```
ip link add <vxlan_name> type vxlan id <vxlan_id> group
<multicast_group> dstport <port> dev <ethernet_interface>
```

- Example for RHEL 7.0

```
ip link add vxlan0 type vxlan id 100 group 239.1.1.1
dstport 0 dev eth1
```

This example creates the VXLAN function and associates it with the parent physical function (PF). The MTU automatically adjusts itself based on the PF MTU size, and the `dstport` option to 0 defaults to the standard port of 8472. For VXLAN functions to talk to each other, the ID and multicast group must match. You can have multiple VXLAN functions on the same port, but they must have separate IDs. You cannot use the same ID multiple times in one system.

   b. Verify that the switch you are using supports multicasting. You might have to configure it to do so. Multicast address ranges can be from 224.0.0.0 to 239.255.255.255. Make sure the range matches up with the one you specified in the VXLAN network earlier.

c. Associate an IP address to the tenant and activate it. Type the following:

```
ip address add <ip/subnet> dev <vxlan_name>
ip link set <vxlan_name> up
```

Example:

```
ip address add 20.0.0.1/24 dev eth1
ip link set vxlan0 up
```

You can now ping between IP addresses from within the VXLAN tenant.

4. For VM functionality, you must link the VXLAN function to a macvtap function and connect that to a VM as a NIC or network bridge. The following is an example of creating the macvtap and attaching it directly to a VM from the KVM hypervisor.

a. From a console, type the following to create a macvtap and link it to your VXLAN function:

```
ip link add link <vxlan_name> name <macvtap_name> type macvtap
```

Example:

```
ip link add link vxlan0 name macvtapvxlan0 type macvtap
```

b. Next, you must give the macvtap a MAC address and enable it:

```
ip link set <macvtap_name> address <MAC_address> up
```

Example:

```
ip link set macvtapvxlan0 address 00:11:22:33:44:55 up
```

c. Install your VM using your desired operating system and shut it down.

d. From the Virtual Machine Manager, right-click your VM and select **open**.

- Select **View and Details**.
- At the bottom, click **Add Hardware**.
- Select **Network** and choose the macvtap interface you created from the Host Device drop-down menu and select **virtio** from the Device Model drop-down.
- Click **Finish** to save the changes.

e. Click the newly created NIC, select **Bridge** from the Source Mode drop-down menu and click **Apply**.

f. Power on your VM. The new NIC is detected and you can configure it as normal.

NOTE You must adjust the MTU size of your virtual NIC to accommodate the MTU size of the VXLAN network. For instance, if the VXLAN network is set to 1450, then you must adjust your virtual NIC MTU to 1450.

g. Complete steps 4a to 4f on both systems. The two VMs can now ping each other.

Appendix H: License Notices

H.1 Open-iSCSI GPLv2 Notice

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

51 Franklin Street, Fifth Floor, Boston, MA 02110-1301, USA

Everyone is permitted to copy and distribute verbatim copies

of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Lesser General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licen is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licens extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You

may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licen cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT

WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.


BROADCOM

Emulex® Poll Mode Driver for OneConnect® Adapters

User Guide

Version 11.2
December 30, 2016

DRVPMD-OCA-UG112

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

Broadcom, the pulse logo, Connecting everything, Avago Technologies, Avago, the A logo, Emulex, and OneConnect are among the trademarks of Broadcom in the United States, certain other countries and/or the EU.

Copyright © 2015–2017 Broadcom. All Rights Reserved.

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

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.

Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Chapter 1: Overview

1.1 Purpose

Broadcom has partnered with 6WIND to provide an Emulex® poll mode driver (PMD) for the OCe14000-series 10Gb and 40Gb Ethernet network adapters and converged network adapters.

The Emulex PMD delivers the power of the data plane development kit (DPDK) to next generation network functions virtualization (NFV) workloads, and enables customers to choose the best networking connectivity for their server deployments.

The Emulex PMD enables Ethernet (oce) adapters to interface with the DPDK, giving customers the flexibility to choose their Ethernet connectivity for standard Intel x86 platforms.

This document provides instructions for installing and configuring the DPDK, the Emulex PMD, and the Emulex SLI User Ready Framework (SURF) application programming interface (API). It also lists limitations and provides operational considerations for the Emulex PMD.

1.2 Features

The Emulex PMD includes the following features:

- Concurrent Internet Small Computer Systems Interface (iSCSI) support
- Universal Multichannel (UMC) configuration support
  - PMD installation is available on all four channels per port provided that the PMD is installed on three or four channels.
  - Four physical functions (PFs) exist per port.
- Single root input/output virtualization (SR-IOV) configuration support
  - On Emulex OCe14000B-series adapters, DPDK SR-IOV is not supported on the last port of the adapter.
    - With receive side scaling (RSS) enabled, up to nine virtual functions (VFs) are supported with one PF in an SR-IOV configuration. In this configuration, the PMD running on the VFs has access to two RSS rings each.
    - With RSS disabled, up to 14 VFs are supported with one PF in an SR-IOV configuration. For more information on specifying command line parameters on the `dpdk_surf.ko` module, see Section 2.3.4, max_rxq_per_port <int>.
    - In a 4-port adapter, up to nine VFs are supported on each of the first three ports. The fourth port does not support any VFs because of limited resources.
    - SR-IOV-enabled systems support up to 14 VFs on each port (does not apply to HP Flex-10 systems).
    - HP Flex-10 systems with SR-IOV-enabled support up to eight VFs on each port.

1.3 References

- *Emulex Drivers for Linux for OneConnect Adapters User Guide.*
  This manual includes SR-IOV configuration information and is available on the Broadcom website.
- Emulex PMD for DPDK program information. Refer to the Broadcom website.
- Open source DPDK project information:
  http://www.dpdk.org/

1.4 Abbreviations

API	application programming Interface
ARI	alternative requester ID interpretation
CPU	central processing unit
DPDK	data plane development kit
LLDP	Link Layer Discovery Protocol
NFV	network functions virtualization
NIC	network interface card
NPAR	NIC partitioning
oce	OneConnect Ethernet
OS	operating system
PF	physical function
PMD	poll mode driver
RHEL	Red Hat Enterprise Linux
RSS	receive side scaling
RX	receive
SR-IOV	single root I/O virtualization
SURF	SLI User Ready Framework
TX	transmit
UMC	universal multichannel
VF	virtual function
VLAN	virtual local area network
VM	virtual machine

Chapter 2: Configuring the Emulex PMD

2.1 Basic Setup

NOTE The Emulex file names in the following table and instructions are examples and may not correspond to the latest file names that may be in your OEM package or on the Broadcom website. For the latest Emulex product numbers, versions, and files, please contact your Broadcom systems engineer.

2.1.1 Download the Required Files

Download the following files:

- `dpdk-1.7.1.tar.gz`, `dpdk-1.8.0.tar.gz`, or `dpdk-2.0`

  You can download these files directly from the dpdk.org website, for example:

  http://www.dpdk.org/browse/dpdk/snapshot/dpdk-1.8.0.tar.gz

  Alternatively, you can access the files by going to http://www.dpdk.org/download and clicking **Other Versions** to see additional available files and formats from the browsing interface.
- Emulex Linux Ethernet (`be2net`) source driver files (`be2net-<build version>.src.rpm`)

  These files are available in your OEM package or at the Broadcom website.
- Emulex firmware (`OneConnect-Flash-11.x.xxx.x-x64.iso`)

2.1.2 Flash the Firmware

To flash the firmware, perform these steps:

1. Flash the Emulex firmware to your OCe14000-series adapter using the Emulex OneConnect® Flash Utility.
2. Configure the OCe14000-series adapter to use the "NIC-Only" profile by using the PXE Boot BIOS.
3. Install Red Hat Enterprise Linux (RHEL) 6.5 64-bit, RHEL 7.0, or RHEL 7.1 64-bit with the **Software Development Workstation** option.
4. Log in as root, and remove the in-box `be2net.ko` driver:

   ```
   rm /lib/modules/2.6.32-431.el6.x86_64/kernel/drivers/net/benet/be2net.ko
   ```
5. Reboot, log in as root, and verify that `be2net.ko` is not loaded:

   ```
   lsmod | grep be2net
   ```

2.1.3 Install Linux Ethernet (be2net) Source Driver Files

To install the Linux Ethernet (`be2net`) source driver files, perform this step:

On the system under test, with the appropriate file (`be2net-<build version>.src.rpm`) downloaded as described in Section 2.1, Basic Setup, install the source driver files using the `rpm` command to a directory of your choice. This directory corresponds to the <`be2net driver directory`> variable in the `export` command line in Section Section 2.2.2, Build the Emulex SURF and PMD Binaries, and the `cp` command line in Section 2.2.3, Copy the Binaries.

2.2 Build and Copy the Binaries

2.2.1 Build the DPDK Binaries

To build the DPDK binaries, perform this step:

On the system under test, with the `dpdk-<version>.tar.gz` file downloaded as described in Section 2.1.1, Download the Required Files, build the DPDK binaries. For example:

```
cd /
cp dpdk-1.8.0.tar.gz .
gunzip dpdk-1.8.0.tar.gz
tar –xvf dpdk-1.8.0.tar
cd dpdk-1.8.0
make config T=x86_64-native-linuxapp-gcc
export RTE_SDK=`pwd`
export RTE_TARGET=build
make
```

2.2.2 Build the Emulex SURF and PMD Binaries

To build the Emulex SURF and PMD Binaries, perform this step:

Use the `Palau_11.x.xxx.x_DPDK_PMD_Internal.zip` file downloaded as described in Section 2.1.1, Download the Required Files, to build the SURF and PMD binaries. For example:

```
cd /
mkdir 11.x.xxx.0
cd 11.x.xxx.0
cp Palau_11.x.xxx.x_DPDK_PMD_Internal.zip .
unzip Palau_11.x.xxx.x_DPDK_PMD_Internal.zip
cd packages
cd DPDK_PMD
cd surf_linux
export ELX_SURF_HUB=`pwd`/surf_hub
export BE2NET=<be2net driver directory>
chmod 777 compile_elx_pmd_dpdk
./compile_elx_pmd_dpdk
```

2.2.3 Copy the Binaries

```
cd /
mkdir DPDK_RH65.218.0
cd /11.x.xxx.0
cd packages
cd DPDK_PMD
cp surf_linux/surf_hub/surf_hub.ko                    /DPDK_RH65.218.0
cp surf_linux/surf_provider/surf_provider.ko          /DPDK_RH65.218.0
cp surf_linux/PMD_DRIVER/librte_pmd_oce.so            /DPDK_RH65.218.0
cp surf_linux/PMD_DRIVER/dpdk_oce_surf/dpdk_surf.ko   /DPDK_RH65.218.0
cp /<be2net driver directory>/be2net.ko               /DPDK_RH65.218.0
cp /dpdk-1.8.0/build/app/testpmd                      /DPDK_RH65.218.0
```

2.3 Parameters for the dpdk_surf.ko Module

When loading the `dpdk_surf.ko` module, the following command line parameters may be specified.

2.3.1 dpdk_oce_trace <int>

The `dpdk_oce_trace` parameter specifies the display of debug traces by the `dpdk_surf.ko` module:

- 0x0 – Module tracing disabled (default value).
- 0x1 – Displays debug messages.
- 0x2 – Displays error messages.
- 0x3 – Displays all messages.

2.3.2 lldp_mode <int>

The `lldp_mode` parameter specifies the behavior of the Link Layer Discovery protocol (LLDP) on the OCe14000-series adapter ports:

- 0x0 – LLDP disabled (default value).
- 0x1 – Enables periodic transmission of LLDP frames.
- 0x2 – Enables processing of input LLDP frames.

2.3.3 max_dpdk_oce_port <int>

The `max_dpdk_oce_port` parameter is the maximum number of OCe14000-series adapter ports to support. The default is 8. The maximum value is 32.

NOTE This parameter must be increased to a number greater than 8 when the total number of `be2net` network interface card (NIC) functions in a system exceed 8. The total number of `be2net` NIC functions present in a system must not exceed 32 to use DPDK.

2.3.4 max_rxq_per_port <int>

The `max_rxq_per_port` parameter specifies the maximum number of receive (RX) queues per port provided to the Emulex PMD. The default value is 0.

NOTE Because of the limited resources on an OCe14000-series adapter, creating a single RSS-enabled RX queue might fail in SR-IOV with more than nine VFs enabled per port. To support 10 to 14 VFs per port, you must request a single non-RSS RX queue by setting `max_rxq_per_port` to 1. This action forces the Emulex PMD to create only one non-RSS RX queue per OCe14000-series adapter port.

2.4 Creating the load_drivers and run_testpmd Scripts

Create the following `load_drivers` and `run_testpmd` scripts to simplify the loading and execution processes.

2.4.1 Creating the load_drivers Script

To create the `load_drivers` script, perform these steps:

1. Gather the Linux Ethernet interface names for the adapter ports:

```
cd /DPDK_RH65.218.0
insmod be2net.ko
ifconfig | grep Ethernet
rmmod be2net.ko
```

2. Create the `/DPDK_RH65.218.0/load_drivers` script, and modify the Linux Ethernet interface names (such as `$INTERFACE_1` and `$INTERFACE_2`) for your environment:

```
#!/bin/bash

# Linux SkyHawk interface names
INTERFACE_1=em1     ## Modify this to match your environment!
INTERFACE_2=em2     ## Modify this to match your environment!

# Directory containing the files
BINARIES=/DPDK_RH65.218.0

# Huge page allocation
mkdir -p /mnt/huge
mount -t hugetlbfs nodev /mnt/huge
echo 128 >
/sys/devices/system/node/node0/hugepages/hugepages-2048kB/nr_hugepages
echo 128 >
/sys/devices/system/node/node1/hugepages/hugepages-2048kB/nr_hugepages
echo 128 >
/sys/devices/system/node/node2/hugepages/hugepages-2048kB/nr_hugepages
echo 128 >
/sys/devices/system/node/node3/hugepages/hugepages-2048kB/nr_hugepages

# Load the be2net driver:
insmod $BINARIES/be2net.ko

# for NPAR configurations add the following delay after loading the be2net
driver:
sleep 2

/* For SR-IOV configurations, install the be2net driver with the required
number of VFs. For RSS support, <num_vfs> is 6 or less. If RSS is disabled,
<num_vfs> is 12 or less. */

/* Check for the file to set sriov_numvfs. */
# find /sys -name sriov_numvfs
/sys/devices/pci0000:00/0000:00:02.2/0000:04:00.0/sriov_numvfs
/sys/devices/pci0000:00/0000:00:02.2/0000:04:00.1/sriov_numvfs
```

```
/* Set the VF number for each PF interface. */
# echo <num_vfs> >
/sys/devices/pci0000:00/0000:00:02.2/0000:04:00.0/sriov_numvfs
# echo <num_vfs>
>/sys/devices/pci0000:00/0000:00:02.2/0000:04:00.1/sriov_numvfs

# Load the SURF hub and provider modules:
insmod $BINARIES/surf_hub.ko
insmod $BINARIES/surf_provider.ko

ip link set $INTERFACE_1 up
ethtool -s $INTERFACE_1 speed 10000 duplex full
ip link set $INTERFACE_2 up
ethtool -s $INTERFACE_2 speed 10000 duplex full

/* Reduce the number of RX & TX queues that the be2net driver uses. This
frees up the queues and allows the PMD to use them */
ethtool -L $INTERFACE_1 combined 1
ethtool -L $INTERFACE_2 combined 1

/*Since all traffic will be handled by the PMD, bring down the NIC driver
interface */
ifconfig $INTERFACE_1 down
ifconfig $INTERFACE_2 down

/* Assign a trivial and unique address to the NIC interfaces to prevent a
MAC address collision in the adapter's tables */
ifconfig $INTERFACE_1 hw ether 00:00:00:00:01
ifconfig $INTERFACE_2 hw ether 00:00:00:00:02

# start the DPDK SURF kernel module with RSS capability enabled by default
# Up to 6 VFs are supported with RSS enabled
insmod $BINARIES/dpdk_surf.ko
# Up to 12 VFs are supported with RSS disabled using the following command:
# insmod $BINARIES/dpdk_surf.ko max_rxq_per_port=1
```

3. Mark the `load_drivers` script executable:

```
chmod 777 /DPDK_RH65.218.0/load_drivers
```

2.4.2 Creating the run_testpmd Script

To create the `run_testpmd` script, perform these steps:

1. Determine the PCI device IDs for your two adapter ports:

```
lspci | grep Emulex
```

2. Create the /DPDK_RH65.218.0/run_testpmd script:

```
#!/bin/bash

BINARIES=/DPDK_RH65.218.0

P0=0000:01:00.00       ## Modify this to match your environment!
P1=0000:01:00.01       ## Modify this to match your environment!

./testpmd -c ff -n 3 -d /$BINARIES/librte_pmd_oce.so -w $P0 -w $P1 -- -i
--nb-cores=2 --nb-ports=2
```

3. Mark the `run_testpmd` script executable:

```
chmod 777 /DPDK_RH65.218.0/run_testpmd
```

2.5 Running a Test

To run a test with the scripts, perform these steps:

1. Verify the correct cabling:
   a. Connect port 0 directly to port 1 on your adapter.
   b. If possible, verify that both link LEDs are on.
2. Generate the DPDK-PMD traffic:

```
cd /DPDK_RH65.218.0
./load_drivers
./run_testpmd
start tx_first
<<< NOTE: Wait for about 10-15 seconds >>>
stop
quit
```

Sample Output

```
EAL: Detected lcore 0 as core 0 on socket 0
EAL: Detected lcore 1 as core 0 on socket 1
EAL: Detected lcore 2 as core 1 on socket 0
EAL: Detected lcore 3 as core 1 on socket 1
EAL: Detected lcore 4 as core 2 on socket 0
EAL: Detected lcore 5 as core 2 on socket 1
EAL: Detected lcore 6 as core 3 on socket 0
EAL: Detected lcore 7 as core 3 on socket 1
EAL: Detected lcore 8 as core 4 on socket 0
EAL: Detected lcore 9 as core 4 on socket 1
EAL: Detected lcore 10 as core 5 on socket 0
EAL: Detected lcore 11 as core 5 on socket 1
EAL: Detected lcore 12 as core 8 on socket 0
EAL: Detected lcore 13 as core 8 on socket 1
```

```
EAL: Detected lcore 14 as core 9 on socket 0
EAL: Detected lcore 15 as core 9 on socket 1
EAL: Detected lcore 16 as core 10 on socket 0
EAL: Detected lcore 17 as core 10 on socket 1
EAL: Detected lcore 18 as core 11 on socket 0
EAL: Detected lcore 19 as core 11 on socket 1
EAL: Detected lcore 20 as core 12 on socket 0
EAL: Detected lcore 21 as core 12 on socket 1
EAL: Detected lcore 22 as core 13 on socket 0
EAL: Detected lcore 23 as core 13 on socket 1
EAL: Support maximum 64 logical core(s) by configuration.
EAL: Detected 24 lcore(s)
EAL: Setting up memory...
EAL: Ask a virtual area of 0xc800000 bytes
EAL: Virtual area found at 0x7f9d18c00000 (size = 0xc800000)
EAL: Ask a virtual area of 0x1000000 bytes
EAL: Virtual area found at 0x7f9d17a00000 (size = 0x1000000)
EAL: Ask a virtual area of 0x2000000 bytes
EAL: Virtual area found at 0x7f9d15800000 (size = 0x2000000)
EAL: Ask a virtual area of 0x400000 bytes
EAL: Virtual area found at 0x7f9d15200000 (size = 0x400000)
EAL: Ask a virtual area of 0x400000 bytes
EAL: Virtual area found at 0x7f9d14c00000 (size = 0x400000)
EAL: Ask a virtual area of 0x1800000 bytes
EAL: Virtual area found at 0x7f9d13200000 (size = 0x1800000)
EAL: Ask a virtual area of 0x400000 bytes
EAL: Virtual area found at 0x7f9d12c00000 (size = 0x400000)
EAL: Ask a virtual area of 0x400000 bytes
EAL: Virtual area found at 0x7f9d12600000 (size = 0x400000)
EAL: Ask a virtual area of 0x400000 bytes
EAL: Virtual area found at 0x7f9d12000000 (size = 0x400000)
EAL: Ask a virtual area of 0x400000 bytes
EAL: Virtual area found at 0x7f9d11a00000 (size = 0x400000)
EAL: Ask a virtual area of 0xe00000 bytes
EAL: Virtual area found at 0x7f9d10a00000 (size = 0xe00000)
EAL: Ask a virtual area of 0x1000000 bytes
EAL: Virtual area found at 0x7f9d0f800000 (size = 0x1000000)
EAL: Ask a virtual area of 0x1c00000 bytes
EAL: Virtual area found at 0x7f9d0da00000 (size = 0x1c00000)
EAL: Ask a virtual area of 0x4800000 bytes
EAL: Virtual area found at 0x7f9d09000000 (size = 0x4800000)
EAL: Ask a virtual area of 0x800000 bytes
EAL: Virtual area found at 0x7f9d08600000 (size = 0x800000)
EAL: Ask a virtual area of 0xc00000 bytes
EAL: Virtual area found at 0x7f9d07800000 (size = 0xc00000)
EAL: Ask a virtual area of 0x400000 bytes
EAL: Virtual area found at 0x7f9d07200000 (size = 0x400000)
EAL: Ask a virtual area of 0x1000000 bytes
EAL: Virtual area found at 0x7f9d06000000 (size = 0x1000000)
EAL: Ask a virtual area of 0x2000000 bytes
EAL: Virtual area found at 0x7f9d03e00000 (size = 0x2000000)
EAL: Ask a virtual area of 0x800000 bytes
EAL: Virtual area found at 0x7f9d03400000 (size = 0x800000)
EAL: Ask a virtual area of 0x400000 bytes
```

```
EAL: Virtual area found at 0x7f9d02e00000 (size = 0x400000)
EAL: Ask a virtual area of 0x200000 bytes
EAL: Virtual area found at 0x7f9d02a00000 (size = 0x200000)
EAL: Requesting 128 pages of size 2MB from socket 0
EAL: Requesting 128 pages of size 2MB from socket 1
EAL: TSC frequency is ~2300000 KHz
EAL: open shared lib //DPDK_RH65.218.0/librte_pmd_oce.so
PMD: librte_pmd_oce by 6WIND registered
EAL: Master core 0 is ready (tid=26419800)
EAL: Core 4 is ready (tid=fffec700)
EAL: Core 5 is ready (tid=ff5eb700)
EAL: Core 6 is ready (tid=febea700)
EAL: Core 7 is ready (tid=fe1e9700)
EAL: Core 3 is ready (tid=9ed700)
EAL: Core 2 is ready (tid=13ee700)
EAL: Core 1 is ready (tid=1def700)
EAL: PCI device 0000:01:00.0 on NUMA socket 0
EAL:   probe driver: 10df:720 librte_pmd_oce
EAL: PCI device 0000:01:00.1 on NUMA socket 0
EAL:   probe driver: 10df:720 librte_pmd_oce
Interactive-mode selected
Configuring Port 0 (socket 0)
Port 0: 00:90:FA:30:97:D6
Configuring Port 1 (socket 0)
Port 1: 00:90:FA:30:97:DA
Checking link statuses...
Port 0 Link Up - speed 10000 Mbps - full-duplex
Port 1 Link Up - speed 10000 Mbps - full-duplex
Done

testpmd> start tx_first

  io packet forwarding - CRC stripping disabled - packets/burst=32
  nb forwarding cores=2 - nb forwarding ports=2
  RX queues=1 - RX desc=128 - RX free threshold=0
  RX threshold registers: pthresh=8 hthresh=8 wthresh=0
  TX queues=1 - TX desc=512 - TX free threshold=0
  TX threshold registers: pthresh=32 hthresh=0 wthresh=0
  TX RS bit threshold=0 - TXQ flags=0x0

<<< NOTE: Wait for about 10-15 seconds >>>
testpmd> stop
Telling cores to stop...
Waiting for lcores to finish...

  ---------------------- Forward statistics for port 0  ----------------------
  RX-packets: 24013794       RX-dropped: 0             RX-total: 24013794
  TX-packets: 24004546       TX-dropped: 0             TX-total: 24004546
  ----------------------------------------------------------------------------

  ---------------------- Forward statistics for port 1  ----------------------
  RX-packets: 24004514       RX-dropped: 0             RX-total: 24004514
  TX-packets: 24013826       TX-dropped: 0             TX-total: 24013826
  ----------------------------------------------------------------------------
```

```
  +++++++++++++++ Accumulated forward statistics for all ports+++++++++++++++
  RX-packets: 48018308       RX-dropped: 0             RX-total: 48018308
  TX-packets: 48018372       TX-dropped: 0             TX-total: 48018372
  ++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++++

Done.
testpmd> quit
Stopping port 0...done
Stopping port 1...done
bye...
```

Chapter 3: Limitations

The following limitations apply when installing and using the Emulex PMD.

3.1 Multichannel Limitations

If you install the Emulex PMD on only one or two channels on the port, it must be installed on the third channel, fourth channel, or both. This installation is necessary because freeing up the resources for the first or second channel does not release the required RSS resources.

3.2 SR-IOV Limitations

The Emulex PMD has the following limitations when using SR-IOV:

- With RSS enabled, up to nine VFs are supported with one PF in an SR-IOV configuration. In this configuration, the PMD running on the VFs has access to two RSS rings each. RSS is enabled by default, and the following command is typically used in the `load_drivers` script:

  ```
  # insmod $BINARIES/dpdk_surf.ko.
  ```

  For additional information, see Section 2.4.1, Creating the load_drivers Script.
- With RSS disabled, up to 14 VFS are supported with one PF. To disable RSS, use the following command in the `load_drivers` script:

  ```
  # insmod $BINARIES/dpdk_surf.ko max_rxq_per_port=1.
  ```

  For additional information, see Section 2.4.1, Creating the load_drivers Script.
- In a 4-port adapter, up to four VFs are supported on each of the first three ports. The fourth port does not support any VFs because of limited resources.
- Alternative requester ID interpretation (ARI) has not been fully qualified. For SR-IOV configuration information, refer to the *Emulex Drivers for Linux for OneConnect Adapters User Guide*.

3.3 RHEL 7.2 with DPDK 2.0 Limitations

RHEL 7.2 does not compile with DPDK 2.0.

Chapter 4: Tuning and Operational Considerations

4.1 Tunable Parameters Using DPDK Applications

DPDK applications include multiple tunable parameters. For details about these parameters, refer to the documentation at the DPDK project site:

http://www.dpdk.org/

4.2 Recommended RX/TX Queues

For best performance, the Emulex PMD requires at least three pairs of receive/transmit (RX/TX) queues.

4.3 Receive and Transmit Tunings

The DPDK API includes a set of RX/TX configuration thresholds that tune the Emulex PMD receive and transmit functions.

4.3.1 Tuning the Receive Function of the Emulex PMD

The Emulex PMD manages the RX free threshold that is supplied in the RX queue configuration data structure at RX queue creation.

The RX Free Threshold parameter drives the notification of the replenished RX queue descriptors to the adapter by the receive function of the PMD that manages the adapter.

The receive function of the PMD notifies only accumulated RX queue entries after their total number is greater than or equal to the RX free threshold, and notifies a number of RX entries that is equal to this threshold when this situation occurs.

The value of the RX free threshold must be a multiple of the number of queue entries per central processing unit (CPU) cache line (eight RX queue entries on Intel CPUs with a 64-byte cache line).

This method minimizes the number of expensive notification operations that perform a 32-bit write access to an adapter doorbell register.

4.3.2 Tuning the Transmit Function of the Emulex PMD

The Emulex PMD manages the following TX configuration parameters supplied in the DPDK API to tune the behavior of the transmit function:

- **TX Completion Threshold** – Can be set with the `--txrst=N` parameter of the testpmd application.
- **TX Free Threshold** – Can be set with the `--txfreet=N` parameter of the testpmd application.
- **TX Write-Back Threshold** – Can be set with the `--txwt=N` parameter of the testpmd application.

4.3.2.1 Transmit Completion Threshold

The TX completion threshold drives the rate at which completed TX queue entries are notified by the adapter.

The transmit function of the Emulex PMD does not systematically set the Completion bit in the TX Queue Header entry (that precedes Queue data entries) of each output packet in the TX Queue. Instead, it only sets the Completion bit when the total number of TX Queue entries supplied to the adapter with their Completion bit unset reaches the value of the TX completion threshold.

The TX completion threshold must meet the following requirements:

- Be less than the number of TQ Queue entries
- Be a multiple of the number of TX Queue entries per CPU cache line (four TX queue entries on Intel CPUs with a 64-byte cache line).

These settings minimize the number of write-back memory accesses performed by the adapter in the TX Completion Queue, and the actual number of Completion Queue entries exchanged for a given output packet rate.

4.3.2.2 Transmit Free Threshold

The configuration of TX queues includes a TX Free Threshold parameter that is the number of used TX queue entries that must be reached before the transmit function of the PMD first checks for so-called "free" TX queue entries whose completion has been notified by the adapter.

The transmit function of the PMD only looks for valid entries in a TX Completion Queue associated with a TX queue when the number of used TX queue entries reaches the value of the TX free threshold.

The value of the TX Free threshold must satisfy the following requirements:

- Be less than the number of TX Queue entries.
- Be a multiple of the number of TX Completion Queue entries per CPU cache line.
- Be a multiple of the TX Completion Threshold times the number of TX Completion Queue entries per CPU cache line.

4.3.2.3 Transmit Write-Back Threshold

The TX write-back threshold drives the rate at which processed TX completion queues entries are notified to the adapter by the transmit function.

The transmit function of the PMD accumulates processed TX Completion Queue entries and notifies them only when their total number is greater than or equal to the TX Write-Back Threshold parameter. The notification includes the threshold value (minimum) when this situation occurs.

The value of the TX Write-Back Threshold parameter must be a multiple of the number of TX Completion Queue entries per CPU cache line (four 16-byte TX Completion entries on Intel CPUs with a 64-byte cache line).

This method minimizes the number of expensive notification operations that perform a 32-bit write access to an adapter doorbell register.

4.4 Disabling VLAN Tag Stripping

OCe14000-series adapters enable the VLAN tag stripping option by default. This option might cause some DPDK applications that expect virtual LAN (VLAN) tags to be included in received packets to not work properly, such as the 6WINDGate application.

This option is a port-wide setting, and disabling it requires the PMD driver to issue a command by a PF on the port. If an application requires VLAN tag stripping to be disabled on a VF, load the PMD on a PF on the port (which will issue the command).

NOTE The Windows operating system expects VLAN tags to be stripped, so disabling the VLAN tag stripping option on a port that is running Windows virtual machines (VMs) causes issues in the Windows VM.

4.5 Disabling LLDP

Some DPDK applications may require Link Layer Discovery Protocol (LLDP) to be disabled. LLDP can only be disabled by a PF, because it is a port-wide setting. The `hbacmd` utility can be invoked on the host with a PF to disable LLDP on the port if it is required by a VM.


BROADCOM

Emulex® Drivers for VMware ESXi for OneConnect® Adapters

User Guide

Version 11.2
December 30, 2016

DRVVM-ESXi-OCA-UG112

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

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

Copyright © 2003–2017 by Broadcom. All Rights Reserved.

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

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.

Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Chapter 1: Introduction

This product supports the Emulex OneConnect family of adapters.

The VMware ESXi drivers support the FC over Ethernet (FCoE), Ethernet or network interface card (NIC), and Internet Small Computer System Interface (iSCSI) protocols.

This document explains how to install the VMware ESXi drivers on your system and how to configure the drivers' capabilities. Below is a partial list of configurable parameters by protocol:

- FCoE
  - Adding logical unit numbers (LUNs) and targets
  - Configuring driver parameters
  - Creating an FC remote boot disk
  - Managing devices through the CIM interface
  - Working with virtual ports (vPorts)
  - Troubleshooting FCoE issues
- NIC
  - Configuring NIC driver parameters
  - Improving driver performance
  - Using vmxnet emulation, including TCP segmentation offload (TSO) and jumbo frames
  - Troubleshooting NIC issues
- iSCSI
  - Configuring iSCSI driver parameters, iSCSI error handling, and link down timeout (LDTO) and extended timeout (ETO)
  - Implementing Multipath I/O (MPIO) support for non-boot and boot targets
  - Configuring iSCSI through Dynamic Host Configuration Protocol (DHCP)
  - Troubleshooting iSCSI issues

Refer to the *Emulex OneCommand CNA Manager for OneConnect Adapters User Guide* and the *Emulex OneCommand CNA Manager for VMware vCenter for OneConnect Adapters User Guide* for complete lists of supported technology.

1.1 ESXi Compatibility

For compatible firmware versions, go to the Documentation and Downloads page at http://www.broadcom.com for the specific adapter.

1.2 Supported Guest Operating Systems

The Emulex driver supports the following guest operating systems:

- CentOS 6.x
- CentOS 7.x
- RHEL 6.6, 6.7, and 6.8
- RHEL 7.1, 7.2, and 7.3
- SLES 11 SP3 and SP4
- SLES 12 SP2

- XenServer® Citrix® 7.0
- Ubuntu® 14.04.03, 15.04, and 16.04
- Debian® 8.1 and 8.4
- FreeBSD® 11.1 and 11.1
- OVM 3.3.3
- OVM 3.4.1
- Oracle® UEK R3 and R4
- Oracle S10 and S11
- Windows Server® 2008, Windows Server 2008 R2, and Windows Server 2016

1.3 Abbreviations

ACL	Access Control List
AL_PA	Arbitrated Loop Physical Address
API	application programming interface
ARI	alternative routing-ID interpretation
BIOS	basic input/output system
CHAP	Challenge Handshake Authentication Protocol
CLI	command line interface
CPU	central processing unit
CRC	cyclic redundancy check
CT	Common Transport
DCBX	Data Center Bridging Capabilities Exchange
DCC	device control channel
DHCP	Dynamic Host Control Protocol
DID	device ID
DMA	direct memory access
DRSS	default queue receive-side scaling
ELS	extended link service
ETO	extended timeout
FC	Fibre Channel
FCF	Fibre Channel over Ethernet Forwarder
FCoE	Fibre Channel over Ethernet
FCP	Fibre Channel Protocol
FDMI	Fabric-Device Management Interface
FIP	FCoE Initialization Protocol
FLOGI	fabric login
GB	gigabyte
Gb	gigabit
GbE	gigabit Ethernet
Gbps	gigabits per second
GFC	gigabit Fibre Channel
GUI	Graphical User Interface
HBA	host bus adapter
HTTP	Hypertext Transfer Protocol
HW	hardware
IHV	Independent Hardware Vendors
I/O	input/output
IMA	iSCSI Management API
IOCB	input/output control block

IOCTL	input/output control
IODM	I/O Device Management
IOPs	I/O operations per second
IP	Internet Protocol
IQN	iSCSI qualified name
iSCSI	Internet Small Computer System Interface
KB	Kilobyte (1024 bytes)
LACP	Link Aggregation Control Protocol
LDTO	link down timeout
LUN	logical unit number
MAC	Media Access Control
MB	megabyte
Mb	megabit
Mbps	megabits per second
MPIO	multipath I/O
MSI	message signaled interrupt
MSI-X	message signaled interrupt-extended
MTU	maximum transmission unit
NIC	network interface card
NFS	Network File System
NLP	Natural Language Processing
NPar	NIC partitioning
NPIV	N_Port ID Virtualization
N_Port	node port
PCI	Peripheral Component Interconnect (interface)
PCIe	Peripheral Component Interconnect Express
PE	Protocol Endpoints
PF	physical function
PLOGI	port login
POST	power-on self-test
R_A_TOV	resource allocation timeout
RFC	Request for Comments
RHEL	Red Hat® Enterprise Linux®
ROM	read-only memory
RPI	remote port indicator
RRQ	Reinstate Recovery Qualifier
RSCN	registered state change notification
RSS	receive-side scaling
SAN	storage area network
SCSI	Small Computer System Interface

SFP	small form-factor pluggable
SLES	SUSE Linux Enterprise Server
SLI	Service Level Interface
SR-IOV	single root I/O virtualization
TCP	Transmission Control Protocol
TSO	TCP segmentation offload
UEFI	Unified Extensible Firmware Interface
UMC	Universal Multichannel
VASA	vSphere® APIs for Storage Awareness
VF	virtual function
VGT	virtual guest tagging
VLAN	virtual local area network
VLAN ID	VLAN identifier
VM	virtual machine
VPD	vital product data
vPort	virtual port
vSwitch	virtual switch
VST	VLAN Switch Tagging
VVols	Virtual Volumes
VXLAN	Virtual extensible LAN
WLAN	wireless LAN
WWN	World Wide Name
WWNN	World Wide Node Name
WWPN	World Wide Port Name
XRI	extensible resource indicator

Chapter 2: Installation

This section provides information for installing the ESXi driver for the FCoE, NIC, and iSCSI protocols.

2.1 Overview

With the release of ESXi 5.5 (vSphere 2013) and subsequent operating systems, VMware introduced a new driver model called *native mode*. Emulex has endorsed the native mode driver model with the following drivers:

- For FCoE functionality, the out-of-box driver is the native mode *brcmfcoe* drive.
- For Ethernet (NIC) functionality, the out-of-box driver is the native mode *elxnet* driver.

NOTE For iSCSI functionality, the out-of-box driver is vmklinux-based *be2iscsi* driver (for ESXi 5.5 and 6.0 versions) and native driver *elxiscsi* (for ESXi 6.5 version).

ESXi 5.5 and subsequent operating systems also require changes to the installation process and tools:

- **Driver parameters migration** – If you upgrade to ESXi 5.5 or later, you must document the existing driver parameter values for the initial ram disk for the server being upgraded. After upgrading, you must then manually reprogram those initial ram disk values. See Section 2.1.1, Migrating Driver Parameters When Upgrading to ESXi 5.5 or Later, for additional information.
- **Command line tool transition** – For the ESXi 5.5 release and all subsequent operating system releases, VMware has transitioned from esxcfg commands to esxcli commands. Both sets of commands are supported in the ESXi 5.5 and 6.0 release. See Section 3.1, ESXi Command Line Tool Transition, for additional information.
- **Native mode driver management support** – With the introduction of the new native mode driver, ethtool is no longer supported. Instead, you can use either the esxcli commands or the VmkMgmtKeyVal interface to provide troubleshooting support. See Section 4.4, Native Mode NIC Driver Troubleshooting Support, and Section 4.5, Native Mode NIC Driver Support for ethtool Commands, for additional information.

2.1.1 Migrating Driver Parameters When Upgrading to ESXi 5.5 or Later

VMware's ESXi 5.5 and later releases have two driver models:

- **vmklinux** – the legacy driver model
- **native mode** – the new driver model

vmklinux is a Linux compatibility layer for device drivers in all ESX releases from ESX 2.x to ESXi 5.1. The vmklinux layer allows IHVs to use their Linux drivers in ESXi with only a few modifications. This consistent driver model usage along the ESXi release train ensures that upgrades automatically copy the initial RAM disk driver parameters from the original operating system to the upgraded operating system. However, this is not the case for ESXi 5.5 and later systems.

VMware recommends loading the drivers by default in native mode because this is the preferred driver model for all subsequent releases. However, because native mode is different than vmklinux and because there is no native mode in prior ESXi releases, there is no driver parameter migration from vmklinux to native mode when migrating to ESXi 5.5 or later.

If you install ESXi 5.5 or later over your existing ESXi system, you will notice after updating and rebooting the system that the driver parameter settings did not migrate. For example, if you set the lpfc_lun_queue_depth driver parameter to 8 on ESXi 5.1 and then update from ESXi 5.1 to ESXi 5.5, you will see the following when the driver parameters are verified after the update:

```
~ # esxcfg-module -g lpfc
Unknown module name lpfc
```

In summary, when updating an existing ESXi system to ESXi 5.5 or later, you must manually set all of the driver parameter values.

2.1.2 Preparing to Upgrade to ESXi 5.5 or Later

Upgrading your server to ESXi 5.5 or later requires an additional planning step. You must evaluate the initial ram disk for the server that will be upgraded and document the existing driver parameter values. After the upgrade, you must then manually reprogram those initial ram disk values.

2.1.2.1 FCoE Driver Example

After upgrading to ESXi 5.5 or later, the server administrator must run the following command.

NOTE The driver binary name is now the native mode Emulex FCoE driver.

```
~ # esxcfg-module -s "brcmfcoe_lun_queue_depth=8" lpfc
```

And to verify that the value has been reprogrammed, type:

```
~ # esxcfg-module -g brcmfcoe
lpfc enabled = 1 options = 'brcmfcoe_lun_queue_depth=8'
```

After the verification is complete, reboot the server to activate the command and burn the new driver parameters into the initial ram disk.

2.1.2.2 Ethernet Driver Example

This example assumes a single OCe11100-series adapter in the host, since the max_vfs parameter receives the number of VFs for each PF.

Before upgrading to ESXi 5.5 or later, you would enter the following command:

```
~ # esxcfg-module -g elxnet
elxnet enabled = 1 options = ‘max_vfs=2,2’
```

After upgrading to ESXi 5.5 or later, when the server administrator runs this command for the out-of-box native mode Ethernet driver for the OCe11100-series adapter, there are no module parameters.

NOTE The driver binary name is now elxnet for the native mode Emulex Ethernet driver for the OCe11100-series adapter.

```
~# esxcfg-module -g elxnet
elxnet enabled = 1 options = ‘’
```

The server administrator must then run the following command to set the driver module parameters:

```
~ # esxcfg-module -s max_vfs=2,2 elxnet
```

And to verify that the value has been reprogrammed, type:

```
~ # esxcfg-module -g elxnet
elxnet enabled = 1 options = 'max_vfs=2,2'
```

After verification has completed, reboot the server to activate the command and burn the new driver parameters into the initial RAM disk.

2.2 Installing the FCoE Driver

This section provides installation information for the driver and the Emulex CIM Provider using the FCoE interface on ESXi systems. Before using this product, you need a working knowledge of FCoE and network storage devices.

2.2.1 Installing the FCoE Driver and Management Software

The Emulex driver is available through the VMware support site. Go to the VMware support website for further details.

For VMware 5.5, 6.0, and 6.5 hosts, you can manage adapters using the Emulex OneCommand® Manager application on Windows® or the OneCommand CNA Manager application for VMware vCenter application, but you must install and use the appropriate Emulex CIM Provider. Refer to the *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide* and the *Emulex OneCommand CNA Manager Application for VMware vCenter for OneConnect Adapters User Guide* for instructions on installing the respective applications. The installation requires that the CIM Provider be installed. For more information on installing the CIM Provider, refer to the *Emulex CIM Provider for OneConnect Adapters Installation Guide*.

NOTE Before installing the management application, you must install the FCoE driver from the VMware website and then reboot the server.

Before installing the FCoE drivers and CIM Provider, verify that the firmware version is correct. If it is, proceed with the installation. If it is not, update the firmware using the OneCommand CNA Manager application from a Windows server or the OneCommand CNA Manager application for VMware vCenter, and reboot the system before proceeding with the installation.

2.2.2 Uninstalling the FCoE Driver

Go to the VMware Patch Download page for instructions.

2.2.3 Installing the Native mode FCoE Driver esxcli Plug-in

This section describes the installation and usage of the esxcli plug-in for the native mode FCoE (brcmfcoe) driver on ESXi systems. This is a diagnostic tool similar to procfs for the FCoE driver module.

NOTE You can download the esxcli plug-in from the Broadcom® website. Currently the vib requires a no-signature-check to install for ESXi 6.0 and ESXi 6.5.

To install the esxcli elxfc plug-in for ESXi5.5, perform these steps:

1. Copy the following file to `/var/log/vmware/` on the ESXi host:
   `vmware-esx-elxcli-fc-<version>.vib`
2. On the ESXi host, install the vibs as follows:

`esxcli software vib install -v /<pathname>/ vmware-esx-elxcli-fc-<version>.vib`

3. Restart the hostd using the following command:
   `/etc/init.d/hostd restart`
   After the driver has restarted, the elxnet namespace appears under elxfc.

Supported commands:

`esxcli elxfc`

Usage:

`esxcli elxfc <cmd> <cmd options>`

Table 1 Available Namespaces

Namespace	Description
adapter	Emulex FCoE HBA instances adapter info
bginjerr	Emulex FCoE HBA instances blockguard error injection page info
event	Emulex FCoE HBA instances events info
eventsnap	Emulex FCoE HBA instances events info with snap after display
fcf	Emulex FCoE HBA instances fcf table info
heap	Emulex FCoE HBA instances heap usage stats and info
iostat	Emulex FCoE HBA instances i/o stats info
kv	Emulex FCoE HBA instances kv page info
lun	Emulex FCoE HBA instances SAN lun info
memlog	Emulex FCoE HBA instances memlog info
memory	Emulex FCoE HBA instances memory stats and info
mgmt	Emulex FCoE HBA instances mgmt stats and info
node	Emulex FCoE HBA instances nodes in SAN info
param	Emulex FCoE HBA instances dynamic config param info
port	Emulex FCoE HBA instances port info and stats
rsnn	Emulex FCoE HBA instances symbolic node name info
sli4q	Emulex FCoE HBA instances sli4q info
target	Emulex FCoE HBA instances SAN targets info
xlane	Emulex FCoE HBA instances xlane info

Table 2 Available Commands

Command	Description
list	Lists Emulex FCoE adapter instance names.
logmessage	Logs `WARNING: lpfc: lpfc_mgmt_log_msg:xxx` message to `vmkernel.log`
modinfo	Returns brcmfcoe module information for list of FCoE boards.

Use the same or higher version of the plug-in as the corresponding driver (brcmfcoe) version installed.

For ESXi 6.0 and ESXi 6.1 follow the same procedure as above, but use the following command in step 2.

`esxcli software vib install -v /<pathname>/ EMULEX-ESX-elxlpfccli-<version>.vib –no-sig-check`

2.3 Installing the NIC Driver

This section provides installation information for the NIC driver. Before using this product, you need a working knowledge of NIC and network-storage devices.

2.3.1 Installing the NIC Driver and Management Software

The latest Emulex NIC driver is available out-of-box. Refer to the operating system guide for installation instructions.

NOTE Before installing the OneCommand CNA Manager application for VMware vCenter software plug-in, you must install the NIC driver from the VMware software website.

Before installing the driver and CIM Provider, verify that the firmware version is correct. If it is, proceed with the installation. If it is not, you must update the firmware using the OneCommand CNA Manager application for VMware vCenter software plug-in. Refer to the *Emulex OneCommand CNA Manager Application for VMware vCenter for OneConnect Adapters User Guide* for additional information. The installation requires that the CIM Provider be installed. For more information on installing the CIM Provider, refer to the *Emulex CIM Provider for OneConnect Adapters Installation Guide*.

2.3.2 Uninstalling the NIC Driver

Go to the VMware Patch Download page for instructions.

2.3.3 Installing the Native Mode Ethernet Driver esxcli Plug-in

This section describes the installation and usage of the esxcli plug-in for the native mode Ethernet (elxnet) driver.

NOTE You can download the esxcli plug-in from the Broadcom website.

To install the esxcli elxnet plug-in, perform these steps:

1. Copy the following file to `/var/log/vmware/` on the ESXi host:
   The vib filename on ESXi 5.5 version is:
   `vmware-esx-elxcli-fc<version>.vib`
   The vib filename on ESXi 6.x version is:
   `EMULEX-ESX-elxnetcli-<version>.vib`
2. On the ESXi host, install the vibs as follows:
   `esxcli software vib install -v /<pathname>/ vmware-esx-elxcli-fc<version>.vib`
   `esxcli software vib install -v /<pathname>/ EMULEX-ESX-elxnetcli-<version>.vib`
3. Restart `hostd`:
   `/etc/init.d/hostd restart`
4. After the driver has restarted, the elxnet namespace appears under elxnet.

Supported commands:

NOTE The following command can be used to get the <*pci dev name*> value:
`#esxcli network nic list`

- **stats** – the stats command is invoked as follows:
  `esxcli elxnet stats get -p <pci dev name>`
  For example:
  `elxnet stats get -p 0000:003:00.1`

- **regdump (fatdump)** – the regdump command is invoked as follows:

  ```
  esxcli elxnet regdump get -f <absolute file pathname> -p <pci device name>
  ```

  For example:

  ```
  esxcli elxnet regdump get -f /fat.txt -p 0000:002:00.0
  ```

- **debugMask** – the debugMask command is invoked as follows:
  - To retrieve the debugMask, type:

    ```
    esxcli elxnet dbgmask get -p <pci dev name>
    ```

  - To set the debugMask, type:

    ```
    esxcli elxnet dbgmask set -p <pci dev name> -m <Debug Mask>
    ```

2.3.4 Virtualization Support

For the best performance, you must install VMware tools in each guest operating system. For information on installing VMware Tools in a Linux or Windows guest operating system, refer to the appropriate VMware ESXi Server documentation.

2.3.5 Obtaining Information on the Installed NIC Driver

To get information on the installed NIC driver, type:

```
esxcli software vib list | grep elxnet
```

2.4 Installing the iSCSI Driver

This section provides installation information for the Emulex iSCSI driver. Before using this product, you need a working knowledge of iSCSI and network-storage devices.

2.4.1 Installing the iSCSI Driver and Management Software

The Emulex iSCSI driver is available out-of-box. A direct link to the iSCSI driver available for download from VMware is provided on the Downloads page on the Broadcom website.

NOTE

Before installing the OneCommand CNA Manager for VMware vCenter Server application, you must install the iSCSI driver from the VMware software website.

Before installing the driver and CIM Provider, verify that the firmware version is correct. If it is, proceed with the installation. If it is not, you must update the firmware using the OneCommand CNA Manager for VMware vCenter Server application or the OneCommand CNA Manager application on a Windows server. Refer to the *Emulex OneCommand CNA Manager for VMware vCenter for OneConnect Adapters User Guide* and the *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide* for additional information. The installation requires that the CIM Provider be installed. For more information on installing the CIM Provider, refer to the *Emulex CIM Provider for OneConnect Adapters Installation Guide*.

The driver kit includes the IMA plug-in; refer to the documentation from VMware for information.

2.5 Updating Drivers with VMware Offline Bundle Media

VMware recommends using the offline bundle to update software on VMware ESXi platforms.

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

To update a driver with the offline bundle media, perform the following steps:

1. Run the following command:

   ```
   esxcli software vib install --maintenance-mode -d
   <pathname>/esx-5.5.0-<build_version>-<protocol>-<driver-version>.zip
   ```

   where <*build_version*> represents the FCoE, NIC, or iSCSI driver.

   For example, to update the iSCSI driver, type the following command:

   ```
   esxcli software vib install --maintenance-mode -d
   <pathname>/esx-5.5.0-472629-be2iscsi-11.2.190.0.zip
   ```

2. Reboot the VMware ESXi Server to activate the drivers.

Chapter 3: Configuration

This section provides configuration information for the ESXi driver for FCoE, NIC, and iSCSI protocols.

3.1 ESXi Command Line Tool Transition

For the ESXi 5.5 release and subsequent operating system releases, VMware has transitioned from esxcfg commands to esxcli commands. This section describes the transition to the esxcli commands.

NOTE Both sets of commands are supported in the ESXi 5.5 and 6.0 releases.

For ESXi 5.5 and subsequent operating systems, Broadcom uses the esxcli version of the command sequence, which has a different command syntax in earlier releases.

For example:

The command syntax to list the supported parameters by a driver:

```
~ # esxcli system module parameters list -m <driver binary name>
```

The command syntax for setting a parameter to a driver module:

```
~ # esxcli system module parameters set -p <parameter_string> -m <driver binary name>
```

3.1.1 FCoE Driver Examples

To set extended logging for the Emulex FCoE driver:

```
~ # esxcli system module parameters set -p lpfc_log_verbose=0x10c3 -m brcmfcoe
```

To list the parameter values assigned to a driver module:

```
~ # esxcli system module parameters list -m brcmfcoe
```

NOTE All the parameters of the new brcmfcoe driver show as lpfc parameters. This issue will be fixed in a future release to represent the parameters as brcmfcoe.

This command is the same as listing parameters. The parameter set has the *value* column completed if applicable. This command lists all instance and global parameters, which makes it necessary to watch for any altered driver parameters.

3.1.2 Ethernet Driver Examples

To list the module parameters supported by the elxnet driver:

```
~ # esxcli system module parameters list -m elxnet
```

To set the number of virtual functions for the Emulex elxnet driver:

```
~# esxcli system module parameters set -p max_vfs=2,2 -m elxnet
```

To verify the altered parameter value in the elxnet driver module:

```
~ # esxcli system module parameters list -m elxnet
```

To disable MSI-X support for the Emulex elxnet driver:

```
~ # esxcli system module parameters set -p msix=0 -m elxnet
```

To verify the altered parameter value in the elxnet driver module:

```
~ # esxcli system module parameters list -m elxnet
```

To enable emi_canceller for the Emulex elxnet driver:

```
~ # esxcli system module parameters set -p emi_canceller=1 -m elxnet
```

To verify the altered parameter value in the elxnet driver module:

```
~ # esxcli system module parameters list -m elxnet
```

To enable VxLAN Offload feature for the Emulex elxnet driver:

```
~ # esxcli system module parameters set -p vxlan_offload=1 -m elxnet
```

NOTE Among Emulex family, only OCe14000-series supports VxLAN Offload feature. LPe16000-series, LPe31000-series, LPe32000-series, BE2, or OCe11000-series do not support VxLAN Offload feature. Even if you specify vxlan_offload = 1 for BE2, OCe11000-series, or for LPe16000-series, LPe31000-series, and LPe32000-series adapters, the driver ignores it.

To disable vxlan_offload for the Emulex elxnet driver:

```
~ # esxcli system module parameters set -p vxlan_offload=0 -m elxnet
```

NOTE This parameter applies to OCe14000-series adapters only.

To enable RSS for the Emulex elxnet driver:

```
# esxcli system module parameters set -p RSS=4,4,4,4 -m elxnet
```

NOTE If RSS is to be disabled for a function, it must be specified as zero for that function.

To enable Dynamic Netqueue feature for the Emulex elxnet driver:

```
~ # esxcli system module parameters set -p dyn_netq=1 -m elxnet
```

To change the global DebugMask value of the Emulex elxnet driver:

```
# esxcli system module parameters set -p debugMask=0x0120 -m elxnet
```

See Section 4.2.3, NIC Informational Log Groups, for more information.

To verify the altered parameter value in the elxnet driver module:

```
~ # esxcli system module parameters list -m elxnet
```

3.2 FCoE Driver Configuration

You can configure driver parameters using native ESXi tools, the Emulex OneCommand CNA Manager application (for use in non-lockdown mode only), or the OneCommand CNA Manager for VMware vCenter Server application (for use in both lockdown and non-lockdown modes).

This document describes how to configure parameters using native ESXi tools. For a more comprehensive description of ESXi tools, go to VMware's public website. If you have further questions, contact a VMware technical support representative.

Refer to the *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide* and the *Emulex OneCommand CNA Manager Command Line Interface for OneConnect Adapters User Guide* for information about the OneCommand CNA Manager application.

Refer to the *Emulex OneCommand CNA Manager for VMware vCenterfor OneConnect Adapters User Guide* for information about the OneCommand CNA Manager for VMware vCenter Server application.

3.2.1 FCoE Driver Parameters Configuration Methods

Configure the driver parameters using any of the following methods:

- Permanent (global)
- Permanent (per adapter)
- Temporary (global)
- Temporary (per adapter)

The OneCommand CNA Manager application supports all four ways to configure driver parameters. This is the preferred method of setting configuration parameters. Refer to the *Emulex OneCommand CNA Manager Application for OneConnect Adapters User Guide* or the *Emulex OneCommand CNA Manager for VMware vCenter for OneConnect Adapters User Guide* for more information.

The native ESXi tools only support permanent configuration methods for the driver parameters. The following section provides further information on permanent configuration methods.

3.2.1.1 Permanent FCoE Configuration Methods Using Native ESXi Tools

Permanent configuration requires that the new values be saved in the ESXi environment. These changes are considered permanent because they persist across system reboots.

See Section 3.2.4, FCoE Driver Configuration Parameters, for parameter names and values. Parameter values are in both hexadecimal and decimal.

NOTE For ESXi systems, the following steps must be executed from the Troubleshooting Administrative Shell environment. If your configuration does not provide access to this shell, refer to VMware's vSphere or VMware's vCenter server guide for enabling driver logging. Alternatively, refer to the *Emulex CIM Provider for OneConnect Adapters Installation Guide* for driver logging.

To make changes that impact all adapters in the system (global changes):

1. From the Troubleshooting Administrative Shell environment's terminal window, type

   ```
   esxcli system module parameters set -p "param1=value param2=value ..." -m lpfc
   ```

   NOTE Use quotation marks around the parameter values only when listing two or more.

2. To reboot the server, type:

   ```
   reboot
   ```

   NOTE VMware does not officially support unloading the driver via `vmkload_mod -u`. If you must unload the driver, contact VMware technical support.

NPIV port creation and deletion are performed by the VMware vSphere client or Virtual Center service. Refer to the VMware documentation for more information.

3.2.1.1.1 Example of Permanent Global Configuration

The following example sets lun_queue_depth (the maximum number of commands that can be sent to a single LUN) to 20 (the default is 30) for all OneConnect adapters in your system.

1. Locate the parameter `lpfc_lun_queue_depth` in Table 3.
2. Set the permanent value by typing:
   ```
   esxcli system module parameters set -p lpfc_lun_queue_depth=20 -m brmcfcoe
   ```
3. To reboot the server, type:
   ```
   reboot
   ```

The new setting is used when the driver reloads.

To verify the setting, type:

```
esxcli system module parameters list -m brmcfcoe
```

3.2.1.1.2 Examples of Permanent Per-Adapter Configuration

The following example sets lun_queue_depth to 20 (the default is 30) for adapter #1.

1. Set the adapter-specific value by typing:
   ```
   esxcli system module parameters set -p lpfc1_lun_queue_depth=20 -m brcmfcoe
   ```
2. To reboot the server, type:
   ```
   reboot
   ```

The new setting is used when the driver is reloaded.

To verify the setting, type:

```
esxcli system module parameters list -m brcmfcoe
```

The following example sets `lun_queue_depth` to 20 (the default is 30) for adapter #1 and `lun_queue_depth` to 10 (the default is 30) for adapter #2.

1. Set the adapter-specific value by typing:
   ```
   esxcli system module parameters set -p "lpfc1_lun_queue_depth=20,
   lpfc2_lun_queue_depth=10" -m brcmfcoe
   ```

   NOTE Type the command all on one line without a carriage return.
2. To reboot the server, type:
   ```
   reboot
   ```

The new settings are used when the driver is reloaded.

To verify the settings, type:

```
esxcli system module parameters list -m brcmfcoe
```

3.2.2 Dynamically Adding LUNs

For instructions on dynamically adding LUNs, refer to the instructions for using rescan in the VMware SAN configuration documentation.

3.2.3 Dynamically Adding Targets

VMware does not provide a native mechanism for dynamically adding targets. After all target/LUN configuration steps have been successfully completed, add the target to the intended fabric zone.

To get the driver to log into the target, either the target or initiator link must be bounced. If the target is configured with security ACLs, the same link bounce requirement applies after the security ACLs are corrected.

To force the ESXi server to rescan all devices:

- Run the following command:

  ```
  esxcfg-rescan vmhbaX
  ```

-Or-

- From the vSphere Client, click **Configuration Tab** > **Storage Adapters**, and then click **Rescan All**.

3.2.4 FCoE Driver Configuration Parameters

Table 3, lists the FCoE driver module parameters, their descriptions, and their corresponding values in native mode.

Dynamic parameters do not require a system reboot for changes to take effect.

Table 3 FCoE Driver Parameters

Module Parameter	Description	ESXi Native Mode Driver Model Values	Comments
throttle_log_cnt	Do not exceed this number of messages logged within 'throttle_log_time'	Default (Def) = 10 Minimum (Min) = 1 Maximum (Max) = 1000	Logging mechanism intended to speed up issue diagnosis by reducing the need to enable driver logging.
throttle_log_time	Do not exceed 'throttle_log_cnt' number of logs within this time limit (seconds)	Def = 1 Min = 1 Max = 60	Works with throttle_log_cnt.
compression_log	Define how often the compression logs are written (in seconds)	Def = 300 Min = 5 Max = 86400	The driver uses this parameter to periodically write status messages to the vmkernel log. The messages provide state analysis on the paths, targets, and adapter. It differs from throttle in that throttle stops the driver from spamming the logs on a very high frequency failure.
suppress_link_up	Suppress link up at initialization: 0x0 = bring up link 0x1 = do not bring up link 0x2 = never bring up link	Def = 0 Min = 0 Max = 2	Enable this parameter to assist with SAN issues during ESX boot.
max_targets	The maximum number of discovered targets allowed.	Def = 256 Min = 0 Max = 4096	Driver parameter to adjust the supported target count.
max_multiq	Sets how many completion queues the driver is requesting from ESXi for each HBA instance. Each completion queue uses a MSI-X vector. 0 indicates disabled Multiqueue.	Def = 4 Min = 0 Max = 8	

Table 3 FCoE Driver Parameters (Continued)

Module Parameter	Description	ESXi Native Mode Driver Model Values	Comments
lpfc_delay_ discovery	Delay NPort discovery when the Clean Address bit is cleared.	Def = 0 Min = 0 Max = 1	
enable_fcp_priority	Enable (1) or disable (0) FCP priority.	Def = 0 Min = 0 Max = 1	
iocb_cnt	IOCBs allocated for extended link service, common transport, and abort sequence in 1024 increments.	Def = 1 Min = 1 Max = 5	
devloss_tmo	The number of seconds the driver holds I/O waiting for a loss device to return.	Def = 10 Min = 1 Max = 255	
log_verbose	Verbose logging bit-mask.	Def = 0 Min = 0 Max = 0x7fffffff	
lun_queue_depth	The maximum number of FCP commands that can queue to a specific LUN. **NOTE** The driver dynamically limits the runtime lun_queue_depth setting to 1/8th of the hba_queue_depth to prevent I/O starvation. An attempt to set the lun_queue_depth higher than the 1/8th setting results in a failure. The console logs and the adapter KeyVal page reflects the failure.	Def = 30 Min = 1 Max = 254	
tgt_queue_depth	The maximum number of FCP commands queued to a specific target port.	Def = 65535 Min = 10 Max = 65535	
hba_queue_depth	The maximum number of FCP commands queued to an FC/FCoE adapter. The driver automatically adjusts the hba_queue_depth to match adapter capabilities. This setting may be overridden.	Def = 8192 Min = 32 Max = 8192	
enable_qfull	Enable the driver's lun_queue_depth ramp down/up functionality when the SCSI device status is Task Set Full (x28). When enabled (value = 1), the Emulex driver gradually reduces the LUN queue depth with each Task Set Full status completion, and then ramps back up as the /IO successfully completes. When disabled (value = 0), the driver takes no action when an I/O completes with a Task Set Full status.	Def = 1 Min = 0 Max = 1	Disabling this parameter has a potential impact on the overall VM performance. Consult your target vendor for guidance.
scan_down	Start scanning for devices from highest AL_PA to lowest.	Def = 1 Min = 0 Max = 1	

Table 3 FCoE Driver Parameters (Continued)

Module Parameter	Description	ESXi Native Mode Driver Model Values	Comments
topology	Select FCoE topology. Valid values are: 0x0 = attempt loop mode then point-to-point 0x01 = internal loopback mode 0x02 = attempt point-to-point mode only 0x04 = attempt loop mode only 0x06 = attempt point-to-point mode then loop	Def = 0 Min = 0 Max = 6	
link_speed	Select link speed: 0 = auto select 1 = 1 Gigabaud 10 = 10 Gigabaud	Def = 0 Min = 0 Max = 32	**NOTE** Setting this option incorrectly can cause the adapter to fail to initialize. **NOTE** If you configure the link speed in a BIOS utility, the link speed may be overwritten by the operating system according to its own configuration settings. To avoid this issue, configure the link speed in both the operating system driver and the Boot BIOS or UEFI driver.
fcp_class	Select FC class of service for FCP sequences.	Def = 3 Min = 2 Max = 3	
use_adisc	Use address discovery on rediscovery, initiated by RSCN, to authenticate FCP devices instead of port login.	Def = 0 Min = 0 Max = 1	
first_burst_size	First burst size for targets that support first burst.	Def = 0 Min = 0 Max = 65536	
max_scsicmpl_time	Use SCSI command completion time to control queue depth to the device. 0 – SCSI command completion time is not used for controlling I/O queue depth. N – I/O queue depth is controlled to limit the I/O completion time to N ms.	Def = 0 Min = 0 Max = 60000	

Table 3 FCoE Driver Parameters (Continued)

Module Parameter	Description	ESXi Native Mode Driver Model Values	Comments
fdmi_on	Enable FDMI support. 0 = disable FDMI support (default setting) 1 = enable FDMI support 3 = enable FDMI with 60-second delay, use a subset of adapter and port attributes 5 = enable FDMI without 60-second delay, use all adapter and port attributes 7 = enable FDMI with 60-second delay, use all adapter and port attributes	Def = 0 Min = 0 Max = 7	FDMI support will default to FDMI-2. If that fails it will failback to FDMI-1. Values 2, 4, and 6 are reserved. Values 3 and 7 are for compatibility with older switches. If issues arise using values 0, 1, or 5, contact Broadcom technical support. If FDMI is not operational, use a value of 7 to introduce a delay. A port reset or adapter reboot is required for the new setting to take effect. The 60-second delay starts after link up. Port attributes subset are attributes 1 to 6. Adapter attributes subset are attributes 1 to 0xB. All port attributes include attributes: 1–0xD, 0x101, 0x102, and 0x103. All adapter attributes include attributes: 1 to 0xC. See Table 4 for a list of the adapter port attributes. See Table 5 for a list of the adapter attributes.
discovery_threads	The maximum number of extended link service commands that can be outstanding during discovery.	Def = 32 Min = 1 Max = 64	
max_luns	The maximum number of LUNs allowed.	Def = 256 Min = 1 Max = 256 (ESXi5.5) Max = 65535 (ESXi 6.0)	
task_mgmt_tmo	The maximum time to wait for task management commands to complete.	Def = 60 Min = 5 Max = 180	
use_msi	Use preferred MSI-X interrupt mode if possible. 0 = MSI disabled (INTx mode) 1 = MSI enabled 2 = MSI-X enabled	Def = 2 Min = 0 Max = 2	
fcf_failover_ policy	FCF Fast failover = 1 Priority failover = 2	Def = 1 Min = 1 Max = 2	

Table 3 FCoE Driver Parameters (Continued)

Module Parameter	Description	ESXi Native Mode Driver Model Values	Comments
enable_rrq	Enable Reinstate Recovery Qualifier functionality. 0x0: disabled, XRI/OXID use not tracked. 0x1: XRI/OXID reuse is timed with R_A_TOV, Reinstate Recovery Qualifier sent. 0x2: XRI/OXID reuse is timed with R_A_TOV, No Reinstate Recovery Qualifier sent.	Def = 2 Min = 0 Max = 2	
sg_seg_cnt	The maximum Scatter Gather Segment Count for DMA. The maximum data allowed in one SG element is 0x80000000.	Def = 64 Min = 64 Max = 4096	
nlp_slab_cnt	NLP slab entries.	Def = 128 Min = 32 Max = 256	Controls the size of the driver's node table. This table in turn limits the driver's ability to discover remote ports, fabrics, initiators, and targets in a zone.
rb_slb_cnt	Receive Buffer slab entries.	Def = 256 Min = 32 Max = 256	Controls the maximum number of Receive Buffers that will be posted to the adapter.
lpfc_max_heap_size	Maximum allowable memory consumption per server for the LPFC module.	Def = 128 MB Min = 64 MB Max = 512 MB	
lpfc_mem_tracker	This turns memory tracking on (1) or off (0).	0 = disabled (default) 1 = enabled	Keeps track of driver heap and slab alloc/free's. Only used when debugging driver memory issues.

NOTE The values in Table 4 and Table 5 are taken from the FC-GS Fibre Channel Standard documents and are passed to the switch through FC-CT commands. Adapter port attributes are provided for each adapter port. Adapter attributes are provided once for each adapter, no matter the number of ports.

Table 4 Adapter Port Attributes for fdmi_on Parameter

Hexadecimal Value	Information Type
0x001	Supported FC-4 Types
0x002	Supported Speed
0x003	Current Port Speed
0x004	Maximum Frame Size
0x005	Operating System Device Name
0x006	Host Name
0x007	Node Name
0x008	Port Name
0x009	Port Symbolic Name
0x00A	Port Type

Table 4 Adapter Port Attributes for fdmi_on Parameter (Continued)

Hexadecimal Value	Information Type
0x00B	Supported Classes of Service
0x00C	Port Fabric Name
0x00D	Port Active FC-4 Types
0x101	Port State
0x102	Number of Discovered Ports
0x103	Port Identifier

Table 5 Adapter Attributes for fdmi_on Parameter

Hexadecimal Value	Information Type
0x001	Node Name
0x002	Manufacturer
0x003	Serial Number
0x004	Model
0x005	Model Description
0x006	Hardware Version
0x007	Driver Version
0x008	Option ROM Version (boot code)
0x009	Firmware Version
0x00A	Operating System Name and Version
0x00B	Maximum CT Payload Length
0x00C	Node Symbolic Name

3.2.5 Creating an FC Remote Boot Disk

For instructions on creating an FC remote boot disk, refer to the VMware SAN Configuration documentation.

3.2.6 Managing Devices through the CIM Interface

VMware on the Visor-based ESXi platforms uses the CIM interface as the only standard management mechanism for device management.

3.2.6.1 Using the OneCommand CNA Manager

For VMware ESXi 5.5, 6.0, and 6.5 hosts, you can manage adapters using the OneCommand CNA Manager application on Windows, but you must install and use the appropriate Emulex CIM Provider.

NOTE If advanced adapter management capabilities are required (for example, iSCSI Management and port disable), use the OneCommand CNA Manager application for VMware vCenter.

3.2.6.2 Using the OneCommand CNA Manager Application for VMware vCenter

The OneCommand CNA Manager for VMware vCenter Server application uses the standard CIM interface to manage the adapters and supports CIM-based device and adapter management. The OneCommand CNA Manager for VMware vCenter Server application also supports existing adapter management functionality based on its proprietary management stack and the standard HBAAPI interface. To manage OneConnect adapters (including updating the

firmware) on an ESXi 5.5, 6.0, or 6.5 host using the OneCommand CNA Manager for VMware vCenter Server application, you must install the out-of-box Emulex CIM Provider on the host.

For more information on installing the OneCommand CNA Manager for VMware vCenter Server application and enabling the CIM Provider, refer to the *Emulex OneCommand CNA Manager for VMware vCenter for OneConnect Adapters User Guide*.

3.2.7 Installing the Emulex CIM Provider

Refer to the *Emulex CIM Provider for OneConnect Adapters Installation Guide* for instructions on installing the Emulex CIM provider.

3.2.8 Creating, Deleting, and Displaying VPorts

The Emulex driver for VMware supports NPIV by default. The only management API for creating and deleting a vPort and creating an NPIV-enabled virtual machine comes from ESXi. VPorts in the driver discover the fabric just like physical ports do, and are subject to the same SAN delays. As the number of VPorts increases, the amount of time it takes to complete remote port discovery increases. This is because the VPorts are created sequentially and each vPort executes discovery synchronously. If your NPIV-enabled virtual machines power-on automatically, powering on could take longer than usual. This is normal for NPIV virtual machines.

The following note applies to VPorts:

NOTE Ensure you are using the latest recommended firmware for vPort functionality. Check the Broadcom website for the latest firmware.

Loop devices and NPIV are not supported on the same port at the same time. If you are running a loop topology and you create a vPort, the vPort's link state is offline. VMware ESX supports fabric mode only.

The OneCommand CNA Manager application sees all VPorts created by the driver, but the application has read-only access to them.

3.2.9 Configuring VVols in ESXi 6.0 and 6.5

The Emulex native mode FCoE driver supports the VVol feature released with ESXi 6.0 and 6.5. VMware's VVol feature allows for dynamic provisioning of storage, based upon the needs of a VM. VM disks, also called VVols, allow VMware administrators to manage storage arrays through the API. Arrays are logically partitioned into storage containers. VVols, are stored natively in the storage containers. I/O from ESX to the array is managed through an access point or Protocol Endpoint (PE) and the storage provider.

3.2.9.1 Storage Containers

Storage containers are a logical abstraction and hold groups of VVols that are physically in the storage array. Storage containers are an alternative to traditional storage based upon LUNs or NFA shares. Storage containers are set up by a storage administrator. Storage container capacity is based on physical storage capacity. The minimum is one storage container per array and the maximum depends upon the array. One storage container can be simultaneously accessed via multiple PEs. Once the storage provider and PEs are in place, the storage container is visible to ESXi hosts.

3.2.9.2 Protocol Endpoints

A PE is an access point that enables communication between an ESXi host and a storage array system. A PE is not a datastore; it is the I/O transport mechanism to access the storage container. A PE is part of the physical storage fabric. A PE is created by a storage administrator.

3.2.9.3 Storage Providers

Storage providers are also referred to as VASA providers. Out-of-band communication between vCenter and the storage array is achieved through the storage provider. The storage provider creates the VVols.

For more information about VVols and instructions on configuring VVols, refer to the VMware and target vendor–supplied documentation.

3.3 NIC Driver Configuration

This section describes how to configure parameters for the ESXi NIC driver.

Table 6, NIC Ethernet Driver Parameters, lists the Ethernet driver module parameters, their descriptions, and their corresponding values in previous ESXi environments and in ESXi 5.5 and 6.0 native mode.

Table 6 NIC Ethernet Driver Parameters

Module Parameter	Description	ESXi Native Mode Driver Model Values	Comments
debugMask	The DebugMask is a bit-vector (uint32) and each bit represents a group. The DebugMask value is in eight-digit hexadecimal format (for example, 0x00000101). DRIVER: 0x1 UPLINK: 0x2 QUEUE: 0x4 INTR: 0x8 MCC: 0x10 TX: 0x20 RX: 0x40 MGMT: 0x80 WORKER: 0x100 SRIOV: 0x200 EVENT: 0x400 VLAN: 0x800 VXLAN: 0x1000	Def = 0x1217 (includes DRIVER, UPLINK, MCC, QUEUE, SRIOV, and VXLAN groups)	
emi_canceller	Enable or disable the EMI Canceller.		
max_vfs	The number of PCI VFs to initialize. 0 = disabled 1–63 = enable this many VFs (depends on the ESXi version and adapter)	Def = 0 Min = 0 Max = 63 (depends on adapter)	The driver supports a maximum of 30 VFs per PF for the OCe11100-series adapters, 32 VFs per PF for the OCe14100-series adapters, and 63 VFs per PF for the OCe14400-series adapters. **NOTE** HP Virtual Connect adapters support a maximum of 24 VFs per PF.
msix	Enable or disable MSI-X.	Def = 1	
RSS	Enable or disable RSS.	Def = 0	

Table 6 NIC Ethernet Driver Parameters (Continued)

Module Parameter	Description	ESXi Native Mode Driver Model Values	Comments
vxlan_offload	Enable or disable vxlan_offload.	Def = 1	
DRSS	Enable or disable DRSS.	Def = 0	
dyn_netq	Enable or disable Dynamic Netqueue.	Def = 1, for elxnet version >= 11.2.1123.0 Def = 0 , for all other older versions of elxnet **NOTE** Use the appropriate value based on the elxnet version.	

In the elxnet driver, there are two new module parameters:

- **Enable or disable MSI-X support**– The driver default is to have MSI-X enabled.
- **Enable or disable vxlan_offload support** – The driver default is to have vxlan_offload enabled.

NOTE The following output is for illustrative purposes only. The actual output may vary depending on the adapter installed in the system.

The server administrator must run the following command to disable MSI-X:

```
~ # esxcli system module parameters set -p msix=0 -m elxnet
```

And to verify that the value has been reprogrammed:

```
~ # esxcli system module parameters list -m elxnet
```

The server administrator must run the following command to disable vxlan_offload:

```
~ # esxcli system module parameters set -p vxlan_offload=0 -m elxnet
```

NOTE This parameter applies to OCe14000-series adapters only.

And to verify that the value has been reprogrammed:

```
~ # esxcli system module parameters list -m elxnet
```

3.3.1 NetQueue Support

- NetQueue enables support for multiple transmit and receive rings for improved performance. NetQueue is enabled by default.
- To determine whether NetQueue is enabled, run the following command from the ESXi Server console:

  ```
  # esxcfg-advcfg -j netNetqueueEnabled
  ```

  If it returns 'netNetqueueEnabled = FALSE', NetQueue is disabled.
- To enable NetQueue, run the following command:

  ```
  # esxcfg-advcfg -k TRUE netNetqueueEnabled
  ```

  Reboot the ESXi server for the change to take effect.

3.3.2 Receive-Side Scaling

RSS enables the ESXi 5.5 or 6.0 NIC driver to distribute incoming TCP traffic across multiple CPU cores for improved performance.

RSS can be selectively enabled for a particular vNIC in a VM. There can be VMs that require RSS and VMs that do not require RSS. This is configurable in the .vmx file. RSS is disabled by default.

To use RSS, the VM must be configured properly:

- The VM must have at least four cores and 4 to 8GB of RAM.
- The vNIC on which RSS is enabled must use VMXNET3 drivers.
- The VMXNET3 drivers must be updated to the latest version.
- The following entry must be added to the .vmx file in the datastore for each VM that requires RSS support for its traffic.

  `ethernet<x>.pnicFeatures="4"`

  Where <x> is the vNIC interface number on which the RSS feature is required.

  The absence of the above line for the vNIC in the .vmx file will not force the VM or ESX to use the RSS feature, even though RSS is enabled and queues are created.

  RSS limitations include:

  — MSI-X must be enabled.
  — When Multichannel is enabled on OCe11100-series adapters, RSS is supported on the first functions of port 0 and port 1 only.
  — The number of regular NetQueues is reduced by 1 when RSS is enabled.
  — RSS is not supported on one-port OCe11100-series adapters.
- RSS is supported when the ESXi host has at least 8 CPUs for UMC or 12 CPUs for a physical NIC (the number of CPUs must be greater than or equal to the number of NetQueues plus the number of RSS queues).

Table 7 Available Queues for Each Adapter Model

Adapter/Mode	Number of NetQueues (including the default queue)	Number of RSS Queues
OCe11100-series physical NIC	8	4
OCe11100-series UMC (function 0 or 1)	3	4
OCe11100-series UMC (functions 2 to 7)	3	0 (RSS is not supported)
OCe14000-series physical NIC	8	4
OCe14000-series, UMC (functions 0 to 7)	4	4

NOTE In cases where the ESXi host does not have the required number of CPUs, the number of NetQueues is reduced by the required number to support four RSS queues.

- RSS support in ESX 5.5 is limited and leaves the configuration of RSS parameters in the hardware completely to the driver. The driver configures the RSS policy in the hardware during driver initialization. The kernel will only request the driver to allocate an RSS queue when an RSS–enabled vNIC (VM) starts receiving packets.

To enable RSS:

1. Run the following command:

   `# esxcli system module parameters set -p "RSS=1" -m elxnet`
2. Reboot the ESXi server for the change to take effect.

To disable RSS:

1. Run the following command:

   `# esxcli system module parameters set -p "RSS=0" -m elxnet`
2. Reboot the ESXi server for the change to take effect.

For ESXi 6.0 systems only, vsish can be used to view the contents of the new node created for RSS in each of the uplinks (when RSS is enabled) once the MAC filter is applied. The following information is displayed:

- The number of RSS queues in the hardware
- The RSS hash key
- The RSS indirection table with load factor
- The RSS indirection table size

To view this information, run the following command:

```
vsish -e get /net/pNics/vmnicX/rxqueues/queues/<qid#>/RSS
```

To view the CPU indirection table, run the following command:

```
vsish -e get /net/pNics/vmnicX/rxqueues/queues/<RSS_qid>/RSS/indTable
```

NOTE This node is available only when there is at least one filter placed on the RSS queue (<*qid#*>). If there is no filter configured on <*qid#*>, the node <*qid#*> is not listed under `/net/pNics/<vmnic#>/rxqueues/queues/`.

The receive traffic on multiple RSS queues can be verified by viewing the driver statistics for receive traffic.

To view the driver statistics for receive traffic, run the following command:

```
esxcli elxnet stats get -p <<device_id>> | grep rx_compl >
```

where <*device_id*> can be obtained from the `esxcli network nic list` command.

NOTE In normal configurations, the RSS queues will be queues with ID - 8, 9, 10, and 11. When Multichannel is enabled, the RSS queues will be queues with ID - 4, 5, 6, and 7.

Multiple core usage can be seen by running the `esxtop` command in the ESX shell. When there is traffic, you can see that more than one CPU is being used even though there is only one active VM.

3.3.3 How an ESXi Server Creates and Names Interfaces

The NIC driver supports a maximum of four adapters per system. For dual-channel adapters running in standard operating mode, the driver creates two interfaces (one for each physical port). The first and second interfaces are respectively named vmnic0 and vmnic1 (assuming there are no other network interfaces in your configuration). The same applies to vNIC-capable adapters if vNIC mode is disabled in the adapter BIOS (for those boards that support vNIC).

If vNIC mode is enabled in the adapter BIOS and the adapter is in NIC-only mode, the driver creates eight interfaces (four for each physical port). The interfaces are labeled vmnic0 through vmnic7 (assuming there are no other network interfaces in your configuration).

NOTE If vNIC mode is enabled in the adapter BIOS and the adapter is in FCoE or iSCSI mode, the driver creates only six NIC interfaces. The other two interfaces are reserved for FCoE and iSCSI.

All vNICs are fully functional and support the same capabilities as a standard NIC. The vNICs can also be linked to a virtual switch in the same way:

```
# esxcfg-nics -l //list recognized nics
# esxcfg-vswitch -l //list available vswitches
# esxcfg-vswitch -a vSwitch0 //create vSwitch0
```

```
# esxcfg-vswitch -A VMNet0 vSwitch0 //create virtual machine network, VMNet0 and
add it to vSwitch0
# esxcfg-vswitch -L vmnic0 vSwitch0 //link vmnic0 to vSwitch0
```

The only difference being that in vNIC mode, each of the four vNICs tied to a physical port shares the port's 10 GbE bandwidth.

3.3.4 Enabling SR-IOV

SR-IOV capability can be enabled for OneConnect adapters if your system BIOS supports SR-IOV.

ESXi 5.5, 6.0, and 6.5 fully support configuration networking options of a virtual function that is assigned to a virtual machine with compatibility for ESXi 5.5 or later, including default VLAN Tagging, VGT, and configuring a static MAC for a virtual function using the ESXi driver.

You can also perform these configuration tasks using the vSphere Client. For more information regarding the vSphere Client, go to the support section of the VMware website.

ESXi 5.5, 6.0, and 6.5 have limited support for extended configuration networking options of a virtual function that is assigned to a virtual machine with compatibility for ESXi 5.5 or later.

Supported extended SR-IOV options:

- Enabling VGT
- Enabling VST mode

Extended SR-IOV option limitations:

- Cannot change the size of the MTU (cannot enable jumbo frames)
- Cannot accept or drop incoming frames for a new address with the MAC address change option
- Cannot enable global promiscuous mode for virtual machine network adapters

For more information regarding supported networking configurations, go to the support section of the VMware website.

NOTE SR-IOV cannot be enabled with other Multichannel technologies, such as UMC or vNIC.

To enable SR-IOV, perform the following steps:

1. Enable IOV capability in the system BIOS.
2. Press <Ctrl>+p to enter the BIOS.
3. Enable SR-IOV for each port of the adapter from the BIOS.
4. If you want to use a newer driver, install the driver .vib file after booting the ESXi host.
5. To load the Ethernet driver with options to enable SR-IOV, run the following command:

   ```
   # esxcli system module parameters set -p "max_vfs=x,y" -m elxnet
   ```

   where *x* and *y* are the number of VFs to be enabled on each of the two NIC PFs. A value of 0 indicates that there are no VFs for the corresponding PF.

   NOTE The driver currently supports a maximum of 30 VFs per PF for OCe11100-series adapters , 32 VFs per PF for OCe14100-series adapters, and 63 VFs per PF for OCe14400-series adapters.
   HP Virtual Connect adapters support a maximum of 24 VFs per PF.

   NOTE The command can be extended to enable more adapters.
6. Reboot the host.

7. To confirm the number of configured VFs, run the following command:

   ```
   # esxcfg-module -g elxnet
   ```

8. To list the SR-IOV enabled pNICs, run the following command:

   ```
   # esxcli network sriovnic list
   ```

9. To list the status of the VFs enabled on a PF, run the following command:

   ```
   # esxcli network sriovnic vf list -n vmnic<X>
   ```

   where `vmnic<X>` is the interface corresponding to the PF.

   NOTE The vSphere Client option Configuration > Advanced Settings also lists the VFs configured with the above command.

10. To assign a VF to a VM, right-click the VM in the vSphere Client and select Edit Settings. Under the Hardware tab, click Add and select PCI Device to attach a VF to the VM.

    NOTE A maximum of six VFs can be assigned to a VM.

3.3.5 Configuring Emulex Universal Multichannel

Emulex Universal Multichannel, or UMC, enables the capability to divide a 10GbE port into multiple physical functions, with flexible bandwidth capacity allocation, that appear to the operating system and network as separate physical devices.

Multichannel can be configured on OCe14000-series adapters through the adapter BIOS or the OneCommand CNA Manager for VMware vCenter application.

- To configure Multichannel using the adapter BIOS, refer to the *Emulex Boot for NIC, iSCSI, and FCoE Protocols User Guide*.
- To configure Multichannel using the OneCommand CNA Manager for VMware vCenter application, see the *Emulex OneCommand CNA Manager for VMware vCenter for OneConnect Adapters User Guide*.

Refer to the *Emulex Universal Multichannel Reference Manual* for additional information on Multichannel.

3.3.5.1 Using ARI

If Multichannel is enabled on an OCe14000-series network adapter, each port can be partitioned into isolated PFs (channels).You can configure up to 16 functions on a one-port OCe14400-series adapter, up to 8 functions per port on a one or two-port OCe14100-series adapter, and up to 4 functions per port on a four-port OCe14100-series adapter.

The maximum number of functions allowed on an adapter is controlled by the adapter's IPL file and the system's support for ARI.

ARI must be available to support more than eight functions on an adapter. OCe14000-series adapters automatically support ARI. However, the following requirements must be met to support more than eight functions on an adapter.

- The system hardware (the motherboard and BIOS) must support ARI.
- ARI must be enabled in the system BIOS.
- The host or guest operating system must support ARI:
  - Windows Server 2012 and newer versions
  - RHEL 6.4 and newer versions
  - SLES 11 SP2 and newer versions
  - ESXi 5.5 and newer versions
- The application management tools, including the OneCommand CNA Manager for VMware vCenter application, must support ARI.
- ARI must be enabled in the firmware using the OneCommand CNA Manager for VMware vCenter application.

If these conditions are not met, you may be able to configure more than eight functions, but only up to eight functions will be running and discovered after a reboot.

3.3.6 Configuring VLANs

VLAN filtering is not supported in the current hardware. To configure VLANs in OCe11000-series adapters, create the vSwitch with the required VLAN ID and use this interface as an adapter to this vSwitch. A native VLAN can also be configured in the guest operating system in VGT mode. (For example, using vconfig in the Linux guest operating system.)

NOTE You cannot run LACP if UMC is enabled.

Refer to the *Emulex Universal Multichannel Reference Manual* for additional information on UMC.

3.3.6.1 Default VLAN Tagging

If default VLAN tagging is configured on a VF, a VLAN ID is automatically added to and removed from every packet sent and received over the VF.

To configure a default VLAN tag:

1. Power off the VM before proceeding to the next step.
2. Edit the <*vm-name*>.vmx file corresponding to the VM from the vmfs Datastore under `/vmfs/volumes/Datastore/<VM_DIR>/`.
3. Append the following line:

   `pciPassthru<X>.defaultVlan = "<vlan_id>"`

   where <*vlan_id*> is the required VLAN ID value for the pass-through VF <*X*> assigned to the VM.
4. Power on the VM.

3.3.6.2 Virtual Guest Tagging

If VGT is configured, a guest operating system can configure any VLAN ID using the native VLAN configuration method in the guest operating system.

To enable VGT:

1. Configure the special VLAN ID 4095 using the steps listed in Section 3.3.6.1, Default VLAN Tagging.
2. Use the native method in the guest operating system to configure the required VLAN ID (for example, vconfig in the Linux guest operating system).

3.3.6.3 Configuring a Static MAC for a VF

A MAC address is automatically configured by ESXi for each VF. This default MAC address can be modified using the following steps:

1. Power off the VM before proceeding to the next step.
2. Edit the <*vm-name*>.vmx file corresponding to a VM from the vmfs Datastore under `/vmfs/volumes/Datastore/<VM_DIR>/`.
3. To configure the MAC address `00:50:56:00:00:02` on the first VF assigned to this VM, append the following lines:

   ```
   pciPassthru0.MACAddressType = "static"
   pciPassthru0.MACAddress = "00:50:56:00:00:02"
   ```
4. Power on the VM.

3.3.7 Configuring NPar (Dell Only)

NOTE This section is specific to Dell.

NPar allows you to divide a 10Gb port into multiple physical functions, with flexible bandwidth capacity allocation, that appear to the operating system and network as separate physical devices.

3.3.7.1 Adapter Configuration

NPar can be configured on OCe14000-series adapters through the adapter BIOS or the OneCommand CNA Manager for VMware vCenter Server application.

- To configure NPar using the adapter BIOS, refer to the *Emulex Boot for NIC, iSCSI, and FCoE Protocols User Guide*.
- To configure NPar using the OneCommand CNA Manager for VMware vCenter Server application, refer to the *Emulex OneCommand CNA Manager for VMware vCenter for OneConnect Adapters User Guide*.

On the host operating system side, NPar provides up to 16 physical functions per device, if NParEP is enabled, using the standard PCI configuration space. The number of physical functions that can be mapped to a physical port depends on the adapter:

- Four-port adapter = four physical functions
- Two-port adapter = eight physical functions

Each physical function or partition is assigned a unique MAC address.

Partitions are available for virtual function assignment and for application segmentation via VLAN or IP subnets.

3.3.7.2 Adapter Requirements

- The partitions can be on separate subnets or VLANs
- Bandwidth allocation is flexible
- No operating system or BIOS changes are required
- No external switch changes are required
- Each partition should have standard NIC properties for stateless offload

The following items are supported on a per-partition basis:

- Per-partition statistics
- TSO required per partition
- MTU per partition
- Support for NetQueues

3.3.7.3 Using NParEP

- NParEP is available on OCe14000-series adapters only.
- NParEP support is available on Dell 13Gb or newer systems only.
- NParEP can be configured on OCe14000-series adapters using the adapter BIOS utility or the OneCommand CNA Manager for VMware vCenter Server application.
- To configure NParEP using the adapter BIOS utility, refer to the *Boot for NIC, iSCSI, and FCoE Protocols User Guide*.
- To configure NParEP using the OneCommand CNA Manager for VMware vCenter application, refer to the *OneCommand CNA Manager for VMware vCenter for OneConnect Adapters User Guide*.

3.3.8 Performance Tuning

3.3.8.1 Using vmxnet Emulation

Using vmxnet3 or vmxnet2 as the emulation driver in guest operating systems is crucial for optimal network performance. To configure vmxnet3 as the emulation driver in guest operating systems, you must install VMware Tools in the guest operating systems. For information on installing VMware Tools in a guest operating system, refer to the appropriate VMware ESXi Server documentation.

After VMware Tools are installed, if you add a network adapter for a guest operating system, select vmxnet3 or vmxnet2 as the adapter type.

3.3.8.2 Enabling TSO

Some Emulex adapters support TSO, which is necessary to achieve optimal transmit throughput performance with low CPU utilization with the adapter. TSO is enabled by default in ESXi servers.

To view the current TSO configuration in the vSphere client, perform these steps:

1. Select the Configuration tab.
2. Under Software, click Advanced Settings.
3. Under Net Features, view the current value of Net.UseHwTSO.
   - If the value is 1, TSO is enabled.
   - If the value is 0, TSO is disabled.

3.3.8.3 Enabling Jumbo Frames

Enabling jumbo frames reduces CPU utilization and is a recommended practice. To use jumbo frames, you must increase the MTU size in the vSwitch and also in the guest operating system. Emulex adapters support MTU sizes between 64 bytes and 9000 bytes. For optimal performance, set the MTU size to the maximum value supported by your network environment. The desired MTU size must be configured in the vSwitch as well as the guest operating systems. Jumbo frames are not enabled by default in ESXi servers.

To configure the MTU size of a vSwitch using the vSphere client, perform the following:

1. Select the Configuration tab.
2. Under Hardware, click Networking.
3. Under Networking, select Properties.
4. Select the vSwitch you want to edit and click Edit.
5. Under vSwitch Properties on the General tab, set the MTU value to 9000.

To configure the MTU size of a vSwitch using the CLI, enter the following command:

```
esxcfg-vswitch -m 9000 vSwitch1
```

3.3.8.3.1 Setting the MTU Size for a Linux Guest Operating System

To set the MTU size in each Linux guest operating system to 9000, run the following command:

```
ifconfig eth<N> mtu 9000
```

where <*N*> is the number of the Ethernet interface on which you are working.

3.3.8.3.2 Setting the MTU Size for a Windows Guest Operating System

To set the MTU size in each Windows guest operating system, perform the following:

1. Go to the Start menu and select Control Panel > System.
2. Select the Hardware tab and open Device Manager.

3. Expand the Network Adapters heading.
4. Right-click the appropriate NIC, and select Properties.
5. Select the Advanced tab and set the MTU value.

3.3.8.4 Using the Port Statistics Counters

You can view all of the port statistics counters maintained by the adapter's Ethernet driver for potential performance issues. Excessive drop or error counters are an indication of a bad link or defective hardware.

To view the statistics of the vmnics on the ESXi host, run the following command:

```
# vsish -e get /net/pNics/<vmnicx>/stats
```

where <*vmnicx*> is the corresponding interface.

To view the port statistics counters on ESXi 5.5 and 6.0, run the following command:

```
# esxcli network nic stats get -n vmnic0
```

See Table 22 for the commands to get driver priv stats.

3.3.8.5 VXLAN Configuration

The elxnet driver supports hardware offloads for VXLAN on the OCe14000-series adapters. These offloads are enabled by default and can be disabled using the module parameter vxlan_offload.

To view the VXLAN statistics, run the following command:

```
# esxcli elxnet stats get -p <pcidevname> | grep vxlan
```

For more information on configuring VXLAN through VMware vCloud Networking and Security Manager and VMware vSphere Distributed Switch (vDS), refer to the VMware VXLAN deployment guide on the VMware website.

3.4 Emulex Drivers for VMware NSX 6.2

This section describes the new functionality associated with the Emulex drivers for VMware NSX 6.2.

3.4.1 Module Parameters

RSS	Creates x RSS (receive-side scaling) queues on one of the Net queues ($x \leq 8$)
DRSS	Creates y RSS queues on the default queue ($y \leq 8$, excluding the default queue)

3.4.2 Functionality

- RSS and DRSS (default queue receive-side scaling) are enabled by default, unless explicitly disabled. This is an NSX requirement.
- RSS=1 – creates four Net queue RSS rings.
- DRSS=1 – creates four default queue RSS rings.
- If system or adapter resources are:
  - — If DRSS is not specified

    -and-

    RSS is not specified or is ≤ 4

    The Net queues are given priority and the RSS rings are reduced as required. The RSS rings are used only if at least two RSS rings can be created with the available resources.

— If RSS is specified and the value is > 4

-or-

DRSS is specified

The RSS queues are given priority and the maximum possible RSS queues are created. The Net queue count is reduced to a minimum of four.

If both Net queue and default queue based RSS rings are to be created, DRSS rings are given priority.

The following table shows the number of RSS and DRSS rings created, corresponding to the value of RSS and DRSS module parameters.

Table 8 RSS and DRSS Rings

RSS	DRSS	No. of RSS Rings Created[1]	No. of DRSS Rings Created[2]
U[3] | 1	U[c] | 1	4	4
0	0	0	0
1	0	4	0
0	1	0	4
x ≥ 2	0	x, max 8	0
x ≥ 2	U[c] | 1	x, max 8	4
U[c] | 1	y ≥ 2	4	y, max 8
0	y ≥ 2	0	y, max 8
x ≥ 2	1	x, max 8	4
x ≥ 2	y ≥ 2	x, max 8	y, max 8

1. Net queue RSS queues include leading Net queue.
2. DRSS rings exclude default Q.
3. Unspecified.

3.4.3 Usage

If two Emulex adapters are installed in the host, then the commands to enable or disable RSS and DRSS queues are as follows:

NOTE If RSS or DRSS is to be disabled for a function, it must be specified as zero for that function.

If the value is unspecified for an adapter, then by default RSS and DRSS are enabled with four rings.

To enable RSS with four rings on each function and disable DRSS, type the following command and press Enter:

```
esxcfg-module elxnet -s "RSS=4,4,4,4 DRSS=0,0,0,0"
```

To enable DRSS with four rings on each function and disable RSS, type the following command and press Enter:

```
esxcfg-module elxnet -s "RSS=0,0,0,0 DRSS=4,4,4,4"
```

To enable DRSS and RSS with four rings each, type the following command and press Enter:

```
esxcfg-module elxnet -s "RSS=4,4,4,4 DRSS=4,4,4,4"
```

To enable DRSS and RSS with the default number of rings (four rings), type the following command and press Enter:

```
esxcfg-module elxnet -s "RSS=1,1,1,1 DRSS=1,1,1,1"
```

To enable DRSS with four rings and RSS with eight rings, type the following command and press Enter:

```
esxcfg-module elxnet -s "RSS=8,8,8,8 DRSS=4,4,4,4"
```

NOTE In case of insufficient resources, the maximum possible RSS queues are created after reducing Net queues to a minimum of four.

To disable both DRSS and RSS, type the following command and press Enter:

```
esxcfg-module elxnet -s "RSS=0,0,0,0 DRSS=0,0,0,0"
```

3.4.4 Limitations

RSS with dynamic Net queue is not supported.

3.4.5 Limitations Exclusive to OCe11000-Series Adapters

- RSS and DRSS features do not support UDP traffic.
- Only one RSS pool is supported per port. Therefore, either DRSS or RSS is supported, but not both.
- The leading queue does not participate in RSS traffic.
- Only the first physical function on the port supports RSS in Multichannel mode, due to limited resources.
- RSS is not supported when SR-IOV is enabled on the adapter.

3.5 iSCSI Driver Configuration

This section describes how to configure parameters for the ESXi iSCSI driver.

3.5.1 Configuring iSCSI Driver Options

The following table lists the user configurable iSCSI driver options. It includes a description of the parameters, default values, and the limits within which they can be configured.

NOTE If the value given for a parameter is outside the supported range (Minimum and Maximum values), the driver will log an error in the

Event Log and continue to load using the default value of the parameter.

Table 9 User-configurable iSCSI Driver Options

Parameter	Default Value	Minimum Value	Maximum Value	Description
LDTO	20 seconds	0 seconds	3600 seconds	Link Down Timeout (in seconds) This determines the amount of time the initiator driver waits for the controller's physical link to become available before reporting that the LUNs are unavailable to the operating system.
ETO	30 seconds	0 seconds	3600 seconds	Extended Timeout (in seconds) This determines the amount of time the initiator driver waits for the target to become available after it has lost connection to the target during an I/O operation. **NOTE** If the minimum value is set between 0 and 19, the driver assumes a value of 20 internally. No modifications are seen in the esxcfg-module -g be2iscsi command output.
im_policy	2	0	4	Controls the rate of interrupts for the adapter. For more information, see Section 3.5.2, Interrupt Moderation Policy Settings.
large I/O	128	128	512	Maximum transfer size in a single I/O request in kilobytes. By default, the iSCSI driver supports a maximum of 128 KB of data in a single I/O request. This option can be used to enable support for 512 KB of data in a single I/O request.
log_level	0x00	—	0xff	Enables the logging of debug information in system logs for a specific path, or all paths, in the driver.

The following command line shows how to configure the driver with an LDTO value of 25 seconds:

```
# vmkload_mod be2iscsi ldto=25
```

To configure the ESXi Server to load the iSCSI driver with this value after each reboot, run the following commands and reboot the system:

```
# esxcfg-module -s "ldto=25" be2iscsi
# /usr/sbin/esxcfg-boot -r
# reboot
```

The im_policy (Interrupt Moderation policy) parameter configures the Emulex iSCSI driver to use different settings for Interrupt Moderation. An im_policy value of 1 achieves the highest interrupt rate, whereas the value 4 provides the lowest interrupt rate. The default value is 2. An im_policy of 0 turns off the Interrupt Moderation algorithm in the driver.

The large_io option can be used to modify the maximum transfer size in a single SCSI command. By default, the Emulex iSCSI driver in ESXi supports up to 128 KB and 32 Scatter Gather entries in a single SCSI command. If applications issue I/O requests larger than 128 KB or need more than 32 Scatter Gather entries, the request will be split into multiple requests by the driver. By specifying large_io=512; the iSCSI driver can support up to 512 KB of data and a total of 128 Scatter Gather entries in a single SCSI command.

NOTE By setting the option to 512, the amount of physical memory consumed by the driver increases. Also, though intermediate values between 64 and 512 will be accepted, the memory used by the driver will be the same as what will be used for large_io=512.

To set the large_io parameter in ESXi, type the following commands:

```
esxcfg-module -s large_io=512 be2iscsi --> to set the parameter
```

```
esxcfg-module -g be2iscsi --> to view the parameter value
esxcfg-module -i be2iscsi --> to view the options
```

Even though the VMware operating system can be tuned to accept larger I/O sizes from guest operating systems, the guest operating systems will sometimes need to be tuned to create those larger I/Os. For example, a Windows 32-bit guest operating system will by default have a maximum transfer size of 64 KB, even if the VMware kernel is tuned to allow 512-KB I/O transfers. You must modify the guest registry to achieve the maximum transfer rate set in the VMware kernel.

To do this, use the registry editor to add or modify the following entry to change the maximum transfer size for the Windows guest:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Symmpi\Parameters\Device\MaximumSGList
```

The formula used to calculate the proper value for MaximumSGList is:

- For a 32-bit Windows guest use the following formula:

  ```
  MaximumSGList = ((Maximum Transfer Size) / 4) + 1
  ```

- For a 64-bit Windows guest use the following formula:

  ```
  MaximumSGList = ((Maximum Transfer Size) / 8) + 1
  ```

For example, to allow a 256 KB transfer size on a 32-bit guest, this would be the formula to use: 256 / 4 = 64 + 1 = 65 (decimal) or 0x41(hexadecimal).

The maximum value allowed for MaximumSGList is 255 or 0xFF. For the particular value of 0xFF, the internal value passed to Windows is increased to 0x101, allowing support for a full 1-MB transfer (2 MB for 64-bit).

The MaximumSGList is 129 for a 512-KB transfer.

Other guest operating systems may also need to be tuned similarly. Consult the tuning guides that accompany those operating systems.

3.5.2 Interrupt Moderation Policy Settings

Interrupt Moderation Policy settings control the rate of interrupts for the adapter. By default, the driver implements an Interrupt Moderation scheme based on the I/O load and the interrupt rate. The default setting for the Interrupt Moderation Policy attempts to vary the interrupt rate between 3500 to 10000 interrupts per second. In addition, the Emulex iSCSI driver allows other configuration settings which are listed in the following table.

Table 10 Interrupt Moderation Policy Settings

Setting	Parameter	Description
Disabled	im_policy=0	Interrupt Moderation algorithm is turned off in the driver.
Aggressive	im_policy=1	Achieves the highest interrupt rate among all available settings.
Moderate	im_policy=2	This is the default value.
Conservative	im_policy=3	Achieves a lower interrupt rate than Moderate.
Very Conservative	im_policy=4	Achieves the minimum interrupt rate among all available settings.

While the default setting of Moderate may work for most configurations, there are instances when the setting may need to be changed. Changing the Interrupt Moderation Policy setting should be based on the initiator system configuration, the number of iSCSI targets that will be connected, the I/O load, and the throughput and latency offered by these iSCSI targets.

On systems capable of sustaining a higher interrupt rate and if the number of targets that will be connected is less (up to 8), the more Aggressive setting will result in lower latency and higher values of I/O operations per second (IOPs).

But the higher interrupt rate could also result in system stalls and freezes, especially during higher queue depth values and smaller sized I/O requests.

On a configuration that involves a large number of iSCSI targets (more than 32 or 64) and higher values of queue depth, the default setting may prove to be too aggressive and the Interrupt Moderation setting may need to be changed to Conservative or Very Conservative. Though this will increase latency of an I/O request, the lower interrupt rate may allow the system to function under a high load.

3.5.2.1 iSCSI Error Handling

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

- Loss of an immediate link to the initiator (for example, a cable disconnect or a port failure).

  The firmware detects and notifies the driver of a loss of the link. When this happens, the driver will queue the I/O requests internally up to a configured timeout period so the operating system does not see I/O errors. This timeout is known as LDTO.
- Loss of connection to the target due to target or network disconnection at the target.

  If the driver has I/O requests pending with the target and the target becomes unavailable (due to target going down or failing over, or network issues at the target), the driver queues up the I/O request internally up to a configured timeout period. This timeout is known as ETO.

When the configured threshold for LDTO and ETO is reached and the initiator is still unable to connect to the target, the driver fails all I/O requests. At this point, I/O errors will be seen by the application and operating system.

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

3.5.3 Configuring LDTO and ETO on ESXi Server

The following table lists the default values of LDTO and ETO on ESXi Server and the limits within which they can be configured.

Table 11 LDTO and ETO Default Values on ESX Server

Parameter	Default Value	Minimum Value	Maximum Value
LDTO	20 seconds	0 seconds	3600 seconds
ETO	30 seconds	0 seconds	3600 seconds

NOTE If the value of ETO is set to a number between 0 and 19, the driver assumes a value of 20 seconds internally. You will not see any modification to the esxcfg-module -g be2iscsi command output.

LDTO and ETO values are configurable during insmod time. The ETO value specified during insmod is the default ETO value applied to all targets.

The following command line shows how to configure the driver with an LDTO value of 25 seconds:

```
# vmkload_mod be2iscsi ldto=25
```

To configure ESXi Server to load the iSCSI driver with this value after each reboot, run the following commands and reboot the system:

```
# esxcfg-module -s "ldto=25" be2iscsi
# /usr/sbin/esxcfg-boot -r
# reboot
```

3.5.4 Multipath I/O Support

This section describes the installation and login processes for multipath I/O support.

3.5.4.1 Configuring and Enabling Support for ESXi MPIO on Non-boot Targets

To configure and enable support for ESXi MPIO on non-boot targets:

1. Connect your configuration for multipath.
2. Log into your targets with all paths using vSphere Client or iSCSISelect. For more information about using vSphere Client, see Section 3.5.4.3, Logging into Targets Using vSphere Client. For information about using iSCSISelect, refer to the *Emulex Boot for NIC, iSCSI, and FCoE Protocols User Guide*.

3.5.4.2 Configuring and Enabling Support for ESXi MPIO on Boot Targets

To configure and enable support for ESXi MPIO on boot targets:

1. Connect your configuration for multipath. Multipath can be configured before or after installation.
2. Log into your targets with all paths using vSphere Client or iSCSISelect. For more information about using vSphere Client, see Section 3.5.4.3, Logging into Targets Using vSphere Client. For information about using iSCSISelect, refer to the *Emulex Boot for NIC, iSCSI, and FCoE Protocols User Guide*.

3.5.4.3 Logging into Targets Using vSphere Client

To log into targets using the vSphere Client:

1. Log into the server you would like to configure from the vSphere Client and select the Configuration tab.
2. In the Hardware section, select the Storage Adapters link.
3. Select the OneConnect host adapter in the Storage Adapters list.
4. Select the initiator port, and then click Properties.
5. From the iSCSI Initiator Properties screen, click the Dynamic Discovery tab.
6. Click Add to display the Add Send Target Server menu.
7. Type in the IP address of the first target portal and click OK.
8. The initiator finds the target machines.
9. Click the Static Discovery tab to find all target portals.
10. Click Close. A dialog box indicates that a rescan is needed.
11. Click Yes. The configured LUNs are listed.
12. Repeat the steps to log into the other target portal to set up MPIO.
13. To check that both paths are connected to the same LUN, follow these steps:
    a. Select the LUN and right-click.
    b. Click Manage Paths. Multipaths on the LUN are shown.

3.5.4.4 Error Handling Under MPIO and Cluster Configurations

In an MPIO or cluster configuration, fault tolerant software is present on the system that makes the iSCSI driver error handling redundant. These configurations also require that I/O errors be reported as soon as they are detected, so the software can failover to an alternate path or an alternative node as quickly as possible.

When the iSCSI driver is run under these configurations, the error handling implemented in the iSCSI driver must be turned off by setting the default value of LDTO and ETO to 0. The changes will take effect during the next driver load.

3.5.5 Reading the Driver Statistics for a Specified Port

You can read the iSCSI driver's statistics for a specific port on the adapter to diagnose potential performance issues.

To read the driver statistics for a particular port, use the following command:

```
cat /proc/be2iscsi/be2iscsi<XX>/driver_stats
```

where XX is the specific port number.

3.5.6 Setting the Log Level

The log level setting enables the logging of debug information in system logs for a specific path, or all paths, in the driver.

Use one of the following methods to set the log level:

1. User configurable driver parameters – use either of the following commands to set the log level.

   NOTE The new log level value will take effect after the next reboot.

   ```
   # esxcfg-module -s "log_level=0xff" be2iscsi
   ```

   -Or-

   ```
   # vmkload_mod be2iscsi "log_level=0xff"
   ```

2. 'log_level' proc entry setting – use the following commands to read and write values to the log_level proc entry.
   - `cat /proc/be2iscsi/log_level –` This command displays the current log level and help.
   - `echo 0xff > /proc/be2iscsi/log_level` – This command sets the log_level value in the driver to enable the desired logging.

See Table 12 for a list of available log level values.

Table 12 Log Level Values

log_level Value	Description
0x01	Enables logging in the iSCSI driver initialization and unload path.
0x02	Enables logging in the iSCSI driver I/O path.
0x04	Enables logging in the iSCSI driver error handling path.
0x08	Enables logging in the iSCSI driver configuration path, such as adding targets, deleting targets, and so on.
0x10	Enables logging in the iSCSI driver IOCTL path.

Chapter 4: Troubleshooting

Your system may operate in an unexpected manner in certain circumstances. This section explains many of these circumstances and offers one or more workarounds for each situation.

4.1 Troubleshooting the FCoE Driver

This section provides troubleshooting information for the FCoE driver.

Table 13, identifies some of the common situations and their potential resolutions.

Table 13 Troubleshooting the FCoE Driver

Situation	Resolution
Port link fails to come up.	Make sure the vfc and vlan are correctly configured for all FCoE links. If the adapter firmware and FCoE switch do not complete FIP correctly, an FCF is not presented to the driver and the virtual link for FCoE will not come up.
The Emulex driver is not loaded and all paths are down.	Use the lspci utility to determine whether the Emulex ports are being properly identified. If not, find out if the driver ISO was correctly installed. You must have the correct driver for the installed adapter because the device PCI IDs are installed with the driver package. Examine the /var/log/vmkernel.log file for brcmfroe log messages indicating an error. If you specified driver logging (see Section 3.2, FCoE Driver Configuration), make sure you spelled the driver parameters correctly. The ESX module subsystem will not load the driver on reboot if the parameters are not spelled correctly. In this case, contact Broadcom support.
All paths are down.	Use the driver's KV pages to get critical information. First check the link state. The KV command shown in the footnote shows the driver's current link, whether or not it has found a fabric and the link speed.[1] If the data shows Link Up -Ready and Mode Online, check the discovered nodes. Fabric, Initiator and Target types show the SAN as it was presented and discovered by the driver. Additionally, if you are experiencing periodic path outage, the command shown in the second footnote in a script loop will show if the node status is changing or if the node is going offline.[2] If your target or initiator is not in the driver's discovered list, check your zone membership and the state of all zone members. Contact Broadcom customer support if you are unable to resolve missing zone members.
The 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.	Download and install the adapter firmware that complies with the minimum supported version (or later) listed on the Broadcom website at http://www.broadcom.com.
System panics when booted with a failed adapter installed.	Remove the failed adapter and reboot.

The FCoE driver does not discover all remote ports in the configuration switch zone. Some initiators or targets may appear to be missing.	Evaluate your switch zone. Count how many entries there are and add at least seven more (to account for fabric logins). If the sum exceeds 128, you must increase the driver's node table size. The following commands increase it to 200 entries. See Section 3.2, FCoE Driver Configuration, for more information on this driver parameter.
	Globally:
	`esxcli system module parameters set -p lpfc_nlp_slab_cnt=200 -m brcmfcoe`
	Per instance:
	`esxcli system module parameters set -p lpfc0_nlp_slab_cnt=200 -m brcmfcoe`
	A reboot is required.

1. Data from KV command 1:

```
[root@lucic:~] /usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i vmhba4/Emulex -k adapter -g
Key 'adapter':
 brcmfcoe Adapter Page

Emulex OneConnect FCoE SCSI 11.2.1079.0
Emulex OneConnect OCe14102-UM 10GbE 2-Port SFP+ PCIe 3.0 Universal CNA on PCI bus 0000:23 device 00 fn 2 port
    0 Link Speed: 10 Gb

BoardNum:2
FW Version:11.0.243.16
HW Version:00000003
ROM Version:11.0.243.16
SerialNum:FC34997302
Vendor Id:72410df

SLI Rev: 4
   MQ: Enabled with 4 queues (4)
   NPIV Supported: VPIs max 255  VPIs used 0
   RPIs max 4096  RPIs used 8   IOCBs inuse 0  IOCB max 14   txq cnt 0  txq max 0  txcmplq 0
   XRIs max 2048  FCP 64  Base 0
   FCP SGEs max 66  DMA buf size 2032

Queue Depth
LUN          30
HBA FCP    1848

PCI read error: 0 retry attempts: 0

Link Up - Ready:
   EDTOV 2000 ms  RATOV 10 sec
   PortID 0x11001
   Fabric
   Current speed 10G

    WWPN 10:00:00:90:fa:5d:2f:29  WWNN 20:00:00:90:fa:5d:2f:29

Mode: Online

Interrupt Type: MSIX
BlockGuard Disabled
```

2. Data from KV command 2:

```
[root@lucic:~] /usr/lib/vmware/vmkmgmt_keyval/vmkmgmt_keyval -i vmhba4/Emulex -k node -g
Key 'node':
 brcmfcoe Node page:

WWNN WWPN ScsiID DID    TypeStatus
10:00:00:05:1e:76:b0:00  20:10:00:05:1e:76:b0:00         xfffffe    Fabric              Node ok
10:00:00:05:1e:76:b0:00  21:fc:00:05:1e:76:b0:00         xffffffc   Fabric              Node ok
00:00:00:00:00:00:00:00  00:00:00:00:00:00:00:00         xffffffd   Fabric              Node logged out
20:00:00:00:c9:ef:3c:a3  10:00:00:00:c9:ef:3c:a3         x010901    Initiator           Node ok
20:00:00:00:c9:ef:3c:a7  10:00:00:00:c9:ef:3c:a7         x010d01    Initiator           Node ok
20:00:00:90:fa:5d:2f:31  10:00:00:90:fa:5d:2f:31         x011101    Initiator           Node ok
20:0e:00:11:0d:59:dc:00  20:0e:00:11:0d:59:dc:00      1  x010200    Target              Node ok
20:10:00:11:0d:59:d0:00  20:10:00:11:0d:59:d0:00      0  x010100    Target              Node ok
10:00:00:05:1e:76:b0:00  21:fa:00:05:1e:76:b0:00         xffffffa   Fabric              Node ok
```

4.1.1 FCoE Driver Log Messages

Log messages have traditionally been organized into logical groups based on code functionality in the FCoE driver. With the introduction of the latest Emulex adapters, that grouping is modified to account for additional behaviors. The traditional grouping is maintained, but recently added messages are no longer grouped together.

The messages provided in this section are unmaskable error conditions. They are automatically added to the system console log.

You can examine the `/var/log/vmkernel file.log` to see any of these messages. If you have concerns, the best policy is to run a vm-support dump and contact VMware or Broadcom support staff.

Log messages are organized into logical groups based on code functionality within the driver. Each group consists of a block of 100 log message numbers. Most groups require a single block of 100 message numbers; however, some groups (INIT, FCP) require two blocks.

Table 14 lists the groups and defines the associated number ranges.

Table 14 Message Log Table

LOG Message Verbose Mask Definition	Verbose Bit	Verbose Description
LOG_ELS	0x1	Extended link service events
LOG_DISCOVERY	0x2	Link discovery events
LOG_MBOX	0x4	Mailbox events
LOG_INIT	0x8	Initialization events
LOG_LINK_EVENT	0x10	Link events
LOG_MGMT_ERROR	0x20	IODM error logging
LOG_FCP	0x40	FCP traffic history
LOG_NODE	0x80	Node table events
LOG_TEMP	0x100	Temperature sensor events
LOG_BG	0x200	BlockGuard events
LOG_MISC	0x400	Miscellaneous and FCoE events
LOG_SLI	0x800	SLI events
LOG_FCP_ERROR	0x1000	Selective FCP events
LOG_LIBDFC	0x2000	IOCTL events

Table 14 Message Log Table (Continued)

LOG Message Verbose Mask Definition	Verbose Bit	Verbose Description
LOG_VPORT	0x4000	NPIV events
LOG_SECURITY	0x8000	Security events
LOG_EVENT	0x10000	IOCTL event
LOG_FIP	0x20000	FIP event
LOG_FCP_UNDER	0x40000	FCP underrun errors
LOG_KVPAGE	0x80000	KV page verbose
LOG_TASKMGMT	0x100000	Task management events
LOG_MGMT_TRACE	0x200000	IODM management trace logging
LOG_ALL_MSG	0x7fffffff	Log all messages

The following is an example of a LOG message:

```
#define LOG_ELS          0x00000001      /* ELS events */
#define LOG_DISCOVERY    0x00000002      /* Link discovery events */
#define LOG_MBOX         0x00000004      /* Mailbox events */
#define LOG_INIT         0x00000008      /* Initialization events */
#define LOG_LINK_EVENT   0x00000010      /* Link events */
#define LOG_MGMT_ERROR   0x00000020      /* IODM management error logging */
#define LOG_FCP          0x00000040      /* FCP traffic history */
#define LOG_NODE         0x00000080      /* Node table events */
#define LOG_TEMP         0x00000100      /* Temperature sensor events */
#define LOG_BG           0x00000200      /* BlockGuard events */
#define LOG_MEM_HEAP     0x00000400      /* Mem tracker heap logging */
#define LOG_SLI          0x00000800      /* SLI events */
#define LOG_FCP_ERROR    0x00001000      /* log errors, not underruns */
#define LOG_LIBDFC       0x00002000      /* Libdfc events */
#define LOG_VPORT        0x00004000      /* NPIV events */
#define LOG_MEM_SLAB     0x00008000      /* Mem tracker slab logging */
#define LOG_EVENT        0x00010000      /* CT,TEMP,DUMP, logging */
#define LOG_FIP          0x00020000      /* FIP events */
#define LOG_FCP_UNDER    0x00040000      /* FCP underruns errors */
#define LOG_KVPAGE       0x00080000      /* KV page verbose */
#define LOG_TASKMGMT     0x00100000      /* Task Management events. */
#define LOG_MGMT_TRACE   0x00200000      /* IODM management trace logging */
#define LOG_SCSI_CMD     0x00400000      /* ALL SCSI commands */
#define LOG_EDIF         0x00800000      /* External DIF events */
#define LOG_KV_ERROR     0x01000000      /* Key-value mgmt error logging */
#define LOG_ESXCLI       0x02000000      /* Esxcli mgmt interface logging */
#define LOG_ALL_MSG      0x7fffffff      /* LOG all messages */
```

In the above LOG message:

- brcmfcoe - driver binary
- lpfc_mbx_cmpl_read_topology - function generating the log
- 1: identifies Emulex HBA1.
- 1305 identifies the LOG message number.

NOTE If the word 'Data:' is present in a LOG message, any information to the right of 'Data:' is intended for Broadcom technical support/engineering use only.

Unless otherwise noted in the ACTION: attribute, report these errors to Broadcom technical support. Broadcom requests that when reporting occurrences of these error messages, you provide a tarball of all vmkernel files in /var/log.

4.1.2 Extended Link Service Events (0100 to 0199)

4.1.2.1 elx_mes0100: FLOGI failure Status:<status>/<extended_status> TMO:<timeout>

DESCRIPTION: An extended link service FLOGI command that was sent to the fabric failed.

DATA: (1) ulpStatus, (2) ulpWord[4], (3) ulpTimeout

ACTION: This error could indicate a fabric configuration error or an internal driver issue. If this issue persists, report the error to Broadcom technical support.

4.1.2.2 elx_mes0111: Dropping received ELS cmd

DESCRIPTION: The driver dropped an extended link service response ring entry.

DATA: (1) ulpStatus, (2) ulpWord[4], (3) ulpTimeout

ACTION: This error could indicate a software driver or a firmware issue. If this issue persists, report the error to Broadcom technical support.

4.1.2.3 elx_mes0113: A FLOGI ELS command <elsCmd> was received from DID <did> in Loop Mode

DESCRIPTION: While in Loop Mode an unknown or unsupported extended link service command was received.

DATA: None

ACTION: Check the device DID.

4.1.2.4 elx_mes0115: Unknown ELS command <elsCmd> received from N_Port <did>

DESCRIPTION: Received an unsupported extended link service command from a remote N_Port.

DATA: None

ACTION: Check the remote N_Port for a potential issue.

4.1.2.5 elx_mes0122 FDISC Failed (value). Fabric Detected Bad WWN

DESCRIPTION: The driver's F_Port discovery failed. The switch reported a bad WWN in the FLOGI request.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.2.6 elx_mes0124 Retry illegal cmd <value> retry:<value> delay:<value>

DESCRIPTION: The port rejected an extended link service command as illegal. The driver is retrying.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.2.7 elx_mes0125: FDISC Failed (value). Fabric out of resources

DESCRIPTION: The fabric rejected an F_Port discovery because the switch cannot support any more virtual ports.

DATA: None

ACTION: Reconfigure the switch to support more NPIV logins. If this issue persists, contact Broadcom technical support.

4.1.2.8 elx_mes0126: FDISC failed (ulpStatus/ulpWord[4])\n

DESCRIPTION: The ELS F_Port discovery command has failed.

DATA: None

ACTION: Check the port and switch configuration.

4.1.2.9 elx_mes0127: ELS timeout

DESCRIPTION: An ELS IOCB command was posted to a ring and did not complete within ULP timeout seconds.

DATA: (1) elscmd, (2) remote_id, (3) ulpcommand, (4) ulpIoTag

ACTION: If the extended link service command is not going through the adapter, reboot the system. If this issue persists, report the error to Broadcom technical support.

4.1.2.10 elx_mes0133: PLOGI: no memory for reg_login

DESCRIPTION: Memory allocation error.

DATA: (1) nlp_DID, (2) nlp_state, (3) nlp_flag, (4) nlp_rpi

ACTION: This is a memory allocation error. Check the system resources and unload any unused modules.

4.1.2.11 elx_mes0134: PLOGI: cannot issue reg_login

DESCRIPTION: The ELS port login mailbox command has failed.

DATA: (1) nlp_DID, (2) nlp_state, (3) nlp_flag, (4) nlp_rpi

ACTION: Check the port and switch configuration.

4.1.2.12 elx_mes0135: cannot format reg_login

DESCRIPTION: The system could not allocate an RPI or DMA buffer for the mailbox command.

DATA: (1) nlp_DID, (2) nlp_state, (3) nlp_flag, (4) nlp_rpi

ACTION: None required.

4.1.2.13 elx_mes0136: PLOGI completes to N_Port <DID> completion

DESCRIPTION: A port login has completed for which there is no NDLP.

DATA: (1) ulpStatus, (2) ulpWord[4]

ACTION: None required.

4.1.2.14 elx_mes0137: No retry ELS command <ELS_CMD> to remote

DESCRIPTION:

DATA: (1) ulpStatus, (2) ulpWord[4]

ACTION: None required.

4.1.2.15 elx_mes0138: ELS rsp: Cannot issue reg_login for <DID>

DESCRIPTION: The REG_LOGIN mailbox command failed.

DATA: (1) nlp_DID, (2) nlp_state, (3) nlp_flag, (4) nlp_rpi

ACTION: None required.

4.1.2.16 elx_mes0140: PLOGI Reject: invalid nname

DESCRIPTION: An invalid node WWN was provided.

DATA: None

ACTION: None required.

4.1.2.17 elx_mes0141: PLOGI Reject: invalid pname

DESCRIPTION: An invalid port WWN was provided.

DATA: None

ACTION: None required.

4.1.2.18 elx_mes0142: PLOGI RSP: Invalid WWN

DESCRIPTION: The port login sent to the port by a remote port had an invalid WWN.

DATA: None

ACTION: None required.

4.1.2.19 elx_mes0144: Not a valid WCQE code: <Completion Code>

DESCRIPTION: The completion queue handler detected an invalid type.

DATA: None

ACTION: None required.

4.1.2.20 elx_mes0147: Failed to allocate memory for an RSCN event

DESCRIPTION: Memory could not be allocated to send the RSCN event to the management application.

DATA: None

ACTION: None required.

4.1.2.21 elx_mes0148: Failed to allocate memory for a LOGO event

DESCRIPTION: Memory could not be allocated to send the N_Port logout event to the FC transport.

DATA: None

ACTION: None required.

4.1.2.22 elx_mes0154: Authentication not complete

DESCRIPTION: Memory could not be allocated to send the N_Port logout event to the FC transport.

DATA: None

ACTION: None required.

4.1.3 Link Discovery Events (0200 to 0299)

4.1.3.1 elx_mes0200: CONFIG_LINK bad hba state <hba_state>

DESCRIPTION: A CONFIG_LINK mailbox command completed and the driver was not in the right state.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.3.2 elx_mes0203: Devloss timeout on WWPN <address> N_Port <nlp_DID>

DESCRIPTION: A remote N_Port that was discovered by the driver disappeared for more than lpfc_devloss_tmo seconds.

DATA: (1) nlp_flag, (2) nlp_state, (3) nlp_rpi

ACTION: If the device generating this message is not a target to which the adapter is connected, this error does not affect the data integrity of the I/O between the adapter and the attached storage and can be ignored.

4.1.3.3 elx_mes0206: Device discovery completion error

DESCRIPTION: This indicates that an uncorrectable error was encountered during device (re)discovery after a link up. FC devices are not accessible if this message is displayed.

DATA: None

ACTION: Reboot the system. If the issue persists, report the error to Broadcom technical support. Run with verbose mode on for more details.

4.1.3.4 elx_mes0207: Device <DID> (<WWN>) sent invalid service parameters. Ignoring device.

DESCRIPTION: Invalid service parameters were received from the DID. Ignoring this remote port.

DATA: DID, WWN

ACTION: Verify the remote port's configuration. If the issue persists, report the error to Broadcom technical support. Run with verbose mode on for more details.

4.1.3.5 elx_mes0222: Initial FLOGI/FDISK timeout

DESCRIPTION: The driver sent the initial FLOGI or FDISK to the fabric and never got a response back.

DATA: None

ACTION: Check the fabric configuration. The driver recovers from this and continues with device discovery.

4.1.3.6 elx_mes0223: Timeout while waiting for NameServer login

DESCRIPTION: The login request to the NameServer was not acknowledged within R_A_TOV.

DATA: None

ACTION: Check the fabric configuration. The driver recovers from this and continues with device discovery.

4.1.3.7 elx_mes0224: NameServer Query timeout

DESCRIPTION: Node authentication timeout, node Discovery timeout. A NameServer Query to the Fabric or discovery of reported remote N_Ports is not acknowledged within R_A_TOV.

DATA: (1) fc_ns_retry, (2) fc_max_ns_retry

ACTION: Check the fabric configuration. The driver recovers from this and continues with device discovery.

4.1.3.8 elx_mes0227: Node Authentication timeout

DESCRIPTION: The driver has lost track of what N_Ports are being authenticated.

DATA: None

ACTION: None required. The driver should recover from this event.

4.1.3.9 elx_mes0228: CLEAR LA timeout

DESCRIPTION: The driver issued a CLEAR_LA that never completed.

DATA: None

ACTION: None required. The driver should recover from this event.

4.1.3.10 elx_mes0230: Unexpected timeout, hba linkstate <link_state>

DESCRIPTION: Discovery has timed out and the adapter state is not ready.

DATA: None

ACTION: None required.

4.1.3.11 elx_mes0231: RSCN timeout

DESCRIPTION: The driver has lost track of what N_Ports have RSCNs pending.

DATA: (1) fc_ns_retry, (2) lpfc_max_ns_retry

ACTION: None required. The driver should recover from this event.

4.1.3.12 elx_mes0233: Nodelist not empty

DESCRIPTION: Driver unloaded or hotplug detected a node still in use.

DATA: None

ACTION: None required.

4.1.3.13 elx_mes0237: Pending Link Event during Discovery: State <state>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.3.14 elx_mes0241: NameServer Rsp Error Data: <data>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.3.15 elx_mes0246: RegLogin failed

DESCRIPTION: The firmware returned a failure for the specified RegLogin.

DATA: (1) Did, (2) mbxStatus, (3) hbaState

ACTION: This message indicates that the firmware could not perform a RegLogin for the specified DID. A limitation on how many nodes an adapter can view may exist.

4.1.3.16 elx_mes0249: Cannot issue Register Fabric login: Err <err>

DESCRIPTION: The firmware could not issue the fabric reg login; the err value is unique for each possible failure.

DATA: None

ACTION: None required.

4.1.3.17 elx_mes0251: NameServer login: no memory

DESCRIPTION: The firmware could not allocate memory for the NDLP structure.

DATA: None

ACTION: None required.

4.1.3.18 elx_mes0252: Cannot issue NameServer login

DESCRIPTION: The firmware could not issue an extended link service port login to the nameserver DID.

DATA: None

ACTION: Check the port connection and the switch configuration.

4.1.3.19 elx_mes0253: Register VPI: Can't send mbox

DESCRIPTION: The system could not issue the REG_LOGIN command for this vPort.

DATA: None

ACTION: None required.

4.1.3.20 elx_mes0254: Register VPI: no memory goto mbox_err_exit

DESCRIPTION: The system could not allocate memory for the REG_LOGIN mailbox command.

DATA: None

ACTION: None required.

4.1.3.21 elx_mes0255: Issue FDISC: no IOCB

DESCRIPTION: All of the pre-allocated IOCBs are in use.

DATA: None

ACTION: None required.

4.1.3.22 elx_mes0256: Issue FDISC: Cannot send IOCB

DESCRIPTION: The system is unable to send the fabric IOCB.

DATA: None

ACTION: Check the switch configuration.

4.1.3.23 elx_mes0257: GID_FT Query error: <ulpStatus> <fc_ns_retry>

DESCRIPTION: The GID_FT common transport request for the nameserver has failed.

DATA: None

ACTION: Check the switch configuration.

4.1.3.24 elx_mes0258: Register Fabric login error: <mbxStatus>

DESCRIPTION: The REG_LOGIN for the fabric has failed.

DATA: None

ACTION: Check the port configuration and the switch configuration.

4.1.3.25 elx_mes0259: No NPIVFabric support

DESCRIPTION: The switch to which the port is connected does not support NPIV.

DATA: None

ACTION: Check the switch configuration.

4.1.3.26 elx_mes0260: Register NameServer error: <mbxStatus>

DESCRIPTION: The REG_LOGIN mailbox command has failed for the nameserver.

DATA: None

ACTION: Check the switch configuration.

4.1.3.27 elx_mes0261: Cannot Register NameServer login

DESCRIPTION: Either a memory allocation issue or an invalid parameter was sent to the REG_LOGIN.

DATA: None

ACTION: At least one message (0142, 0121, 0133, 0134, or 0135) should precede this message.

4.1.3.28 elx_mes0262: No NPIV Fabric support

DESCRIPTION: The switch to which the port is connected does not support NPIV.

DATA: None

ACTION: Check the switch configuration.

4.1.3.29 elx_mes0263: Discovery Mailbox error: state: <port_state>: <sparam_mbox> <cfglink_mbox>

DESCRIPTION: Either the driver could not allocate resources or it could not send sparam_mbox or cfglink_mbox.

DATA: (1) address of sparam_mbox command, (2) address of cfglink_mbox command.

ACTION: Try to unload and reload the driver when it is convenient.

4.1.3.30 elx_mes0264: No NPIV Fabric support

DESCRIPTION: The switch to which the port is connected does not support NPIV.

DATA: None

ACTION: Check the switch configuration.

4.1.3.31 elx_mes0266: Issue NameServer Req <cmdcode> err <rc> Data: <fc_flag> <fc_rscn_id_cnt>

DESCRIPTION: The driver was unable to send the nameserver common transport command.

DATA: (1) VPorts fc_flag, (2) VPorts fc_rscn_id_cnt

ACTION: Check the switch and port configurations.

4.1.3.32 elx_mes0267: NameServer GFF Rsp <did> Error (<ulpStatus> <un.ulpWord[4]>) Data: <fc_flag> <fc_rscn_id_cnt>

DESCRIPTION: The nameServer GFF common transport request failed.

DATA: (1) VPorts fc_flag, (2) VPorts fc_rscn_id_cnt

ACTION: Check the switch and port configurations.

4.1.3.33 elx_mes0268: NS cmd <cmdcode> Error (<ulpStatus> <un.ulpWord[4]>)

DESCRIPTION: The nameServer common transport request failed.

DATA: None.

ACTION: Check the switch and port configurations.

4.1.3.34 elx_mes0271: Illegal State Transition: node <nlp_DID> event <evt>, state <nlp_state> Data:<nlp_rpi> <nlp_flag>

DESCRIPTION: The current node state does not have a handler for this event.

DATA: (1) nlp_rpi, (2) nlp_flag

ACTION: Verify that all targets are still visible to the SCSI mid-layer.

4.1.3.35 elx_mes0272: Illegal State Transition: node <nlp_DID> event <evt>, state <nlp_state> Data: <nlp_rpi> <nlp_flag>

DESCRIPTION: The driver is completing a port login but does not have the rcv_plogi flag set.

DATA: (1) nlp_rpi, (2) nlp_flag

ACTION: Verify that all targets are still visible to the SCSI mid-layer.

4.1.3.36 elx_mes0273: Unexpected discovery timeout,vport State <port_state>

DESCRIPTION: The discovery process has timed out.

DATA: None

ACTION: Verify that all targets are visible.

4.1.3.37 elx_mes0282: did:<value> ndlp:<value> pusgmap:<value> refcnt<value>, ndlp->nlp_DID, (void *)ndlp, lpfc_init.c-ndlp->nlp_usg_map

DESCRIPTION: Driver clean-up has found a node that is still on the node list during driver unload or PCI hotplug removal.

DATA: None.

ACTION: None required.

4.1.3.38 elx_mes0283: Failed to allocate mbox cmd memory

DESCRIPTION: Mailbox allocation error.

DATA: None

ACTION: None required.

4.1.3.39 elx_mes0285: Allocated DMA memory size <alloclen> is less than the requested DMA memorysize <reqlen>

DESCRIPTION: Memory allocation was truncated.

DATA: None

ACTION: None required.

4.1.3.40 elx_mes0286: lpfc_nlp_state_cleanup failed to allocate statistical data buffer <nlp_DID>

DESCRIPTION: Memory allocation failed for the node's statistical data.

DATA: None

ACTION: None required.

4.1.3.41 elx_mes0287: lpfc_alloc_bucket failed to allocate statistical data buffer <nlp_DID>

DESCRIPTION: Memory allocation failed for the node's statistical data.

DATA: None

ACTION: None required.

4.1.3.42 elx_mes0288: Unknown FCoE event type <event_type> event tag <event_tag>

DESCRIPTION: The firmware has detected an unknown FCoE event.

DATA: None

ACTION: Check the FCoE switch configuration and the adapter DCBX mode.

4.1.3.43 elx_mes0289: Issue register VFI failed: Err <rc>

DESCRIPTION: The driver could not register the virtual fabric index for the FCFI.

DATA: None

ACTION: Check the switch and port configurations.

4.1.3.44 elx_mes0290: The SLI4 DCBX asynchronous event is not handled yet

DESCRIPTION: The SLI-4 DCBX asynchronous event is not handled yet.

DATA: None

ACTION: None required.

4.1.3.45 elx_mes0291: Allocated DMA memory size <alloc_len> is less than the requested DMA memorysize <req_len>

DESCRIPTION: The asynchronous DCBX events are not handled in the driver.

DATA: None

ACTION: Check the switch configuration.

4.1.3.46 elx_mes0293: PM resume failed to start workerthread: error=<error>

DESCRIPTION: The PCI resume (hotplug) could not start the worker thread for the driver.

DATA: None

ACTION: Unload and reload the driver.

4.1.3.47 elx_mes0294: PM resume failed to enable interrupt

DESCRIPTION: The PCI resume (hotplug) could not get an interrupt vector.

DATA: None

ACTION: Unload and reload the driver.

4.1.3.48 elx_mes0297:invalid device group <pci_dev_grp>

DESCRIPTION: While unloading the driver, the driver detect a PCI device that it should not have claimed.

DATA: None

ACTION: None required.

4.1.3.49 elx_mes0299: Invalid SLI revision <sli_rev>

DESCRIPTION: While processing a host attention error or an unrecoverable error, the driver detected an invalid SLI revision.

DATA: None

ACTION: None required.

4.1.4 Mailbox Events (0300 to 0339)

4.1.4.1 elx_mes0300: LATT: Cannot issue READ_LA: Data: <rc>

DESCRIPTION: The link attention handler could not issue a READ_LA mailbox command.

DATA: None

ACTION: None required.

4.1.4.2 elx_mes0303: Ring <ringno> handler: portRspPut <portRspPut> is bigger then rsp ring <portRspMax>

DESCRIPTION: The port rsp ring put index is larger than the size of the rsp ring.

DATA: None

ACTION: This error could indicate a software driver, firmware, or hardware issue. Report the error to Broadcom technical support.

4.1.4.3 elx_mes0304: Stray mailbox interrupt, mbxCommand <mbxcommand> mbxStatus <mbxstatus>

DESCRIPTION: Received a mailbox completion interrupt and there are no outstanding mailbox commands.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

4.1.4.4 elx_mes0306: CONFIG_LINK mbxStatus error <mbxStatus> HBA state <hba_state>

DESCRIPTION: The driver issued a CONFIG_LINK mbox command to the adapter that failed.

DATA: None

ACTION: This error could indicate a firmware or hardware issue. Report the error to Broadcom technical support.

4.1.4.5 elx_mes0310: Mailbox command <mbxcommand> timeout

DESCRIPTION: A mailbox command was posted to the adapter and did not complete within 30 seconds.

DATA: (1) hba_state, (2) sli_flag, (3) mbox_active

ACTION: This error could indicate a software driver or firmware issue. If no I/O is going through the adapter, reboot the system. If the issue persists, report the error to Broadcom technical support.

4.1.4.6 elx_mes0311 Mailbox command <value> cannot issue data: <value> <value>

DESCRIPTION: The driver detected an HBA error and can't issue the mailbox.

DATA: (1) sli flags (2) hba flags

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.4.7 elx_mes0312: Ring <ringno> handler: portRspPut <rspPutInx> is bigger then rsp ring <numRiocb>

DESCRIPTION: The IOCB command ring put pointer is ahead of the get pointer.

DATA: None

ACTION: None required.

4.1.4.8 elx_mes0315: Ring <ringno> issue: portCmdGet <local_getidx> is bigger then cmd ring <max_cmd_idx>

DESCRIPTION: The port cmd ring get index is greater than the size of cmd ring.

DATA: None

ACTION: This error could indicate a software driver, firmware, or hardware issue. Report the error to Broadcom technical support.

4.1.4.9 elx_mes0317: iotag <ulp_IoTag> is out of range: max iotag <max_iotag> wd0 <wd0>

DESCRIPTION: The IoTag in the completed IOCB is out of range.

DATA: None

ACTION: This error could indicate a software driver, firmware, or hardware issue. Report the error to Broadcom technical support.

4.1.4.10 elx_mes0319: READ_SPARAM mbxStatus error <mbxStatus> hba state <hba_state>

DESCRIPTION: The driver issued a READ_SPARAM mbox command to the adapter that failed.

DATA: None

ACTION: This error could indicate a firmware or hardware issue. Report the error to Broadcom technical support.

4.1.4.11 elx_mes0320: CLEAR_LA mbxStatus error <mbxStatus> hba state <hba_state>

DESCRIPTION: The driver issued a CLEAR_LA mbox command to the adapter that failed.

DATA: None

ACTION: This error could indicate a firmware or hardware issue. Report the error to Broadcom technical support.

4.1.4.12 elx_mes0323: Unknown Mailbox command <mbxCommand> Cmpl

DESCRIPTION: A unknown mailbox command completed.

DATA: None

ACTION: This error could indicate a software driver, firmware, or hardware issue. Report the error to Broadcom technical support.

4.1.4.13 elx_mes0324: Config port initialization error, mbxCmd <mbxCommand> READ_NVPARM, mbxStatus <mbxStatus>

DESCRIPTION: A read nvparams mailbox command failed during port configuration.

DATA: None

ACTION: This error could indicate a software driver, firmware, or hardware issue. Report the error to Broadcom technical support.

4.1.4.14 elx_mes0330: IOCB wake NOT set

DESCRIPTION: The completion handler associated with the IOCB was never called.

DATA:(1) timeout, (2) timeleft/jiffies

ACTION: This error could indicate a software driver, firmware, or hardware issue. If the issue persists, report the error to Broadcom technical support.

4.1.4.15 elx_mes0332: IOCB wait issue failed, Data <value>

DESCRIPTION: A driver-issued I/O failed to complete in polling mode.

DATA: (1) error value.

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.4.16 elx_mes0334: Unknown IOCB command

DESCRIPTION: Received an unknown IOCB command completion.

DATA: (1) type, (2) ulpCommand, (3) ulpStatus, (4) ulpIoTag, (5) ulpContext

ACTION: This error could indicate a software driver or firmware issue. If this issue persists, report the error to Broadcom technical support.

4.1.4.17 elx_mes0335: Unknown IOCB command

DESCRIPTION: Received an unknown IOCB command completion.

DATA: (1) ulpCommand, (2) ulpStatus, (3) ulpIoTag, (4) ulpContext

ACTION: This error could indicate a software driver or firmware issue. If this issue persists, report the error to Broadcom technical support.

4.1.4.18 elx_mes0338: IOCB wait timeout error - no wake response Data <value> <value>

DESCRIPTION: Driver issued I/O did not get a wake signal in polling mode.

DATA: (1) wait time (2) wake value

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.4.19 elx_mes0340: Adapter temperature is OK now

DESCRIPTION: The adapter temperature has reverted to normal range.

DATA: Temperature in Celsius

ACTION: No action needed, informational.

4.1.4.20 elx_mes0341: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <un.ulpWord[3]>

DESCRIPTION: No additional pre-allocated buffers are available to handle unsolicited buffers.

DATA: None

ACTION: Verify that this port is not being managed by multiple ports.

4.1.4.21 elx_mes0342: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <unsli3.sli3Words>

DESCRIPTION: There was a multiple IOCB unsolicited command and sufficient buffer space cannot be allocated for it.

DATA: None

ACTION: None required.

4.1.4.22 elx_mes0343: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <un.ulpWord[3]>

DESCRIPTION: No additional pre-allocated buffers are available to handle unsolicited buffers.

DATA: None

ACTION: None required.

4.1.4.23 elx_mes0344: Ring <ringno> Cannot find buffer for an unsolicited iocb tag <unsli3.sli3Words[7]>

DESCRIPTION: No additional pre-allocated buffers are available to handle unsolicited buffers.

DATA: None

ACTION: None required.

4.1.4.24 elx_mes0345: Resetting board due to mailbox timeout

DESCRIPTION: A mailbox command failed to complete. The driver is resetting the port.

DATA: None

ACTION: If the mailbox command fails again, set the lpfc_log_verbose to LOG_MBOX and retry.

4.1.4.25 elx_mes0346: Ring <ring number> handler: unexpected ASYNC_STATUS evt_code <evtcode>

DESCRIPTION: The adapter received an asynchronous event that was not a temperature event.

DATA: None

ACTION: None required.

4.1.4.26 elx_mes0347: Adapter is very hot, please take corrective action. Temperature: <value> Celsius

DESCRIPTION: The adapter temperature is above normal range

DATA: Temperature in Celsius

ACTION: Shutdown and remove the adapter. Contact customer support.

4.1.4.27 elx_mes0348: NameServer login: node freed

DESCRIPTION: The enable mode failed to free up the nameserver login.

DATA: None

ACTION: None required.

4.1.4.28 elx_mes0349: rc should be MBX_SUCCESS

DESCRIPTION: The next mailbox command on the mailbox queue has failed.

DATA: None

ACTION: None required.

4.1.4.29 elx_mes0350: rc should have been MBX_BUSY

DESCRIPTION: Attempting to unregister a default RPI from an interrupt context and the mailbox state is not busy.

DATA: None

ACTION: None required.

4.1.4.30 elx_mes0352: Config MSI mailbox command failed, mbxCmd <u.mb.mbxCommand>, mbxStatus <u.mb.mbxStatus>

DESCRIPTION: The mailbox command sent to the firmware to configure the adapter to use MSI-X has failed.

DATA: None

ACTION: Ensure the hardware platform supports MSI-X.

4.1.4.31 elx_mes0359: Not a valid slow-path completion event: majorcode=<value>, minorcode=<value>

DESCRIPTION: SLI-4: The EQE is not valid.

DATA: None

ACTION: None required.

4.1.4.32 elx_mes0360: Unsupported EQ count. <entry_count>

DESCRIPTION: The firmware cannot create an event queue of this size.

DATA: None

ACTION: None required.

4.1.4.33 elx_mes0361: Unsupported CQ count. <entry_count>

DESCRIPTION: The firmware cannot create a completion queue of this size.

DATA: None

ACTION: None required.

4.1.4.34 elx_mes0362: Unsupported MQ count. <entry_count>

DESCRIPTION: The firmware cannot create MQ count of this size.

DATA: None

ACTION: None required.

4.1.4.35 elx_mes0364: Invalid param

DESCRIPTION: SLI-4: The post SGL function was passed an invalid XRI.

DATA: None

ACTION: None required.

4.1.4.36 elx_mes0365: Slow-path CQ identifier <cqid> does not exist

DESCRIPTION: The Completion Queue ID passed in the Event Queue entry does not reference a valid completion queue.

DATA: None

ACTION: None required.

4.1.4.37 elx_mes0366: Not a valid fast-path completion event: majorcode=<major code hex>, minorcode=<minor code hex>

DESCRIPTION: The major or minor code in the Event Queue field is not valid.

DATA: None

ACTION: None required.

4.1.4.38 elx_mes0367: Fast-path completion queue does not exist

DESCRIPTION: The fast path completion queue referenced by the CQID does not exist.

DATA: None

ACTION: None required.

4.1.4.39 elx_mes0368: Miss-matched fast-path completion queue identifier: eqcqid=<cqid>, fcpcqid=<queue_id>

DESCRIPTION: The CQID in the event queue entry does not match the fcp_cqid that was passed into the routine.

DATA: None

ACTION: None required.

4.1.4.40 elx_mes0369: No entry from fast-path completion queue fcpcqid=<queue_id)

DESCRIPTION: No completions exist in the completion queue referenced by fcpcqid.

DATA: None

ACTION: None required.

4.1.4.41 elx_mes0370: Invalid completion queue type <type>

DESCRIPTION: The event queue entry is not for a mailbox or a work queue entry.

DATA: None

ACTION: None required.

4.1.4.42 elx_mes0371: No entry from the CQ: identifier <queue_id>, type <type>

DESCRIPTION: No completion queue event exists for this event queue entry.

DATA: None

ACTION: None required.

4.1.4.43 elx_mes0372: iotag <iotag> is out of range: max iotag (<sli.last_iotag>)

DESCRIPTION: The IOCB lookup cannot be performed because the iocb_tag is out of range.

DATA: None

ACTION: None required.

4.1.4.44 elx_mes0376: READ_REV Error. SLI Level <sli_rev> FCoE enabled <hba_flag & HBA_FCOE_SUPPORT>

DESCRIPTION: This SLI-4 only adapter setup function was called for a non-SLI-4 device.

DATA: None

ACTION: None required.

4.1.4.45 elx_mes0377: Error <rc> parsing vpd. Using defaults.

DESCRIPTION: Could not parse the VPD data, so the driver is using the default values.

DATA: None

ACTION: None required.

4.1.4.46 elx_mes0381: Error <rc> during queue setup.

DESCRIPTION: Could not set up all the queues that the driver requires to exchange I/Os with the adapter.

DATA: None

ACTION: Reload the driver.

4.1.4.47 elx_mes0382: READ_SPARAM command failed status <issue status>, mbxStatus <mailbox status>

DESCRIPTION: The READ_SPARAM mailbox command has failed during initialization. The adapter has been set to error state.

DATA: None

ACTION: Perform a dump with hbacmd and then try reloading the driver.

4.1.4.48 elx_mes0383: Error <error> during scsi sgl post operation

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.4.49 elx_mes0384: There is pending active mailbox cmd

DESCRIPTION: The mailbox commands have overlapped. This command should have been added to the mailbox queue.

DATA: None

ACTION: None required.

4.1.4.50 elx_mes0385: rc should have been MBX_BUSY

DESCRIPTION: The completion handler for REG_LOGIN detected the IMMED_UNREG flag and tried to issue the unreg_login command from an interrupt level. The mailbox status should still be busy.

DATA: None

ACTION: None required.

4.1.4.51 elx_mes0387: Failed to allocate an iocbq

DESCRIPTION: Failed to get an IOCBQ from the list of available IOCBQs.

DATA: None

ACTION: None required.

4.1.4.52 elx_mes0388: Not a valid WCQE code: <hex cqe_code>

DESCRIPTION: The event code is invalid. This event is dropped.

DATA: None

ACTION: Verify that the adapter's firmware is current.

4.1.4.53 elx_mes0391: Error during rpi post operation

DESCRIPTION: The driver was trying to post pages to the firmware to keep target login information and encountered a failure.

DATA: None

ACTION: Unload and reload the driver.

4.1.4.54 elx_mes0393: Error <rc> during rpi post operation

DESCRIPTION: The driver was trying to post pages to the firmware to keep target login information and encountered a failure.

DATA: None

ACTION: Unload and reload the driver.

4.1.4.55 elx_mes0394: Failed to allocate CQ_EVENT entry

DESCRIPTION: The asynchronous event handler was unable to allocate an event queue entry to which to transfer the asynchronous event.

DATA: None

ACTION: This could be a V-LINK clear from the switch or a fatal error from the firmware. Perform a dump from the OneCommand CNA Manager application.

4.1.4.56 elx_mes0395: The mboxq allocation failed

DESCRIPTION: The asynchronous link event handler could not allocate a mailbox command to issue the READ_LA (read link attention) mailbox command.

DATA: None

ACTION: None required.

4.1.4.57 elx_mes0396: The lpfc_dmabuf allocation failed

DESCRIPTION: The asynchronous link event handler could not allocate a DMA buffer for the mailbox command to issue the READ_LA (read link attention) mailbox command.

DATA: None

ACTION: None required.

4.1.4.58 elx_mes0397: The mbuf allocation failed

DESCRIPTION: The asynchronous link event handler could not allocate DMA-able memory for the READ_LA mailbox command.

DATA: None

ACTION: None required.

4.1.4.59 elx_mes0398: Invalid link fault code: < hex link_fault>

DESCRIPTION: The attempt to read the link attention register has returned an unknown value.

DATA: None

ACTION: None required.

4.1.4.60 elx_mes0399: Invalid link attention type: <hex link_type>

DESCRIPTION: The READ_LA mailbox command has returned an invalid link type.

DATA: None

ACTION: None required.

4.1.5 Initialization Events (0400 to 0599)

4.1.5.1 elx_mes0400: Phys Attribute Count Exceeded, Max <value>, Actual <value>

DESCRIPTION: Too many driver configuration parameters have been set. The limit is given as Max.

DATA: (1) Maximum number (2) Actual number

ACTION: Reduce the number of actual parameters.

4.1.5.2 elx_mes0402: Cannot find virtual addr for buffer tag on ring <ringno>

DESCRIPTION: A DMA buffer is not available for this unsolicited command.

DATA: (1) tag, (2) next, (3) prev, (4) postbufq_cnt

ACTION: None required.

4.1.5.3 elx_mes0403: lpfc_nodev_tmo attribute cannot be set to <val>, allowed range is [<LPFC_MIN_DEVLOSS_TMO>, <LPFC_MAX_DEVLOSS_TMO>]

DESCRIPTION: Attempt to set the nodev timeout value is outside the range of the devloss timeout range.

DATA: None

ACTION: Set the nodev timeout between the minimum and maximum of the devloss timeout range.

4.1.5.4 elx_mes0404: Config Param <value> set to <value>

DESCRIPTION: The driver is setting a persistent vPort parameter to a different value.

DATA: (1) New value

ACTION: None. This message is notification only.

4.1.5.5 elx_mes0405: Config Param <value> set to <value>

DESCRIPTION: The driver is setting a persistent vPort parameter to a different value.

DATA: (1) New value

ACTION: None. This message is notification only.

4.1.5.6 elx_mes0406: Adapter maximum temperature exceeded (<temperature>), taking this port offline

DESCRIPTION: The driver has received an error from the adapter indicating that the maximum allowable temperature has been exceeded.

DATA: (1) work_hs, (2) work_status[0], (3) work_status[1]

ACTION: Make sure that the server fans are not blocked. Shut down the server if the airflow is restricted.

4.1.5.7 elx_mes0408: Cannot create debugfs root

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.5.8 elx_mes0409: Cannot create debugfs nodelist

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.5.9 elx_mes0410: Cannot find virtual addr for mapped buf on ring <ringno>

DESCRIPTION: The driver cannot find the specified buffer in its mapping table. Thus it cannot find the virtual address needed to access the data.

DATA: (1) phys, (2) next, (3) prev, (4) postbufq_cnt

ACTION: This error could indicate a software driver or firmware issue. If the issue persists report the error to Broadcom technical support.

4.1.5.10 elx_mes0411: Cannot create debugfs hbqinfo

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.5.11 elx_mes0412: Cannot create debugfs hba

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.5.12 elx_mes0413: Cannot create debugfs dumpHBASlim

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.5.13 elx_mes0414: Cannot create debugfs dumpHostSlim

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.5.14 elx_mes0415: Cannot create debugfs slow_ring trace

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.5.15 elx_mes0416: Cannot create debugfs slow_ring buffer

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.5.16 elx_mes0417: Cannot create debugfs

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.5.17 elx_mes0418: Cannot create debugfs disc trace buffer

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.5.18 elx_mes0419: Cannot create debugfs discovery trace

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.5.19 elx_mes0423: Vport Attribute Instance Error. Defaulting lpfc_#attr to <value>, error value <value>, allowed range is [min, max]

DESCRIPTION: A vPort attribute was set out of range. The driver reset the parameter to its default.

DATA: None

ACTION: Set the module parameter between the minimum and maximum values.

4.1.5.20 elx_mes0424: Vport Attribute Count Exceeded, Max <value>, Actual <value>

DESCRIPTION: The total number of vPort attributes set exceeded the max allowed.

DATA: None

ACTION: Reduce the number of set attributes to below the maximum.

4.1.5.21 elx_mes0425: lpfc_"#attr" attribute cannot be set to <value>, allowed range is [min, max]

DESCRIPTION: Driver attribute lpfc_#attr was defined with an out-of-range value.

DATA: None

ACTION: Set the parameter between the minimum and maximum values.

4.1.5.22 elx_mes0426: lpfc_enable_auth attribute cannot be set to <value>, allowed range is [<min>, <max>]

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.5.23 elx_mes0427: Cannot re-enable interrupt after slot reset.

DESCRIPTION: The driver was not able to enable the interrupt after an adapter reset.

DATA: None

ACTION: Unload and reload the driver.

4.1.5.24 elx_mes0430: PM resume Failed to enable interrupt

DESCRIPTION: The driver's power management resume function could not enable the interrupt.

DATA: None

ACTION: Perform another PM suspend and resume or adapter reset.

4.1.5.25 elx_mes0431: Failed to enable interrupt.

DESCRIPTION: The driver failed to start the interrupt.

DATA: None

ACTION: Unload and reload the driver.

4.1.5.26 elx_mes0433: Wakeup on signal: rc=<rc>

DESCRIPTION: A signal other than the LPFC_DATA_READY was received on the worker thread.

DATA: None

ACTION: Unload and reload the driver.

4.1.5.27 elx_mes0434: PM resume failed to start worker thread: error=<error>.

DESCRIPTION: The driver's power management resume function could not start the worker thread.

DATA: None

ACTION: Unload and reload the driver.

4.1.5.28 elx_mes0435: Adapter failed to get Option ROM version status <rc>.

DESCRIPTION: The driver could not read the adapter's option ROM.

DATA: None

ACTION: Reset the adapter. Make sure that the adapter's firmware is current.

4.1.5.29 elx_mes0436: Adapter failed to init, timeout, status reg <status>

DESCRIPTION: The adapter failed during power-up diagnostics after it was reset.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

4.1.5.30 elx_mes0437: Adapter failed to init, chipset, status reg <status>

DESCRIPTION: The adapter failed during power-up diagnostics after it was reset.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

4.1.5.31 elx_mes0438: Adapter failed to init, chipset, status reg <status>

DESCRIPTION: The adapter failed during power-up diagnostics after it was reset.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

4.1.5.32 elx_mes0439: Adapter failed to init, mbxCmd <mbxCommand> READ_REV, mbxStatus <mbxStatus>

DESCRIPTION: Adapter initialization failed when issuing a READ_REV mailbox command.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

4.1.5.33 elx_mes0440: Adapter failed to init, READ_REV has missing revision information

DESCRIPTION: A firmware revision initialization error was detected.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. Update the firmware. If the issue persists, report the error to Broadcom technical support.

4.1.5.34 elx_mes0442: Adapter failed to init, mbxCmd <mbxCommand> CONFIG_PORT, mbxStatus <mbxStatus>

DESCRIPTION: Adapter initialization failed when issuing a CONFIG_PORT mailbox command.

DATA: (1) hbainit

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

4.1.5.35 elx_mes0445: Firmware initialization failed.

DESCRIPTION: The driver was unable to initialize the hardware.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

4.1.5.36 elx_mes0446: Adapter failed to init, mbxCmd <mbxCommand> CFG_RING, mbxStatus <mbxStatus>, ring <num>

DESCRIPTION: Adapter initialization failed when issuing a CFG_RING mailbox command.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

4.1.5.37 elx_mes0448: Adapter failed to init, mbxCmd <mbxCommand> READ_SPARM, mbxStatus <mbxStatus>

DESCRIPTION: Adapter initialization failed when issuing a READ_SPARM mailbox command.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

4.1.5.38 elx_mes0449: Phys attribute Instance Error. Defaulting to lpfc_#attr to <value>. Allowed range is [min, max]

DESCRIPTION: A physical device attribute has an out-of-range value. The driver is correcting it.

DATA: (1) value written, (2) minimum value, (3) maximum value

ACTION: Write the default value.

4.1.5.39 elx_mes0450: lpfc_%attr attribute cannot be set to <value>, allowed range is [%min, %max]

DESCRIPTION: Sysfs attribute value written exceeds attribute range.

DATA: (1) attribute name, (2) value written, (3) minimum value, (3) maximum value

ACTION: Write a value within the supported range.

4.1.5.40 elx_mes0451: Failed to enable interrupt

DESCRIPTION:

DATA: None.

ACTION: None required.

4.1.5.41 elx_mes0453: Adapter failed to init, mbxCmd <mbxCommand> READ_CONFIG, mbxStatus<mbxStatus>

DESCRIPTION: Adapter initialization failed when issuing a READ_CONFIG mailbox command.

DATA: None

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

4.1.5.42 elx_mes0456: Adapter failed to issue ASYNCEVT_ENABLE mbox status <rc>.

DESCRIPTION: The mailbox command to enable an asynchronous event notification failed.

DATA: None

ACTION: Make sure the adapter firmware is current. Reload the driver.

4.1.5.43 elx_mes0457: Adapter Hardware Error

DESCRIPTION: The driver received an interrupt indicating a possible hardware issue.

Data: (1) status, (2) status1, (3) status2

ACTION: This error could indicate a hardware or firmware issue. If the issue persists, report the error to Broadcom technical support.

4.1.5.44 elx_mes0462: Too many cmd / rsp ring entries in SLI2 SLIM Data: <values> /home/pely/svn_linux_lpfc_upstream/branches/vmware-8.2.1-ucna/kernel/lpfc/lpfc_sli.c:

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.5.45 elx_mes0472: Unknown PCI error state: <value>

DESCRIPTION: The PCI bus has detected an error.

DATA: (1) state value

ACTION: The driver resets the adapter and attempts recovery. If the issue persists, contact Broadcom technical support.

4.1.5.46 elx_mes0474: Unable to allocate memory for issuing "MBOX_CONFIG_MSI command”

DESCRIPTION: Mailbox memory pool allocation error.

DATA: None

ACTION: None required.

4.1.5.47 elx_mes0475: Not configured for supporting MSI-X cfg_use_msi: <cfg_use_msi>.

DESCRIPTION: The lpfc_use_msi module parameter should have been set to 2.

DATA: None

ACTION: Set module parameter lpfc_use_msi=2.

4.1.5.48 elx_mes0476: HBA not supporting SLI-3 or later SLI Revision: <sli_rev>.

DESCRIPTION: The adapter does not support SLI-3 or SLI-4.

DATA: None

ACTION: This adapter does not support MSI. Set lpfc_use_msi=0.

4.1.5.49 elx_mes0479: Deferred Adapter Hardware Error

DESCRIPTION: An adapter hardware error was sent to the driver.

DATA: (1) work_hs, (2) work_status[0], (3) work_status[1]

ACTION: Perform a dump using hbacmd.

4.1.5.50 elx_mes0482: Illegal interrupt mode

DESCRIPTION: The driver could not set MSI-X, MSI or INTx interrupt modes.

DATA: None

ACTION: This could be a server issue. Reboot. If this issue persists, report the error to Broadcom technical support.

4.1.5.51 elx_mes0483: Invalid link-attention link speed: <value>, bf_get(lpfc_acqe_link_speed, acqe_link).

DESCRIPTION: The link speed reported in the link attention interrupt is invalid.

DATA: None

ACTION: Check the switch configuration.

4.1.5.52 elx_mes0492: Unable to allocate memory for issuing SLI_CONFIG_SPECIAL mailbox command

DESCRIPTION: A memory allocation fault occurred when issuing a mailbox.

DATA: None

ACTION: This could be a transient error. If this issue persists, report the error to Broadcom technical support.

4.1.5.53 elx_mes0493: SLI_CONFIG_SPECIAL mailbox failed with status <rc>.

DESCRIPTION: The mailbox command failed.

DATA: None

ACTION: Make sure the adapter’s firmware is current. Unload and reload the driver.

4.1.5.54 elx_mes0494: Unable to allocate memory for issuing "SLI_FUNCTION_RESET mailbox command"

DESCRIPTION: Mailbox memory pool allocation error.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.55 elx_mes0495: SLI_FUNCTION_RESET mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>.

DESCRIPTION: Mailbox command failed.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.56 elx_mes0496: Failed allocate slow-path EQ

DESCRIPTION: The event queue for the slow path was not allocated.

DATA: None

ACTION: Unload and reload the driver.

4.1.5.57 elx_mes0497: Failed allocate fast-path EQ

DESCRIPTION: The event queue for the fast path was not allocated.

DATA: None

ACTION: Unload and reload the driver.

4.1.5.58 elx_mes0498: Adapter failed to init, mbxCmd <cmd> INIT_LINK, mbxStatus <status>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.5.59 elx_mes0499: Failed allocate fast-path FCP CQ (<fcp_cqidx>).

DESCRIPTION: The completion queue event for the fast path could not be allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.60 elx_mes0500: Failed allocate slow-path mailbox CQ

DESCRIPTION: The driver failed to allocate the slow-path mailbox completion queue.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.61 elx_mes0501: Failed allocate slow-path ELS CQ

DESCRIPTION: The driver failed to allocate the slow-path extended link service completion queue.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.62 elx_mes0503: Failed allocate fast-path FCP

DESCRIPTION: The driver failed to allocate the fast-path FCP.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.63 elx_mes0504: Failed allocate slow-path ELS WQ

DESCRIPTION: The driver failed to allocate the slow-path extended link service WQ.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.64 elx_mes0505: Failed allocate slow-path MQ

DESCRIPTION: The driver failed to allocate the slow-path MQ.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.65 elx_mes0506: Failed allocate receive HRQ

DESCRIPTION: The driver failed to allocate the receive HRQ.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.66 elx_mes0507: Failed allocate receive DRQ

DESCRIPTION: The driver failed to allocate the receive DRQ.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.67 elx_mes0520: Slow-path EQ not allocated

DESCRIPTION: The slow-path event queue is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.68 elx_mes0521: Failed setup of slow-path EQ rc = <value>

DESCRIPTION: The slow-path event queue setup failed with status rc.

DATA: (1) status code

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.69 elx_mes0522: Fast-path EQ <fcp_eqidx> not allocated

DESCRIPTION: The fast-path event queue is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.70 elx_mes0523: Failed setup of fast-path EQ <fcp_eqidx>, rc = <rc>

DESCRIPTION: The fast-path event queue setup failed.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.71 elx_mes0526: Fast-path FCP CQ <fcp_cqidx> not allocated

DESCRIPTION: The fast-path FCP is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.72 elx_mes0527: Failed setup of fast-path FCP CQ <fcp_cqidx>, rc = <rc>

DESCRIPTION: The fast-path FCP completion queue setup failed.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.73 elx_mes0528: Mailbox CQ not allocated

DESCRIPTION: The mailbox completion queue is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.74 elx_mes0529: Failed setup of slow-path mailbox CQ: rc = <value>

DESCRIPTION: The driver failed to setup the Completion Queue. Failure code is reported.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.5.75 elx_mes0530: ELS CQ not allocated

DESCRIPTION: The extended link service completion queue is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.76 elx_mes0531: Failed setup of slow-path ELS CQ: rc = <value>

DESCRIPTION: The extended link service completion queue is allocated, but failed initial setup.

DATA: (1) status

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.77 elx_mes0534: Fast-path FCP WQ <fcp_eqidx> not allocated

DESCRIPTION: The fast-path FCP WQ is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.78 elx_mes0535: Failed setup of fast-path FCP WQ <fcp_wqidx>, rc = <rc>

DESCRIPTION: The fast-path FCP WQ setup failed.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.79 elx_mes0536: Slow-path ELS WQ not allocated

DESCRIPTION: The slow-path extended link service WQ is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.80 elx_mes0537: Failed setup of slow-path ELS WQ: rc = <value>

DESCRIPTION: The driver failed to setup the Work Queue. A failure code is reported.

DATA: (1) (2) (3)

ACTION: Software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.5.81 elx_mes0538: Slow-path MQ not allocated

DESCRIPTION: The slow-path MQ is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.82 elx_mes0539: Failed setup of slow-path MQ: rc = <value>

DESCRIPTION: The slow-path MQ is allocated, but failed initial setup.

DATA: (1) status

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.83 elx_mes0540: Receive Queue not allocated

DESCRIPTION: The Receive Queue is not allocated.

DATA: None

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.84 elx_mes0541: Failed setup of Receive Queue: rc = <value>

DESCRIPTION: The Receive Queue is allocated, but failed setup.

DATA: (1) status

ACTION: The driver fails to load. Contact Broadcom technical support.

4.1.5.85 elx_mes0542: lpfc_create_static_vport failed to allocate mailbox memory

DESCRIPTION: The driver failed to allocate the mailbox memory for vPort creation.

DATA: None

ACTION: Static VPorts does not load. Contact Broadcom technical support.

4.1.5.86 elx_mes0543: lpfc_create_static_vport failed to allocate vport_info

DESCRIPTION: The driver failed to allocate the VPort_info.

DATA: None

ACTION: Static VPorts does not load. Contact Broadcom technical support.

4.1.5.87 elx_mes0545: lpfc_create_static_vport bad information header <value> <value>, le32_to_cpu(vport_info->signature), le32_to_cpu(vport_info->rev) & VPORT_INFO_REV_MASK);

DESCRIPTION: Invalid information header; the signature or revision is invalid.

DATA: None

ACTION: Static VPorts does not load. Contact Broadcom technical support.

4.1.5.88 elx_mes0582: Error <rc> during sgl post operation

DESCRIPTION: The SGL post operation failed.

DATA: None

ACTION: None required.

4.1.5.89 elx_mes0602: Failed to allocate CQ_EVENT entry

DESCRIPTION: The driver failed to allocate a CQ_EVENT entry.

DATA: None

ACTION: None required.

4.1.5.90 elx_mes0603: Invalid work queue CQE subtype <subtype>

DESCRIPTION: Invalid work queue CQE.

DATA: None

ACTION: None required.

4.1.6 FCP Traffic History (0700 to 0799)

4.1.6.1 elx_mes0700: Bus Reset on target <i> failed

DESCRIPTION: The bus reset for the specified target failed.

DATA: None

ACTION: None required.

4.1.6.2 elx_mes0706: IOCB Abort failed - outstanding <value> failed <value>

DESCRIPTION: The driver did not recover all I/O following a reset task management command.

DATA: (1) outstanding I/O count (2) number of unrecovered I/O

ACTION: The reset call fails. ESXi tries to recover.

4.1.6.3 elx_mes0713: SCSI layer issued Device Reset (<value>, <value>) reset status <value> flush status <value>

DESCRIPTION: A device reset has completed on (tgt, lun). Status values are displayed.

DATA: (1) tgt (2) lun (3) task mgmt status (4) flush status

ACTION: None required.

4.1.6.4 elx_mes0714: SCSI layer issued bus reset

DESCRIPTION: The SCSI layer is requesting the driver to abort all I/Os to all targets on this adapter.

DATA: (1) ret

ACTION: Check the state of the targets in question.

4.1.6.5 elx_mes0717: FCP command <value> residual underrun converted to error Data: <values>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.6.6 elx_mes0718: Unable to dma_map single request_buffer: <value>

DESCRIPTION: The driver could not map a single virtual address to a DMA address.

DATA: (1) DMA mapping error

ACTION: None. The driver fails the I/O back to ESXi.

4.1.6.7 elx_mes0720: FCP command <value> residual overrun error. Data: <values>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.6.8 elx_mes0721: Device Reset rport failure: rdata <rdata>

DESCRIPTION: The reset of the Rport failed.

DATA: None

ACTION: None required.

4.1.6.9 elx_mes0724: I/O flush failure for context <cntx> on <tgt:lun> cnt <value>

DESCRIPTION: The I/O flush to the {LUN, TARGET, or HOST} has failed.

DATA: (1) count of unrecovered I/O

ACTION: None required. The reset is retried.

4.1.6.10 elx_mes0727: TMF <cmd> to TGT <TGT#> LUN <LUN#> failed (<ulpStatus>, <ulpWord[4]>)

DESCRIPTION: The task management function command failed.

DATA: None

ACTION: None required. The task management function command gets retried.

4.1.6.11 elx_mes0748: Abort handler timed out waiting for abort to complete:ret <status> ID <target id> LUN <lun id> snum <serial number>

DESCRIPTION: The abort handler timed out waiting for abort to complete.

DATA: None

ACTION: None required.

4.1.6.12 elx_mes0798: Device Reset rport failure: rdata <value>

DESCRIPTION: Driver failed a device reset - no rdata buffer.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.7 Node Table Events (0900 to 0999)

4.1.7.1 elx_mes0915: Register VPI failed: <mbxStatus>

DESCRIPTION: Could not register the VPI.

DATA: None

ACTION: None required.

4.1.8 Security Events (1000 to 1099)

4.1.8.1 elx_mes1000: Authentication is enabled but authentication service is not running

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.2 elx_mes1005: AUTHENTICATION_FAILURE Nport: <port>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.3 elx_mes1006: Bad Name tag in auth message

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.4 elx_mes1007: Bad Name length in auth message

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.5 elx_mes1008: Bad Number of Protocols

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.6 elx_mes1009: Bad param type

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.7 elx_mes1010: Bad Tag 1

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.8 elx_mes1011: Auth_neg no has function chosen.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.9 elx_mes1012: Auth_negotiate Bad Tag 2

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.10 elx_mes1013: Auth_negotiate no DH_group found.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.11 elx_mes1014: dhchap challenge bad name tag.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.12 elx_mes1015: dhchap challenge bad name length.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.13 elx_mes1016: dhchap challenge Hash ID not Supported.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.14 elx_mes1017: dhchap challenge could not find DH Group.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.15 elx_mes1018: dhchap challenge No Public key for non-NULL DH Group.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.16 elx_mes1021: ERROR: attempted to queue security work, when no workqueue created.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.17 elx_mes1028: Start Authentication: No buffers

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.18 elx_mes1029: Reauthentication Failure

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.19 elx_mes1031: Start Authentication: Get config failed.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.20 elx_mes1032: Start Authentication: get config timed out.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.21 elx_mes1034: Not Expecting Challenge - Rejecting Challenge.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.22 elx_mes1035: Transport ID does not math - Rejecting Challenge.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.23 elx_mes1036: Authentication transaction reject - re-auth request reason <value> exp <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.24 elx_mes1037: Authentication transaction reject - restarting authentication. reason <value> exp <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.25 elx_mes1039: Not Expecting Reply - rejecting. State <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.26 elx_mes1040: Bad Reply trans_id - rejecting. Trans_id: <value> Expecting: <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.27 elx_mes1043: Authentication LS-RJT

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.28 elx_mes1045: Issue AUTH_NEG failed. Status: <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.29 elx_mes1048: Issue AUTH_REJECT failed.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.30 elx_mes1049: Authentication is enabled but authentication service is nor running

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.31 elx_mes1050: Authentication mode is disabled, but is required by the fabric.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.32 elx_mes1053: Start Authentication: Security service offline.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.33 elx_mes1055: Authentication parameter is disabled, but is required by the fabric.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.34 elx_mes1056: Authentication mode is disabled, but is required by the fabric.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.8.35 elx_mes1057: Authentication transaction reject. reason <value> exp <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9 Miscellaneous and FCoE Events (1200 to 1299)

4.1.9.1 elx_mes1201: Failed to allocate dfc_host

DESCRIPTION: The driver failed to allocate a DFC host and bind it to the management stack.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.9.2 elx_mes1209: C_CT Request error Data: <value> <value>

DESCRIPTION: IOCTL common transport response error - the driver is failing the IOCTL request.

DATA: (1) response buffer flag (2) Data Size

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.9.3 elx_mes1210: Invalid cmd size: <cmd value> <cmdsz value> <rspsz value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.4 elx_mes1211: genreq alloc failed: <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.5 elx_mes1213: FCoE cmd overflow: <off value> + <cnt value> > <cmdsz value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.6 elx_mes1214: Cannot issue FCoE cmd, SLI not active: <off value> rc = -EACCESS

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.7 elx_mes1215: Cannot issue FCoE cmd: not ready or not in maint mode <off value> <rc value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.8 elx_mes1216: FCoE IOCB failed: <off value> <rc value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.9 elx_mes1223: menlo_write: couldn't alloc genreq <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.10 elx_mes1230: Could not find buffer for FCoE cmd: <off value> <indmp value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.11 elx_mes1231: bad bpl

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.12 elx_mes1235: Could not find buffer for FCoE cmd: <off value> poff: <value> cnt: <value> mlastcnt: <value> addl: <value> addh: <value> mdsz: <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.13 elx_mes1238: FCoE IOCB failed: <off value> <rc value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.14 elx_mes1240: Unable to allocate command buffer memory.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.15 elx_mes1243: Menlo command error. code=<value>.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.16 elx_mes1244: Unable to allocate response buffer memory.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.17 elx_mes1246: FCoE chip is running golden firmware. Update FCoE chip firmware immediately

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.18 elx_mes1247: FCoE chip is running diagnostic firmware. Operational use suspended.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.19 elx_mes1248: FCoE chip is running unknown firmware.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.20 elx_mes1249: Invalid FRU data found on adapter. Return adapter to Emulex for repair

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.21 elx_mes1250: Menlo command error. code=<value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.22 elx_mes1251: Menlo command error. code=<value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.23 elx_mes1252: Menlo command error. code=<value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.9.24 elx_mes1262: Failed to allocate dfc_host

DESCRIPTION: The driver could not allocate memory to the dfc_host_struct.

DATA: None

ACTION: None required.

4.1.10 Link Events (1300 to 1399)

4.1.10.1 elx_mes1300: Link Down Event in loop back mode

DESCRIPTION: The driver received a link down event while in loopback mode - unexpected event.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.10.2 elx_mes1302: Invalid speed for this board: Reset link speed to auto: <value>

DESCRIPTION: The driver detected an invalid link speed. Resetting the Link to Auto mode.

DATA: (1) Invalid speed detected

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.10.3 elx_mes1303: Link Up Event <eventTag> received Data: <value> <value> <value> <value> <value> <value> <value>

DESCRIPTION: A link up event was received. It is possible for multiple link events to be received together.

DATA:(1) fc_eventTag, (2) granted_AL_PA, (3) UlnkSpeed, (4) alpa_map[0]

ACTION: If numerous link events are occurring, check the physical connections to the Fibre Channel network.

4.1.10.4 elx_mes1305: Link Down Event <eventTag> received Data: <value> <value> <value>

DESCRIPTION: A link down event was received.

DATA: (1) fc_eventTag, (2) hba_state, (3) fc_flag

ACTION: If numerous link events are occurring, check the physical connections to the Fibre Channel network.

4.1.10.5 elx_mes1306: Link Up Event in loop back mode<eventTag> received Data: <value> <value> <value> <value>

DESCRIPTION: Link up notification; configured for loopback.

DATA: (1) fc_eventTag, (2) granted_AL_PA, (3) UlnkSpeed, (4) alpa_map[0]

ACTION: None required.

4.1.10.6 elx_mes1308: Menlo Maint Mode Link up Event <value> rcvd Data: <value> <value> <value>

DESCRIPTION: Link up notification in Menlo maintenance mode.

DATA: (1) fc_eventTag, (2) port_state, (3) vPort fc_flag

ACTION: None required.

4.1.10.7 elx_mes1309: Link Down Event <value> received Data <value> <value> <value>

DESCRIPTION: The port generated a link down event to the host.

DATA: (1) fc_eventTag (2)port_state (3) vPort fc_flag

ACTION: None required.

4.1.10.8 elx_mes1310: Link Up Event npiv not supported in loop topology

DESCRIPTION: Loop topologies are not supported when NPIV is enabled.

DATA: None

ACTION: Put the link into Fabric mode.

4.1.11 Port Setup Events (1400 to 1499)

4.1.11.1 elx_mes1400: Failed to initialize sgl list.

DESCRIPTION: Failed to initialize the SGL list during initialization.

DATA: None

ACTION: Reboot the server. If the issue persists, contact Broadcom technical support.

4.1.11.2 elx_mes1401: Failed to enable pci device.

DESCRIPTION: Failed to enable a PCI device during initialization.

DATA: None

ACTION: Reboot the server. If the issue persists, contact Broadcom technical support.

4.1.11.3 elx_mes1402: Failed to set up pci memory space.

DESCRIPTION: PCI initialization failed.

DATA: None

ACTION: Reboot the server. If the issue persists, contact Broadcom technical support.

4.1.11.4 elx_mes1403: Failed to set up driver resource.

DESCRIPTION: Driver resource initialization failed.

DATA: None

ACTION: None required.

4.1.11.5 elx_mes1404: Failed to set up driver resource.

DESCRIPTION: Driver resource initialization failed.

DATA: None

ACTION: None required.

4.1.11.6 elx_mes1405: Failed to initialize iocb list.

DESCRIPTION: IOCB initialization failed.

DATA: None

ACTION: None required.

4.1.11.7 elx_mes1406: Failed to set up driver resource.

DESCRIPTION: Initialization failed to set up a driver resource.

DATA: None

ACTION: None required.

4.1.11.8 elx_mes1407: Failed to create scsi host.

DESCRIPTION: Initialization failed to create a SCSI host.

DATA: None

ACTION: None required.

4.1.11.9 elx_mes1408: Port Failed POST – portsmphr=<value>, perr=<port error>, sfi=<sfi reg>, nip=<nip reg>, ipc=<ipc reg>, scr1=<value>, scr2=<value>, hscratch=<value>, pstatus=<port status>

DESCRIPTION: The adapter's power-on self-test has failed.

DATA: None

ACTION: Make sure the adapter firmware is up to date. Contact Broadcom technical support if the issue persists after a system reboot.

4.1.11.10 elx_mes1410: Failed to set up pci memory space.

DESCRIPTION: Initialization failed to set up the PCI memory space.

DATA: None

ACTION: None required.

4.1.11.11 elx_mes1411: Failed to set up driver resource.

DESCRIPTION: Initialization failed to set up a driver resource.

DATA: None

ACTION: None required.

4.1.11.12 elx_mes1412: Failed to set up driver resource.

DESCRIPTION: Initialization failed to set up a driver resource.

DATA: None

ACTION: None required.

4.1.11.13 elx_mes1413: Failed to initialize iocb list.

DESCRIPTION: Initialization failed to initialize the IOCB list.

DATA: None

ACTION: None required.

4.1.11.14 elx_mes1414: Failed to set up driver resource.

DESCRIPTION: Initialization failed to set up a driver resource.

DATA: None

ACTION: None required.

4.1.11.15 elx_mes1415: Failed to create scsi host.

DESCRIPTION: Initialization failed to create a SCSI host.

DATA: None

ACTION: None required.

4.1.11.16 elx_mes1416: Failed to allocate sysfs attr

DESCRIPTION: Initialization failed to allocate a sysfs attribute.

DATA: None

ACTION: None required.

4.1.11.17 elx_mes1418: Invalid HBA PCI-device group: <dev_grp>

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

4.1.11.18 elx_mes1419: Invalid HBA PCI-device group: <dev_grp>

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

4.1.11.19 elx_mes1420: Invalid HBA PCI-device group: <dev_grp>

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

4.1.11.20 elx_mes1421: Failed to set up hba

DESCRIPTION: Initialization failed to set up the adapter.

DATA: None

ACTION: None required.

4.1.11.21 elx_mes1422: Unrecoverable Error Detected during POST uerr_lo_reg=<ue lo>, uerr_hi_reg=<ue hi>

DESCRIPTION: The adapter has notified the driver that it has encountered an unrecoverable error.

DATA: None

ACTION: Perform a dump from the OneCommand CNA Manager application. Then, unload and reload the driver.

4.1.11.22 elx_mes1423: HBA Unrecoverable error: uerr_lo_reg=<ue lo>, uerr_hi_reg=<ue hi>, ue_mask_lo_reg=<ue mask lo>, ue_mask_hi_reg=<ue mask hi>

DESCRIPTION: The adapter has notified the driver that it has encountered an unrecoverable error.

DATA: None

ACTION: Perform a dump from the OneCommand CNA Manager application. Then, unload and reload the driver.

4.1.11.23 elx_mes1424: Invalid PCI device group: <pci_dev_grp>

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

4.1.11.24 elx_mes1425: Invalid PCI device group: <pci_dev_grp>

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

4.1.11.25 elx_mes1426: Invalid PCI device group: <pci_dev_grp>

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

4.1.11.26 elx_mes1427: Invalid PCI device group: <pci_dev_grp>

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

4.1.11.27 elx_mes1428: Invalid PCI device group: <pci_dev_grp>

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

4.1.11.28 elx_mes1429: Invalid PCI device group: <pci_dev_grp>

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

4.1.11.29 elx_mes1430: Failed to initialize sgl list.

DESCRIPTION: The driver failed to initialize the SGL list.

DATA: None

ACTION: None required.

4.1.11.30 elx_mes1431: Invalid HBA PCI-device group: <dev_grp>

DESCRIPTION: An invalid adapter PCI-device group was detected.

DATA: None

ACTION: None required.

4.1.11.31 elx_mes1432: Failed to initialize rpi headers.

DESCRIPTION: The RPI headers required by the firmware failed to initialize.

DATA: None

ACTION: None required.

4.1.11.32 elx_mes1476: Failed to allocate sysfs attr.

DESCRIPTION: The driver failed to allocate a sysfs attribute.

DATA: None

ACTION: None required.

4.1.11.33 elx_mes1477: Failed to set up hba

DESCRIPTION: The driver failed to set up an adapter.

DATA: None

ACTION: None required.

4.1.12 IOCTL Events (1600 to 1699)

None.

4.1.13 VPort Events (1800 to 1832)

4.1.13.1 elx_mes1800: Could not issue unreg_vpi

DESCRIPTION: The driver's attempt to unregister a VPI failed.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.13.2 elx_mes1801: Create vport work array FAILED: cannot do scsi_host_get

DESCRIPTION: The driver failed to create a working list of VPorts.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.13.3 elx_mes1802: HBQ <index>: local_hbqGetIdx <index> is > than hbqp->entry_count <count>

DESCRIPTION: An error occurred when processing a queue related to an adapter in a particular slot.

DATA: (1) hbqno, (2) local_hbqGetIdx, (3) entry_count

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.13.4 elx_mes1803: Bad hbq tag. Data: <tag> <count>

DESCRIPTION: An error occurred when processing queue-related tags for an adapter in a particular slot.

DATA: (1) tag, (2) buffer_count

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.13.5 elx_mes1804: Invalid asynchronous event code: <evt code>

DESCRIPTION: The asynchronous event code that the firmware passed to the driver is invalid.

DATA: None

ACTION: None required.

4.1.13.6 elx_mes1805: Adapter failed to init.Data: <command> <status> <queue num>

DESCRIPTION: An error occurred when processing queue-related tags for an adapter in a particular slot.

DATA: (1) mbxCommand, (2) mbxStatus, (3) hbaqno

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.13.7 elx_mes1806: Mbox <command> failed. No vport.

DESCRIPTION: A mailbox command could not be communicated because there was no vPort associated with the mailbox command.

DATA: (1) mbxCommand

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.13.8 elx_mes1807: IOCB <value> failed. No vport

DESCRIPTION: An IOCB command could not be communicated because there was no vPort associated with the mailbox command.

DATA: (1) ulpCommand

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.13.9 elx_mes1808: Create VPORT failed: NPIV is not enabled: SLImode <mode>

DESCRIPTION: The driver failed to create a port because the adapter was in the wrong mode or was not capable of NPIV.

DATA: (1) sli_rev

ACTION: Load the driver with NPIV enabled on an adapter that supports SLI-3.

4.1.13.10 elx_mes1809: Create VPORT failed: Max VPORTs (<vpi>) exceeded.

DESCRIPTION: The driver failed to create a port because the maximum number of ports supported by the driver is exceeded.

DATA: (1) max_vpi

ACTION: No Action. The driver cannot create any more VPorts.

4.1.13.11 elx_mes1810: Create VPORT failed: Cannot get instance number.

DESCRIPTION: The driver failed to allocate resources for an adapter and could not assign an instance number.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.13.12 elx_mes1811: Create VPORT failed: vpi x<vpi>

DESCRIPTION: The driver failed to create a port and had to eliminate all its resources.

DATA: (1) vpi

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.13.13 elx_mes1812: vport_delete failed: Cannot delete physical host

DESCRIPTION: An attempt to delete a port failed because it was to delete a physical port and not a virtual port. Only VPorts on physical ports can be deleted on an NPIV system.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.13.14 elx_mes1813: Create VPORT failed. Cannot get sparam.

DESCRIPTION: The port could not be created because it could not be initialized, possibly due to unavailable resources.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.13.15 elx_mes1814: Mbox <u.mb.mbxCommand> failed, no vport

DESCRIPTION: The vPort field of this mailbox command was not completed.

DATA: None

ACTION: None required.

4.1.13.16 elx_mes1815 Could not issue unreg_did (default rpis)

DESCRIPTION: An attempt to unregister RPI failed.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.13.17 elx_mes1818: VPort failed init, mbxCmd <mailbox command> READ_SPARM mbxStatus <mailbox status>, rc = <status>

DESCRIPTION: A pending mailbox command issued to initialize a port failed.

DATA: (1) mbxCommand, (2) mbxStatus, (3) rc

ACTION: This is a software driver error. If this issue persists, report these issues to Broadcom technical support.

4.1.13.18 elx_mes1820: Unable to select SLI-3. Not supported by adapter.

DESCRIPTION: The adapter is not capable of operating in a given mode.

DATA: None

ACTION: This is an informational message. SLI-3 mode is only available on some adapters. Do not attempt to force an adapter to run in SLI-3 mode if that adapter does not support SLI-3 mode. Adapters that do not support SLI-3 are configured to run in SLI-2 mode. Nevertheless, Broadcom recommends you use the auto setting (0).

4.1.13.19 elx_mes1821: Create VPORT failed. Invalid WWN format

DESCRIPTION: The port could not be created due to an invalid WWNN or WWPN format.

DATA: None

ACTION: Provide a valid WWN when creating VPorts.

4.1.13.20 elx_mes1822: Invalid <name>: <xx: xx: xx: xx: xx: xx: xx: xx>

DESCRIPTION: An invalid WWN was used when creating a vPort.

DATA: (1) type_name, (2) wwn[1], (3) wwn[3], (3) wwn[5], (4) wwn[7]

ACTION: When creating a vPort you must furnish a valid WWN.

4.1.13.21 elx_mes1823: Create VPORT failed. Duplicate WWN on HBA.

DESCRIPTION: The port could not be created because it would duplicate an existing WWNN adapter address. The resources for the port had to be discarded.

DATA: None

ACTION: Provide a unique WWN.

4.1.13.22 elx_mes1825: Vport Created.

DESCRIPTION: This message indicates that a port was created in the system. It is provided at this level to ensure that it always appears at all log levels.

DATA: None

ACTION: No action, informational.

4.1.13.23 elx_mes1826: Vport Disabled.

DESCRIPTION: The port had to be disabled in the system.

DATA: None

ACTION: No action, informational.

4.1.13.24 elx_mes1827: Vport Enabled.

DESCRIPTION: The port was enabled after possible recovery from some errors.

DATA: None

ACTION: No action, informational.

4.1.13.25 elx_mes1828: Vport Deleted.

DESCRIPTION: A vPort was deleted.

DATA: None

ACTION: No action, informational.

4.1.13.26 elx_mes1830: Signal aborted mbxCmd <command>

DESCRIPTION: A pending mailbox command was aborted because the thread received a signal.

DATA: None

ACTION: The command is retried.

4.1.13.27 elx_mes1831: Create VPORT Interrupted.

DESCRIPTION: The port creation process was unexpectedly interrupted at a critical time and the operation was unsuccessful.

DATA: None

ACTION: The process was interrupted while creating a vPort. Retry the command.

4.1.13.28 elx_mes1832: No pending MBOX command to handle.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.14 Extended Link Service Events (1833 to 2800)

4.1.14.1 elx_mes1835: Vport discovery quiesce failed: state <port_state> fc_flags <fc_flag> wait msecs <jiffies_to_msecs(jiffies - start_time)>

DESCRIPTION: Could not pause discovery on this vPort.

DATA: None

ACTION: None required.

4.1.14.2 elx_mes1836: Could not issue unreg_login(all_rpis) status <rc>

DESCRIPTION: The unreg_login cannot be issued.

DATA: None

ACTION: None required.

4.1.14.3 elx_mes1837: vport_delete failed: Cannot delete static vport.

DESCRIPTION: Static VPorts cannot be deleted.

DATA: None

ACTION: None required.

4.1.14.4 elx_mes1838: Failed to INIT_VPI on vpi <vpi> status <rc>

DESCRIPTION: Failed to INIT_VPI.

DATA: None

ACTION: None required.

4.1.14.5 elx_mes1839: Create VPORT failed. vname allocation failed.

DESCRIPTION: The driver failed to allocate a buffer for Virtual Machine name.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.6 elx_mes1840: Delete VPORT cannot proceed at this time due to SCSI layer busy.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.14.7 elx_mes1920: Exec format error, Dropping Link state event

DESCRIPTION: There was no dfchba instance available for a Link State event - dropping.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.8 elx_mes1923: Exec format error, Dropping rscn event

DESCRIPTION: There was no dfchba instance available for a RSCN event - dropping.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.9 elx_mes1926: Exec format error

DESCRIPTION: There was no dfchba instance available for an IOCTL loopback test - dropping.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.10 elx_mes1927: Exec format error, Dropping temp event

DESCRIPTION: There was no dfchba instance available for a Temperature event - dropping.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.11 elx_mes1928: Exec format error, Dropping dump event

DESCRIPTION: There was no dfchba instance available for a Dump event - dropping.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.12 elx_mes1929: Exec format error

DESCRIPTION: There was no dfchba instance available for an IOCTL loopback XRI read - dropping.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.13 elx_mes1934: ENOMEM DMA coherent resource unavailable

DESCRIPTION: The driver failed to allocate a DMA buffer for an IOCTL request.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.14 elx_mes1935: Loopback test did not receive any data

DESCRIPTION: The driver ran a loopback test, but did not receive a response.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.15 elx_mes1936: ENOMEM Kernel resource unavailable

DESCRIPTION: The driver failed to allocate a DMA buffer during a loopback test.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.16 elx_mes1944: ENOMEM kernel memory resource unavailable

DESCRIPTION: The driver failed to allocate a kernel buffer for a timed out I/O request.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.17 elx_mes1949: ENOEXEC NULL parameter passed to function

DESCRIPTION: The driver tried to post a receive buffer, but no receive buffers are available.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.18 elx_mes1950: ENOMEM IOCB resource not available

DESCRIPTION: The driver could not allocate the IOCBs needed to post loopback receive buffers.

DATA: (1) (2) (3)

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.19 elx_mes1951: ENOMEM MBUF resource not available

DESCRIPTION: The driver failed to get the memory buffer needed for a loopback test.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.20 elx_mes1952: ENOMEM DMA resource not available

DESCRIPTION: The driver failed to get the DMA buffers needed for a loopback test

DATA: (1) (2) (3)

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.21 elx_mes1957: EPERM Illegal BDE count [<value>]

DESCRIPTION: The driver received too many receive buffers for a loopback operation.

DATA: (1) receive buffer count

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.22 elx_mes2000: Failed to allocate mbox for READ_FCF cmd

DESCRIPTION: The driver failed to allocate a mailbox for the READ_FCF command.

DATA: None

ACTION: None required.

4.1.14.23 elx_mes2001: Unable to allocate memory for issuing SLI_CONFIG_SPECIAL mailbox command

DESCRIPTION: Unable to allocate memory for issuing the SLI_CONFIG_SPECIAL mailbox command.

DATA: None

ACTION: None required.

4.1.14.24 elx_mes2002: Error Could not grow rpi count

DESCRIPTION: An error occurred because the RPI count could not be increased.

DATA: None

ACTION: None required.

4.1.14.25 elx_mes2007: Only Limited Edition cmd Format supported <iocb.ulpCommand>

DESCRIPTION: SLI-4 only supports the Limited Edition command format.

DATA: None

ACTION: None required.

4.1.14.26 elx_mes2008: Error <rc> posting all rpi headers

DESCRIPTION: The RPI headers could not be posted to the firmware.

DATA: None

ACTION: None required.

4.1.14.27 elx_mes2009: Failed to allocate mbox for ADD_FCF cmd

DESCRIPTION: The driver failed to allocate a mailbox for the ADD_FCF command.

DATA: None

ACTION: None required.

4.1.14.28 elx_mes2010: Resume RPI Mailbox failed status <status>, mbxStatus <mbx status>.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.14.29 elx_mes2011: Unable to allocate memory for issuing SLI_CONFIG_SPECIAL mailbox command

DESCRIPTION: Unable to allocate memory for issuing SLI_CONFIG_SPECIAL mailbox command.

DATA: None

ACTION: None required.

4.1.14.30 elx_mes2012: Mailbox failed, mbxCmd <mbx_cmd> READ_CONFIG, mbxStatus <mbx status>.

DESCRIPTION: The READ_CONFIG mailbox command failed.

DATA: None

ACTION: None required.

4.1.14.31 elx_mes2013: Could not manually add FCF record 0, status <rc>

DESCRIPTION: Could not add an FCF record to the FCF list.

DATA: None

ACTION: None required.

4.1.14.32 elx_mes2014: Invalid command <iocb.ulpCommand>

DESCRIPTION: The IOCB command is invalid.

DATA: None

ACTION: None required.

4.1.14.33 elx_mes2015: Invalid CT command <iocb.ulpCommand>

DESCRIPTION: An Invalid Command-Type in the IOCB is not supported.

DATA: None

ACTION: None required.

4.1.14.34 elx_mes2017: REG_FCFI mbxStatus error <mbx status> HBA state <port_state>.

DESCRIPTION: The REG_FCFI mailbox command has failed.

DATA: None

ACTION: None required.

4.1.14.35 elx_mes2018: REG_VFI mbxStatus error <mbx status> HBA state <port_state>.

DESCRIPTION: The REG_VFI mailbox command has failed.

DATA: None

ACTION: None required.

4.1.14.36 elx_mes2022: INIT VPI Mailbox failed status <status>, mbxStatus <mbxStatus>

DESCRIPTION: The INIT VPI mailbox command has failed.

DATA: None

ACTION: None required.

4.1.14.37 elx_mes2400: Failed to allocate xri for ELS sgl

DESCRIPTION: Initialization failed to allocate XRI for the extended link service SGL.

DATA: None

ACTION: None required.

4.1.14.38 elx_mes2500: EQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>

DESCRIPTION: The mailbox command sent to create the event queue has failed.

DATA: None

ACTION: None required.

4.1.14.39 elx_mes2501: CQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>

DESCRIPTION: The mailbox command sent to create the completion queue has failed.

DATA: None

ACTION: None required.

4.1.14.40 elx_mes2502: MQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>

DESCRIPTION: The mailbox command sent to create the mailbox queue has failed.

DATA: None

ACTION: None required.

4.1.14.41 elx_mes2503: WQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>

DESCRIPTION: The mailbox command sent to create the work queue has failed.

DATA: None

ACTION: None required.

4.1.14.42 elx_mes2504: RQ_CREATE mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>

DESCRIPTION: The mailbox command sent to create the receive queue has failed.

DATA: None

ACTION: None required.

4.1.14.43 elx_mes2505: EQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>

DESCRIPTION: The mailbox command sent to delete the event queue has failed.

DATA: None

ACTION: None required.

4.1.14.44 elx_mes2506: CQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>

DESCRIPTION: The mailbox command sent to delete the completion queue has failed.

DATA: None

ACTION: None required.

4.1.14.45 elx_mes2507: MQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>

DESCRIPTION: The mailbox command sent to delete the mailbox queue has failed.

DATA: None

ACTION: None required.

4.1.14.46 elx_mes2508: WQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>

DESCRIPTION: The mailbox command sent to delete the work queue has failed.

DATA: None

ACTION: None required.

4.1.14.47 elx_mes2509: RQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>

DESCRIPTION: The mailbox command sent to delete the receive queue has failed.

DATA: None

ACTION: None required.

4.1.14.48 elx_mes2510: RQ_DESTROY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>

DESCRIPTION: The mailbox command sent to delete the receive queue has failed.

DATA: None

ACTION: None required.

4.1.14.49 elx_mes2511: POST_SGL mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>

DESCRIPTION: The mailbox command sent to post the SGL pages to the firmware has failed.

DATA: None

ACTION: None required.

4.1.14.50 elx_mes2512: MODIFY_EQ_DELAY mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>

DESCRIPTION: The mailbox command sent to modify the event queue delay in the firmware has failed.

DATA: None

ACTION: None required.

4.1.14.51 elx_mes2513: POST_SGL_BLOCK mailbox command failed status <shdr_status> add_status <shdr_add_status> mbx status <rc>

DESCRIPTION: The mailbox command sent to post the SGL pages to the firmware has failed.

DATA: None

ACTION: None required.

4.1.14.52 elx_mes2514: POST_RPI_HDR mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx status <rc>

DESCRIPTION: The mailbox command sent to post the RPUI header pages to the firmware has failed.

DATA: None

ACTION: None required.

4.1.14.53 elx_mes2515: ADD_FCF_RECORD mailbox failed with status <rc>

DESCRIPTION: The mailbox command to add the FCF record has failed.

DATA: None

ACTION: None required.

4.1.14.54 elx_mes2521: READ_FCF_RECORD mailbox failed with status <shdr_status> add_status <shdr_add_status>, mbx

DESCRIPTION: The READ_FCF_RECORD mailbox command has failed.

DATA: None

ACTION: None required.

4.1.14.55 elx_mes2522: Synchronous READ_FCF_RECORD mailbox failed with status <value> add_status <value>

DESCRIPTION: Driver failed to read the active FCF Record on an FCoE link - FCF may not be available.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.56 elx_mes2523: Allocated DMA memory size (<alloc_len>) is less than the requested DMA memory size (<req_len>)

DESCRIPTION: The ADD_FCF_RECORD mailbox command failed to retrieve the length required from the firmware.

DATA: None

ACTION: None required.

4.1.14.57 elx_mes2524: Failed to get the non-embedded SGE virtual address

DESCRIPTION: The READ_FCF_RECORD mailbox command could not retrieve the scatter/gather entry that was requested.

DATA: None

ACTION: None required.

4.1.14.58 elx_mes2527: Failed to allocate non-embedded SGE array.

DESCRIPTION: The driver failed to allocate the non-embedded scatter/gather entry array.

DATA: None

ACTION: None required.

4.1.14.59 elx_mes2528: Mailbox command <vpi> cannot issue

DESCRIPTION: The mailbox command could not be issued because the mailbox interrupt is disabled.

DATA: (1) mbxCommand, (2) sli_flag, (3) flag

ACTION: None required.

4.1.14.60 elx_mes2529: Mailbox command <vpi> cannot issue

DESCRIPTION:

DATA: (1) mbxCommand, (2) sli_flag, (3) flag

ACTION: None required.

4.1.14.61 elx_mes2530: Mailbox command <vpi> cannot issue

DESCRIPTION: The SLI layer in the driver is inactive.

DATA: (1) mb.mbxCommand, (2) sli_flag, (3) flag

ACTION: None required.

4.1.14.62 elx_mes2531: Mailbox command <cpi> cannot issue

DESCRIPTION:

DATA: (1) mb.mbxCommand, (2) sli_flag, (3) flag

ACTION: None required.

4.1.14.63 elx_mes2532: Mailbox command <vpi> (<mbxCommand>) cannot issue

DESCRIPTION: The mailbox bootstrap code detected that the SLI layer is active.

DATA: (1) sli4_mbox_opcode, (2) sli_flag, (3) MBX_POLL

ACTION: None required.

4.1.14.64 elx_mes2533: Mailbox command <vpi> (<mbxCommand>) cannot issue

DESCRIPTION:

DATA: (1) sli4_mbox_opcode, (2) sli_flag, (3) MBX_NOWAIT

ACTION: None required.

4.1.14.65 elx_mes2535: Unsupported RQ count. (<entry_count>).

DESCRIPTION: The receive queue ring can only be 512, 1024, 2048, or 4096.

DATA: None

ACTION: None required.

4.1.14.66 elx_mes2536: Unsupported RQ count. (<entry_count>).

DESCRIPTION: The receive queue ring can only be 512, 1024, 2048, or 4096.

DATA: None

ACTION: None required.

4.1.14.67 elx_mes2537: Receive Frame Truncated!

DESCRIPTION: The receive unsolicited handler detected a truncated frame.

DATA: None

ACTION: None required.

4.1.14.68 elx_mes2540: Ring <value> handler: unexpected Rctl <value> Type <value> received

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.14.69 elx_mes2543: Mailbox command <vpi> (<mbxCommand>) cannot issue

DESCRIPTION: The mailbox command does not have all of the fields set correctly.

DATA: (1) sli4_mbx_opcode, (2) sli_flag, (3) flag

ACTION: None required.

4.1.14.70 elx_mes2544: Mailbox command <vpi> (<mbxCommand>) cannot issue

DESCRIPTION: The adapter cannot be accessed on the PCI bus.

DATA: (1) sli4_mbx_opcode, (2) sli_flag, (3) flag

ACTION: None required.

4.1.14.71 elx_mes2546: New FCF found index <index> tag <event_tag>

DESCRIPTION: A new FCF has been found.

DATA: None

ACTION: None required.

4.1.14.72 elx_mes2547: Issue FCF scan read FCF mailbox command failed

DESCRIPTION: Could not read the FCF mailbox command from the firmware.

DATA: None

ACTION: None required.

4.1.14.73 elx_mes2548: FCF Table full count <count> tag <event_tag>

DESCRIPTION: The FCF table is full.

DATA: None

ACTION: None required.

4.1.14.74 elx_mes2549: FCF disconnected from network index <index> tag <event_tag>.

DESCRIPTION: The FCF has disconnected from the network.

DATA: None

ACTION: None required.

4.1.14.75 elx_mes2550: UNREG_FCFI mbxStatus error <u.mb.mbxStatus> HBA state <port_state>.

DESCRIPTION: The unregistered FCFI has failed.

DATA: None

ACTION: None required.

4.1.14.76 elx_mes2551: UNREG_FCFI mbox allocation failed HBA state <port_state>.

DESCRIPTION: The allocation for the UNREG_FCFI mailbox command has failed.

DATA: None

ACTION: None required.

4.1.14.77 elx_mes2552: Unregister FCFI command failed rc <rc> HBA state <port_state>.

DESCRIPTION: The unregister FCFI mailbox command has failed.

DATA: None

ACTION: None required.

4.1.14.78 elx_mes2553: lpfc_unregister_unused_fcf failed to read FCF record HBA state.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.14.79 elx_mes2554: Could not allocate memory for fcf record

DESCRIPTION: The driver was unable to allocate memory for the fcf record.

DATA: None

ACTION: None required.

4.1.14.80 elx_mes2555: UNREG_VFI mbxStatus error <u.mb.mbxStatus> HBA state <port_state>

DESCRIPTION: The unregister virtual fabric index mailbox command has failed.

DATA: None

ACTION: None required.

4.1.14.81 elx_mes2556: UNREG_VFI mbox allocation failed HBA state <port_state>

DESCRIPTION: The driver could not allocate memory for UNREG_VFI mailbox command.

DATA: None

ACTION: None required.

4.1.14.82 elx_mes2557: UNREG_VFI issue mbox failed rc <rc> HBA state <port_state>

DESCRIPTION: The driver could not issue the UNREG_VFI mailbox command.

DATA: None

ACTION: None required.

4.1.14.83 elx_mes2558: ADD_FCF_RECORD mailbox failed with status<shdr_status> add_status <shdr_add_status>

DESCRIPTION: The ADD_FCF_RECORD mailbox command has failed.

DATA: None

ACTION: None required.

4.1.14.84 elx_mes2560: Failed to allocate xri for scsi buffer

DESCRIPTION: The driver failed to allocate an XRI for the SCSI buffer.

DATA: None

ACTION: None required.

4.1.14.85 elx_mes2561: Allocated DMA memory size (<alloclen>) is less than the requested DMA memory size (<reqlen>)

DESCRIPTION: The driver could not get the memory required for the number of XRIs that are trying to be posted.

DATA: None

ACTION: None required.

4.1.14.86 elx_mes2562: Failure to allocate an ELS sgl entry: <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.14.87 elx_mes2563: Failure to allocate an ELS mbuf: <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.14.88 elx_mes2564: POST_SGL_BLOCK mailbox command failed status <shdr_status> add_status <shdr_add_status> mbx status <rc>

DESCRIPTION: The list of XRI SGEs failed to register with the firmware.

DATA: None

ACTION: None required.

4.1.14.89 elx_mes2566: Failed to allocate connection table entry

DESCRIPTION: The driver failed to allocate the connection table entry.

DATA: None

ACTION: None required.

4.1.14.90 elx_mes2567: Config region 23 has bad signature

DESCRIPTION: Configuration region 23 has an invalid signature.

DATA: None

ACTION: None required.

4.1.14.91 elx_mes2568: Config region 23 has bad version

DESCRIPTION: Configuration region 23 has an invalid version.

DATA: None

ACTION: None required.

4.1.14.92 elx_mes2572: Failed allocate memory for fast-path per-EQ handle array

DESCRIPTION: The driver failed to allocate memory for the fast-path per-event queue handle array.

DATA: None

ACTION: None required.

4.1.14.93 elx_mes2573: Failed allocate memory for msi-x interrupt vector entries

DESCRIPTION: The driver failed to allocate memory for MSI-X interrupt vector entries.

DATA: None

ACTION: None required.

4.1.14.94 elx_mes2576: Failed allocate memory for fast-path EQ record array

DESCRIPTION: The driver failed to allocate memory for the fast-path event queue record array.

DATA: None

ACTION: None required.

4.1.14.95 elx_mes2577: Failed allocate memory for fast-path CQ record array

DESCRIPTION: The driver failed to allocate memory for the fast-path completion queue record array.

DATA: None

ACTION: None required.

4.1.14.96 elx_mes2578: Failed allocate memory for fast-path WQ record array

DESCRIPTION: The driver failed to allocate memory for the fast-path WQ record array.

DATA: None

ACTION: None required.

4.1.14.97 elx_mes2598: Adapter Link is disabled.

DESCRIPTION: The adapter link is disabled.

DATA: None

ACTION: None required.

4.1.14.98 elx_mes2599: Adapter failed to issue DOWN_LINK mbox command rc <rc>.

DESCRIPTION: The adapter failed to issue a DOWN_LINK mailbox command.

DATA: None

ACTION: None required.

4.1.14.99 elx_mes2600: failed to allocate mailbox memory

DESCRIPTION: The adapter failed to allocate mailbox memory.

DATA: None

ACTION: None required.

4.1.14.100 elx_mes2605: lpfc_dump_static_vport: memory allocation failed

DESCRIPTION: Memory allocation failed.

DATA: None

ACTION: None required.

4.1.14.101 elx_mes2606: No NPIV Fabric support

DESCRIPTION: No NPIV Fabric support.

DATA: None

ACTION: None required.

4.1.14.102 elx_mes2607: Failed to allocate init_vpi mailbox

DESCRIPTION: The adapter failed to allocate the init_vpi mailbox.

DATA: None

ACTION: None required.

4.1.14.103 elx_mes2608: Failed to issue init_vpi mailbox

DESCRIPTION: The driver failed to issue the init_vpi mailbox.

DATA: None

ACTION: None required.

4.1.14.104 elx_mes2609: Init VPI mailbox failed <u.mb.mbxStatus>

DESCRIPTION: Initialization of the VPI mailbox has failed.

DATA: None

ACTION: None required.

4.1.14.105 elx_mes2610: FCF <value> reached driver's book keeping dimension: <value>.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.14.106 elx_mes2619: Config region 23 has bad signature

DESCRIPTION: Configuration region 23 has an invalid signature.

DATA: None

ACTION: None required.

4.1.14.107 elx_mes2620: Config region 23 has bad version

DESCRIPTION: Configuration region 23 has an invalid version.

DATA: None

ACTION: None required.

4.1.14.108 elx_mes2707: Ring <Ring#> handler: Failed to allocate iocb Rctl <fh_rctl> Type <fh_type> received

DESCRIPTION: The driver could not allocate an IOCB with which to associate this received frame.

DATA: None

ACTION: None required.

4.1.14.109 elx_mes2710: PCI channel disable preparing for reset

DESCRIPTION: The driver is resetting the PCI slot for this port - starting preparations.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.110 elx_mes2711: PCI channel permanent disable for failure

DESCRIPTION: The driver has detected a fatal port error - disabling the PCI channel.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.111 elx_mes2712: lpfc_aer support attribute value <value> out of range, allowed values are 0|1, setting it to default value of 1

DESCRIPTION: The lpfc_aer support attribute value was out of range. The adapter is resetting it to the default value of 1.

DATA: None

ACTION: None required.

4.1.14.112 elx_mes2718: Clear Virtual Link Received for VPI <index> tag <event_tag>

DESCRIPTION: A Clear virtual link command was received from the Fabric for this VPI.

DATA: None

ACTION: None required.

4.1.14.113 elx_mes2719: Invalid response length: tgt <TGT_ID> lun <LUN> cmnd <CMD> rsplen <RSPLEN>

DESCRIPTION: The response length for this FCP command is not supported.

DATA: None

ACTION: None required.

4.1.14.114 elx_mes2723 PCI channel I/O abort preparing for recovery

DESCRIPTION: The driver is preparing the port PCI channel for reset/recovery after an I/O error.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.115 elx_mes2726: READ_FCF_RECORD Indicates empty FCF table

DESCRIPTION: The driver requested the firmware provide a list of FCF entries to connect to and the firmware responded that the FCF table is empty.

DATA: None

ACTION: None required.

4.1.14.116 elx_mes2729: Unable to dma_map_single request_buffer: <value>

DESCRIPTION: The driver was unable to map the SCSI command scatter-gather buffer.

DATA: (1) dma mapping error.

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.117 elx_mes2731: Cannot find fabric controller node.

DESCRIPTION: The driver was unable to find the Fabric controller node in its data base.

DATA: None

ACTION: None required.

4.1.14.118 elx_mes2732: Failed to issue INIT_VPI mailbox command.

DESCRIPTION: The driver wanted to send a INIT_VPI mailbox command to initialize a vPort, but failed to send the mailbox command due to the state of the adapter.

DATA: None

ACTION: None required.

4.1.14.119 elx_mes2745: Failed to allocate mbox for requesting FCF rediscover

DESCRIPTION: The driver is trying to rediscover the FCF table, but failed to allocate the memory needed.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.120 elx_mes2746: Requesting for FCF rediscovery failed status <value> add_status <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.14.121 elx_mes2747: Issue FCF scan read FCF mailbox command failed

DESCRIPTION: The driver wanted to send a read FCF record mailbox command to start fast FCF failover on FCF scan, but failed to send the mailbox command due to the state of the adapter.

DATA: None

ACTION: None required.

4.1.14.122 elx_mes2748 Failed to prepare for unregistering HBA's FCF record: rc=<value>

DESCRIPTION: The driver encountered an initialization error when preparing to rescan the FCF tables and needed to unregister an old FCF record.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.123 elx_mes2749 Failed to prepare for unregistering HBA's FCF record: rc=<value>

DESCRIPTION: The driver encountered an initialization error when preparing to unregister an FCF and needed to prepare the command.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.124 elx_mes2751: Adapter failed to restart, status reg <status>, FW Data: A8 <0xA8> AC <0xAC>

DESCRIPTION: The adapter has failed to restart.

DATA: None

ACTION: If the issue persists, report the error to Broadcom technical support.

4.1.14.125 elx_mes2752: KILL_BOARD command failed retval <retval>

DESCRIPTION: The KILL BOARD mailbox command failed to complete.

DATA: None

ACTION: If the issue persists, report the error to Broadcom technical support.

4.1.14.126 elx_mes2753: PLOGI failure DID:<DID> Status:<Status>/<Extended Status>.

DESCRIPTION: A port login to <DID> was failed either by the driver, the firmware, or the target. The <status> and <extended status> indicate why the port login failed.

DATA: None

ACTION: If the issue persists, report the error to Broadcom technical support.

4.1.14.127 elx_mes2754: PRLI failure DID:<DID> Status:<Status>/<Extended Status>.

DESCRIPTION: A process login to <DID> failed either by the driver, the firmware, or the target. The <status> and <extended status> indicate why the process login failed.

DATA: None

ACTION: If the issue persists, report the error to Broadcom technical support.

4.1.14.128 elx_mes2755: ADISC failure DID:<DID> Status:<Status>/<Extended Status>.

DESCRIPTION: An address discovery to <DID> failed either by the driver, the firmware, or the target. The <status> and <extended status> indicate why the address discovery failed.

DATA: None

ACTION: If the issue persists, report the error to Broadcom technical support.

4.1.14.129 elx_mes2756: LOGO failure DID:<DID> Status:<Status>/<Extended Status>.

DESCRIPTION: An N_Port logout to <DID> failed either by the driver, the firmware, or the target. The <status> and <extended status> indicate why the N_Port logout failed.

DATA: None

ACTION: If the issue persists, report the error to Broadcom technical support.

4.1.14.130 elx_mes2757: Protocol failure detected during processing of FCP I/O op: tgt <tgt ID> lun <LUN> cmnd <CMD> rspInfo3 <rspInfo3>

DESCRIPTION: The FCP response from a target indicated that the response length is valid, but rspInfo3 indicates that there is no Failure. This is an FCP specification violation by the target.

DATA: None

ACTION: If the issue persists, report the error to Broadcom technical support.

4.1.14.131 elx_mes2758: Failed to allocate mbox for READ_FCF command.

DESCRIPTION: The driver failed to allocate memory from the mempool for issuing an FCF read mailbox command during the round robin FCF bmask update.

DATA: None

ACTION: None required.

4.1.14.132 elx_mes2759: Failed to allocate memory for round robin FCF failover bmask.

DESCRIPTION: The driver failed to allocate memory for the round robin FCF failover bmask.

DATA: None

ACTION: Make sure the system has enough kernel memory, you may need to reload the driver after the memory issue is resolved.

4.1.14.133 elx_mes2762: FCF <value> reached driver's book keeping dimension: <value>

DESCRIPTION:

DATA: None

ACTION: Inform Broadcom about this.

4.1.14.134 elx_mes2763: Failed to allocate mbox for READ_FCF cmd.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.14.135 elx_mes2765 Mailbox command READ_FCF_RECORD failed to retrieve a FCF record

DESCRIPTION: The driver failed to find an FCF record when the FCF table scan completed.

DATA:

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.14.136 elx_mes2772: Issue FCF rediscover mailbox command failed, failt through to FCF dead event

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.14.137 elx_mes2774: Issue FCF rediscover mailbox command failed, through to CVL event

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15 New Events

4.1.15.1 elx_mes2796: Mailbox memory allocation failed

DESCRIPTION: The driver failed to get memory resources to release an RPI.

DATA: None.

ACTION: None. The driver's heap is exhausted. A server reboot is required to fix the exhaustion. Contact Broadcom technical support if the issue persists.

4.1.15.2 elx_mes2798: Unreg_vpi failed vpi <value>, mb status = <value>

DESCRIPTION: The driver attempted to unregister a virtual port index and failed. The failure status is shown.

DATA: None.

ACTION: This condition is not catastrophic, but is unexpected. If the issue persists, contact Broadcom technical support.

4.1.15.3 elx_mes2813: Mgmt IO is Blocked <value> - mbox cmd <value> still active.

DESCRIPTION: The HBA's management interface is marked as blocked in preparation for an online or offline state transition. All user space access to the HBA via libdfc interface will be blocked.

DATA: None.

ACTION: None. Notification of a run-state change only.

4.1.15.4 elx_mes2822: IOCB failed <value> iotag <value> xri <value>

DESCRIPTION: The driver attempted to drain an internal queue and failed. The failure reason and some state variables are written to the console.

DATA: None

ACTION: None required. This should be a transient condition. If not, contact Broadcom technical support.

4.1.15.5 elx_mes2823: txq empty and txq_cnt is <value>

DESCRIPTION: The driver has detected a discrepancy between the elements queued to the txq and the counter tracking the number or items.

DATA: None

ACTION: None required. The driver has nothing to do except correct the counter - the txq is empty.

4.1.15.6 elx_mes2824: Cannot re-enable interrupt after slot reset.

DESCRIPTION: The driver failed to re-enable interrupts following a PCI slot reset command.

DATA: None

ACTION: A system reboot may be required to fully recover. Contact Broadcom technical support if the issue persists.

4.1.15.7 elx_mes2825: Unknown PCI error state: <value>

DESCRIPTION: The driver writes this message to the console if the PCI subsystem has detected an error on a Broadcom port and called the driver. The driver reacts by resetting the port.

DATA: None

ACTION: None required. The driver resets the device in an attempt to recover. Contact Broadcom technical support if the issue persists.

4.1.15.8 elx_mes2826: PCI channel disable preparing for reset

DESCRIPTION: The driver writes this message to the console if it is preparing the port for a reset operation.

DATA: None

ACTION: None required. This message is notification of a corrective measure. Contact Broadcom technical support if the issue persists.

4.1.15.9 elx_mes2827: PCI channel permanent disable for failure

DESCRIPTION: The driver writes this message to the console if a recovery mechanism has failed and the driver wants to mark the port with a permanent failure.

DATA: None

ACTION: A system reboot may correct the failure. If not, contact Broadcom technical support.

4.1.15.10 elx_mes2828: PCI channel I/O abort preparing for recovery

DESCRIPTION: The driver writes this message to the console when it is preparing the port for a recovery operation.

DATA: None

ACTION: None required. This is a notification message for the recovery action.

4.1.15.11 elx_mes2831: FLOGI response with cleared Fabric bit fcf_index <value> Switch Name <value> Fabric Name <value>

DESCRIPTION: When the driver completed a FLOGI, the common service parameters did not indicate an FPort or NPort remote node. The driver treats this as an error.

DATA: None

ACTION: Validate the external cable connection and FPort/Nport configuration. Contact Broadcom technical support if the issue persists.

4.1.15.12 elx_mes2856: Config Port Security Crypto Error: <value>, pmb->u.mb.un.varCfgPort.sec_err

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.13 elx_mes2858: FLOGI failure Status:<value>/<value> TMO:<value>

DESCRIPTION: The driver issued a FLOGI, but never received any completion within the timeout period. The driver is marking the FLOGI as failed and stops discovery.

DATA: None

ACTION: Check your Fabric to verify that it is operating correctly. Contact Broadcom technical support if the issue persists.

4.1.15.14 elx_mes2860: SLI authentication is required for INIT_LINK but has not been done yet

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.15 elx_mes2862: FCF (<value>) matches property of in-use FCF (<value>)

DESCRIPTION: The driver has found an FCF record that matches the properties of the current FCF record, except for the VLAN ID and Index. The driver will attempt to use this FCF.

DATA: None

ACTION: None required. The driver is in its FCF discovery phase and is trying to recover a match to its in-use FCF.

4.1.15.16 elx_mes2863: New FCF (<value>) matches property of in-use FCF (<value>)

DESCRIPTION: The driver has found a new FCF record that matches the properties of the current FCF record, but the record instance numbers do not match.

DATA: None

ACTION: None required. The driver is in its FCF discovery phase and is trying to recover a match to its in-use FCF.

4.1.15.17 elx_mes2877: FCP XRI exchange busy wait time: <value> seconds

DESCRIPTION: An FCP exchange cannot be released - no port completion. The driver is waiting.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.15.18 elx_mes2878: ELS XRI exchange busy wait time: <value> seconds

DESCRIPTION: An extended link service exchange cannot be released - no port completion. The driver is waiting.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.15.19 elx_mes2881: RRQ failure DID:<value> Status:<value>/<value>

DESCRIPTION: A Driver Reinstate Recovery Qualifier request failed - driver write target DID and status values.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.15.20 elx_mes2882: RRQ completes to NPort <value> with no ndlp. Data: <value> <value> <value>

DESCRIPTION: The driver completes a Reinstate Recovery Qualifier, but there is no node association.

DATA: (1) Status (2) Reason (3) IoTag

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.15.21 elx_mes2884: Vport array allocation failed

DESCRIPTION: The driver could not create a buffer list of VPorts.

DATA: None

ACTION: This is a software driver error. If this issue persists, report the error to Broadcom technical support.

4.1.15.22 elx_mes2885: Port Status Event: port status reg <value>, port smphr reg <value>, error 1=<value>, error 2=<value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.23 elx_mes2886: HBA Error Attention on unsupported if type <value>.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.24 elx_mes2887: Reset Needed: Attempting Port Recovery

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.25 elx_mes2888: Unrecoverable port error following POST: port status reg <value>, port smphr reg <value>, error 1=<value>, error 2=<value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.26 elx_mes2889: Port overtemperature event, taking port

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.27 elx_mes2890: Port error detected during port reset: wait_tmo: <value>, port status reg <value>, error 1=<value>, error 2=<value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.28 elx_mes2891: Init VFI mailbox failed <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.29 elx_mes2892: Failed to allocate init_vfi mailbox

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.30 elx_mes2893: Failed to issue init_vfi mailbox

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.31 elx_mes2894: SLI_INTF reg contents invalid sli_intf reg <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.32 elx_mes2895: Non FC link Event detected.(<value>)

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.33 elx_mes2897: The mboxq allocation failed

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.34 elx_mes2898: The lpfc_dmabuf allocation failed

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.35 elx_mes2899: The mbuf allocation failed

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.36 elx_mes2919: Failed to release resource extents for type <value> - Status <value> Add'l Status <value>. Resource memory not released.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.37 elx_mes2920: Failed to alloc Resource IDs rc = <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.38 elx_mes2925: Failed to issue CT ABTS RSP <value> on xri <value>, Data <value>

DESCRIPTION: The driver tried and failed to issue a response to an unsolicited abort sequence from the SAN.

DATA: None

ACTION: None required.

4.1.15.39 elx_mes2929: Resource Extent Opcode <value> is unsupported

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.40 elx_mes2930: Failed to get resource extents Status <value> Add'l Status <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.41 elx_mes2936: Could not find Vport mapped to vpi <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.42 elx_mes2945: SLI_CONFIG(mse) rd, ext_buf_cnt(<value>) out of range(<value>)

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.43 elx_mes2946: SLI_CONFIG(hbd) rd, ext_buf_cnt(<value>) out of range(<value>)

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.44 elx_mes2950: Failed SLI_CONFIG(hbd) rd (<value>)

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.45 elx_mes2953: SLI_CONFIG(mse) wr, ext_buf_cnt(<value>) out of range(<value>)

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.46 elx_mes2954: SLI_CONFIG(hbd) wr to interface type:<value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.47 elx_mes2999: Unsupported SLI4 Parameters Extents and RPI headers enabled.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.48 elx_mes3008: No available Resource Extents for resource type <value>: Count: <value>, Size <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.49 elx_mes3010: Run link diag test mailbox failed with mbx_status <value> status <value>, add_status <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.50 elx_mes3029: SLI_CONFIG(hbd) rd to interface type:<value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.51 elx_mes3042: Failed SLI_CONFIG(hbd) wr (<value>)

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.52 elx_mes3045: SLI_CONFIG(hbd) wr, ext_buf_cnt(<value>) out of range(<value>)

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.53 elx_mes3061: Last IDX <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.54 elx_mes3069: Clearing FCP rules

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.55 elx_mes3070: lpc_clr_ndlps_pri: fcp_priority sz = <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.56 elx_mes3071: no memory for priority rules

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.57 elx_mes3084: Allocated DMA memory size (<value>) is less than the requested DMA memory size (<value>)

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.58 elx_mes3089: Failed to allocate queues

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.59 elx_mes3103: Adapter Link is disabled.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.60 elx_mes3104: Adapter failed to issue

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.61 elx_mes3105: failed to allocate mailbox memory

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.62 elx_mes3112: lpfc_link_speed attribute cannot be set to <value>. Speed is not supported in loop mode.

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.63 elx_mes3113: Loop mode not supported at speed <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.64 elx_mes3125: Not receiving unsolicited event

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.65 elx_mes3141: Loopback mode: <value> not supported

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.66 elx_mes3142: Failed loopback test issue iocb: <value> iocb_stat: <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.67 elx_mes3143: Port Down: Firmware Restarted

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.68 elx_mes3144: Port Down: Debug Dump

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.69 elx_mes3145: Port Down: Provisioning

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.70 elx_mes3147: Fast-path EQs not allocated

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.71 elx_mes3148: Fast-path FCP CQ array not allocated

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.72 elx_mes3149: Fast-path FCP WQ array not allocated

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.73 elx_mes3150: No privilege to perform the requested access: <value>

DESCRIPTION: The driver lacks the necessary privileges for the requested action.

DATA: Requested action

ACTION: Contact Broadcom technical support.

4.1.15.74 elx_mes3151: PCI bus read access failure: <value>

DESCRIPTION: A read to the adapter's PCI registers failed.

DATA: Contents read from requested register.

ACTION: Contact Broadcom technical support.

4.1.15.75 elx_mes3152: Unrecoverable error, bring the port offline

DESCRIPTION: Unable to recover the adapter port post hardware error. Taking the port offline.

DATA: None

ACTION: Contact Broadcom technical support.

4.1.15.76 elx_mes3153: Fail to perform the requested access: <value>

DESCRIPTION: The driver failed to execute the requested management action.

DATA: Register action.

ACTION: Contact Broadcom technical support.

4.1.15.77 elx_mes3154: BLS ABORT RSP failed, data <value/value>

DESCRIPTION: The driver issued BLS ABORT Response failed to complete.

DATA: I/O status and I/O reason

ACTION: Contact Broadcom technical support.

4.1.15.78 elx_mes3161: Failure to post els sgl to port.

DESCRIPTION: The driver failed to port the scatter gather list to the adapter.

DATA: None

ACTION: Contact Broadcom technical support.

4.1.15.79 elx_mes3172: SCSI layer issued Host Reset Data: <value>

DESCRIPTION: The SCSI layer issued a host reset request to the driver.

DATA: Reset result

ACTION: Contact Broadcom technical support.

4.1.15.80 elx_mes3175: Failed to enable interrupt

DESCRIPTION: The driver failed to get interrupts re-enabled after an adapter reset.

DATA: None

ACTION: Contact Broadcom technical support.

4.1.15.81 elx_mes3176: Misconfigured Physical Port - Port Name <value>

DESCRIPTION: The driver has detected an unknown firmware name.

DATA: Detected firmware name

ACTION: Contact Broadcom technical support.

4.1.15.82 elx_mes3177: Nport <value>, DID <value> in wrong state for discovery to start. Failing request Data <value>

DESCRIPTION: The driver is trying to restart SAN discovery and detected an NPort is in the wrong initial state.

DATA: NPort ID, FCID, NPort state

ACTION: Reset the link. If this fails, contact Broadcom technical support.

4.1.15.83 elx_mes3249: Unable to allocate memory for QUERY_FW_CFG mailbox command

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.84 elx_mes3250: QUERY_FW_CFG mailbox failed with status <value> add_status <value>, mbx status <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.85 elx_mes3252: WQ doorbell offset not supported

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.86 elx_mes3262: RQ doorbell format not supported

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.87 elx_mes3263: WQ failed to memmap pci barset: <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.88 elx_mes3265: WQ doorbell format not supported: <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.89 elx_mes3269: RQ failed to memmap pci barset: <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.90 elx_mes3270: RQ doorbell format not supported: <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.91 elx_mes3279: Invalid provisioning of rpi

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.92 elx_mes3280: Invalid provisioning of vpi

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.93 elx_mes3281: Invalid provisioning of xri

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.94 elx_mes3282: Invalid provisioning of vfi

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.95 elx_mes3300: In-use FCF modified, perform FCF rediscovery

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.96 elx_mes3303: Failed to obtain vport vpi

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.97 elx_mes3317: HBA not functional: IP Reset Failed after <value> retries, try: echo fw_reset > board_mode

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.98 elx_mes3321: Recovered mailbox <value>

DESCRIPTION:

DATA: None

ACTION: None required.

4.1.15.99 elx_mes9000: Allocated DMA memory size (<value>) is less than the requested DMA memory size (<value>)

DESCRIPTION:

DATA: None

ACTION: None required.

4.2 Troubleshooting the NIC Driver

This section provides troubleshooting information for the NIC driver.

Table 15 identifies some of the common situations and their potential resolutions..

Table 15 Troubleshooting the NIC Driver

Issue	Resolution
Unable to ping from one VM to another VM.	The NIC driver creates two vmnic interfaces, one for each port. If these interfaces are configured as uplinks in two separate vSwitches, the VMs in each of these switches are in separate networks with no network path between them. Thus, pinging between the VMs in the two groups fails. If you want all these VMs in the same network, configure them as teaming uplinks to one vSwitch option. Each of the vmnics, vmnic1 to vmnic16, must be configured in a separate vSwitch. In this configuration, there is no network path between the vSwitches and pinging between these VMs does not work.
When inserting or removing a 1Gb SFP-RJ45 module on an adapter without RJ45 copper cables attached, the operating system indicates link up or down status. When inserting or removing copper cables attached to a switch to the module, link up or down events are not reported to the operating system.	No solution to this issue exists as the link status is not reported due to a PHY limitation on the card.
Flow control setting is not stored per port after rebooting the system.	With flow control, there is no persistence across reboot. TX and TX are always on when the system starts. For persistence, run a config command from an RC file at reboot.

4.2.1 ESXi Server NIC Event Log Entries

Table 16 below provides a list of ESXi Server network event log error messages. It includes the severity of the error, the message displayed, and the message description. When reporting a problem with the adapter to Broadcom, check the message log (/proc/vmware/log) and report any of these entries that may be present.

NOTE In Table 16, <D>, <DD>, or <DDD> in the 'Message Displayed' column refers to decimal values that appear in the error messages.

Table 16 ESXi Server NIC Event Log Entries

Severity	Displayed Message	Description
Error	OneConnect POST failed	The power-on self-test of the adapter failed. This indicates either a hardware or a firmware problem. Try rebooting the system after a reset.
Error	OneConnect initialization failed	Either the initialization of the adapter or the allocation of some resource for initializing the driver failed. In most cases, this message is accompanied by another more specific error message. Try rebooting the system after a power cycling. If the problem persists, this could indicate a hardware problem or corrupted firmware.
Error	RSS cannot be supported on this interface as SRIOV is enabled in the BIOS	RSS cannot be enabled when SR-IOV is enabled in the system BIOS.
Error	RSS cannot be supported on this interface as VFs are created	On LPe16202/OCe15100 and LPe16000-series adapters, RSS is not supported when VFs are already created on the interface.
Error	RSS cannot be supported when msix is disabled	MSI-X must be enabled to enable RSS.
Error	RSS is not supported on this interface	RSS is not supported on this interface.
Error	RSS cannot be supported on 1G interface	RSS cannot be enabled on 1GbE interfaces.
Error	RSS cannot be supported as the host does not have minimum required CPUs.	The host currently does not have enough CPUs to enable RSS. The host must have at least four cores and 4 to 8GB of RAM.
Warning	Using INTx interrupts. NetQueues feature will be disabled	The driver could not allocate the MSI-X vector for interrupts. The driver might continue to work, but the performance may be impacted.
Warning	WARNING: Found a OneConnect card in Gen <D> x<D> PCI-e slot. Should be in Gen 2, x8 slot for best performance	The adapter is an x8, Gen2 PCIe® device. For best performance, the adapter should be installed in a Gen2 PCIe slot 8 or 16 channels wide. The driver displays this warning if it finds the device in a slower or narrower PCIe slot. The device continues to work with reduced performance.
Warning	Command to get pause frame settings failed	The firmware command to get the PAUSE settings failed.
Warning	Command to set pause frame settings failed	The firmware command to change the PAUSE settings failed.
Warning	Command to apply MAC address filter failed	The driver could not set the MAC address filter on the hardware. The device continues to work. A performance impact may occur.
Warning	Command to delete MAC address filter failed	The firmware command to delete a MAC address filter failed. The device should continue to work.
Warning	Unable to get Firmware Version	The command to get the firmware revision number failed. The version number is not shown. The device will continue to work.

Table 16 ESXi Server NIC Event Log Entries (Continued)

Severity	Displayed Message	Description
Warning	Did not receive completions for all TX requests	While the driver was unloading, some outstanding transmit requests were found. This is an indication that the hardware is not functioning properly.
Warning	Failed to register char device	Could not create the char device used for certain management functions. The driver will continue to work. You may not be able to use HBACMD to interact with the device.
Warning	Invalid MTU requested. MTU must be between 64 and 9000 bytes.	There was an invalid MTU size in the MTU configuration IOCTL. The MTU will not be changed.
Warning	Invalid vlan priority labeled. Must be 0 - 7	A request to set a VLAN priority tag was made with an invalid value.
Warning	Failed to allocate memory for pass through command	The memory allocation for a pass through command failed. The driver will continue to function. The configuration utility that issued the pass through IOCTL will fail.
Warning	Pass through command failed. opcode <DDD>, status 0x<XXX>	The pass through firmware command with the indicated opcode failed. The driver should continue to function. The configuration utility that issued the pass through IOCTL will fail.
Warning	Command to modify EQ delay failed	The firmware command to change the event queue delay failed. The driver will continue to function. Adaptive interrupt coalescing does not function correctly.
Warning	Unqualified SFP+ module detected on Port <adapter port number> from <vendor name> part number: <vendor part number>	The specified port has an unqualified SFP+ module inserted into it.

4.2.2 NIC Adapter Firmware Error

The following POST message appears if you have loaded firmware on the adapter that the controller does not support:

```
POST Error : Firmware halted. This firmware does not support this controller.
```

4.2.3 NIC Informational Log Groups

The elxnet driver informational logs are categorized into groups. You can enable or suppress the logs of a specific group by setting or clearing the corresponding bit on DebugMask. The DebugMask is a bit-vector (uint32) and each bit represents a group.

Table 17 lists the elxnet driver informational log categories.

Table 17 NIC Informational Log Groups

Hexadecimal Value	Description
0x1	DRIVER
0x2	UPLINK
0x4	QUEUE
0x8	INTR
0x10	MCC

Table 17 NIC Informational Log Groups (Continued)

Hexadecimal Value	Description
0x20	TX
0x40	RX
0x80	MGMT
0x100	WORKER
0x200	SRIOV
0x400	EVENT
0x800	VLAN
0x1000	VXLAN

The default enabled groups include DRIVER, UPLINK, MCC, QUEUE, SRIOV, and VXLAN.

Log groups reduce informational log clutter at the default log level and aid in debugging by enabling or disabling group logs. Warning and error logs are always logged, regardless of the debug mask. See Table 16, ESXi Server NIC Event Log Entries, for more information on the warning and error logs.

The debug masks control which groups are enabled or disabled, and they are not meant to be a replacement of the log level adjustments supported by vsish. The vmkernel log levels control the verbosity of logs at each group. For example, level 0 results in less logging and level 4 results in more verbose logs.

There are two ways to modify the debug mask, with the ModuleParam debugMask command or the esxcli-Plugin dbgmask command.

4.2.3.1 ModuleParam debugMask

NOTE You must reload (or reboot) the driver after modifying the debugMask module parameter.

This is a global debugMask command; therefore, the same value applies to all Emulex NIC PFs.

Use either of these commands:

```
esxcfg-module -s "debugMask=0x0120" elxnet
```

- Or -

```
esxcli system module parameters set -p "debugMask=0x0120" -m elxnet
```

4.2.3.2 esxcli-Plugin

NOTE A driver reload (or reboot) is not required after modifying the dbgmask parameter.

This is a per-PF debugMask command; therefore, the value applies to a specific Emulex NIC PF (PF with SBDF is specified with "-p").

Command to Get:

```
esxcli elxnet dbgmask get -p 0000:03:00.1
```

Command to Set:

```
esxcli elxnet dbgmask set -p 0000:03:00.1 0x000000ff
```

4.3 Troubleshooting the iSCSI Driver

This section provides troubleshooting information for the iSCSI driver.

Table 18 identifies some of the common situations and their potential resolutions. r.

Table 18 Troubleshooting the iSCSI Driver

Issue	Resolution
Mutual CHAP authentication cannot be configured with the vSphere client.	One-way CHAP is the only authentication method that can be configured with the vSphere Client. To set Mutual CHAP authentication, you must use iSCSISelect.
IPV6 support is not available in vSphere client for hardware iSCSI adapters.	To set the IPv6 address, you must use iSCSISelect.

4.3.1 Retrieving iSCSI Error Log Codes

The iSCSI driver generates error codes to the /var/log/vmkernel log file. The vmkernel log file is an ASCII text file that can be viewed and searched with a text editor such as vim. The vmkernel log file is automatically rotated as it gets larger, and the rotated log files are named vmkernel.x, where x is an integer.

To search the log file for error messages, at the command prompt type:

```
#cd /var/log
#vim vmkernel.log
```

For example, you may see the following message:

```
Sep 9 19:48:04 esx-server vmkernel: OneConnect iSCSI Driver: The be2iscsi driver
received a Task Management Function that is not supported and rejected this
request. The error log entry immediately following this entry will indicate the
TMF function code that was rejected.
```

4.3.2 iSCSI Error Log Code Entries

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

Table 19 iSCSI Error Log Code Entries

Severity	Message	Recommended Resolution
Error	The be2iscsi driver failed to load because initialization failed during a power management boot-up.	This failure may be due to the firmware not being present or not running currently. This failure may also indicate a hardware problem.
Error	The be2iscsi driver was unable to map one or more PCI Base Address Register and hence failed to load.	This failure may indicate a low memory condition or a hardware error.
Error	The be2iscsi driver ignored a configuration entry since the entry was invalid.	Check the registry configuration for any recent entries added for driver parameters. The invalid entry must be removed or corrected.
Error	The be2iscsi driver failed to load due to memory allocation failure.	This failure occurred due to a failed memory allocation in the driver. Check low memory conditions.
Error	The be2iscsi driver failed to load because initialization failed during normal boot-up.	This failure may be due to the firmware not being present or not running currently. This failure may also indicate a hardware problem.

Table 19 iSCSI Error Log Code Entries (Continued)

Severity	Message	Recommended Resolution
Error	An internal API failed in be2iscsi driver during initialization.	This failure may indicate a low memory condition.
Error	An Unrecoverable Error was detected by the be2iscsi driver. Following this error log entry, the next 3 entries will indicate the error codes.	This may be due to hardware errors or due to unhandled exceptions in the hardware or firmware.
Error	The be2iscsi driver failed an IOCTL request since the number of scatter gather elements required for the IOCTL buffer exceeded the adapter's firmware limit. Following this error log entry, the next entry will indicate the IOCTL opcode and the payload length requested.	This error may indicate an incorrect configuration option for the iSCSI driver. It may also indicate a low memory condition.
Error	The be2iscsi driver detected an error during offloading the iSCSI connection. The operation will be retried again. Following this error log entry, the next entry will indicate the session handle and the firmware error code.	This may indicate a target is in error or may point to transient network connectivity issues. It may also indicate a firmware error.
Error	The be2iscsi driver did not receive an iSCSI command window update up to 25 seconds during I/O operations. Following this error log entry, the next entry will indicate the session handle where this error occurred. The be2iscsi driver will trigger a session recovery on the session and continue.	Check for any errors reported at the target. The iSCSI initiator is only supported with certified targets. Check for the latest software at the target vendor's website. Check for the latest software at www.broadcom.com. If the above fails, contact Broadcom technical support.
Warning	The be2iscsi driver received an invalid iSCSI Command Sequence Number update from the target. Following this error log entry, the next three entries will indicate the session handle and the iSCSI parameters - MaxCmdSN and ExpCmdSN respectively.	Check for any errors reported at the target. The iSCSI initiator is only supported with certified targets. Check for the latest software at the target vendor's website. Check for the latest software at www.broadcom.com If the above fails, contact Broadcom technical support.
Warning	A connection to the target was lost for a period exceeding the Extended Timeout (ETO). The error log entry immediately following this entry will indicate the session ID of the target that lost the connection. Event log entries from the disk subsystem indicating that the drives were lost will exist. If any I/Os were in progress, the system may see I/O errors or failures.	Check the connection to the target or the state of the target device. If the target is made available, any sessions that existed previously will be reestablished and the devices will be available for I/O.
Warning	The be2iscsi driver received a Task Management Function that is not supported and rejected this request. The error log entry immediately following this entry will indicate the task management function code that was rejected.	The operating system version is not supported.
Warning	The be2iscsi driver received a Task Management Function Abort request for an I/O request that is not present with the driver.	This message may indicate a slow connection to the target. Check network connectivity to the target for any errors.
Error	The be2iscsi driver encountered a mismatched version of the firmware running on the board. This error may be followed by more error codes 0x31840001 or 0x31880001 indicating that the be2iscsi driver failed to load.	This failure indicates that the driver version that is running on the system does not match the version of the firmware flashed on the board. Fix this by installing the desired version.

Table 19 iSCSI Error Log Code Entries (Continued)

Severity	Message	Recommended Resolution
Error	The be2iscsi driver detected a failure in the hardware during initialization. This error may be followed by more error codes 0x31840001 or 0x31880001 indicating that the be2iscsi driver failed to load.	This failure indicates that the hardware has not been initialized or is malfunctioning. This may also indicate that the firmware is not running correctly.
Warning	Both Port 0 and Port 1 links were down for a period exceeding the Link Down Timeout (LDTO). If the initiator has connection to the target, there will be event log entries from the disk subsystem indicating that the drives were lost. If any I/Os were in progress, the system may see I/O errors or failures.	Check the links to the adapter. If the link is reestablished, any sessions that existed previously will be reestablished and the devices will be available for I/O.
Warning	Both Port 0 and Port 1 links are down.	Check the links to the adapter.

4.3.3 Additional iSCSI Driver Messages

The following iSCSI error messages are returned if you specify illegal options when loading the driver:

- be2iscsi <*PCI_DEVICE_LOCATION*>: WARNING: dic value = %d out of range. Valid range is 0 - 1. Using default value = 1
- be2iscsi <*PCI_DEVICE_LOCATION*>: WARNING: eto value = %d out of range. Valid range is 0 - 3600. Using default value = 30
- be2iscsi <*PCI_DEVICE_LOCATION*>: WARNING: ldto value = %d out of range. Valid range is 0 - 3600. Using default value = 20
- be2iscsi <*PCI_DEVICE_LOCATION*>: WARNING: ios_per_ctrl value = %d out of range. Valid range is 1 - 512. Using default value = 512
- be2iscsi <*PCI_DEVICE_LOCATION*>: WARNING: max_io_size value = %d out of range. Valid range is 4 - 512 (KByte). Using default value = 128
- be2iscsi <*PCI_DEVICE_LOCATION*>: WARNING: tmf_reset value= %d out of range. Valid range is 1 - 3. Using default value= 1

NOTE For the <PCI_DEVICE_LOCATION>, the driver returns a value in the following format – 0000:00:00.0 (for example, 0000:05:00.4).

%d represents a signed integer outside of the valid range. The driver returns a number in the Warning message.

Any other messages will be returned in the following form:

`be2iscsi <PCI_DEVICE_LOCATION>: FUNCTION_NAME:LINE:MESSAGE`

where the <PCI_DEVICE_LOCATION> is a value returned by the driver in the following format – 0000:00:00.0.

For example:

```
be2iscsi 0000:05:00.4: 1088: kmalloc failed.
```

You are asked to find the function <FUNCTION_NAME> at line <LINE> in the source.

4.3.3.1 Description of Mandatory and Optional Parameters

The following table describes the parameters used in the data string for option 43.

Table 20 Data String Parameters for Option 43

Parameter	Description	Field Type
<TargetIP>	Replace it with a valid IPv4 address in dotted decimal notation.	Mandatory
<TargetTCPPort>	Replace it with a decimal number ranging from 1 to 65535 (inclusive). The default TCP port 3260 is assumed, if not specified.	Optional
<LUN>	It is the hexadecimal representation of the Logical Unit Number of the boot device. If not provided, LUN 0 is assumed to be the boot LUN. It is an eight-byte number which should be specified as a hexadecimal number consisting of 16 digits, with an appropriate number of 0's padded to the left, if required.	Optional
<TargetName>	Replace it with a valid iSCSI target IQN name of up to 223 characters.	Mandatory
<InitiatorName>	Replace it with a valid iSCSI IQN name of up to 223 characters. If not provided the default Initiator name (generated by the adapter based on the board's MAC address) will be used.	Optional
<HeaderDigest>	Replace it with either E or D. ■ E denotes that the header digest is enabled. ■ D denotes that it is disabled.	Optional
<DataDigest>	Replace it with either E or D. ■ E denotes that the data digest is enabled. ■ D denotes that it is disabled. If not provided it is assumed that Data Digest is disabled by default.	Optional
<AuthenticationType>	If applicable replace it with D, E, or M. ■ D denotes that authentication is disabled. ■ E denotes that one-way CHAP is enabled - the username and secret to be used for oneway CHAP must be specified by non-DHCP means ■ M denotes that MutualCHAP is enabled - user name and passwords required for mutual CHAP authentication must be specified by non-DHCP means. If not specified, this field defaults to authentication disabled.	Optional

4.3.3.1.1 Examples

The following is an example of the Default Initiator name and Data Digest Settings:

```
iscsi:"192.168.0.2":"3261":"000000000000000E":"iqn.2009-4.com:1234567890"::"E"::"
E"
```

where the values are:

- Target IP address: 192.168.0.2
- Target TCP port: 3261
- Target boot LUN: 0x0E
- Target IQN name: iqn.2009-04.com:1234567890
- Initiator name: Not specified. Use the Initiator name already configured. Use the default name if none was configured.
- Header Digest: Enabled
- Data digest: Not specified. Assume disabled.
- Authentication Type: 1-way CHAP.

The following is an example of Default TCP Port and Mutual CHAP Settings:

```
iscsi:"192.168.0.2"::"000000000000000E":"iqn.2009-4.com:1234567890"::"E":"D":"M"
```

where the values are:

- Target IP address: 192.168.0.2
- Target TCP port: Use default from RFC 3720 (3260)
- Target boot LUN: 0x0E
- Target IQN name: iqn.2009-04.com:1234567890
- Initiator name: Not specified. Use the Initiator name already configured. Use the default name if none was configured.
- Header Digest: Enabled
- Data digest: Data Digest disabled
- Authentication Type: Mutual CHAP

4.4 Native Mode NIC Driver Troubleshooting Support

Table 21 lists the troubleshooting support provided by the earlier vmklinux driver in terms of the proc interface and the equivalent support provided by the native mode NIC driver through either the esxcli plug-in or the VmkMgmtKeyVal interface.

Table 21 Native Mode NIC Driver Troubleshooting Support

Troubleshooting Items from Earlier vmklinux Driver using the proc Interface		Native Mode Driver Support		Usage in ESXi 5.5 and ESXi 6.0
	Description	(VmkMgmtKeyVal interface)	(esxcli plug-in)	Native Mode Driver
csr_read	Read the 32-bit register value from the CSR space at the offset set through the csr_read file.	X	—	Not implemented.
csr_write	Set the CSR space offset for csr_write.	X	—	Not implemented.
drvr_stat	Read the driver statistics.	—	X	Available in custom esxcli plug-in in the asynchronous driver: `esxcli elxnet stats get -p <pci dev name>`
eth_ring	Get the various Ethernet ring properties.	—	—	Not implemented.
misc_stat	Read the miscellaneous counters, format them with description. These counters are clear on read.	—	Partially implemented	Error stats available in a custom esxcli plug-in in the asynchronous driver: `esxcli elxnet stats get -p <pci dev name>`
pci_read	Read the 32-bit register value from the PCI space at the offset set through the pci_read file.	X	—	Not implemented.

Table 21 Native Mode NIC Driver Troubleshooting Support (Continued)

Troubleshooting Items from Earlier vmklinux Driver using the proc Interface		Native Mode Driver Support		Usage in ESXi 5.5 and ESXi 6.0
	Description	(VmkMgmtKeyVal interface)	(esxcli plug-in)	Native Mode Driver
pci_write	Set the PCI space offset for pci_write.	X	—	Not implemented.
port_stat	Read the counters for port 0 and port 1 and formats them with descriptions. These counters are clear-on-read.	—	Partially implemented	From vsish node /net/pNics/vmnicXX/stats. Also available in a custom esxcli plug-in in the asynchronous driver: `esxcli elxnet stats get -p <pci dev name>`
vlan_stat	Get the driver-specific VLAN status.	—	Not implemented	Not relevant since the driver does not use VLAN filtering.

4.5 Native Mode NIC Driver Support for ethtool Commands

Table 22 lists the ethtool support provided by the earlier vmklinux driver and the equivalent support provided through the esxcli and vsish commands.

Table 22 ESXi 5.5 and 6.0 Native Mode NIC Driver Support for ethtool Commands

ethtool Command	Description	vmklinux Driver Support (VmkMgmtKeyVal interface)	ESXi 5.5 and 6.0 Native Mode Driver Support (esxcli/ vsish command)
ethtool -a|--show-pause DEVNAME	Show pause options.	X	✓ (`esxcli network nic get -n vmnic0`)

ethtool -A|--pause DEVNAME	Set pause options.	X	✓
			To enable RX/TX pause:
			`vsish -e set /net/pNics/vmnic1/firmware/ pauseParams < autonegsupport> <txpauseenabled> <rxpauseenabled>`
			Enable:
			`vsish -e set /net/pNics/vmnic1/firmware/ pauseParams 0 1 1`
			Disable:
			`vsish -e set /net/pNics/vmnic1/firmware/ pauseParams 0 0 0`

ethtool -c|--show-coalesce DEVNAME	Show coalesce options.	X	✓ (`esxcli network nic coalesce get -n vmnicX`)
ethtool -C|--coalesce DEVNAME	Set coalesce options.	X	✓ (`esxcli network nic coalesce set -n vmnicX`)
ethtool -g|--show-ring DEVNAME	Query RX/TX ring parameters.	X	esxcli does not support it

Table 22 ESXi 5.5 and 6.0 Native Mode NIC Driver Support for ethtool Commands (Continued)

ethtool Command	Description	vmklinux Driver Support (VmkMgmtKeyVal interface)	ESXi 5.5 and 6.0 Native Mode Driver Support (esxcli/ vsish command)
ethtool -G|--set-ring DEVNAME	Set RX/TX ring parameters.	Not implemented	esxcli does not support it
ethtool -k|--show-offload DEVNAME	Get protocol offload information.	X	X (`esxcli network nic sg/tso/cso get`)

ethtool -K|--offload DEVNAME	Set protocol offload.	Not implemented	Not implemented (`esxcli network nic sg/tso/cso set -n vmnicX`)
			Supported using:
			`vsish:vsish -e get /net/pNics/vmic<>/hwCapabilities`
			`vsish -e set /net/pNics/vmnic0/hwCapabilities/<CAP> <1/0>`

ethtool -i|--driver DEVNAME	Show driver information.	✓	✓ (`esxcli network nic info get`)
ethtool -d|--register-dump DEVNAME	Dump device registers.	✓	✓ (`esxcli elxnet regdump get -p <pcidevname> -f <filepath>`)
ethtool -e|--eeprom-dump DEVNAME	Dump device EEPROM.	✓	Not implemented (`esxcli network nic eeprom dump -n vmnicX`)
ethtool -E|--change-eeprom DEVNAME	Change bytes in device EEPROM.	Not implemented	Not implemented (`esxcli network nic eeprom change -n vmnicX`)
ethtool -r|--negotiate DEVNAME	Restart N-WAY negotiation.	Not implemented	Not implemented (`esxcli network nic negotiate restart -n vmnicX`)
ethtool -p|--identify DEVNAME	Show visible port identification (for example, blinking).	✓	esxcli does not support it
ethtool -t|--test DEVNAME	Execute adapter self test.	Not implemented	Not implemented (`esxcli network nic selftest run -n vmnicX`)
ethtool -S|--statistics DEVNAME	Show adapter statistics.	✓	✓ (`esxcli network nic stats get`) The standard esxcli statistics only expose a limited set of statistic counters. More extensive statistics are available through the esxcli plug-in available in the asynchronous driver: `esxcli elxnet stats get -p <pci dev name>`

Appendix A: Configuring iSCSI Through DHCP

This section provides information for configuring the iSCSI driver using DHCP.

A.1 Recommendations

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

A.1.1 Vendor-Specific Option 43

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

A.1.1.1 Format of Vendor-Specific Option 43

The data returned in DHCP vendor-specific option 43 uses the following format:

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

The guidelines for creating the data string include:

- Strings shown in quotation marks are part of the syntax and are therefore mandatory.
- Fields enclosed in angular brackets (including the angular brackets) must be replaced with their corresponding values. Some of these fields are optional and may be skipped.
- If an optional field is skipped, a colon must be used as a placeholder to indicate the default value for that field.
- If specified, the value of each parameter must be enclosed in double quotes.

All options are case sensitive.

For more information on the Option 43 values, see Section 4.3.3.1, Description of Mandatory and Optional Parameters.


BROADCOM

Emulex® CIM Provider Package for OneConnect® Adapters

Installation Guide

Version 11.2
December 30, 2016

CIMPROVOCA-IG112

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

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

Copyright © 2011–2016 by Broadcom. All Rights Reserved.

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

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.

Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Chapter 1: Introduction

The Emulex® Common Information Model (CIM) provider enables comprehensive management of Emulex converged network adapters (CNAs). It uses an industry standard Application Programming Interface (API), Common Manageability Programming Interface (CMPI) v2.0, to manage various Emulex adapters.

The CIM provider supports the basic CNA (Fibre Channel over Ethernet [FCoE], Internet Small Computer System Interface [iSCSI], and network interface card [NIC]) profile, and it supports Emulex OneConnect® OCe11000-series and OCe14000-series adapters.

During installation, the Broadcom Emulex CIM provider registers with a Web-based Enterprise Management (WBEM) server running a CIM Object Manager (CIMOM) service. The CIM provider performs the following functions:

- Internally communicates with the Emulex management API.
- Internally communicates with the Emulex drivers.
- Handles inquiries and requests from various CIM clients.

1.1 Supported Platforms

The following table shows the platforms supported with the Emulex CIM provider kits.

Table 1 Broadcom Emulex CIM Provider Supported Platforms

Operating Systems	Broadcom Emulex CIM Provider Kits
VMware ESXi 5.5	`vmware-esx-provider-brcm-cim-provider-5.5.0-<kit version>.zip`
VMware ESXi 6.0	`vmware-esx-provider-brcm-cim-provider-6.0.0-<kit version>.zip`
VMware ESXi 6.5	`vmware-esx-provider-brcm-cim-provider-6.5.0-<kit version>.zip`

NOTE The CIM Provider kits are available in both `.zip` and `.vib` formats.

1.2 Supported CIM Provider Profiles

- **Access Points** – Storage Networking Industry Association (SNIA) Storage Management Initiative (SMI)-S 1.5 Part 2 Clause 23
- **FCoE Initiator Ports** – SNIA SMI-S 1.5 Part 2 Clause 22
- **Host Discovered Resources** – SNIA SMI-S 1.5 Part 6 Clause 7
- **Indication Profile (partial)** – Distributed Management Task Force (DMTF) DSP1054
- **Peripheral Component Interconnect (PCI) Device** – DMTF DSP1075
- **Physical Assets Profile** – DMTF DSP1011
- **Physical Package** – SNIA SMI-S 1.5 Part 2 Clause 31
- **Profile Registration** – DMTF DSP1033
- **Record Log** – DMTF DSP1010
- **Software Inventory** – DMTF DSP1023
- **Software Update** – DMTF DSP1025
- **Storage HBA** – SNIA SMI-S 1.5 Part 6 Clause 6
- **Diagnostics (CDMv2)** – DMTF DSP1002
- **Ethernet Port** – DMTF DSP1014

- **Host local area network (LAN) Network Port** – DMTF DSP1035
- **Internet Protocol (IP) Interface** – DMTF DSP1036
- **iSCSI Initiator** – SNIA SMI-S 1.5 Part 6 Clause 9
- **iSCSI Initiator Ports Profile** – SNIA SMI-S 1.5 Part 2 Clause 16

NOTE The iSCSI profiles are only instantiated when one of the PCI function instances on the CNA is acting as an iSCSI initiator.

1.3 Abbreviations

API	Application Programming Interface
CIM	Common Information Model
CIMOM	CIM Object Manager
CMPI	Common Manageability Programming Interface
DMTF	Distributed Management Task Force
CNA	converged network adapter
FC	Fibre Channel
FCoE	Fibre Channel over Ethernet
HBA	host bus adapter
IP	Internet Protocol
iSCSI	internet Small Computer System Interface
LAN	local area network
NIC	network interface card (or controller)
PCI	Peripheral Component Interconnect
SFCB	Small Footprint CIM Broker
SMI	Storage Management Initiative
SNIA	Storage Networking Industry Association
URI	Uniform Resource Identifier
VIB	vSphere Installation Bundle
vNIC	virtual network interface card
WBEM	Web-based Enterprise Management

Chapter 2: Installing the Emulex CIM Provider

The following items must be installed before you can install the Emulex CIM provider.

- One of the following adapters:
  - OCe11000-series adapter
  - OCe14000-series adapter
- The appropriate adapter drivers

NOTE CNAs on an ESXi host running Emulex CIM providers can be managed by a Windows server using the following applications (installed on Windows operating systems):

- The Emulex OneCommand® CNA Manager application for Windows
- The Emulex OneCommand CNA Manager for VMware vCenter

Go to the download page on the Broadcom website, at http://www.broadcom.com, or to the vendor website to verify the driver version or the OneCommand CNA Manager application version that must be installed on your system.

2.1 Installing the VIB in the VMware Operating System

Use one of the standard `esxcli` commands to install the VIB or the offline bundle.

- To install the VIB, type:

```
esxcli software vib install -v=<provider.vib> --maintenance-mode
```

- To install the signed offline bundle, type:

```
esxcli software vib install -d=<offline-bundle.zip> --maintenance-mode
```

- To install an unzipped file, type:

```
esxcli software vib install --viburl=<file:/vmware-esx-provider-brcm.vib>
--maintenance-mode
```

2.2 Using Adapters

This section describes updating firmware, enabling logs, running diagnostics, and discovering vNICs on Emulex adapters.

If you require additional information, contact an authorized Broadcom technical support representative at ccx-tech.support@broadcom.com or request assistance online at https://oemsupportportal.emulex.com/web2tech/ccx.html.

2.2.1 Updating Firmware on Broadcom Emulex Adapters

NOTE If a secure version of firmware (version 11.0 or later) is installed on an OCe14000B-series adapter and you want to revert to an earlier unsecure version of firmware, you must remove the secure firmware jumper block before performing the update. Refer to the installation guide for the adapter for more information.

Use the Software Update profile to update the firmware on Emulex adapters. The following methods are implemented in the Emulex-specific Software Update profile:

- **InstallfromByteStream** – Requires a custom CIM client that can read the firmware file and create a ByteStream used to update the firmware.
- **InstallfromURI** – Supports two different types of URI:
  - The firmware file to be updated is available locally on the machine hosting the Emulex adapter.
  - The firmware file to be updated is available on a remote machine, such as an http or https server. In this case, the Emulex CNA CIM Provider uses the `libcurl` library available on the host machine, where the CIM provider is running, to download the firmware file.

The Emulex `CIM_SoftwareInstallationServiceCapabilities` class has the attribute `SupportedURISchemes`, which identifies the supported URI schemes.

To update the firmware, perform these steps:

1. List the `CIM_SoftwareInstallationServices` in the Emulex namespace. Select the `CIM_SoftwareInstallationService` specific to the adapter on which the firmware is to be updated. For example, if the CNA needs a firmware update, select the `ELXUCNA_SoftwareInstallationService` class instance:

```
wbemcli  -noverify ein
'https://root:<password>@<IP>/root/brcmcna:elxucna_softwareinstallationservice' -nl
```

2. List the `CIM_EthernetPorts` and select the desired port from the listed instances. For example:

```
wbemcli  -noverify ein 'https://root:<password>@<IP>/root/brcmcna:elxucna_ethernetport'
-nl
```

3. Run `InstallFromURI` using the output from step 1 and step 2.

   For example:

```
wbemcli cm -noverify 'https://root:<password>@<IP>/root/brcmcna:<Output of step 1>
InstallFromURI.URI=<Full path of fw file>,Target=<Output of step 2>
```

2.2.2 Enabling Logs and Collecting Symptoms

To enable provider logs, perform these steps:

1. List the `ELXUCNA_RecordLog` (for a CNA in FCoE+NIC mode) class and note the provider log instance. If the list command fails, perform the steps detailed in Section 2.2.3, Generating Provider Logs if Listing Them Fails.
2. Set the provider log settings:

```
wbemcli -noverify cm 'https://root:<password>@<IP>/root/brcmcna:<Instance from step 1> '
SetLogParams.LogLevel=5,logmode=2,tracepath='<file path>'
```

3. Set the provider log state:

```
wbemcli -noverify cm 'https://root:<password>@<IP>/root/brcmcna:<Instance from step 1> '
RequestStateChange.RequestedState=2
```

4. Perform the operation that is not working as expected. The provider logs are available in the *<file path>* specified in step 2.
5. Contact a Broadcom technical support representative with the provider logs and the system logs (`/var/log/syslog.log`).

To disable provider logs, enter the following command:

```
wbemcli -noverify cm 'https://root:<password>@<IP>/root/brcmcna:<Instance from step 1> '
RequestStateChange.RequestedState=3
```

2.2.3 Generating Provider Logs if Listing Them Fails

To enable provider logs if listing them fails, perform these steps:

1. Stop the CIMOM.
2. Create a `.dmp` file. To create a provider log, create a `brcm_ucna_provider.dmp` file in the `/etc/cim/brcmcna` location.
3. Manually enter the following two lines (without spaces):

   ```
   2,5,1
   /tmp/providerlogs.txt
   ```

4. Start the CIMOM and list the Emulex classes in the `root/brcmcna` namespace.
5. Perform the operation that is not working as expected. The provider logs are available in *</tmp/providerlogs.txt>* specified in step 3.
6. Contact a Broadcom technical support representative with the provider logs and the system logs (`/var/log/syslog.log`).

NOTE Dead dump is not supported.

2.2.4 Running Diagnostic Tests

To run a diagnostic test on the Emulex CIM provider, perform these steps:

1. Get the instance of the managed element (Ethernet port). For example:

   ```
   wbemcli  -noverify ein
   'https://root:<password>@<IP>/root/brcmcna:elxucna_ethernetport' -nl
   ```

2. Get the instance of the `ELXBRCMUCNA_DiagnosticTest` class. For example:

   ```
   wbemcli  -noverify ein
   'https://root:<password>@<IP>/root/brcmcna:elxucna_diagnostictest' -nl
   ```

3. Run the `RunDiagnosticservice` on the `brcmucna_diagnostictest`. For example:

   ```
   wbemcli -noverify cm 'https://root:<password>@<IP>/root/brcmcna:<Diagnostic
   Test instance from step 2>'
   RunDiagnosticservice.ManagedElement=<ManagedElement instance from step 1>'
   ```

   A `CIM_ConcreteJob` instance is created for each diagnostic test run. For example:

   ```
   wbemcli  -noverify ein
   'https://root:<password>@<IP>/root/brcmccx:brcmucna_concretejob' -nl
   ```

   Results of the diagnostic test runs are available in `BRCMUCNA_DiagnosticCompletionRecord` class instances. For example:

```
wbemcli  -noverify ein
'https://root:<password>@<IP>/root/brcmcna:elxucna_diagnosticcompletionrecord' -nl
```

The diagnostic logs can be cleared using the `ClearLog` function of `BRCMUCNA_DiagnosticsLog` class. For example:

```
wbemcli -noverify cm 'https://root:<password>@<IP>/root/brcmcna:<Corresponding
Diagnostic log instance>' ClearLog
```

2.2.5 Discovering vNICs

You can discover vNICs in three ways:

- Determine vNICs using the CNA CIM provider

- Determine vNICs hosted by the CNA in FCoE+NIC mode in a specific host
- Determine vNICs operating over the same physical port

2.2.5.1 Determining vNICs Using the CNA CIM Provider

1. Find all the `CIM_VLANEndpoint` instances whose `SystemName` property has a value equal to the name of the desired host. These instances are vNIC endpoints.
2. Find the `CIM_EndpointIdentity` instance associating the `CIM_VLANEndpoint` instance to a `CIM_LANEndpoint` instance. Follow the association to the `CIM_LANEndpoint`.
3. Find the `CIM_DeviceSAPImplementation` instance that associates the `CIM_LANEndpoint` to an instance of `CIM_EthernetPort`. Follow the associations to the `CIM_EthernetPort` instance representing the vNIC.

2.2.5.2 Determining vNICs Hosted by the CNA in FCoE+NIC Mode in a Specific Host

1. Find the `CIM_Card` instance for the CNA in FCoE+NIC mode.
2. Find the vNICs in the host as detailed earlier, based on the version of the CIM provider.
3. For each `CIM_EthernetPort` representing a vNIC, perform these steps:
   a. Find the `CIM_ControlledBy` instance associating the `CIM_EthernetPort` to a `CIM_PortController` instance. Follow the association to the `CIM_PortController` instance.
   b. Find the `CIM_Realizes` instance associating the `CIM_PortController` to a `CIM_PhysicalConnector` instance. Follow the association to the `CIM_PhysicalConnector` instance.
   c. Find the `CIM_Container` instance associating the `CIM_PhysicalConnector` to a `CIM_Card` instance.

      Each `CIM_EthernetPort` with a `CIM_Container` instance referencing the `CIM_Card` instance found in step 1 is hosted by the targeted CNA.

2.2.5.3 Determining vNICs Operating Over the Same Physical Port

1. Find the `CIM_PhysicalConnector` instance representing the port in question. The key value for `CIM_PhysicalConnector` instances contains the serial number of the CNA hosting the port and the port number assigned to the physical port.
2. Find the vNICs in the host as detailed earlier, based on the version of the CIM provider. For each `CIM_EthernetPort` representing a vNIC:
   a. Find the `CIM_ControlledBy` instance associating the `CIM_EthernetPort` to an instance of `CIM_PortController`. Follow the association to the `CIM_PortController` instance.
   b. Find the `CIM_Realizes` instance associating the `CIM_PortController` to an instance of `CIM_PhysicalConnector`.

      Each `CIM_EthernetPort` with a `CIM_Realizes` instance referencing the `CIM_PhysicalConnector` instance is running over the targeted physical port.

Chapter 3: Troubleshooting

The following error message may appear if the CIM hosts are not properly added to the OneCommand CNA Manager application:

```
"Unknown or invalid host specified"
```

There could be instances in which the drivers, the CIM provider, and the CIM Client on a Windows machine are all properly installed, but the CIM hosts are still not added to the OneCommand CNA Manager application. The following table shows the most common reasons for this problem.

Table 2 Problems Adding a CIM Host

Situation	Resolution
The machine with the specified IP is not reachable.	Verify whether the machine is reachable from the CIM Client.
The specified protocol (HTTP or HTTPS) is not supported by the CIMOM.	Most often the CIMOM is configured to use HTTPS. Therefore, if you are trying to connect with HTTP, you might get an error. Try using HTTPS instead.
The namespace is invalid.	Verify that the namespace for the Emulex provider is `root/brcmccx`.
The username or password is invalid.	Verify that the username is correct, and retype the password.

The CIMOM is not running on the ESXi host.	You can check whether the CIMOM, a small footprint CIM broker (SFCB), is running by typing one of these commands.
	`/etc/init.d/sfcbd-watchdog status`
	or
	`ps -ef | grep sfcb`
	If the CIMOM is listening to a port other than 5988 or 5989, the connection might not take place. You can configure the SFCB CIMOM settings by editing `/etc/sfcb/sfcb.cfg`.
	In ESXi 6.5, `sfcb` is disabled by default. To enable and run `sfcb`, make the following change in `/etc/sfcb/sfcb.cfg`
	`enabled: false`
	to
	`enabled: true`

The CIM provider is running, but enumerations are not occurring properly.	Verify that the correct CIM provider for the ESXi host is installed.

If you still experience problems when adding the host, run the following commands on the ESXi host and send the output to the Broadcom Technical Support team.

```
vm-support
esxcfg-module -l
esxcfg-scsidevs -a
esxcfg-nics -l
lspci
esxcli software vib list | grep -i be2
esxcli software vib list | grep -i lpfc
esxcli software vib list | grep -i emu
esxcli software vib list | grep -i elxnet
```

Send the `/var/log/syslog.log` file for all of the above operations.


BROADCOM

Emulex® OneCommand® CNA Manager for VMware vCenter for OneConnect Adapters

User Guide

Version 11.2
December 30, 2016

OCM-VM-OCA-UG112

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

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

Copyright © 2011–2017 Broadcom. All Rights Reserved.

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

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.

Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Chapter 1: Introduction

Emulex OneCommand CNA Manager for VMware vCenter is a comprehensive management utility for Emulex adapters that provides a powerful, centralized adapter management suite for the VMware vCenter management console.

This comprehensive solution builds upon standard Emulex Common Information Model (CIM) management providers and advanced functionality delivered with the OneCommand CNA Manager application to present native configuration management, status monitoring, and online maintenance of Emulex adapters in VMware ESXi environments, using a graphical interface (GUI) or a command line interface (CLI).

1.1 Compatibility

For supported versions of operating systems, platforms, and adapters, go to http://www.broadcom.com.

OneCommand CNA Manager for VMware vCenter can be hosted and registered on the following servers:

- vCenter Server 6.5
- vCenter Server 6.0
- vCenter Server 5.5

OneCommand CNA Manager for VMware vCenter is supported for the following Windows operating systems:

- Windows 7 -x64
- Windows 8 -x64
- Windows 8.1 -x64
- Windows Server 2012
- Windows Server 2016

Refer to Table 1 to determine the support provided by the Emulex CIM Provider.

NOTE Illustrations in this guide are for illustrative purposes only. Your system information can vary.

OneCommand CNA Manager for VMware vCenter supports Emulex OneConnect adapters, including OCe14100-series and OCe11000-series adapters. The features available in OneCommand CNA Manager for VMware vCenter are determined by the Emulex CIM Provider version installed.

Table 1 lists the Broadcom Converged Network Adapter (CNA) OneCommand Manager application support provided by the Emulex CIM Provider package and the individual FCoE and CNA providers in each package. The Emulex CIM Provider packages can be downloaded from http://www.broadcom.com.

Table 1 Support Provided by Emulex CNA CIM Provider Versions

Emulex OneCommand CNA Manager Application Features	ELX CIM Provider Package v11.2		OneCommand CNA Manager for VMware vCenter v 11.2
	NIC/iSCSI	FCoE	
Discover virtual ports connected to an FCoE port	x	x	x
View virtual port information in a cluster (host-centric mode)	N/A	N/A	x
Discover hosts, adapters (FCoE, NIC, and Internet Small Computer System Interface [iSCSI]), targets, and logical unit numbers (LUNs) for selected ESXi hosts	x	x	x
Discover FCoE hosts and adapters for selected ESXi fabrics	x	x	x

Table 1 Support Provided by Emulex CNA CIM Provider Versions (Continued)

Emulex OneCommand CNA Manager Application Features	ELX CIM Provider Package v11.2		OneCommand CNA Manager for VMware vCenter v 11.2
	NIC/iSCSI	FCoE	
View firmware version	x	x	x
View boot code version	x	x	x
Update firmware and boot code on a single adapter	x	x	x
Change World Wide Port Name (WWPN) or World Wide Name (WWN) for FCoE adapters	x	x	x
Locate OneConnect adapters with beaconing	x	x	x
View Peripheral Component Interconnect Express (PCIe) registers	x	x	x
PHY loopback test	x	x	x
MAC loopback test	x	N/A	x
Time-domain reflectometer (TDR) tests for 10GBASE-T-series adapters	x	N/A	x
Enable/disable ports	x	x	x
Get driver parameters (global and port)	x	x	x
Set global driver parameters to adapters	x	x	x
Set port driver parameters to adapters	x	x	x
Target and LUN information	x	x	x
Reset Port (FCoE only)	x	x	x
View vital product data (VPD)	x	x	x
View physical data for 10GBASE-T adapters	x	N/A	x
iSCSI management	x	N/A	x
Enable/disable FCoE Initialization Protocol (FIP) data	x	N/A	x
DCB configuration	x	N/A	x
Export storage area network (SAN) configuration information at cluster level and at host level	N/A	N/A	x
Perform diagnostic dump at the adapter and port level	x	x	x
Personality (protocol) change and licensing	x	N/A	x
iSCSI boot firmware table (iBFT) support	x	N/A	x
Channel Management for OCe11100–series adapters	x	N/A	x
Single root I/O virtualization (SR-IOV)	x	N/A	x
Virtual Ethernet port aggregator (VEPA) for OCe14000-series adapters	x	N/A	x

1.2 Abbreviations

AL_PA	Arbitrated Loop Physical Address
ARI	alternative routing-ID interpretation
ASIC	application-specific integrated circuit
BIOS	basic input/output system
BOFM	Blade Open Firmware Management Protocol
BW	bandwidth

CA	certificate authority
CDB	command descriptor block
CHAP	Challenge Handshake Authentication Protocol
CIM	Common Information Model
CIN	Cisco, Intel, Nuova (Data Center Bridging Exchange)
CLI	command line interface
CLP	Command Line Protocol
CNA	converged network adapter
CPU	central processing unit
CRC	cyclic redundancy check
CSR	Certificate Signing Request
CSV	comma-separated value
DAC	direct-attach copper
DCB	Data Center Bridging
DCBX	Data Center Bridging Capabilities Exchange
DHCP	Dynamic Host Control Protocol
DIF	Data Integrity Field
DMA	direct memory access
DNS	Domain Name System or Domain Name Server
DOCSIS	Data Over Cable Service Interface Specification
EC	engineering change
ETO	extended timeout
ETS	Enhanced Transmission Selection
FC	Fibre Channel
FCF	Fibre Channel over Ethernet Forwarder
FCoE	Fibre Channel over Ethernet
FCP	Fibre Channel Protocol
FIP	FCoE Initialization Protocol
FPMA	fabric-provided MAC address
FW	firmware
Gb	gigabit
Gb/s	gigabits per second
GUI	graphical user interface
HBA	host bus adapter
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
HW	hardware
iBFT	iSCSI boot firmware table
IEEE	Institution of Electrical and Electronics Engineers
I/O	input/output
IP	Internet Protocol
IPL	initial program load
IPv6	Internet Protocol version 6

iSCSI	Internet Small Computer System Interface
ISID	initiator session identifier
iSNS	Internet Storage Name Service
KB	Kilobyte (1024 bytes)
LAN	local area network
LDPC	low-density parity check
LDTO	link down timeout
LIP	Loop Initialization Primitive
LLDP	Link Layer Discovery Protocol
LPVID	logical port VLAN ID
LUN	logical unit number
MAC	Media Access Control
MB	megabyte
Mb	megabit
Mb/s	megabits per second
MDI	Medium Dependent Interface
MN	manufacturer ID
MSE	mean squared error
MTU	maximum transmission unit
N/A	not applicable
NCSI	Network Communication Services Interface
NIC	network interface card
NOS	network operating system
NPar	NIC partitioning
NVP	nominal velocity of propagation
NVRAM	nonvolatile random-access memory
OS	operating system
OUI	organizationally unique identifier
PCI	Peripheral Component Interconnect (interface)
PCIe	Peripheral Component Interconnect Express
PDU	protocol data unit
PFC	priority flow control
PGID	priority group ID
PLL	phase-locked loop
POST	power-on self-test
PXE	Preboot Execution Environment
QoS	quality of service
R2T	request to transmits
RFC	Request for Comments
Rx	receive
SAN	storage area network
SCSI	Small Computer System Interface
SFP	small form-factor pluggable

SLI	Service Level Interface
SNMP	Simple Network Management Protocol
SNR	signal-to-noise ratio
SR-IOV	single root I/O virtualization
SSL	secure sockets layer
TCP	Transmission Control Protocol
TCP/IP	TCP over Internet Protocol
TDR	time-domain reflectometer
Tx	Transmit
UFP	Unified Fabric Port
ULP	Upper Layer Protocol
UMC	Universal Multichannel
vCSA	VMware for vCenter Server Appliance
VEPA	virtual Ethernet port aggregator
VF	virtual function
VLAN	virtual local area network
VLAN ID	VLAN identifier
VM	virtual machine
VPD	vital product data
vPort	virtual port
WLAN	wireless LAN
WWN	World Wide Name
WWNN	World Wide Node Name
WWPN	World Wide Port Name
XML	Extensible Markup Language

Chapter 2: Installing and Enabling OneCommand CNA Manager for VMware vCenter

This section describes installing, uninstalling, reinstalling, and enabling OneCommand CNA Manager for VMware vCenter. This section also describes registering and unregistering OneCommand CNA Manager for VMware vCenter.

NOTE OneCommand CNA Manager for VMware vCenter provides real-time management as a plug-in through the VMware vSphere console and VMware Web Client.

The web client allows you to run OneCommand CNA Manager for VMware vCenter from the vSphere Web client. System performance is directly influenced by the speed and efficiency of the underlying network infrastructure.

2.1 Best Installation Practices

- To avoid network latencies, ensure that all the managed ESXi hosts, the vCenter Server, and the VI infrastructure client are in the same network.
- To avoid configuration conflicts, install the vCenter Server and the VI infrastructure client on different hosts within the same network.

2.2 Hardware Requirements

- Physical or virtual (x86/x86_64) servers with a minimum RAM of 2 MB and 250 MB of disk space.

2.3 Software Requirements

- Adobe Flash Player 11.2 or later
- Driver and Firmware Requirements

  Go to http://www.broadcom.com for the latest compatible driver and firmware versions.
- Operating system – Windows 7 (64 bit), Windows 8 (64 bit), Windows Server 2012 (64 bit), Windows Server 2012 R2, or Windows Server 2016
- vCenter Server 5.5, 6.0, and 6.5
- VMware ESXi 5.5, 6.0, and 6.5 environments
- vSphere Client 5.5, 6.0, and 6.5
- vSphere Web Client to manage the VMware vCenter server instead of the vSphere client

  NOTE On the system where OneCommand CNA Manager for VMware vCenter is installed, make sure the port numbers configured during the installation are open and dedicated to the OneCommand CNA Manager for VMware vCenter server only. No other service should be listening on this port.
- Emulex CIM Provider Package 11.2

NOTE 11.2 packages are not compatible with the 11.1 or older version of Emulex software.

2.4 Installing OneCommand CNA Manager for VMware vCenter

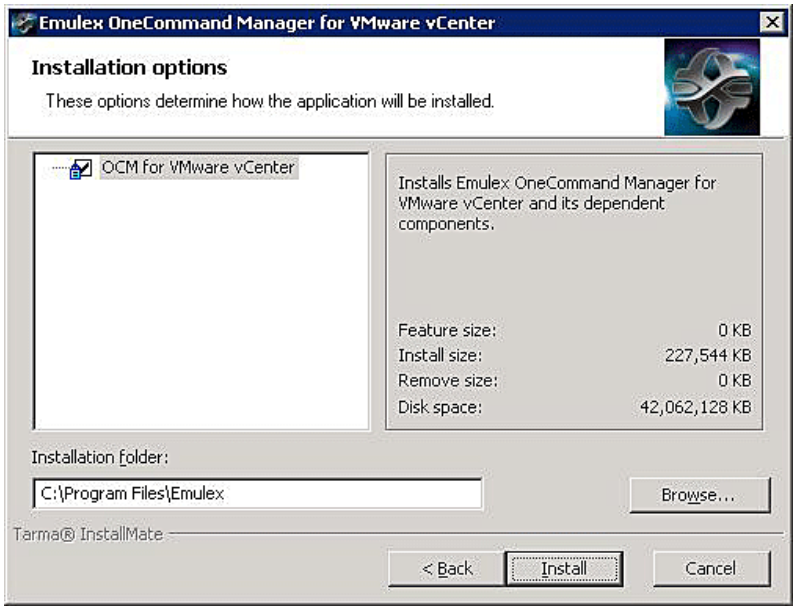
The Emulex CIM Provider must be installed on your ESXi host before you install OneCommand CNA Manager for VMware vCenter. For more information on installing the CIM Provider, see the *Emulex CIM Provider Package Installation Guide for OneConnect Adapters* available on http://www.broadcom.com.

NOTE As part of the OneCommand CNA Manager for VMware vCenter installation, the OneCommand CNA Manager for VMware vCenter Web Client components are installed and registered.

To install OneCommand CNA Manager for VMware vCenter in Windows, perform the following steps:

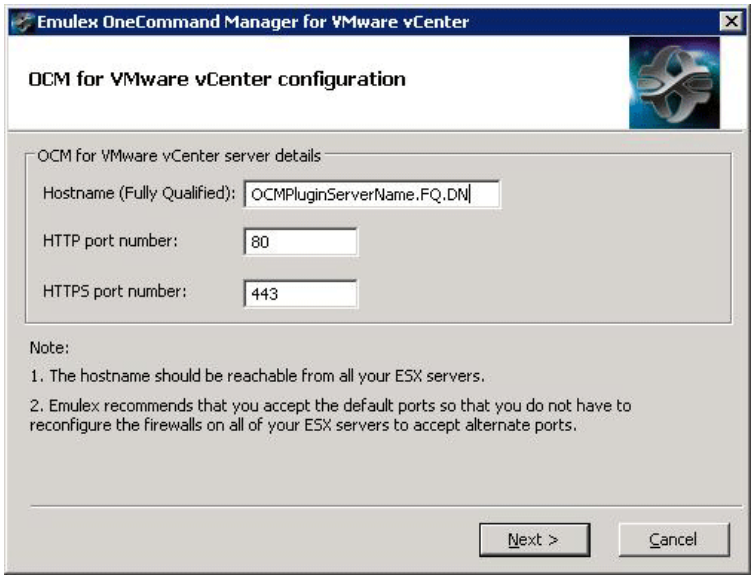
1. Go to http://www.broadcom.com to download the **BrcmOCM-VMware-vCenter-**<*version*>**-Setup.exe** installation file to your system.
2. Navigate to the system directory to which you downloaded the file.
3. Double-click **BrcmOCM-VMware-vCenter-**<*version*>**-Setup.exe**. The OneCommand CNA Manager for VMware vCenter window appears.
4. Click **Next**. The License agreement window appears.
5. Read the agreement and select the **I agree to these terms and conditions** checkbox.
6. Click **Next**. The Installation options window with the default Installation folder appears (Figure 1).

Figure 1 Installation Options Window



7. Ensure that the **OCM for VMware vCenter** checkbox is selected.
8. Program files install by default to `C:\Program Files\Brcm`. To change this location, click **Browse** and navigate to where you want the program files to reside.
9. Click **Install**. The Operation in progress screen appears. When the process is complete, the OCM for VMware vCenter configuration window appears (Figure 2).

Figure 2 OCM for VMware vCenter configuration Window

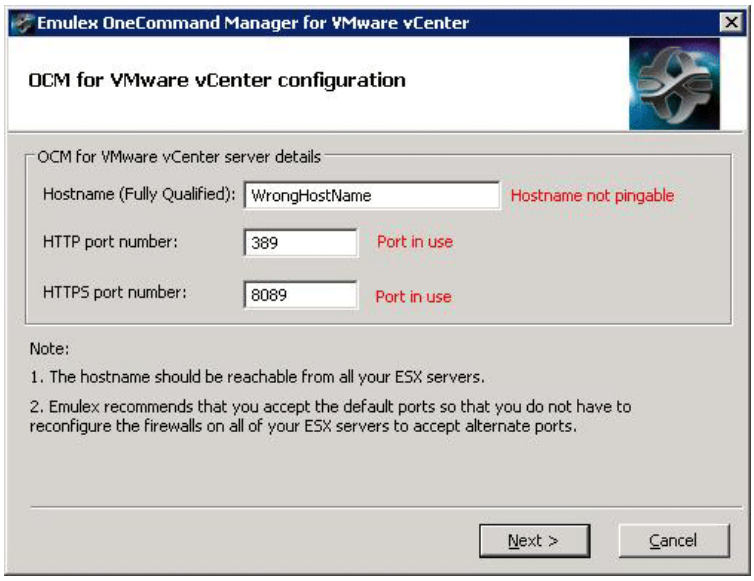


Default port numbers for OneCommand CNA Manager for VMware vCenter server are displayed.

NOTE The Windows firewall setting must allow incoming connections on the HTTP and HTTPS ports that you configure here.

If the port numbers are already in use, a warning appears next to the port number (Figure 3).

Figure 3 OCM for VMware vCenter Configuration Window with Port in use Warning



10. Click **Next**. The Operation in progress screen appears. When the installation process is complete, a message prompts you to launch the registration utility.
11. Click **Yes**. The Register/Unregister dialog appears in a new browser window (Figure 4).

Figure 4 Register/Unregister Dialog



12. Enter the following details of the vCenter Server:
    - vCenter Server Name – The IP address of the vCenter Server.
    - vCenter Server HTTPS Port – The HTTPS port number of the vCenter Server.

      NOTE The vCenter Server HTTPS port is 443 by default. You can change this value if you have configured a different HTTPS port while installing the vCenter.
    - Username – The user name with required privileges.
    - Password – The user password.
13. Click **Register** to register OneCommand CNA Manager for VMware vCenter with a new vCenter Server.

    NOTE You can unregister an existing OneCommand CNA Manager for VMware vCenter by clicking **Unregister**.

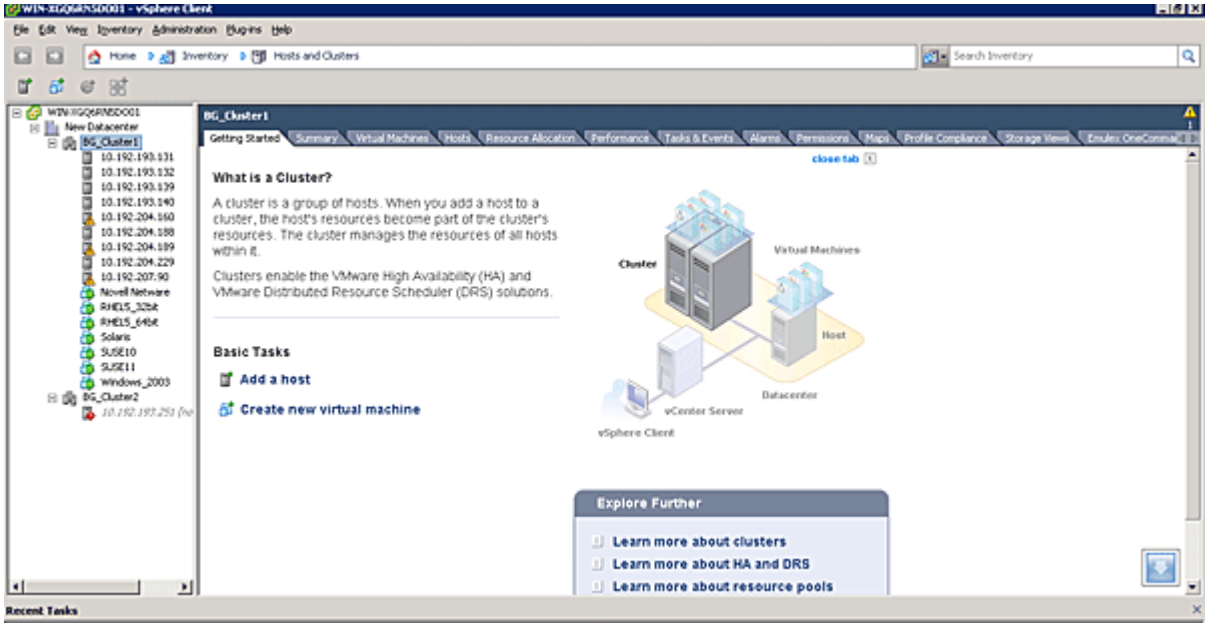
    NOTE If you change the host name of the machine hosting the vCenter server, you must re-install the vCenter server and re-register.
14. When the operation is successful, a message is displayed. Click **OK**.
15. Close the browser window. The Installation completed window appears.
16. Click **Finish**. The OneCommand CNA Manager for VMware vCenter Registration icon is created on the desktop. You do not need to reboot the system.

2.4.1 Verifying the OneCommand CNA Manager for VMware vCenter Installation

To verify the OneCommand CNA Manager for VMware vCenter installation, perform the following steps:

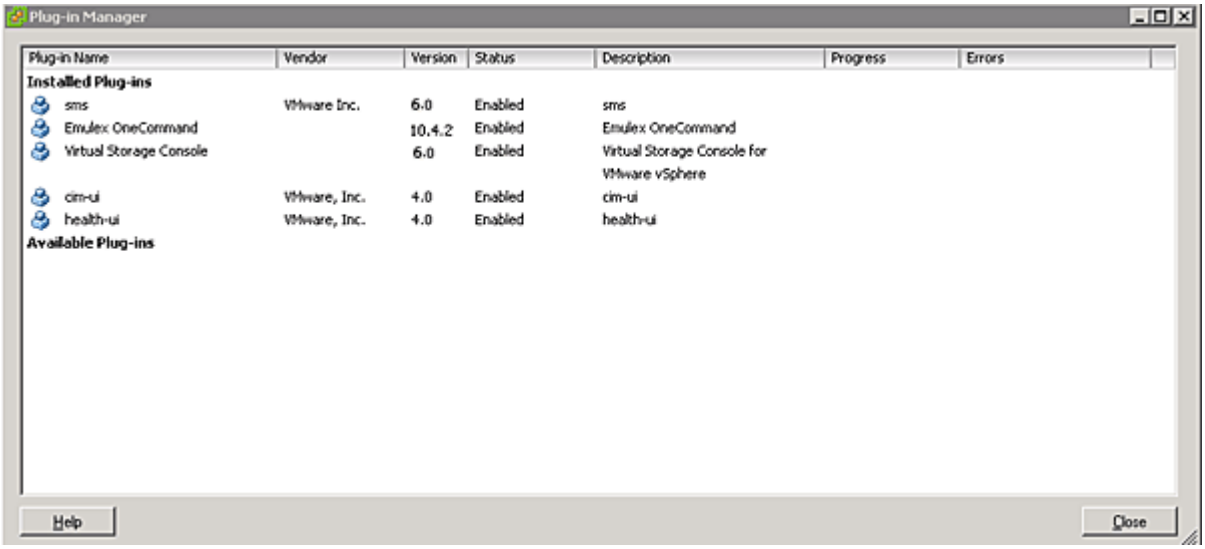
1. Log in to the vCenter Server through the VMware vSphere console.
2. Enter the IP address and credentials of the vCenter Server where OneCommand CNA Manager for VMware vCenter is registered. The vSphere console appears (Figure 5).

Figure 5 vSphere Console, Getting Started Tab



3. In the vSphere console, select **Plug-ins** from the menu bar and select **Manage Plug-ins**. The Plug-in Manager window appears (Figure 6).

Figure 6 Plug-in Manager Window



4. In the Plug-in Manager window, note the status of OneCommand Manager for VMware vCenter (Emulex CNA OneCommand). If the OneCommand CNA Manager for VMware vCenter installation is complete, the status of Emulex OneCommand is enabled by default.

2.5 Activating and Enabling ESXi Management

This section describes activating OneCommand CNA Manager for VMware vCenter in the web client, enabling OneCommand Manager for vCSA, and enabling OneCommand CNA Manager for VMware vCenter in the web client.

NOTE Refer to the vSphere guide on the VMware website for information on creating a user with required privileges and changing access permissions for a user in the Active Directory.

2.5.1 Requirements

Only a user with these specific privileges can read and perform active management in OneCommand CNA Manager for VMware vCenter:

- **Extension.Register extension** to register OneCommand CNA Manager for VMware vCenter using the registration utility.
- **Extension.Unregister extension** to unregister OneCommand CNA Manager for VMware vCenter using the registration utility.
- **Host.CIM.CIM Interaction** to read and manage data through the OneCommand CNA Manager for VMware vCenter

All other users, including the root user, of the ESXi host cannot perform any actions including reading data. If a user without the required privileges attempts to perform an action in OneCommand CNA Manager for VMware vCenter, an error message is displayed.

NOTE To configure user roles and assign privileges, refer to the *VMware vCenter Server Manual* on the VMware website.

2.5.2 Lockdown Mode Feature

Refer to the vSphere manual on the VMware website for information on enabling and disabling lockdown mode.

If lockdown mode is enabled for an ESXi host, then only a user with the required privileges can access the ESXi host and manage adapters using OneCommand CNA Manager for VMware vCenter. All other users, including the root user, will not have access to the ESXi host.

2.5.3 Activating OneCommand CNA Manager for VMware vCenter in the Web Client

To activate OneCommand CNA Manager for VMware vCenter, which allows it to be accessed from the **Manage** tab, you must modify the webclient.properties file on the server where the vSphere web client is installed.

1. Stop the vSphere web client service.
2. Locate the webclient.properties file in the vSphere Web Client install directory, which is typically the following:

   ```
   %ProgramData%\VMware\vSphere Web Client
   ```

3. Add the following line:

   ```
   scriptPlugin.enabled = true
   ```

4. Save and close the file.
5. Restart the vSphere web client service.

2.5.4 Enabling OneCommand CNA Manager for VMware for vCSA

vCSA is a preconfigured Linux-based virtual machine that runs the vCenter Server and associated services. vCSA is an alternative for centralized management of VMware vSphere instead of Windows vCenter Server.

To enable OneCommand CNA Manager for VMware vCenter on vCSA:

1. Edit webclient.properties file at:

   ```
   /var/lib/vmware/vsphere-client/webclient.properties
   ```

   and add the statement:

   ```
   scriptPlugin.enabled = true
   ```

2. Save the file.
3. Restart the vSphere Web Client service on the web console.
4. Log in to vCSA again.

The **Manage** tab for the OneCommand CNA Manager for VMware vCenter is now displayed.

2.5.5 Enabling OneCommand CNA Manager for VMware vCenter Web Client

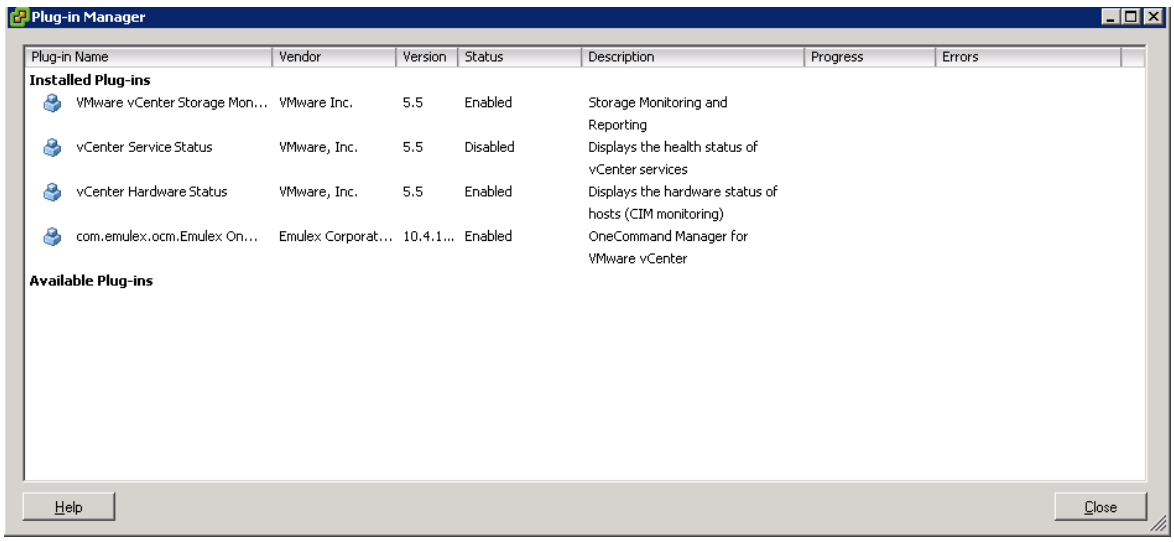
OneCommand CNA Manager for VMware vCenter can be enabled or disabled from the Plug-In Management section in the vSphere Web Client.

NOTE You must have sufficient privileges to access the Plug-In Management section. Refer to the VMware documentation for information on configuring users and privileges.

To enable OneCommand CNA Manager for VMware vCenter in the vSphere Web Client, perform the following steps:

1. From the navigation panel on the left-side of the vSphere Web Client Home page, click **Plugins > Manage Plug-ins**. The Plug-in Manager page is displayed (Figure 7).

Figure 7 vSphere Web Client Plug-In Management Page



2. In the Plug-In Manager page, select **com.emulex.Emulex OneCommand** and right-click under the **Status** column. A context menu opens.

3. Select **Enable** to enable OneCommand CNA Manager for VMware vCenter or select **Disable** to disable OneCommand CNA Manager for VMware vCenter.

2.6 Enabling and Disabling OneCommand CNA Manager for VMware vCenter with the Plug-in Manager

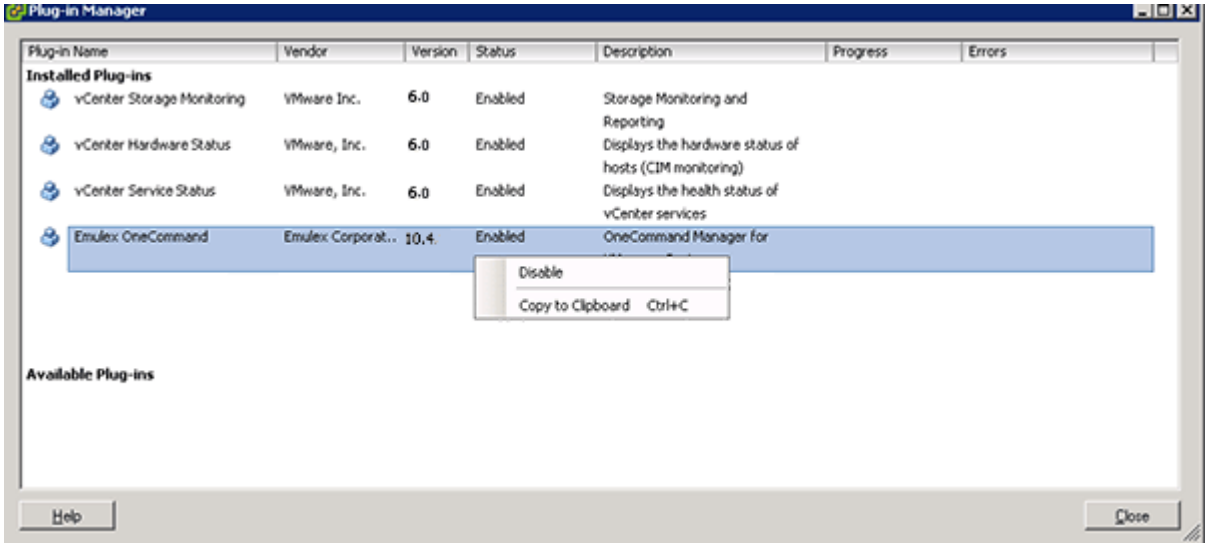
If you enable OneCommand CNA Manager for VMware vCenter, the OneCommand CNA Manager for VMware vCenter functionality immediately appears on the vSphere console. If two or more plug-ins are registered, the **Broadcom CNA OneCommand** sub-tab (under the **Classic Solutions** tab) immediately appears on the vSphere console. If only one plug-in is registered, the OneCommand CNA Manager for VMware vCenter functionality immediately appears on the **Classic Solutions** tab.

If you disable OneCommand CNA Manager for VMware vCenter, the OneCommand CNA Manager for VMware vCenter functionality immediately disappears. If two or more plug-ins are registered, the **Broadcom CNA OneCommand** sub-tab immediately disappears from the vSphere console. If only one plug-in is registered, the OneCommand CNA Manager for VMware vCenter functionality immediately disappears from the **Classic Solutions** tab.

To change the status of OneCommand CNA Manager for VMware vCenter, perform the following steps:

1. Log in to the vCenter Server through the VMware vSphere console.
2. Enter the IP address and credentials of the vCenter Server where OneCommand CNA Manager for VMware vCenter is registered. The vSphere console appears.
3. In the vSphere console, select **Plug-ins** from the menu bar and select **Manage Plug-ins**. The Plug-in Manager window appears.
4. Click the Emulex OneCommand row. A context menu appears (Figure 8).

Figure 8 Plug-in Manager with Selected Row



5. Select **Enable** or **Disable**.

2.7 Registering and Unregistering OneCommand CNA Manager for VMware vCenter

OneCommand CNA Manager for VMware vCenter can be registered with more than one vCenter Server. If you are using the vSphere Web Client, you must verify that the vSphere Web Client and the vCenter server are registered to the same vCenter Single Sign On server to ensure that the vSphere Web Client can access the vCenter server inventory.

To register or unregister OneCommand CNA Manager for VMware vCenter with a new vCenter Server, perform the following steps:

1. Double-click the **OCM for VMware vCenter Registration** icon on the desktop. This icon is created when OneCommand CNA Manager for VMware vCenter is successfully installed. The register/unregister dialog is displayed (Figure 4).
2. Enter the following details of the vCenter Server:
   - vCenter Server Name – The IP address of the vCenter Server.
   - vCenter Server HTTPS Port – The HTTPS port number of the vCenter Server.
   - The vCenter Server HTTPS port is 443 by default. You can change this value if you have configured a different HTTPS port while installing the vCenter.
   - Username – The user name with required privileges.
   - Password – The user password.
3. Do one of the following:
   - Click **Register** to register OneCommand CNA Manager for VMware vCenter with a new vCenter Server.

   or

   - Click **Unregister** to unregister an existing OneCommand CNA Manager for VMware vCenter with a vCenter Server.

   NOTES
   - If you change the host name of the machine hosting the vCenter server, you must re-install the vCenter server and re-register.
   - If the vCenter server is already registered with another instance of OneCommand CNA Manager for VMware server, it is replaced with this server instance.
4. When the operation is successful, a message is displayed. Click **OK**.
5. Close the window.

2.8 Uninstalling OneCommand CNA Manager for VMware vCenter

Before you uninstall OneCommand CNA Manager for VMware vCenter, you must unregister it from the vCenter Server. For more information, refer to Section 3, Using OneCommand CNA Manager for VMware vCenter.

CAUTION When you uninstall OneCommand CNA Manager for VMware vCenter, ensure that you do not delete the default configuration and log files that are stored in the `%TEMP%\Emulex\OCM for VMware` directory. If these files are deleted, all historical information of active management performed from the host is permanently lost.

To uninstall OneCommand CNA Manager for VMware vCenter in a Windows system, perform the following steps:

1. Navigate to the system directory to which you downloaded the `BrcmOCM-VMware-vCenter-<version>-Setup.exe` file.

2. Double-click the `BrcmOCM-VMware-vCenter-<version>-Setup.exe` file. The OneCommand CNA Manager for VMware vCenter window prompts you to reinstall/uninstall the application. Select **Uninstall the application completely** and click **Next**. A progress window is displayed. The screen indicating the detection of OneCommand CNA Manager for VMware vCenter appears.
3. Click **OK**. When uninstallation is complete, the Uninstallation Completed window is displayed.
4. Click **Finish**. You do not need to reboot the system.

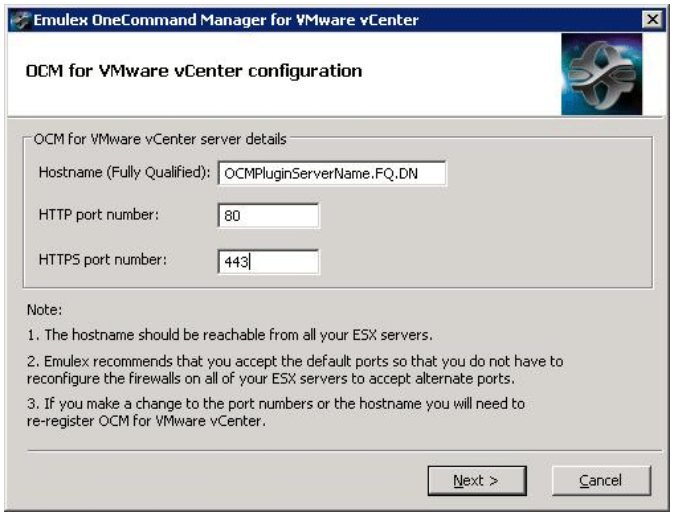
NOTE You can also uninstall the OneCommand CNA Manager application from the Programs and Features window.

2.9 Reinstalling OneCommand CNA Manager for VMware vCenter

To reinstall OneCommand CNA Manager for VMware vCenter in a Windows system, perform the following steps:

1. Navigate to the system directory to which you downloaded the `BrcmOCM-VMware-vCenter-<version>-Setup.exe` file.
2. Double-click `BrcmOCM-VMware-vCenter-<version>-Setup.exe`. The OneCommand CNA Manager for VMware vCenter window prompts you to reinstall/uninstall the application.
3. Select **Re-install the application** and click **Next**. A progress window is displayed. The OCM for VMware vCenter window appears.
4. Click **Next**. The License agreement screen appears.
5. Read the agreement and select the **I agree to these terms and conditions** checkbox.
6. Click **Next**. The Installation Options window with the previous installation folder location appears (see Figure 1).
7. Ensure that the **OCM for VMware vCenter** checkbox is selected.
8. To change the installation folder location, click **Browse** and navigate to where you want the program files to reside.
9. Click **Install** on the Installation Options window. The operation in progress window appears. When the installation process is complete, the OneCommand CNA Manager for VMware vCenter configuration window appears (Figure 9).

Figure 9 Reinstallation Configuration Dialog



If OneCommand CNA Manager for VMware vCenter was installed earlier with port numbers other than the defaults provided, those configured ports are shown. If the port numbers are already in use, a warning appears next to the port number.

NOTE The Windows firewall setting must allow incoming connections on the HTTP and HTTPS ports that you configure here.

10. Follow the on-screen instructions and complete the installation with steps 10 to 15 of Section 1.4, Installing OneCommand CNA Manager for VMware vCenter.

Chapter 3: Using OneCommand CNA Manager for VMware vCenter

OneCommand CNA Manager for VMware vCenter is available at the host level and the cluster level in the vSphere inventory list.

NOTE To increase the size of the OneCommand CNA Manager for VMware screen, the Recent Tasks panel on the right can be unpinned and collapsed.

3.1 Accessing with VMware Web Client 6.0 and 6.5

After you are logged into the VMware Web Client, the OneCommand CNA Manager for VMware vCenter version 6.0 or version 6.5 is under the **Manage** tab for a particular host or cluster that you select in the vSphere client. When the **Broadcom CNA OneCommand** tab is loaded, the functionality and features available are the same as those for the OneCommand CNA Manager for VMware vCenter available for vSphere console.

To launch the OneCommand CNA Manager for VMware vCenter Web Client, perform the following steps:

1. Log in to the vSphere Web Client. The vSphere Web Client home page is displayed.
2. Navigate to an ESXi host or cluster in the Navigation pane.
3. From the Host level view, select the host you want to display.

   or

   From the Cluster level view, select the cluster you want to display.
4. Go to the **Manage** tab to access the OneCommand CNA Manager for VMware vCenter.

3.2 Accessing with VMware Web Client Version 5.5

After you are logged into the vSphere client, the OneCommand CNA Manager for VMware vCenter version 5.5 is under the **Classic Solutions** tab for a particular host or cluster that you select in the vSphere client. When the **Classic Solutions** tab is loaded, the functionality and features available are the same as those for the OneCommand CNA Manager for VMware vCenter available for vSphere console.

To launch the OneCommand CNA Manager for VMware vCenter Web Client, perform the following steps:

1. Log in to the vSphere Web Client. The vSphere Web Client home page is displayed.
2. Navigate to an ESXi host or cluster in the Navigation pane.
3. From the Host level view, select the host you want to display.

   or

   From the Cluster level view, select the cluster you want to display.
4. Go to the **Classic Solutions** tab to access the OneCommand CNA Manager for VMware vCenter.

NOTE If OneCommand CNA Manager for VMware vCenter is the only plug-in registered on the vCenter server, the tab title Broadcom CNA OneCommand does not appear under the **Classic Solutions** tab. OneCommand CNA Manager for VMware vCenter directly loads under the **Classic Solutions** tab with no title. The **Emulex OneCommand** plug-in is shown as a sub-tab only if two or more plug-ins are registered.

3.3 Accessing the vSphere Console

To access the vSphere console, perform the following steps:

1. Log in to the VMware vSphere Client:
   a. Double-click the **VMware vSphere Client** icon.
   b. Enter the credentials and click **Login**.
   c. Select a cluster or host in the vSphere console tree-view, and if applicable, click the **Broadcom CNA OneCommand** tab to perform active management on a cluster or a host.

   NOTE If OneCommand CNA Manager for VMware vCenter is the only plug-in registered on the vCenter server, the tab title Broadcom CNA OneCommand does not appear under the **Classic Solutions** tab. OneCommand CNA Manager for VMware vCenter directly loads under the **Classic Solution**s tab with no title. The Broadcom CNA OneCommand plug-in is shown as a sub-tab only if two or more plug-ins are registered.
2. A Security Alert is displayed. Click **Yes**. The vSphere console is displayed.

3.3.1 vSphere Console Views

The vSphere console supports two views, cluster view and host view.

The difference between the two views is the Broadcom CNA Device Management area, which may contain options or a tree-view.

In a cluster view, you can view fabric-centric, or host-centric information, about all the hosts in a cluster. Figure 10 shows the Broadcom CNA Device Management options available in a host-centric cluster view. You can perform active management simultaneously on a fabric or a group of hosts belonging to a cluster.

In a host view, you can only view information about a single host. Figure 11 shows the Broadcom CNA Device Management tree-view elements for a selected host. In the host view, in addition to selecting an element in the tree-view, you must also select a tab in the information pane to view the information for the element selected in the tree-view.

NOTE A host can be managed by only one vCenter Server at a time.
- The OneCommand CNA Manager for VMware vCenter user interface is auto refreshed every five minutes.

To access the cluster view, select a cluster in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab. The cluster view appears (Figure 10).

To access the host view, select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab. The host view appears (Figure 11).

In both cluster and host views, the OneCommand CNA Manager for VMware vCenter window consists of three major parts:

- The Broadcom CNA Device Management area.
- The Information pane corresponding to the option or element selected in the Broadcom CNA Device Management area.
- The About area shows version information about the Broadcom CNA OneCommand CNA Manager for VMware vCenter.

3.3.2 OneCommand CNA Manager for VMware vCenter Window Elements

The OneCommand CNA Manager for VMware vCenter window (Figure 10 and Figure 11) contains four basic components: the Broadcom CNA Device Management area, the information pane, the console buttons, and the filter options menu (if applicable).

Figure 10 Cluster View with Callouts

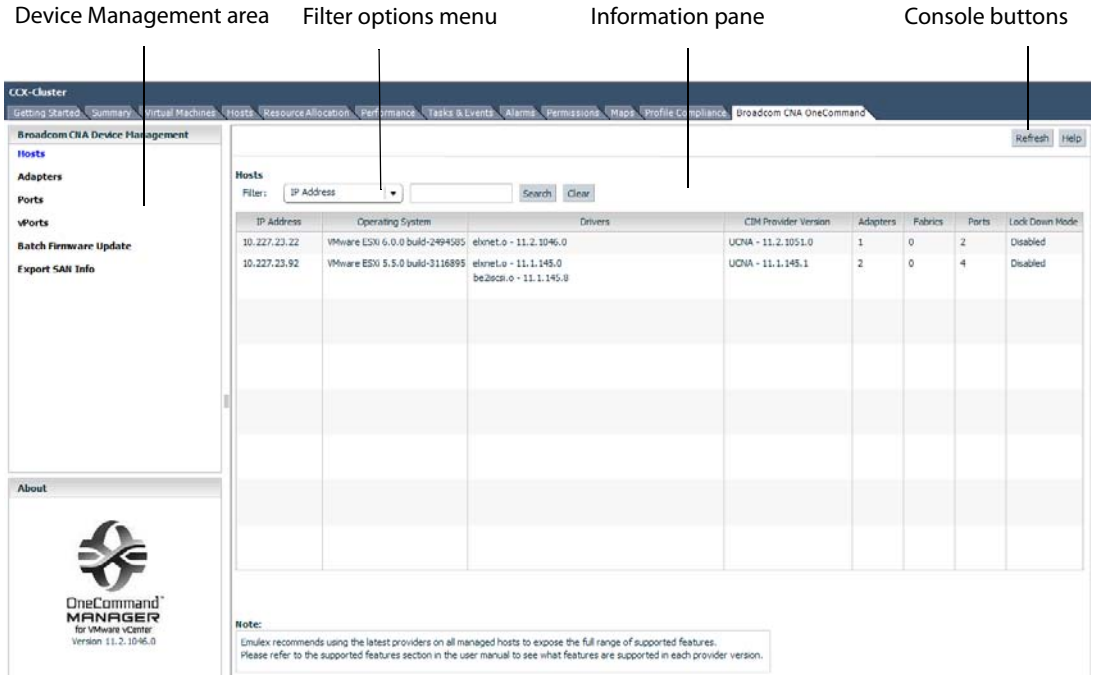
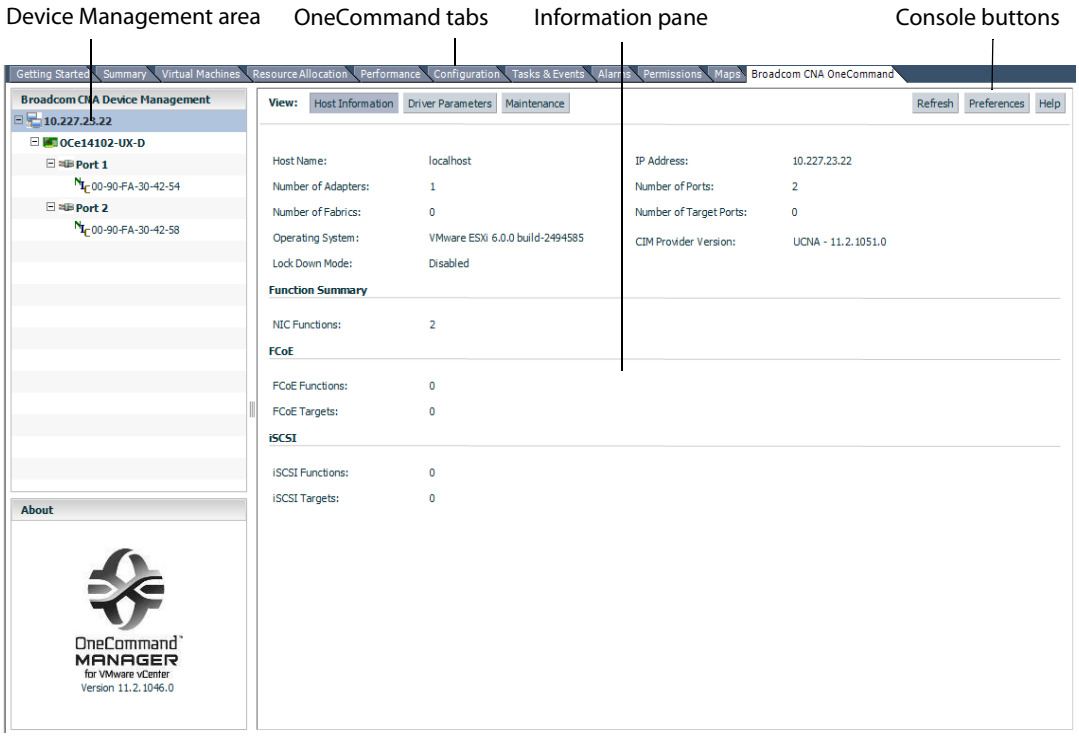


Figure 11 Host View with Callouts



3.3.2.1 Broadcom CNA Device Management Area

In a cluster view, the Broadcom CNA Device Management area contains links that determine what is displayed in the Information pane.

In a host view, the Broadcom CNA Device Management area is a discovery-tree with icons that represent discovered hosts, adapters, ports, virtual ports, fabrics, targets, and LUNs.

3.3.2.2 OneCommand Tabs

In a host view, the **OneCommand** tabs display configuration, statistical, and status information for network elements.

3.3.2.3 Information Pane

In a cluster view, the information pane shows information based upon what is selected in the Broadcom CNA Device Management area.

In a host view, the information pane shows information based upon the **OneCommand** tab that is selected.

3.3.2.4 Filter Options Menu

In a cluster view, selecting an item from the Filter Options menu shows information that is sorted and displayed based upon by the item selected.

3.3.2.5 Console Buttons

- Help – Click to load the complete indexed online help for OneCommand CNA Manager for VMware vCenter. You can search for information for all OneCommand CNA Manager for VMware vCenter tabs and functions.
- Preferences – Click to access the User Preferences window. The User Preference window is available in host view only. In the User Preferences window, select the Event Logging checkbox to show up or down events for FCoE and

iSCSI ports logged into the vSphere console. NIC port events are not logged, as VMware already does this. FCoE and iSCSI port events are limited to the active vCenter client. If the same user login is used from another vCenter client, the User Preference window does not show these events.

- Refresh – Click to refresh cluster or host information.

NOTE The User Preferences window logs only up and down events for the port. Other events, such as temperature, are not posted.

Chapter 4: Managing Clusters and Hosts

This section pertains to viewing cluster and host information.

4.1 Managing Clusters

From within a cluster, you can view information about:

- Hosts in a cluster
- Adapters belonging to hosts in a cluster
- Physical ports in host-centric mode
- Virtual ports in host-centric mode
- Physical ports in fabric-centric mode

Figure 10 shows the main elements of the cluster view.

4.1.1 Viewing Hosts in a Cluster

NOTE Hosts in a cluster with different provider versions support features as listed in Table 1.

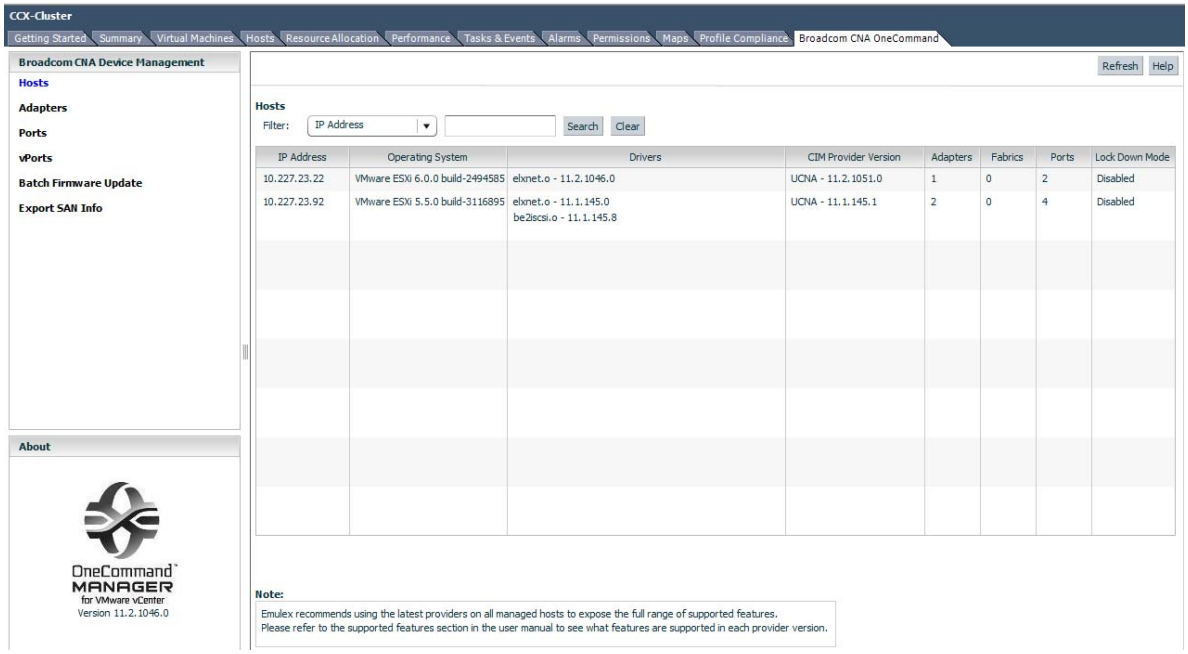
To view information about hosts in a cluster, perform the following steps:

1. Select a cluster in the vSphere console tree-view. The **Getting Started** tab is selected by default.
2. If applicable, select the **Broadcom CNA OneCommand** tab. The Hosts information pane is displayed (Figure 12).

To filter information for a host within a cluster field, perform the following steps:

1. Use the Filter list to filter information.
   a. Search by any column title by selecting the title from the Filter list.
   b. To further narrow the search, enter a value in the field to the right of the Filter list. You can also enter a wildcard for this field.
   c. Click **Search**.
2. Click **Clear** to clear the search criteria.
3. Click **Refresh** to refresh the information. If a host is added to the currently selected cluster, the host information is refreshed.

Figure 12 Hosts within a Cluster



The following Hosts within a Cluster fields are displayed:

- IP Address – The IP address of the host in the cluster.
- Operating System – The operating system and version installed on the host.
- Drivers – The drivers and their versions installed on the host.
- CIM Provider Version – The version of the Emulex CIM Provider that is running on the ESXi host.
- Adapters – The number of adapters installed in the host.
- Fabrics – The number of fabrics to which the host is connected.
- Ports – The number of discovered physical ports that can be managed by the host.
- Lock Down Mode – Specifies if lockdown mode is enabled or disabled.

4.1.2 Viewing Adapters in a Cluster

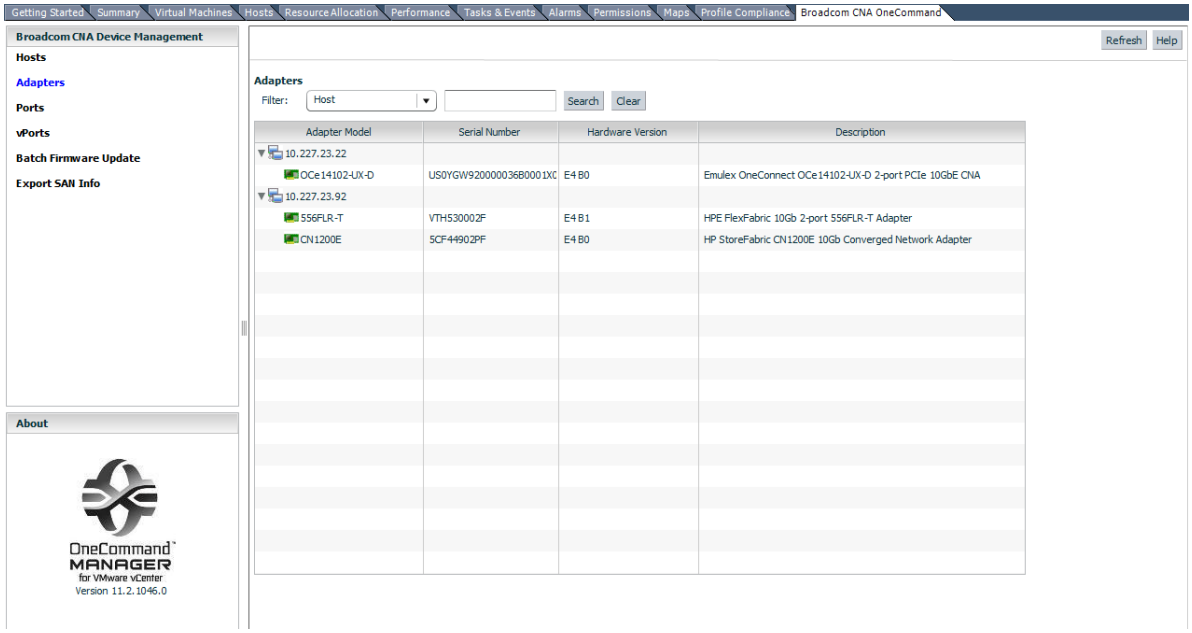
To view information about adapters belonging to a host in a cluster, perform the following steps:

1. Select a cluster in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. From the Broadcom CNA Device Management options, select **Adapters**. Adapters information is displayed (Figure 13).

To filter adapter information within a cluster:

1. Use the Filter list to filter the adapter information. Choose any of the available adapter information fields from the list and enter a wildcard for the field and click **Search**.
2. Click **Clear** to clear the search criteria and the corresponding adapter information.
3. Click **Refresh** to refresh the information. If an adapter is added to any of the hosts, the adapter information is refreshed.

Figure 13 Adapters within a Cluster



The following Adapters within a Cluster fields are displayed:

- Adapter Model – The model of the adapter.
- Serial Number – The serial number of the adapter.
- Hardware Version – The board type and code version.
- Description – The type of adapter.

4.1.3 Viewing Physical Ports in a Cluster (Host-centric mode)

To view information about a physical port that is in a cluster, in host-centric mode, perform the following steps:

1. Select a cluster in the vSphere console tree-view.
2. Select the **Broadcom CNA OneCommand** tab. Hosts information is displayed.
3. From the Broadcom CNA Device Management options, select **Ports**. The host-centric information pane is displayed (Figure 14 and Figure 15).

NOTE Make sure the **Group by Fabric** checkbox is cleared.

Figure 14 Physical Ports within a Cluster, Host-centric View

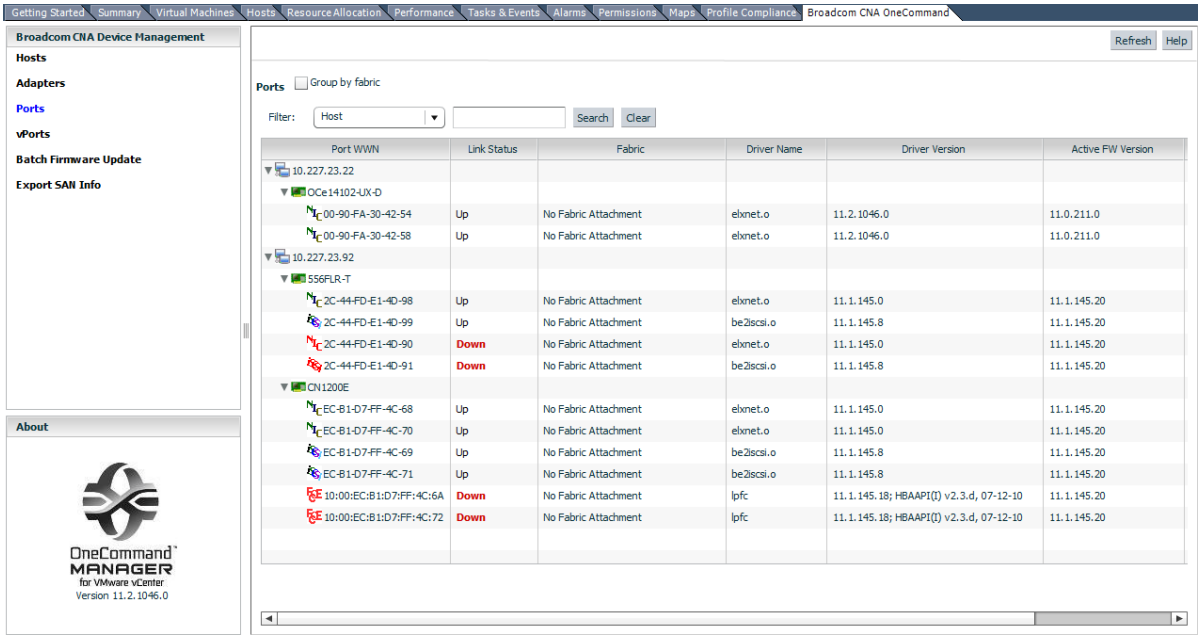
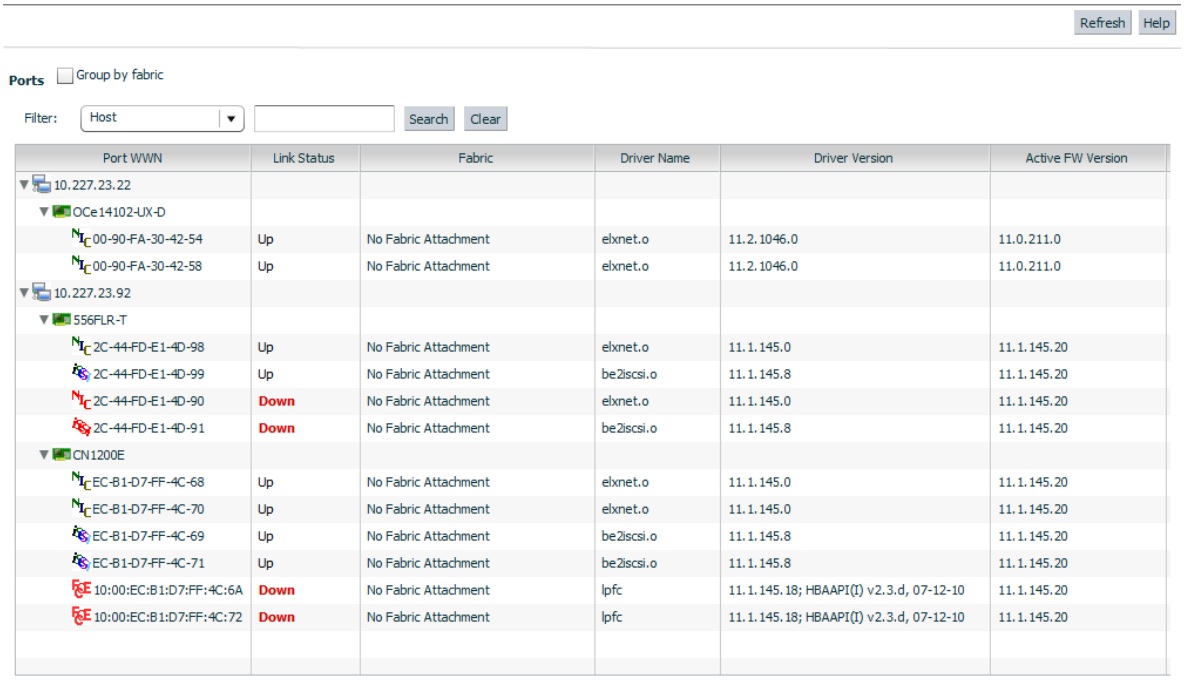


Figure 15 Close-Up of Physical Ports within a Cluster, Host-centric View



The following Physical Ports within a Cluster, Host-centric fields are displayed:

- Port WWN – The port World Wide Name.
- Link Status – The status of the link on the selected port.
- Fabric – This is a 64-bit worldwide unique identifier assigned to the fabric.
- Driver Name – The executable file image name for the driver as it appears in the Emulex driver download package.

- Driver Version – The version of the installed driver.
- Active FW Version – The active FW version.
- Flash FW Version – The FW version that becomes active after the system is rebooted.

4.1.4 Viewing Virtual Ports in a Cluster (Host-centric mode)

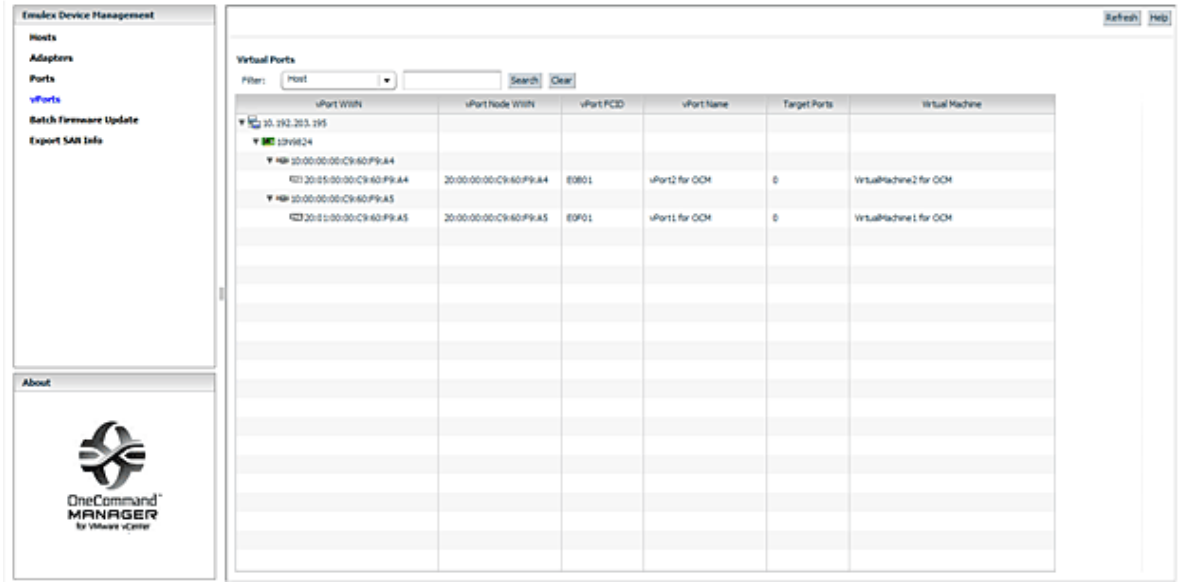
To view information about a virtual port that is in a cluster, in host-centric mode, perform the following steps:

1. Select a cluster in the vSphere console tree-view.
2. If applicable, select the **Broadcom CNA OneCommand** tab. Port host information is displayed.
3. From the Broadcom CNA Device Management options, select **vPorts**. The virtual port information is displayed (Figure 16).

To filter virtual port information in a cluster:

1. Use the Filter list to filter the port information.
2. Choose any of the available host information fields from the list and enter a wildcard for the field and click **Search**.
3. Click **Clear** to clear the search criteria.

Figure 16 Virtual Ports in a Cluster, Host-centric View



The following Virtual Ports in a Cluster, Host-centric fields are displayed:

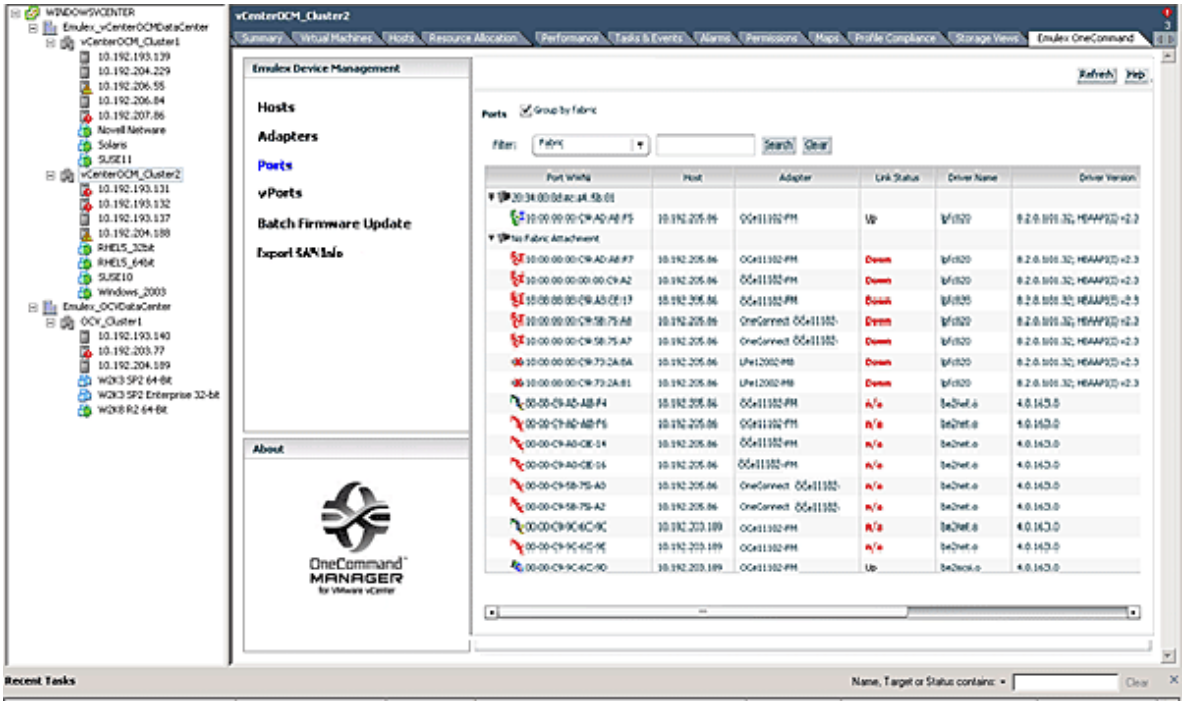
- vPort WWN – The virtual port World Wide Name.
- vPort Node WWN – The virtual port node World Wide Name.
- vPort FCID – The virtual port FC ID.
- vPort Name – The virtual port name.
- Target Ports – The number of target ports.
- Virtual Machine – Virtual machine information.

4.1.5 Viewing Physical Port Information in a Cluster (Fabric- centric mode)

To view physical port information in a cluster, in fabric-centric mode, perform the following steps:

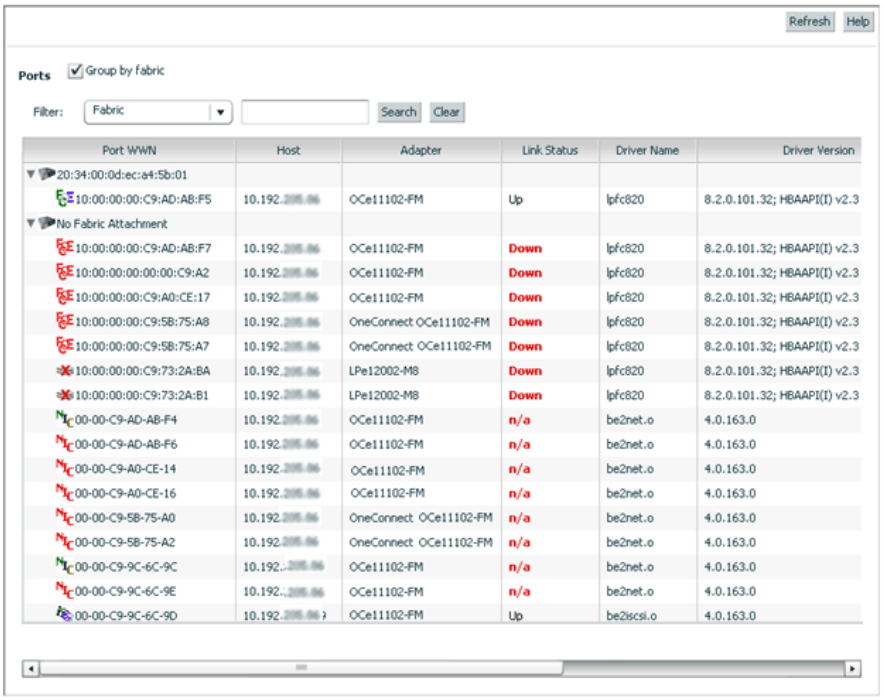
1. Select a cluster in the vSphere console tree-view.
2. Select the **Broadcom CNA OneCommand** tab, host information is displayed.
3. From the Broadcom CNA Device Management options, select **Ports**. The host-centric information pane is displayed.
4. Select the **Group by Fabric** checkbox. Fabric information is displayed (Figure 17 and Figure 18).

Figure 17 Information for a Physical Port in a Cluster, Fabric-Centric View



NOTE iSCSI and NIC ports are not attached to any fabric, so they are listed under No Fabric Attachment.

Figure 18 Close-up of Information for a Physical Port in a Cluster, Fabric-Centric View



The following Fabric-centric Information fields are displayed:

- Port WWN – The port World Wide Name.
- Host – The host IP address.
- Adapter – The adapter model.
- Link Status – The status of the link on the selected port.
- Driver Name – The executable file image name for the driver as it appears in the Emulex driver download package.
- Driver Version – The version of the installed driver.
- Active FW Version – The active firmware version.
- Flash FW Version – The firmware version that becomes active after the system is rebooted.

4.2 Managing Hosts

Host information includes:

- Information for a single host
- Driver parameters for all adapters in a host
- Firmware information for all adapters in a host

Figure 11 shows the main elements of the host view.

4.2.1 Viewing Host Information for a Single Host

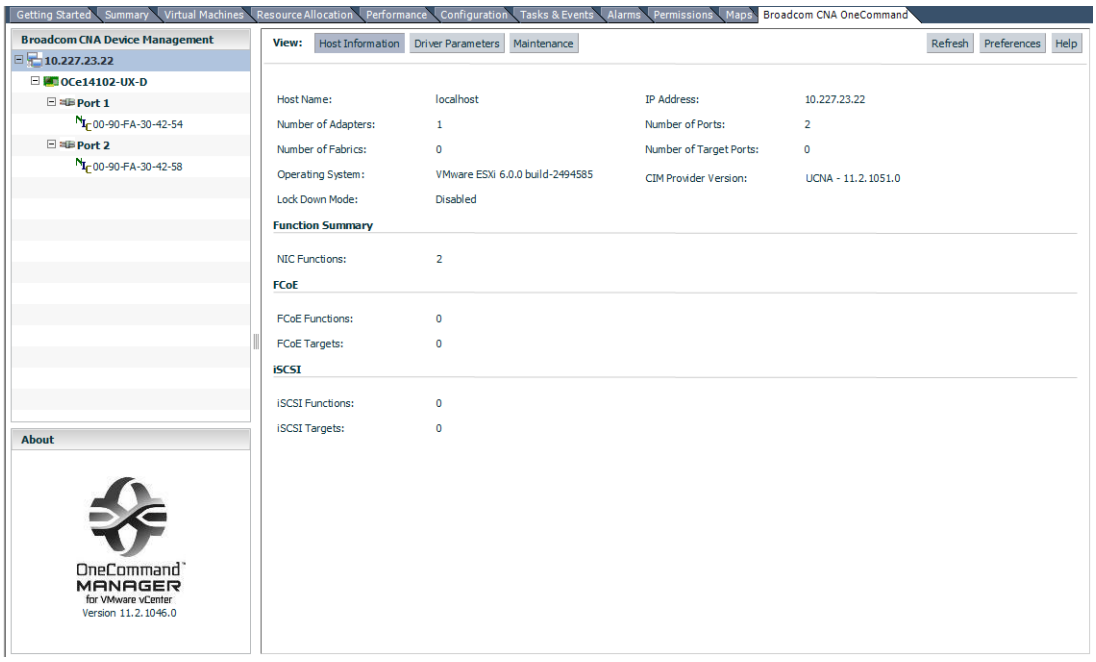
To view host information for a single host, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab. The **Host Information** tab is selected by default and the information of the selected host appears (Figure 19).

To filter the host information:

1. Use the Filter list to filter the fabric information.
2. Choose any of the available fabric information fields from the list and enter a wildcard for the field and click **Search**.
3. Click **Clear** to clear the search criteria.

Figure 19 Information for a Single Host



The following Host Information fields are displayed:

- Host Name – The host identifier.
- Number of Adapters – The number of adapters installed in the host.
- Number of Fabrics – The number of fabrics to which the host is connected.
- Operating System – The operating system and version installed on the selected host.
- Lock Down Mode – Indicates whether lockdown mode is enabled or disabled.
- IP Address – The IP address of the host.
- Number of Ports – The number of discovered physical ports that can be managed by this host.
- Number of Target Ports – The number of targets discovered across the ports.
- CIM Provider Version – The versions of the Emulex CIM Providers that are running on the ESXi host.
- Function Summary area (NIC):
  - NIC Functions – The number of NIC functions running on the discovered adapters on this host.
- Function Summary area (FCoE):
  - 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 FCoE functions on this host.

- Function Summary area (iSCSI):
  - 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.

4.2.2 Viewing Driver Parameters of All Adapters in a Host

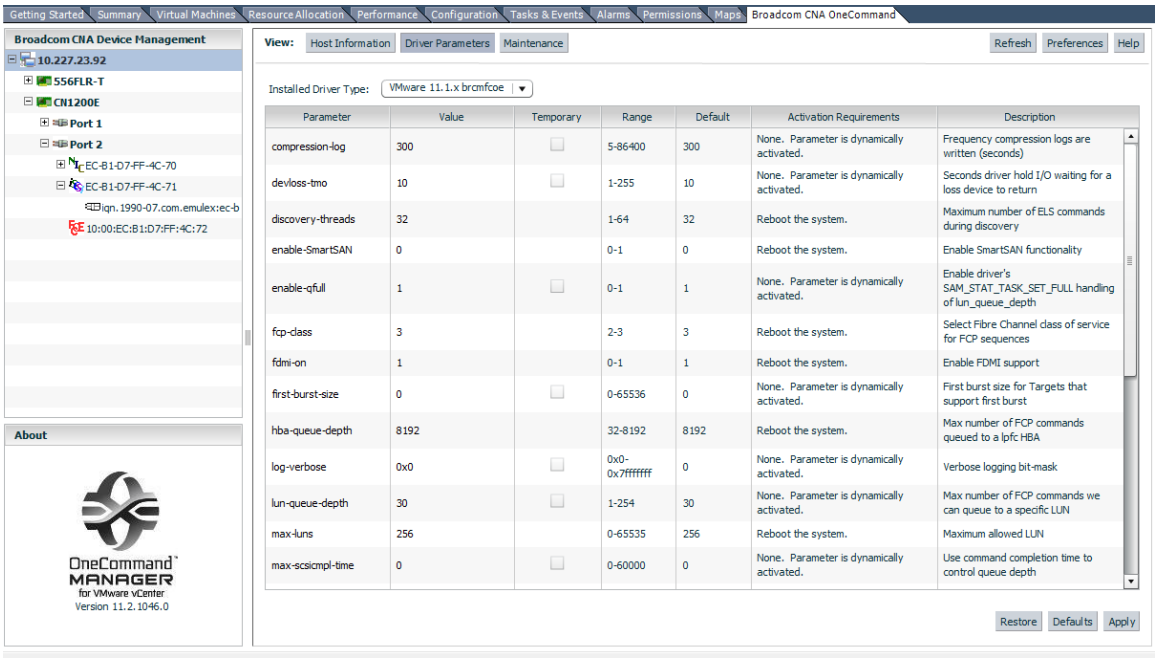
The host **Driver Parameters** tab enables you to view and edit the adapter driver parameter settings for a specific host. The host driver parameters are global values and apply to all adapters in that host unless they are overridden by parameters assigned to a specific adapter using the port **Driver Parameters** tab. For each parameter, the information pane 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 the host, refer to Section 6.2.5, Configuring FCoE Port Driver Parameters.

NOTE If there are no discovered FCoE adapters, the driver parameters table is empty. This occurs because there are no adapters to which the host driver parameters apply.

To view driver parameters for all adapters in a host, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. Select the **Driver Parameters** tab (Figure 20).

Figure 20 Host Driver Parameters



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

- Installed Driver Type – The current driver installed on this host.
- Driver Parameter table – A list of adapter driver parameters and their current values.

  Driver parameter specific information is displayed in this area. This includes value, range, default, activation requirements and description.

- Parameter – The name of the driver parameter.
- Value – The value of the driver parameter.
- Temporary – Indicates if the value can be set temporarily at port level.
- Range – The range of acceptable values for the driver parameter.
- Default – The default value of the parameter.
- Activation Requirements – The steps required to activate the changed value of the driver parameter.
- Description – The description of the driver parameter.

To change the driver parameters for all adapters in a host, perform the following steps:

1. From the vSphere console tree-view, select the host whose adapter driver parameters you want to change. If applicable, select the **Broadcom CNA OneCommand** tab.
2. Select the host **Driver Parameters** tab (Figure 20). 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.
3. In the driver parameter table, click the **Value** field of a parameter that you want to change. The range for the value is displayed. Enter a value in decimal or hexadecimal format, depending on how the current value is presented. If the value is in hexadecimal format, it is prefixed by 0x (for example -0x2d).
4. Parameters that can be changed temporarily can only be changed from the corresponding port. Such parameters are represented with a checkbox next to them.
5. Click **Apply**.

If you changed some parameters and did not click Apply, you can restore the parameters back to the value they had before you made the changes. To restore the parameters, click **Restore**.

To reset all parameters back to their default values, click **Defaults**.

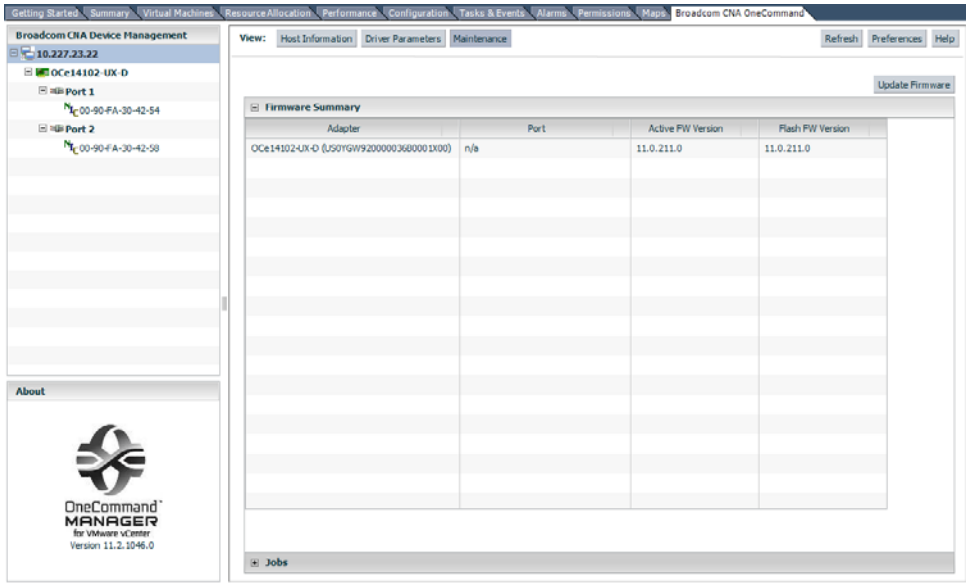
4.2.3 Viewing Firmware Information for All Adapters in a Host

The host **Maintenance** tab enables you to view and update firmware on multiple adapters in a specific host.

To view firmware for all adapters in a host, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. Select the **Maintenance** tab (Figure 21).

Figure 21 Host Maintenance Tab



The following Host **Maintenance** tab Fields are displayed:

- Adapter – The model of the adapter.
- Port – The port WWN.
- Active FW Version – The current firmware on the adapter.
- Flash FW Version – The flashed firmware on the adapter.

Chapter 5: Managing Adapters and Ports

This section describes the various adapter and port management functions you can perform using OneCommand CNA Manager for VMware vCenter.

5.1 OneConnect Adapters

This section pertains to viewing general OneConnect adapter information, as well as installing licenses, configuring personalities, configuring parameters, and managing channels.

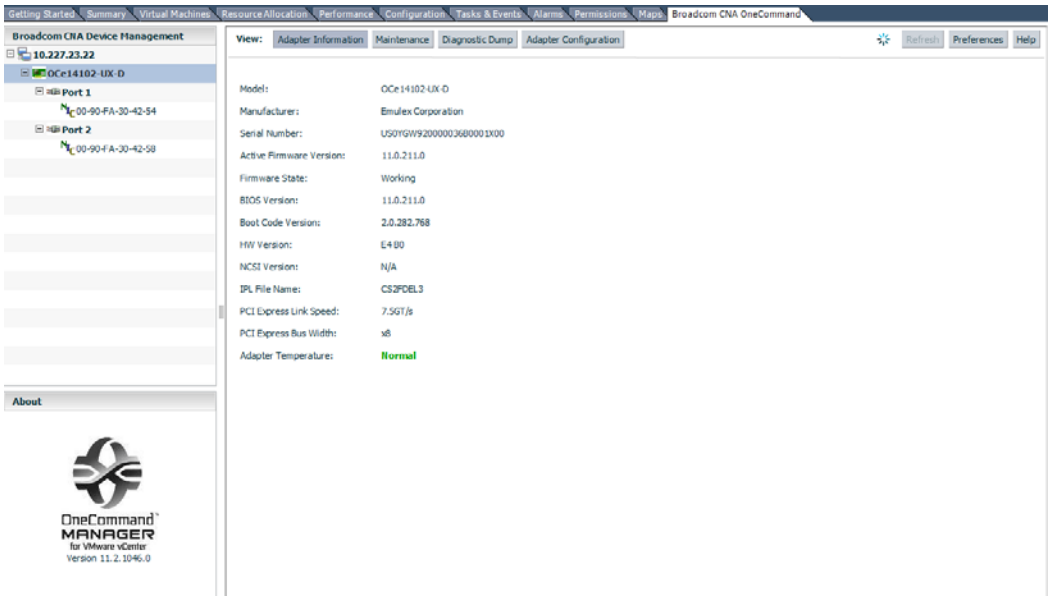
5.1.1 Viewing OneConnect Adapter Information

The **Adapter Information** tab contains general attributes associated with the selected CNA.

To view OneConnect adapter information, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select an adapter. The **Adapter Information** tab is displayed (Figure 22).

Figure 22 Adapter Information Tab



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

- Model – The model of the selected adapter.
- Manufacturer – The manufacturer of the adapter.
- Serial Number – The serial number of the selected adapter.
- Active Firmware Version – The firmware version running on the adapter.
- Firmware State – The condition of the firmware.
- BIOS Version – The version of the BIOS in use.
- Boot Code Version – The version of the boot code 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 PCI bus speed 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 shows 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, inoperative 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, inoperative fans, and air conditioning issues that cause high ambient air temperatures.

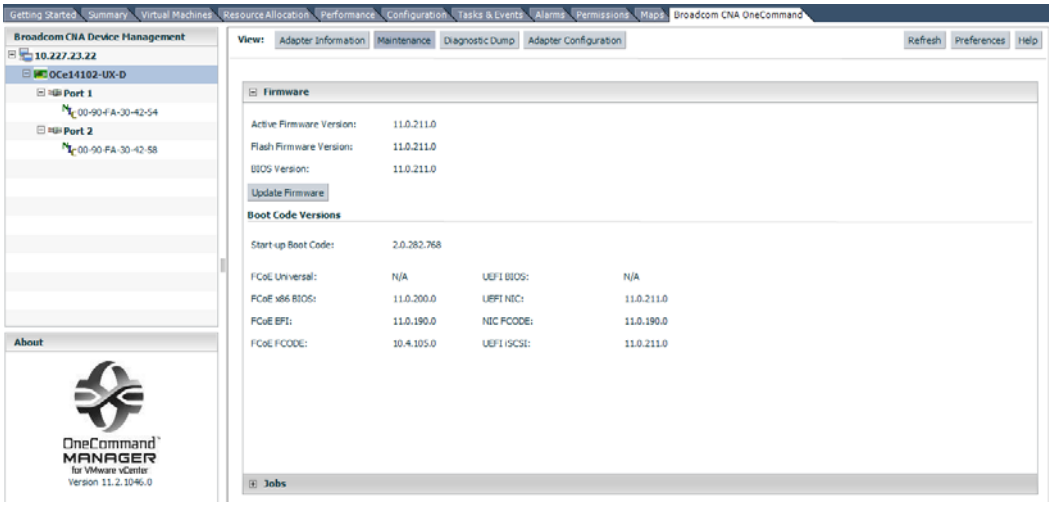
  After a system overheating issue is resolved and the adapter is cooled down, reboot the system or, if the system supports hot swapping, cycle the power of the adapter slot.
- Personality area (OCe11000-series adapters only) – See Section 5.2.2, Changing Personalities for instructions on changing personalities.
- Channel Management area (OCe11100–series adapters only) – For instructions on configuring channel management for OCe11000-series adapters, see Section 5.2.4, Managing Port Channels Using Channel Management. For OCe14000-series adapters, see Section 5.3.1, Configuring OCe14000–Series Adapters.
- License Features area (OCe11000-series adapters only) – For instructions on installing licenses, see Section 5.2.1, Viewing Licenses.

5.1.2 Viewing Firmware Information

To view firmware information for an adapter, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the adapter for which you want to view firmware information.
3. Select the **Maintenance** tab. The **Maintenance** tab enables you to view and update firmware on an adapter in a specific host (Figure 23).

Figure 23 Maintenance Tab



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

- Active Firmware Version – The firmware version currently being used by the adapter.
- Flash Firmware Version – The flash firmware version flashed on the adapter.
- BIOS Version – The version of the BIOS currently being used by the adapter.
- Boot Code Versions – The adapter's boot code versions currently being used by the adapter.

NOTE The Start-up Boot Code is the boot code that boots the adapter. It has no relation to the FC or PXE boot code versions.

For instructions on updating firmware, see Section 7.1, Updating Firmware on a Single Adapter.

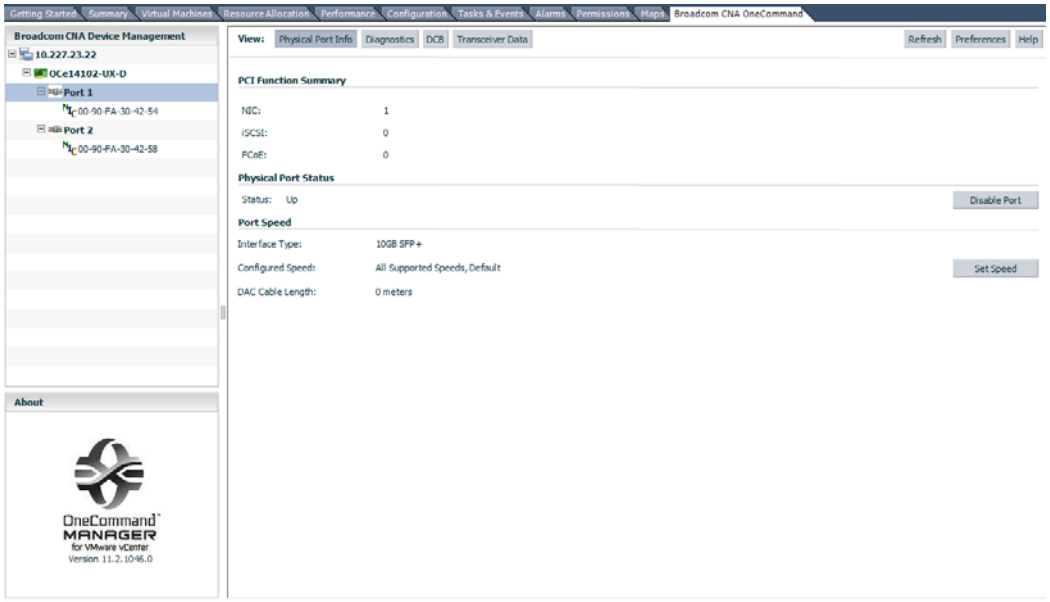
5.1.3 Viewing Physical Ports

The **Physical Port Info** tab shows general information about the PCI functions under the selected physical port and the current physical port status.

To view physical port information, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the physical port for which you want to view information.
3. Select the **Physical Port Info** tab (Figure 24).

Figure 24 Physical Port Info Tab



The following **Physical Port Info** tab fields are displayed:

- PCI Function Summary area:
  - NIC – The number of NIC functions running on the adapter.
  - iSCSI – The number of iSCSI functions running on the adapter.
  - FCoE – The number of FCoE functions running on the adapter.
- Physical Port Status area:
  - Status:
    - Down – Not operational
    - Up – Operational
    - n/a – Not available
- Port Speed area:
  - Interface Type – The module type and speed.
  - Configured Speed – The configured port mode (default, force or auto-negotiate) and speed.
  - DAC Cable Length – The DAC cable length in meters.
- Line Rate Negotiation area (10GBASE-T adapters only):
  - Cable Length – The cable length in meters.
  - Auto Negotiation – Indicates whether speed auto-negotiation is supported on the local port.
  - Advertised Line Rate – The default advertised link speed.
  - Link Partner Auto Negotiation – Indicates whether speed auto-negotiation is supported on the remote port.
  - Auto Negotiation Status – Indicates whether negotiation is completed or is in progress.

5.1.3.1 Enabling or Disabling a Physical Port

To change the physical port status, on the **Physical Port Info** tab, click **Enable Port/Disable Port.** This button toggles between Disable Port and Enable Port, depending on the current status of the physical port.

5.1.3.2 Changing a Physical Port's Speed

The Change Port Speed dialog allows you to configure port mode, speed and DAC cable length.

To configure port speed for a physical port, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the physical port whose port speed you want to configure.
3. Select the **Physical Port Info** tab (Figure 24).
4. Click **Set Speed**. The Change Port Speed dialog appears (Figure 25).

Figure 25 Change Port Speed Dialog



The following Change Port Speed dialog fields are displayed:

- Mode:
  - Default – Speed and DAC Cable Length fields are not active.
  - Force
  - Auto-Negotiate
- Speed – This field is active if the mode is set to Force or Auto-Negotiate.
  - If the mode is set to Force, you can set the speed to one of the following, depending on the interface type:
    - 10 Mb/s
    - 100 Mb/s
    - 1 Gb/s
    - 10 Gb/s
  - If the mode is set to Auto-Negotiate, you can set the speed setting to a single or multiple speed value, depending on the interface type.

    Single speed values:
    - 10 Mb/s
    - 100 Mb/s
    - 1 Gb/s
    - 10 Gb/s

    Multiple speed values combinations:
    - 10 Gb/s and 1 Gb/s
    - 1 Gb/s and 100 Mb/s
    - 1 Gb/s, 100 Mb/s, and 10 Mb/s
    - 100 Mb and 10 Mb/s
- DAC Cable Length – This field is required only if Force mode is selected and the speed is set to 10 Gb/s. Select from 1 to 10 meters. This field is not active in all other cases.

5. Make any applicable mode, speed, and DAC cable length changes.
6. Click **OK**.

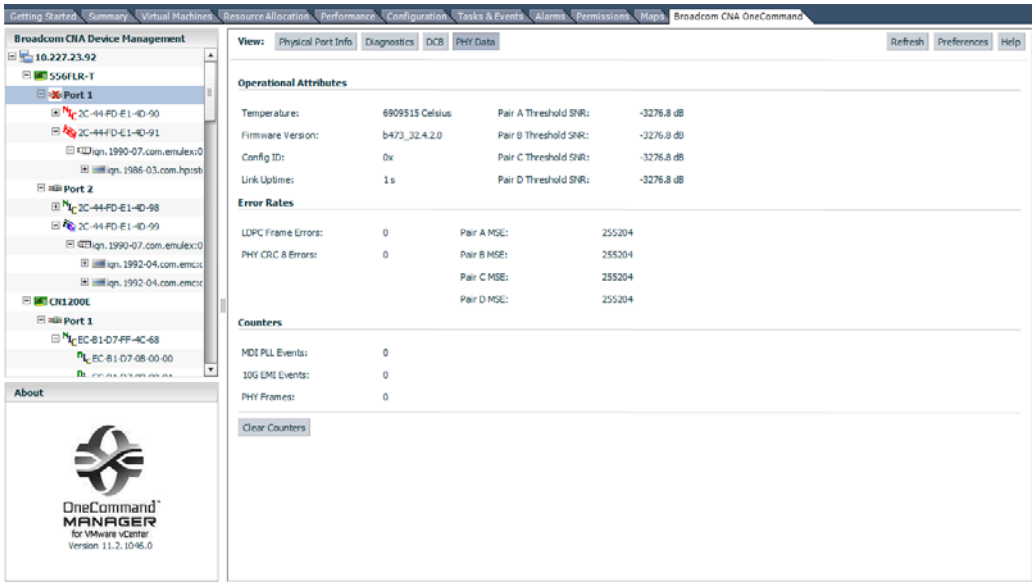
5.1.4 Viewing PHY Data for 10GBASE–T Adapters

The **PHY Data** tab shows PHY data information for 10GBASE-T–series adapters.

To view physical data information, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the physical port for which you want to view physical data information.
3. Select the **PHY Data** tab (Figure 26).

Figure 26 PHY Data Tab



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

- Operational Attributes area:
  - Temperature – The internally measured adapter temperature.
  - Firmware – The version of Emulex firmware currently active on the adapter port.
  - Config ID – The CRC configuration details.
  - Link Uptime – The time in seconds that the link has been up. Pair A/B/C/D SNR Margin – The port's average SNR margin for twisted pairs A, B, C, and D.
- Error Rates area:
  - LDPC Frame Errors – The number of LDPC frame errors.
  - PHY CRC8 Errors – The number of physical CRC errors.
  - Pair A/B/C/D MSE – The port's average MSE relative to the transmitted code words for twisted pairs A, B, C, and D.
- Counters area:
  - MDI PLL Events – The number of MDI PLL events.
  - 10G EMI Events – The 10G EMI Event counter tracks the number of single-tone interference detected by the 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 physical 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.

To clear the 10GBASE-T counters (OCe14000-series 10GBASE-T adapters only), click **Clear Counters**.

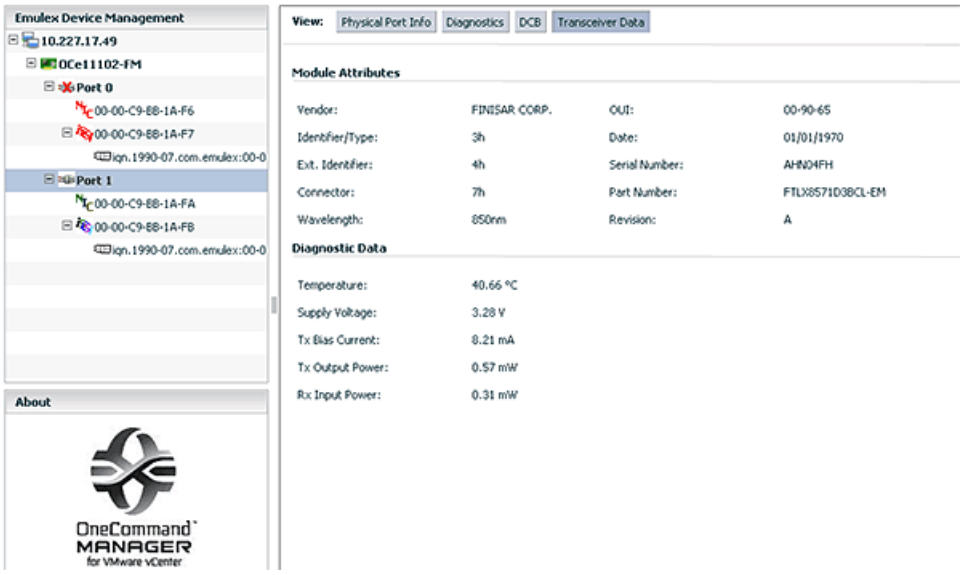
5.1.5 Viewing Port Transceiver Information

The **Transceiver Data** tab enables you to view transceiver information such as vendor name, serial number, part number, and diagnostic data. If the adapter/transceiver does not support some or all of the transceiver data, the fields show N/A.

To view transceiver information, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the physical port whose transceiver information you want to view.
3. Select the **Transceiver Data** tab (Figure 27).

Figure 27 Transceiver Data Tab



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

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

5.2 OCe11100-Series Adapters

5.2.1 Viewing Licenses

To view available licenses for an OCe11100-series adapter, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select an adapter whose licenses you want to view.
3. Select the **Adapter Information** tab (Figure 22).

   License information is displayed in the License Features area. The table shows the available licenses and whether they are licensed or not.

NOTE An empty feature list means the adapter has no licensable capabilities.

Installing Licenses

OneCommand CNA Manager for VMware vCenter allows you to view available licenses and install licenses to enable capabilities such as FCoE or iSCSI personalities on OCe11100-series adapters.

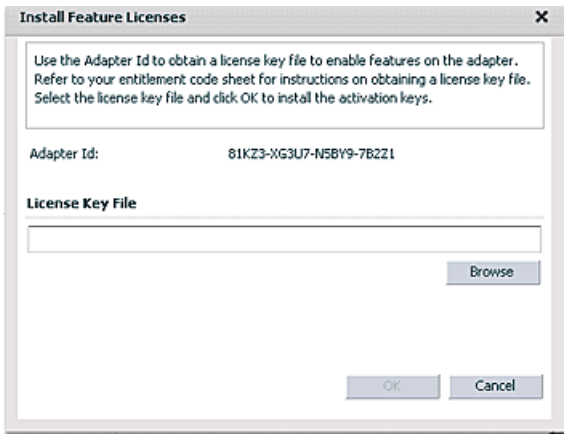
NOTES

- When licenses or licensable features are not available for the adapter, license information is not displayed.
- Contact Broadcom for the required license files.

To install licenses for an OCe11100-series adapter, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the adapter for which you want to install licenses.
3. Select the **Adapter Information** tab (Figure 22).
4. Click **Install Feature Licenses**. The Install Feature Licenses dialog shows the Adapter ID (Figure 28).

Figure 28 Install Feature Licenses Dialog



5. Following the instructions you received with the Entitlement Code, go to the License website and enter the Adapter ID and Entitlement Code.

6. When the Adapter ID and Entitlement Code are successfully validated, download a License Key File containing one or more activation keys.
7. In the Install Feature Licenses dialog, click **Browse**. The file selection dialog appears. Select the file you want to use and click **OK**.
8. A dialog confirms that you want to install the licenses. Click **Yes**.
9. A dialog notifies you that the installation was successful or why it failed. Click **OK**.

5.2.2 Changing Personalities

OneCommand CNA Manager for VMware vCenter allows you to change the personality running on an OCe11100-series adapter. There are three personalities: NIC-Only, FCoE, and iSCSI.

In some cases, the adapters come pre-configured to support multiple personalities. In other cases an enablement license must be installed before the adapter can support multiple personalities. Contact Broadcom for the required license files.

To change the personality of an adapter, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select an adapter whose personality you want to change.
3. Select the **Adapter Information** tab (Figure 22). From the Personality area of the **Adapter Information** tab, select a personality type and click **Apply**.
    - FCoE – Assigns the FCoE personality.
    - iSCSI – Assigns the iSCSI personality.
    - NIC Only – Assigns the NIC-only personality.

    NOTE If the adapter does not support personalities, personality controls are not displayed. If the adapter does not support a particular personality type that control is disabled.
4. Reboot the host for the personality change to take effect.

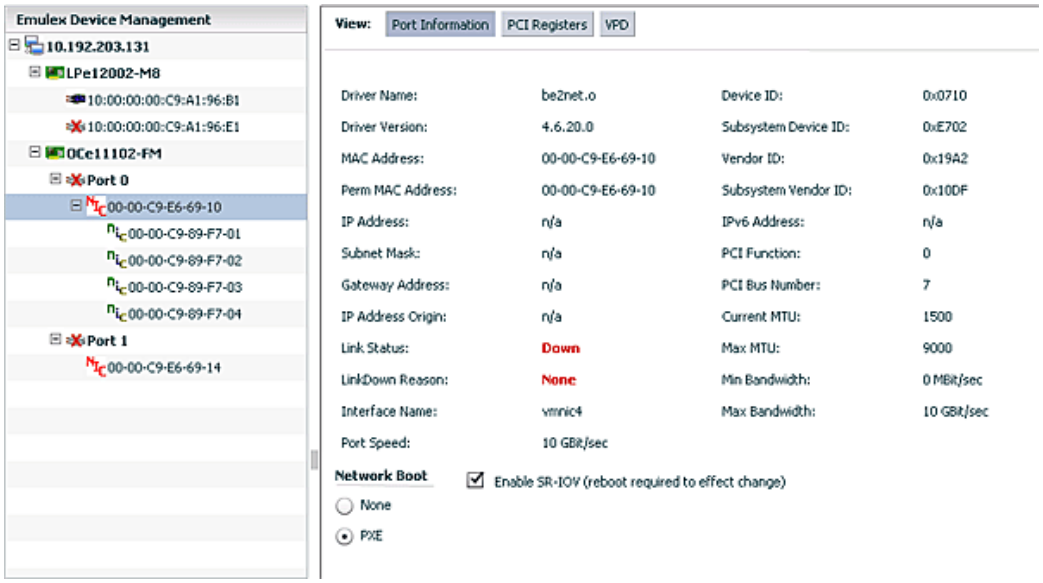
5.2.3 Configuring SR-IOV

If SR–IOV is enabled, multiple VFs can be created on a NIC PF for an adapter port. These VFs are assigned by the ESXi host to the 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 it is available for network I/O. The discovery-tree displays the VFs running on the discovered PFs. It also shows the selected VFs MAC Address, VLAN ID, Link Status, and Transmit Rate.

To enable or disable SR-IOV, do the following, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select a NIC port function on an adapter in the discovery-tree.
3. Select the **Port Information** tab (Figure 29).
4. Click **Enable SR-IOV** or clear the **Enable SR-IOV** checkbox.

Figure 29 Enabling SR-IOV



NOTES

- Channel management and SR-IOV cannot be used together; they are mutually exclusive. If you set channel management, then you must disable SR-IOV. If you set SR-IOV, then you must disable channel management.
- To use SR-IOV, you must enable SR-IOV in the OneCommand CNA Manager for VMware vCenter application and the system BIOS. Refer to the *Emulex Drivers for VMware ESXi User Guide for One Connect Adapters* for information on enabling SR-IOV in the system BIOS.
- If you disable SR-IOV from the **Port Information** tab by clearing the Enable SR-IOV checkbox, the operating system or port stops supporting SR-IOV after reboot.
- A system reboot is required to change the SR-IOV state.

5.2.4 Managing Port Channels Using Channel Management

Channel management allows each supported OCe11100-series adapter's physical port to 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 capabilities such as enabling or disabling send and receive traffic, minimum and maximum transmit rates, and VLAN IDs for untagged packets.

Ports are managed using the **Channel Management** tab. The **Channel Management** tab is displayed when an adapter port is selected in the discovery-tree and Channel Management is enabled on the **Adapter Information** tab. See Section 5.2.4.2, Configuring Channel Management.

The **Channel Management** tab for UMC or SIMode shows the function number, type of protocol, MAC address, and more for each channel. You can also configure and enable or disable channels, and assign LPVIDs.

Refer to the *Emulex Universal Multi-Channel Reference Guide* for more information about channel management.

NOTES

- You must reboot the host for channel management to activate.

- When the channel management mode in the **Adapter Information** tab is set to None, the **Channel Mgmt** tab is not displayed.
- Channel management and SR-IOV cannot be used together. If you set channel management, then you must disable SR-IOV. If you set SR-IOV, then you must disable UMC.
- Properties for all channels on a port can be viewed and modified at the time channel management is enabled on the **Adapter Information** tab, before the required reboot to activate channel management on the adapter.
- SIMode or UMC mode is required to dynamically control CNA ports using channel management.

5.2.4.1 Enabling and Disabling Channel Management

To enable channel management, select a management type from the Management Type After Reboot list in the Channel Management area when in the Adapter Information view (Figure 22).

To disable channel management, select **None** from the Management Type After Reboot list.

5.2.4.2 Configuring Channel Management

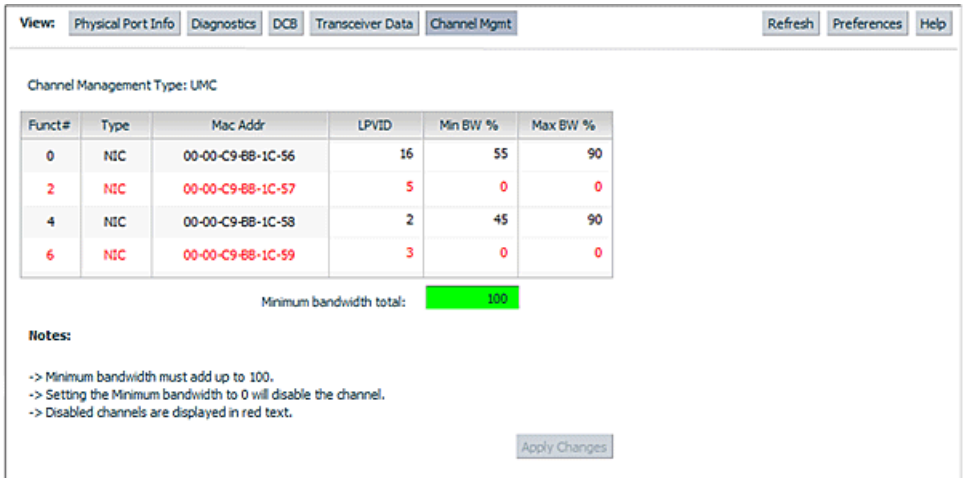
When you select an OCe11102 adapter port with channel management enabled, the **Channel Mgmt** tab is available.

NOTE The **Channel Mgmt** tab shows the settings for the configured channel management mode after reboot. This may not match the mode that is currently running on the adapter if you have changed the mode in the **Adapter Information** tab, but have not rebooted.

To manage the ports for an OCe11102 adapter, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the adapter port for which you want to manage channels.
3. Select the **Channel Mgmt** tab. The **Channel Mgmt** tab is displayed (Figure 30).

Figure 30 Channel Mgmt Tab for the OCe11102 Adapter



4. Make appropriate changes to the LPVID, Min BW%, and Max BW% values in accordance with the guidelines that follow.

5. Click **Apply**. A reboot is not required.

The following **Channel Mgmt** tab fields are displayed:

- Channel Management Type – Channel management protocol type in use.
- Funct# – Function number used to identify a specific channel.
- Type – Protocol personality of a specific channel.
- MAC Address – MAC address assigned to a specific channel.
- LPVID – Logical port VLAN ID.
- Minimum and maximum bandwidths (in percent) – Percentage of available link bandwidth allocated to a particular function.
- Minimum bandwidth total – Sum of the minimum bandwidths, which must total 100% in order to apply any changes made in the **Channel Mgmt** tab.
- If you change any properties, you must click **Apply Changes** to save them. A reboot is not required for changes to take effect.

5.2.4.3 Disabling a Specific Channel

To disable a channel, click the Minimum Bandwidth field for the specific channel and enter a zero value.

To enable a channel, click the Minimum Bandwidth field for the specific channel and enter a non-zero value.

For more information on setting bandwidths, see Section 5.2.4.5, Assigning Bandwidth for a Specific Channel.

5.2.4.4 Assigning an LPVID for a Specific Channel

LPVID values range between 2 and 4094. Any enabled channel must be assigned an LPVID. LPVIDs must be unique across all channels on a port.

NOTE LPVID does not apply to iSCSI and FCoE channel types. LPVID is set to zero for these protocols.

To assign an LPVID value, perform the following steps:

1. Click the LPVID field for the channel you want to change.
2. Enter a valid LPVID value.
3. Click **Apply Changes**.

NOTE If you click **Apply Changes** and the LPVID values are not between 2 and 4094 or if they are not unique, an error message is displayed indicating the issue and that the changes were not saved.

5.2.4.5 Assigning Bandwidth for a Specific Channel

Minimum bandwidth values of the enabled channels can be set between 1% and 100%, but must all add up to exactly 100%. The total of these bandwidths is displayed and updated each time the minimum bandwidth field is changed. If the total is 100%, the background in the Minimum bandwidth total: is green. If not, the background is red to indicate changes must be made.

The maximum bandwidth of a channel must be greater than or equal to the minimum bandwidth for that channel.

To assign a bandwidth value, perform the following steps:

1. Click the appropriate Bandwidth field for the channel you want to change.
2. Enter a valid bandwidth value.
3. Click **Apply Changes**.

NOTE If you click **Apply Changes** and the total minimum bandwidths do not add up to 100% or one of the maximum bandwidths is not valid, an error message is displayed indicating the issue and that the changes were not saved.

5.2.4.6 Viewing the Channel Management Tab for UFP

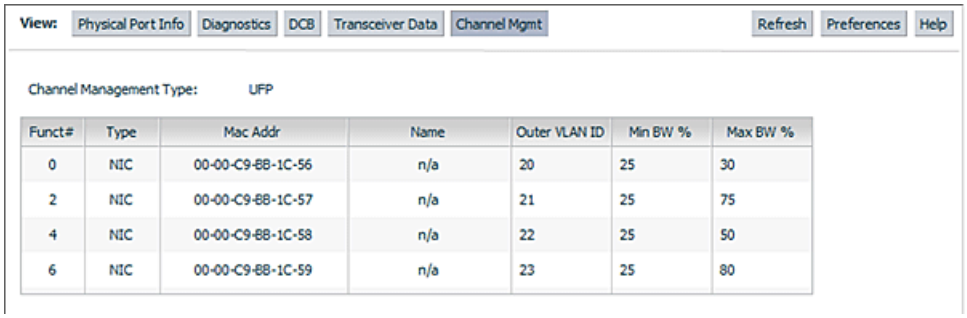
For adapters that support UFP, the UFP properties are displayed in the Channel Management (**Channel Mgmt**) tab when UFP is enabled. The **Port Channel Management** tab for UFP is read-only with specific attributes listed in the table as shown in Figure 31.

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 http://www.broadcom.com.

To view UFP channel management, perform the following steps:

1. From the discovery-tree, select the UFP adapter port whose channel information you want to view.
2. Select the **Channel Management** tab (Figure 31).

Figure 31 Channel Management Tab in UFP Mode



The following Channel Management fields for UFP are displayed:

- Channel Management Type – The channel management protocol type in use.
- Function Number – The function number assigned by the system to the channel.
- Type – The protocol type in use.
- MAC address – The MAC address assigned to the channel.
- Name – The name assigned to the UFP 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.
- Min. Bandwidth – The minimum bandwidth at which the port is guaranteed to run.
- Max. Bandwidth – The maximum bandwidth at which the port is guaranteed to run.

5.2.4.7 Viewing the Channel Management Tab for vNIC1

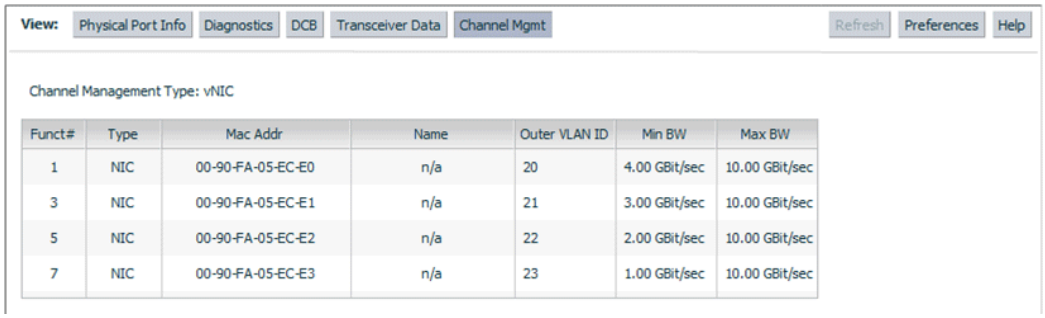
For adapters that support vNIC1, the vNIC properties are displayed in the **Channel Management** tab when vNIC is enabled. The **Port Channel Management** tab for vNIC is read-only with the vNIC specific attributes listed in the table as shown in Figure 32.

NOTE vNIC is only supported on certain virtual fabric adapters. For specific information as to whether it is supported on a specific adapter, see the release notes that are available on http://www.broadcom.com.

To view vNIC channel management, perform the following steps:

1. From the discovery-tree, select the vNIC adapter port whose channel information you want to view.
2. Select the **Channel Management** tab.

Figure 32 Channel Management Tab in vNIC1 Mode



The following Channel Management fields for vNIC1 are displayed:

- Channel Management Type – Channel management protocol type in use.
- Funct# – Function number used to identify a specific channel.
- Type – Protocol personality of a specific channel.
- MAC Addr – MAC address assigned to a specific channel.
- Name – Name assigned to the vNIC by an administrator during switch configuration.
- Outer VLAN ID – VLAN identifier used between the NIC port and the switch. The switch maps this value into the VLAN ID used on the network.
- Mn BW – Minimum bandwidth at which the port is guaranteed to run.
- Max BW – Maximum bandwidth at which the port is guaranteed to run.

5.3 OCe14000–Series Adapters

This section pertains to viewing and configuring OCe14000-series adapter and port information.

5.3.1 Configuring OCe14000–Series Adapters

When you select an OCe14000-series adapter from the discovery-tree, the **Adapter Configuration** tab is displayed. The type of information that appears depends on the what protocols and capabilities are available on the adapter. (The IPL defines the protocols and capabilities.)

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. Custom configuration is available if the adapter is capable of mixed or concurrent mode storage.
- Multichannel – View the channel properties and change the configurable properties (based on the multichannel type).

For NPar, the adapter is either running with NPar disabled (the adapter is running as a single NIC function on each port), or with NPar enabled, which can be similar to a multichannel configuration or a storage channel configuration.

Figure 33 OCe14000-series Adapter Configuration Tab without Link Reconfiguration Support

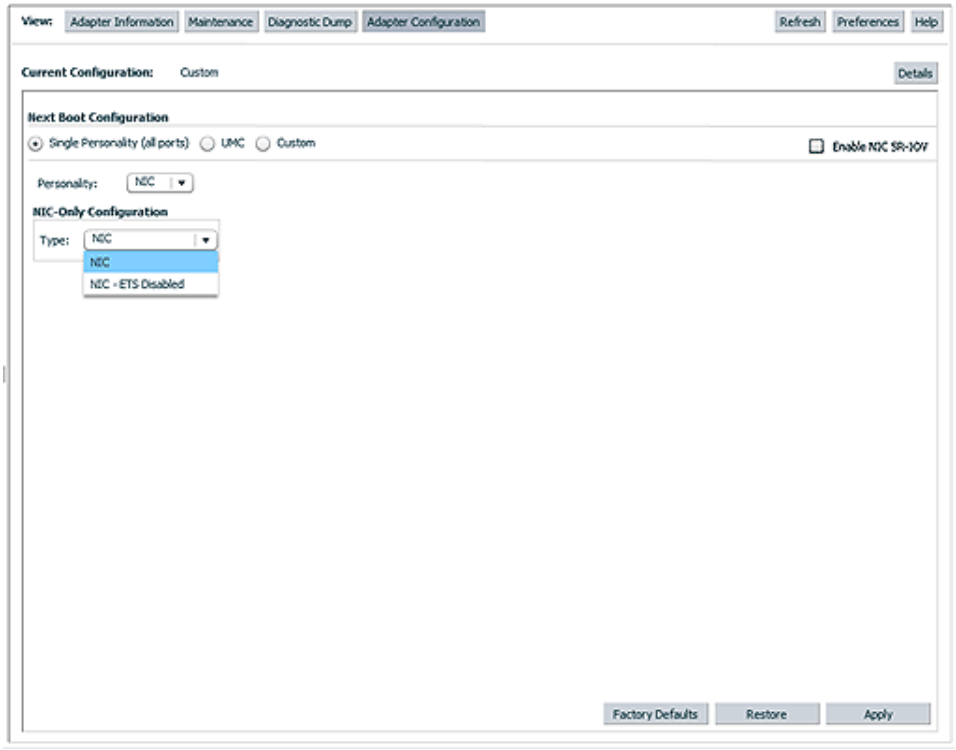
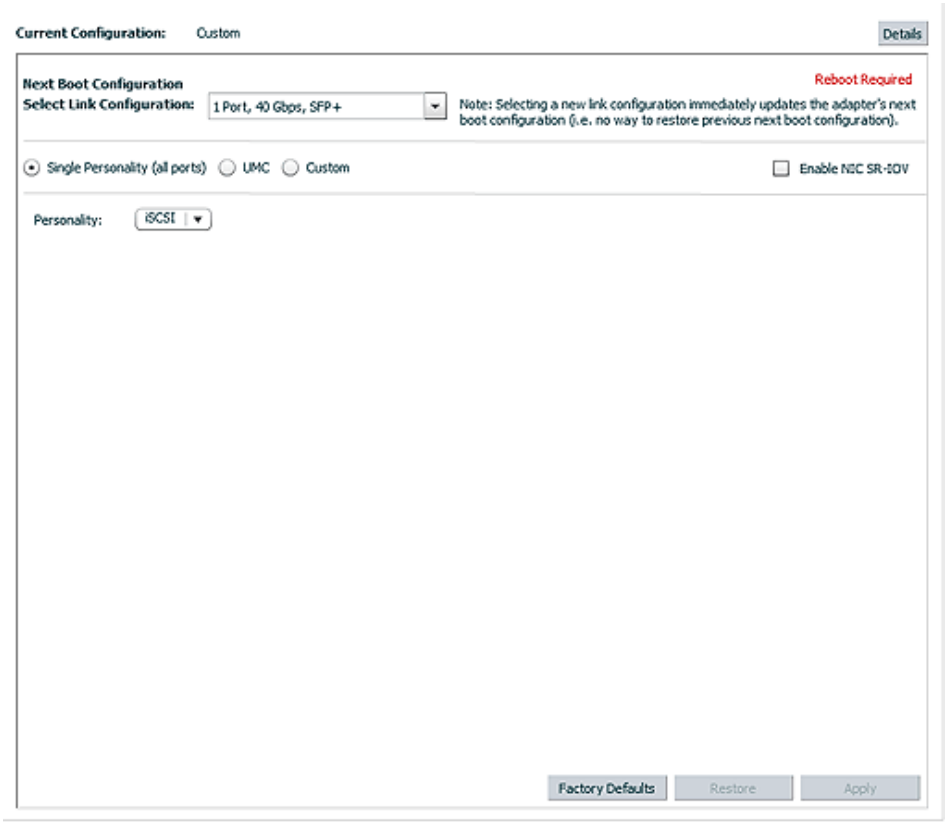


Figure 34 Adapter Configuration Tab with Link Reconfiguration Support



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

- Current Configuration – The protocol running on the adapter.
- Next Boot Configuration
  - Shows the Select Link Configuration drop-down menu, if the adapter supports link reconfiguration (Figure 34).
  - Shows 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.
- The display in the lower area of the tab depends upon the selected view of the next boot configuration. For non-NPar configurations, depending upon configuration options available on the adapter, you can configure the adapter for a single personality (for example, FCoE), multichannel, or a custom configuration (for mix and concurrent modes).

To apply any changes to the next boot configuration of the adapter, click **Apply**.

To restore any changes back to the currently saved settings on the adapter, click **Restore**.

NOTE If the adapter supports link reconfiguration, **Restore** is inactive after link reconfiguration is applied; the previous link configuration cannot be restored.

To restore the adapter to its factory default profile and settings, click **Factory Defaults**. A reboot is required. If the adapter supports link reconfiguration, clicking **Factory Defaults** sets the adapter to its default link configuration.

To view more information about the configuration currently running on the adapter (Figure 35), click **Details**. If the adapter supports link reconfiguration, the current link configuration is displayed above the Current Configuration field (Figure 36).

Figure 35 OCe14000-series Adapter Current Configuration Details Without Link Reconfiguration Support

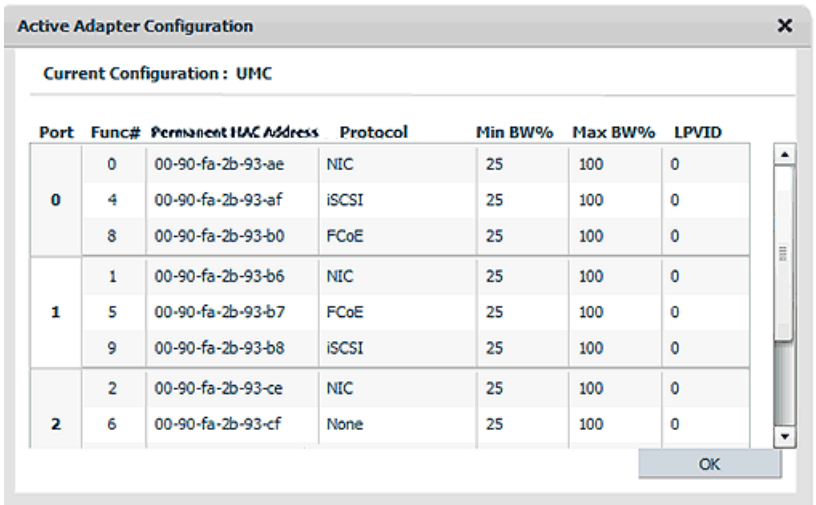
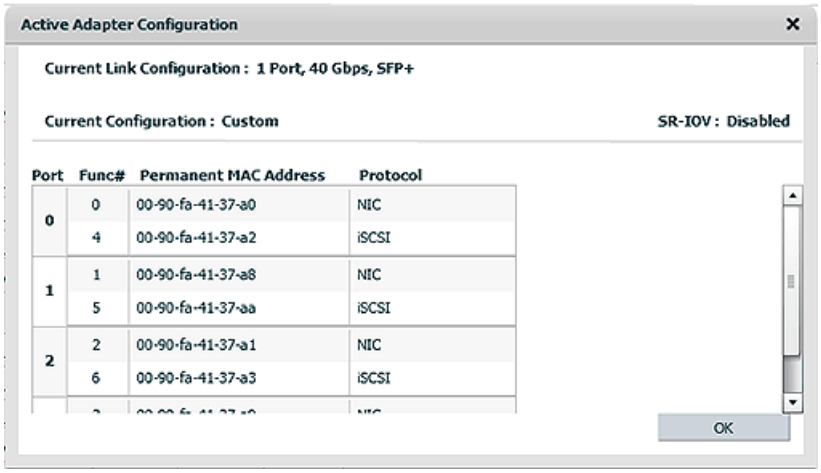


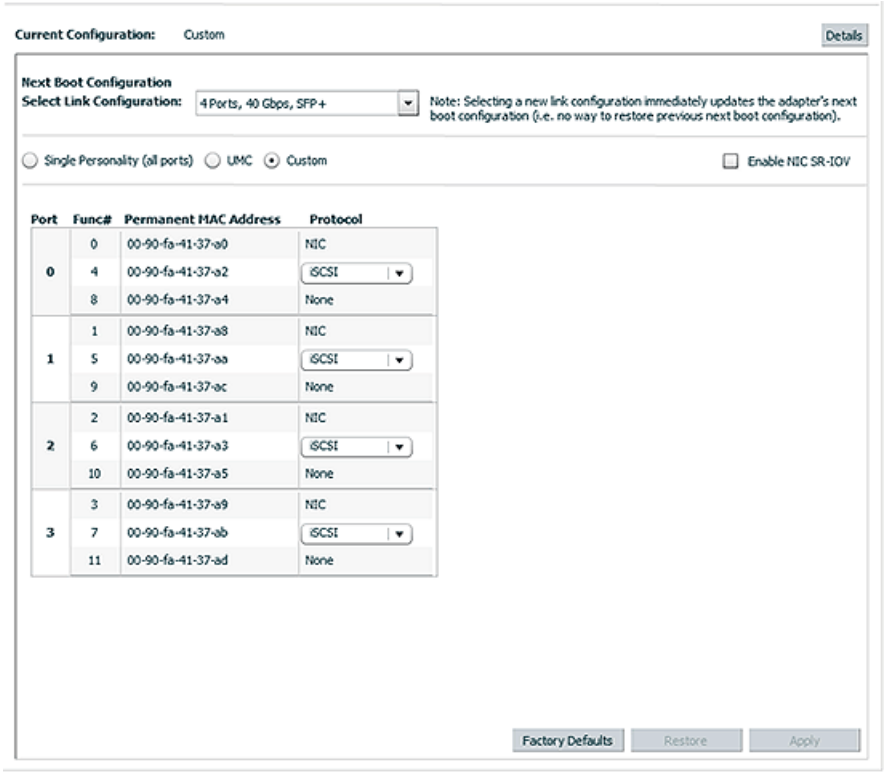
Figure 36 OCe14000-series Adapter Current Configuration Details With Link Reconfiguration Support



5.3.1.1 Link Reconfiguration

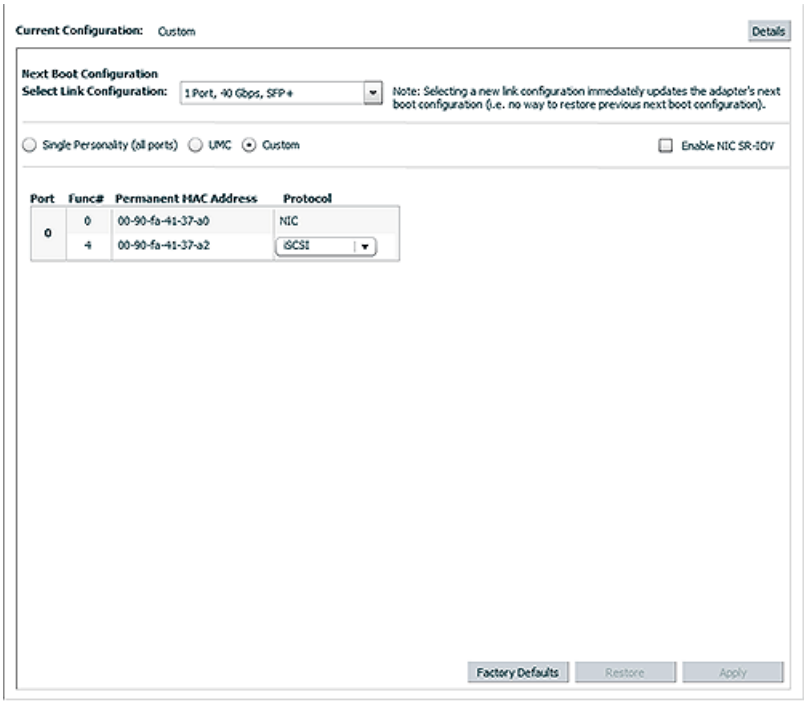
If the adapter supports link configuration, you can adjust the number of ports and port speeds on the adapter by selecting a specific link reconfiguration setting.

Figure 37 OCe14000-series Adapter Configuration Before Link Reconfiguration



After changing the link configuration, the new port configuration is displayed.

Figure 38 OCe14000-series Adapter Configuration After Link Reconfiguration



You can change the functions running on that port before performing a reboot to activate the new link and port configuration.

5.3.1.2 Enabling or Disabling SR-IOV

The **Adapter Configuration** tab (Figure 33) allows you to enable or disable SR-IOV on an adapter. When SR–IOV is enabled, multiple VFs can be created on a NIC PF for an adapter port. These VFs are assigned by the ESXi host to the 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 shows the VFs running on the discovered PFs. It also shows the selected VFs MAC Address, VLAN ID, Link Status, and Transmit Rate.

NOTE You can enable or disable SR-IOV for operating systems that do not support SR-IOV. Enabling or disabling SR-IOV for operating systems that do not support SR-IOV may be applicable in some cases. Examples include:

- An adapter is configured to support SR-IOV in an operating system that supports SR-IOV, and the adapter is moved to a system that doesn't support SR-IOV.
- The OneCommand CNA Manager for VMware vCenter Command Line Interface is used in an offline environment, such as WinPE and CentOS.

To enable or disable SR-IOV, perform the following steps:

1. Select **Host** view.
2. Select a NIC port function on an adapter in the discovery-tree.
3. Select the **Adapter Configuration** tab.
4. Select or clear the **Enable SR-IOV** checkbox (Figure 33).

5. Click **Apply**.

NOTES

- To use SR-IOV, you must enable SR-IOV in OneCommand CNA Manager for VMware vCenter and the system BIOS. Refer to the *Emulex Drivers for VMware ESXi User Guide for One Connect Adapters* for information on enabling SR-IOV in the system BIOS.
- If you disable SR-IOV from the **Adapter Configuration** tab by clearing the Enable SR-IOV checkbox, the operating system or port stops supporting SR-IOV after reboot.
- SR-IOV must be enabled at the system BIOS level.

5.3.1.3 Configuring a Single Personality

The single personality runs the same protocol on all ports of the adapter.

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

SR-IOV can be enabled or disabled for all adapter types that support SR-IOV, except multichannel configurations.

NOTE SR-IOV can be enabled on NPar adapters only when all ports are configured to run NIC functions.

To configure a single personality, perform the following steps:

1. From the discovery-tree, select the OCe14100-series adapter whose personality you want to configure.
2. Select the **Adapter Configuration** tab (Figure 33).
3. Check the **Single personality (all ports)** radio button.
4. Select the personality you want to apply from the Personality drop-down menu.
   - — If you choose a NIC configuration, there may be multiple NIC configurations available (for example; NIC or NIC- ETS Disabled). If the adapter supports multiple NIC configurations (profiles), you must select a NIC configuration from the NIC configuration list.
   - — If you choose the NIC- ETS Disabled configuration, SR-IOV is automatically enabled on each NIC function and allows up to 63 virtual functions. If you choose the NIC configuration, you must check the Enable NIC SR-IOV checkbox to enable SR-IOV.
5. Click **Apply**. A message appears notifying you of the profiles activation requirements. Not all selections require a reboot.

5.3.1.4 Setting Up a Custom Configuration

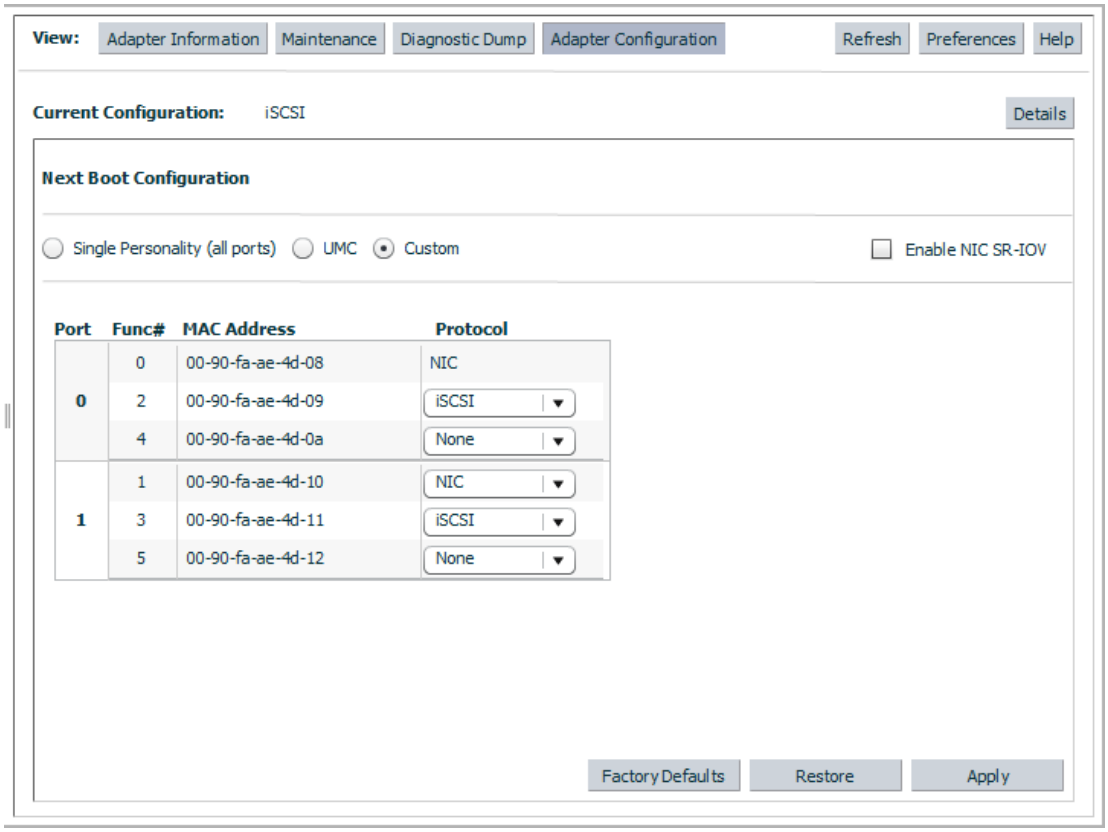
The Custom configuration view allows you to customize the protocols running on each port of an adapter. The Port Configuration table only shows the available ports (Figure 39).

NOTES

- Custom configurations are only available if the adapter supports mixed or concurrent mode storage.
- If you configure all ports the same, you essentially configure a single personality and the **Adapter Configuration** tab shows the **Single Personality** radio button automatically selected.

SR-IOV can be enabled or disabled for all adapter types that support SR-IOV, except multichannel configurations, and can be enabled on NPar adapters only when all ports are configured to run NIC functions.

Figure 39 OCe14100-series Adapter Configuration Tab, Custom Configuration



NOTE Enabling and configuring more than four functions per physical port requires that Custom mode is enabled on the adapter. The options to enable and configure more than four functions per physical port are not available unless Custom mode is selected.

If you enable Custom mode but do not create a Custom configuration, the personality might revert to NIC, iSCSI, or FCoE. If Custom mode is enabled, the following conditions apply:

- If you leave all functions set to NIC, the configured personality is automatically switched from Custom to NIC.
- If you leave the functions configured the same as they would be for iSCSI, the configured personality is automatically switched from Custom to iSCSI.
- If you leave the functions configured the same as they would be for FCoE, the configured personality is automatically switched from Custom to FCoE

To set up a custom configuration, perform the following steps:

1. From the discovery-tree, select the OCe14100-series adapter whose personality you want to configure.
2. Select the **Adapter Configuration** tab.
3. Check the **Custom** radio button.
4. Select the personality you want to apply from the Protocol drop-down menu for each port.

- If you choose a NIC configuration, there may be multiple NIC configurations available (for example, NIC or NIC-ETS Disabled). If the adapter supports multiple NIC configurations (profiles), you must select a NIC configuration from the NIC configuration list.
- If you choose the NIC- ETS Disabled configuration, SR-IOV is automatically enabled on each NIC function and allows up to 63 virtual functions. If you choose the NIC configuration, to enable SR-IOV, you must check the Enable NIC SR-IOV checkbox.

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

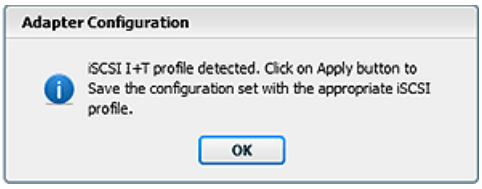
5.3.1.5 iSCSI Initiator + Target (I + T) Profile Detection in Single Personality or Custom Views

OneCommand CNA Manager for VMware vCenter does not allow you to select an iSCSI initiator + target profile on an OCe14100-series adapter. Other tools exist that allow you to run some OCe14100-series adapter models as iSCSI targets. If you take one of these adapters and re-purpose it in a server to be used as an iSCSI initiator, OneCommand CNA Manager for VMware vCenter detects that the iSCSI initiator + target profile is actively running on the OCe14100-series adapter.

To adjust the number of ports and ports speeds on the adapter, and save the configuration with the appropriate iSCSI profile, perform the following steps:

1. In the Broadcom CNA Device Management area, select the adapter running the iSCSI I +T profile.
2. Select the **Adapter Configuration** tab (Figure 33).
3. If the **Single Personality** or **Custom** radio button is selected, the following message is displayed.

Figure 40 iSCSI I + T Profile Message



4. Click **OK** on the message.
5. If the **Single Personality** radio button is selected, the new configuration is selected automatically. You must click **Apply** on the **Adapter Configuration** tab to save the configuration.
6. If the **Custom** radio button is selected, you must click **Apply** on the **Adapter Configuration** tab.
7. For both single and custom personalities, you must reboot the systems to activate the changes.

The adapter is reconfigured.

NOTE OneCommand CNA Manager for VMware vCenter does not discover and display the iSCSI functions implemented by this profile. These iSCSI functions are running with a target iSCSI driver that is not supported by OneCommand CNA Manager for VMware vCenter for iSCSI management.

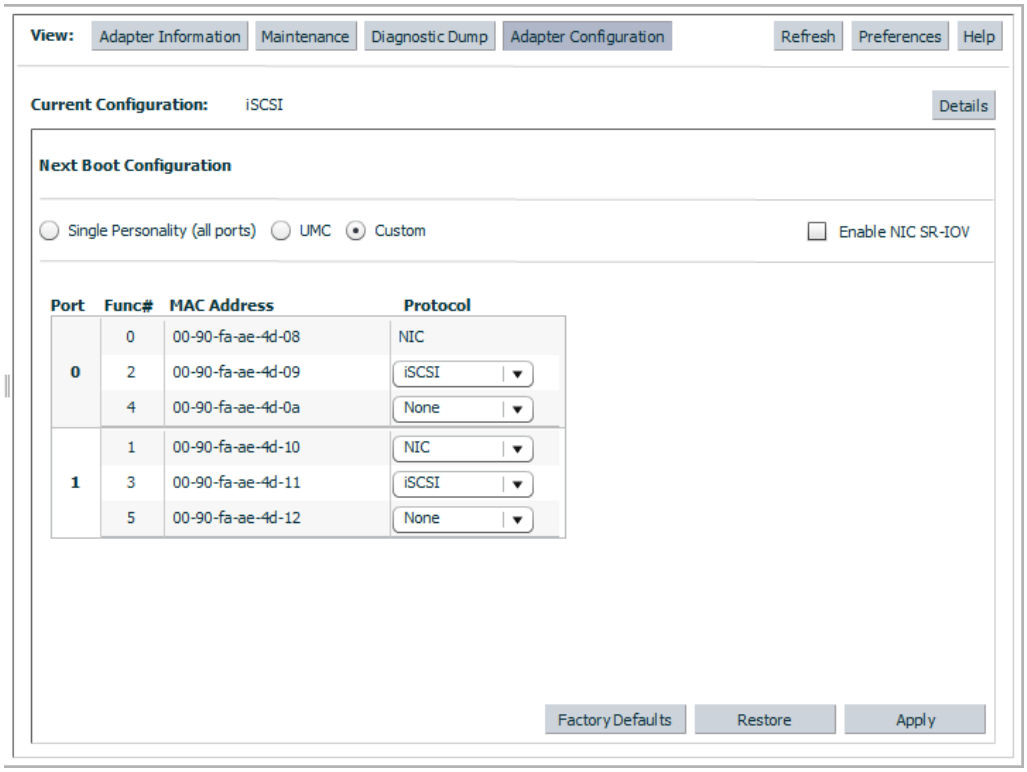
5.3.1.6 Setting Up a Mixed Mode Configuration

When mixed mode is available, you can configure up to two functions per port. The first function is NIC. The second function can be a storage protocol or None. When you click on the second function's drop-down menu, the available storage protocols for the function are displayed (Figure 41).

SR-IOV can be enabled or disabled for all adapter types that support SR-IOV, except multichannel configurations.

SR-IOV can be enabled on NPar adapters only when all ports are configured to run NIC functions.

Figure 41 Mixed Mode Protocol Drop-Down Menu



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

NOTES

- 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 drop-down menu.
- If NIC protocol is selected for all functions, a drop-down menu appears with available NIC-Only profiles. One of the profiles must be selected.

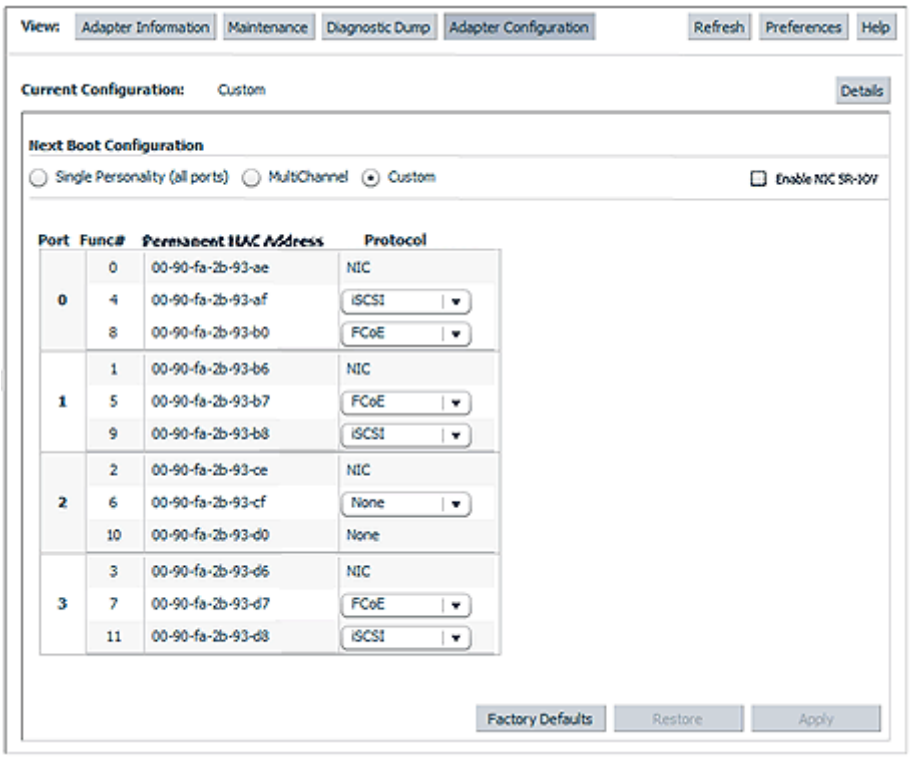
5.3.1.7 Setting Up a Concurrent Mode Configuration

When concurrent mode is available, you can configure up to three functions per port (Figure 42).

SR-IOV can be enabled or disabled for all adapter types that support SR-IOV, except multichannel configurations.

SR-IOV can be enabled on NPar adapters only when all ports are configured to run NIC functions.

Figure 42 Concurrent Storage Configuration View



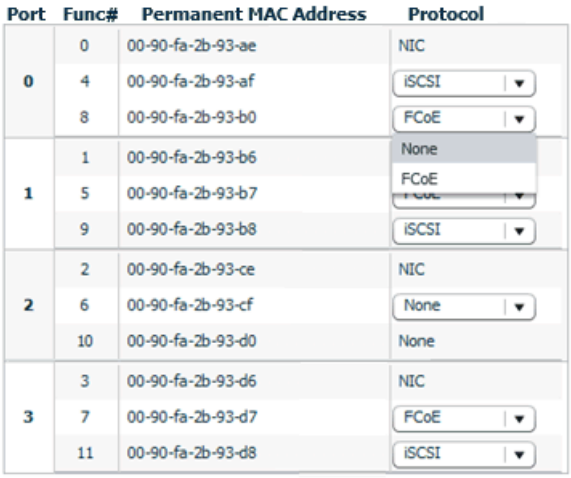
The drop-down menu for selecting a storage protocol to run on the second and third functions is the same as the drop-down menu for mixed mode storage.

Concurrent Storage Configuration is subject to the following rules:

- The first function must be NIC.
- If two storage functions are configured on a port, they must be different storage protocols (for example, concurrent mode).
- When you select the storage protocol for the second function, the choices in the drop-down menu for the third function are limited to the other storage protocol or None (Figure 43).

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 **Apply**, the tab switches to the single personality view and the selected personality is shown in the personality drop-down menu.

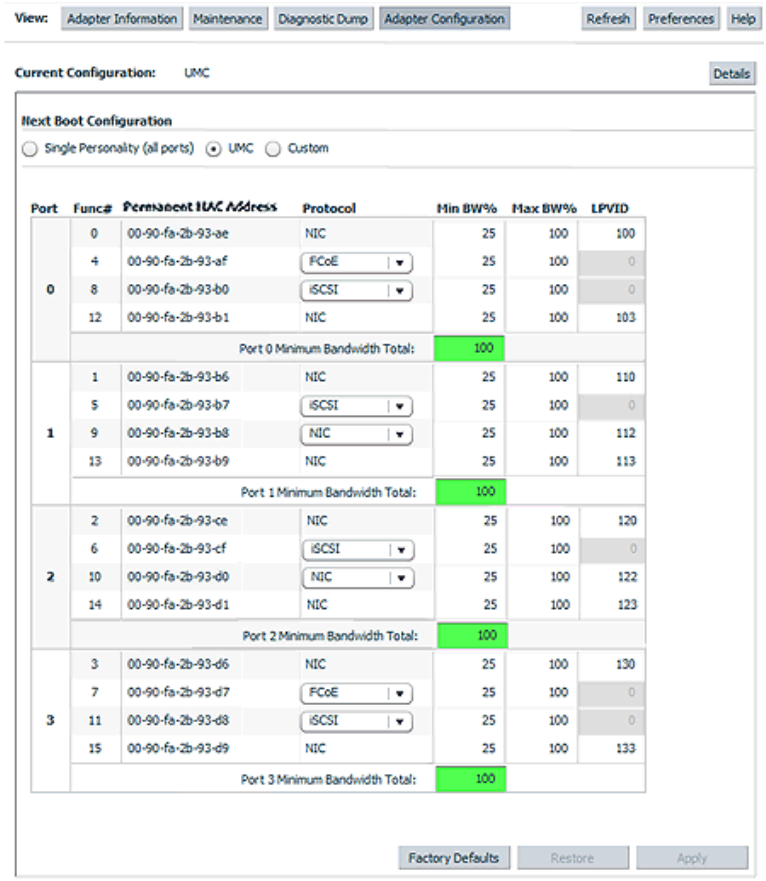
Figure 43 Concurrent Storage Configuration Choices for the Third Function



5.3.2 Setting Up a UMC Configuration

When UMC is available and the **UMC** radio button is selected, the **Adapter Configuration** tab is similar to Figure 44.

Figure 44 UMC View (Four Ports)



UMC protocol assignments are subject to the following rules:

- The first channel can only be assigned NIC protocols. The first channel cannot be assigned None (disabled).
- The second and third channels can be assigned the following:
  - NIC
  - None (disabled)
  - Storage protocol (third channel only when 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 (if available) can only be assigned NIC or None (disabled).
- A channel's logical link is set to down by setting both the minimum and maximum bandwidths to 0.
- The total minimum bandwidth for all enabled channels on a port must add up to 100, except when all the maximum bandwidths are 0 (for example, setting a 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 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.

5.3.2.1 Default UMC Settings

Most hosts today do not support AR; therefore, the default UMC settings are designed to prevent enabling functions that require ARI (when possible). This prevents the possibility of enabling channels that are not seen on the PCI bus when the system boots. There is one case where functions requiring ARI are enabled when UMC is enabled. On a four port adapter, when UMC is disabled, if the third function on each port is already enabled when UMC is enabled, that function continues to be enabled. In this case, it is assumed that you previously enabled the third function because ARI was available.

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

For each port:

- All enabled functions (that is, functions not set to None) continue to be enabled, running the same protocol before UMC was enabled even if its function number is greater than seven (only possible on 4 four-port adapters).
- Any channel with a function number less than eight is enabled with NIC.
- Any channel with a function number greater than seven is set to None (disabled).
- The minimum bandwidths are assigned by dividing up 100% as evenly as possible between the enabled channels.
- The maximum bandwidths for the enabled channels are set to 100.
- Both the minimum and maximum bandwidth for the channels assigned None are set to 0.
- The LPVIDs are set to 0. You must assign any NIC channel a valid LPVID.

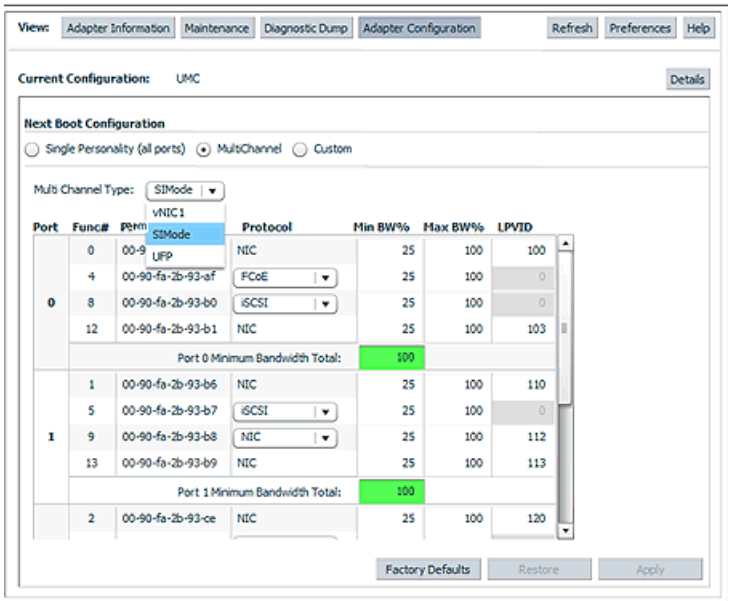
NOTE These rules create the default settings for the channels. You can update these rules before saving them.

5.3.3 Setting Up an Lenovo System X Multichannel Configuration

NOTE Lenovo System X multichannel currently only allows four channels per port even in one and two port configurations.

For Lenovo System X multichannel configuration, there are three different multichannel types available: vNIC, SIMode, and UFP. When you check the **MultiChannel** radio button, the **Adapter Configuration** tab is displayed (Figure 45).

Figure 45 Multichannel View (Showing the Multi Channel Type Drop-Down Menu)

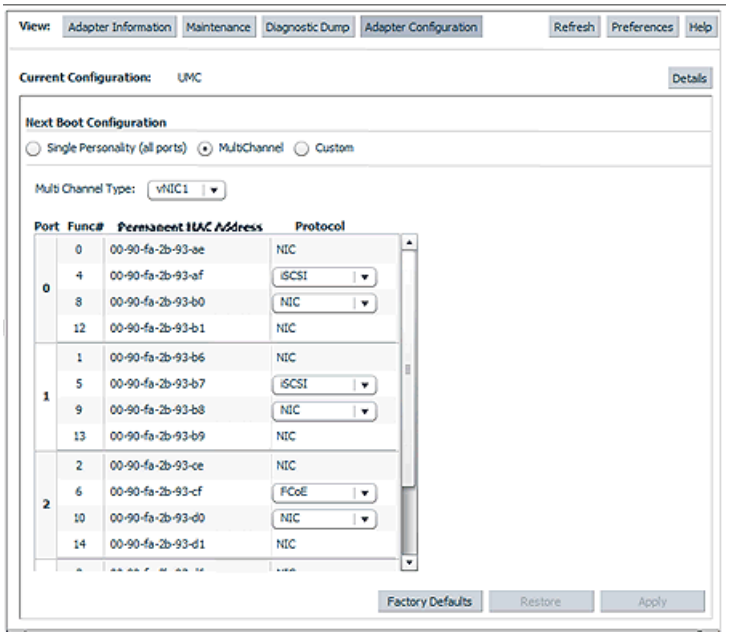


A MultiChannel Type drop-down menu is available to select the type of multichannel configuration on the adapter. The contents of the Channel Configuration table depend upon the selected multichannel type.

5.3.3.1 Setting Up a Multichannel vNIC Configuration

When vNIC is selected on the adapter, the **Adapter Configuration** tab is displayed (Figure 46).

Figure 46 Multichannel vNIC View (Mixed 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).

NOTE RoCE is available as Technical Preview only. Do not use RoCE in a production environment.

It can be assigned any value from 2 to 4094. The outer VLAN ID is assigned by the switch and can be viewed by clicking Details.

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

- The first channel can only be assigned a NIC protocol. The first channel cannot be assigned None.
- The second and third channels can be assigned the following:
  - NIC
  - None (disabled)
  - Storage protocol (On the third channel only when concurrent mode storage is available)
- The fourth and higher channels (if available) can only be assigned NIC or None (disabled).
- 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 2.

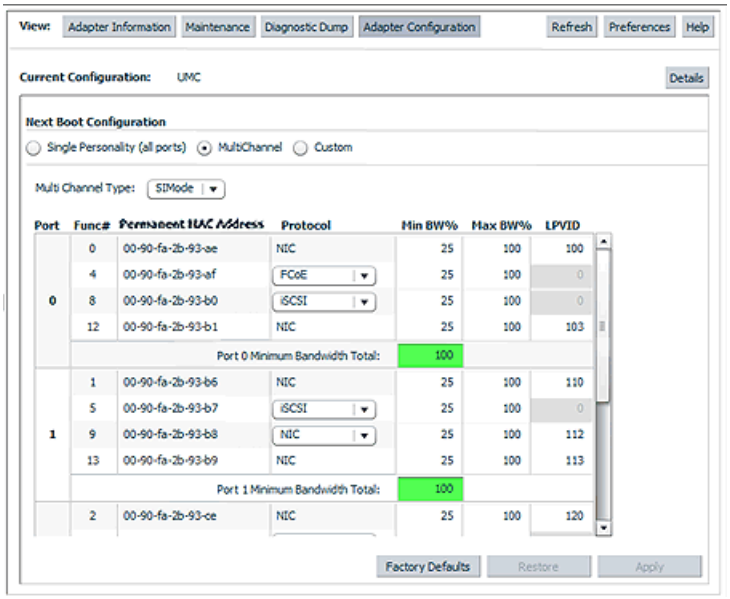
Table 2 vNIC Maximum Number of Functions per Port

Number of Ports	Port Speed	Default Functions per Port	Maximum Functions per Port
1	10Gb	4	4
2	10Gb	4	4
4	10Gb	4	4

5.3.3.2 Setting Up a Multichannel SIMode Configuration

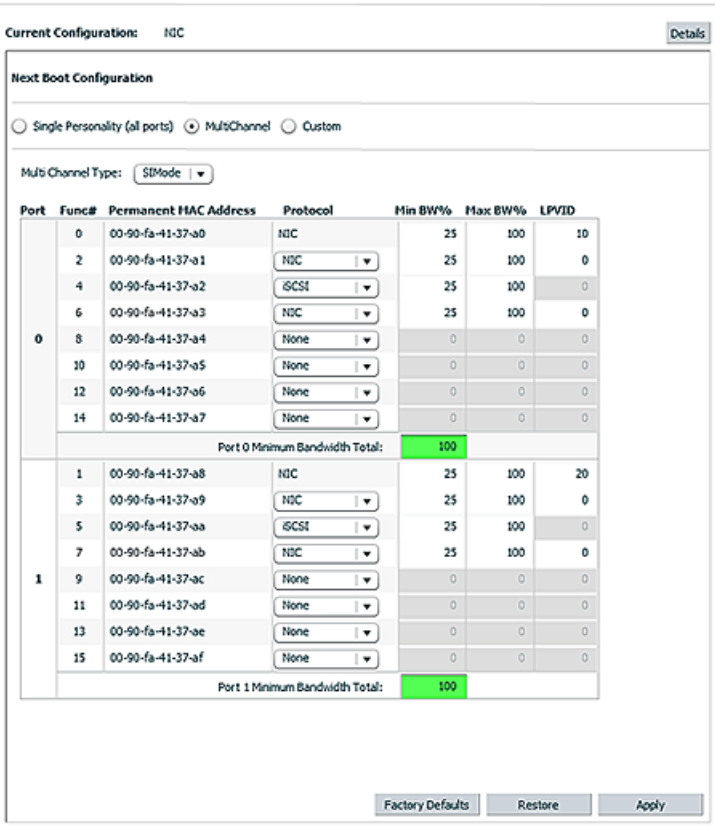
When Multichannel SIMode is selected, the **Adapter Configuration** tab is displayed (Figure 47). Other than the names, SIMode and UMC configuration is exactly the same. See Section 8.3, Configuring DCB Parameters for NIC + FCoE + iSCSI Adapter Ports.

Figure 47 Multichannel SIMode View



For two-port adapters configured to run SIMode, up to eight functions (channels) can be configured on each port. In the **Adapter Configuration** tab, when the **MultiChannel** radio button is selected, you must select a multichannel type from the Multi Channel type menu. For vNIC and UFP multichannel, up to four functions per port can be configured. However, when SIMode is selected in the menu, eight functions per port are displayed and can be configured as shown in Figure 48.

Figure 48 2-Port Lenovo System X Adapter Configured for SIMODE (with 8 Functions per Port)



When switching from a non-multichannel to an 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.

The maximum number of functions per port is shown in Table 3.

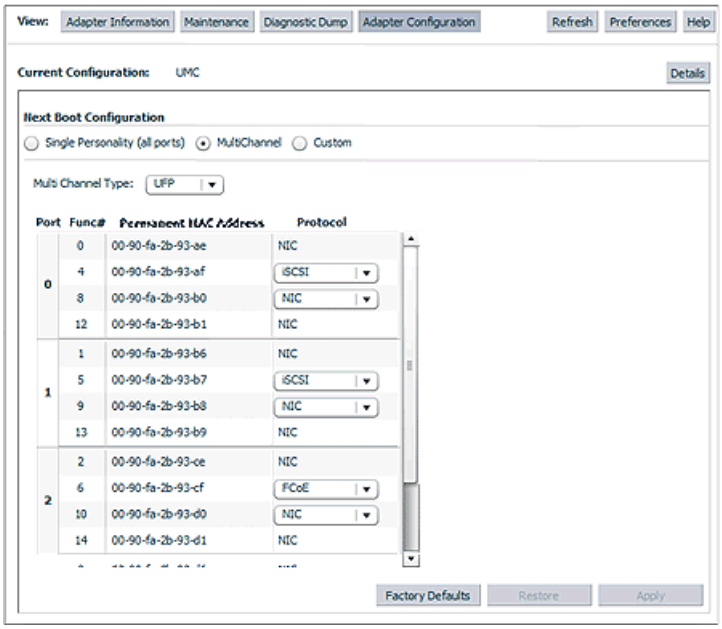
Table 3 SIMode Maximum Number of Functions per Port

Number of Ports	Port Speed	Default Functions per Port	Maximum Functions per Port
1	10Gb	4	8
2	10Gb	4	8
4	10Gb	4	4
1	40Gb	4	16

5.3.3.3 Setting Up a Multichannel UFP Configuration

When UFP is selected, the **Adapter Configuration** tab shows the Channel Configuration table, which shows the channel properties for multichannel UFP (Figure 49).

Figure 49 Multichannel UFP View (Concurrent Mode Storage)



The Channel Configuration table shows the channel properties for UFP.

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 1 up to the maximum number supported by the port. Different number of functions per port can be configured on the same adapter.

For example, you can configure four functions per port on one port and six functions per port on another. If you configure six functions per port on both ports of a 2-port adapter, then 12 functions are enumerated on the PCI bus (not 16).

When UFP is enabled on an adapter, the default is four functions per port. The maximum number of functions per port is shown in Table 4.

Table 4 UFP Maximum Number of Functions per Port

Number of Ports	Port Speed	Default Functions per Port	Maximum Functions per Port
1	10Gb	4	8
2	10Gb	4	8
4	10Gb	4	4
1	40Gb	4	16

There are drop-down menus for each port's second and third channels. Multichannel UFP channel protocol assignments are subject to the following rules:

- The first channel can only be assigned NIC protocols. The first channel cannot be assigned None.
- NIC+RoCE cannot be configured with UFP.
- The second and third channels can be assigned the following:
  - NIC

- None (disabled)
- Storage protocol (third channel only when concurrent mode storage is available)
- Storage protocols can only be assigned to the second channel when mix-mode storage is available.
- The third channel can be assigned a storage protocol when concurrent mode storage is available under the following restrictions.
  - Can only be assigned a storage protocol if the second channel is assigned a storage protocol.
  - The storage protocol must be different than the second channel's storage protocol.

- The fourth and higher channels (if available) can be assigned the NIC protocol or None (disabled).

5.3.4 Setting Dell NPar Bandwidth Settings

Similar to UMC, a minimum and maximum bandwidth is assigned to each enabled (that is, not None) partition. The sum of the minimum bandwidths must add up to 100. The maximum bandwidth for each partition must be equal to or greater than the minimum bandwidth on that partition and no larger than 100. Unlike UMC, if the minimum and maximum bandwidths are set to 0, the logical link on the adapter is not brought down. A small of amount of data trickles through the port for that partition when bandwidth is available. Bandwidth changes take effect immediately (no reboot required).

If DCB is enabled on a switch, the switch settings for bandwidth are applied instead of the adapter settings. These can be viewed in the ETS Priority Group Properties area of the **DCB** tab. See Section 8, DCB Configuration, for more information.

5.3.4.1 Two-Port Adapters

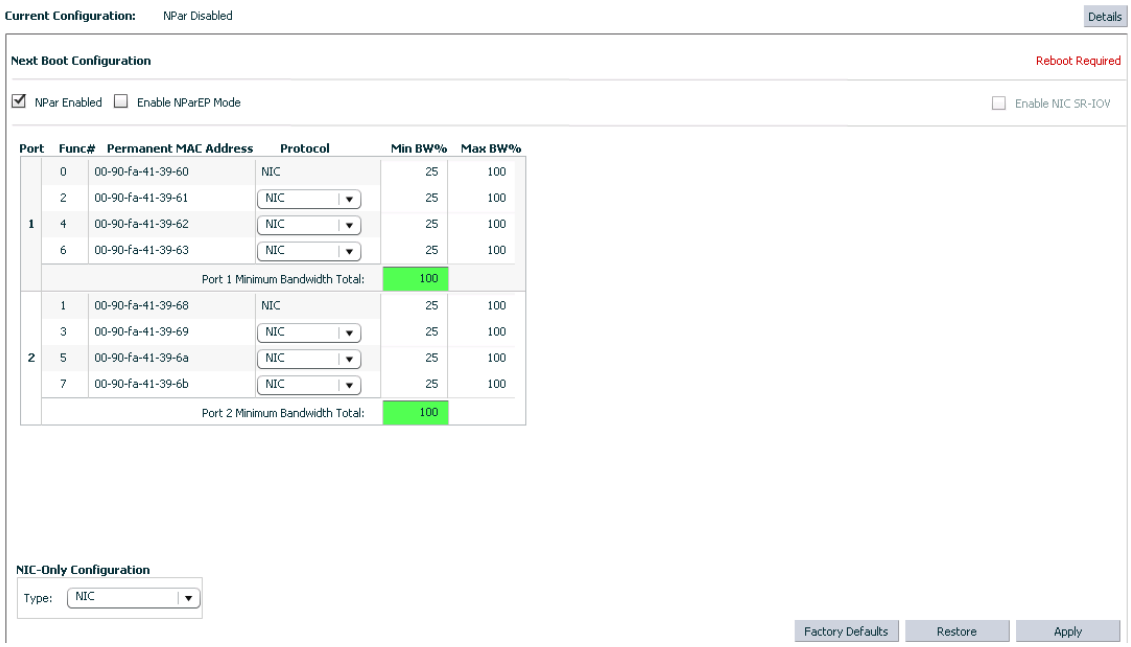
Dell two-port adapters allow up to four or eight partitions per port, depending upon whether the NParEP Mode checkbox is checked or cleared.

5.3.4.1.1 NParEP Mode Disabled

For Dell NPar adapters, the **Adapter Configuration** tab is different than for the other adapters. There are two choices in the Next Boot Configuration: NPar Enabled and Enable NParEP Mode. When NPar is checked (enabled) and NParEP Mode is cleared (disabled), the adapter is basically a NIC-only adapter (Figure 50).

If the NParEP Mode checkbox is cleared, only four partitions per port can be configured. FCoE, iSCSI, or None can be configured on the second through fourth partitions, but only one FCoE and one iSCSI can be configured. None cannot be configured on the first partition.

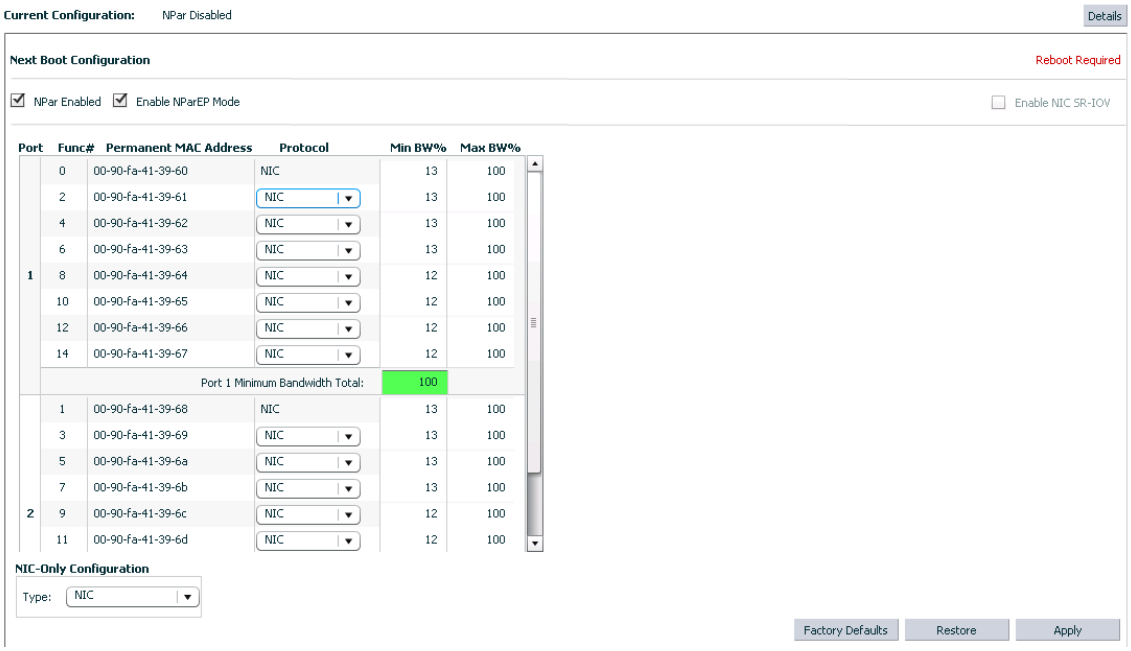
Figure 50 Dell Two-Port NPar Configuration, NParEP Mode Disabled



5.3.4.1.2 NParEP Mode Enabled

If the NParEP Mode checkbox is selected (enabled), up to eight partitions per port can be configured (Figure 51). FCoE or iSCSI can be configured on the second through fourth partitions, but only one FCoE and one iSCSI can be configured. Partitions higher than the fourth partition can only configured for NIC or None. None cannot be configured on the first partition.

Figure 51 Dell Two-Port NPar Configuration, NParEP Mode Enabled



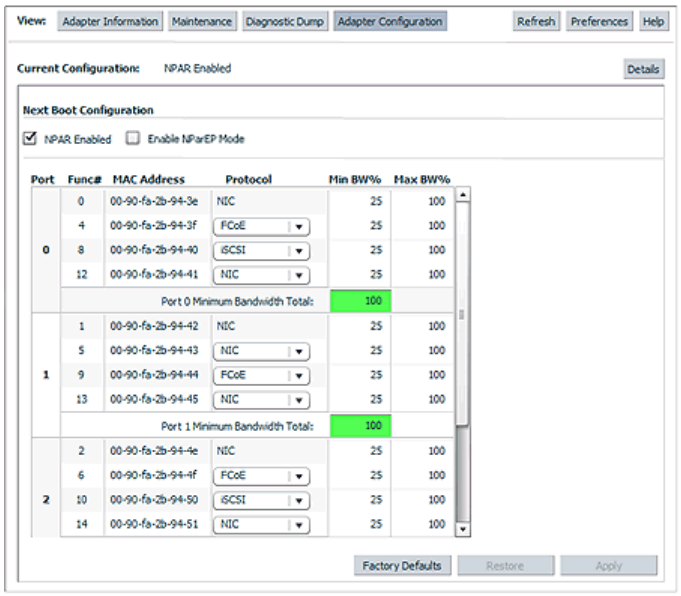
5.3.4.2 Four-Port Adapters

Dell four-port adapters allow up to two or four partitions per port, depending upon whether the NParEP Mode checkbox is checked or cleared.

5.3.4.2.1 Dell NPar Enabled

When NPar is enabled on the **Adapter Configuration** tab (Figure 52), you can switch between showing eight functions and showing sixteen functions for a given adapter.

Figure 52 Adapter Configuration Tab with NPar Enabled



You can display the appropriate number of functions for an adapter in an NParEP Mode configuration.

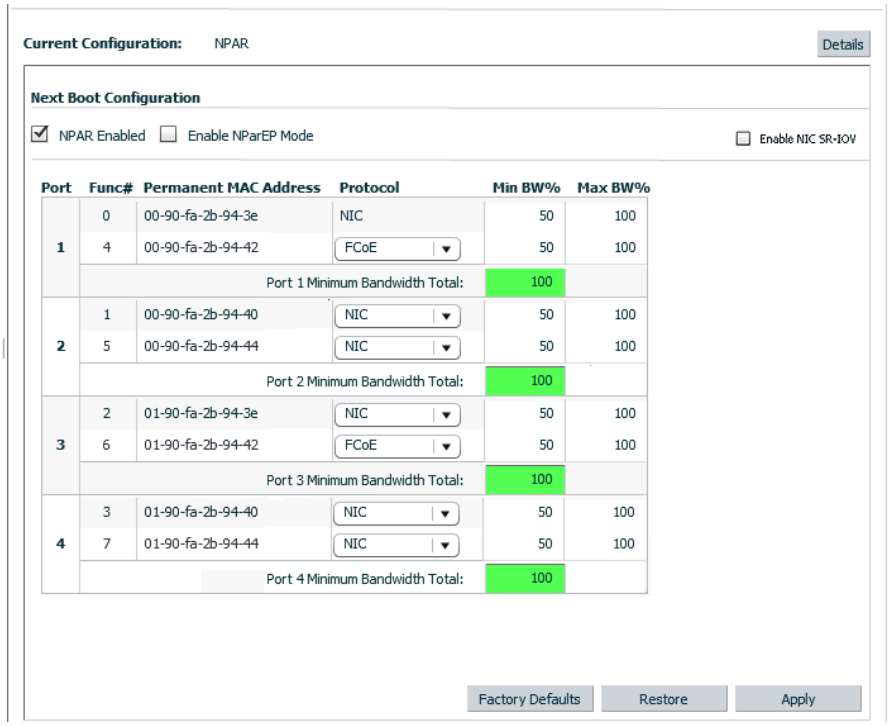
NPar protocol assignments are subject to the following rules:

- The first channel can only be assigned NIC protocols. The first channel cannot be assigned None.
- The second and third channels can be assigned the following:
  - NIC
  - None (disabled)
  - Storage protocol (Note: on 3rd channel only when concurrent mode storage is available.)
- The fourth and higher channels (if available) can only be assigned NIC or None (disabled).
- The total of the minimum bandwidths of the enabled functions for each port must add up to 100 percent. The maximum bandwidth for each function must be greater than, or equal to, the minimum bandwidth for that channel.
- After changing the NPar configuration and applying the changes, except for bandwidth changes, a reboot is required to activate the new configuration.
- NParEP support is available only on Dell 13G or newer systems.

5.3.4.2.2 NParEP Mode Disabled

When NPar is enabled and NParEP Mode is disabled on the **Adapter Configuration** tab (Figure 53) only two partitions per port can be configured. FCoE, iSCSI, or None can be configured on the second function, None cannot be configured on the first function.

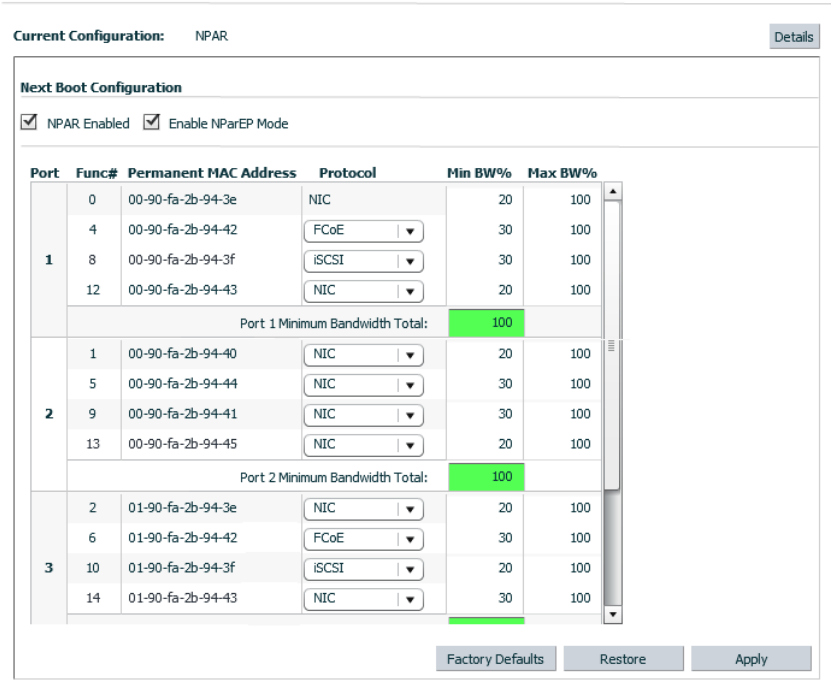
Figure 53 Dell Four Port NPar Configuration, NParEP Mode Disabled



5.3.4.2.3 NParEP Mode Enabled

If the NParEP Mode checkbox is selected (enabled), up to four partitions per port can be configured (Figure 54). FCoE or iSCSI can be configured on the second through fourth partitions, but only one FCoE and one iSCSI can be configured. None cannot be configured on the first partition.

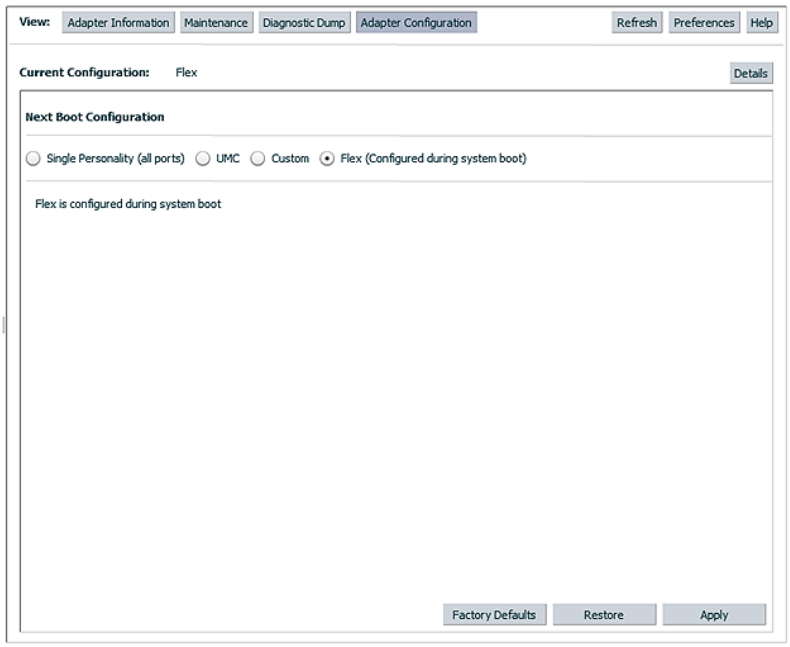
Figure 54 Dell Four Port NPar Configuration, NParEP Mode Enabled



5.3.5 Setting Up HP Flex Configurations

In the Next Boot Configuration area, the **Flex** or **FlexQbg** radio button is available, depending upon the type of Flex adapter installed in your system. Selecting this radio button causes Flex to be configured during the next system boot.

Figure 55 HP Flex Adapter Configuration Tab, Flex Selected

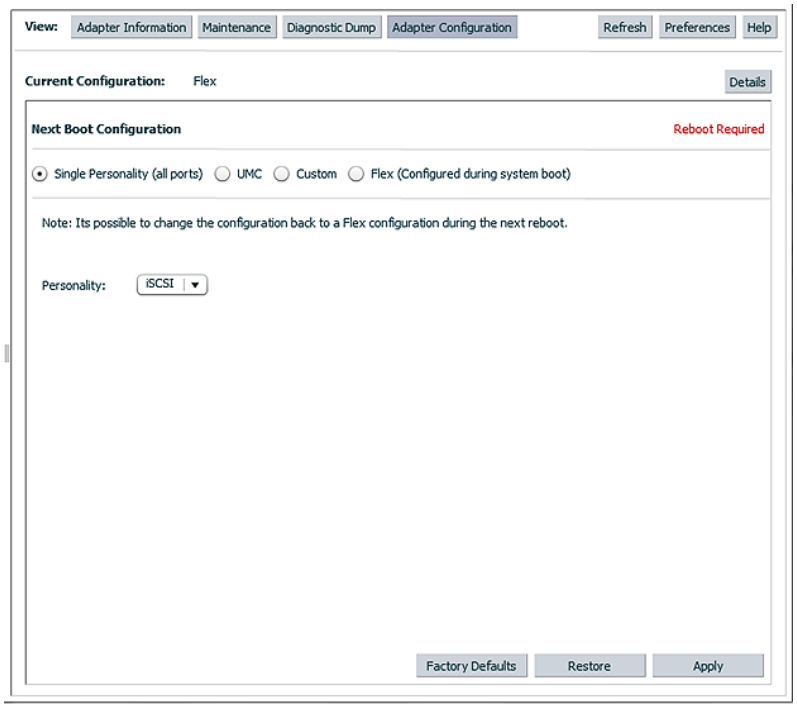


To configure an adapter for Flex or FlexObg, perform the following steps:

1. Select the **Flex** or **FlexQbg** radio button. The **Flex** or **FlexQbg** button title is displayed depending upon the type of Flex adapter installed in your system.
2. Click **Apply**. Once you click **Apply**, the **Flex** or **FlexQbg** radio button is inactive.

Flex is configured during the next system boot.

Figure 56 HP Flex Adapter Configuration Tab, Single Personality Selected



5.3.6 Removing a Port

Port removal allows you to prevent the appearance of the port and its functions on the PCiE bus. It is supported for non-NPar OCe14000-series adapters. A port can only be removed in multichannel and Custom modes. This option is not available in Single Personality mode, since this type of configuration does not allow configuring ports and functions individually or separately.

NOTE The first port of an adapter cannot be disabled. This is because function 0 (which is always on the first port) cannot be set to None.

To remove a port:

Set the first function to None in all except the first port in Multichannel and Custom modes.

When all functions belonging to a port are set to None, the port is removed and is not discovered after reboot.

To add a port back:

Perform one of the following:

- Set the first function in the removed port to NIC or NIC+RoCE protocol.
- Configure the adapter to single personality, which causes all the removed ports to be added back after reboot.

5.3.6.1 Custom View Configuration

The Custom view configuration provides a drop-down menu, which includes None as one of the items in the list for all the first functions belonging to ports other than the first port. Function 0 (on the first port) is not able to change its protocol. By setting the first function on the port to None, all of the other functions on that port are automatically set to None.

5.3.6.2 Multichannel Configuration View (UMC, SIMode, vNIC, UFP)

The Multichannel view configuration provides a drop-down menu, which includes None as one of the items in the list, for all the first functions belonging to ports other than the first port. Function 0 (on the first port) is not able to change its protocol.

If the first function belonging to ports other than the first port is set to None, then all the remaining functions from that port are forced to None, and the protocol drop-down menus become disabled to prevent changing protocols. The drop-down menus become enabled when the first function is reset to NIC.

5.3.6.3 Adding Ports in Multichannel Mode (UMC, SIMode, vNIC, UFP)

When you add a port while running in multichannel mode, you cannot set any multichannel properties for the port being added until after a system reboot. The Min BW, Max BW, and LPVID (for UMC and SIMode) text field boxes (or any other multichannel property box) remain unavailable. A pop-up box is also displayed immediately after clicking **Apply**, indicating that multichannel properties can only be set on this port after a system reboot has been performed.

Chapter 6: Managing Protocols

This section pertains to managing NIC, FCoE, and iSCSI protocols.

6.1 NIC

This section pertains to viewing NIC adapter and port information. This section also includes booting from the iSCSI SAN and enabling and disabling VEPA.

6.1.1 Viewing NIC Port Information

When you select a NIC port from the Broadcom CNA Device Management tree-view, the NIC **Port Information** tab shows general attributes associated with the selected NIC port. If the adapter supports vNIC, vNIC data is also displayed.

NOTES

- NIC ports not only exist on NIC-only adapters, but also exist on FCoE and iSCSI adapters.
- vNIC is supported only on Lenovo System X virtual fabric adapters. For specific information as to whether it is supported on a specific adapter, refer to the release notes that are available on the Lenovo System X adapter pages on http://www.broadcom.com.

To view general NIC port information, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the NIC port for which you want to view information.

NOTE In Virtual Ports view, NIC ports appear only on FCoE adapters. They do not appear on NIC-only adapters.

3. Select the **Port Information** tab (Figure 57).

Figure 57 NIC Port Information Tab with SR-IOV Disabled on the Adapter



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

- Driver Name – The NIC driver file name.
- Driver Version – The NIC driver version.
- MAC Address – The NIC MAC address currently assigned to the port.
- Perm MAC Address – The original factory assigned NIC MAC address.
  - For Dell, the virtual MAC address is displayed along with the permanent MAC address.
- IP Address – The IP address for the NIC port.
- Subnet Mask – The subnet mask for the NIC port.
- Gateway Address – The gateway address for the NIC port.
- IP Address Origin – The origin of the IP address (DHCP or static).
- Link Status – The link status is always displayed as n/a (not available).
- LinkDown Reason – The reason the link is down on the selected port.
- Interface Name – The interface assigned to this port by the host operating system.
- Port Speed – The current port speed of the NIC port.
- 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.
- IPv6 Address – The IPv6 address, if applicable.
- PCI Function – The PCI function number assigned to the NIC function.
- PCI Bus Number – The PCI bus number to which the port is mapped.
- Current MTU – The current transmission unit for Ethernet traffic.
- Max MTU – The maximum transmission unit for Ethernet traffic.
- Min Bandwidth – The minimum SoftCap bandwidth assigned.
- Max Bandwidth – The maximum SoftCap bandwidth assigned.

- Enable VEPA checkbox (OCe14000-series adapters only) - Select this checkbox to enable VEPA. For more information on VEPA, see Section 6.1.1.1, Enabling and Disabling VEPA (OCe14000-Series Adapters Only).

  NOTE This checkbox is not enabled if SR-IOV is disabled, or if multichannel is enabled.

- **Network Boot** radio buttons - For more information on network boot, see Section 6.1.1.2, Network Boot. The **Network Boot** radio buttons do not appear if SR-IOV is enabled on the adapter.
  - None – Select to disable the network boot functions.
  - PXE – Select to boot from a network interface using a preboot execution environment.
  - iBFT – Boots from an iSCSI SAN using the iBFT format. The iBFT boot option is only available for adapters that support the NIC or iSCSI protocol.

6.1.1.1 Enabling and Disabling VEPA (OCe14000-Series Adapters Only)

VEPA allows the network traffic at the VFs to bypass the hypervisor's virtual switch, and allows the PFs to transmit the traffic to the external switch. This relieves the CPU of the switching burden, and puts the traffic on the physical network, where you can monitor it.

To enable VEPA, check **Enable VEPA** (Figure 57). To disable VEPA, clear the Enable VEPA checkbox.

NOTE SR-IOV must be enabled to enable VEPA. For more information on SR-IOV, see Section 5.3.1.2, Enabling or Disabling SR-IOV.

VEPA is only available on OCe14000-series adapters, and it is only available for NIC ports.

6.1.1.2 Network Boot

NOTES

- Booting from an iSCSI SAN using the iSCSI Boot Firmware Table format (iBFT) is not supported in this release.
- The iBFT boot option is only available for adapters that support the NIC or iSCSI protocol.

To boot from the iSCSI SAN, the host must have an iSCSI boot capable network adapter that supports the iBFT format. The iBFT is a method of communicating parameters about the iSCSI boot device to an operating system. To boot the host from the network interface, use PXE.

NIC functions include iBFT boot configuration and PXE boot configuration, which are displayed under the Network Boot section of the NIC **Port Information** tab. If the currently selected NIC port does not support network boot, the Network Boot section is not displayed.

NOTE Network boot is supported by CIM Provider 3.6.12 and above.

Network boot options are only available on ports with a PCI function of 0 or 1.

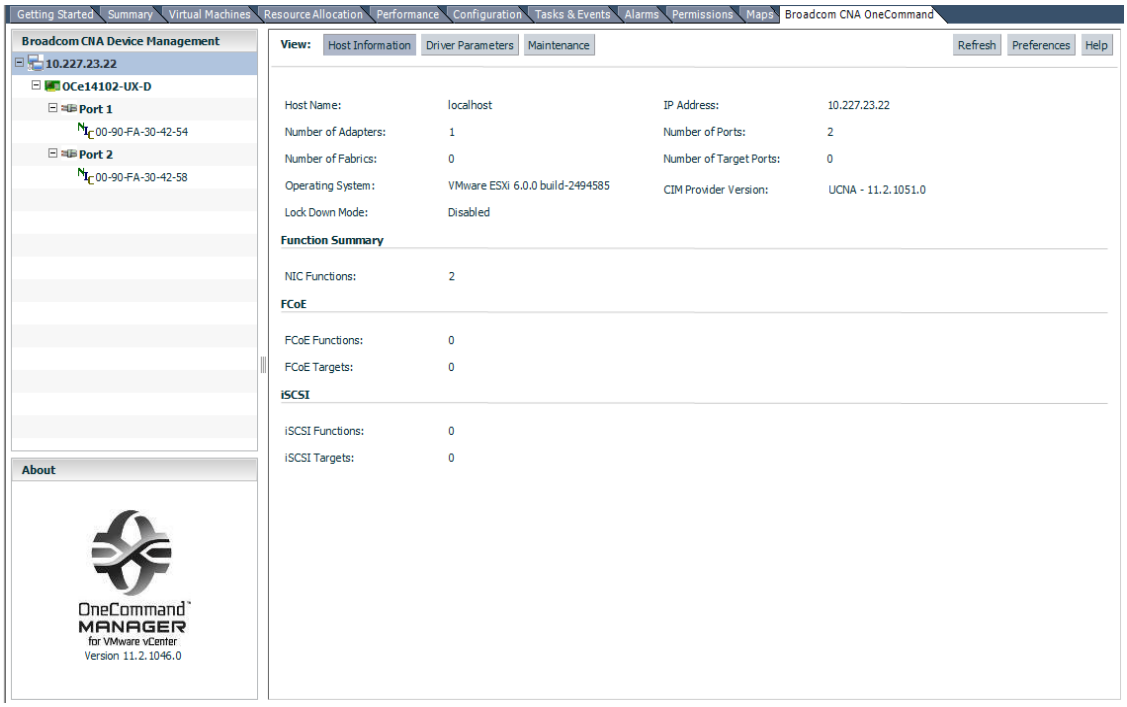
To select a network boot configuration on the selected port, click the desired option from the Network Boot section:

- iBFT – To boot from an iSCSI SAN using the iSCSI Boot Firmware Table format. This feature is not supported in this release.
- PXE – To boot from a network interface using a preboot execution environment.
- None – To disable the network boot functions.

6.1.2 Viewing Virtual Functions on a Host Operating System

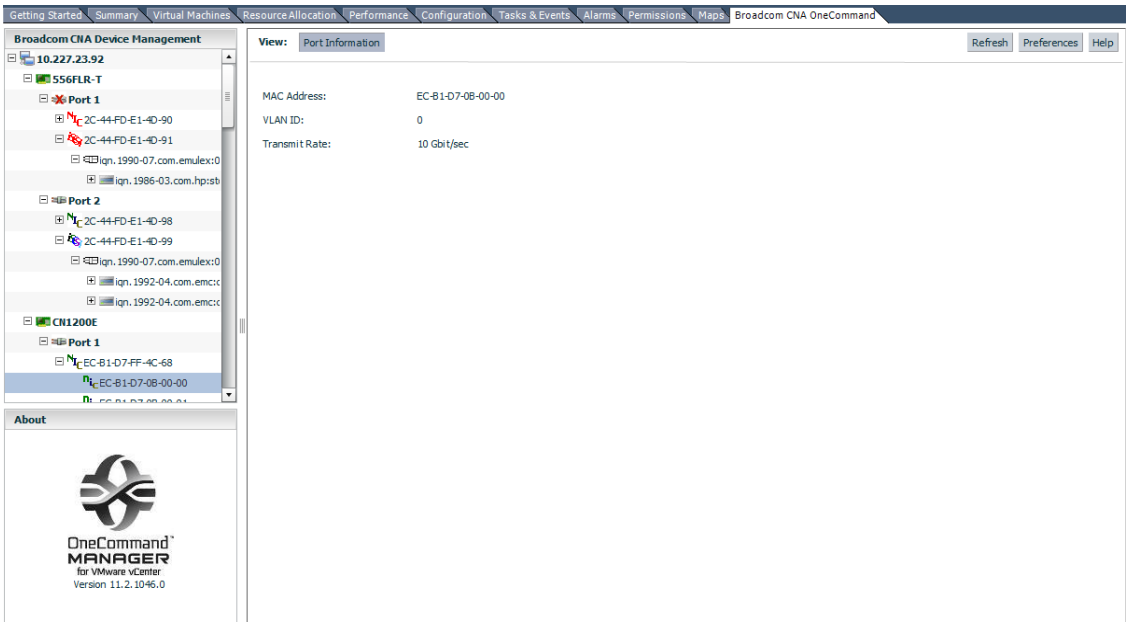
OneCommand Manager for VMware vCenter can discover the VFs running on the discovered physical functions and show them in the discovery-tree (Figure 58).

Figure 58 Virtual Functions Discovered on the Physical Function on the Host Operating System



Select the VF under the host operating system in the discovery-tree to display the **Port Information** tab (Figure 59).

Figure 59 Port Information tab for NIC Virtual Function



The following Port Information fields are displayed:

- MAC Address – The NIC MAC address currently assigned to the port.
- VLAN-ID – The VLAN identifier used by the NIC port.
- Link Status – The status of the link on the selected port.
- Transmit Rate – The rate at which data is transmitted over the port in Mbps.

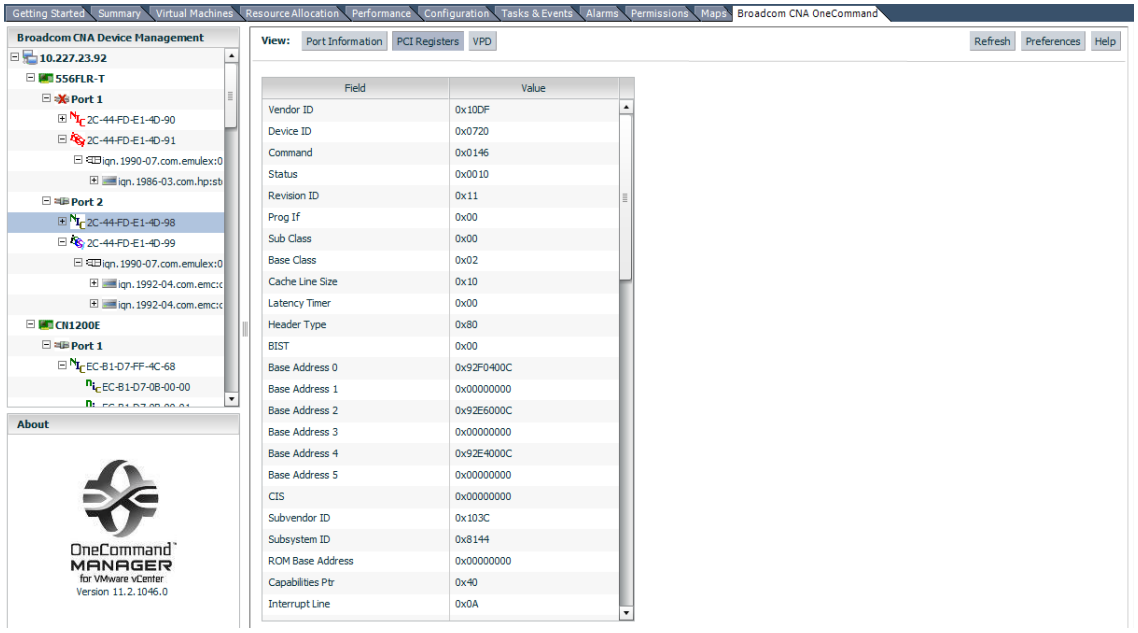
6.1.3 Viewing PCI Registers for a NIC Port

The **PCI Registers** tab contains PCI information about the selected NIC function.

To view a NIC port's PCI registers, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the NIC port for which you want to view PCI information.
3. Select the **PCI Registers** tab (Figure 60).

Figure 60 NIC Port PCI Registers Tab



6.1.4 Viewing NIC Port Vital Product Data (VPD)

The **VPD** tab shows vital product data (if available) for the selected NIC port such as the product name, part number, and serial number for example.

To view VPD information of an NIC port, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the NIC port for which you want to view VPD information.
3. Select the **VPD** tab (Figure 61).

Figure 61 NIC Port VPD Tab



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

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

NOTE Some adapters may show additional VPD information such as EC and MN.

6.2 FCoE

This section pertains to managing FCoE adapter and port information.

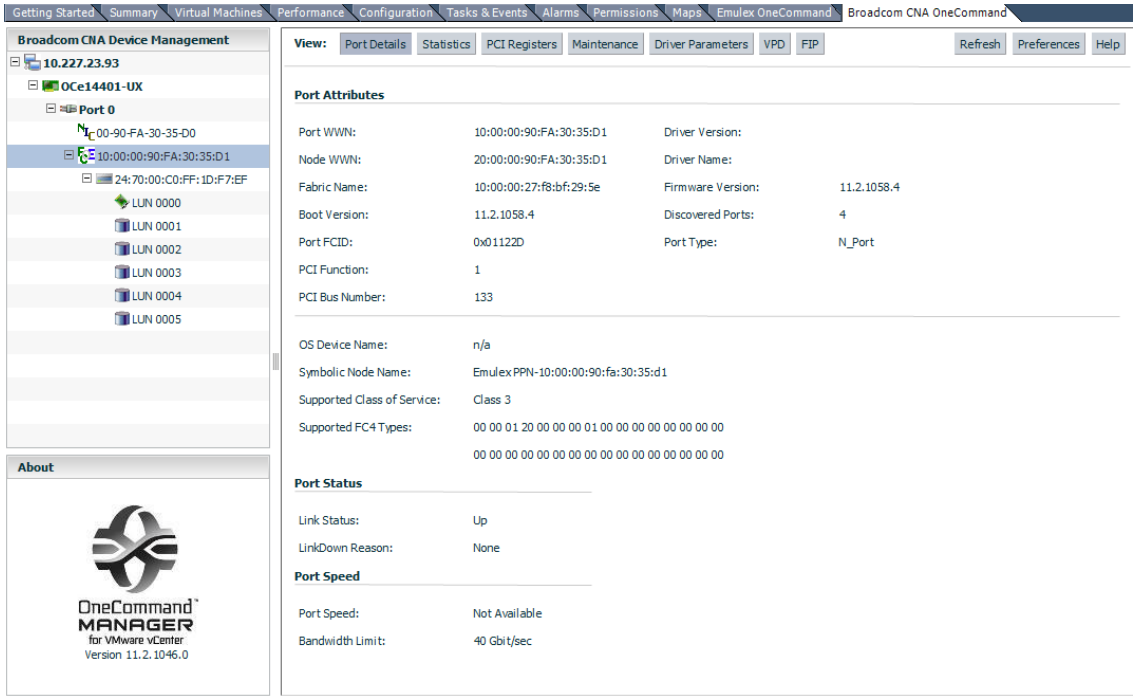
6.2.1 Viewing FCoE Port Details

Use the Port Details tab to view general attributes associated with the selected FCoE port.

To view port details for an FCoE port, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the FCoE port whose details you want to view.
3. Select the **Port Details** tab (Figure 62).

Figure 62 FCoE Port Details Tab



The following FCoE **Port Details** tab fields are displayed:

- Port Attributes Area
  - Port WWN – The port World Wide Name of the adapter.
  - Node WWN – The node World Wide Name of the selected adapter.
  - Fabric 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 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 shows **Disabled**.
  - Port FCID – The FCoE ID for the selected adapter port.
  - PCI Function – The PCI function number assigned by the system.
  - PCI Bus Number – The PCI bus number.
  - 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 current operational mode of the selected adapter's port.
  - 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. There are three classes of service:
    - 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:
  - Link Status – The status of the link on the selected adapter port.
  - LinkDown Reason – The reason for the Link Status on the selected port.
- Port Speed area:
  - Port Speed – The current port speed of the selected adapter port.
  - Bandwidth Limit – The QoS bandwidth restriction on the port.

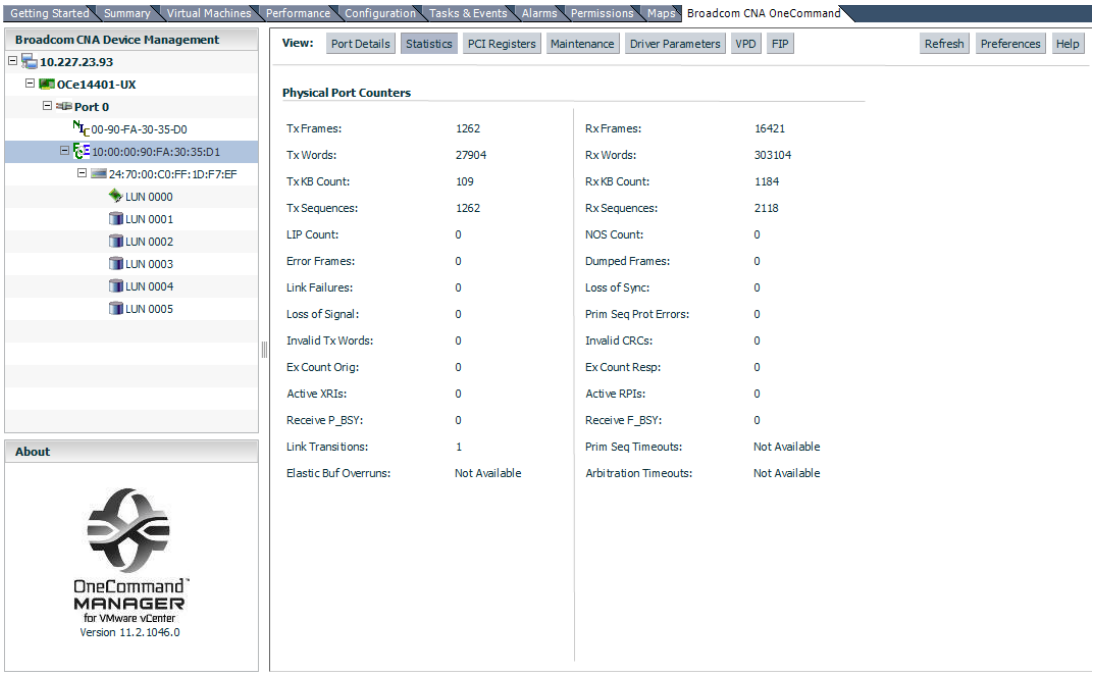
6.2.2 Viewing FC/FCoE Port Statistics

When you select an FC/FCoE port from the discovery-tree, the **Port Statistics** tab shows cumulative totals for error events and statistics on the port. Some statistics are cleared when the adapter is reset.

To view statistics for a port, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the port whose statistics you want to view.
3. Select the **Statistics** tab (Figure 63).

Figure 63 FC/FCoE Port Statistics Tab



The following FC/FCoE **Port Statistics** tab fields are displayed:

- Tx Frames – The FC frames transmitted by this adapter port.
- Tx Words – The FC words transmitted by this adapter port.
- Tx KB Count – The FC kilobytes transmitted by this adapter port.

- Tx Sequences – The FC sequences transmitted by this adapter port.
- LIP count – The number of LIP events that have occurred for the port. This field is supported only if the topology is 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 can cause 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 adapter port.
- Ex Count Orig – The number of FC exchanges originating on this port.
- Active XRIs – The number of active exchange resource indicators.
- 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 adapter port.
- Rx Words – The number of FC words received by this adapter port.
- Rx KB Count – The received kilobyte count by this adapter port.
- Rx Sequences – The number of FC sequences received by this adapter port.
- NOS count – The number of NOS events that have occurred on the switched fabric (not supported for an arbitrated loop).
- Dumped Frames – The number of frames that were lost due to a lack of host buffers available.
- 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 port.
- Active RPIs – The number of remote port indicators.
- 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.
- 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.

6.2.3 Viewing PCI Registers

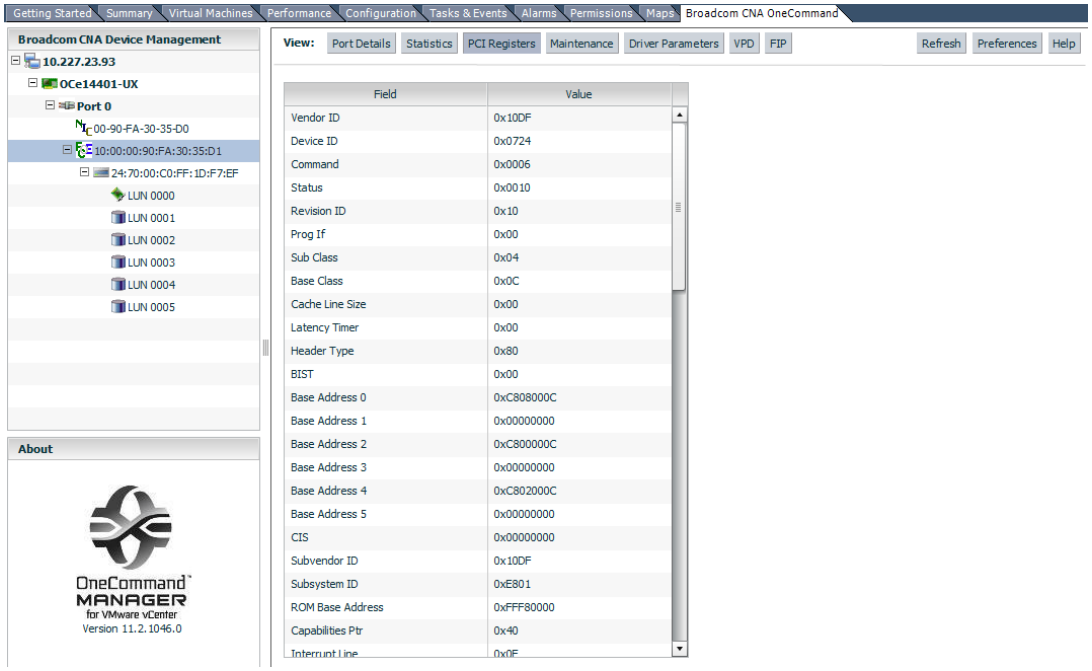
The **PCI Registers** tab shows PCI information, including PCIe details, about the selected port.

NOTE The PCI fields may vary with the type of adapter installed.

To view PCI registers for an FCoE port, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the FCoE port whose PCI information you want to view.
3. Select the **PCI Registers** tab (Figure 64)

Figure 64 FCoE Port PCI Registers Tab



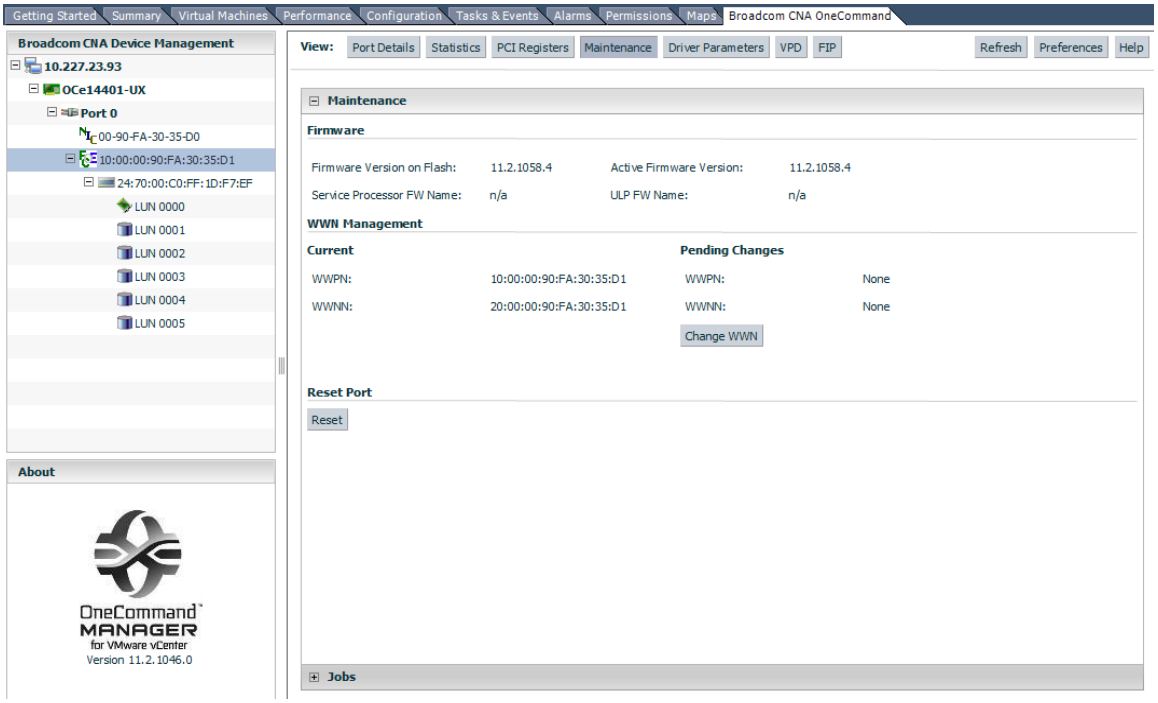
6.2.4 Viewing FCoE Port Maintenance and Firmware Information

The **Maintenance** tab shows firmware information for an FCoE port.

To view firmware information for a port, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the FCoE port whose firmware information you want to view.
3. Select the **Maintenance** tab (Figure 65).

Figure 65 Port Maintenance Tab



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

- Firmware area:
  - Firmware Version on Flash – The firmware version stored on the adapter's nonvolatile storage. When the system restarts, this version becomes the active firmware version.
  - Service Processor FW Version – The operational firmware version on the adapter.
  - Active Firmware Version – The firmware version running on the selected adapter.
  - ULP FW Name – The firmware version running on the ULP processors within the ASIC.
- WWN Management area:
  - Current
    - WWPN – The WWPN for the selected adapter port.
    - WWNN – The WWNN for the selected adapter port.
  - Pending Changes
    - WWPN – If the WWPN has been changed, the new WWPN is displayed in this list. After rebooting, the new WWPN is displayed in the Current list.
    - WWNN – If the WWNN has been changed, the new WWNN is displayed in this list. After rebooting, the new WWNN is displayed in the Current list.

  See Section 6.2.4.1, Changing WWN Configuration for an FCoE Port for more information about changing the WWN configuration.

  For instructions on resetting a port, see Section 6.2.4.2, Resetting an FCoE Port for more information.

6.2.4.1 Changing WWN Configuration for an FCoE Port

The **Maintenance** tab enables you to change the WWPN and the WWNN of a selected adapter port. For example, you might want to use an installed adapter as a standby if another installed adapter fails. By changing the standby adapter's WWPN or WWNN, it can assume the identity and configuration (such as driver parameters and persistent binding settings) of the failed adapter.

There are three options for referencing WWNs:

- Factory default WWN – The WWN as shipped from the factory.
- Nonvolatile WWN – The values that are saved in the nonvolatile adapter's flash memory that survive a reboot or a power outage.
- Volatile WWN – A temporary value that is saved in the volatile memory on the flash. If volatile WWNs are set, they are used instead of the nonvolatile WWNs. Volatile WWN changes require a warm system reboot to take effect. Volatile WWN changes are lost on systems that power cycle the adapters during the reboot.

CAUTION Changing volatile WWNs result in taking the selected adapter offline. Ensure that this adapter is not controlling a boot device and that all I/O activity on this adapter has stopped before proceeding. This could result in data loss or corruption.

6.2.4.1.1 Considerations When Changing WWN Configuration for an FCoE Port

- To avoid address conflicts, do not assign a WWNN or WWPN with OneCommand Manager for VMware vCenter if you also use another address management tool.
- The WWPN and WWNN in the Pending Changes list can show **n/a** instead of **None**. This happens when the remote host is busy processing some critical task and WWN Management cannot obtain the current state of WWN management.
- In an environment where preboot management exists, a WWPN or WWNN modified by OneCommand Manager for VMware vCenter can be overridden by preboot management such as Lenovo System X BOFM and industry standard CLP.

  For example:

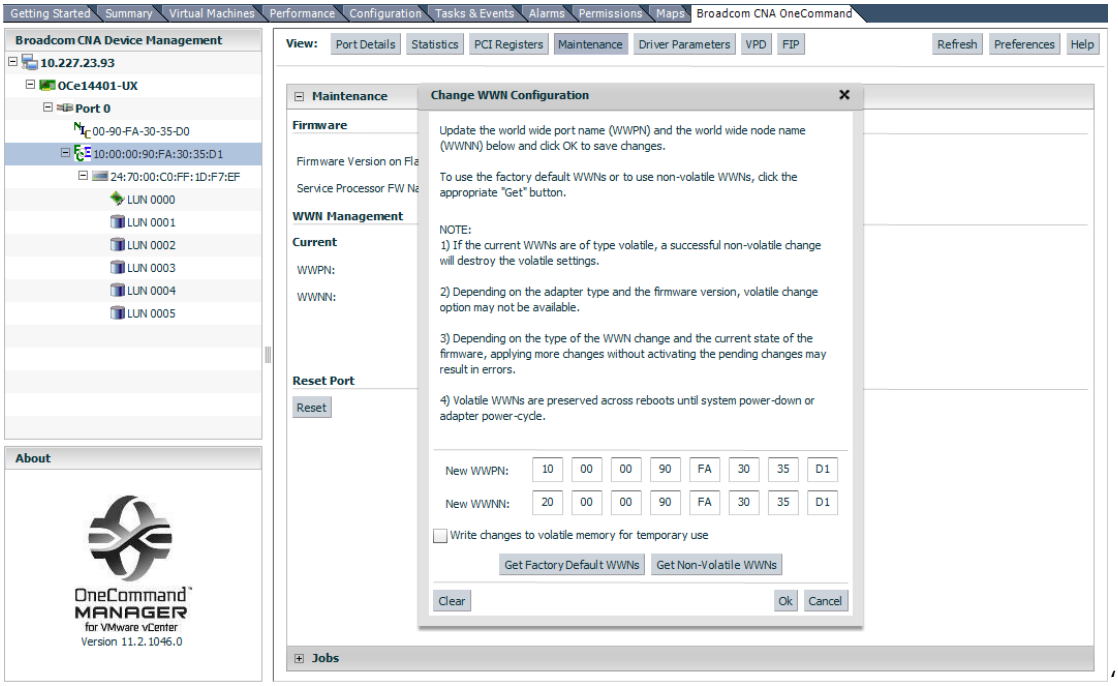
  In an environment with CLP/BOFM, OneCommand Manager for VMware vCenter modifies the WWNN or WWPN. OneCommand Manager for VMware vCenter requires a reboot to complete the change. After a reboot, the CLP string is sent during the system boot and rewrites the WWNN or WWPN, or EFIBoot finds the BOFM protocol and uses the default WWNN or WWPN per the BOFM's command.

  In an environment without CLP/BOFM, OneCommand Manager for VMware vCenter modifies the WWNN or WWPN. OneCommand Manager for VMware vCenter requires a reboot to complete the change. The system boots and the OneCommand Manager for VMware vCenter-modified WWNN or WWPN is used.
- On a system where OneCommand Manager for VMware vCenter is installed, make sure the port numbers configured during the installation are open and dedicated to the OneCommand Manager for VMware vCenter server only. No other service should be listening on this port.

To change a port's WWPN or WWNN, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the port for which you want to change the WWN information.
3. Select the **Maintenance** tab (Figure 65).
4. Click **Change WWN**. The Change WWN Configuration dialog is displayed (Figure 66).

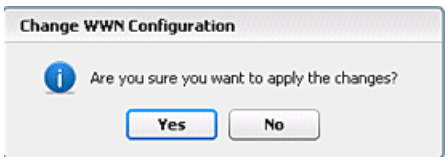
Figure 66 Change WWN Configuration Dialog



5. Do one of the following:
   - Enter a new WWPN or WWNN.
   - Click **Get Factory Default WWNs** to load the settings that were assigned when the adapter was manufactured. These values can then be modified and saved as volatile or nonvolatile WWNs.
   - Click **Get Non-Volatile WWNs** to load the current nonvolatile WWN. These values can be modified and saved to volatile or nonvolatile memory. You can edit the data returned from the button.
6. Check **Write changes to volatile memory for temporary use** to save the New WWPN and New WWNN settings as volatile WWNs. If cleared, the New WWPN and New WWNN settings are saved as nonvolatile WWNs.

   NOTE If the adapter or firmware does not support volatile WWNs, the **Write changes to volatile memory for temporary use** checkbox is disabled. This type of change is supported locally and via TCP/IP connections. This checkbox is disabled for remote in-band adapters, regardless of adapter models and firmware version.
7. Click **OK**. The following warning is displayed (Figure 67).

Figure 67 Warning About Changing WWN



8. Click **Yes**. The New WWPN and new WWNN values are saved. The new WWPN and WWNN appear in the Pending Changes list in the WWN Management area of the **Maintenance** tab.

9. Reboot the system for the changes to take effect (the new WWPN and WWNN appear in the Pending Changes list of the **Maintenance** tab until the system is rebooted). After rebooting, the changes are applied and appear in the Current section of the **Maintenance** tab.

NOTE After changing the WWN of an adapter, you must reboot the system before trying to access the adapter on that system.

6.2.4.2 Resetting an FCoE Port

For FCoE ports, a reset is only necessary to activate changed driver parameters or FIP settings. It does not actually perform an adapter level reset of the port.

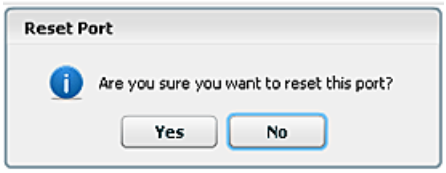
CAUTION Do not reset an adapter port while copying or writing files. This could result in data loss or corruption.

NOTE When you reset a port or change WWN configuration on OneCommand Manager for VMware vCenter, do not perform any active management operations on the ESXi host.

To reset an FCoE port, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the FCoE port you want to reset.
3. Select the **Maintenance** tab (Figure 65).
4. Click **Reset**. The following warning is displayed (Figure 68):

Figure 68 Reset Port Warning



5. Click **Yes**. The adapter port is reset. The reset can require several seconds to complete. While the adapter port is resetting, the message Operation is in progress is displayed. When the reset is finished, the message Reset Port Completed is displayed.

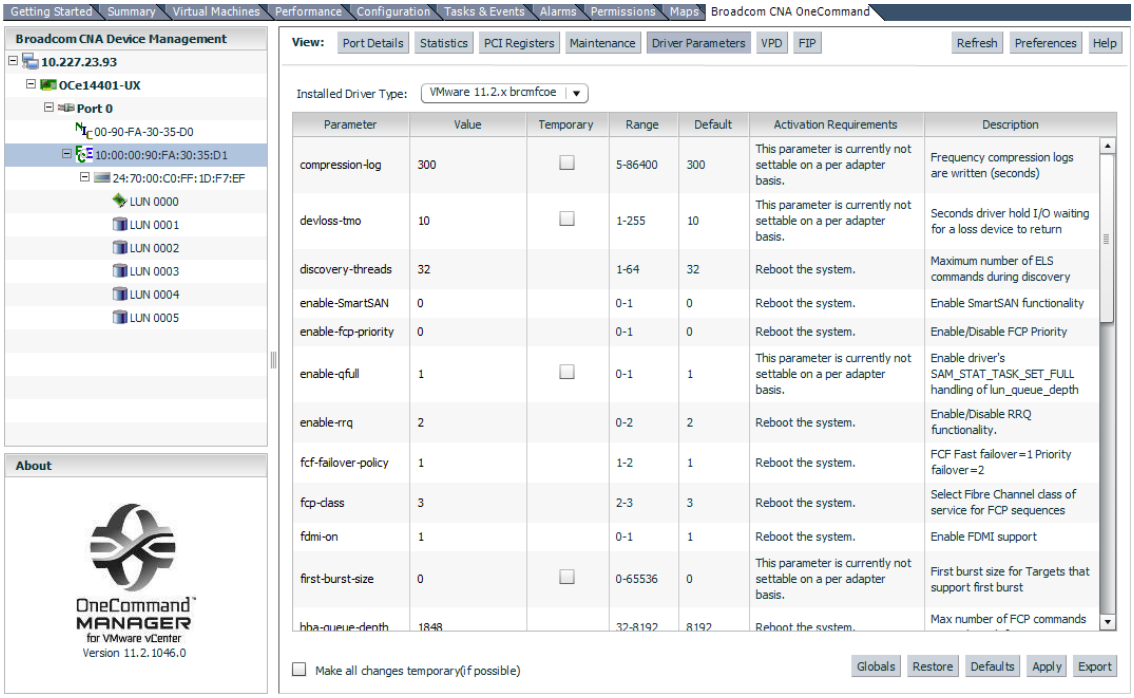
6.2.5 Configuring FCoE Port Driver Parameters

The **Driver Parameters** tab shows driver parameters for a port.

To view driver parameters for a port, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the port whose driver parameters you want to view.
3. Select the **Driver Parameters** tab (Figure 69).

Figure 69 FCoE Port Driver Parameters Tab



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

- Installed Driver Type – The current driver installed on this host.
- Port Parameter table – A list of port driver parameters and their current values.
  - Parameter – The driver parameter's name.
  - Value – The driver parameter's value.
  - Temporary – Indicates if the value is temporary.
  - Range – The range of acceptable values for the driver parameter.
  - Default – The driver parameter's default value.
  - Activation Requirements – The steps required to activate the changed value of the driver parameter.
  - Description – The driver parameter's description.

To change the driver parameters for a port using the **Value** field, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the port for which you want to change the driver parameters.
3. Select the **Driver Parameters** tab (Figure 69).
4. In the driver parameter table, click the **Value** field of a parameter that you want to change. The range for the value is displayed. Enter a value 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). You can enter a hexadecimal value without the 0x. For example, if you enter ff10, this value is interpreted and displayed as 0xff10.
5. 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 **Temporary** box. This option is available only for dynamic parameters.
6. 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.
7. Click **Apply**.

To set a port parameter value to the corresponding host parameter value, click **Globals**. All parameter values are set the same as the global, or host, values.

To apply the global values, click **Apply**.

If you changed parameters but did not click **Apply,** and you want to restore the parameters to their last saved values, click **Restore**.

To reset all parameter values to their default (factory) values, click **Defaults**.

To save driver parameters to a file:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the FCoE port for which you want to change the driver parameters.
3. Select the **Driver Parameters** tab (Figure 69).
4. Click **Export** to create and save a desired port parameter configuration. Each definition is saved in a comma-delimited file with the following format:

   `<parameter-name>=<parameter-value>`
5. Click **Apply** to apply your configuration changes.

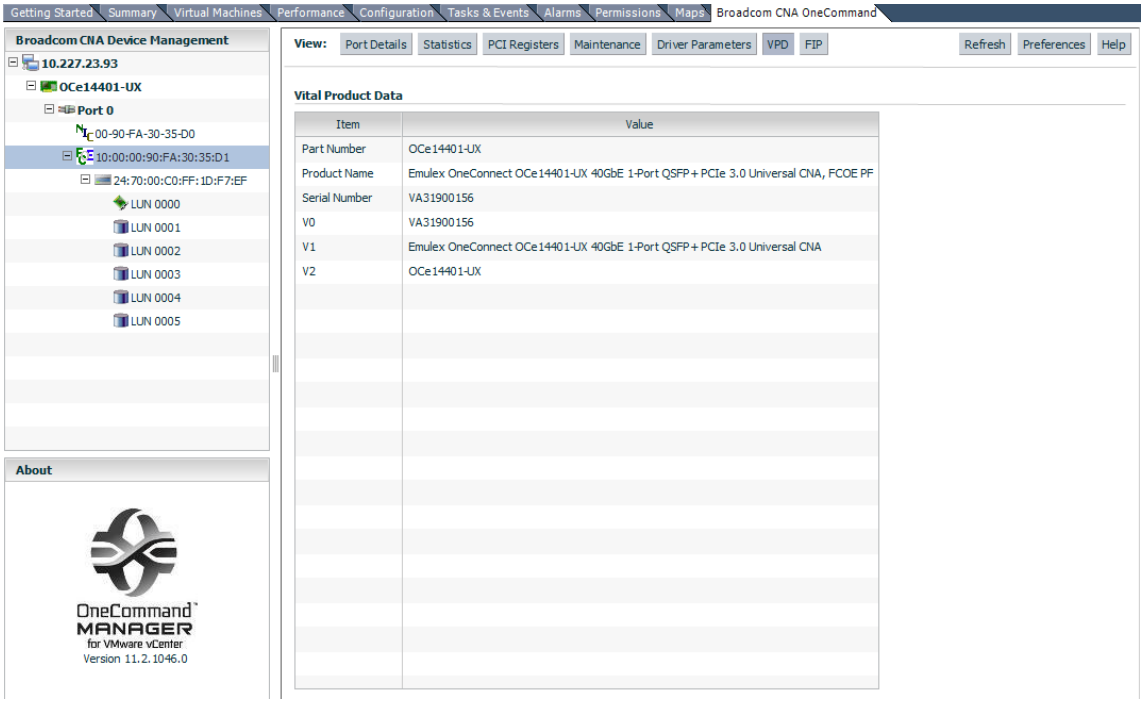
6.2.6 Viewing FCoE Port Vital Product Data (VPD)

The **VPD** tab shows vital product data (if available) for the selected FCoE port such as the product name, part number, serial number, and so on.

To view VPD information for a port, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the port whose VPD information you want to view.
3. Select the **VPD** tab (Figure 70).

Figure 70 Port VPD Tab



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

- Part Number – The adapter's part number.
- Product Name – The product information about the selected adapter port.
- 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 V*x* (where *x* is a number).

NOTE Some adapters may show additional VPD information such as EC and MN.

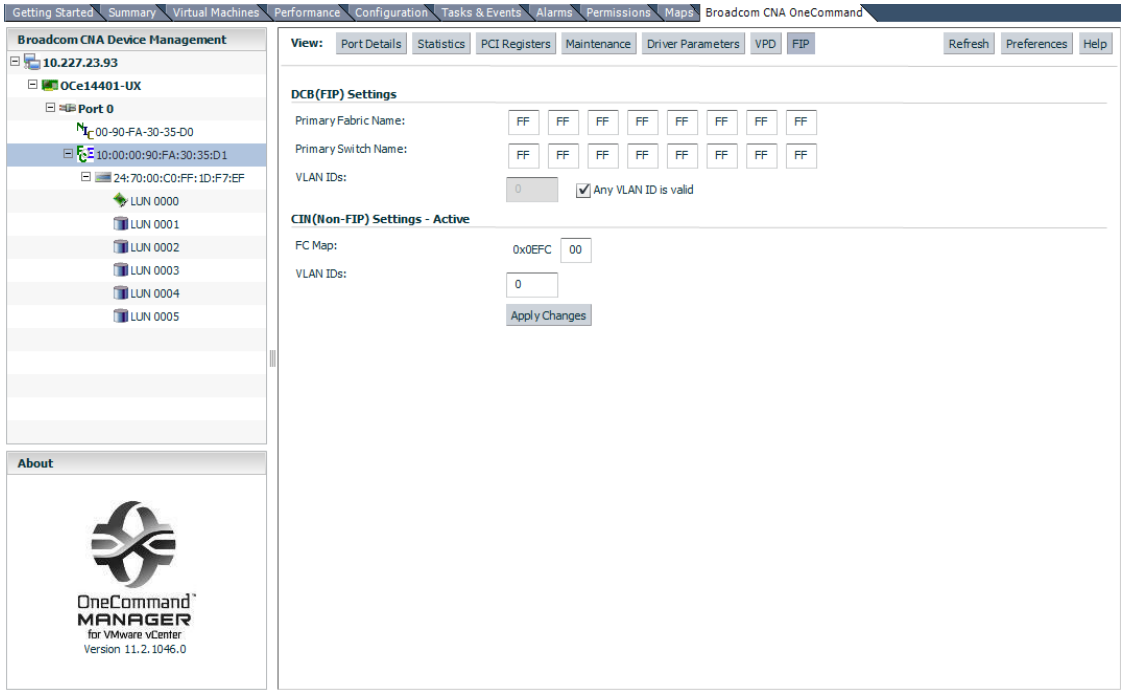
6.2.7 Configuring FIP

The **FIP** tab allows you to view configure FIP for an FCoE adapter port.

To configure FIP, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the FCoE adapter port whose FIP properties you want to configure.
3. Select the **FIP** tab (Figure 71).
4. Set the parameters you want and click **Apply Changes**.

Figure 71 FCoE Port FIP Tab



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

- DCB (FIP) Settings area: DCB settings are only applicable when the DCBX Mode in the **DCB** tab is set to DCB. See Section 8.3, Configuring DCB Parameters for NIC + FCoE + iSCSI Adapter Ports, for more information.

  NOTE If the DCB (FIP) Settings area has Active in its title, this indicates that DCB (FIP) is the current DCBX mode (DCBX mode is set from the **DCB** tab). If CIN (Non-FIP) is the current DCBX mode, Active is in the CIN (Non-FIP) Settings area.

  - Primary Fabric Name – Indicates the FC fabric's WWN to which to connect. If the Primary Fabric Name is wild (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 (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 to 4095 and supports wild card values if the Any VLAN ID is valid checkbox is selected.
  - Any VLAN ID is valid checkbox – When checked, the VLAN ID field of the FCoE forwarder can be any valid value.
- CIN Settings area:
  - FC Map – Enter the bit value that completes the FPMA.
  - VLAN ID – Determines the VLAN where the adapter FCoE services are available. It can have a value from 0 to 4095.

5. Click **Apply Changes** to set the new FIP parameters.

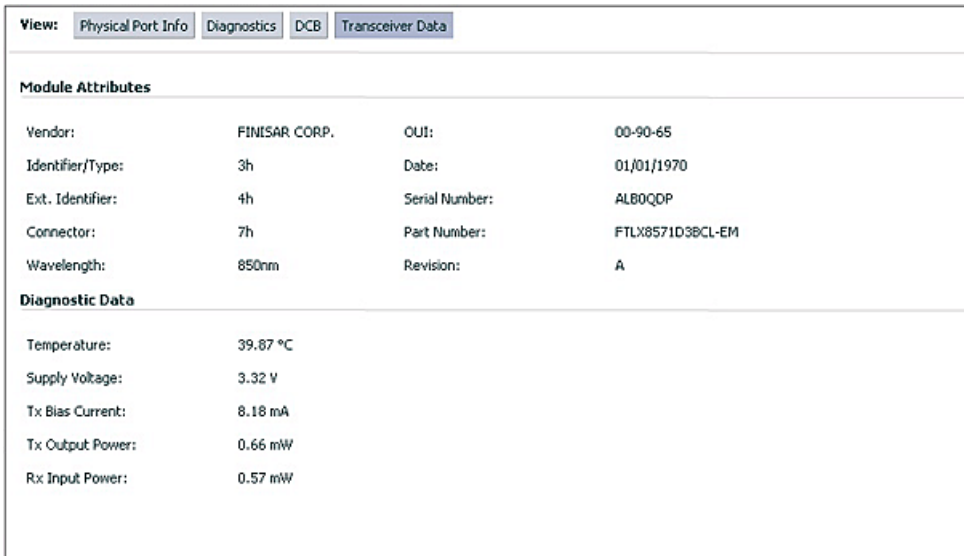
6.2.8 Viewing FCoE Port Transceiver Information

The **Transceiver Data** tab shows transceiver information such as vendor name, serial number, and part number. If the adapter/transceiver does not support some or all of the transceiver data, the fields display N/A.

To view transceiver information for an FCoE port, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the FCoE port whose transceiver information you want to view.
3. Select the **Transceiver Data** tab (Figure 72).

Figure 72 Transceiver Data Tab



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

- Module Attributes area:
  - Vendor – The name of the vendor.
  - Identifier/Type – A 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 Organizationally Unique Identifier. 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.

6.2.9 Viewing FC/FCoE Target Information

When you select a port target associated with an FC/FCoE adapter from the Broadcom CNA Device Management tree-view, the **Target Information** tab shows information associated with that target.

To view FC/FCoE target information, perform the following steps:

1. Select a host in the vSphere console tree-view, and select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the FCoE target port whose information you want to view. **The Target Information** tab appears (Figure 73).

Figure 73 Target Information Tab



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

- 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.

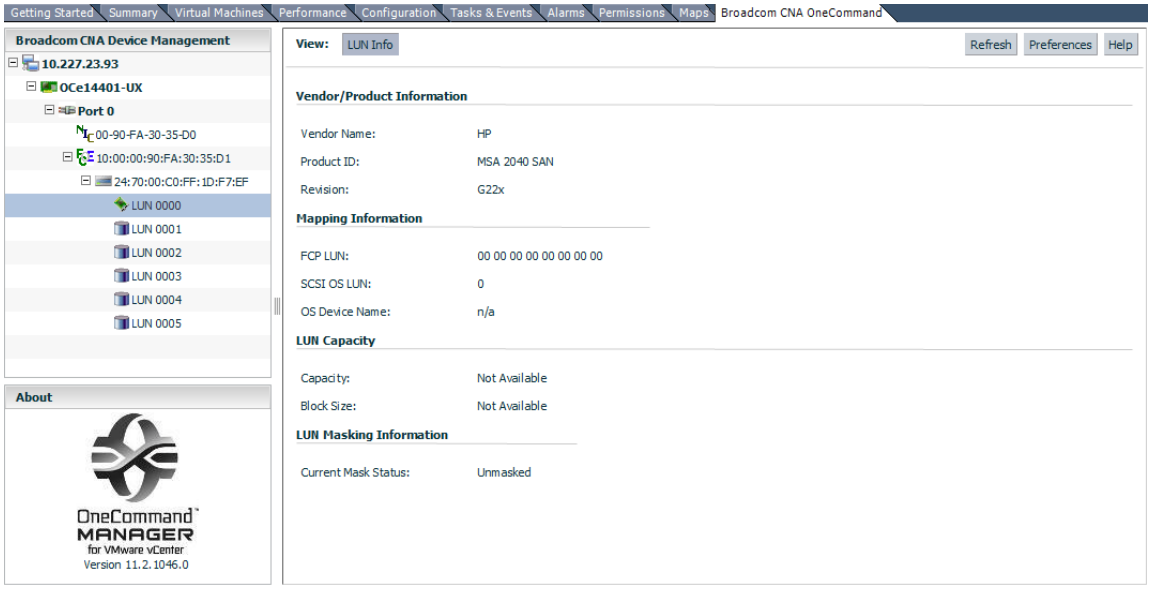
6.2.10 Viewing FCoE LUN Information

When you select a LUN associated with an FCoE adapter from the Broadcom CNA Device Management tree-view, the LUN Info pane shows information associated with that LUN.

To view LUN information, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the FCoE LUN whose information you want to view. The corresponding LUN Info pane is displayed (Figure 74).

Figure 74 LUN Info Tab



The following FCoE **LUN Info** tab fields are displayed:

- Vendor Product Information area:
  - Vendor Name – The name of the vendor of the LUN.
  - Product ID – The vendor-specific ID for the LUN.
  - Revision – The vendor-specific revision number for the LUN.
- Mapping Information area:
  - FCP LUN – The FC identifier used by the adapter to map to the operating system 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: LUN capacity information is only provided when 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 Information area:
  - Current Mask Status - The current status is masked or unmasked.

6.3 iSCSI

This section pertains to viewing iSCSI information and managing iSCSI ports.

6.3.1 Viewing iSCSI Port Information

When you select an iSCSI port from the Broadcom CNA Device Management tree-view, the iSCSI **Port Information** tab shows general attributes associated with the selected iSCSI port.

To view general iSCSI port information, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.

2. In the Broadcom CNA Device Management tree-view, select the iSCSI port for which you want to view information.
3. Select the **Port Information** tab (Figure 75).

Figure 75 iSCSI Port Info Tab



The following **iSCSI Port Info** tab fields are displayed:

- Driver Name – The iSCSI driver file name.
- Driver Version – The iSCSI driver version.
- MAC Address – The iSCSI MAC address currently assigned to the port.
- 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 – The maximum transmission unit for Ethernet traffic.
- Default ETO – The default extended timeout.
- Max CDB Length – The maximum SCSI command descriptor block size.
- Subsystem 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 to which the port is mapped.
- TCP/IP Info area:
  - DHCP Enabled – The DHCP authentication status of the selected port.
  - IP Address – The IP address for the NIC port.
  - Subnet Mask – The subnet mask for the NIC port.
  - Gateway Address – The gateway address for the NIC port.
  - VLAN Enabled – The VLAN enabled state for the iSCSI interface.

- VLAN ID – The VLAN identifier to use, from 0 to 4094 (only valid when VLAN is enabled). A value of 0 indicates the VLAN is disabled.
- VLAN Priority – The VLAN priority for the iSCSI interface.

- Network Info area:
  - Link Status – The status of the link on the selected port.
  - LinkDown Reason – The reason the link is down on the selected port.
  - Port Speed – The current port speed of the iSCSI port.
  - Min Bandwidth – The minimum SoftCap bandwidth assigned.
  - Max Bandwidth – The maximum SoftCap bandwidth assigned.

6.3.1.1 Enabling iSCSI Boot

To enable iSCSI boot, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the iSCSI port for which you want to view information.
3. Select the **iSCSI Port Info** tab (Figure 75).
4. To enable iSCSI boot, check **Enable iSCSI Boot**. To disable iSCSI boot, clear **Enable iSCSI Boot**.

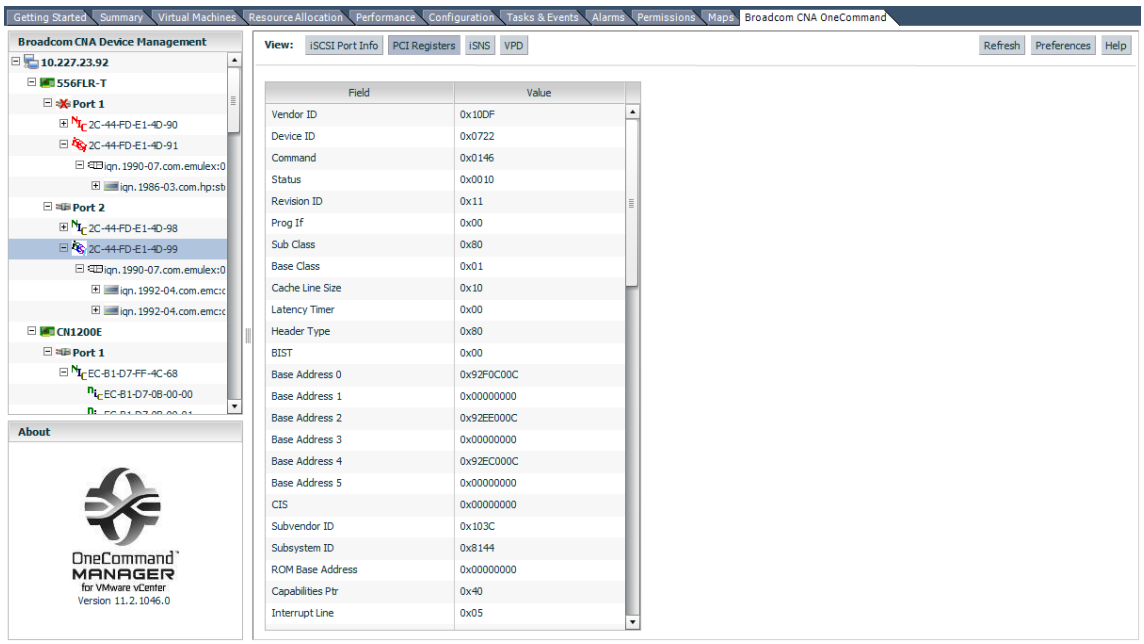
6.3.2 Viewing PCI Registers for an iSCSI Port

The **PCI Registers** tab contains PCI information about the selected iSCSI port.

To view PCI registers for an iSCSI port, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the iSCSI port for which you want to view PCI information.
3. Select the **PCI Registers** tab (Figure 76).

Figure 76 iSCSI Port, PCI Registers Tab



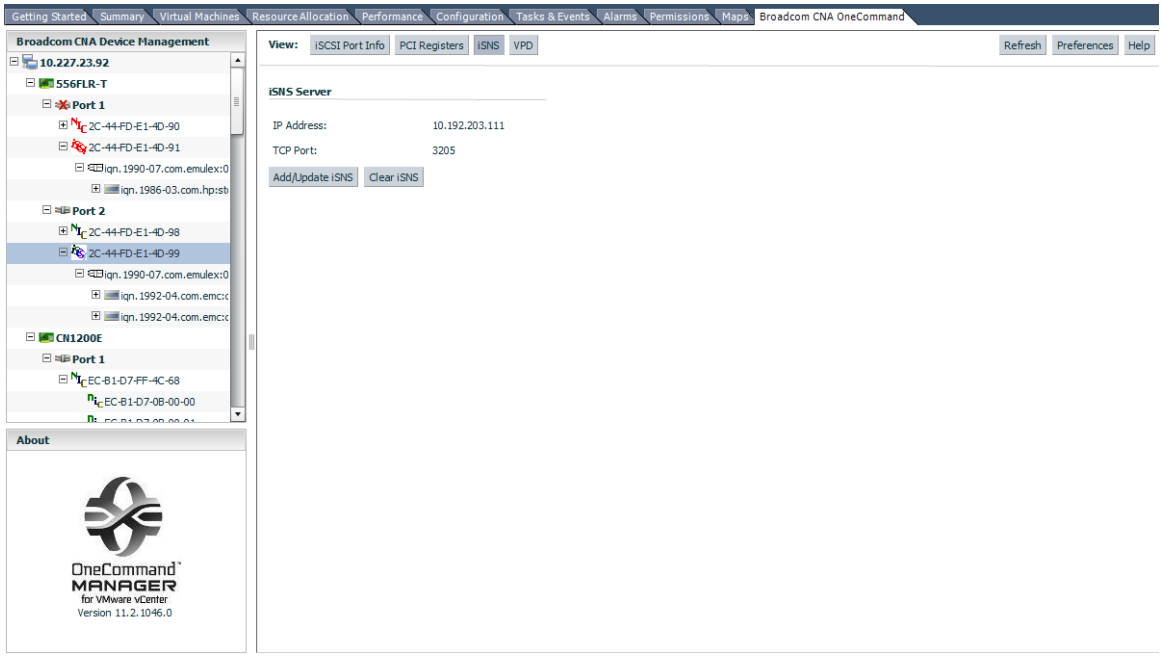
6.3.3 Viewing and Changing iSCSI ISNS Server Information

If the server has been set up as an ISNS, the **iSNS** tab shows the IP address and TCP Port for the selected iSCSI port.

To view ISNS server information, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the iSCSI port for which you want to view ISNS information.
3. Select the **iSNS** tab (Figure 77).

Figure 77 iSCSI Port, iSNS Tab



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

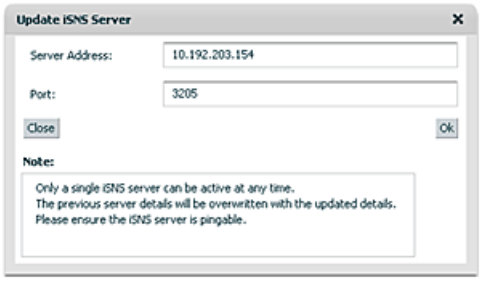
- iSNS Server area:
  - IP Address – The IP address of the ISNS server.
  - TCP Port – The TCP port number for the connection to the ISNS server.

6.3.3.1 Changing the ISNS Server

To change the ISNS server, perform the following steps:

1. On the **iSNS** tab, click **Add/Update iSNS**. The Update iSNS Server dialog is displayed (Figure 78).

Figure 78 iSCSI Port, Update iSNS Server Dialog



2. Enter the new server address and Port ID.
3. Click **OK** to commit the new ISNS server. The new ISNS information is displayed on the **iSNS** tab (Figure 77). The new or updated ISNS server is also queried for iSCSI targets. Any discovered targets are added to the i**SCSI Target Discovery** tab.

6.3.3.2 Disabling the ISNS Server

To disable the ISNS server, click **Clear iSNS** on the **iSNS** tab (Figure 77).

6.3.3.3 Configuring ISNS for iSCSI Target Discovery

An ISNS maintains a database of storage network elements that can be queried by other hosts within the SAN. iSCSI storage devices in particular can register targets with the ISNS for efficient discovery by iSCSI clients. Use the **iSCSI SNS** tab to configure the ISNS server or to discover the server. For more information about the **iSCSI iSNS** tab, refer to Section 6.3.3.1, Changing the ISNS Server.

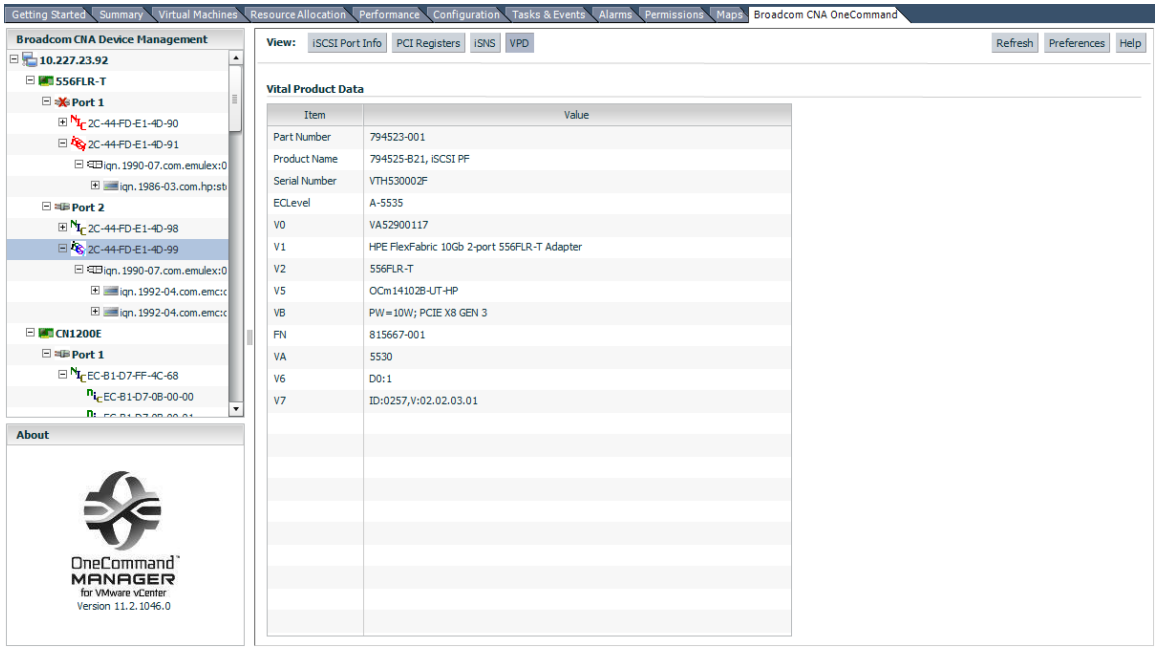
6.3.4 Viewing iSCSI Port VPD

The **VPD** tab shows vital product data (if available) for the selected iSCSI port such as the product name, part number, serial number, and so on.

To view VPD information of an iSCSI port, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the iSCSI port for which you want to view VPD information.
3. Select the **VPD** tab (Figure 79).

Figure 79 iSCSI Port VPD Tab



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

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

NOTE Some adapters may show additional VPD information such as EC and MN.

6.3.5 Viewing iSCSI Initiator Statistics

The **iSCSI Statistics** tab shows information for the selected iSCSI port.

To view iSCSI initiator statistics, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the iSCSI initiator for which you want to view Statistic information.
3. Click the **iSCSI Statistics** tab (Figure 80).

Figure 80 iSCSI Initiator, iSCSI Statistics Tab



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

- Node Roles – The node role for this iSCSI initiator.
- Portal Count – The number of rows in the iscsiPortaltypeTable that are currently associated with this iSCSI instance.
- Node Count – The number of rows in the iscsiNodetypeTable that are currently associated with this iSCSI instance.
- Session Count – The number of rows in the iscsiSessiontypeTable that are currently associated with this iSCSI instance.
- Session Failure Count – The number of times a session belonging to this port has failed.
- Last Session Failure Type – The type of failure encountered in the last session failure.
- Last Session Remote Node – The iSCSI name of the remote node from the failed session.
- Session Digest Errors – The count of PDUs that were received on the session and contained header or data digest errors.
- Session Connection Timeout – The count of connections within this session that have been terminated due to a timeout.
- 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 of 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 a normal reason code.
- 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 via SNMP.
- Portal Role – The role of a portal. A portal can operate as a target portal and/or an initiator portal.

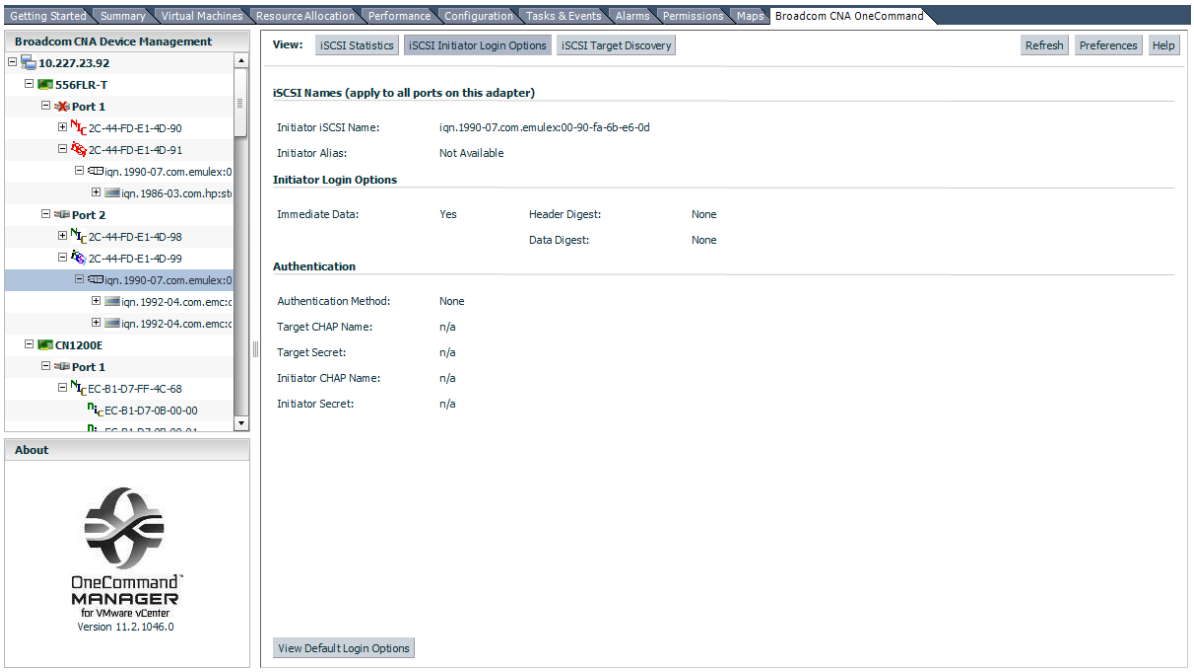
6.3.6 Viewing iSCSI Initiator Login Options

The **iSCSI Initiator Login Options** tab shows the set of login options used by the initiator when logging into a target.

To view iSCSI initiator login information, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the iSCSI initiator for which you want to view initiator login information.
3. Click the **iSCSI Initiator Login Options** tab (Figure 81).

Figure 81 iSCSI Initiator Logon Options Tab



The following **iSCSI Initiator Login Options** tab fields are displayed:

- iSCSI Names area:

- Initiator iSCSI Name – The iSCSI qualifier name of the initiator. The initiator name must be followed by a colon and the port number.
- Initiator Alias – An optional non-unique string used to identify the initiator.

- Initiator Login Options area:
  - Immediate Data – If set to **Yes**, the initiator can append unsolicited data to an iSCSI command.
  - Header Digest – If set to **CRC32C**, the integrity of an iSCSI PDU's header segments are protected by a CRC32C checksum. If set to **None**, the PDU's header segments are not protected.
  - Data Digest – If set to **CRC32C**, the integrity of an iSCSI PDU's data segment is protected by a CRC32C checksum. If set to **None**, the PDU's data segments are not protected.
- Authentication area:
  - Authentication Method
    - One-Way CHAP requires only that the authenticator (iSCSI target) authenticate the iSCSI initiator.
    - Mutual CHAP requires that both the iSCSI target and iSCSI initiator authenticate each other.
    - None indicates that no authentication is performed.
  - Target CHAP Name – The iSCSI login name sent by the initiator to the target for authentication. This name is required for both One-Way CHAP and Mutual CHAP authentication. The target CHAP name is also known as the username.
  - Target Secret – The iSCSI login secret sent by the initiator to the target for authentication. A target secret is required for both One-Way CHAP and Mutual CHAP authentication.
  - Initiator CHAP Name – The iSCSI login name sent by the target to the initiator for authentication. The initiator CHAP name is only required for Mutual CHAP authentication. The initiator CHAP name is also known as the username.
  - Initiator Secret – The iSCSI login secret sent by the target to the initiator for authentication. The initiator secret is only required for Mutual CHAP authentication.

To view iSCSI initiator default login options, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the iSCSI initiator for which you want to view Initiator Login information.
3. Click the **iSCSI Initiator Login Options** tab (Figure 81).
4. Click **View Default Login.** The Initiator Default Login Options window is displayed (Figure 82).

Figure 82 iSCSI Initiator Default Login Options Window



The following Initiator Default Login Options window fields are displayed:

- InitialR2T – The initial request to transmit. When set to **Yes**, the initiator has to wait for the target to solicit SCSI data before sending it. When set to **No**, it allows the initiator to send a burst of unsolicited FirstBurstLength bytes.
- ImmediateData – If set to **Yes**, allows the initiator to append unsolicited data to a command.

- 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.
- DefaultTimeTo Retain – 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.
- 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.
- DataPPDUInOrder – The order of data PDUs within a sequence.
- DataSequenceInOrder – The order between sequences.
- MaxRecDataSegmentLength – 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.

6.3.7 Viewing iSCSI Target Discovery Information

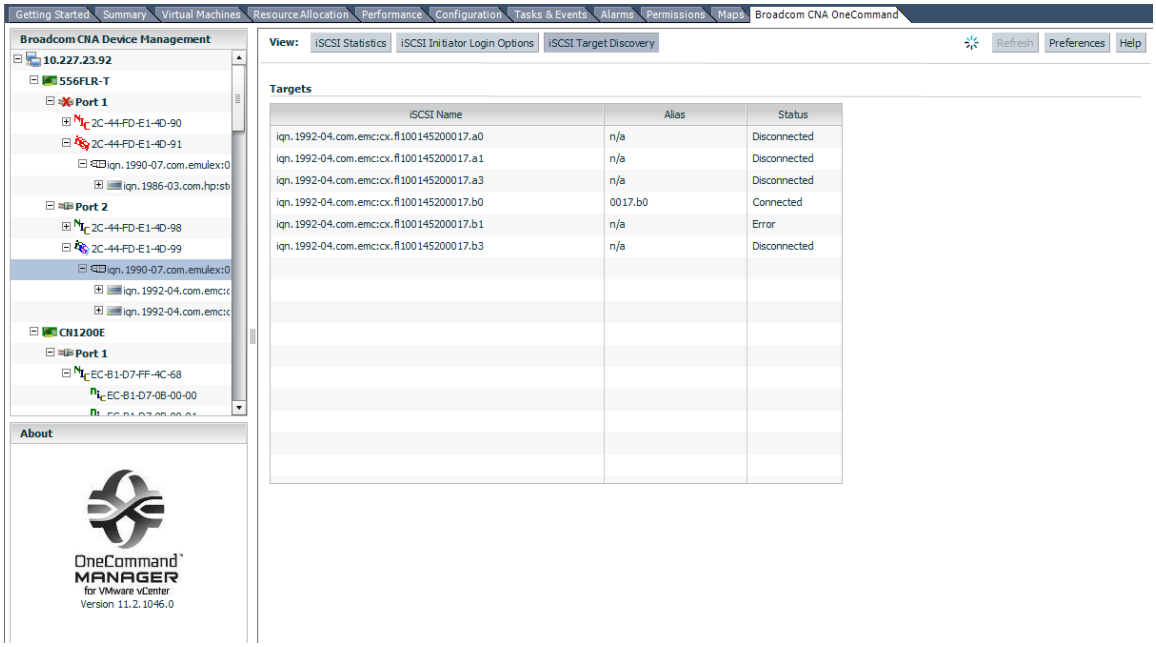
The **iSCSI Target Discovery** tab shows all currently discovered targets. Targets on this tab come from one of three sources:

- The target was added manually.
- The target was discovered via a target portal.
- The target was found through an ISNS server query.

To view iSCSI target discovery information, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the iSCSI initiator for which you want to view target discovery information.
3. Click the **iSCSI Target Discovery** tab (Figure 83).

Figure 83 iSCSI Target Discovery Tab



The following **iSCSI Target Discovery** tab fields are displayed:

- iSCSI Name – The iSCSI name assigned to the target.
- Alias – The iSCSI alias assigned to the target. This is assigned at the target portal, not by OneCommand Manager for VMware vCenter.
- Status – The status column indicates if the target is connected or disconnected.

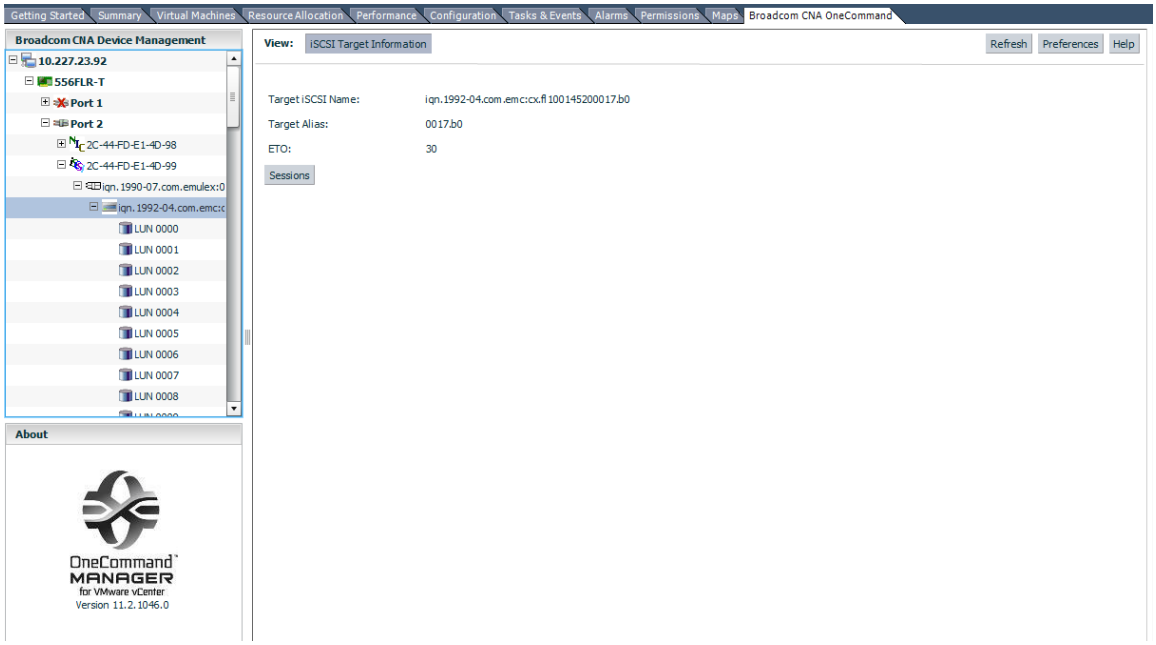
6.3.8 Viewing iSCSI Target Information

The **iSCSI Target Information** tab shows information about the target. You can also view information about a target session.

To view iSCSI target discovery information, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the iSCSI target for which you want to view target discovery information. The **iSCSI Target Information** tab is displayed (Figure 84).

Figure 84 iSCSI Target Information Tab



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

- 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 OneCommand CNA Manager for VMware vCenter.
- ETO – The extended timeout for the target. The driver ensures that I/Os are not timed out until this time has expired (from the time that the target stopped responding).

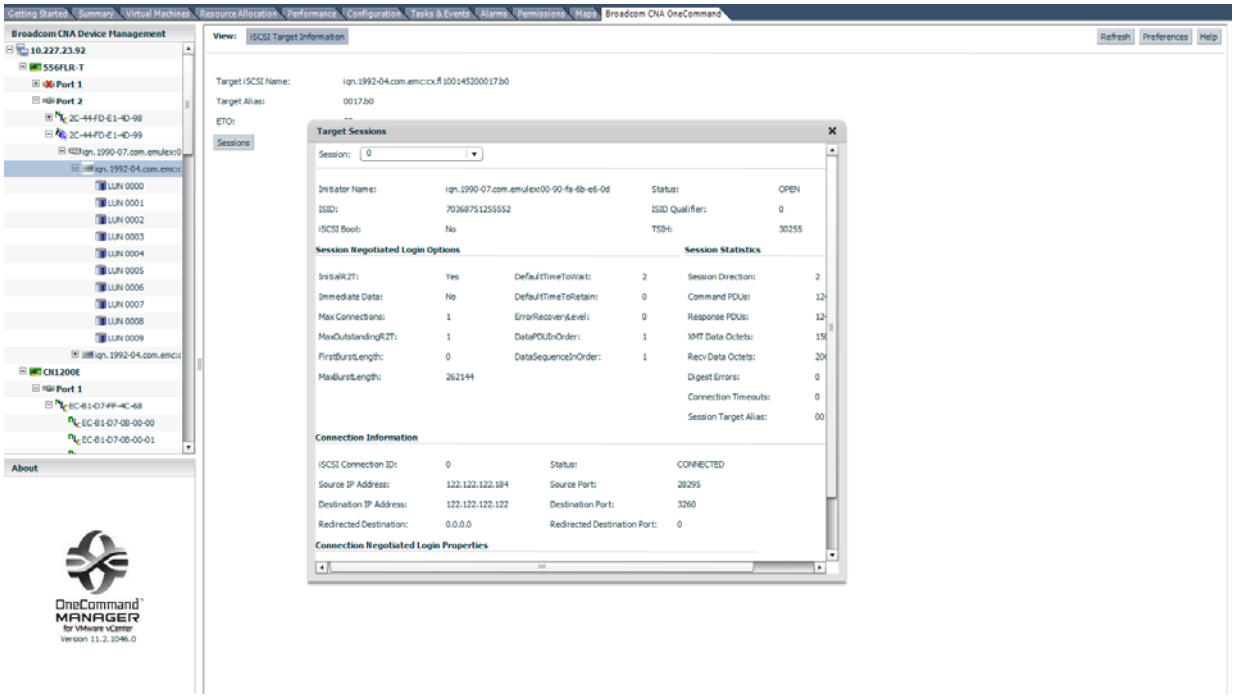
6.3.8.1 Viewing Target Session Information

There may be multiple active sessions for a connected target.

To view iSCSI target session information, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the iSCSI target for which you want to view target discovery information. The **iSCSI Target Information** tab is displayed (Figure 84).
3. Click **Sessions (**Figure 85). You can toggle between sessions by using the drop-down menu in the upper left corner of the window.

Figure 85 iSCSI Target Sessions Window



The following Target Sessions window fields are displayed:

- Initiator Name – The initiator named used to log into the session.
- ISID – The initiator session identifier (unique for each session).
- iSCSI Boot – If set to **Yes**, indicates a boot session. Logout is not possible from a boot session.
- Status – The session status (logged in, login in progress, login failed, recovery, or unknown).
- 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.
- 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.
  - MaxConnections – 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 can 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** - recovery only by session restart.

- **1** - recovery by reissuing commands, data, or status.
- **2** - connection failure recovery.

— DataPDUInOrder – The order of PDUs within a sequence.
— DataSequenceInOrder – The order between sequences.

- Session Statistics area:
  — Session Direction – The direction of the 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 which 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 Properties area:
  — 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 – When set to **CRC32C**, the integrity of an iSCSI PDU's header segments is protected by a CRC32C checksum.
  — Data Digest – When 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.

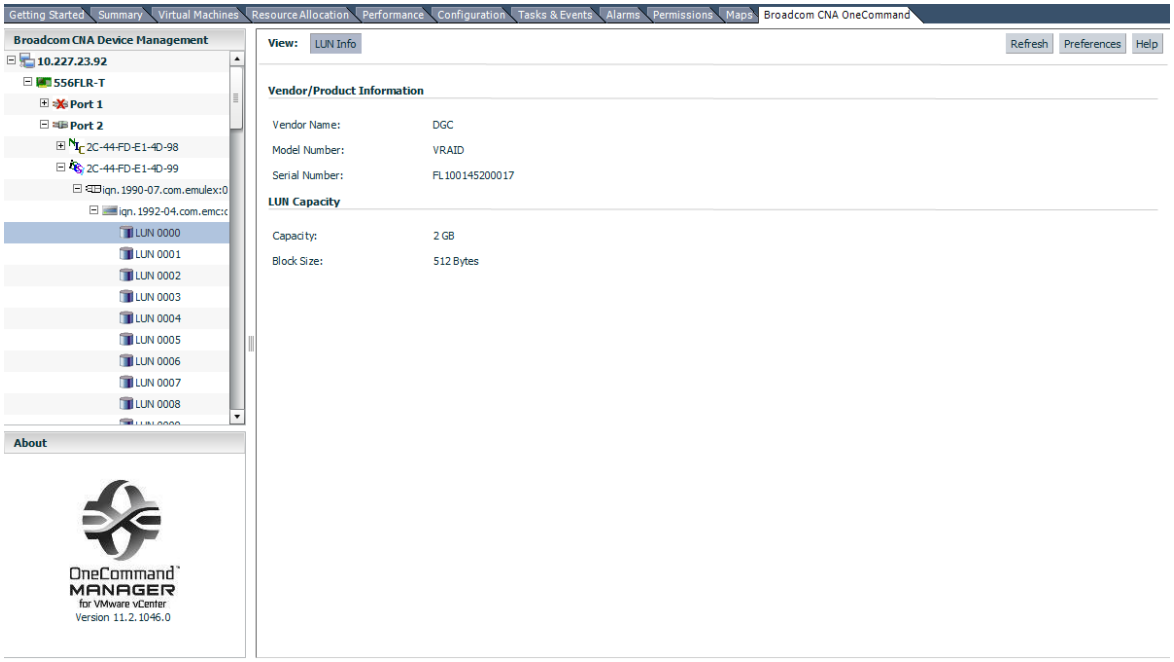
6.3.9 Viewing iSCSI LUN Information

The **iSCSI LUN Info** tab shows basic vendor, product and LUN masking information.

To view iSCSI LUN information, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the LUN for which you want to view information. The **LUN Info** tab is displayed (Figure 86).

Figure 86 iSCSI LUN Info Tab



The following **iSCSI LUN Info** tab fields are displayed:

- Vendor/Product Information area
  - 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.
  - Serial Number – The vendor's serial number for the LUN.
- LUN Capacity area:
  - Capacity – The unformatted size of the LUN, in megabytes.
  - Block Size – The length of a logical unit block in bytes.

Chapter 7: Updating Firmware

OneCommand CNA Manager for VMware vCenter enables you to update firmware for a single adapter or simultaneously for multiple adapters.

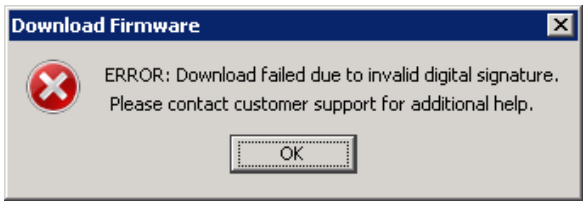
The submitted firmware update job can be tracked in the VMware tasks window.

Multiple firmware update jobs can be submitted for different adapters on the same or different ESXi hosts simultaneously to OneCommand CNA Manager for VMware. However, only a single job is processed on a given ESXi host. The remaining jobs on that host will be queued and processed sequentially.

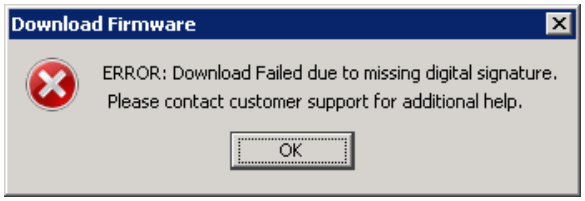
The firmware update progresses in the background until all the jobs are completed. During this period, you can still browse through the other tabs. However, if you start a firmware update and log out from the vSphere console before the firmware update is completed, all pending jobs fail.

NOTE If a secure version of firmware (version 11.0 or later) is installed on an OCe14000B-series adapter and you want to update to an earlier unsecured version of firmware, you must remove the secure firmware jumper block before performing the update. See the installation guide for the adapter for more information.

If you attempt to update unauthenticated firmware for a secure OCe14000B-series adapter, the following error message will be displayed:



If you attempt to update unsecured firmware for a secure OCe14000B-series adapter, the following error message will be displayed:



Contact customer support for more information.

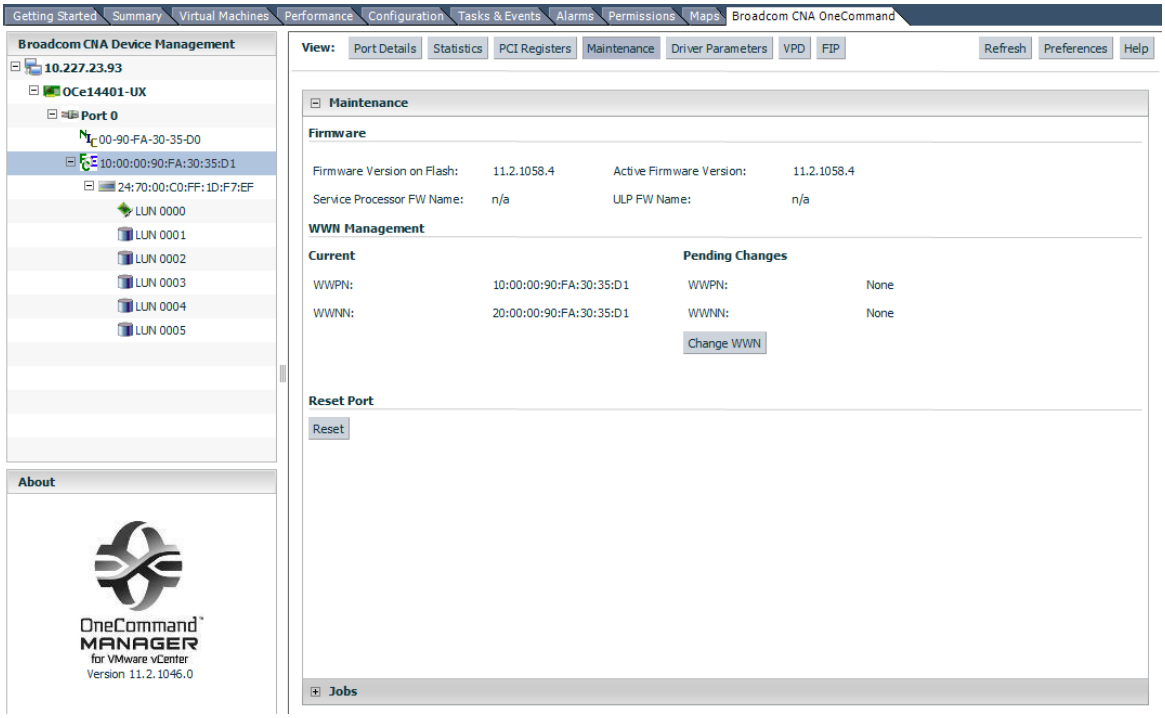
7.1 Updating Firmware on a Single Adapter

Firmware is updated for the entire adapter and not for individual ports.

To update firmware for an adapter, perform the following steps:

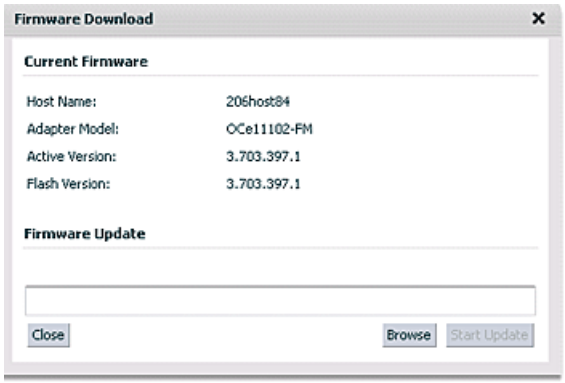
1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the adapter for which you want to update firmware information.
3. Select the **Maintenance** tab (Figure 87).

Figure 87 FCoE Maintenance Tab



4. Click **Update Firmware**. If a warning message appears, click **Yes**. The Firmware Download dialog is displayed (Figure 88).

Figure 88 Firmware Download Dialog



5. Click **Browse** and navigate to the unzipped, extracted image file you want to download.
6. On the browse window, select the file and click **OK**.
7. Click Start **Update**.

   When the update begins, the Jobs window is displayed (Figure 93). A status message in the Active Job list shows the progress of the download. The ports on which the firmware is being downloaded have the status **Job is in progress**, the other ports display the **Waiting in queue to start** status.

   When the download has completed, the entry moves to the Completed Jobs window. The status message field in the Completed Jobs window shows the status of the completed job.

   You can monitor jobs submitted through the OneCommand CNA Manager for VMware vCenter Command Line Interface. If you start a firmware update and log out from the vSphere console before the firmware update is

completed, all pending jobs fail. The firmware update progresses in the background until all the jobs are completed. During this period, you can still browse through the other tabs. The firmware update job submitted can also be tracked in the VMware tasks window.

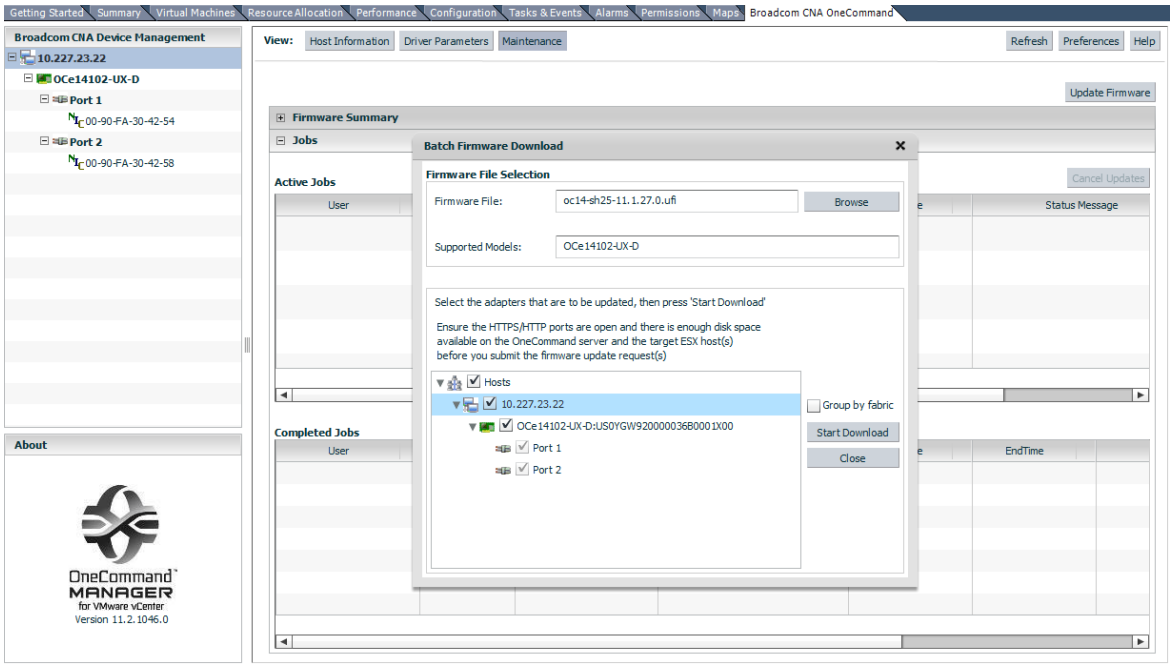
8. Open **Firmware Summary** and the updated firmware information for the selected adapter is displayed.

7.2 Updating Firmware on Multiple Adapters in a Host

To update firmware for multiple adapters on a single host, perform the following steps:

1. Select a host in the vSphere console tree-view and if applicable, select the **Broadcom CNA OneCommand** tab.
2. Select the **Maintenance** tab and click **Update Firmware**. If a warning message appears, click **Yes**. The Batch Firmware Download dialog appears (Figure 89).

Figure 89 Batch Firmware Download Dialog

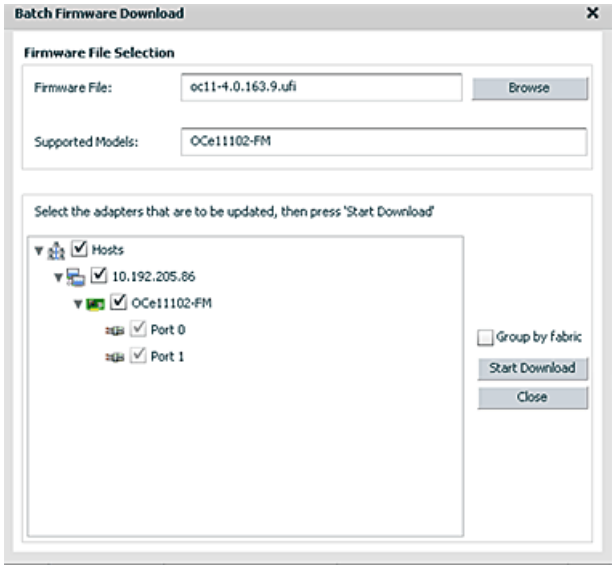


3. On the Batch Firmware Update Information pane, click **Update Firmware** (Figure 92). A Batch Firmware Download dialog appears with blank fields.

NOTE Do not select a particular tree element for this operation.

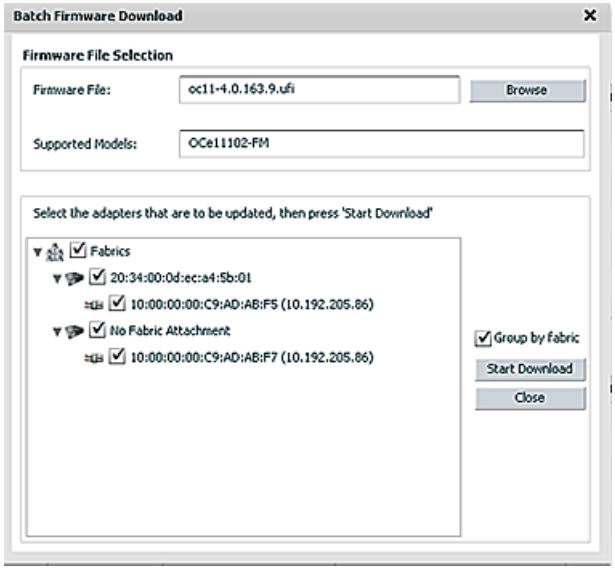
4. Click **Browse**, and a search dialog appears. Select the file you want to use and click **OK**. A status message appears indicating that OneCommand CNA Manager for VMware vCenter is searching for compatible adapters. After compatible adapters are found, the following is displayed in the Batch Firmware Download dialog (Figure 90):
   - The Firmware File field shows the selected image file name.
   - The Supported Models field shows a list of all adapter models that are compatible with the selected image file.
   - The set of compatible adapters appears in the dialog's tree-view.
5. To view the compatible adapters in host-centric mode, make sure the **Group by Fabric** checkbox is cleared.

Figure 90 Batch Firmware Download Dialog, Host-Centric View



6. To view the compatible adapters in a fabric-centric mode, select the **Group by Fabric** checkbox.

Figure 91 Batch Firmware Download Dialog, Fabric-Centric View



The tree-view shows all adapters and their corresponding hosts or fabrics for which the selected firmware file is compatible. Use the checkboxes next to the entries to select or deselect a host, fabric, adapter or port (if the adapters where the individual port can have new firmware downloaded).

NOTE iSCSI and NIC ports are not attached to any fabric, they are not listed in the Batch Firmware Download dialog.

7. Make your selections and click **Start Download**. A message prompting you to confirm the firmware update appears, click **Yes**.

7.3 Performing a Batch Firmware Update in Cluster View

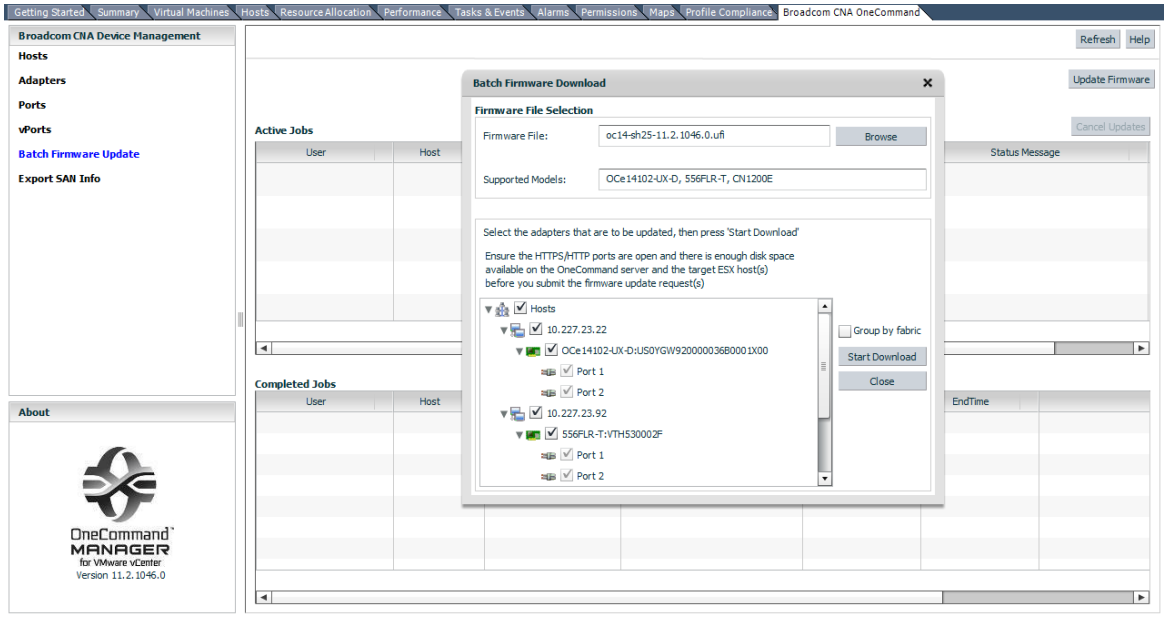
Before you can perform a batch update, you must download the firmware file from http://www.broadcom.com and extract it into a directory on your local drive.

NOTE On the system where OneCommand CNA Manager for VMware vCenter is installed, make sure the port numbers configured during the installation are open and dedicated to the OneCommand CNA Manager for VMware vCenter server only. No other service should be listening on this port.

To perform a batch firmware update on a host, perform the following steps:

1. Select a cluster in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. From the Broadcom CNA Device Management options, select **Batch Firmware Update**. The Batch Firmware Download dialog appears (Figure 92).

Figure 92 Batch Firmware Update Information

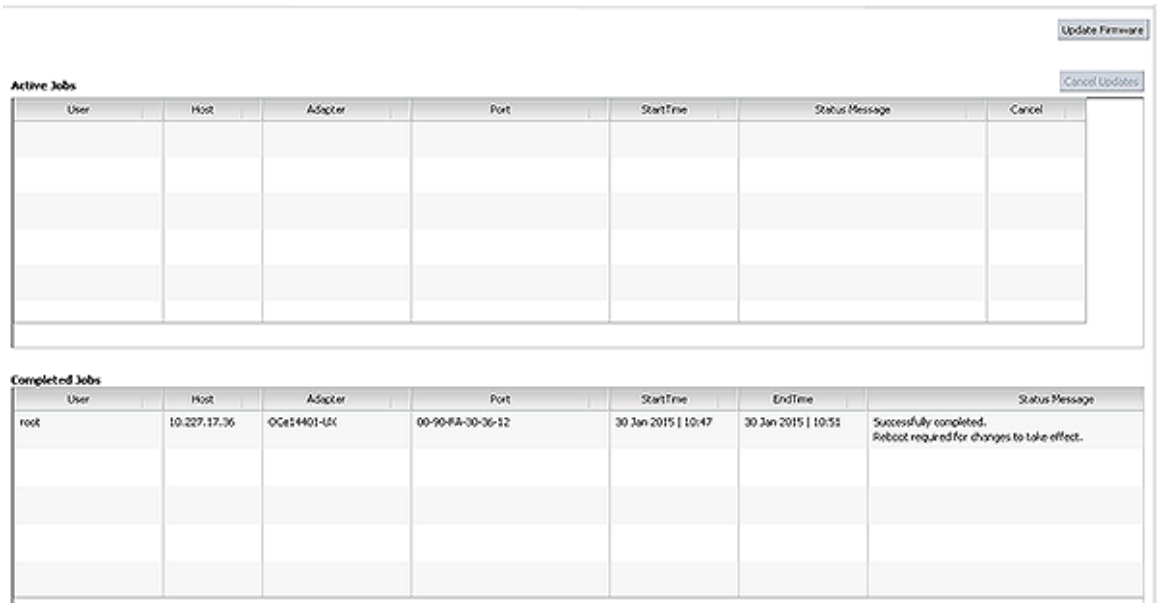


3. Select the hosts on which you want to update firmware in batch mode.
4. Click **Start Download** to install firmware on multiple adapters in a single step. Batch firmware loading is restricted to a single firmware file and to all accessible adapters for which that file is compatible.

7.3.1 Jobs Window

When the download begins, the Jobs window is displayed. A status message in the Active Job list shows the progress of the firmware download. The ports on which the firmware is being downloaded have the status **Job is in progress**, the other ports have **Waiting in queue to start**. When the download is completed, the entry moves to the Completed Jobs window. The status message field in the Completed Jobs window shows the status of the completed job (Figure 93).

Figure 93 Jobs Window



The following Jobs window fields are displayed:

- Active Jobs column:
  - User – The user who updated the firmware.
  - Host – The IP address of the host to which the adapter/port belongs.
  - Adapter – The model of the adapter selected for the firmware update request.
  - Port – The Port WWN for an FC adapter. This field is blank for a CNA.
  - Start Time – The time the job is submitted.
  - Status Message – The status of the job. This field also indicates if a reboot is required to activate the firmware.
- Completed Jobs column:
  - User – The user who updated the firmware.
  - Host – The IP address of the host to which the adapter/port belongs.
  - Adapter – The model of the adapter.
  - Port – The port WWN for an FC adapter. This field is blank for a CNA.
  - Start Time – The time the job is submitted.
  - End Time – The time the job is completed.
  - Status Message – The status of the job.

To cancel the jobs with the status "Waiting in queue to start", click **Cancel Updates**.

Chapter 8: DCB Configuration

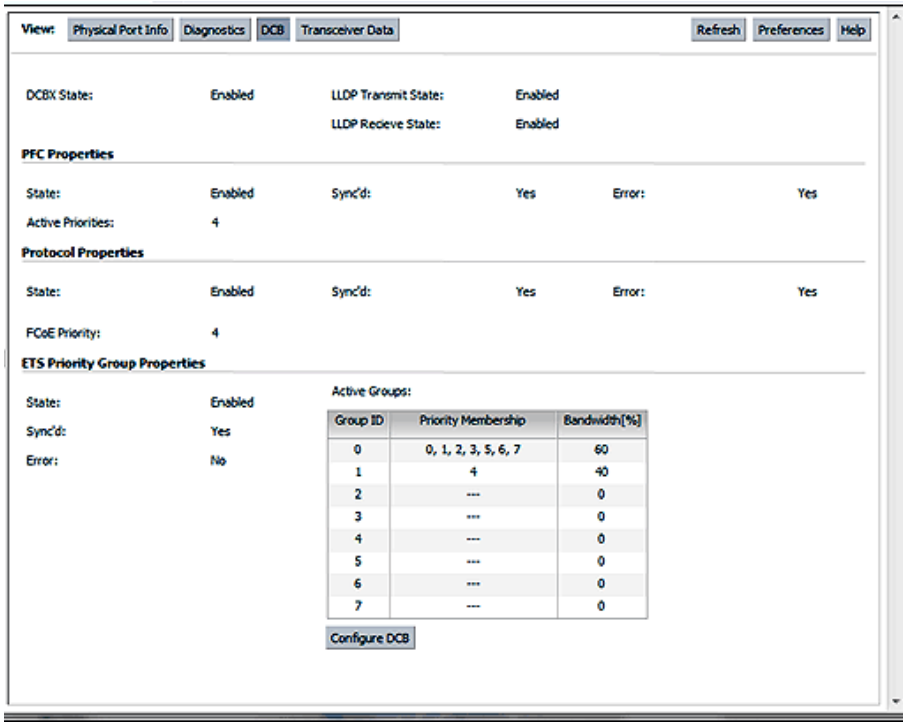
The following section describes how to configure DCB parameters on OneConnect adapter ports.

8.1 Configuring DCB Parameters for NIC + FCoE Adapter Ports

To view the DCB parameters for a NIC + FCoE port, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the port for which you want to view DCB parameters.
3. Select the **DCB** tab (Figure 94).

Figure 94 DCB Tab (NIC + FCoE)



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

- 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 receive and transmit states must be enabled. If either the receive or transmit state 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 receive and transmit states must be enabled. If either the receive or transmit states is disabled, DCBX is disabled.
- PFC Properties area:
  - State – Enabled indicates that there is flow control in both directions (receive and transmit).
  - Sync'd – If Yes, the PFC priorities have been set by the peer. This parameter cannot be set.

- Error – Indicates whether an error occurred during the configuration exchange with the peer. Error is also set to Yes when the Compatible method fails.
- Active Priorities – Lists the number of priorities with PFC set to enabled.

- Protocol Properties area:
  - State – Enabled indicates that there is flow control in both directions (receive and transmit).
  - Sync'd – If Yes, the PFC priorities have been set by the peer. This parameter cannot be set.
  - Error – Indicates whether an error occurred during the configuration exchange with the peer. Error is also set to Yes when the Compatible method fails.
  - FCoE Priority – The configured FCoE priority.
- ETS Priority Group Properties area:

  NOTE The ETS Priority Group Properties area is not displayed if multichannel is enabled on the adapter, with the exception of NPar and UFP (if enabled by the switch).

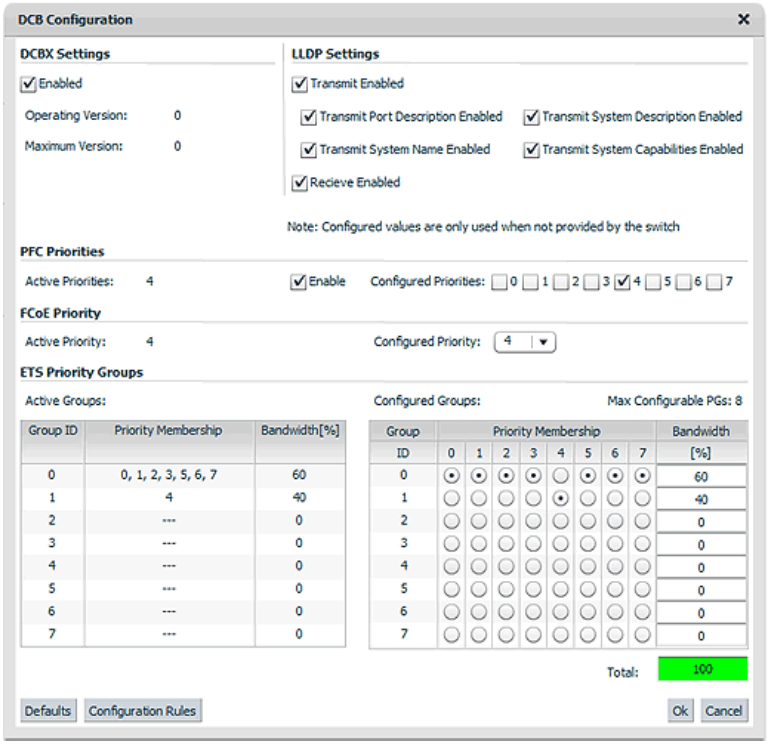
  - State – The protocol state. It can be enabled or disabled.
  - Sync'd – If Yes, the protocol priority has been set by the peer. This parameter cannot be set.
  - Error – Indicates whether an error occurred during the configuration exchange with the peer. Error is also set to Yes when the Compatible method fails.
  - Active Groups table:
    - Group ID – The Priority Group number. It can be 0 to 7.
    - Priority Membership – 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.

The DCB Configuration dialog allows you to configure physical port DCB parameters. Although there are eight ETS Priority Groups in which the priorities can be assigned, only two of those groups can be assigned priorities and bandwidths. The FCoE priority is assigned to one Priority Group and the other seven (NIC) priorities must be assigned to one and only one other Priority Group. If you are configuring Priority Groups on a DCBX enabled switch, in order to work properly with the adapter port, you must only configure two Priority Groups in the same way.

To configure DCB for NIC + FCoE adapter ports, perform the following steps:

1. From the **DCB** tab, click **Configure DCB**. The DCB Configuration dialog appears (Figure 95).

Figure 95 DCB Configuration Dialog (NIC + FCoE)



2. Configure the settings. The following Configure DCB dialog fields are displayed:

- 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 – 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. If the LAN device supports RFC 2863, the value equals the ifDescr object.
  - 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. If the LAN device supports RFC 3418, the value equals the sysDescr object.
  - Transmit System Name Enabled – Provides the system's assigned name in an alpha-numeric format. If the LAN device supports RFC 3418, the value equals the sysName object.
  - Transmit System Capabilities Enabled – Indicates the primary functions of the device and whether 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 – Enables PFC priorities.
  - To configure a NIC port, clear this checkbox.
  - To configure an FCoE port, select this checkbox.
- Configured Priorities – If the Enable checkbox is selected, these are the priorities that are configured but that might not yet be active. You can select two PFC priority checkboxes, one of which must match the FCoE priority. The other PFC priority is for Ethernet traffic. The PFC priority for Ethernet traffic must be assigned to a Priority Group that has no other priorities.
- FCoE Priority area:
  - Active Priority – The active FCoE priority.
  - Configured Priority – The configured FCoE priority.

- ETS Priority Groups area:
  - Active Groups – This area displays the settings that are currently active:
    - Max Configurable PGs – The maximum number of Priority Groups that can be configured.
    - Group ID – The active Priority Group ID.
    - Priority Membership – The priorities that are currently assigned to the various Priority Groups.
    - Bandwidth – The bandwidths that are currently assigned to various Priority Groups.
  - Configured Groups – This area allows you to create a new configuration:
    - Group ID – The configured Priority Group ID.
    - Priority Membership – The configured priority membership grouping.
    - Bandwidth% – The configured bandwidth value for the various Priority Groups.

    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.

3. Click **OK** to apply and save your changes.

To see the configuration rules that apply to configuring priorities, click **Configuration Rules**. You must observe these rules when configuring Priority Groups for adapter ports. Configuration rules are displayed based on the type of adapter ports you are configuring.

- FCoE configuration rules:
  - Only one priority is configured for the priority.
  - You can select a maximum of two PFC priorities and one of them must match the priority.
  - The PFC Priority Group to which the priority is assigned must contain no other priorities.
  - The additional PFC priority must be assigned to a Priority Group that has no other priority.
  - Bandwidths of all the Priority Groups must total 100%.
- NIC configuration rules:
  - Only one priority is configured for the priority.
  - The PFC priority must be assigned to a Priority Group that has no other priorities.
  - Bandwidths of all the Priority Groups must total 100%.

To reset the DCB configuration to its default values, click **Defaults**.

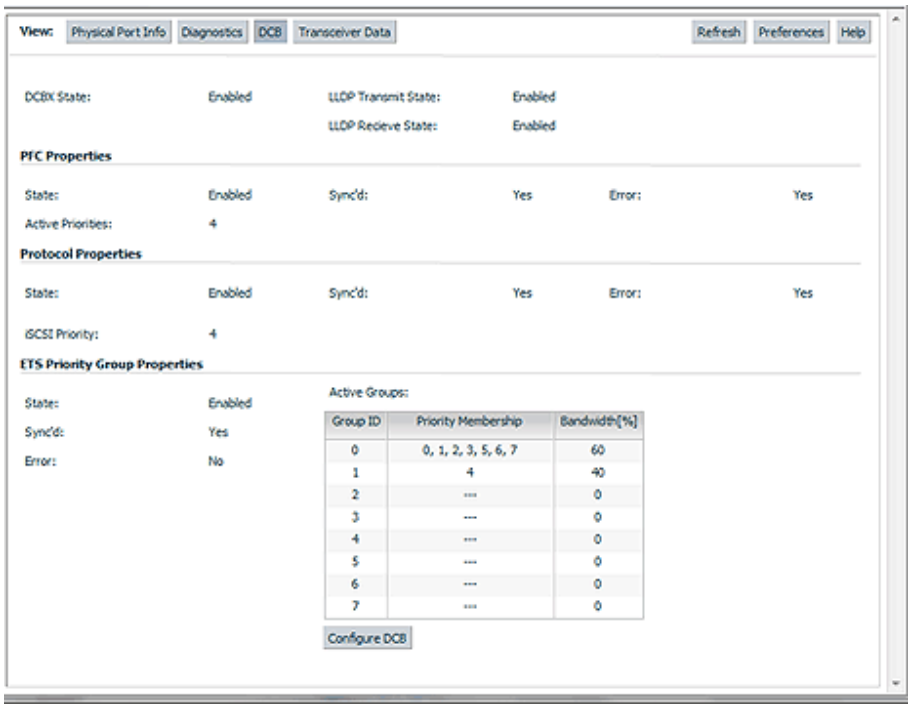
8.2 Configuring DCB Parameters for NIC + iSCSI Adapter Ports

To view the DCB parameters for a NIC + iSCSI port, perform the following steps:

1. Select a host in the vSphere console tree-view. If applicable, select the **Broadcom CNA OneCommand** tab.

2. In the Broadcom CNA Device Management tree-view, select the port for which you want to view DCB parameters.
3. Select the **DCB** tab (Figure 96).

Figure 96 DCB Tab (NIC + iSCSI)



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

- 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:

  NOTE PFC is not supported on all iSCSI adapter ports.

  — State – If Enabled, there is flow control in both directions (Tx and Rx). If Disabled, flow control is currently disabled. The priority value, if shown, is not applicable. This can be caused by:
    - The switch port PFC is set to on instead of automatic.
    - The switch port is using port flow control instead of PFC.
    - PFC is disabled at the adapter or the 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 – 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.
- Protocol 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 – Indicates whether an error has occurred during the configuration exchange with the peer.

- ETS Priority Group Properties area:

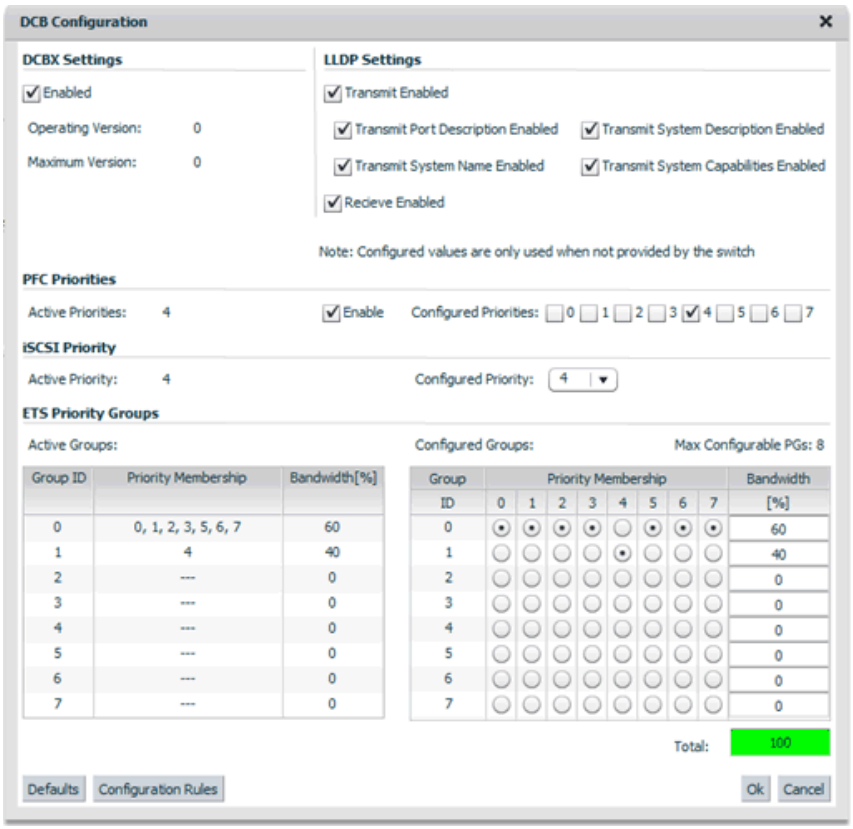
  NOTE Not displayed if multichannel is enabled on the adapter, with the exception of NPar and UFP (if enabled by the switch).

  - 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 – Indicates whether an error has occurred during the configuration exchange with the peer.
- Active Groups table:
  - 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.

To configure DCB for NIC + iSCSI, perform the following steps:

1. From the **DCB** tab, click **Configure DCB**. The Configure DCB dialog appears (Figure 97).

Figure 97 Configure DCB Dialog for iSCSI Adapter Ports (DCBX Enabled)



2. Configure the settings. The following Configure DCB dialog fields are displayed:

- 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 – 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. If the LAN device supports RFC 2863, the value equals the ifDescr object.
  - Transmit System Name Enabled – Provides the system's assigned name in an alpha-numeric format. If the LAN device supports RFC 3418, the value equals the sysName object.
  - 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 system's name and versions of hardware, operating system, and networking software supported by the device. If the LAN device supports RFC 3418, the value equals the sysDescr object.
  - Transmit System Capabilities Enabled – Indicates the primary functions of the device and whether 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 – When checked, PFC is enabled.
  - Configured Priorities – The priorities that are configured but that might not yet be active. A maximum of two PFC priority checkboxes can be selected, one of which must match the iSCSI priority. The additional PFC priority is for Ethernet traffic. The PFC priority for Ethernet traffic must be assigned to a Priority Group that has no other priorities.
  - iSCSI Priority area:
    - Active Priority – The active iSCSI priority.
    - Configured Priority – The configured iSCSI priority.
- ETS Priority Groups area:

    NOTE This area is not available if UMC is enabled on the adapter.
  - Active Groups – This area displays the settings that are currently active:
    - Group ID – The active Priority Group ID.
    - Priority Membership – The priorities that are currently assigned to the various Priority Groups.
    - Bandwidth – The bandwidths that are currently assigned to various Priority Groups.
  - Configured Groups – This area allows you to create a new configuration:
    - Group ID – The configured Priority Group ID.
    - Priority Membership – The configured priority membership grouping.
    - Bandwidth% – The configured bandwidth value for the various Priority Groups.

      NOTE An error message is displayed if you try to configure more priority groups than the adapter supports. The Max Configurable PGs field displays the number of priority groups supported by the adapter.

To see the configuration rules that apply to configuring priorities, click **Configuration Rules**. You must observe these rules when configuring Priority Groups for adapter ports. Configuration rules are displayed based on the type of adapter ports you are configuring.

- iSCSI configuration rules:
  - Only one priority is configured for the priority.
  - You can select a maximum of two PFC priorities and one of them must match the priority.

- The PFC Priority Group to which the priority is assigned must contain no other priorities.
- The additional PFC priority must be assigned to a Priority Group that has no other priority.
- Bandwidths of all the Priority Groups must total 100%.

- NIC configuration rules:
  - Only one priority is configured for the priority.
  - The PFC priority must be assigned to a Priority Group that has no other priorities.
  - Bandwidths of all the Priority Groups must total 100%.

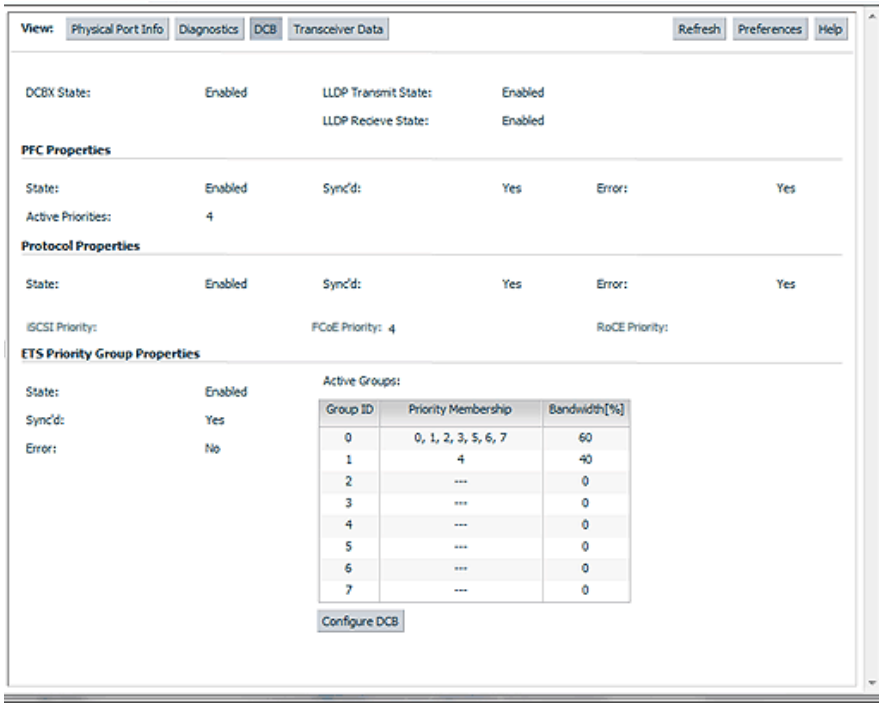
To reset the DCB configuration to its default values, click **Defaults**.

8.3 Configuring DCB Parameters for NIC + FCoE + iSCSI Adapter Ports

To view the DCB parameters for a NIC + FCoE +iSCSI port, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the port whose DCB properties you want to configure.
3. Select the **DCB** tab (Figure 98).

Figure 98 DCB Tab (NIC + FCoE +iSCSI Adapter Port Selected)



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

- 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.
  - Active Priorities – Lists the priorities with PFC that are set to Enabled.
  - Sync'd – If set to Yes, the PFC priorities have been set by the peer. This parameter cannot be set.
  - Error – 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.
- Protocol Properties area:
  - State – The FCoE or iSCSI state. It can be enabled or disabled.
  - Sync'd – If Yes, the FCoE or iSCSI priority has been set by the peer. This parameter cannot be set.
  - Error – 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.
  - iSCSI Priority and FCoE Priority – Protocol properties, as applicable.
- ETS Priority Group Properties area:

  NOTE This area is not displayed if multichannel is enabled on the adapter, with the exception of NPar and UFP (if enabled by the switch).

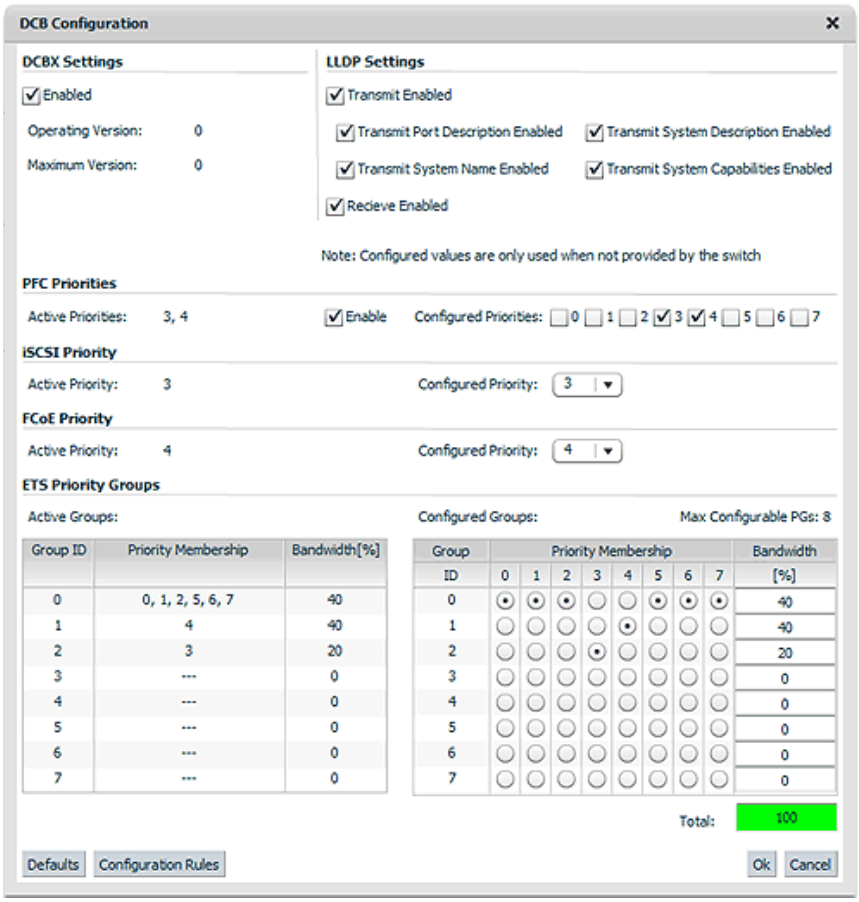
  - State – The FCoE, iSCSI, or NIC-only state. It can be enabled or disabled.
  - Sync'd – If Yes, the Priority Groups have been set by the peer. This parameter cannot be set.
  - Error – Indicates whether an error has occurred during the configuration exchange with the peer. Error is also set to Yes when the Compatible method for the capability fails.
  - Active Groups:
    - Group ID–The Priority Group number. It can be 0 to 7.
    - Priority Membership – 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.

The DCB Configuration dialog allows you to configure physical port DCB parameters. Although there are eight ETS Priority Groups in which the priorities can be assigned, only two of those groups can be assigned priorities and bandwidths. The FCoE priority is assigned to one Priority Group and the other seven (NIC) priorities must be assigned to one and only one other Priority Group. If you are configuring Priority Groups on a DCBX enabled switch, in order to work properly with the adapter port, you must only configure two Priority Groups in the same way.

To configure DCB parameters for a physical port, perform the following steps:

1. From the **DCB** tab, click **Configure DCB**. The DCB Configuration dialog is displayed (Figure 99).

Figure 99 DCB Configuration Dialog (NIC + FCoE + iSCSI)



2. Configure the settings you want and click **OK**. The following DCB Configuration dialog fields are displayed:

- 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.
  - DCBX Mode – The DCBX mode can be set to DCB or CIN. Changes to the DCBX mode 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. If the LAN device supports RFC 3418, the value equals the sysName object.
  - 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. If the LAN device supports RFC 3418, the value equals the sysDescr object.

- Transmit System Capabilities Enabled – Indicates the primary function(s) of the device and whether 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 – This checkbox enables PFC priorities.
    - If you are configuring a NIC port, clear this checkbox.
    - If you are configuring an iSCSI or FCoE port, select this checkbox.
  - Configured Priorities – If the Enable checkbox is selected, these are the priorities that are configured, but might not yet be active. You can select two PFC priority checkboxes. One of these selections must match the FCoE or iSCSI priority. The other PFC priority is for Ethernet traffic. The PFC priority for Ethernet traffic must be assigned to a Priority Group which has no other priorities.
  - 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:
  - Active Groups – This area displays the settings that are currently active:
    - Group ID – The active Priority Group ID.
    - Priority Membership – The priorities that are currently assigned to the various Priority Groups.
    - Bandwidth – The bandwidths that are currently assigned to various Priority Groups.
  - Configured Groups – This area allows you to create a new configuration:
    - Group ID – The configured Priority Group ID.
    - Priority Membership – The configured priority membership grouping.
    - Bandwidth% – The configured bandwidth value for the various Priority Groups.

3. Click **OK** to apply and save your changes.

If you want to see the configuration rules that apply to configuring priorities, click **Configuration Rules**. You must observe these rules when configuring Priority Groups for adapter ports. Configuration rules are displayed based on the type of adapter ports you are configuring.

- FCoE and iSCSI configuration rules:
  - Only one priority is configured for the priority.
  - You can select a maximum of two PFC priorities and one of them must match the priority.
  - The PFC Priority Group to which the priority is assigned must contain no other priorities.
  - The additional PFC priority must be assigned to a Priority Group which has no other priority.
  - Bandwidths of all the Priority Groups must total 100%.
- NIC configuration rules:
  - Only one priority is configured for the priority.
  - The PFC priority must be assigned to a Priority Group which has no other priorities.
  - Bandwidths of all the Priority Groups must total 100%.

To reset the DCB configuration to its default values, click **Defaults**.

Chapter 9: Exporting SAN Information in Cluster View

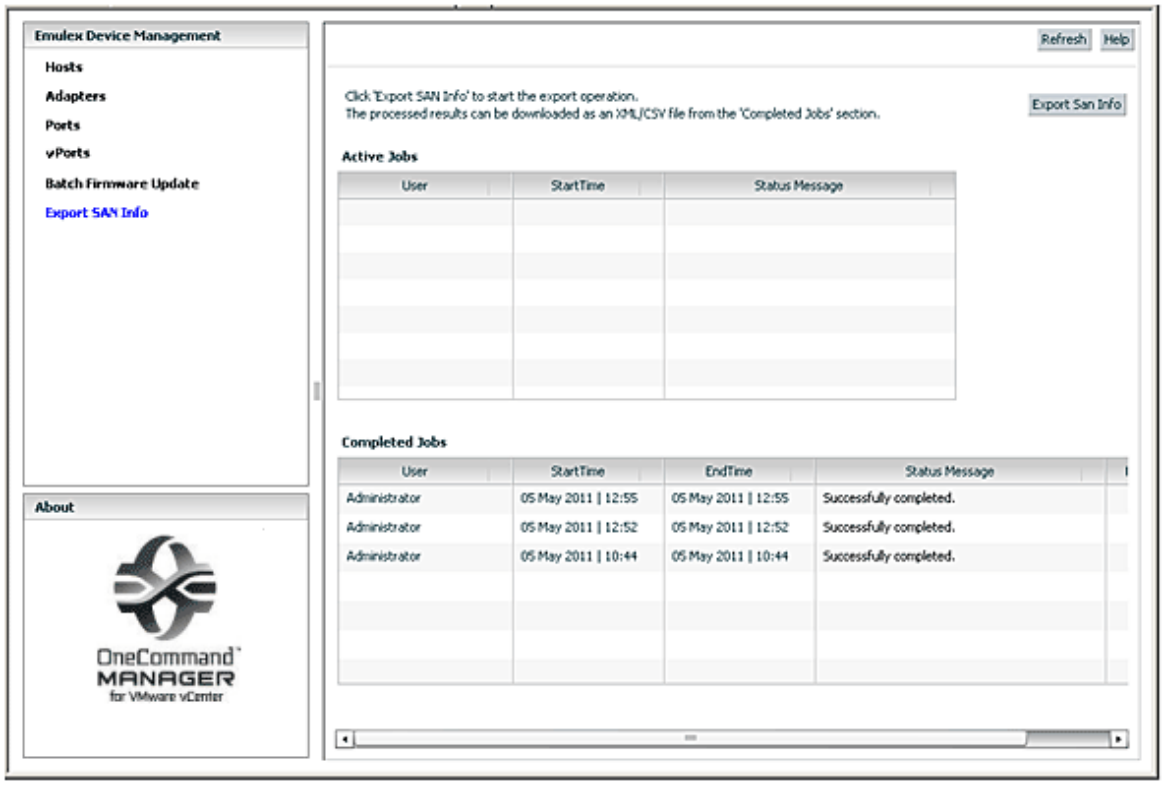
Exporting SAN information (creating a SAN report) progresses in the background until all the jobs are completed. During this period you can still browse through the other tabs.

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

To export SAN information, perform the following steps:

1. Select a cluster in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. From the Broadcom CNA Device Management tree, select **Export San Info**. Export SAN Information is displayed (Figure 100).

Figure 100 Export SAN Info



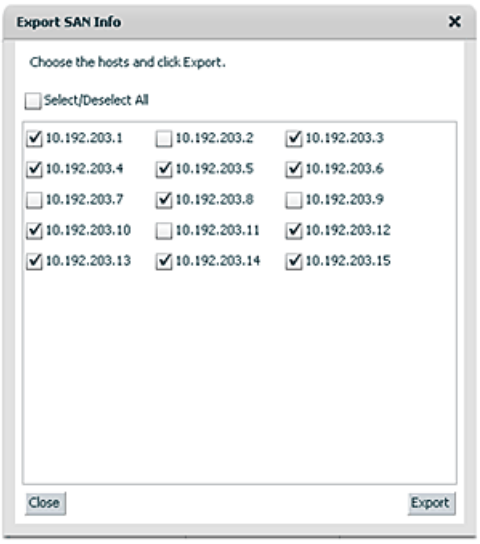
The following Export SAN Info window fields are displayed:

— Active Jobs column:
- User – The user who updated the firmware.
- Start Time – The time the job is submitted.
- Status Message – The status of the job.

— Completed Jobs column:
- User – The user who updated the firmware.
- Start Time – The time the job is submitted.
- End Time – The time the job is completed.
- Status Message – The status of the job.
- Download – Click **XML** or **CSV** to download the file with SAN information in the specified format.

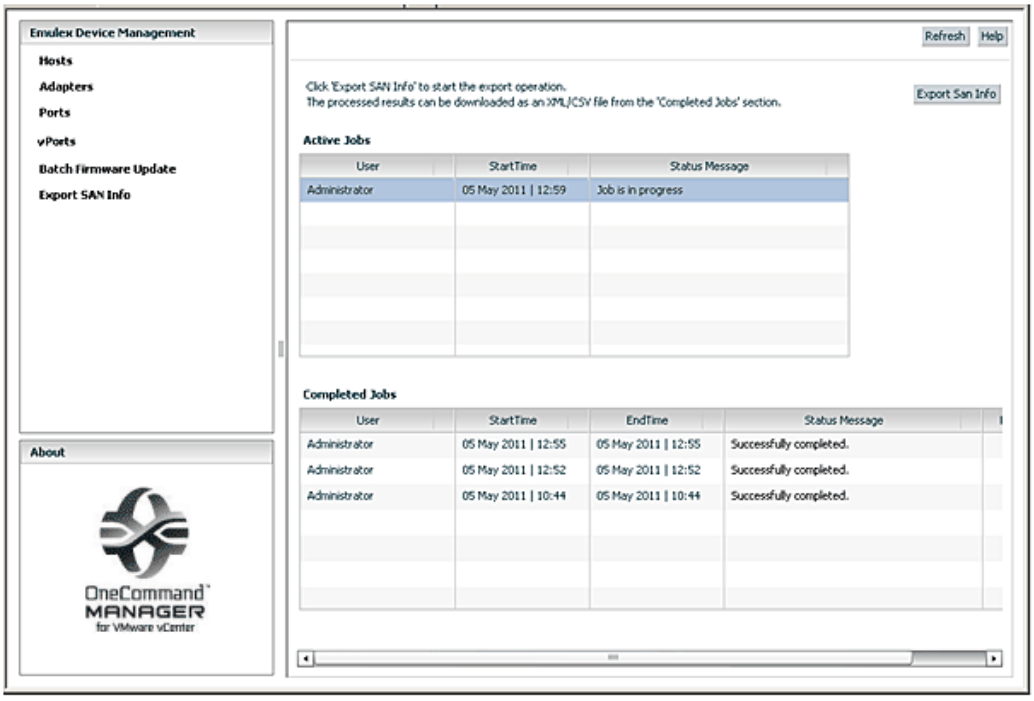
3. Click **Export San Info**. The Export SAN Info dialog is displayed (Figure 101).

Figure 101 Export SAN Info Dialog



4. Select host or hosts to export. Optionally, select the **Select/Deselect All** checkbox to select all hosts. Clear the **Select/Deselect All** checkbox to clear all host checkboxes.
5. Click **Export**.
6. The **Export SAN Info Jobs** window is displayed. The status is displayed in the Active Jobs list (Figure 102).

Figure 102 Export SAN Info Jobs Window, Active Jobs List



7. When the export job is completed, the entry is displayed in the Completed Jobs list with a **Successfully completed** status.

9.1 Capturing SAN Information in XML or CSV Format

1. Click the row that represents the information you want to capture.
2. Scroll to the right until the Download column appears.
3. In the Download column, click **XML** or **CSV** to capture the information.

9.2 Considerations When Exporting SAN Information in a Cluster View

- At any time only ten completed jobs are available to be exported. If more than ten jobs are completed, the first completed jobs are not available for download.
- If you click **Export SAN Info** and log out from the vSphere console before the export is completed, all pending jobs fail.

Chapter 10: Diagnostics

This section describes diagnostics for OneConnect adapters.

When running port diagnostic tests using OneCommand CNA Manager for VMware vCenter, do not perform any active management operations on the ESXi host.

10.1 Running Loopback Tests

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.

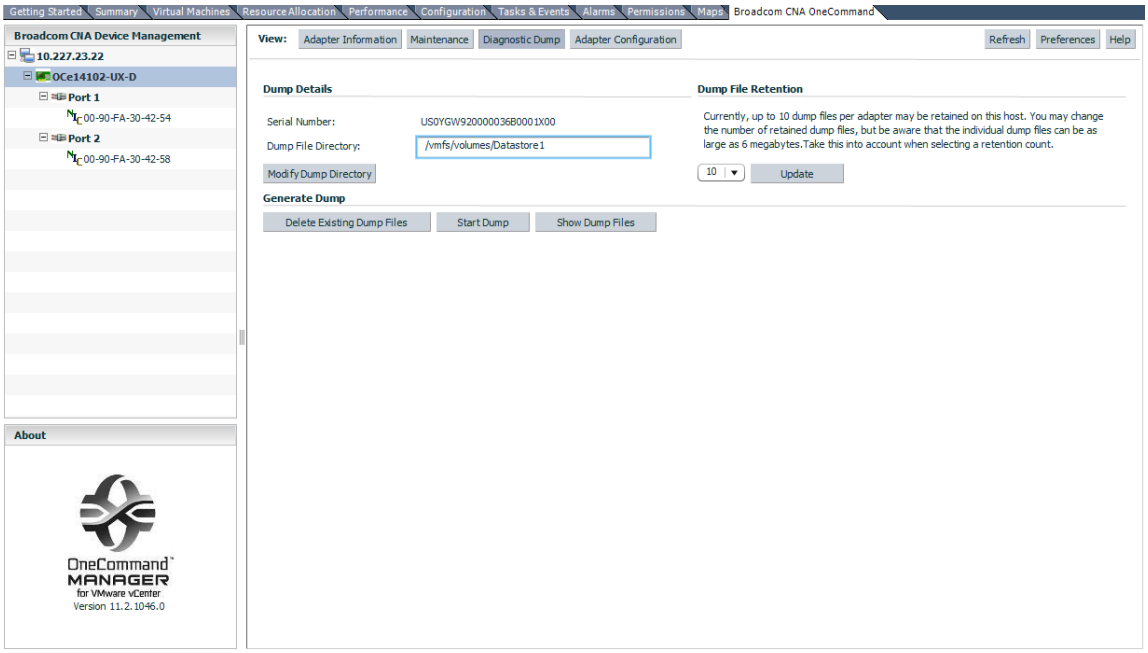
Run the following loopback tests:

- PHY Loopback – 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.
- MAC Loopback – The MAC loopback test connects the transmit output of the MAC controller to the receive input of the MAC controller (bypassing the PHY).

To run loopback tests, perform the following steps:

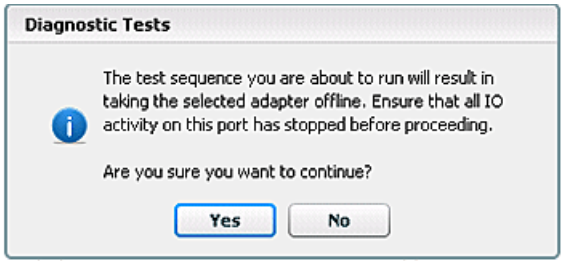
1. From the Broadcom CNA Device Management tree-view, select the physical port on which you want to run the loopback test.
2. Select the **Diagnostics** tab (Figure 103). In the Loopback Test area of the dialog, choose the type of loopback test you want to run and define the loopback test parameters.

Figure 103 Emulex OneConnect Adapter Diagnostics Tab



3. Click **Start Test**. The following warning is displayed (Figure 104).

Figure 104 Diagnostic Tests Warning



4. Click **Yes**. A progress bar shows 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. Click **Show Test Logs** to view and save the log file.

10.2 Using Beaconing

Beaconing can force an 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.

When you enable beaconing for OneConnect adapters, the two LEDs blink rapidly in unison until you disable beaconing.

NOTE Beaconing is disabled if the selected adapter does not support beaconing.

To enable or disable beaconing, perform the following steps:

1. From the Broadcom CNA Device Management tree-view, select the adapter port whose LEDs you want to set.
2. Select the **Diagnostics** tab (Figure 103) and click **Beacon On** or **Beacon Off**.

10.3 Running a TDR Test (10GBASE–T–Series Adapters Only)

To run a TDR test, perform the following steps:

1. From the Broadcom CNA Device Management tree-view, select the adapter port on which you want to run a TDR test.
2. Select the **Diagnostics** tab (Figure 103).
3. In the TDR area, enter an NVP value. This value compares the cable speed of reflection versus the speed of light.
4. Click **Start TDR Test**.

10.4 Setting Up Diagnostic Test Options

Setting up test options include error actions, test cycle counts, and test patterns.

10.4.1 Setting Up a Test Failure Error Action

There are two error action options in the event of a test failure:

- Stop Test – Does not log the error and aborts the test. No further tests are run.
- Ignore – Logs the error and proceeds with the next test cycle.

10.4.2 Setting Up Test Cycles

Specify one of the following test cycles:

- 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.

10.4.3 Setting Up a 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.

10.5 Test Status

The Test Status area shows how many completed cycles of each test ran, as well as the number of errors. 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 area. Click **Show Test Logs** to view and save the log file.

10.6 Saving the Log File

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

- The default location is the OneCommand CNA Manager for VMware vCenter install directory on your local drive.
- In the VMware Server: There is no default directory for ESXi.

After writing an entry into the log, you are prompted to clear the display.

To save the log file, perform the following steps:

1. After running a test from the **Diagnostics** tab, click **Show Test Logs**. The Diagnostic Test Logs dialog appears. The default name of a saved file is `DiagTest.log`.
2. Browse to the desired directory, change the log file name if you want and click **Save to File**.

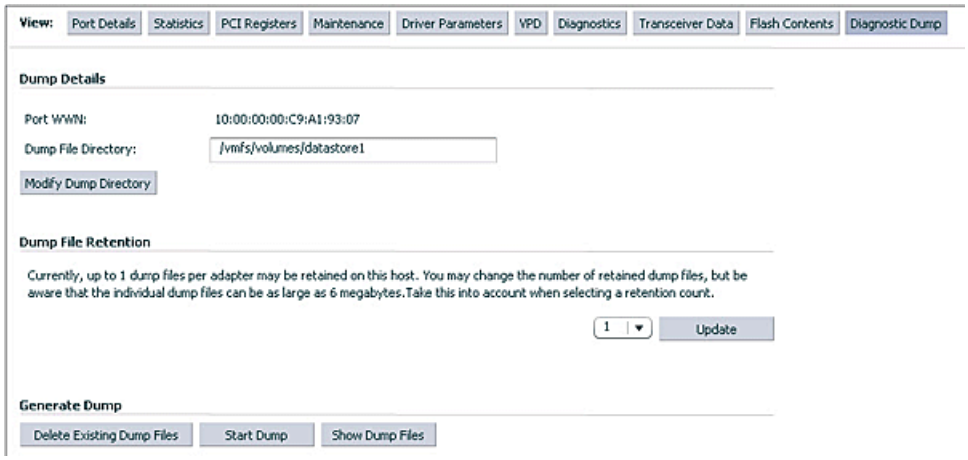
10.7 Creating OneConnect Adapter Port Diagnostic Dumps

Diagnostic dump enables you to create and manage a diagnostic dump for a selected OneConnect adapter port. Dump files contain information such as firmware version and driver version that is particularly useful when troubleshooting an adapter.

To start a diagnostic dump, perform the following steps:

1. Select a host in the vSphere console tree-view, and if applicable, select the **Broadcom CNA OneCommand** tab.
2. In the Broadcom CNA Device Management tree-view, select the port.
3. Select the **Diagnostic Dump** tab (Figure 105).

Figure 105 Port Diagnostic Dump Tab



4. Enter a name in the **Dump File Directory** field in the Dump Details area to set the dump file directory. A prefix of /vmfs/volumes is added to the location, if it is not specified.
5. To specify up to 20 files to retain using the Dump Files Retention counter, enter the number of files and click **Update**.
6. Click **Start Dump** to initiate a diagnostic dump on the selected port. Click **Delete Existing Dump Files** to remove existing dump files for the selected Emulex OneConnect port.

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

The following **Diagnostic Dump** tab fields are displayed:

- Port WWN– The port WWN.
- Dump File Directory – The location where the dump files are created.
- Dump File Retention – The number of dump files per port to retain on the host.

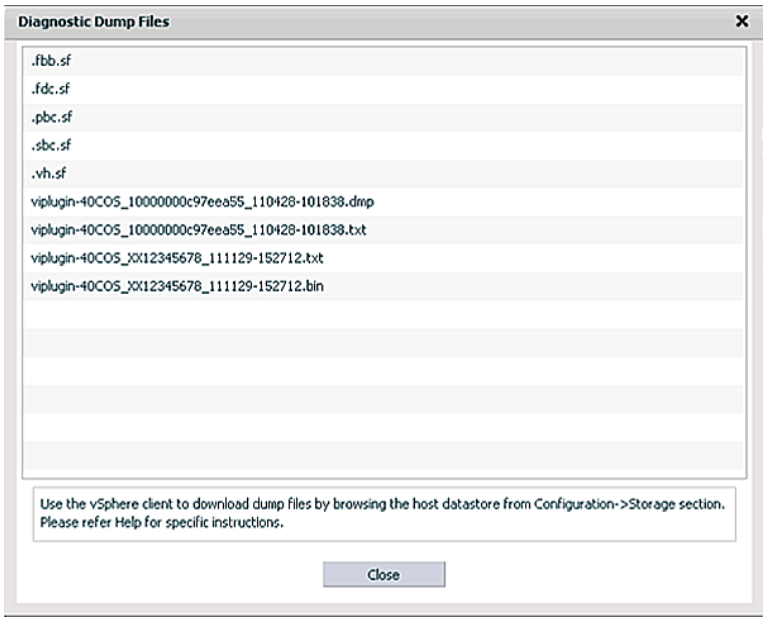
10.7.1 Viewing Diagnostic Dump Files

Diagnostic dump file names can be viewed using the OneCommand CNA Manager for VMware vCenter. The dump files are stored on the host's data store and the vSphere client can be used to download dump files by browsing the host data store.

To view the diagnostic dump in the OneCommand CNA Manager for VMware vCenter, perform the following steps:

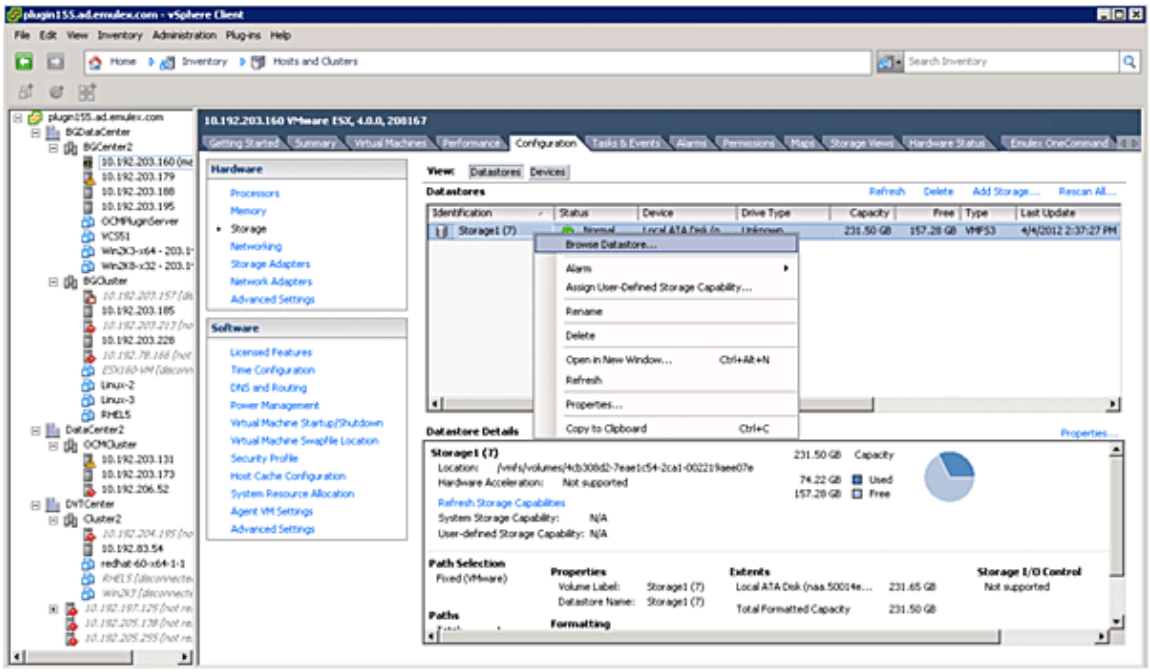
1. On the **Diagnostic Dump** Tab (Figure 105), click **Show Dump Files**. The Diagnostic Dump Files window is displayed, showing the diagnostic dump files currently on your system (Figure 106). These files are available in the dump directory configured from the **Diagnostic Dumps** tab. You can extract these files using the vSphere client.

Figure 106 Diagnostic Dump Files Window



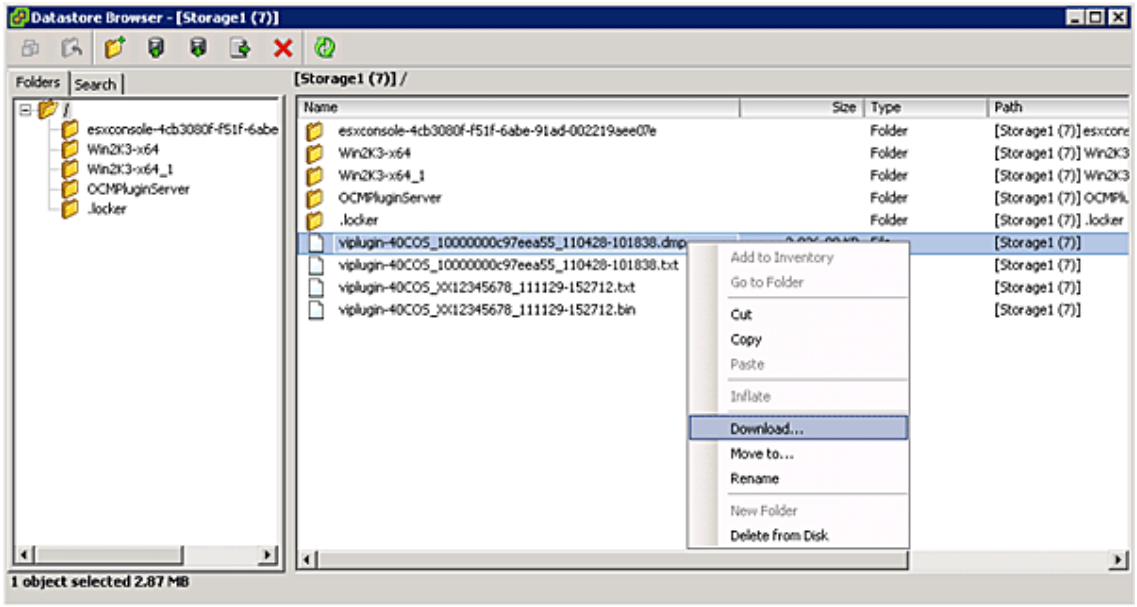
2. Extract the dump files by using the vSphere client to download the dump files by browsing the data store.
   a. Click the **Configuration** tab in the vSphere client. The Datastores view is displayed (Figure 107).

Figure 107 Datastores View of vSphere Client



b. Right-click the datastore where your dump file is located, and select **Browse Datastore...** from the context menu. A Datastore Browser window opens (Figure 108).

Figure 108 vSphere Datastore Browser



c. In the Datastore Browser, right-click the dump file that you want to review and select **Download** from the context menu. A Browse for Folder window opens.

d. Select the desired location for the dump files in the Browse for Folder window and press **OK**. The file is down loaded to the location you select.

You can view the dump file in any text editor.

Chapter 11: Generating and Installing Secured Certificates

OCMNG is a web application, based on a client-server model, running on the Apache Tomcat Web Server. Data is exchanged between the client (browser) and the server (on a remote machine), which requires a secure user login to manage Emulex adapters on different and multiple hosts.

11.1 Secure Sockets Layer (SSL) Certificate

An SSL certificate establishes an encrypted connection between the web server and the web browser on a remote machine. This connection allows private information to be transmitted without eavesdropping, data tampering, or message forgery.

An SSL certificate provides security through encryption and authentication. Encryption is ensured by accessing the remote server using the HTTPS protocol and an SSL certificate.

NOTE If OneCommand CNA Manager for VMware vCenter is running, the server must be configured to support HTTPS protocol access and provide a self-signed certificate.

The OneCommand CNA Manager for VMware vCenter server is authenticated to the browser by a Public Key in the self-signed certificate.

11.1.1 Generating an SSL Certificate

To allow secured communication between the client and server, perform the following steps:

1. Generate a self-signed certificate with a keystore for each server providing the server's domain name and company details. See Section 11.1.2, Generating a Self-Signed Certificate for instructions. For more information, see the X.509 attributes list on the International Telecommunications Union website.
2. Use this certificate to create a request to the customer's trusted certificate authority (CA). The request certificate is referred as a Certificate Signing Request (CSR). The CA issues a new SSL certificate. See Section 11.2.1, Generating a CSR for a Server Using the Java Tool for instructions.
3. Import the new SSL certificate to the application server, and install the SSL certificate on the client's browser. See Section 11.2.4.1, Installing the Certificates to the Keystore of OneCommand CNA Manager for VMware vCenter for instructions.
4. Configure the server to use the keystore file. See Section 11.2.4.2, Configuring a Web Server for instructions.
5. Access the server's content through the browser using HTTPS protocol.
6. The browsers understand the certificate and the browsers allow access to and from the remote server.

11.1.2 Generating a Self-Signed Certificate

A self-signed certificate is a certificate that is signed by itself (the server hosting OneCommand CNA Manager for VMware vCenter) rather than a trusted CA. This self-signed certificate includes a public or private key that is distributed by the SSL to verify the identity of the server.

A self-signed certificate can also be used as an alternative to SSL certificates if the server is not running in a public domain.

If a self-signed certificate is used in place of an SSL, a warning is shown in the browser before accessing the server content.

For Java-based applications, a self-signed certificate can be generated using the tools provided by Java. This creates a keystore file which must be installed on the web-server. This keystore includes a private key specific to the server used for generating a CSR and authenticating the server.

As the OneCommand CNA Manager for VMware vCenter server is developed using Java, it leverages the keytool tool provided by Java to generate the self-signed certificates at no cost.

NOTE The self-signed certificate for the OneCommand CNA Manager for VMware vCenter server is generated and installed on its server as part of the OneCommand CNA Manager for VMware server installation on a Windows machine. This self-signed certificate is generated with Broadcom organization details using RSA algorithm and private key of size 2048 bits.

To generate a self-signed certificate, perform the following steps:

1. In the OCM for VMware installation directory, go to `ApacheTomcat\conf`.

```
>>cd /d  "C:\Program Files\Brcm\OCM for VMware\ApacheTomcat\conf"
```

2. Run the following command:

```
>> ..\..\JRE\bin\keytool.exe -genkey -alias <new-alias> -keyalg RSA
-keystore emulex.vcplugin.jks  -keysize 2048
```

NOTE You can change alias, keysize, and keystore name.

Example

```
Enter keystore password: (Enter "emulex" if using the same keystore name)
Re-enter new password:
What is your first and last name?
  [Unknown]:  pluginserver.ad.emulex.com  (Give the complete domain name of the
server [FQDN])
What is the name of your organizational unit?
  [Unknown]:  ocm
What is the name of your organization?
  [Unknown]:  elx
What is the name of your City or Locality?
  [Unknown]:  bg
What is the name of your State or Province?
  [Unknown]:  ka
What is the two-letter country code for this unit?
  [Unknown]:  in
Is CN=pluginserver.ad.emulex.com, OU=ocm, O=elx, L=bg, ST=ka, C=in correct?
  [no]:  yes

Enter key password for <elxocm>:
        (RETURN if same as keystore password)
```

11.2 Generating a Certificate Signing Request (CSR)

A CSR is a block of encrypted text that is generated on the server on which the certificate is used. A CSR contains information to be included in the SSL certificate such as the organization name, common name (domain name), locality, country, and other X.509 attributes. It also contains the public key that is included in the certificate. The CA uses the CSR to create a new SSL certificate.

11.2.1 Generating a CSR for a Server Using the Java Tool

To generate a CSR for a server, use the Java tool available in the jre/bin folder. The syntax using the Java tool is:

```
keytool -certreq -keyalg <algorithm>  -alias <alias-name>  -file <csr-name>
-keystore <keystore-name>
```

Example

```
keytool -certreq -keyalg RSA -alias selfsigned -file elxocmreq.csr -keystore
emulex.vcplugin.jks
```

11.2.2 Generating and Validating a CSR

To generate a CSR, perform the following steps:

1. Generate a self-signed certificate (see Section 11.1.2, Generating a Self-Signed Certificate for instructions).
2. Generate a CSR using the following syntax:

```
>>..\..\JRE\bin\keytool -certreq -v -alias <new-alias> -file elxocmreq.csr
-keypass elxocm -keystore emulex.vcplugin.jks
Enter the keystore password: (Enter "emulex" if using the default keystore
name)
Certification request stored in file <elxocmreq.csr>
```

To validate a CSR for its completeness, perform the following steps:

You can validate the generated CSR for its completeness before submitting (with the help of the CA). Copy the CSR content from the following link for validation.

```
http://www.sslshopper.com/csr-decoder.html
```

NOTE The CSR must begin and end with the following tags:

```
-----BEGIN NEW CERTIFICATE REQUEST-----
-----END NEW CERTIFICATE REQUEST-----
```

Example

```
-----BEGIN NEW CERTIFICATE REQUEST-----
MIIC2TCCAcECAQAwZDELMAkGA1UEBhMCaW4xCzAJBgNVBAgTAmthMQswCQYDVQQHEwJiZzEMMAoG
A1UEChMDZWx4MQwwCgYDVQQLEwNvY20xHzAdBgNVBAMTFmJnc3N5ZWQxLmFkLmVtdWxleC5jb20w
ggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQCDhovljXhfjNPM/5eBsX4280AIl3YARn0p
R6Z7eOqs1r5Qh07kT58M6T8fER+NpIN7WhOOF/TbsFsS0gmfYwJQqttvtq1dtxUGpvFe9lywbP+l
kY+w6GOyPTG2qnXgILtX5ArZbC2UBbz+J8WJ3SjPHXiSY35EZWnyZZmIN8v1vOe9e21f8vwRkn/4
fdfFrpoQa3H+GcAJMRSBRTd5H6mXQv6HaA5Z0BbsABisFx4scqSuM/HJKLP6GcSHR61bzHfiO/NH
4qU/s6I2LC5DvGs1hIW3PPbmb1rxBiEFpjPtWhfzPxPMKSU8uey+lE0UIPMS0FMTxo63oYnMeiSX
X5mxAgMBAAGgMDAuBgkqhkiG9w0BCQ4xITAfMB0GA1UdDgQWBBSvpKLBF3lY03Jin9kI4ym94bJi
zjANBgkqhkiG9w0BAQsFAAOCAQEAfs94wzEUlDAMq0jITi6fiD7YxK2KFWJgMBfjxZIGex2zxlHL
mOS14BGSWk5dvSMwqDBC14l4C79rUOlTUwwWs9zFqMHynndQ2Ze2vuJNTWUlnFyFb37/rEvbFufB
QVvFXgycaKRgUpWo2x5sekRJRAPxXI/vLWOFRLLrzcVykgZ/sg3QrO4ezlKFc49put0vKpvI1dY9
l9BN2REuWrlmq5y3L8nx8mKX9dRmP6CKzHBaVrvY+nVju+Vf/ikfTtQIDEXAIW2Q7AObpcOaudnf
Nsaey+u27vGy77gAv7092xBHsDyOrD7COy/83b194igmVBVY4dt0496oXkOHCA0txA==
-----END NEW CERTIFICATE REQUEST----
```

11.2.3 Getting an SSL Certificate

The CSR can be submitted to the trusted CA (as chosen by you). The CA validates the CSR and issues a new SSL certificate.

11.2.4 Installing the SSL into the Web Server

When you receive the SSL certificate from the CA, you must install the SSL certificate on the server to accept the secure connections.

NOTE The CSR must be generated on the same machine that the server is running on. The SSL certificate must also be installed on this same server.

11.2.4.1 Installing the Certificates to the Keystore of OneCommand CNA Manager for VMware vCenter

The Root Certificate file, the Intermediate Certificate file, and the Primary Certificate file must all be installed in the keystore.

To install the certificates to the keystore of OneCommand CNA Manager for VMware vCenter, perform the following steps:

1. Download the SSL certificate file from the CA. Save the SSL certificate file to the same directory as the keystore (self-signed certificate) that was created for the CSR.

   NOTE The certificate only works with the same keystore that was initially created for the CSR. The certificates must be installed to your keystore in the correct order.

2. Install the Root Certificate file.

   Every time you install a certificate to the keystore, you must enter the keystore password that you chose when you generated it. Enter the following command to install the Root Certificate file:

   ```
   keytool -import -trustcacerts -alias root -file RootCertFileName.crt
   -keystore keystore.key
   ```

   If the following message is displayed, select Yes:

   ```
   Certificate already exists in system-wide CA keystore under alias <...> Do
   you still want to add it to your own keystore?
   ```

   If successful, the following message is displayed:

   ```
   Certificate was added to keystore.
   ```

3. Install the Intermediate Certificate file.

   If the CA provided an Intermediate Certificate file, you must install it here using the following command:

   ```
   keytool -import -trustcacerts -alias intermediate -file
   IntermediateCertFileName.crt -keystore keystore.key
   ```

   If successful, the following message is displayed:

   ```
   Certificate was added to keystore.
   ```

4. Install the Primary Certificate file.

   Use following command to install the Primary Certificate file (for your domain name):

   ```
   keytool -import -trustcacerts -alias tomcat -file PrimaryCertFileName.crt
   -keystore keystore.key
   ```

   If successful, the following message is displayed:

   ```
   Certificate reply was installed in keystore.
   ```

All the certificates are now installed to the keystore file. You must configure your server to use the keystore file.

11.2.4.2 Configuring a Web Server

NOTE These configuration changes are not required if the default keystore name and password are used. If they are different, you must change the configuration as needed.

1. Copy the keystore file (or Self-signed Certificate) to a directory (preferably, the conf folder) of the Web Server.
2. Open the file `${CATALINA_HOME}/conf/server.xml` in a text editor.
3. Uncomment the SSL Connector Configuration.
4. Make sure the keystorePass matches the password for the keystore, and the keystoreFile contains the path and filename of the keystore.

   Your connector should be displayed similar to the following:

   ```
   <Connector className="org.apache.catalina.connector.http.HttpConnector"
   port="8443" minProcessors="5" maxProcessors="75" enableLookups="true"
   acceptCount="10" debug="0" scheme="https" secure="true">

   <Factory className="org.apache.catalina.net.SSLServerSocketFactory"
   clientAuth="false" protocol="TLS" keystoreFile="./conf/emulex.vcplugin.jks"
   keystorePass="emulex"/>
   ```
5. Save the changes to `server.xml`.
6. Restart the Web Server.

   If you launch the OneCommand CNA Manager for VMware vCenter URL in the browser, the application should be launched without any security warnings.

NOTE Use the hostname with the domain name that you used to generate the CSR.

Chapter 12: Troubleshooting

This section includes information about certificate or insecure-content warnings that may be displayed on the machine. This section also describes unexpected circumstances and some proposed solutions.

12.1 Security

The vSphere Web Client server and the OneCommand CNA Manager for VMware vCenter can be installed on different machines. As a result, certificate or insecure-content warnings could occur. There are two ways to remedy the issue:

- Accept the blocked content – temporary solution
- Install a security certificate – permanent solution

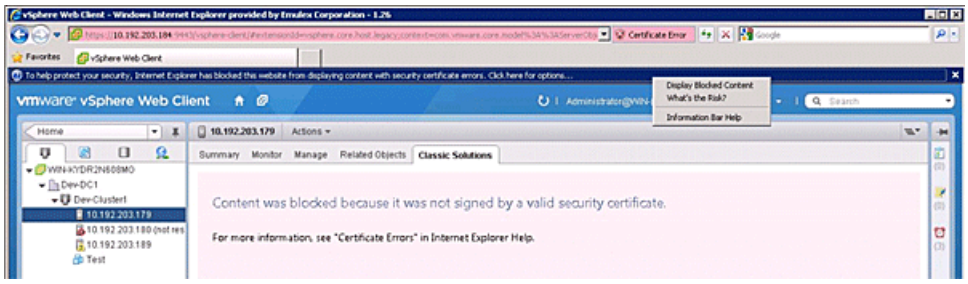
12.1.1 Accepting the Blocked Content

The procedure for accepting blocked content depends on the type of browser you are using. This is not a permanent solution and you must repeat this procedure every time you load the web client. If you want a permanent solution, you must install the correct security certificate. See Section 12.1.2, Installing a Security Certificate.

12.1.1.1 Internet Explorer

Accept the blocked content (Figure 109).

Figure 109 Blocked Content in Internet Explorer



12.1.1.2 Chrome and Firefox

1. Load the plug-in URL in a separate tab or window.

   The plug-in URL format is:

   ```
   https://<plugin-server>:<https-port>/brcmvcplugin
   ```

   For example:

   ```
   https://pluginserverhostFQDN:443/brcmxvcplugin
   ```

   NOTE You can extract the plug-in server, IP address, host name, and port number from the browser warning message.

2. Confirm or accept the certificate warning (Figure 110).
3. Refresh the vSphere web client tab or window.

Figure 110 Blocked Content in the Firefox or Chrome Browser



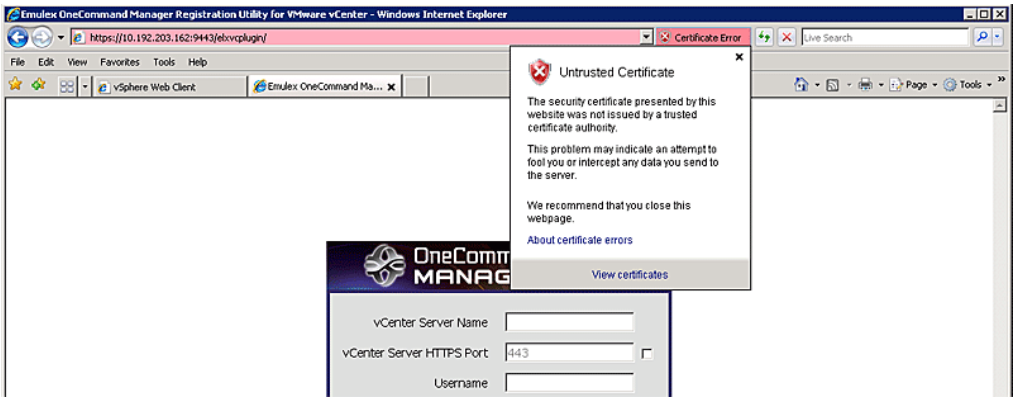
12.1.2 Installing a Security Certificate

A permanent solution to the security warnings is to install the correct security certificate.

To install a security certificate, perform the following steps:

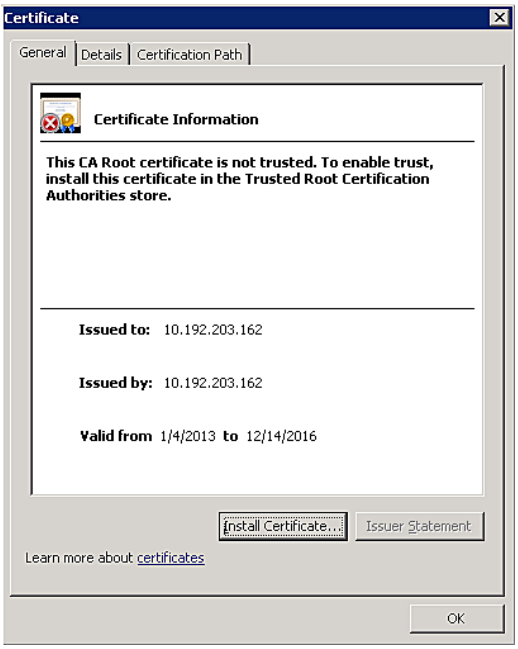
1. Open Internet Explorer in Administrative Mode.
2. Load the plug-in URL and accept the certificate warning.
   The plug-in URL must have the following format:
   `https://<plugin-server>:<https-port>/brcmvcplugin`
   The page loads with a certificate error.
3. Click the **Certificate Error**. A list appears showing the untrusted certificate (Figure 111).

Figure 111 Untrusted Certificate



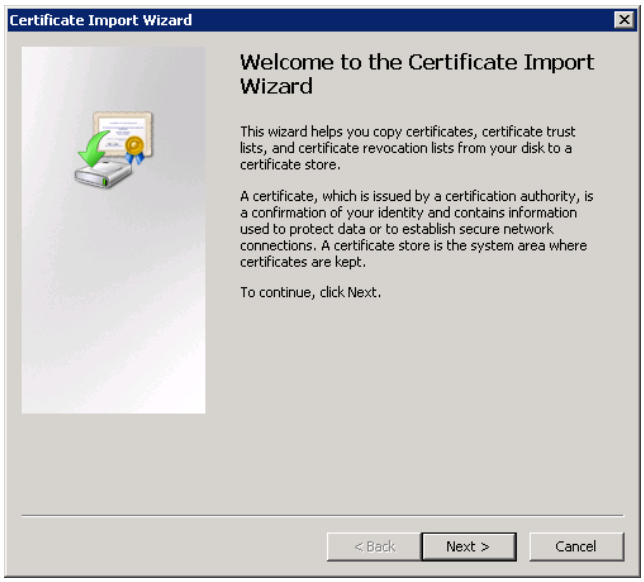
4. Click **View certificates** in the Certificate Error list. The Certificate dialog is displayed (Figure 112).

Figure 112 Certificate Dialog



5. Click **Install Certificate...**. The Certificate Import Wizard is displayed (Figure 113).

Figure 113 Certificate Import Wizard



6. Follow the wizard instructions and install the certificate to the Trusted Root Certification Authorities location.

12.1.3 Issues and Resolutions

There are several circumstances in which your system may operate in an unexpected manner. Table 5 explains some of these circumstances and offers one or more solutions for each issue.

Table 5 Troubleshooting Issues and Resolutions

Issue	Resolution
The **Broadcom CNA OneCommand** tab is not visible in the vSphere console.	In the vSphere console, select the **Plug-in** menu and choose **Manage Plug-ins**. In the Plug-in Manager window, check the status of the Broadcom CNA OneCommand Manager for VMware vCenter (Broadcom CNA OneCommand). The status must be Enabled. If it is not, enable it. On the machine where OneCommand CNA Manager for VMware vCenter is installed, make sure the port numbers configured during the installation are open and dedicated to the plug-in server only. No other service should be listening on this port.
When you select the **Broadcom CNA OneCommand** tab in the vSphere console, a message appears indicating that the Adobe Flash player is required.	Ensure the version of Adobe Flash player installed is 11.2 or later. If you have not installed the Adobe Flash player, you can download it from the Adobe website.
When you select the **Broadcom CNA OneCommand** tab in the vSphere console, the Broadcom CNA Device Management tree-view does not display any elements.	Ensure that you have the required privileges to view information in the vSphere console.
There is slow response from OneCommand CNA Manager for VMware vCenter.	Ensure that the following are on the same network: ■ ESXi servers managed by the OneCommand CNA Manager VMware for vCenter server ■ Systems hosting the OneCommand CNA Manager for VMware vCenter server ■ OneCommand CNA Manager for VMware vCenter
Firmware update fails or installing licenses fails on an OCe11000-series adapter.	On the ESXi host, check the firewall settings and ensure that the HTTP/HTTPS ports are open. Use the following command to disable the firewall: `esxcli network firewall unload`
Firmware update fails with the error message `Error reading resource`.	Check the following: **NOTE** Make sure that you can ping the hostname, on which the OneCommand CNA Manager for VMware vCenter is installed, from the ESXi host. If you cannot ping the hostname, either reinstall providing the reachable IP/hostname (with domain) or add the host name to the DNS. ■ Check the memory space in the ESXi host and clean up the old logs.
When you make any changes to the ESXi host, such as plugging cables, unplugging cables, or storage references, OneCommand CNA Manager for VMware vCenter does not reflect the change immediately.	Click **Refresh** in the GUI. If the change is not reflected, restart `sfcb` on the ESXi host using the command: `/etc/init.d/sfcbd-watchdog restart` Click **Refresh** again in the GUI.

Table 5 Troubleshooting Issues and Resolutions (Continued)

Issue	Resolution
On a Windows 7 x64 operating system, executing the CLI commands using the executables results in unnecessary error traces.	The C disk is highly protected; even the administrator account has limited privileges. For example, the contents in the directory `C:\Program Files\` have no write and full control privileges. To remedy the problem: ■ Assign your account write and full control privileges to `C:\Program Files \`. *or* ■ Install OneCommand CNA Manager for VMware vCenter on another disk, for example, D.
When OneCommand CNA Manager for VMware vCenter loads within the vSphere console, it shows a Security warning.	See Section 12.1, Security.
When the VMware Web Client is launched, it requires OneCommand CNA Manager for VMware vCenter to load in another browser tab and accept the Broadcom Certificate.	See Section 12.1, Security.

Chapter 13: Using OneCommand CNA Manager for VMware vCenter Command Line Interface

The CLI client component of OneCommand CNA Manager for VMware vCenter is installed as part of OneCommand CNA Manager for VMware vCenter installation.

13.1 brcmvcpcmd Syntax Usage

- OneCommand CNA Manager for VMware vCenter CLI runs only in TCP/IP mode.
- OneCommand CNA Manager for VMware vCenter CLI can manage OneConnect adapters in systems with VMware ESXi 5.5, 6.0, and 6.5 environments.
- CLI Client commands are supported for Windows only.
- All commands must start with `brcmvcpcmd.exe`. The `brcmvcpcmd.exe` file is available in the OneCommand CNA Manager for VMware vCenter installation directory (which is by default `C:\Program Files\Brcm`). This component is intended for use in scripted operations within batch files. Each time you run this script from the command line, a single operation is performed.
- Most operations retrieve information about an entity on the SAN and show that information on the console.
- Most of the CLI client commands require one or more additional parameters that specify the nature of the command.
- The requested operation must contain at least three characters, or as many as needed to distinguish it from any other operation.
- The parameters must be specified in the order indicated in the syntax.
- Parameters that are not required and can be omitted are in square brackets [ ].
- To run a command at the cluster level, use:

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>
c=<clustername> <ocm_cmd> [<ocm_cmd_arg>...]
```

- To run a command at the host level, use:

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> h=<esx_host>
<ocm_cmd> [<ocm_cmd_arg>...]
```

- For FCoE ports, the WWPN of the adapter must be specified.

NOTE When a WWPN is specified, individual fields are separated by colons (:).

For example, run the following command to display the port attributes for the adapter with the specified WWPN:

```
brcmvcpcmd.exe  v=10.120.121.122 u=Administrator p=password
h=10.120.121.123 portattributes 10:00:00:00:C9:39:6C:9D
```

- For NIC ports, the MAC address must be specified.

NOTE When a MAC address is specified, the fields are separated by a dash (-).

For example, run the following command to show the port attributes for the CNA port with the specified MAC address:

```
brcmvcpcmd.exe  v=10.120.121.122 u=Administrator p=password
h=10.120.121.123 portattributes 00-11-22-33-44-55
```

- For help purposes only, commands are grouped together.

13.2 Help Commands

Help commands include help for a single command or a group of commands.

13.2.1 help (single command)

Syntax

```
brcmvcpcmd.exe [CmdName] help
```

Description

This command shows the help for a specific command.

Parameters

CmdName Any CLI command.

13.2.2 help (group)

Syntax

```
brcmvcpcmd.exe help
```

Description

This command shows the help for a group. These help groups are categorized based on the functionality of the commands. Table 7 shows all supported groups, definitions, and supported commands. You can specify the group name in the help command to find the commands supported for a group.

Example Command

```
brcmvcpcmd.exe  attributes help
```

Parameters

None.

13.3 CLI Command Reference Tables

Table 6 lists CLI commands in alphabetical order, with the corresponding section number for details.

Table 7 lists the CLI commands within a help group. These help groups are categorized based on the functionality of the commands.

Table 6 CLI Commands

Command	Section
changepersonality	13.4.16.1
changewwn	13.4.20.1
clearisnsserver	13.4.14.1
cmgetparams	13.4.6.1
cmmode	13.4.6.2
cmsetbw	13.4.6.3

Table 6 CLI Commands (Continued)

Command	Section
deletedumpfiles	13.4.11.1
driverconfig	13.4.10.1
dump	13.4.11.2
echotest	13.4.9.1
enablebootcode	13.4.5.1
exportsaninfo	13.4.7.1
firmwareupdate	13.4.7.2
getadapterportconfig	13.4.3.1
getbeacon	13.4.9.2
getdcbparams	13.4.8.1
getdriverparams	13.4.10.2
getdriverparamsglobal	13.4.10.3
getdumpdirectory	13.4.11.3
getdumpfilenames	13.4.11.4
getfcfinfo	13.4.12.1
getfipparams	13.4.12.2
getinitiatorproperties	13.4.14.2
getiscsiluns	13.4.14.3
getiscsiportstats	13.4.14.4
getlinkconfig	13.4.15.1
getlunlist	13.4.12.3
getpcidata	13.4.4.1
getpginfo	13.4.8.2
getportstatistics	13.4.4.2
getretentioncount	13.4.11.5
getsessioninfo	13.4.14.5
getvpd	13.4.4.3
getwwncap	13.4.4.4
getxcvrdata	13.4.9.3
hbaattributes	13.4.4.5
installadapterlicense	13.4.2.1
listhbas	13.4.4.6
listprofiles	13.4.14.6
listsessions	13.4.14.7
listvfuncs	13.4.17.1
listvms	13.4.19.1
listvports	13.4.19.2
loopbacktest	13.4.9.4
pcidata	13.4.4.7
portattributes	13.4.4.8

Table 6 CLI Commands (Continued)

Command	Section
readwwn	13.4.20.3
reset	13.4.1.2
restorewwn	13.4.20.4
saveconfig	13.4.10.4
serverattributes	13.4.4.9
setadapterportconfig	13.4.3.2
setbeacon	13.4.9.5
setcableenvp	13.4.9.6
setcnapgbw	13.4.8.3
setdcbpriority	13.4.8.5
setdriverparam	13.4.10.5
setdriverparamdefaults	13.4.10.6
setdumpdirectory	13.4.11.6
setfipparams	13.4.12.4
setlinkconfig	13.4.15.2
setphyportspeed	13.4.4.10
setportenabled	13.4.4.11
setretentioncount	13.4.11.7
setvcred	13.4.1.1
showisnsserver	13.4.14.8
showlicenseadapterfeatures	13.4.2.2
showlicenseadapterid	13.4.2.3
showpersonalities	13.4.16.2
showtarget	13.4.14.9
sriovenable	13.4.17.2
targetmapping	13.4.12.5
tdrtest	13.4.9.7
umcsetlpvid	13.4.6.4
updateisnsserver	13.4.14.10
VEPAenable	13.4.18.1
version	13.4.1.3

Table 7 Help Groups and Supported CLI Commands

Help Group	Supported CLI Commands	Command Descriptions
General group - General commands that can be run on the OneCommand CNA manager application for VMware servers.	■ setvccred ■ reset ■ version	`setvcred` - An optional command that saves vCenter credentials and executes subsequent commands without specifying vCenter credentials. `reset` - Resets the adapter. `version` - Shows the version of the installed CLI.
AdapterLicense group - Commands to install and read personality licenses for an OCe11000-series adapter.	■ installadapterlicense ■ showlicenseadapterfeatures ■ showlicenseadapterid	`installadapterlicense` - Installs the license keys from a license file to enable specific features on the adapter. `showlicenseadapterfeatures` - Shows the licenses available for the adapter to which the specified port belongs. `showlicenseadapterid` - Shows the license ID of the adapter to which a specified port belongs.
Adapter Management group (OCe14000 Adapters only) - Commands to retrieve and set adapter port configuration.	■ getadapterportconfig ■ setadapterportconfig	`getadapterportconfig` - Lists the active and configured adapter port configurations. `setadapterportconfig` - Configures the function protocol assignments for all ports on an adapter.
Attributes group - Commands to read and manage information about an adapter.	■ getpcidata ■ getportstatistics ■ getvpd ■ getwwncap ■ hbaattributes ■ listhbas ■ pcidata ■ portattributes ■ serverattributes ■ setphyportspeed ■ setportenabled	`getpcidata` - Lists PCI register details for FCoE, NIC, and iSCSI ports. `getportstatistics` - Lists statistics for an FCoE port. `getvpd` - Shows the VPD details for the specified port on the adapter. `getwwncap` - Shows the WWN capabilities of the specified port. `hbaattributes` - At the host level, shows adapter attributes for a port on the adapter. `listhbas` - For a cluster or a host, shows a list of manageable adapters. `pcidata` - Lists PCI attributes for a port on the adapter. `portattributes` - Shows a list of all port attributes for the port on the adapter. `serverattributes` - Lists basic information about the host. `setphyportspeed` -Sets the port speed. `setportenabled` - Enables or disables the port status for an FCoE, NIC, or iSCSI port on a host.
Boot Group - Commands to enable an ESXi host to manage the boot environment.	■ enablebootcode	`enablebootcode` - Enables or disables the bootBIOS state on a given port.
Channel Management group - Commands that enable and manage an ESXi host to handle Channel Management.	■ cmgetparams ■ cmmode ■ cmsetbw ■ umcsetlpvid	`cmgetparams` - Lists channel management information for the adapter. `cmmode` - Enables or disables channel management on the adapter. `cmsetbw` - Sets bandwidth for channels on the physical port. `umcsetlpvid` - Sets the logical port VLAND IDs for each channel on a physical port.

Table 7 Help Groups and Supported CLI Commands (Continued)

Help Group	Supported CLI Commands	Command Descriptions
Cluster group - Commands that can be run on a cluster.	■ exportsaninfo ■ firmwareupdate ■ listhbas	`exportsaninfo` - Exports SAN information related to adapters in all the hosts in a cluster. `firmwareupdate` -Updates the firmware on the adapters found in a VMware cluster or the ESXi host. `listhbas` - Lists manageable adapters.
DCB group - Commands to enable an ESXi host to read and manage data center bridging parameters for an adapter.	■ getdcbparams ■ getpginfo ■ setcnapgbw ■ setdcbparam ■ setdcbpriority	`getdcbparams` - Extracts the DCB parameters associated with a given FCoE, NIC, or iSCSI port. `getpginfo` - Lists the priorities and bandwidth percentages for all the priority groups for the port. `setcnapgbw` - Sets the bandwidth percentage of a priority group on a port. `setdcbparam` - Configures DCB and LLDP settings. `setdcbpriority` - Sets the priorities for a priority group.
Diagnostic group - Commands to run diagnostic tests for an adapter or port.	■ echotest ■ getbeacon ■ loopbacktest ■ setbeacon ■ setcablenvp ■ tdrtest	`echotest` - Runs an echo test on a port. `getbeacon` - Shows the current beacon state for a port on an adapter. `loopbacktest` - Runs a loopback test on a port. `setbeacon` - Turns the beacon state on or off for a port on an adapter. `setcablenvp` - On a 10GBASE-T adapter, sets the NVP (required for the `TDRTest` command) for the cable that connects to the physical port associated with the WWPN or MAC. `tdrtest` - On a 10GBASE-T adapter, runs a test to verify the integrity of the link.
Link Reconfiguration group- If the adapter supports this, commands to manage link reconfiguration.	■ getlinkconfig ■ setlinkconfig	`getlinkconfig` - Lists the active and configured configuration IDs and lists the available link configurations. `setlinkconfig` - Sets the configured link configuration ID.
DriverParameters group - Commands to enable an ESXi host to read and manage driver parameters for the host adapters.	■ driverconfig ■ getdumpdirectory ■ getdriverparamsglobal ■ saveconfig ■ setdriverparam ■ setdriverparamdefaults	`driverconfig` - Sets all driver parameters to the values in the `.dpv` file on an ESXi host. `getdumpdirectory` - Lists driver parameters for a port. `getdriverparamsglobal` - Lists global driver parameters for a port. `saveconfig` - Saves an FCoE adapter's driver parameters to a file. `setdriverparam` - Sets a driver parameter for a port and designates the scope of the change. `setdriverparamdefaults` - Restores all driver parameters to the default value, at the port or global level (temporarily or permanently).

Table 7 Help Groups and Supported CLI Commands (Continued)

Help Group	Supported CLI Commands	Command Descriptions
Dump group - Commands to manage dump files for a selected adapter. A dump files is useful when troubleshooting.	■ deletedumpfiles ■ dump ■ getdumpdirectory ■ getdumpfilenames ■ getretentioncount ■ setdumpdirectory ■ setretentioncount	`deletedumpfiles` - Deletes all diagnostic dump files for a port. `dump` - Performs a core dump of a local port. `getdumpdirectory` - Shows a dump file directory for a port in the host. `getdumpfilenames` - Lists all dump filenames for a port. `getretentioncount` - Shows the maximum number of diagnostic dump files to keep for a port. `setdumpdirectory` - Sets the dump directory for all adapters in the server. `setretentioncount` - Specifies the maximum number of diagnostic dump files for the adapter.
FCoE group - Commands to enable an ESXi host to read and manage FCoE functions for an adapter.	■ getfcfinfo ■ getfipparams ■ getlunlist ■ setfipparams ■ targetmapping	`getfcfinfo` - Lists Fibre Channel Forwarding information for an adapter in FCoE mode. `getfipparams` - Lists FIP parameters for an adapter in FCoE mode. `getlunlist` - Lists the LUNs attached to the FCoE target of the specified port. `setfipparams` - Sets FIP parameters for an adapter in FCoE mode. `targetmapping` - Lists the targets attached to the FCoE of the specified port.
Firmware group - A command for an ESXi host to update the firmware for an adapter or a port.	■ firmwareupdate	`firmwareupdate` - Updates the firmware on adapters in a VMware cluster or ESXi host.
iSCSI group - Commands to enable an ESXi host to read and manage iSCSI functions for an adapter.	■ clearisnsserver ■ getinitiatorproperties ■ getiscsiluns ■ getiscsiportstats ■ listprofiles ■ listsessions ■ showisnsserver ■ showtarget ■ updateisnsserver	`clearisnsserver` - Clears the configured ISNS server and disables ISNS target discovery. `getinitiatorproperties` - Lists all the initiator login options for the specified port. `getiscsiluns` - Lists the LUNs attached to the iSCSI target of an iSCSI port. `getiscsiportstats` - Lists the port statistics for an iSCSI port. `listprofiles` - Lists available profile IDs and their descriptions for an adapter, including the currently active and configured profile IDs. `listsessions` - Lists all the sessions on a specified target. `showisnsserver` - Shows the currently configured Internet Storage name Server. `showtarget` - Lists the iSCSI target details of the adapter to which the iSCSI port belongs. `updateisnsserver` - Updates the Internet Storage name Server.

Table 7 Help Groups and Supported CLI Commands (Continued)

Help Group	Supported CLI Commands	Command Descriptions
Personality group - Commands to enable an ESXi host to read and change personalities for an adapter **NOTE** Not supported for OCe14000–series adapters. Use the setadapterportconfig command for these adapters.	■ changepersonality ■ showpersonalities	`changepersonality` - Changes the valid personality of the adapter to which the specified port belongs. `showpersonalities` - Lists the personalities of the adapter to which the specified port belongs.
SR-IOV group - Commands that allow multiple VFs to be created on a NIC PF for an adapter port. **NOTE** SR-IOV is not supported with UMC channel management.	■ listvfuncs ■ sriovenable	`listvfuncs` - Lists all of the virtual functions discovered. `sriovenable` - Enables or disables SR-IOV support for an adapter's physical port.
VEPA group - A command that enables or disables the VEPA state on an OCe14000-series NIC port. **NOTE** VEPA commands require that SR-IOV is enabled.	■ VEPAenable	VEPA`enable` - Enables or disabled VEPA on an OCe14000-series NIC port.
Virtual Machines group - Commands to enable an ESXi host to find Virtual Machines and their attached ports.	■ listvms ■ listvports	`listvms` - Lists all virtual machines and their information for all manageable ports. `listvports` - Lists all virtual ports on the specified physical port.
WWWNManagement group - Commands to enable ESXi host to read and manage the WWN details for the port on an adapter.	■ changewwn ■ getwwncap ■ readwwn ■ restorewwn	`changewwn` - Changes the WWN of the port. `getwwncap` - Lists the WWN capabilities of the port. `readwwn` - Lists the WWN details of the port and category. `restorewwn` - Restores the WWN value of the port.

13.4 Group Commands and CLI Command Descriptions

This section provides syntax and descriptions for group and CLI commands.

13.4.1 General Group Commands

The general group commands save vCenter credentials, reset the adapter, and show the version of the installed CLI.

13.4.1.1 setvccred

13.4.1.1.1 Syntax

```
brcmvcpcmd.exe setvccred v=<vCenter IP/Name> u=<username> p=<password>
```

13.4.1.1.2 Description

While executing a set of CLI commands, you must enter vCenter credentials repeatedly. By executing the `setvccred` command first, you can save these credentials including vCenter server name/IP, username and password to a file in an encrypted format and execute subsequent commands without use of vCenter credentials.

Using this command is optional. You can continue to execute the commands providing all credentials.

NOTE This command does not apply to ports or adapters.

13.4.1.1.3 Parameters

v The vCenter server IP address.

u The user name for the vCenter server.

p The user password for the vCenter server.

13.4.1.1.4 Examples

Execute `setvccred` first:

```
brcmvcpcmd.exe setvccred v=12.345.678.901 u=username p=password
```

Subsequent commands can be:

New format:

```
brcmvcpcmd.exe h=12.345.678.123 listhbas
```

or

Old format:

```
brcmvcpcmd.exe v=12.345.678.901 u=username p=password h=12.345.678.123 listhbas
```

13.4.1.2 reset

This command resets the adapter. An adapter reset can require several seconds to complete, especially for remote devices. When the reset is completed, the system command prompt is displayed.

NOTES

- Supported only for FCoE ports. Not supported for NIC and iSCSI ports.
- This command only resets the driver to update changed driver parameters that require a driver reset. It does not cause a hardware reset of the adapter port.

13.4.1.2.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
reset <WWPN>
```

13.4.1.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.

13.4.1.3 version

This command shows the version of the CLI installed.

13.4.1.3.1 Syntax

```
brcmvcpcmd.exe version
```

13.4.1.3.2 Parameters

None.

13.4.2 AdapterLicense Commands

The Adapter License commands install license keys and show license information.

NOTE These commands apply to OCe11000-series adapters only.

13.4.2.1 installadapterlicense

This command installs the license keys from a license file to enable specific features on the adapter. You must specify either the WWPN or the MAC address.

13.4.2.1.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
installadapterlicense <WWPN | MAC> [Licence File Path]
```

13.4.2.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.
License File Path	The local file path of the license file to be installed.

13.4.2.2 showlicenseadapterfeatures

This command shows the licenses available for the adapter to which the specified port belongs.

13.4.2.2.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
showlicenseadapterfeatures <WWPN | MAC>
```

13.4.2.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.

13.4.2.3 showlicenseadapterid

This command shows the license ID of the adapter to which the specified port belongs.

13.4.2.3.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
showlicenseadapterid <WWPN | MAC>
```

13.4.2.3.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.

13.4.3 Adapter Management Commands

The Adapter Management commands show adapter port configurations, set adapter port configurations, and enable or disable multichannel.

NOTE Adapter Management Commands support only OCe14000-series adapters.

13.4.3.1 getadapterportconfig

After rebooting, the command shows the active and configured adapter port configurations (including the total number of ports, the maximum number of FCoE and iSCSI functions that can be defined, and whether the adapter allows mixed mode storage and concurrent mode storage), and the SR-IOV state. The number of functions per port depends upon the current setup of the adapter (for example, mixed or concurrent storage mode, multichannel, NPar). Depending upon the presence of NParEP on a host, this command may show more functions per port than are actually running on the host.

NOTES

- For multichannel and NPar configurations, this command only shows the port protocol assignments. For the entire multichannel configuration, use the cmgetparams command.
- If you run the `getadapterportconfig` command after changing the link configuration ID, the configured Reboot Configuration displayed reflects the default port configuration for the newly configured link configuration ID.

13.4.3.1.1 Syntax

```
brcmvcpcmd.exe v=<vcenter server> u=<vc username> p=<vc password> h=<esx host>
GetAdapterPortConfig <MAC|WWPN>[table]
```

13.4.3.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC or iSCSI port.
table	The table option displays the output in tabular form, including all function and available protocols for each function.

13.4.3.2 setadapterportconfig

The setadapterportconfig command configures the function protocol assignments for all ports on an adapter. Additionally, you can use this command to restore the adapter to its factory default profile and settings. The number of ports and function protocol assignments specified in this command depends upon the available ports, storage mode (mixed or concurrent), and multichannel/NPar state. Depending upon the presence of NParEP on a host, this command may allow the configuration of more functions per port than is available on the PCI bus. This means a lesser number of functions per port than are configured by this command may be noticed after rebooting the host. The SR-IOV state can be enabled or disabled using the optional SR-IOV parameter.

13.4.3.2.1 Syntax

```
brcmvcpcmd.exe v=10.192.000.000 u=root p=vmware h=10.192.203.195
SetAdapterPortConfig <MAC|WWPN> defaults
```

or

```
brcmhbacmd.exe SetAdapterPortConfig <MAC|WWPN> p0=P0Config [p1=P1Config]
[p2=P2Config][p3=P3Config] [pid=ProfileID] [mctype=Type]
[NParEpMode=state][sriov=SriovState]
```

13.4.3.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
MAC	MAC address of any iSCSI or NIC function on the adapter.
WWPN	World Wide Port Name of any FCoE function on the adapter.
tab	The table option displays the output in tabular form, including all function and available protocols for each function.
p0=fcnCfg	1st port's function(s)
p1=fcnCfg	2nd port's function(s) (required for 2 and 4 port adapters)
p2=fcnCfg	3rd port's function(s) (required for 4 port adapters)
p3=fcnCfg	4th port's function(s) (required for 4 port adapters)
fcnCf	f0[,f1[,...fM]] where f0: 1st function protocol on port (must be NIC) f1: 2nd function protocol on port (optional) fM: Mth function protocol on port (optional) The number of functions (that is, the value of M) per port are model and multichannel type specific.
pid	Optionally specifies the profile ID to use for the personality. This parameter is required for NIC only.
Type	Optional multichannel type specification (specify none to disable multichannel). This parameter is required to specify multichannel. If not specified, no change to the configured multichannel type is made.
State	NParEP Mode state. This parameter can only be specified when mctype=NPar. Valid states are: 0 = Disabled 1 = Enabled
SriovState	SR-IOV state Valid states are: 0 = Disabled 1 = Enabled

13.4.3.2.3 Syntax Rules for the setadapterportconfig Command

- No spaces are allowed between around the = or commas in when specifying a port (for example, p0 = nic, iscsi would cause an error).
- The only choice for the f0 parameter is NIC. Choices for the f1–fm parameters are None, iSCSI, or FCoE (or NIC for multichannel configurations).
- This command fails if you specify two storage protocols on a port and the concurrent storage mode is not available.
- This command fails if you specify three storage functions on a port.
- This command fails if you specify two of the same storage functions on a port.
- The number of functions (f) you can specify depend upon the adapter model and the configuration type (for example, non-multichannel, multichannel, NPar, etc.). If you specify an incorrect number of functions, the command fails.
- You can only specify NIC as the f0 parameter if no other f parameters are specified. In addition, you cannot specify NIC with multichannel configurations.
- The pid=profile_id option is required for NIC only configurations. It you specify an incorrect number of functions, the command fails.

- If the mctype parameter is specified, it must be an available multichannel type (see CMGetParams command for available types) or set to None to disable multichannel. If the mctype parameter is not specified, the currently configured multichannel type (including None) remains unchanged.
- If one or more functions are set to None and they are the last functions on a port, they do not need to be specified (for example, p0=nic,iscsinone can be shortened to p0=nic,iscsi).
- You can only specify NParEPMode with the mctype-NPar parameter or when the current configured MC type is NPar.
- When you disable NPar (for example mctype=none), the NParEP Mode is automatically disabled. You must reboot to activate a new port configuration.
- SR-IOV can be enabled or disabled only when the mctype is a non-multichannel type such as a single personality, NPar, or Custom.
- Port removal is not supported on Dell NPar adapters. An error message is displayed if port removal is attempted.

13.4.3.2.4 Example Commands

Mixed Mode

```
brcmvcpcmd.exe v=10.192.000.000 u=root p=vmware h=10.192.000.000
setadapterportconfig 00-00-c9-12-34-56 p0=nic,iscsi p1=nic,fcoe p2=nic,fcoe
p3=nic
```

Concurrent Mode with Two Ports

```
brcmvcpcmd.exe v=10.192.000.000 u=root p=vmware h=10.192.000.000
setadapterportconfig 00-00-c9-12-34-56 p0=nic,iscsi,fcoe p1=nic,fcoe
```

13.4.3.3 Enabling and Disabling Multichannel

Use the Type parameter to enable and disable multichannel. When enabling multichannel, the additional function protocols must be specified for the extra channels that are not available when multichannel is disabled.

13.4.3.3.1 Examples for Enabling and Disabling Multichannel

To enable multichannel:

```
brcmvcpcmd.exe v=10.192.000.000 u=root p=vmware h=10.192.000.000
setadapterportconfig 00-00-c9-12-34-56 p0=nic,iscsi,nic,nic p1=nic,fcoe,nic,nic
p2=nic,fcoe,nic,nic p3=nic,iscsi,nic,nic mctype=umc
```

To disable multichannel:

```
brcmvcpcmd.exe v=10.192.000.000 u=root p=vmware h=10.192.000.000
setadapterportconfig 00-00-c9-12-34-56 p0=nic,iscsi p1= nic,fcoe p2=nic,fcoe
p3=nic,iscsi mctype=none
```

13.4.4 Attribute Commands

The Attribute commands show port information, adapter attributes, pcidata, and server attributes. These commands also enable a port on a host and set physical port speed.

13.4.4.1 getpcidata

This command shows PCI register details for FCoE, NIC, and iSCSI ports.

13.4.4.1.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password> h=<esx_host>
getpcidata <WWPN | MAC>
```

13.4.4.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.

13.4.4.2 getportstatistics

This command extracts the statistics for a designated FCoE port. The optional clear parameter clears the 10GBASE-T counters for an OCe14000-series port.

NOTE If the clear parameter is specified for other than an OCe14000-series adapter, an error message is displayed.

13.4.4.2.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password> h=<esx_host>
getportstatistics <WWPN> [clear]
```

13.4.4.2.1.1 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
clear	An optional parameter that clears the counters for an OCe14000-series 10GBASE-T port.

13.4.4.3 getvpd

This command shows the VPD details for the specified port on the adapter.

13.4.4.3.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
getvpd <WWPN | MAC>
```

13.4.4.3.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.

13.4.4.4 getwwncap

This command shows the WWN capabilities of the specified port.

13.4.4.4.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
getwwncap <WWPN>
```

13.4.4.4.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.

13.4.4.5 hbaattributes

This command shows a list of all adapter attributes for the specified port on the adapter. This command is supported only at the host level.

13.4.4.5.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> <h=ESXhostIP>
hbaattributes  <WWPN | MAC>
```

13.4.4.5.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.

13.4.4.6 listhbas

This command shows a list of the manageable adapters found by remote discovery.

NOTE The output of this command also displays the current and the permanent MAC address (Dell only).

13.4.4.6.1 Syntax

For a cluster:

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> <c=clustername>
listhbas
```

For a host:

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> <h=ESXhostIP>
listhbas
```

13.4.4.6.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
c	The cluster name in the vSphere console.
h	The IP address of the ESXi host.

13.4.4.7 pcidata

This command shows the PCI attributes for the port specified on the adapter.

13.4.4.7.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> <h=ESXhostIP>
pcidata <WWPN | MAC>
```

13.4.4.7.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.

13.4.4.8 portattributes

This command shows a list of all port attributes for the port on the adapter.

NOTE The output of this command also displays the current and the permanent MAC address (Dell only).

13.4.4.8.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
PortAttributes <WWPN | MAC>
```

13.4.4.8.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port. If the port supports VEPA, the output contains the line VEPA State: <enabled | disabled> (OCe14000-series adapters only).

13.4.4.9 serverattributes

This command is used to retrieve basic information about the host such as the operating system version and CIM provider version.

13.4.4.9.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
serverattributes
```

13.4.4.9.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.

13.4.4.10 setphyportspeed

This command sets the port speed.

13.4.4.10.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
SetPhyPortSpeed <WWPN | MAC> <Mode> [Speed [Length]]
```

13.4.4.10.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the CNA port.
Mode	Mode number: 0 = Default 1 = Auto-Negotiate 2 = Force When Mode = 0, the Speed and Length arguments are ignored. When Mode = 1, the Length arguments is ignored.
Speed	Speed string that is required for Mode = 1 or 2 (s=10Gb/1Gb/100Mb).
Length	DAC cable length in meters (0-10, required for 10Gb SFP Plus interface type).

13.4.4.10.3 Example

The following command configures the port to a forced speed of 1 Gb/s with a cable length of 4 meters.

```
brcmvcpcmd.exe v=11.190.705.200 u=username p=password h=11.190.206.300
setphyportspeed 2 1Gb 4
```

13.4.4.11 setportenabled

This command is used to enable or disable an FCoE, NIC, or iSCSI port on a host.

NOTE Ensure that all I/O traffic on the port is stopped before disabling the port.

13.4.4.11.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setportenabled <WWPN | MAC> <Flag>
```

13.4.4.11.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.
Flag	0 = Disabled 1 = Enabled

13.4.5 Boot Command

The enablebootcode command enables or disables the bootBIOS state on a port.

13.4.5.1 enablebootcode

This command enables or disables the bootBIOS state on a given port by enabling or disabling the boot code on the FCoE adapter, the PXE BIOS or the iBFT.

13.4.5.1.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
enablebootcode <WWPN | MAC> [Flag][NetworkBootMethod]
```

13.4.5.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.
Flag	0 = disable BootBIOS state. 1 = enable BootBIOS state.
NetworkBootMethod	Network boot method to be used by the NIC. iBFT PXE

13.4.6 Channel Management Commands

The Channel Management commands show channel management information, enable or disable channel management on an adapter, set a channel's minimum and maximum bandwidth, and set the Logical Port VLAN ID for each channel.

13.4.6.1 cmgetparams

This command shows the following information for the adapter:

- The active (booted) mode (NPar or disabled) for the adapter
- The channel management state of the adapter
- The configured state of the adapter

- The configured channel management mode, that is N/A if disabled
- The available channel management modes
- The port's channels and their channel properties:
  - The protocol type (Because the channels are based on all physical ports, some adapters show only two or three channels instead of four because that is all the physical port allows.)
  - The MAC address
  - The LPVID
  - The minimum and maximum bandwidth

13.4.6.1.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
cmgetparams <WWPN | MAC>
```

13.4.6.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	WWPN of an FCoE function on the physical port.
MAC	MAC address of any NIC function on the physical port.

13.4.6.1.3 Examples

All NIC

```
brcmvcpcmd.exe v=10.20.30.40 u=username p=password h=10.11.12.13 cmgetparams
00-00-c9-12-34-56
```

Multichannel enabled, UMC, Mixed Mode

```
brcmvcpcmd.exe v=10.192.000.000 u=root p=vmware h=10.192.000.000 cmgetparams
00-00-c9-12-34-56
brcmvcpcmd.exe v=10.192.000.000 u=root p=vmware h=10.192.000.000 cmgetparams
00-00-c9-12-34-59
```

NPar enabled, 4 functions per port

```
brcmvcpcmd.exe v=10.192.000.000 u=root p=vmware h=10.192.000.000 cmgetparams
00-00-c9-12-34-56
```

13.4.6.2 cmmode

This command enables or disables channel management on an adapter.

NOTE A system reboot is required for the change to take effect.

13.4.6.2.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password> h=<esx_host>
cmmode  <WWPN | MAC Address> <None | Type>
```

13.4.6.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	WWPN of an FCoE function on the physical port.
MAC	MAC address of any NIC function on the physical port.
None|Type	Specify None to disable channel management, or specify a channel management type to enable channel management.

13.4.6.2.3 Example

This command shows that the channel management on an adapter containing a NIC function with a MAC address of [00-00-c9-ad-ad-aa] is disabled.

```
brcmvcpcmd.exe v=10.20.30.40 u=username p=password h=10.11.12.13 CMMode
00-00-c9-bb-cc-aa None
```

13.4.6.3 cmsetbw

This command sets the minimum and maximum bandwidth for up to N channels each on the physical port (where N is number of channels enabled on the adapter). The total of the minimum bandwidths must add up to 100%. The total of the maximum bandwidths must be greater than or equal to the minimum bandwidths and have a maximum of 100%. When only three channels are available, specifying Min3,Max3 causes an error. When only two channels are available, specifying Min2,Max2 and Min3,Max3 causes an error.

When in channel management mode, you can disable a channel's logical link by setting the minimum bandwidth and the maximum bandwidth to 0. If a channel's protocol is set to None, the minimum and maximum bandwidth must be 0. Setting the minimum bandwidth to a value greater than 0 enables the channel if it was previously disabled.

NOTES

- This command is not available for 1 Gb ports.
- For OCe14000-series adapters (with NPar disabled), the total minimum bandwidth must add up to 0 or 100%. If the adapter is enabled for NPar, the total minimum bandwidth must add up to 100%.
- The maximum bandwidth of a channel must be greater than or equal to the minimum bandwidth for that channel up to a maximum of 100.
- This command fails if the current channel management mode is vNIC1 or UFP, or None.
- When channel management is enabled, a system reboot is not required for the change take effect.

13.4.6.3.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password> h=<esx_host>
cmsetbw <WWPN | MAC> <Min0,Max0> <Min1,Max1>...[MinN,MaxN]
```

13.4.6.3.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FCoE function on the physical port.
MAC	The MAC address of a NIC function on the physical port.
Bandwidth Value	
	NOTE The number of channels and the maximum value of N are adapter model specific.
Min0,Max0	The minimum and maximum bandwidths for channel 0.
Min1,Max1	The minimum and maximum bandwidths for channel 1.
MinN,MaxN	The minimum and maximum bandwidths for channel 2.

13.4.6.3.3 Example

```
brcmvcpcmd.exe v=10.20.30.40 u=username p=password h=10.11.12.13CMSetBW
00-00-c9-bb-cc-aa 25,50 50,75 0,0 25,100
```

13.4.6.4 umcsetlpvid

This command sets the Logical Port VLAN IDs (LPVID) for each channel on a physical port. LPVID values are in the range of 2-4094. When only three channels are available, specifying LPVID3 causes an error. When only two channels are available, specifying LPVID2 and LPVID3 causes an error.

NOTES

- LPVID values are in the range of 2 to 4094.
- This is the only command available to set LPVID values.
- Every NIC channel on a physical port must have a unique LPVID.
- This command is not available for FCoE, and iSCSI channels because they are storage channels. LPVIDs for FCoE and iSCSI channels must be set to zero.
- This command is not available for 1 Gb ports.
- If UMC is disabled when this command is executed, a warning message is displayed stating that UMC is disabled and must be enabled for any changes to take effect.
- A system reboot is not required for the change take effect.
- LPVID properties are only associated with UMC. The `umcsetlpvid` command fails if the current channel management mode is not UMC.
- SR-IOV is not supported with UMC channel management.

13.4.6.4.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>
h=<esx_host>umcsetLPVID <WWPN | MAC> <LPVID0> <LPVID1> [LPVIDN]
```

13.4.6.4.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
MAC	MAC address of any NIC function on the physical adapter port.
LPVID State	
LPVID0	LPVID for channel 0
LPVID1	LPVID for channel 1
LPVIDN	LPVID for channel N

NOTE The number of channels and the maximum value of N are adapter model specific.

13.4.6.4.3 Example

```
brcmvcpcmd.exe v=10.192.000.000 u=root p=vmware h=10.192.000.000 umcsetlpvid
00-00-c9-12-34-56 1001 1002 1003 1004
```

13.4.7 Cluster Commands

Cluster commands export SAN information, update firmware, and list adapters.

13.4.7.1 exportsaninfo

This command exports all SAN information related to adapters in all the hosts in a cluster. This command is supported only at the cluster level.

NOTE Due to the amount of information that must be obtained and reported, this command can take a long time on large SAN configurations. You can redirect this output to a file with a proper extension, `.xml` for XML-formatted files and `.csv` for CSV-formatted files.

13.4.7.1.1 Syntax

```
brcmvcpcmd.exe v=<vcenter server> u=<vc_username> p=<vc_pwd>  <c=clustername>
exportsaninfo [format]
```

NOTE [format] is optional. If the format parameter is specified as csv, the adapter information is shown in CSV format. If the format parameter is specified as xml, the adapter information is shown in XML format. Leaving the format parameter blank shows the data in XML format.

13.4.7.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
c	The cluster name.
Format	csv – The output information in CSV format. xml – The output information in XML format (default).

13.4.7.2 firmwareupdate

This command updates the firmware on the adapters found in the specified VMware cluster or the ESXi host.

13.4.7.2.1 Procedure

1. Run the firmwareupdate command. A list of adapter serial numbers are displayed, applicable to the firmware file specified.
2. Select the adapter or port option to use for the update. **List**, **Range**, **All** or **Choice** options are displayed.
3. The firmware update process begins and returns the result for each adapter or port.
4. If you press **<Ctrl> + <C>** and the firmware update process has started on any adapter or port, the update continues. But if the firmware update process is queued, the update is canceled.

   NOTE The optional argument [all|MACs] updates the firmware without any user prompt. Either all or a combination of MACs can be given as an option.

   You can view the status of submitted firmware jobs on the OneCommand CNA Manager VMware for vCenter, **Maintenance** tab of the cluster or host.

13.4.7.2.2 Syntax

For a cluster:

```
brcmvcpcmd.exe v=<vcenter server> u=<vc_username> p=<vc_pwd> c=<clustername>
firmwareupdate [all|MACs] <filelocation>
```

For a host:

```
brcmvcpcmd.exe v=<vcenter server> u=<vc_username> p=<vc_pwd> h=<ESXhostIP>
firmwareupdate [all|MACs] <filelocation>
```

13.4.7.2.3 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
c	The cluster name in the vSphere console.
h	The IP address of the ESXi host.
all|MACs (optional)	The all optional argument updates all compatible adapters or ports without any user prompt. The MACs optional argument updates the port MACs belonging to a specified cluster or host without any user prompt.
file_location	The firmware file path on the local disk.

13.4.7.3 listhbas

See Section 13.4.4.6, listhbas.

13.4.8 DCB Commands

DCB commands show DCB parameter information, show priority and bandwidth percentages for all priority groups, configure the DCB and LLDP settings on an adapter, and set bandwidth percentage of a priority group on a port.

13.4.8.1 getdcbparams

This command extracts the DCB parameters associated with a given FCoE, NIC, or iSCSI port.

13.4.8.1.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getdcbparams <WWPN | MAC>
```

13.4.8.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC or iSCSI port.

13.4.8.2 getpginfo

This command shows the priorities and bandwidth percentages for all the priority groups for the port. Additionally, this command shows the number of priority groups supported by an adapter.

13.4.8.2.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getpginfo <WWPN>
```

13.4.8.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.

13.4.8.3 setcnapgbw

This command sets the bandwidth percentage of a priority group on a given FCoE, NIC, or iSCSI port according to the following rules:

- BW0 to BW7 must add up to 100%.
- Bandwidth can be assigned to a priority group that has priorities.

13.4.8.3.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setcnapgbw <WWPN | MAC> <BW0>...<BW7>
```

13.4.8.3.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.
BW0 to BW7	The bandwidth for priority group 0 to 7.

13.4.8.3.3 Example

This command sets the bandwidth of PGID0 to 50 PGID1 to 50 and the rest to 0%.

```
brcmvcpcmd.exe setcnapgbw 10:00:00:00:c9:3c:f7:88 50 50 0 0 0 0 0 0
```

13.4.8.4 setdcbparam

This command configures the DCB and LLDP settings.

13.4.8.4.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setdcbparam <WWPN | MAC> <Param> <Value>
```

13.4.8.4.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC or iSCSI port.

Param and Value Parameters

DCB Parameters and Values (1 = Enabled and 0= Disabled)

DCBXState	The current DCBX (Data Center discovery and Capability exchange protocol) state.
PFCEnable	If set to 1, there is flow control in both directions (Tx and Rx).
FCoEPriority	You must specify a single priority (0 to 7). [only for FCoE ports]
iSCSIPriority	You must specify a single priority (0 to 7). [only for iSCSI ports]
PFCPriorities	Specify PFCPriorities as a single priority or as a list of comma separated values (up to 7 values ranging from 0 to 7).
Default	Setting the SetDCBParam arguments to default sets all CNA DCB parameters (including priority groups) to their default values.

Link Layer Discovery Protocol (LLDP) Parameters and Values

TxState	Transmit State – 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.
RxState	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.
TxPortDesc	Transmit Port Description – Provides a description of the port in an alpha-numeric format. If the LAN device supports RFC 2863, the value equals the ifDescr object.
TxSysDesc	Transmit System Description – 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. If the LAN device supports RFC 3418, the value equals the sysDescr object.
TxSysName	Provides the system's assigned name in an alpha-numeric format. If the LAN device supports RFC 3418, the value equals the sysName object.
TxSysCap	Transmit System Capabilities - Indicates the primary functions of the device and whether 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. Where multiple values are possible, specify them with a comma-separated list.

13.4.8.4.3 Example Command

```
brcmvcpcmd v=12.345.678.901 u=username p=password h=12.345.678.123 setdcbparam
00-00-c9-3c-f7-88 fcoepriority 3
```

13.4.8.5 setdcbpriority

This command sets the priorities for a priority group.

13.4.8.5.1 General Rules

- The priorities range from 0 to 7.
- The PGIDs range from 0 to 7.
- A priority can exist in only one priority group.
- All priorities must appear once in any of the eight (PG0-PG7) priority groups.
- Each set of priorities for a group must be separated by a space.

- Specify multiple priorities for the same group by a comma-separated list.
- To specify none, use - for the argument.
- The same priority values cannot be specified to different groups.
- All priorities (0 to 7) must be assigned to a PGID.
- Not all PGIDs must be assigned a priority.
- Not all adapters support two PFC priorities and eight priority groups. For adapters, if you exceed the PFC priorities or priority groups, an error message appears.

13.4.8.5.2 FCoE Port Specific Rules

- Minimum of one and a maximum of two PFC priorities can be configured.
- One of the PFC priorities must match FCoE priority.
- The additional PFC priority must be assigned to a priority group that has no other priorities.
- The FCoE priority must be assigned to a priority group that has no other priorities.

13.4.8.5.3 iSCSI Port Specific Rules

- Minimum of one and a maximum of two PFC priorities can be configured.
- One of the PFC priorities must match iSCSI priority.
- The additional PFC priority must be assigned to a priority group that has no other priorities.
- The iSCSI priority must be assigned to a priority group that has no other priorities.

13.4.8.5.4 NIC Port Specific Rules

- Only one PFC priority can be configured.
- The PFC priority must be assigned to a priority group that has no other priorities.

13.4.8.5.5 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setdcbpriority <WWN | MAC> <PFC Priorities> <Priorities of PGID0> < Priorities of
PGID1>...<Priorities of PGID7>
```

13.4.8.5.6 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC or iSCSI port.
PFCPriorities	PFC priority that is a comma separated list of up to 7 values ranging from 0-7.
Priorities of PGID	Priority group membership that is a comma separated list of priorities ranging in value for 0-7.

13.4.8.5.7 Example Command

```
brcmvcpcmd v=12.345.678.901 u=username p=password h=12.345.678.123 setdcbpriority
10:00:00:00:c9:3c:f7:88 3 0,1,2,4,5,6,7 3 0 0 0 0 0 0
```

13.4.9 Diagnostic Commands

The Diagnostic commands run diagnostics including POST, loopback, and TDR tests. Diagnostic commands also show and set beacon status, and set the nominal velocity of propagation (N-V-P) required for the TDR test (10GBASE-T adapters only).

13.4.9.1 echotest

This command runs a loopback test on a given FCoE, NIC, or iSCSI port.

13.4.9.1.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
echo <WWN | MAC> <Destination WWPN> <Count> <StopOnError> [Pattern]
```

13.4.9.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC or iSCSI port.
Destination WWPN	WWPN of the destination (echoing) adapter.
Count	The number of times to run the test (0 = run test infinitely, Range = 1 to 99,999).
StopOnError	Checks if the test must be halted on error. 0 = No halt 1 = Halt
Pattern (optional)	One to eight hexadecimal bytes to use for loopback data (for example: 1a2b3c4d).

13.4.9.1.3 Example

```
brcmvcpcmd.exe  v=10.20.30.40 u=user p=password  h=1.2.3.4 echo
10:00:00:c9:12:34:56 10:00:00:c9:ab:cd:ee 100 1 1a2b3c4d5e
```

13.4.9.2 getbeacon

This command shows the current beacon state, on or off.

13.4.9.2.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getbeacon <WWPN | MAC>
```

13.4.9.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC or iSCSI port.

13.4.9.3 getxcvrdata

This command shows transceiver data such as vendor name and serial number.

13.4.9.3.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getxcvrdata <WWN | MAC>
```

13.4.9.3.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC or iSCSI port.

13.4.9.4 loopbacktest

This command is used to run a loopback test on a given FCoE, NIC, or iSCSI port.

13.4.9.4.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
loopbacktest <WWPN | MAC> <Type> <Count> <StopOnError> [Pattern]
```

13.4.9.4.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC or iSCSI port.
Type	The type of loopback test to run: 2 = External LoopBack Test (requires loopback plug) 3 = Host to DDR DMA (OneConnect adapters) Loopback Test 4 = PHY Loopback Test 5 = MAC Loopback Test
Count	The number of times to run the test (0 = run test infinitely, Range = 1 to 99,999)
StopOnError	Checks if the test must be halted on error. 0 = No halt 1 = Halt
Pattern (optional)	1 to 8 hexadecimal bytes to use for loopback data (for example: 1a2b3c4d)

13.4.9.4.3 Example Command

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>
h=<esx_host> loopbacktest 00-00-c9-93-2f-9f 4 120 0
```

13.4.9.5 setbeacon

This command turns the beacon on or off.

13.4.9.5.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setbeacon <WWN | MAC> <BeaconState>
```

13.4.9.5.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	MAC address of the NIC or iSCSI port.
BeaconState	New state of the beacon: 0 = Off 1 = On

13.4.9.6 setcablenvp

This command sets the nominal velocity of propagation (N-V-P) required for the TDR test, for the cable that connects to the physical port associated with the WWPN or MAC

NOTE This command supports only 10GBASE-T adapters.

13.4.9.6.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setcablenvp <WWN | MAC> <NPV>
```

13.4.9.6.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC or iSCSI port.
NVP	A percentage value between 1 and 100 (see your cable documentation to obtain the proper NVP value).

13.4.9.7 tdrtest

This command runs a TDR test on a 10GBASE-T adapter. The TDR test attempts to determine whether any cable faults are compromising the integrity of the link. The test requires that the cable be free of disturbances, down, and quiet. It is best if the cable is not terminated; however, if it is terminated then the link partner must not be active during the test.

For each twisted pair cable (labeled A, B, C, or D):

- If a fault cannot be detected, the test output displays **OK** and an estimated cable length (in meters), if possible. If the length cannot be determined, the estimated length is displayed as invalid.
- If one of two faults (a short or an open connection) is detected, the test output shows the fault type (Shorted or Open) and the distance to the fault (in meters).

NOTE This command only supports 10GBASE-T adapters, which use coax copper cables.

13.4.9.7.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
tdrtest <WWPN | MAC>
```

13.4.9.7.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.

13.4.10 Driver Parameter Commands

The Driver Parameter commands show, set, and save driver parameter values. You can also change the parameters back to factory default values.

13.4.10.1 driverconfig

This command sets all driver parameters to the values in the `.dpv` file on a particular ESXi host. The `.dpv` file's driver type must match the driver type of the host platform FCoE adapter.

13.4.10.1.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
driverconfig <WWPN> <FileName> <Flag>
```

13.4.10.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
FileName	The name of the `.dpv` file (stored in the Broadcom Repository directory)
Flag	G = Makes change global (all adapters on this host) N = Makes change non-global (adapter-specific)

13.4.10.2 getdriverparams

This command shows the driver parameters of the specified port.

13.4.10.2.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
getdriverparams <WWPN | MAC>
```

13.4.10.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi hos.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC or iSCSI port.

13.4.10.3 getdriverparamsglobal

This command shows the global driver parameters of the specified port.

13.4.10.3.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
getdriverparamsglobal <WWPN | MAC>
```

13.4.10.3.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC or iSCSI port.

13.4.10.4 saveconfig

This command saves the specified FCoE adapter's driver parameters to a file on an ESXi host. The resulting file contains a list of driver parameter definitions in ASCII file format with definitions delimited by a comma. Each definition is of the form:

<parameter-name>=<parameter-value>.

The command saves either the values of the global set, or those specific to the adapter in the Broadcom Repository directory.

NOTE Driver parameters that are set temporarily and globally (using the G and T flags) must be read using the getdriverparamsel xvcpcmd command to view the current value of the parameter. The `getdriverparamsglobal brcmvcpcmd` command returns only permanently set driver parameter values. Additionally, if temporary, global values have been set for one or more driver parameters, the `saveconfig brcmvcpcmd` command must be run with the N flag (using the N flag is analogous to the `brcmvcpcmd` command `getdriverparams`) to force the driver parameter values for the specified adapter to be saved. Inaccurate values may be saved if the G flag is used for this command.

13.4.10.4.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
saveconfig <WWN> <FileName> <Flag>
```

13.4.10.4.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
FileName	The name of the local `.dpv` file.
Flag	Valid types are: G = Make change global (all adapters on this host) N = Make change non-global (adapter-specific)

13.4.10.5 setdriverparam

This command sets the driver parameter at the port or global level, either permanently or temporarily for the specified port.

13.4.10.5.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
setdriverparam <WWPN | MAC> <Flag1> <Flag2> <Param> <Value>
```

13.4.10.5.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.
Flag1	L = Local (all adapters on this host) G = Global (all adapters on this host)
Flag2	P = Permanent (persists across re-boot) T = Temporary.
Param	The name of the driver parameter.
Value	The value of the driver parameter.

13.4.10.6 setdriverparamdefaults

This command restores the driver parameter to the default value at the port or global level, either permanently or temporarily for the specified port.

13.4.10.6.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
setdriverparamdefaults <WWPN | MAC> <Flag1> <Flag2>
```

13.4.10.6.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.
Flag1	L = Local (all adapters on this host) G = Global (all adapters on this host)
Flag2	P = Permanent (persists across re-boot) T = Temporary.

13.4.11 Dump Commands

The Dump commands initiate a core dump on a local port, show the dump file for the port on the host, show the diagnostic dump file retention count set on a port, and specify the maximum number of diagnostic dump files for the adapter.

13.4.11.1 deletedumpfiles

This command deletes all diagnostic dump files for a given port.

13.4.11.1.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
deletedumpfiles <WWN | MAC>
```

13.4.11.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC or iSCSI port.

13.4.11.2 dump

This command performs a core dump on a local port. The dump file is placed in the dump directory with the following filename format:

```
<Hostname>_<Adapter serial number>_<datetimestamp>.core
```

If the command is successful, the following message is displayed:

```
Dump Successful.
```

NOTE Since this command is dumping memory, it can take time while generating large files.

13.4.11.2.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
dump <WWPN | MAC> core [Options]
```

13.4.11.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the CNA port.
Options	Additional options are available under the direction of Technical Support.

13.4.11.3 getdumpdirectory

This command shows the dump file directory for the FCoE, NIC, or iSCSI port in the host.

NOTE The dump directory applies to all adapters in the server. There is not a separate dump directory for each adapter.

13.4.11.3.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getdumpdirectory <WWN | MAC>
```

13.4.11.3.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC or iSCSI port.

13.4.11.4 getdumpfilenames

This command shows a list of all the dump filenames for a given port.

13.4.11.4.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getdumpfilenames <WWPN | MAC>
```

13.4.11.4.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC or iSCSI port.

13.4.11.5 getretentioncount

This command shows the diagnostic dump file retention count set on a port.

NOTE The retention count applies to all adapters in the server.

13.4.11.5.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getretentioncount <WWN | MAC>
```

13.4.11.5.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC or iSCSI port.

13.4.11.6 setdumpdirectory

Use the `setdumpdirectory` command to set the dump directory for a given FCoE, NIC, or iSCSI port. To use the `SetDumpDirectory` command, you must have a directory mapped under `/vmfs/volumes/` where the files will be placed.

NOTE The dump directory applies to all adapters in the server. There is no separate dump directory for each adapter.

13.4.11.6.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setdumpdirectory <WWN | MAC> <DumpDirectory>
```

13.4.11.6.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.
DumpDirectory	The directory under `/vmfs/volumes` that you created to store the dump files.

13.4.11.6.3 Example Command

```
brcmvcpcmd v=12.345.678.901 u=username p=password h=12.345.678.123
setdumpdirectory 10:00:00:00:c9:61:f2:64 vcenter-datastore
```

This example shows the dump directory set to `/vmfs/volumes/vcenter-datastore`.

13.4.11.7 setretentioncount

This command specifies the maximum number of diagnostic dump files for the adapter. When the count reaches the limit, the next dump operation deletes the oldest file.

The retention count applies to all adapters in the server.

13.4.11.7.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setretentioncount <WWN | MAC> < RetentionCount >
```

13.4.11.7.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC or iSCSI port.
RetentionCount	The number of dump files to retain.

13.4.12 FCoE Commands

The FCoE commands show the FCF information of an adapter, the FIP parameters of the adapter, and the LUNs attached to the target of the specified port. The FCoE commands also set the FIP parameters of the adapter in FCoE mode and show the targets attached to the specified FCoE port.

13.4.12.1 getfcfinfo

This command shows the FCF information of an FCoE port.

13.4.12.1.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>
h=<esx_host> getfcfinfo <WWPN>
```

13.4.12.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.

13.4.12.2 getfipparams

This command shows the FIP parameters of the FCoE port.

13.4.12.2.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getfipparams <WWPN>
```

13.4.12.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.

13.4.12.3 getlunlist

This command shows the LUNs attached to an FC or FCoE target of the specified FCoE port.

13.4.12.3.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
getlunlist <WWPN> <TargetWWN>
```

13.4.12.3.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
Target WWPN	The WWPN of the FC or FCoE target.

13.4.12.4 setfipparams

This command sets the FIP parameters of the FCoE port.

13.4.12.4.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
setfipparams <WWPN> <Param> <Value>
```

13.4.12.4.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
Param	FIP parameter name (pfabric, pswitch, vlanid, fcmap, or cinvlanid).
Value	FIP parameter value and the valid range for the FIP parameter. pfabric – 8 byte fabric name (format XX:XX:XX:XX:XX:XX:XX:XX) pswitch – 8 byte switch name (format XX:XX:XX:XX:XX:XX:XX:XX) vlanid – 2 byte VLAN_ID [0-4095] OR 'any' for any valid VLAN_ID fcmap – 3-byte FC_map, 0x0EFCxx cinvlanid – 2-byte VLAN_ID [0-4095]

13.4.12.4.3 Example Command

```
brcmvcpcmd  v=12.345.678.901 u=username p=password h=12.345.678.123 setfipparams
10:00:00:00:c9:5b:3a:6d fcmap 0x0efc99
```

13.4.12.5 targetmapping

This command shows the targets attached to the specified FCoE port.

13.4.12.5.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
targetmapping <WWPN>
```

13.4.12.5.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.

13.4.13 Firmware Commands

The firmwareupdate command updates firmware on the adapters found in the specified VMware cluster or the ESXi host.

13.4.13.1 firmwareupdate

See Section 13.4.7.2, firmwareupdate.

13.4.14 iSCSI Commands

iSCSI commands show initiator login options for a port, LUNs attached to the target of the port, statistics for a port, session information, available profile IDs, all sessions for a target, the configured Internet Storage Name, and iSCSI target details of the adapter to which the port belongs. The iSCSI commands also update the configured iSNS server.

13.4.14.1 clearisnsserver

This command clears the configured ISNS server and disables ISNS target discovery for a given port. If there is no ISNS server currently configured, this command does nothing.

13.4.14.1.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
clearisnsserver <MAC>
```

13.4.14.1.2 Parameters

v	The IP address of the vCenter server managing the ESXi host.
u	The user name of the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
MAC	The MAC address of the iSCSI port.

13.4.14.2 getinitiatorproperties

This command shows all the initiator login options for the specified port. These properties are set as the target portal's login properties to be used when discovering the targets on the target portal. The discovered targets inherit these properties.

13.4.14.2.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getinitiatorproperties <MAC>
```

13.4.14.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
MAC	The MAC address of the NIC port.

13.4.14.3 getiscsiluns

This command shows the LUNs attached to the iSCSI target of the specified iSCSI port.

13.4.14.3.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
getiscsiluns <MAC> <Target>
```

13.4.14.3.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
MAC	The MAC address of the NIC port.
Target	The name of the iSCSI target.

13.4.14.4 getiscsiportstats

This command shows all the port statistics for a specified port.

13.4.14.4.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getiscsiportstats <MAC>
```

13.4.14.4.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
MAC	The MAC address of the iSCSI port.

13.4.14.5 getsessioninfo

This command lists all session information for a specified session. You must specify the iSCSI target name <Target>. You must also specify the TSIH <TSIH> of the session, or the session's ISID Qualifier <ISID_Qual> and the target's IP address <Target_IP>. These parameters tell the command to gather the information from the specified target and session. You can find the TSIH and ISID qualifier by running the `ListSessions` command.

13.4.14.5.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
getsessioninfo <MAC> <Target> <TSIH | <ISID_Qual Target_IP>>
```

13.4.14.5.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
MAC	The MAC address of the iSCSI port.
Target	The target's iSCSI name enclosed in quotes (string length:11-255).
TSIH	The TSIH value of the session (value: 1-65535).
ISID_Qual	The ISID qualifier of the session (value: 0-65535).
Target_IP	The Target's IP address.

NOTE You must specify either the TSIH value or the ISID qualifier. If ISID qualifier is specified, you must also specify the Target's ID address.

13.4.14.6 listprofiles

This command shows a list of the available profile IDs and their descriptions for an adapter. This command also shows currently active and configured profile IDs. The ID descriptions shown with the command are useful when you determine the appropriate profile ID to specify for a `ChangePersonality` command. Examples of ID descriptions include the following:

- ISCSI initiator, no DIF
- FCoE initiator + Target, with DIF
- NIC - ETS Disabled
- Mixed Storage, 25% ISCSI initiator only, 75% FCoE, no DIF

13.4.14.6.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd> <h=ESXhost>
ListProfiles <WWPN | MAC address>
```

13.4.14.6.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FCoE function on the physical port.
MAC address	The MAC address of any NIC function on the physical adapter port.

13.4.14.6.3 Example

```
brcmvcpcmd.exe v=10.192.000.000 u=root p=vmware h=10.192.000.000 listprofiles
00-00-c9-12-34-56
```

13.4.14.7 listsessions

This command lists all the sessions on a specified target. The iSCSI target name <Target> instructs the command to gather the information from the listed iSCSI target name.

13.4.14.7.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
listsessions <MAC> <Target>
```

13.4.14.7.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
MAC	The MAC address of the iSCSI port.
Target	The iSCSI target name.

13.4.14.8 showisnsserver

This command shows the currently configured Internet Storage Name Server. This command also indicates whether ISNS discovery is enabled.

13.4.14.8.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
showisnsserver <MAC>
```

13.4.14.8.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
MAC	The MAC address of the iSCSI port.

13.4.14.9 showtarget

This command shows the iSCSI target details of the adapter to which the specified iSCSI port belongs.

13.4.14.9.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
showtarget <MAC>
```

13.4.14.9.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
MAC	The MAC address of the iSCSI port.

13.4.14.10 updateisnsserver

This command updates the configured ISNS server. This command requires the server IP <Server_IP> and port number <Port> of the ISNS server be available to respond to the ISNS requests.

13.4.14.10.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
updateisnsserver <MAC> <Server IP> <Port>
```

13.4.14.10.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
MAC	The MAC address of the iSCSI port.
Server_IP	The IP address of the ISNS server to configure.
Port	The port number of the ISNS server to configure (value: 1024-65535).

13.4.15 Link Reconfigurability Commands

If the adapter supports this functionality, the link reconfigurability commands show and set link configuration settings.

13.4.15.1 getlinkconfig

If supported, this command shows the active and configured link IDs and lists the available link configurations

13.4.15.1.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
getlinkconfig <WWPN | MAC>
```

13.4.15.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of any FCoE function on the adapter.
MAC	The MAC address of the NIC or iSCSI port.

13.4.15.1.3 Example

```
brcmvcpcmd.exe v=10.192.000.000 u=root p=vmware h=10.192.203.195 getlinkconfig
00-90-FA-41-37-A0
```

13.4.15.2 setlinkconfig

If supported, this command changes the configured link configuration ID. It can change it to a specified ID or to the factory default ID. A reboot is required to activate the new link configuration

NOTE If you run the `getadapterportconfig` command after changing the link configuration ID, the configured Reboot Configuration displayed reflects the default port configuration for the newly configured link configuration ID.

13.4.15.2.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
setlinkconfig <WWPN | MAC> <LinkID|default>
```

13.4.15.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of any FCoE function on the adapter.
MAC	The MAC address of the NIC or iSCSI port.

13.4.15.2.3 Example

```
brcmvcpcmd.exe v=10.192.000.000 u=root p=vmware h=10.192.203.195 getlinkconfig
00-90-FA-41-37-A0
```

13.4.16 Personality Commands

Personality commands shows the personalities of the adapter to which the specified port belongs and change the valid personality of the adapter to which the specified port belongs.

13.4.16.1 changepersonality

This command changes the valid personality of the adapter to which the specified port belongs.

NOTES

- The listprofiles command (see Section 13.4.14.6, listprofiles) shows the currently active and configured profile IDs. The descriptions of these IDs are useful for determining the appropriate profile ID to specify for the `changepersonality` command.
- This command is not supported for the OCe14000-series adapters. Use the `setadapterportconfig` command for these adapters.

13.4.16.1.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
changepersonality <WWPN | MAC>  <type> [FCoE\iSCSI\NIC]
```

13.4.16.1.2 Parameters

h	The IP address of the ESXi host.
WWPN	The WWPN of any FCoE function on the adapter.
MAC	The MAC address of the NIC or iSCSI port.
[FCoE\iSCSI\NIC]	The changed personality that is activated upon a system reboot.
pid	A profile ID to use for NIC.

13.4.16.1.3 Example

```
brcmvcpcmd.exe v=10.192.000.000 u=root p=vmware h=10.192.203.195
changepersonality 00-00-c9-12-34-56 nic 21
```

13.4.16.2 showpersonalities

This command shows the personalities of the adapter to which the specified port belongs. It is not supported for the OCe14000-series adapters. Use the `setadapterportconfig` command instead.

NOTE If the active or configured configurations are not single personality, the active or configured personality indicators are not shown.

The personalities shown depend upon the availability of the personality, which is controlled by the profiles that are exposed by an adapter's IPL.

13.4.16.2.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
showpersonalities <WWPN | MAC>
```

13.4.16.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.

13.4.16.2.3 Example

```
brcmvcpcmd.exe v=10.192.000.000 u=root p=vmware h=10.192.000.000
showpersonalities 00-00-c9-12-34-56

Adapter Personalities:
NIC
iSCSI
FCoE (active) (configured)
```

13.4.17 SR-IOV Commands

SR-IOV commands show a list of all the discovered virtual functions and also enables or disables SR-IOV support for an adapter's physical port.

13.4.17.1 listvfuncs

This command shows a list of all the virtual functions discovered.

13.4.17.1.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter_server> u=<vc_username> p=<vc_password> h=<esx_host>
listvfuncs <MAC>
```

13.4.17.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
MAC	MAC address of any NIC function on the physical adapter port

13.4.17.2 sriovenable

This command enables or disables SR-IOV support for an adapter's physical port.

13.4.17.2.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter_server> u=<vc_username> p=<vc_password> h=<esx_host>
sriovenable <MAC>  <0 | 1>
```

13.4.17.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
MAC	The MAC address of any NIC function on the physical adapter port.
SR-IOV Enable state	0 = disable SR-IOV 1 = enable SR-IOV

NOTE SR-IOV is not supported with UMC channel management.

13.4.18 VEPA Commands

13.4.18.1 VEPAenable

This command enables or disables VEPA. VEPA is only supported on OCe14000-series NIC ports. For all other adapter types and port types, this feature is undefined, and this command returns the following error message:

```
Hardware Does Not Support
```

SR-IOV must be enabled on the NIC port for VEPA to be enabled.

13.4.18.1.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_pwd> h=<esx_host>
VEPAenable <MAC Address> <0|1>
```

13.4.18.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
MAC	The MAC address of any NIC on the relevant physical port.
(0/1)	Specifies whether to enable or disable VEPA: 0 is disabled and 1 is enabled.

13.4.19 Virtual Machines Commands

13.4.19.1 listvms

This command lists all virtual machines and their information for all manageable ports.

If you specify the host with the h=<esx_host> option or just give the physical WWPN, only the virtual machines for that host are shown. If you specify the physical port and the virtual port, only the virtual machine for the specified virtual port returns.

The virtual machine name is not displayed.

13.4.19.1.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
listvms <WWPN>
```

13.4.19.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FCoE port.

13.4.19.2 listvports

This command lists virtual ports on the specified physical port. Leaving the physical WWPN parameter blank lists all virtual ports on all manageable hosts that support virtual ports.

13.4.19.2.1 Syntax

```
brcmvcpcmd.exe v=<vcenter_server> u=<vc_username> p=<vc_password>  h=<esx_host>
listvports <WWNP>
```

13.4.19.2.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of an FCoE port.

13.4.20 WWWN Management Commands

13.4.20.1 changewwn

This command changes the WWN of the specified port.

13.4.20.1.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
changewwn <WWPN> <New WWPN> <New WWNN> <ReadType>
```

13.4.20.1.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
ReadType	0 = Volatile, 1 = Nonvolatile.

13.4.20.2 getwwncap

See Section 13.4.4.4, getwwncap.

13.4.20.3 readwwn

This command shows the WWN details of the specified port and category.

13.4.20.3.1 Syntax

```
brcmvcpcmd.exe v=<vcenter server> u=<vc_username> p=<vc_pwd> <h=ESXhostIP>
showpersonalities <WWPN | MAC> [ReadType]
```

13.4.20.3.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
MAC	The MAC address of the NIC port.
ReadType	Valid types are: 0 = Volatile 1 = Nonvolatile 2 = Factory default 3 = Current 4 = Configured

13.4.20.4 restorewwn

This command restores the WWN value of the specified port.

13.4.20.4.1 Syntax

```
brcmvcpcmd.exe  v=<vcenter server> u=<vc_username> p=<vc_pwd>  <h=ESXhostIP>
restorewwn <WWPN> <RestoreType>
```

13.4.20.4.2 Parameters

v	The vCenter server IP address.
u	The user name for the vCenter server.
p	The user password for the vCenter server.
h	The IP address of the ESXi host.
WWPN	The WWPN of the FCoE port.
RestoreType	Valid types are: 0 = Restore Default WWNs 1 = Restore NVRAM WWNs

13.5 Viewing Audit Logs Using the CLI Command

You can use the brcmvcpaudit.exe script to log all the historical active management performed through the vSphere console on the adapters. To see the usage information, run the script with no parameters specified. All supported events are displayed.

NOTE All active management actions performed are saved to a log file specific to the action. The maximum size of a log file is 2 MB. If the size of the log file exceeds this limit, old log entries are deleted for the particular event.

13.5.1 Syntax

```
brcmvcpaudit.exe <event name>
```

NOTE If an event name is not specified, all events are displayed.

13.5.2 Parameter

event name	The command name describing the appropriate active management action performed.

13.5.3 List of Supported Event Names

installadapterlicense	Licenses installed on the CNA.
changepersonality	Personality changes on the CNA.
loopbacktest	Diagnostic tests on the CNA.
setdriverparam	Driver parameters changes at both port and global levels.
setdcbparam	DCB parameter changes on the CNA.
setbeacon	Beacon setting changes on the CNA.
setfipparam	FIP parameter changes on the CNA.
download	Firmware download activities.
reset	Port reset activities.
changewwn	WWN change activities.

13.5.4 Example

```
Audit log  for       : changewwn
--------------------

User Name            : Administrator
Date and Time        : 2011-06-16T19:25:12Z
Operation            : ChangeWWNJobInfo
Host Name            : 10.192.203.179
Adapter Id           : BT11161224
Port Id              : 10:00:00:00:C9:BB:1E:77
Message              : Successfully changed WWN

New State
                     New WWPN = 10:00:00:00:C9:BB:1E:78
                     New WWNN = 20:00:00:00:C9:BB:1E:78
                     Volatile = false

Old State
                     Old WWPN = 10:00:00:00:C9:BB:1E:77
                     Old WWNN = 20:00:00:00:C9:BB:1E:77
```

Appendix A: License Notices

A.1 VI Java SDK

Copyright (c) 2012 Steve Jin. All Rights Reserved.
Copyright (c) 2008 VMware, Inc. All Rights Reserved.
Copyright (c) 2009 Altor Networks. All Rights Reserved.
Copyright (c) 2009 NetApp. All Rights Reserved.
Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:
* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
* Neither the name of VMware, Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.
THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED.
IN NO EVENT SHALL VMWARE, INC. OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

A.2 Java Notice

Copyright(c) 2002 Sun Microsystems, Inc. All Rights Reserved.

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

- Redistribution of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistribution in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of Sun Microsystems, Inc. or 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 "AS IS," without a warranty of any kind. ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE HEREBY EXCLUDED. SUN MICROSYSTEMS, INC. ("SUN") AND ITS LICENSORS SHALL NOT BE LIABLE FOR ANY DAMAGES SUFFERED BY LICENSEE AS A RESULT OF USING, MODIFYING OR

DISTRIBUTING THIS SOFTWARE OR ITS DERIVATIVES. IN NO EVENT WILL SUN OR ITS LICENSORS BE LIABLE FOR ANY LOST REVENUE, PROFIT OR DATA, OR FOR DIRECT, INDIRECT, SPECIAL, CONSEQUENTIAL, INCIDENTAL OR PUNITIVE DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF THE USE OF OR INABILITY TO USE THIS SOFTWARE, EVEN IF SUN HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

You acknowledge that this software is not designed, licensed or intended for use in the design, construction, operation or maintenance of any nuclear facility.


BROADCOM

Emulex Boot for NIC, iSCSI, and FCoE Protocols

User Guide

Version 11.2
December 30, 2016

NIF-Boot-OCA-UG112

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

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

Copyright © 2003–2017 by Broadcom. All Rights Reserved.

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

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.

Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Chapter 1: Introduction

This guide describes installing, enabling, and configuring boot code for Emulex® network interface card (NIC), Internet Small Computer System Interface (iSCSI), and Fibre Channel over Ethernet (FCoE) adapters. This guide also describes the Emulex boot from SAN implementation and its operation with distinct hardware and operating system requirements.

Boot from SAN is the process of booting a server directly from a disk operating system image located on a storage area network (SAN) by way of Emulex adapters using Emulex boot code. When booting from SAN, the storage device is typically identified by its World Wide Port Name (WWPN) and a logical unit number (LUN). By extending the server system boot basic input/output system (BIOS), boot from SAN functionality is provided by the boot BIOS contained on an Emulex adapter in the server. If properly configured, the adapter then permanently directs the server to boot from a logical unit (disk) on the SAN as if it were a local disk.

Broadcom provides the following types of boot code:

- Preboot Execution Environment (PXE) boot for NIC adapters in x86 and x64 systems
- x86 BootBIOS for FCoE adapters in x86 and x64 systems
- iSCSI boot for iSCSI adapters in x86 and x64 systems
- UEFIBoot for NIC, iSCSI, and FCoE adapters in x64 systems

  UEFIBoot provides system boot capability through the use of the UEFI (Unified Extensible Firmware Interface) shell. It also functions on UEFI 2.x-based platforms through the HII (Human Interface Infrastructure).

NOTE Emulex drivers support multipath boot configurations. See your storage vendor's documentation for information on configuring multipath booting.

The Emulex boot code and the following utilities provide a variety of capabilities:

- Boot from SAN across different networking protocols and operating systems
- UEFI configuration using the Emulex NIC, FCoE, and iSCSI configuration utilities
- PXESelect utility
  - Configuring the port and adapter
  - Configuring Universal Multichannel (UMC) support and personality options
- FCoE x86 BootBIOS utility
  - Scanning for target devices
  - Configuring boot devices and advanced adapter parameters
- iSCSISelect utility
  - Setting up a basic configuration
  - Configuring and managing iSCSI initiators and targets

The boot code is distributed in the same image used to flash the firmware. Several methods are available for flashing the firmware and boot code image. See the documentation accompanying each utility for more information on the flash procedure.

NOTE Screenshots in this guide are for illustrative purposes only. Your system information can vary.

1.1 Abbreviation

ACL	Access Control List
API	Application Programming Interface
ARI	Alternative Routing-ID Interpretation
ARP	Address Resolution Protocol
BBS	BIOS Boot Specification
BIOS	Basic Input/Output System
CHAP	Challenge Handshake Authentication Protocol
CLI	Command Line Interface
CNA	Converged Network Adapter
DCBX	Data Center Bridging Capabilities Exchange
DHCP	Dynamic Host Control Protocol
DID	device ID
DMA	Direct Memory Access
DNS	Domain Name System or Domain Name Server
EDD	Enhanced Disk Device
EFI	Extensible Firmware Interface
FC	Fibre Channel
FCF	Fibre Channel over Ethernet Forwarder
FCoE	Fibre Channel over Ethernet
FoD	Feature on Demand
FTP	File Transfer Protocol
FUI	FoD Unique Identifier
GPT	GUID partition table
GUI	Graphical User Interface
GUID	Globally Unique Identifier
HBA	Host Bus Adapter
HII	Human Interface Infrastructure
HTTP	Hypertext Transfer Protocol
iBFT	iSCSI Boot Firmware Table
ICMP	Internet Control Message Protocol
I/O	input/output
IP	Internet Protocol
IQN	iSCSI Qualified Name
iSCSI	Internet Small Computer System Interface
ISID	Initiator Session Identifier
iSNS	Internet Storage Name Service
JBOD	just a bunch of disks
LAN	Local Area Network
LED	Light-emitting diode
LOM	LAN on motherboard

LPVID	logical port VLAN ID
LUN	Logical Unit Number
MAC	Media Access Control
MPIO	Multipath I/O
MTU	Maximum Transmission Unit
NBP	Network Boot Program
NFS	Network File System
NIC	Network Interface Card
NPar	NIC partitioning
NTFS	New Technology File System
NVRAM	Non-volatile Random-access Memory
OS	Operating System
PCI	Peripheral Component Interconnect (interface)
PCIe	Peripheral Component Interconnect Express
PF	physical function
POST	Power-on Self-test
PXE	Preboot Execution Environment
QoS	Quality of Service
RAID	Redundant Array of Independent Disks
RDMA	Remote Direct Memory Access
RHEL	Red Hat Enterprise Linux
RoCE	RDMA over Converged Ethernet
ROM	read-only memory
RSS	Receive-side Scaling
SAN	Storage Area Network
SCSI	Small Computer System Interface
SFS	SAN Foundation Software
SLES	SUSE Linux Enterprise Server
SMB	Server Message Block
SR-IOV	Single Root I/O Virtualization
SVID	service VLAN ID
TCP	Transmission Control Protocol
TCP/IP	TCP over Internet Protocol
TFTP	Trivial File Transfer Protocol
TOE	TCP Offload Engine
UEFI	Unified Extensible Firmware Interface
UFP	Unified Fabric Port
UMC	Universal Multichannel
UNDI	Universal Network Device Interface
USB	Universal Serial Bus
VF	virtual function
VLAN	virtual local area network
VLAN ID	VLAN identifier

VMQ	virtual machine queue
WDS	Windows Deployment Services
WoL	Wake on LAN
WWN	World Wide Name
WWNN	World Wide Node Name
WWPN	World Wide Port Name

Chapter 2: Configuring the PXE Boot for NIC on OCe11100-Series Adapters

NOTE This section applies to Emulex OCe11100-series adapters. For information on configuring Emulex OCe14000-series adapters, see Chapter 3, Configuring PXE Boot for NIC on OCe14000-Series Adapters.

This section describes using and configuring the PXE to boot computers using a network interface independent of available data storage devices (such as hard disks) or installed operating systems.

The PXE protocol is a combination of DHCP and TFTP with subtle modifications to both. DHCP locates the appropriate boot server or servers, and TFTP downloads the initial bootstrap program and additional files.

Network booting enables you to perform the following tasks:

- Boot diskless systems such as thin clients and dedicated systems.
- Deploy software and operating systems for your systems.
- Automate system maintenance, such as backups.
- Automate system checking, such as virus scanning.
- Ensure system security.

2.1 Pre-OS

The PXE can be used in a *pre-OS* environment. Pre-OS is the process of loading a small operating environment to perform a client management task before loading the final operating system from the local hard drive. For example, with a pre-OS you can scan the hard drive for viruses. This guarantees that the client is not infected before it starts. Another example is the WDS using this feature to install operating systems on local disks or boot from SAN disks.

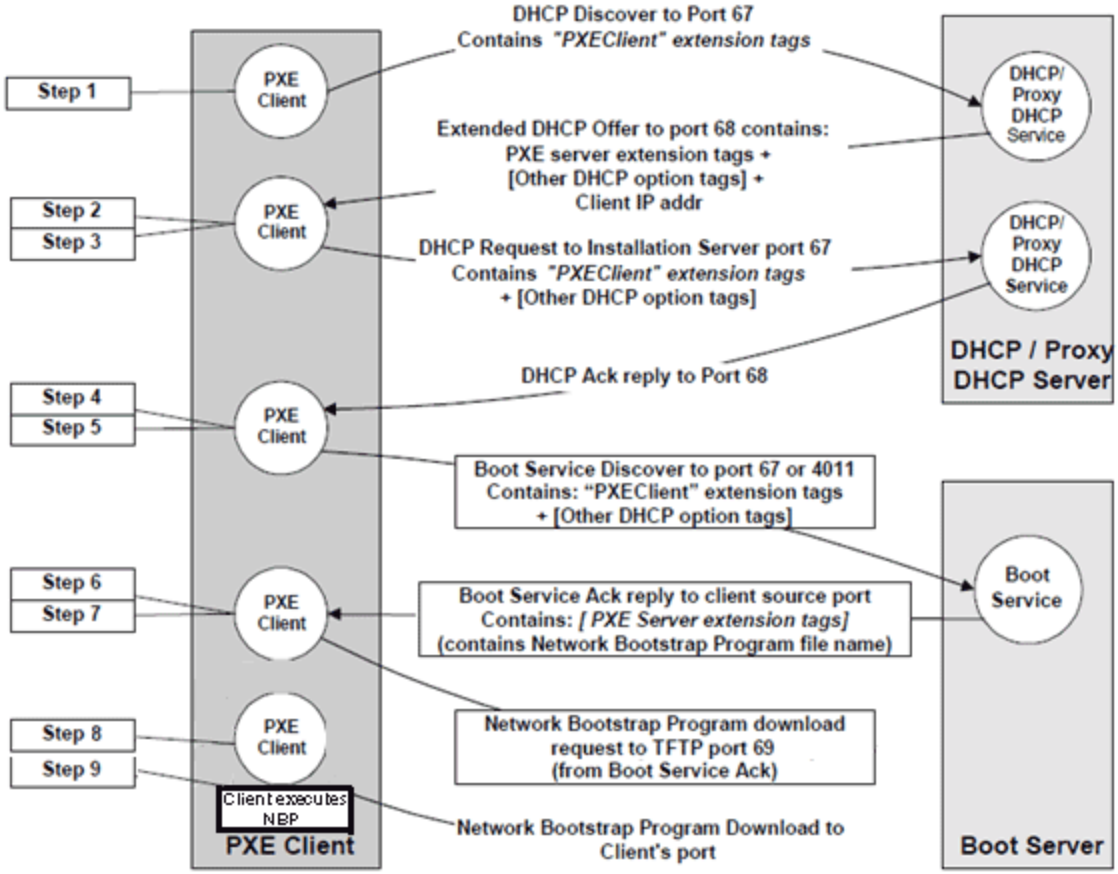
2.2 PXE Boot Process

After the PXE boot is enabled in the system UEFI and BIOS, the PXE client can boot up and start the PXE boot ROM code. This is the boot code physically located on the NIC adapter.

Figure 1 shows the boot process.

NOTE To enable or disable the PXE boot, it must be enabled or disabled in the system UEFI and BIOS; see the documentation that accompanied the server for more information.

Figure 1 PXE Boot Process



1. The PXE boot ROM sends a DHCP request extended with PXE specific options (step 1 in the figure).
2. The DHCP responses contain the DHCP options (DHCPOFFERs) that include the NBP filename and boot server lists (steps 2 through 5).
3. The PXE client attempts to download the specified NBP over TFTP from one of the specified boot servers (steps 6 and 7).
4. The PXE client executes the downloaded NBP (steps 8 and 9).

NOTE If any of these steps fail, the boot process typically continues using the next available device in the boot sequence, depending on the system configuration and boot order.

2.2.1 Remotely Installing with the PXE for Windows Server 2008, 2008R2, 2012, 2012 R2, and Windows Server 2016

For remote installation with PXE, a network driver for the Emulex adapter must be part of the client's installation image on the server. The current versions of Windows Server 2008 and Windows Server 2008 R2 do not include network drivers for the Emulex adapter; however, Windows Server 2012, 2012 R2, and 2016 include the network driver for the Emulex adapter.

To add the image and installation using the driver with the Remote Installation Service Setup:

Select **Start Menu > Programs > Administrative Tools > Remote Installation Services Setup**

In addition to the network driver for the Emulex adapter, you must configure the following services to use the PXE for remote installations:

- DHCP server
- Remote Installation Services
- Windows Deployment Services

Microsoft provides extensive documentation on deploying its operating systems for remote installations, and different setups may be required depending on your individual implementation. Microsoft provides step-by-step guides for its Windows Deployment Services for configuring the server, adding images, and installing an operating system. It also includes instructions for more advanced tasks like creating multicast transmissions, creating custom images, and performing an unattended installation.

For detailed information on deploying and managing remote installations on Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, and Windows Server 2016 see the Microsoft website and visit Microsoft TechNet. Search on the bulleted terms above to access Microsoft's wide-ranging documentation on these subjects.

2.3 Remotely Installing with the PXE for Linux and Citrix Servers

Linux allows for PXE installation over a network using the NFS, TFTP, or HTTP protocols. If the system to be installed contains an Emulex NIC adapter with PXE Boot support, it can be configured to boot from files on another networked system rather than local media.

The Linux distributions provide extensive documentation on deploying and managing remote installations of the Linux operating system using PXE. See the distribution's documentation for instructions on deploying a PXE installation over the network.

For remote installation with PXE, a network driver for the Emulex adapter must be part of the client's installation image on the server. The current Linux distribution may or may not include network drivers for Emulex adapters. If it does, the driver may need to be added to the operating system installation image, or added during installation. See the distribution's documentation for instructions on adding drivers during installation.

2.4 Using the PXESelect Utility

This section describes how to configure the PXE boot options using the PXESelect utility.

2.4.1 Navigating the PXESelect Utility

Use the following methods to navigate the PXESelect utility:

- Press the up and down arrows on the keyboard to move through and select menu options or configuration fields. If multiple adapters are listed, use the up and down arrows to scroll to the additional adapters.
- Press the Tab key to move to the next field or to select a row in a configuration dialog. Press shiftbar-Tab key to move to the previous field.
- Press Enter to accept a selection, select a menu option, or change a configuration default.
- Press Esc to return to the previous menu or page, cancel a selection or dialog box, or exit the utility.

2.4.2 Running the PXESelect Utility

To run the PXESelect utility, start or restart the computer. When prompted, press Ctrl-P. If you are running the PXESelect utility with multiple adapters, all of the adapters are displayed when you start the utility. For example:

```
Press <Ctrl><P> for PXESelect(TM)Utility

OCe14102B-UM NIC#0 Port#1 Base 0x00000380001A0000 @ Bus:1B Dev:00 Fun:00
OCe14102B-UM NIC#0 Port#2 Base 0x00000380001E0000 @ Bus:1B Dev:00 Fun:00
- Initializing ...Done.
```

The PXE Configuration menu appears after the boot BIOS initializes to begin the PXE configuration.

NOTE A UEFI-capable system typically does not display the prompt for running the PXESelect utility unless it is configured for legacy booting. See your system configuration manual for information on performing a legacy boot. All configurations that can be performed in the PXESelect utility can instead be performed in the UEFI configuration utility. For more information on configuring UEFI, see Chapter 11, Configuring UEFI for Ethernet.

2.4.3 Setting Up a PXE Bootable Network

After the PXE boot BIOS is initialized, you can use the PXESelect utility to set up a PXE bootable network by configuring the adapters.

To configure adapters for a PXE boot:

1. At the following Controller Selection Menu), use the Tab key to select the adapter you want to configure and press Enter.

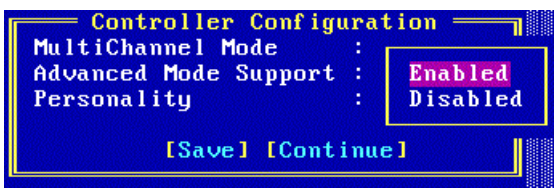
Figure 2 Controller Selection Menu



NOTE The Controller Selection Menu appears only if two or more adapters are connected.

The following Controller Configuration dialog appears.

Figure 3 Controller Configuration Dialog



Depending on the adapter, the following options are available:

— **MultiChannel Mode** – Allows you to access up to eight virtual network interfaces. For more information on Multichannel mode, see Section 2.4.5, Configuring Universal Multichannel.

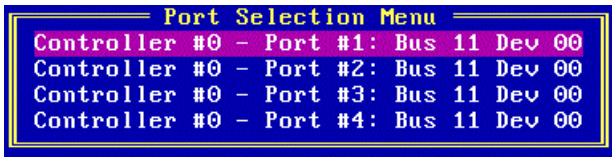
NOTE If you are using a Lenovo System x adapter, see Section 2.4.6, Configuring Multichannel for Lenovo System x Adapters, for more information on Multichannel mode.

- **Advanced Mode Support** – Enables you to run Advanced mode–aware drivers. For more information on configuring Advanced mode, see Section 2.4.7, Advanced Mode Support.
- **Personality** – Allows you to select the protocols that can be configured on an adapter. For more information on adapter personality, see Section 2.4.8, Personality Option.

2. After you set the options in the Controller Configuration dialog, select Save and press Enter.
3. To proceed, select Continue and press Enter.

   The Port Selection Menu (Figure 4) appears.
4. Select the port you want to configure and press Enter.

Figure 4 Port Selection Menu



NOTE The Port Selection Menu appears if two or more ports are connected.

Depending on the installed adapter and the Multichannel mode setting, proceed to the following sections:

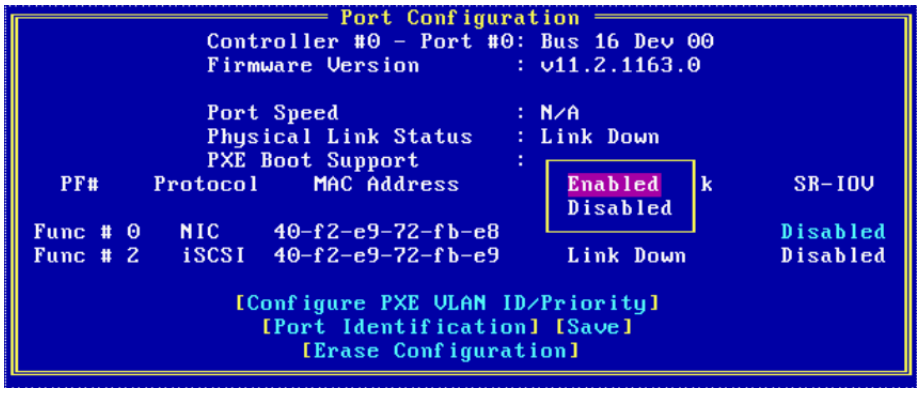
- If MultiChannel Mode is disabled, or it is not supported, see Section 2.4.4, Configuring Ports If UMC Is Disabled or Not Supported, for additional configuration information.
- If MultiChannel Mode is enabled, see Section 2.4.5, Configuring Universal Multichannel, for additional configuration information.
- If MultiChannel Mode is enabled and you are using Lenovo System x adapters, see Section 2.4.6, Configuring Multichannel for Lenovo System x Adapters, for additional configuration information.

2.4.4 Configuring Ports If UMC Is Disabled or Not Supported

NOTE If UMC is enabled, for information on configuring UMC, see Section 2.4.5, Configuring Universal Multichannel. If you are using a Lenovo System x adapter, see Section 2.4.6, Configuring Multichannel for Lenovo System x Adapters.

If UMC is disabled, or if it is not supported, the Port Configuration dialog (Figure 5) appears after you select a port in the Port Selection Menu (Figure 4).

Figure 5 Port Configuration Dialog



NOTE Logical Link Status appears as N/A for storage functions.

The Port Configuration dialog enables you to perform the following tasks:

- Configure the PXE boot support
- Configure SR-IOV
- Configure the PXE VLAN ID and Priority
- Identify ports
- Erase the port and adapter configuration

2.4.4.1 Configuring the PXE Boot Support

To configure the PXE boot support:

1. On the Port Configuration dialog, use the Tab key to select the PXE Boot Support. A pop up appears.
2. Select **Enabled** or **Disabled** and press Enter.
3. Select Save and press Enter.

NOTE During system startup, the PXE contacts the DHCP server for an IP address to boot from the network.

2.4.4.2 Configuring SR-IOV

If the system BIOS supports SR-IOV, you can enable it on the adapter. SR-IOV support can be enabled only if UMC is disabled or not supported. For more information on SR-IOV configuration, see the appropriate Emulex drivers for OneConnect® adapters user guide.

To configure SR-IOV:

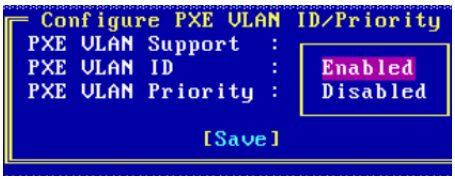
1. On the Port Configuration dialog, use the Tab key to select SR-IOV. A pop up appears.
2. Select **Enabled** or **Disabled** and press Enter.
3. Select Save and press Enter.

2.4.4.3 Configuring the PXE VLAN ID and Priority

To configure a PXE VLAN ID and set the priority level:

1. On the Port Configuration dialog, select Configure PXE VLAN ID/Priority and press Enter. The Configure PXE VLAN ID/Priority menu (Figure 6) appears.

Figure 6 Configure PXE VLAN ID/Priority Menu



2. Use the Tab key to select PXE VLAN Support. A pop up appears.
   a. For PXE VLAN Support, select **Enabled** and press Enter.
   b. For the PXE VLAN ID, select a number from 0 to 4094 and press Enter.
   c. For the PXE VLAN Priority level, select a number from 0 to 7 and press Enter. This unique value assigns a priority to outbound packets containing a specified VLAN ID. Valid values range from 0 to 7, with 0 the highest priority level.
3. Select Save and press Enter.

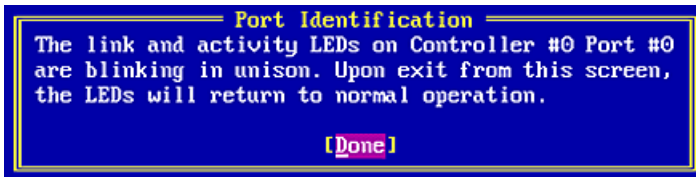
After you exit the PXESelect utility, the system must reboot for the configuration to take effect.

2.4.4.4 Physically Identifying the Port

To physically determine which port you are configuring by causing the link and activity LEDs of that port to blink:

1. On the Port Configuration dialog, select Port Identification and press Enter. The following Port Identification pop up appears.

Figure 7 Port Identification Pop Up



2. The LEDs on the adapter begin blinking. The selected port's LED status indicators blink on the adapter until you select Done on this pop up and press Enter.

NOTE Not all adapters have externally–visible LEDs. If you are using an add-in card in a blade server environment, port identification is not supported.

2.4.4.5 Erasing Ports and Adapter Configurations

NOTE On OCe11100-series adapters, if you select the Erasing Ports and Adapter Configurations setting, all previous configuration settings are returned to their factory default settings except for the current personality selection. Performing this action ensures a clean environment for new configuration settings to take effect.

To erase the ports and adapter configuration:

1. On the Port Configuration dialog, select Erase Configuration and press Enter. The system displays a warning asking if you want to erase the current configuration for all ports of the adapter.
2. Press Y to delete the configuration. You will receive another warning asking you to confirm the permanent removal of the configuration.

3. Press Y to delete the configuration.

To exit the PXESelect utility after erasing the ports and adapter configuration:

1. Follow the instructions on the bottom of the individual menu until you are prompted to exit.
2. Press Y to exit. Depending on the settings that were changed, a reboot may be necessary.

NOTE For older systems, depending on the memory allocation method supported, the PXESelect utility automatically reboots even when there are no changes made to the system.

2.4.5 Configuring Universal Multichannel

UMC provides the ability to configure multiple physical functions or I/O channels for each physical adapter port. For more information on UMC support, refer to the *Emulex Universal Multichannel Reference Manual*.

NOTE If UMC is enabled and the adapter supports Virtual Fabric mode or Unified Fabric Protocol mode, see Section 2.4.6, Configuring Multichannel for Lenovo System x Adapters, for more information. If UMC is disabled or not supported on the adapter, see Section 2.4.4, Configuring Ports If UMC Is Disabled or Not Supported.

If UMC is enabled, you must configure the minimum and maximum bandwidth settings for iSCSI and FCoE storage functions in the NIC BIOS before they can be configured further from their respective utilities (iSCSI or FCoE BIOS). Otherwise, the Logical Link for that function will be shown as down, and you will not be able to log in to targets or find LUNs behind those targets.

NOTE The PXE VLAN is not supported in UMC environments.

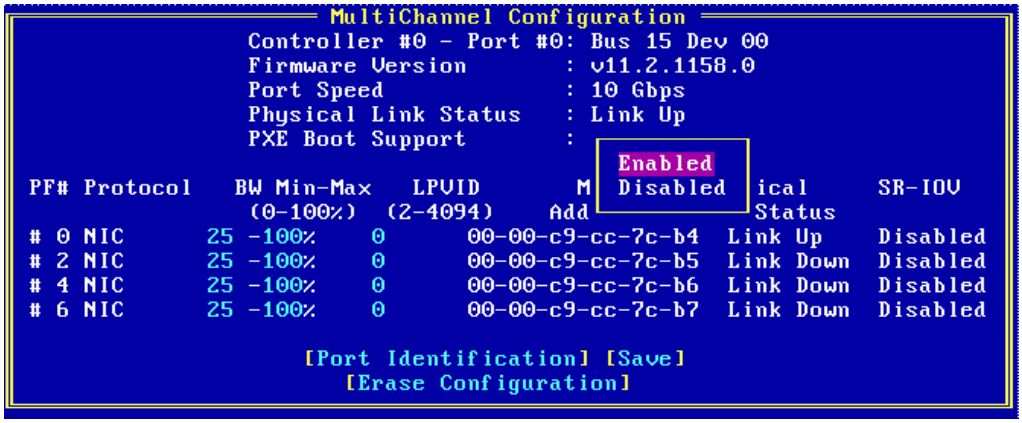
UMC functionality is supported on Emulex OneConnect adapters running in 10Gb Ethernet mode only. UMC is not supported in 1GbE mode.

If the Multichannel mode has changed since the last boot, a reboot will occur when you exit the PXESelect utility.

To view the Multichannel Configuration dialog:

1. From the Port Selection Menu (Figure 4), select a port to configure and press Enter. The following Multichannel Configuration dialog (Figure 8) appears.

Figure 8 Multichannel Configuration Dialog



NOTE The Logical Link Status appears as N/A for storage functions.

SR-IOV support can be enabled only if UMC is disabled.

If the system does not support UMC, it appears as N/A on the Switch Independent Mode Configuration dialog).

2. From the Multichannel Configuration dialog (Figure 8), you can perform the following tasks:
   - Configure PXE boot support – See Section 2.4.4.1, Configuring the PXE Boot Support, for instructions.
   - Configure PXE VLAN ID/Priority – See Section 2.4.4.3, Configuring the PXE VLAN ID and Priority, for instructions.
   - Identify ports – See Section 2.4.4.4, Physically Identifying the Port, for instructions.
   - Erase the port and adapter configuration – See Section 2.4.4.5, Erasing Ports and Adapter Configurations, for instructions.
   - Configure the minimum and maximum bandwidth for each channel – See Section 2.4.5.1, Configuring Minimum and Maximum Bandwidth, for instructions.
   - Configure the LPVID for each channel – See Section 2.4.5.2, Configuring LPVID, for instructions.

NOTE Your adapter or system may not support all UMC options.

2.4.5.1 Configuring Minimum and Maximum Bandwidth

To configure bandwidth:

1. On the MultiChannel Configuration dialog, use the Tab key to select Minimum Bandwidth or Maximum Bandwidth.
   - The Minimum Bandwidth value is the least amount of bandwidth that the function can provide. It is represented as a percentage. The Minimum Bandwidth value must be less than or equal to the Maximum Bandwidth value. The total of the Minimum Bandwidth values for all enabled functions on that port must be equal to 100.

     NOTE A Minimum Bandwidth value of 0 is a valid value. If all of the partitions' Minimum Bandwidth values are zero, the bandwidth is distributed equally among the current active partitions. If a specific partition's Minimum Bandwidth and Maximum Bandwidth values are 0, that partition's logical link will be brought down.
   - The Maximum Bandwidth value is the greatest amount of bandwidth that the function can provide. It is represented as a percentage.
2. Enter the value for the specified option and press Enter.

3. Select Save and press Enter.

2.4.5.2 Configuring LPVID

The LPVID is used to enforce a VLAN ID on all traffic originating from an IP address, channel, or PCI function. If the operating system for that PCI function has set up a VLAN ID, then the operating system–configured VLAN ID takes precedence over the LPVID for transmit packets, while the operating system–configured VLAN ID and LPVID-tagged packets will both be received. If the operating system has not set up any VLAN IDs, then the LPVID is used for tagging.

NOTE LPVID and user-configured VLAN IDs from the operating system must be different.

LPVIDs also must be configured on the switch port.

Each LPVID must be unique and is relevant for NIC traffic only. The LPVID is not supported for storage functions. For iSCSI storage functions, you must configure a VLAN ID through iSCSISelect or through the host. For more information, see Section 9.3.2, Configuring VLAN ID and VLAN Priority.

During the PXE boot, when the UNDI Driver (BIOS) is functional, the PXE VLAN is used. However, after the NIC driver is operational, the LPVID is used.

NOTE The PXE VLAN is not supported in UMC environments.

For example:

```
PXE Install OS
PXE Server configured with VLAN 5
PXE VLAN=5
LPVID for function 0=5
```

To configure LPVID:

1. On the MultiChannel Configuration dialog, use the Tab key to select LPVID. A dialog box appears.
2. Enter a value and press Enter. The LPVID range is 2 to 4094. A value of 0 disables the LPVID.
3. Select Save and press Enter.

2.4.6 Configuring Multichannel for Lenovo System x Adapters

NOTE This section pertains to Emulex OneConnect adapters that support Multichannel for Lenovo System x Virtual Fabric or UFP only.

If you want to enable Multichannel on Lenovo System x adapters, you must select a Multichannel mode.

To configure Multichannel:

1. From the Controller Configuration dialog (Figure 9), use the Tab key to select Multichannel Mode. A pop up appears.

Figure 9 Controller Configuration Dialog - Lenovo System x Adapters



2. Select one of the following settings and press Enter:
   — Switch Independent Mode

— IBM Virtual Fabric Mode
— IBM Unified Fabric Protocol Mode
— Disabled

NOTE For more information on the available Multichannel options, see Section 2.4.6.1, Multichannel Modes.

3. Select Save and press Enter.
4. To proceed, select Continue and press Enter.

With the Multichannel option selected, you can now configure Multichannel support. See Section 2.4.6.2, Multichannel Configuration for Lenovo System x, for more information.

2.4.6.1 Multichannel Modes

NOTE The following modes are available on Lenovo System x adapters and systems that support IBM Virtual Fabric mode and IBM Unified Fabric Protocol mode only.

Multichannel for Lenovo System x is supported on Emulex OneConnect adapters running in 10GbE mode only. Multichannel for Lenovo System x is not supported in 1GbE mode.

Emulex OneConnect OCe11100-series adapters support a maximum of four PFs per port in all Multichannel modes.

A Lenovo System x Virtual Fabric–enabled switch provides the ability to configure an LPVID for a virtual channel or I/O channel on an adapter port. If Multichannel for Lenovo System x is supported on the system, the PXESelect utility enables you to select one of the following modes:

- **IBM Virtual Fabric Mode** – Select this mode if an Emulex OneConnect adapter is attached to an Lenovo System x Virtual Fabric–enabled switch.
- **IBM Unified Fabric Protocol Mode** – Select this mode if an Emulex OneConnect adapter is attached to a Lenovo System x UFP–enabled switch.

NOTE Some Lenovo System x switches support both IBM Unified Fabric Protocol mode and IBM Virtual Fabric mode.

- **Switch Independent Mode** – Select this mode if you are using a switch other than a Lenovo System x Virtual Fabric or UFP-enabled switch. For more information on configuring Multichannel, see Section 2.4.5, Configuring Universal Multichannel.

2.4.6.2 Multichannel Configuration for Lenovo System x

Multichannel for Lenovo System x provides the ability to configure multiple PCI functions or I/O channels for each physical adapter port.

NOTE Setting up Multichannel for IBM Virtual Fabric mode or IBM Unified Fabric Protocol mode depends on cooperation with adjacent switches.

For information on configuring Multichannel for Lenovo System x if Switch Independent mode is enabled, see Section 2.4.5, Configuring Universal Multichannel.

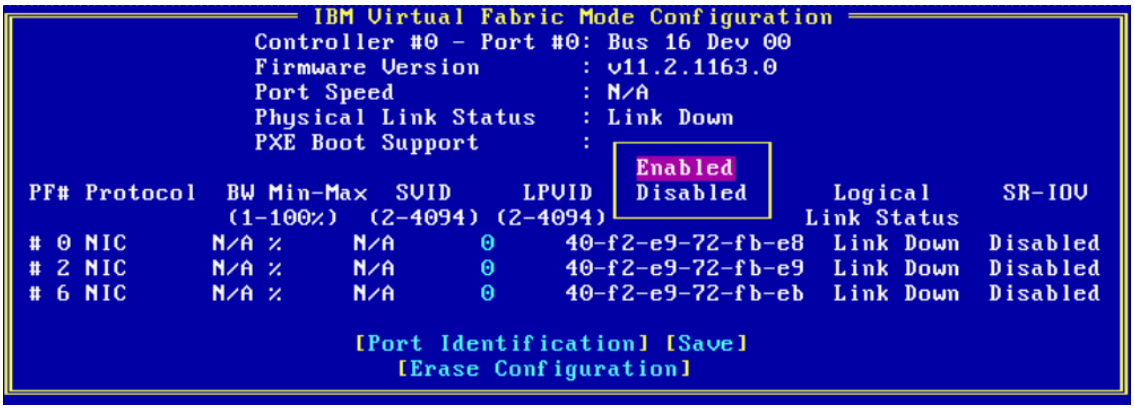
To view the configuration dialog:

1. From the Port Selection Menu (Figure 4), select the port you want to configure and press Enter. The configuration dialog appears.

If you configure Multichannel for Lenovo System x on adapters that support IBM Virtual Fabric mode and IBM Unified Fabric Protocol mode, the configuration dialog will depend on the Multichannel mode that has been selected.

NOTE For illustration purposes, the dialogs in this section are for a system with IBM Virtual Fabric mode enabled.

Figure 10 Virtual Fabric Mode Configuration Dialog



NOTE Logical Link Status appears as N/A for storage functions.

SR-IOV support can be enabled only if Multichannel mode is disabled.

If the system does not support Multichannel for Lenovo System x, it appears as N/A on the configuration dialog.

The SVID, or Outer VLAN ID, appears in IBM Virtual Fabric mode and IBM Unified Fabric Protocol mode. This value is provided by the Lenovo System x switch and will appear if the protocol modes are configured on the switch and the adapter. If the protocol modes are not configured, a value of N/A appears.

2. From the configuration dialog, you can perform the following tasks:
   - Configure PXE boot support – See Section 2.4.4.1, Configuring the PXE Boot Support, for instructions.
   - Configure PXE VLAN ID/Priority – See Section 2.4.4.3, Configuring the PXE VLAN ID and Priority, for instructions.
   - Identify ports – See Section 2.4.4.4, Physically Identifying the Port, for instructions.
   - Erase the port and adapter configuration – See Section 2.4.4.5, Erasing Ports and Adapter Configurations, for instructions.
   - Configure the LPVID for each channel – See Section 2.4.5.2, Configuring LPVID, for instructions.

NOTE An LPVID is optional in IBM Virtual Fabric mode, but it is required for every function in Switch Independent mode.

If IBM Virtual Fabric mode or IBM Unified Fabric Protocol mode is enabled, you must configure the minimum and maximum bandwidth settings on the switch.

2.4.7 Advanced Mode Support

Advanced mode provides driver compatibility. With Advanced mode enabled, you can run Advanced mode-aware drivers that provide the advanced capabilities listed in the following table. With Advanced mode disabled, you can run older legacy inbox drivers that are not Advanced mode-aware with the latest firmware versions.

NOTE Advanced mode support is enabled by default on OCe11100-series two-port and four-port adapters. On four-port adapters, the Advanced mode setting is not provided in the PXESelect utility. The Advanced mode setting on these platforms is implicitly enabled, and Advanced mode-aware drivers must be installed. Compatibility with legacy drivers requires that Advanced mode support be disabled on two-port adapters.

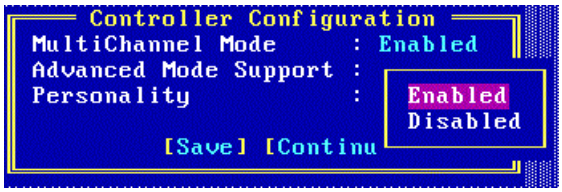
Table 1 Advanced Mode Capabilities (by Operating System)

Operating System	Advanced Mode Enabled	Advanced Mode Disabled (Legacy Mode)
Windows	16 RSS queues **NOTE** Supported on Windows Server 2008 R2 and later versions only.	4 RSS Queues
	VMQ lookahead split **NOTE** VMQs are supported on Windows Server 2008 R2 and later versions only.	Lookahead split is silently ignored. A small performance penalty for VMQs may occur.
Linux and Citrix	16 RSS queues	4 RSS queues
	VFs and PFs can be increased up to 30	
VMware ESXi	For both 1500 and 9000 MTU: 16 NetQueues per PF in non-VFA 4 NetQueues per PF in VFA	**1500 MTU** – 8 NetQueues per PF in non-VFA and 4 NetQueues per PF in VFA. **9000 MTU** – 4 NetQueues per PF in both VFA and non-VFA

To configure Advanced mode support:

1. From the following Controller Configuration dialog (Figure 11), use the Tab key to select Advanced Mode Support. A pop up appears.

Figure 11 Controller Configuration Dialog – Advanced Mode Support Selection



2. Select **Enabled** or **Disabled** and press Enter.
3. Select Save and press Enter.

2.4.8 Personality Option

The *personality* reflects the adapter protocols. This option specifies a list of available protocols that can be configured on an adapter. Depending on the personalities for which the adapter is licensed, one or more of the following selections appear:

- NIC
- iSCSI
- FCoE

The menu displays the available personalities only, including both free and licensed personalities.

The NIC personality implies that all of the enabled functions provide NIC/TOE functionality.

iSCSI and FCoE personalities are enabled on one function per adapter port and include NIC functionality on the other enabled functions. Only one storage protocol is allowed on each port.

To select the personality of the adapter:

1. From the Controller Configuration dialog (Figure 12), use the Tab key to select Personality and press Enter. A menu appears.

Figure 12 Controller Configuration Dialog – Personality Selection



2. Select NIC, iSCSI, or FCoE and press Enter.
3. Select Save and press Enter.

NOTE If the Personality setting has changed since the last boot, a reboot will occur when you exit the PXESelect utility.

2.5 PXE Boot Parameters Default Values

The default settings for the PXE Boot parameters are listed in the following table.

Table 2 PXE Boot Parameter Default Values

Parameter	Default Value	Valid Values
Advanced Mode	Enabled two-port and four-port adapters	Enabled Disabled
PXE Boot Support	The default for this parameter varies depending on the vendor configuration	Enabled Disabled
SR-IOV	Disabled	Enabled Disabled
VLAN Support	Disabled	Enabled Disabled
VLAN ID	0	0 to 4094

Table 2 PXE Boot Parameter Default Values (Continued)

Parameter	Default Value	Valid Values
VLAN Priority	0	0 to 7
Multichannel Mode	The default for this parameter varies depending on the vendor configuration	For Emulex adapters: ■ Enabled ■ Disabled For Lenovo System x adapters: ■ IBM Virtual Fabric Mode ■ IBM Unified Fabric Protocol Mode ■ Switch Independent Mode ■ Disabled
Minimum Bandwidth	0%	Must have a total of 100% across all ports
Maximum Bandwidth	0%	From the minimum bandwidth value up to 100% (inclusive)
LPVID	0	2 to 4094

Chapter 3: Configuring PXE Boot for NIC on OCe14000-Series Adapters

NOTE This section applies only to Emulex OCe14000-series adapters. See Chapter 2, Configuring the PXE Boot for NIC on OCe11100-Series Adapters, for information specific to those adapters.

For a Dell OCe14000-based adapter, see Appendix G, Configuring PXE Boot for NIC on Dell OCe14000-Based Systems.

This section describes using and configuring PXE to boot computers using a network interface independent of available data storage devices (such as hard disks) or installed operating systems.

The PXE protocol is a combination of DHCP and TFTP with subtle modifications to both. DHCP locates the appropriate boot server or servers, and TFTP downloads the initial bootstrap program and additional files.

Network booting enables you to perform the following tasks:

- Boot diskless systems such as thin clients and dedicated systems.
- Deploy software and operating systems for your systems.
- Automate system maintenance, such as backups.
- Automate system checking, such as virus scanning.
- Ensure system security.

3.1 Pre-OS

PXE can be used in a *pre-OS* environment. Pre-OS is the process of loading a small operating environment to perform a client management task before loading the final operating system from the local hard drive. For example, with a pre-OS you can scan the hard drive for viruses. This guarantees that the client is not infected before it starts. Another example is the use of this feature by WDS to install operating systems on local disks or boot from SAN disks.

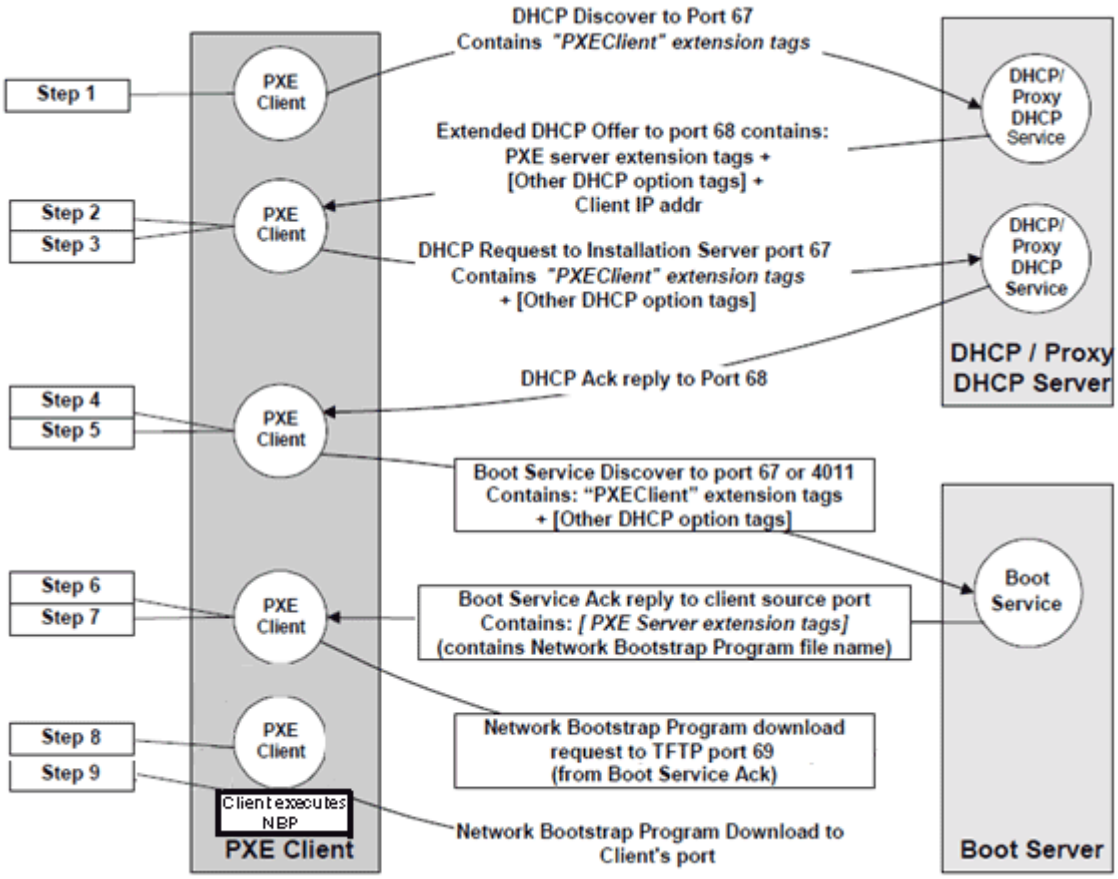
3.2 PXE Boot Process

After PXE boot is enabled in the system UEFI and BIOS, the PXE client can boot up and start up the PXE boot ROM. This is the boot code physically located on the NIC adapter.

Figure 13, PXE Boot Process shows the boot process.

NOTE To enable or disable PXE boot, it must be enabled or disabled in the system UEFI and BIOS; Refer to the documentation that accompanied the server for more information.

Figure 13 PXE Boot Process



1. The PXE boot ROM sends a DHCP request extended with PXE specific options (step 1 in the figure).
2. The DHCP responses contain the DHCP options (DHCPOFFERs) that include the NBP filename and boot server lists (steps 2 through 5).
3. The PXE client attempts to download the specified NBP over TFTP from one of the specified boot servers (steps 6 and 7).
4. The PXE client executes the downloaded NBP (steps 8 and 9).

NOTE If any of these steps fail, the boot process typically continues using the next available device in the boot sequence, depending on the system configuration and boot order.

3.3 Remotely Installing with PXE for Windows Server 2008, 2008 R2, 2012, 2012 R2, and 2016

For remote installation with PXE, a network driver for the Emulex adapter must be part of the client's installation image on the server. The current versions of Windows Server 2008 and Windows Server 2008 R2 do not include network drivers for the Emulex adapter; however, Windows Server 2012, 2012 R2, and 2016 include the network driver for the Emulex adapter.

To add the image and installation using the driver with the Remote Installation Service Setup:

Select **Start Menu > Programs > Administrative Tools > Remote Installation Services Setup**

In addition to the network driver for the Emulex adapter, you must configure the following services to use PXE for remote installations:

- DHCP server
- Remote Installation Services
- Windows Deployment Services

Microsoft provides extensive documentation on deploying its operating systems for remote installations, and different setups may be required depending on your individual implementation. Microsoft provides step-by-step guides for its Windows Deployment Services for configuring the server, adding images, and installing an operating system. It also includes instructions for more advanced tasks like creating multicast transmissions, creating custom images, and performing an unattended installation.

For detailed information on deploying and managing remote installations on Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, and Windows Server 2016, see the Microsoft website and visit Microsoft TechNet. Search on the bulleted terms above to access Microsoft's wide-ranging documentation on these subjects.

3.4 Remotely Installing with PXE for Linux and Citrix Servers

Linux allows for PXE installation over a network using the NFS, TFTP, or HTTP protocols. If the system to be installed contains an Emulex NIC or adapter with PXE Boot support, it can be configured to boot from files on another networked system rather than local media.

The Linux distributions provide extensive documentation on deploying and managing remote installations of the Linux operating system using PXE. See the distribution's documentation for instructions on deploying a PXE installation over the network.

For remote installation with PXE, a network driver for the Emulex adapter must be part of the client's installation image on the server. The current Linux distribution may or may not include network drivers for Emulex adapters. If it does, the driver may need to be added to the operating system's installation image, or added during installation. See the distribution's documentation for instructions on adding drivers during installation.

3.5 Using the PXESelect Utility

This section describes how to configure the PXE boot options using the PXESelect utility.

NOTE For Dell OCe14000-based adapters, see Appendix G, Configuring PXE Boot for NIC on Dell OCe14000-Based Systems, for information on using the PXESelect utility.

3.5.1 Navigating the PXESelect Utility

Use the following methods to navigate the PXESelect utility:

- Press the up or down arrows on the keyboard to move through and select menu options or configuration fields. If multiple adapters are listed, use the up and down arrows to scroll to the additional adapters.
- Press the left or right arrows on the keyboard to move through and select setting options.

- Press Enter to accept a selection, select a menu option, or change a configuration default.
- Press the F7 key to save the changes on the screen.
- Press Esc to return to the previous menu or page, cancel a selection or dialog box, or exit the utility.

3.5.2 Running the PXESelect Utility

To run the PXESelect utility, start or restart the computer. When prompted, press Ctrl-P. If you are running the PXESelect utility with multiple adapters, all of the adapters are displayed when you start the utility. For example:

```
Press <Ctrl><P> for PXESelect(TM)Utility

OCe14102B-UM NIC#0 Port#1 Base 0x00000380001A0000 @ Bus:1B Dev:00 Fun:00
OCe14102B-UM NIC#0 Port#2 Base 0x00000380001E0000 @ Bus:1B Dev:00 Fun:00
- Initializing ...Done.
```

The PXE Configuration menu appears after the boot BIOS initializes to begin the PXE configuration.

NOTE A UEFI-capable system typically does not display the prompt for running the PXESelect utility unless it is configured for legacy booting. See your system configuration manual for information on performing a legacy boot. All configurations that can be performed in the PXESelect utility can instead be performed in the UEFI configuration utility. See Chapter 11, Configuring UEFI for Ethernet, for more information.

3.5.3 Setting Up a PXE Bootable Network

After the PXE boot BIOS is initialized, you can use the PXESelect utility to set up a PXE bootable network by configuring the adapters.

To configure adapters for PXE boot:

1. From the following Controller List (Figure 14), use the up or down arrow keys to select the adapter you want to configure and press Enter.

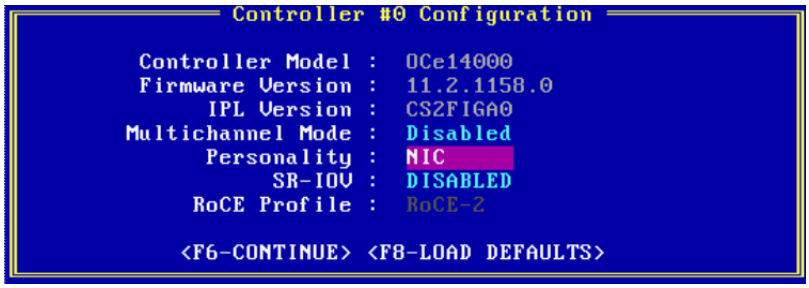
Figure 14 Controller List



NOTE The Controller List appears only if two or more adapters are connected.

The following Controller Configuration dialog (Figure 15) appears.

Figure 15 Controller Configuration Dialog



Depending on the adapter, one or more of the following items are displayed:

- **Controller Name** – Displays the adapter's name.
- **Firmware Version** – Displays the current firmware version.
- **IPL Version** – Displays the IPL version.
- **MultiChannel** – If UMC is enabled, you can access up to 16 virtual network interfaces. For more information on UMC, see Section 3.5.5, Configuring Universal Multichannel.

  NOTE If you are using a Lenovo System x adapter, see Section 3.5.6, Configuring Multichannel for Lenovo System x Adapters, for more information.

  On HP systems, you can select either Disable or Flex-10 for the MultiChannel setting. If Flex-10 is selected, the settings are only displayed and cannot be configured.
- **Personality** – This option specifies a list of available protocols that can be configured on an adapter. For more information on adapter personality, see Section 3.5.4.1, Configuring Personality.
- **SR-IOV** – If the system BIOS supports SR-IOV, you can enable it if UMC is disabled or it is not supported. For more information on SR-IOV configuration, see Section 3.5.4.2, Configuring SR-IOV.
- **RoCE Profile** – This setting is available only if the Personality is set to NIC+RoCE. For more information, see Section 3.5.4.1.2, Configuring the RoCE Profile.

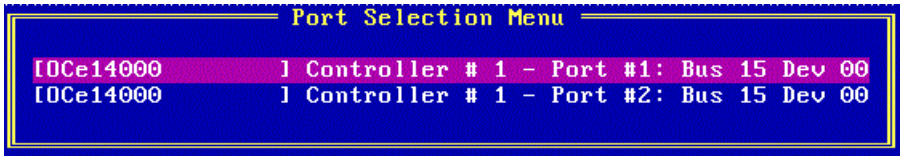
  NOTE RoCE is available as Technical Preview only. Do not use RoCE in a production environment.
- **NIC Mode** – This setting is available only if SR-IOV is enabled, Personality is set to NIC, and a two-port Emulex OCe14000-series adapter is in use. For more information, see Section 3.5.4.3, Configuring NIC Mode.

2. After you set the options in the Controller Configuration dialog, press the F7 key to save the current settings.
3. Press F6 key to proceed.

   NOTE To erase the current configuration, press the F8 key. See Section 3.5.4.9, Erasing Ports and Adapter Configurations, for more information.

   The following Port Selection Menu appears.

Figure 16 Port Selection Menu



NOTE The Port Selection Menu appears only if two or more ports are connected.

4. Select the port you want to configure and press Enter. The Port Menu appears.

   Depending on the installed adapter and the Multichannel setting, proceed to the following sections:

   — If UMC is disabled or not supported, see Section 3.5.4, Configuring Ports If UMC Is Disabled or Not Supported, for additional configuration information.

   — If UMC is enabled, see Section 3.5.5, Configuring Universal Multichannel, for additional configuration information.

   — If UMC is enabled and you are using Lenovo System x adapters, see Section 3.5.6, Configuring Multichannel for Lenovo System x Adapters, for additional configuration information.

3.5.4 Configuring Ports If UMC Is Disabled or Not Supported

NOTE If UMC is enabled, see Section 3.5.5, Configuring Universal Multichannel. If you are using a Lenovo System x adapter, see Section 3.5.6, Configuring Multichannel for Lenovo System x Adapters.

3.5.4.1 Configuring Personality

The *personality* reflects the adapter's protocols. This option specifies a list of available protocols that can be configured on an adapter. Depending on the personalities for which the adapter is licensed, one or more of the following selections may appear:

- **NIC** – Implies that all the enabled functions provide NIC functionality
- **iSCSI or FCoE** – These personalities are enabled on one function per adapter port and include NIC functionality on the other enabled functions

NOTE Only one of each storage protocol is allowed on each port.

NOTE Changing the personality to iSCSI or FCoE on one physical port also changes the personality on all other ports from the same adapter.

- **NIC+RoCE** – Implies that all the enabled functions provide RoCE functionality.
- **Custom** – For information about custom personalities, see Section 3.5.4.1.1, Configuring the Custom Personality Selection.

NOTE If NIC+RoCE is enabled, the RoCE Profile selection is available.
If SR-IOV is enabled, the NIC+RoCE personality setting is unavailable.
RoCE is supported on Windows and Linux only.

The menu displays the available personalities only, including both free and licensed personalities.

NOTE Two iSCSI functions and two FCoE functions are not allowed on a single port.

To select the personality of the adapter:

1. From the Controller Configuration dialog (Figure 15), use the up or down arrow keys to select Personality.
2. Use the left or right arrow keys to select the desired personality.
3. Press the F7 key to save.

NOTE If the Personality setting has changed since the last boot, a reboot will occur when you exit the PXESelect utility.

3.5.4.1.1 Configuring the Custom Personality Selection

The OCe14000-series adapters support the Custom personality mode when UMC is disabled. The OCe14000-series adapters support NIC, iSCSI, and FCoE protocols on a single port. While the normal NIC, iSCSI, and FCoE personalities support only NIC, NIC+iSCSI, or NIC+FCOE configuration on both ports, the Custom mode allows support for NIC, iSCSI, and FCoE on both ports.

NOTE When using the Custom personality mode with UMC disabled, only three functions are available on each port.

With UMC disabled, each port can be configured with one of the following configurations:

- One NIC function
- One NIC function and one iSCSI function
- One NIC function and one FCoE function
- One NIC function, one iSCSI function, and one FCoE function

NOTE Two of the same storage protocols are not supported on a single port.
Two NIC functions are not supported on a single port.

Therefore, on a 2-port controller you can potentially have 6 functions, while a 4-port controller can have up to 12 functions. The function-to-port mapping is detailed in the following two tables.

Table 3 Function-to-Port Mapping for 4-Port OCe14000-Series Adapters

Port 0	Port 1	Port 2	Port 3
0	1	2	3
4	5	6	7
8	9	10	11

Table 4 Function-to-Port Mapping for 2-Port OCe14000-Series Adapters

Port 0	Port 1
0	1
2	3
4	5

If you enable Custom mode, the personality may revert to NIC, iSCSI, or FCoE if an actual Custom configuration is not defined. If Custom mode is enabled, the following conditions apply:

- If you leave all functions set to NIC, the configured personality is automatically switched from Custom to NIC.
- If you leave the functions configured the same as they would be for iSCSI, the configured personality is automatically switched from Custom to iSCSI.
- If you leave the functions configured the same as they would be for FCoE, the configured personality is automatically switched from Custom to FCoE.

NOTE If higher order functions have been enabled and you want to disable them, you must first enable the Custom personality before the option to disable the functions appears.

To select the Custom personality:

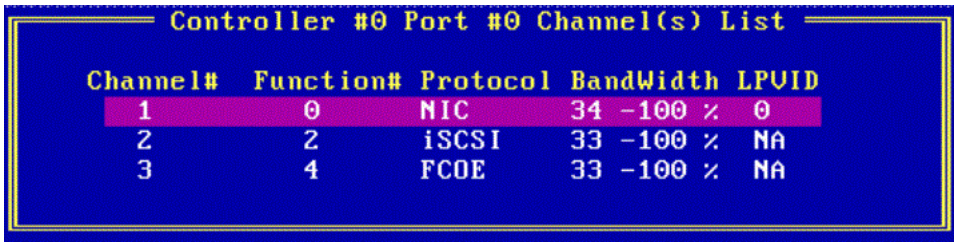
1. From the Controller Configuration dialog (Figure 15), use the up or down arrow keys to select the Personality option.
2. Use the left or right arrow keys to select the Custom personality.
3. Press the F6 key to continue. The Port Selection Menu (Figure 16) appears.
4. On the Port Selection Menu, select the appropriate port. The following Port Menu appears.

Figure 17 Port Menu (UMC Disabled)



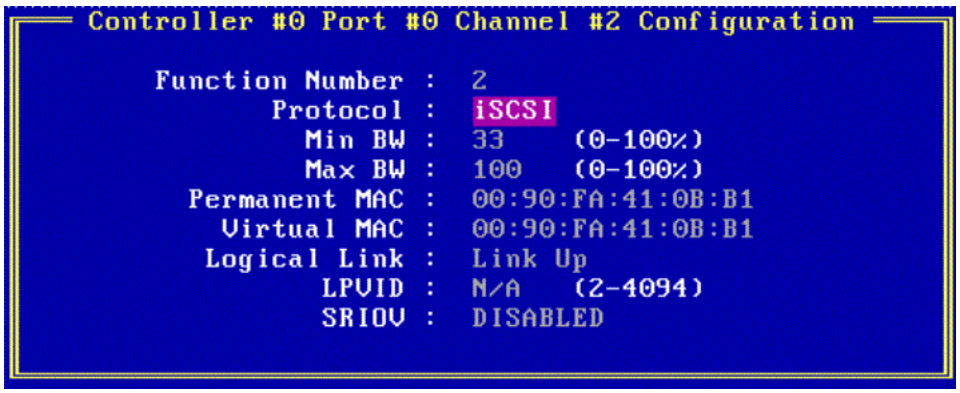
5. On the Port Menu, select NIC Configuration and press Enter. The following Channels List appears.

Figure 18 Channels List (UMC Disabled)



6. On the Channels List, select the channel you want to configure and press Enter. The following Channel Configuration dialog appears.

Figure 19 Channel Configuration Dialog (UMC Disabled)



7. On the Channel Configuration dialog, use the up or down arrow keys to select Protocol.
8. Select the desired personality for the appropriate function using the arrow keys.

NOTE The iSCSI and FCoE personalities can only be enabled on one function per adapter port.

9. Press the F7 key to save.

NOTE If the Personality setting has changed since the last boot, a reboot will occur when you exit the PXESelect utility.

3.5.4.1.2 Configuring the RoCE Profile

NOTE The RDMA application profile setting is available only when the NIC+RoCE personality is selected.

NOTE The RDMA application profile setting is available only when the Virtualization mode is set to None.

The RDMA application profile settings allow you to select the appropriate configuration for the particular system setup. The RDMA application profile choices include:

- **RoCE-1** – Not supported.
- **RoCE-2** – The RoCE-2 profile supports Windows SMB Direct, Linux iSER, NFS-RDMA, and VM Migration.

To select a RoCE Profile:

1. From the Controller Configuration dialog (Figure 15), select the NIC+RoCE personality.
2. Use the up or down arrow keys to select RoCE Profile.
3. Use the left or right arrow keys to select the desired profile and press Enter.
4. Press the F7 key to save.

3.5.4.2 Configuring SR-IOV

If the system BIOS supports SR-IOV, you can enable it on the adapter. SR-IOV support can be enabled only if Multichannel is disabled or not supported.

NOTE If SR-IOV is enabled, the NIC+RoCE personality setting is unavailable.

To configure SR-IOV:

1. From the Controller Configuration dialog (Figure 15), use the up or down arrow keys to select SRIOV.
2. Use the left or right arrow keys to enable or disable SR-IOV.
3. Press the F7 key to save.

For more information on SR-IOV configuration, see the appropriate Emulex driver user guide for OneConnect adapters.

3.5.4.3 Configuring NIC Mode

NOTE This setting is available on two-port OCe14000-series adapters only if SR-IOV is enabled and the personality is set to *NIC*.

This setting allows you to control the VF count on a two-port OCe14000-series adapter. The choices include:

- **NIC** – The VF count is restricted to 31 VFs per port and QoS is supported for those VFs.

NOTE For two-port OCe14000B-series adapters only, this setting supports up to 63 VFs per port.

- **NIC-ETS Disabled** – The VF count is set to 63 VFs per port but QoS is not supported for those VFs.

NOTE This setting is not supported on OCe14000B-series adapters.

To select the NIC mode:

1. From the Controller Configuration dialog (Figure 15), ensure that the NIC personality is selected and SR-IOV is enabled.
2. Use the up or down arrow keys to select the NIC mode.
3. Use the left or right arrow keys to select the desired setting and press Enter.
4. Press the F7 key to save.

3.5.4.4 Loading Default Settings

If you want to erase the current configuration, from the Controller Configuration dialog (Figure 15) press the F8 key.

See Section 3.5.4.9, Erasing Ports and Adapter Configurations, for more information.

3.5.4.5 Selecting a Port

After you set the options in the Controller Configuration dialog (Figure 15), save your changes. To save the current settings, press the F7 key.

To proceed to the Port Selection Menu (Figure 16):

1. From the Controller Configuration dialog (Figure 15), press the F6 key. The Port Selection Menu (Figure 16) appears.

NOTE The Port Selection Menu appears if two or more ports are connected.

2. Select the port you want to configure and press Enter. The Port Menu (Figure 17) appears.

The Port Menu enables you to perform the following tasks:

- **View NIC configuration information** – See Section 3.5.4.6, Viewing NIC Configuration Information, for more information.
- **Configure boot options** – See Section 3.5.4.7, Configuring Boot Options, for instructions.
- **Configure port options** – See Section 3.5.4.8, Configuring Port Options, for instructions.

3.5.4.6 Viewing NIC Configuration Information

To view the NIC configuration information:

1. From the Port Menu (Figure 17), use the up or down arrow keys to select NIC Configuration and press Enter. Depending on the currently selected Personality setting, one of the following dialogs appear.
   - If the Personality setting is set to iSCSI or FCoE, the Channels List (Figure 18) appears.

     Select a channel using the up or down arrow keys and press Enter. The Channel Configuration dialog (Figure 19) appears.
   - If the Personality setting is set to NIC or NIC+RoCE, the Channel Configuration dialog (Figure 24) appears.
2. From the Channel Configuration dialog, you can view the following information:
   - Function number
   - Protocol
   - Minimum bandwidth (Min BW)
   - Maximum bandwidth (Max BW)
   - Permanent and virtual MAC addresses
   - Logical link status
   - Assigned LPVID
   - SR-IOV status

3.5.4.7 Configuring Boot Options

To view the Boot Configuration dialog:

1. From the Port Menu (Figure 17), select Boot Configuration and press Enter. The following Boot Configuration dialog appears.

Figure 20 Boot Configuration Dialog



From the Boot Configuration dialog, you can configure PXE boot support.

3.5.4.7.1 Configuring PXE Boot Support

To configure PXE boot support:

1. On the Boot Configuration dialog (Figure 20), use the left or right arrow keys to enable or disable the PXE Boot setting.
2. Press the F7 key to save.

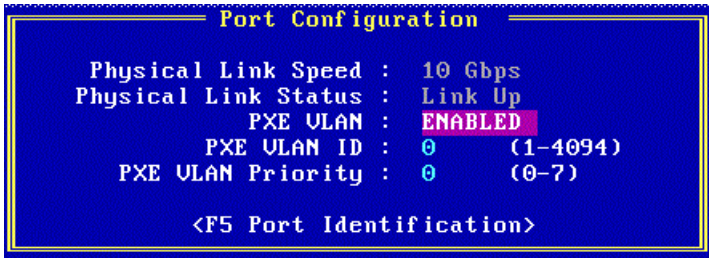
NOTE If PXE boot is enabled, during system startup PXE contacts the DHCP server for an IP address to boot from the network.

3.5.4.8 Configuring Port Options

To view the Port Configuration dialog:

1. From the Port Menu (Figure 17), select Port Configuration and press Enter. The following Port Configuration dialog appears.

Figure 21 Port Configuration Dialog



2. From the Port Configuration dialog, you can perform the following tasks:
   - View the physical link speed and link status.
   - Configure PXE VLAN, including the PXE VLAN ID and priority – See Section 3.5.4.8.1, Configuring PXE VLAN, for instructions.
   - Identify the port – See Section 3.5.4.8.2, Physically Identifying the Port, for instructions.

3.5.4.8.1 Configuring PXE VLAN

To enable PXE VLAN:

1. On the Port Configuration dialog (Figure 21), use the up or down arrow keys to highlight the current PXE VLAN setting.

2. Use the left or right arrow keys to set PXE VLAN to enabled.
3. Press the F7 key to save.

To configure a PXE VLAN ID:

1. On the Port Configuration dialog (Figure 21), use the up or down arrow keys to highlight the current PXE VLAN ID setting.
2. Use the left or right arrow keys to select a number from 1 to 4094.
3. Press the F7 key to save.

To set the priority level:

1. On the Port Configuration dialog (Figure 21), use the up or down arrow keys to highlight the current PXE VLAN Priority setting.
2. Use the left or right arrow keys to select a number from 0 to 7.

   This unique value assigns a priority to outbound packets containing a specified VLAN ID. Valid values range from 0 to 7, with 0 being the highest priority level.
3. Press the F7 key to save.

After you exit the PXESelect utility, the system will reboot for the configuration to take effect.

3.5.4.8.2 Physically Identifying the Port

To physically determine which port you are configuring by causing the link and activity LEDs of that port to blink:

1. On the Port Configuration dialog (Figure 21), press the F5 key.
2. The LEDs on the adapter begin blinking. The selected port's LED status indicators blink on the adapter for approximately 15 seconds.

NOTE Not all adapters have externally–visible LEDs. If you are using an add-in card in a blade server environment, port identification is not supported.

3.5.4.9 Erasing Ports and Adapter Configurations

NOTE If this setting is selected, all previous configuration settings are returned to their factory default settings. Performing this action ensures a clean environment for new configuration settings to take effect.

To erase the ports and adapter configuration:

1. On the Controller Configuration dialog (Figure 15), press the F8 key to restore the factory default settings. The system displays a warning asking if you want to erase the current configuration for all ports of the adapter.
2. Press Y to delete the configuration. The system displays another warning asking you to confirm the permanent removal of the configuration.
3. Press Y to delete the configuration.

To exit the PXESelect utility after erasing the ports and adapter configuration:

1. Follow the instructions on the bottom of the individual menus until you are prompted to exit.
2. Press Y to exit. Depending on the settings that were changed, a reboot may be necessary.

NOTE For older systems, depending on the memory allocation method supported, the PXESelect utility automatically reboots even if no changes were made to the system.

3.5.5 Configuring Universal Multichannel

NOTE If UMC is enabled and the adapter supports Lenovo System x Virtual Fabric mode or Lenovo System x UFP mode, see Section 3.5.6, Configuring Multichannel for Lenovo System x Adapters, for more information. If UMC is disabled or not supported on the adapter, see Section 3.5.4, Configuring Ports If UMC Is Disabled or Not Supported.

3.5.5.1 Enabling UMC

UMC provides the ability to configure multiple physical functions or I/O channels for each physical adapter port. For more information on UMC support, refer to the *Emulex Universal Multichannel Reference Manual*.

NOTE If UMC is enabled, you must configure the minimum and maximum bandwidths settings for iSCSI and FCoE storage functions in the NIC BIOS before they can be configured further from their respective utilities (iSCSI or FCoE BIOS). Otherwise, the Logical Link for that function will still appear as down, and you will not be able to log in to targets or find LUNs behind those targets.

PXE VLAN is not supported in UMC environments.

UMC functionality is supported on Emulex OneConnect adapters running in 10GbE or 40GbE mode only. UMC is not supported in 1GbE mode.

SR-IOV is not available if UMC is enabled.

To enable UMC:

1. From the Controller Configuration dialog (Figure 15), use the up or down arrow keys to select MultiChannel.
2. Use the left or right arrow keys to enable UMC.
3. Press the F6 key to proceed.

3.5.5.1.1 ARI Support

The OCe14000-series adapters support ARI, which provides the ability to configure additional PCI functions or I/O channels for each physical adapter port if UMC is enabled.

You can configure up to 16 functions on a one-port OCe14400-series adapter, up to 8 functions per port on a one or two-port OCe14100-series adapter, and up to 4 functions per port on a four-port OCe14100-series adapter.

The maximum number of functions allowed on an adapter is controlled by the adapter's IPL file and the system's support for ARI.

The following requirements must be met to support more than eight functions on an adapter.

- The system hardware (the motherboard and BIOS) must support ARI.
- ARI must be enabled in the system BIOS.
- The host operating system must support ARI:
  - Windows Server 2012 and newer versions
  - RHEL 6.4 and newer versions
  - SLES 11 SP2 and newer versions
  - SLES 12 and newer versions
  - ESXi 5.5 and newer versions
- The application management tools, including the Emulex OneCommand® CNA Manager application, must support ARI.
- ARI must be enabled in the firmware using the Emulex OneCommand CNA Manager application.

If these conditions are not met, you may be able to configure more than eight functions, but only up to eight functions will be running and discovered after a reboot.

3.5.5.2 Configuring More than Four Functions per Port on OCe14000-Series Adapters

By default on 1-port or 2-port OCe14000-series adapters, only four functions are enabled when UMC is enabled. To enable the other four functions, you must manually enable those functions for them to be operational.

To configure more than four functions per port on 1-port or 2-port OCe14000-series adapters, perform the following steps:

1. Ensure that UMC is enabled. By default, UMC enables four functions per port. See Section 3.5.5.1, Enabling UMC, for information on configuring UMC.
2. Select the Custom mode in the Controller Configuration dialog (Figure 15). This enables you to view all eight functions per port. See Section 3.5.5.3.1, Configuring the Custom Personality Selection, for information on selecting the Custom mode.

   NOTE The options to enable and configure more than four functions per physical port are not available unless Custom mode is selected.
3. To enable each function, change the function's personality type to the desired protocol. See Section 3.5.5.3.1, Configuring the Custom Personality Selection, for information on selecting the personality.
4. Assign the appropriate bandwidth and LPVID. See Section 3.5.5.7, Configuring Minimum and Maximum Bandwidth, for information on configuring the bandwidth and Section 3.5.5.8, Configuring LPVID, for information on configuring the LPVID.

3.5.5.3 Configuring Personality

The *personality* reflects the adapter's protocols. This option specifies a list of available protocols that can be configured on an adapter. Depending on the personalities for which the adapter is licensed, one or more of the following selections appear:

- **NIC** – Implies that all the enabled functions provide NIC functionality
- **iSCSI or FCoE** – These personalities are enabled on one function per adapter port and include NIC functionality on the other enabled functions

  NOTE Only one of each storage protocol is allowed on each port.
  Two iSCSI functions and two FCoE functions are not allowed on a single port.

  NOTE Changing the personality to iSCSI or FCoE on one physical port also changes the personality on all other ports from the same adapter.
- **Custom** – For information about custom personalities, see Section 3.5.5.3.1, Configuring the Custom Personality Selection.

The menu displays the available personalities only, including both free and licensed personalities.

To select the personality of the adapter:

1. From the Controller Configuration dialog (Figure 15), use the up or down arrow keys to select Personality.
2. Use the left or right arrow keys to select the desired personality.
3. Press the F7 key to save.

   NOTE If the Personality setting has changed since the last boot, a reboot will occur when you exit the PXESelect utility.

3.5.5.3.1 Configuring the Custom Personality Selection

The Custom personality allows you to select the protocol type for each function. The first function on the port must always be assigned the NIC personality. The iSCSI and FCoE personalities can be enabled on one function per adapter port only.

NOTE Two iSCSI functions and two FCoE functions are not allowed on a single port.

NOTE If higher order functions have been enabled and you want to disable them, you must first enable the Custom personality before the option to disable the functions appears.

To select the Custom personality:

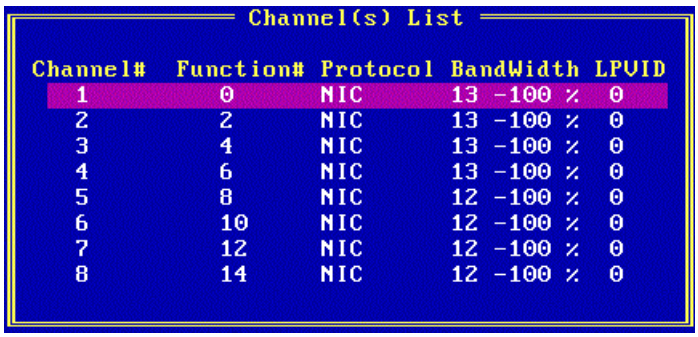
1. From the Controller Configuration dialog (Figure 15), use the up or down arrow keys to select **Personality**.
2. Use the left or right arrow keys to select the Custom personality.
3. Press the F6 key to continue. The Port Selection Menu (Figure 16) appears.
4. On the Port Selection Menu, select the appropriate port. The following Port Menu appears.

Figure 22 Port Menu (UMC Enabled)



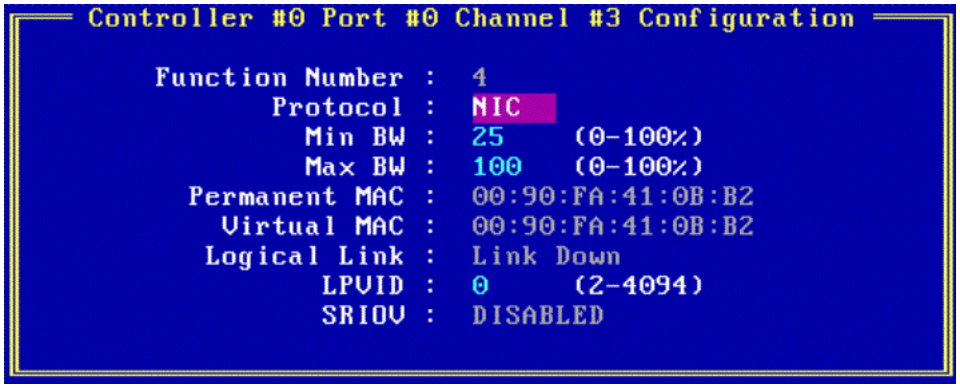
5. On the Port Menu, select MultiChannel Configuration and press Enter. The following Channels List appears.

Figure 23 Channels List (UMC Enabled)



6. On the Channels List, select the channel that you want to configure and press Enter. The following Channel Configuration dialog (Figure 24) appears.

Figure 24 Channel Configuration Dialog (UMC Enabled)



7. On the Channel Configuration dialog, use the up or down arrow keys to select Protocol.
8. Select the desired personality for the appropriate function using the arrow keys.

NOTE The iSCSI and FCoE personalities can only be enabled on one function per adapter port.

9. Press the F7 key to save.

NOTE If the Personality setting has changed since the last boot, a reboot will occur when you exit the PXESelect utility.

3.5.5.4 Loading Default Settings

If you want to erase the current configuration, from the Controller Configuration dialog (Figure 15), press the F8 key.

See Section 3.5.4.9, Erasing Ports and Adapter Configurations, for more information.

3.5.5.5 Selecting a Port

After you set the options in the Controller Configuration dialog, press the F7 key to save your changes.

To proceed to the Port Selection menu:

1. From the Controller Configuration dialog (Figure 15), press the F6 key. The Port Selection Menu (Figure 16) appears.

NOTE The Port Selection Menu appears if two or more ports are connected.

2. Use the up or down arrow keys to select the port you want to configure and press Enter. The Port Menu (Figure 22) appears.

The Port Menu enables you to perform the following tasks:

- Configure UMC options – See Section 3.5.5.6, Configuring UMC Options, for instructions.
- Configure boot options – See Section 3.5.4.7, Configuring Boot Options, for instructions.
- Configure port options – See Section 3.5.4.8, Configuring Port Options, for instructions.

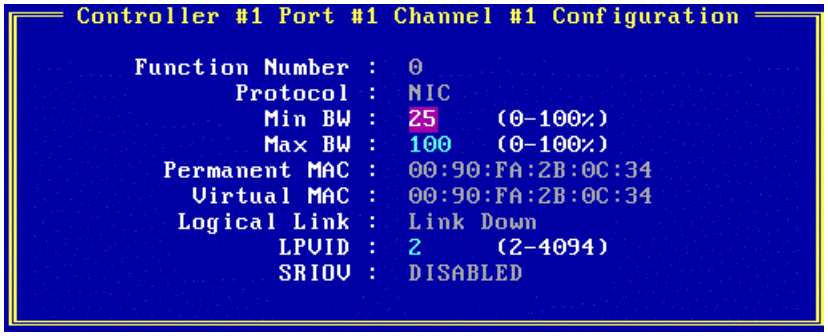
NOTE PXE VLAN configuration options are available only when UMC is disabled.

3.5.5.6 Configuring UMC Options

To configure the UMC options:

1. From the Port Menu (Figure 22), use the up or down arrow keys to select MultiChannel Configuration and press Enter. The Channels List (Figure 23) appears.
2. Select a channel using the up or down arrow keys and press Enter. The following Channel Configuration dialog appears.

Figure 25 Channel Configuration Dialog (UMC Enabled)



3. From the Channel Configuration dialog, you can perform the following tasks:
   - View the function number.
   - Select the protocol for each function, except for the first listed function. See Section 3.5.5.3, Configuring Personality, for instructions.
   - Configure the minimum and maximum bandwidth for each channel. See Section 3.5.5.7, Configuring Minimum and Maximum Bandwidth, for instructions.
   - View the permanent and virtual MAC addresses.
   - View the logical link status.

     NOTE Logical Link Status appears as N/A for storage functions.
   - Configure the LPVID for each channel. For more information, see Section 3.5.5.8, Configuring LPVID.
   - View the SR-IOV status.

     NOTE SR-IOV is automatically disabled if UMC is enabled.
     Your adapter or system may not support all UMC options.

3.5.5.7 Configuring Minimum and Maximum Bandwidth

To configure bandwidth:

1. On the Channel Configuration dialog (Figure 24), use the up or down arrow keys to select Min BW or Max BW.
2. Use the left or right arrow keys to select the desired bandwidth value.
   - The Minimum Bandwidth value is the least amount of bandwidth the function can provide. It is represented as a percentage. The Minimum Bandwidth value must be less than or equal to the Maximum Bandwidth value. The total of the Minimum Bandwidth values for all enabled functions on that port must be equal to 100.

     NOTE A Minimum Bandwidth value of 0 is a valid value. If all of the partitions' Minimum Bandwidth values are zero, the bandwidth is distributed equally among the current active partitions. If a specific partition's Minimum Bandwidth and Maximum Bandwidth values are 0, that partition's logical link will be brought down.

— The Maximum Bandwidth value is the greatest amount of bandwidth the function can provide. It is represented as a percentage.

3. Press the F7 key to save.

3.5.5.8 Configuring LPVID

The LPVID is used to enforce a VLAN ID on all traffic originating from an IP address, channel, or PCI function. If the operating system for that PCI function has set up a VLAN ID, then the operating system–configured VLAN ID takes precedence over the LPVID for transmit packets, while both the operating system–configured VLAN ID and LPVID–tagged packets will be received. If the operating system has not set up any VLAN IDs, then the LPVID is used for tagging.

NOTE LPVID and user-configured VLAN IDs from the operating system must be different.

LPVIDs must be configured on the switch port.

Each LPVID must be unique and is relevant for NIC traffic only. The LPVID is not supported for storage functions. For iSCSI storage functions, you must configure a VLAN ID through iSCSISelect or through the host. For more information, see Section 9.3.2, Configuring VLAN ID and VLAN Priority.

During PXE boot when the UNDI driver (BIOS) is functional, the PXE VLAN is used. However, after the NIC driver is operational the LPVID is used.

NOTE PXE VLAN is not supported in UMC environments.

For example:

```
PXE Install OS
PXE Server configured with VLAN 5
PXE VLAN=5
LPVID for function 0=5
```

To configure LPVID:

1. On the Configure Channel dialog (Figure 25), use the use the up or down arrow keys to select LPVID.
2. Use the left or right arrow keys to select the desired LPVID value. The LPVID range is 2 to 4094. A value of 0 disables the LPVID.
3. Press the F7 key to save.

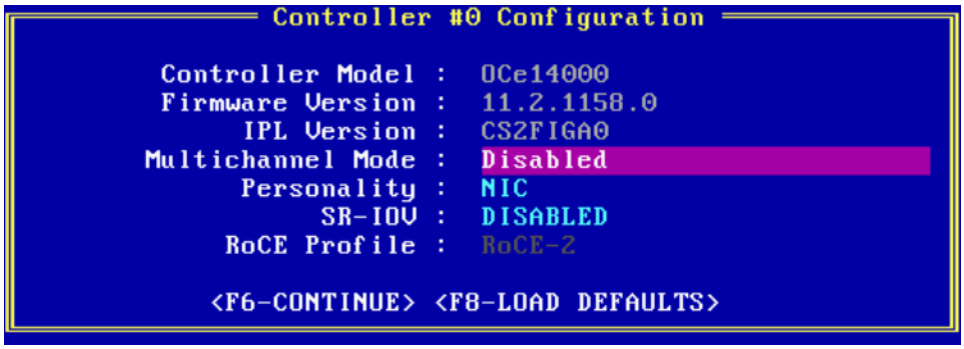
3.5.6 Configuring Multichannel for Lenovo System x Adapters

NOTE This section pertains to certain Emulex OneConnect adapters that support Multichannel for Lenovo System x Virtual Fabric or UFP only.

To select a Multichannel mode:

1. From the following Controller Configuration dialog (Figure 26), use the up or down arrow keys to select Multichannel Mode.

Figure 26 Controller Configuration Dialog - Lenovo System x Adapters



2. Use the left or right arrow keys to select one of the following settings and press Enter:
   - Switch Independent Mode
   - Virtual Fabric Mode (referred to as IBM Virtual Fabric Mode on some OCe14000-series adapters)
   - Unified Fabric Protocol Mode (referred to as IBM Unified Fabric Protocol Mode on some OCe14000-series adapters)
   - Disabled

   NOTE For more information on the available Multichannel modes, see Section 3.5.6.2, Multichannel Modes.
3. Press the F6 key to proceed.

With the Multichannel mode selected, you can now configure Multichannel for Lenovo System x support. See Section 3.5.6.3, Multichannel Configuration for Lenovo System x, for more information.

3.5.6.1 Configuring More than Four vNICs per Port on OCe14000-Series Adapters

By default on 1-port or 2-port OCe14000-series adapters, only four vNICs are enabled when Multichannel is enabled. To enable the other four vNICs, you must manually enable those vNICs for them to be functional.

To configure more than four vNICs per port, perform the following steps:

1. Ensure that Multichannel is enabled. By default, Multichannel enables four vNICs per port. See Section 3.5.6, Configuring Multichannel for Lenovo System x Adapters, for information on configuring Multichannel.
2. Select the Custom mode in the Controller Configuration dialog (Figure 26). This allows you to view all eight vNICs per port. See Section 3.5.5.3.1, Configuring the Custom Personality Selection, for information on selecting the Custom mode.
3. For each vNIC to be enabled, change the personality type to the desired protocol. See Section 3.5.5.3.1, Configuring the Custom Personality Selection, for information on selecting the personality.
4. Assign the appropriate bandwidth and LPVID. See Section 3.5.5.7, Configuring Minimum and Maximum Bandwidth, for information on configuring the bandwidth and Section 3.5.5.8, Configuring LPVID, for information on configuring the LPVID.

3.5.6.2 Multichannel Modes

NOTE The following modes are available only on Lenovo System x adapters and systems that support Virtual Fabric mode and Unified Fabric Protocol mode.

Additional PFs are supported when ARI is enabled on the system. For more information on ARI, see Section 3.5.5.1.1, ARI Support.

A Lenovo System x Virtual Fabric–enabled switch provides the ability to configure an LPVID for a virtual channel or I/O channel on an adapter port. If Multichannel for Lenovo System x is supported on the system, the PXESelect utility enables you to select one of the following modes:

- **Virtual Fabric Mode** – Select this mode if an Emulex OneConnect adapter is attached to an Lenovo System x Virtual Fabric–enabled switch. This mode supports up to four PFs per port.

  NOTE Virtual Fabric mode is not supported on 40GbE switches.

- **Unified Fabric Protocol Mode** – Select this mode if an Emulex OneConnect adapter is attached to a Lenovo System x UFP-enabled switch. In this mode, the maximum PFs per port for each adapter configuration is listed below:
  - 2-port 10GbE adapters – 8 PFs per port
  - 4-port 10GbE adapters – 4 PFs per port
  - 1-port 40GbE adapters – 16 PFs per port

  NOTE Some Lenovo System x switches support both Unified Fabric Protocol mode and Virtual Fabric mode.

- **Switch Independent Mode** – Select this mode if you are using a switch other than a Lenovo System x Virtual Fabric or UFP-enabled switch. In this mode, the maximum PFs per port for each adapter configuration is listed below:
  - 2-port 10GbE adapters – 8 PFs per port
  - 4-port 10GbE adapters – 4 PFs per port
  - 1-port 40GbE adapters – 16 PFs per port

  If this mode is enabled, see Section 3.5.5, Configuring Universal Multichannel, for information on configuring Multichannel for Lenovo System x.

  NOTE Multichannel for Lenovo System x is supported on Emulex OneConnect adapters running in 10GbE mode only. Multichannel for Lenovo System x is not supported in 1GbE mode.

3.5.6.3 Multichannel Configuration for Lenovo System x

Multichannel for Lenovo System x provides the ability to configure multiple PCI functions or I/O channels for each physical adapter port.

NOTE By default on 1-port or 2-port OCe14000-series adapters using either Virtual Fabric mode or Unified Fabric Protocol mode, only four vNICs are enabled when Multichannel is enabled. To enable the other four vNICs, you must manually enable those vNICs using the steps provided in Section 3.5.6.1, Configuring More than Four vNICs per Port on OCe14000-Series Adapters.

Setting up Multichannel for Lenovo System x Virtual Fabric or UFP depends on cooperation with adjacent switches.

See Section 3.5.5, Configuring Universal Multichannel, for information on configuring Multichannel for Lenovo System x if Switch Independent mode is enabled.

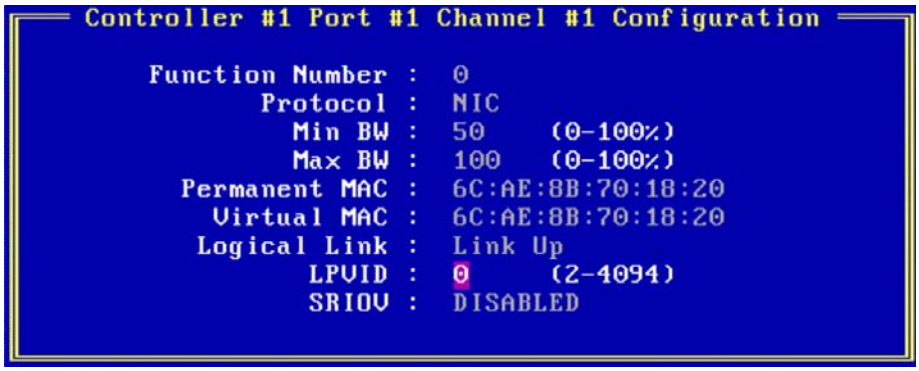
To view the configuration dialog:

1. From Port Selection Menu (Figure 16), select the port you want to configure and press Enter. The configuration dialog appears.

   If you are configuring Multichannel for Lenovo System x, the configuration options will depend on the Multichannel mode selected.

NOTE For illustration purposes, the dialogs in this section are for a system with Virtual Fabric mode selected.

Figure 27 Virtual Fabric Mode MultiChannel Configuration Dialog



NOTE If the system does not support some of the Multichannel options, the unavailable options are disabled on the Channel Configuration dialog.

2. From the Channel Configuration dialog (Figure 27), you can perform the following tasks:
   - View the function number.
   - Select the protocol for each function, except for the first listed function. See Section 3.5.5.3, Configuring Personality, for instructions.
   - Configure the minimum and maximum bandwidth for each channel. See Section 3.5.5.7, Configuring Minimum and Maximum Bandwidth, for instructions.

     NOTE If Switch Independent mode is enabled, you must configure the minimum and maximum bandwidth settings on the switch.
   - View the permanent and virtual MAC addresses.
   - View the logical link status.

     NOTE Logical Link Status appears as N/A for storage functions.
   - Configure the LPVID for each channel. For more information, see Section 3.5.5.8, Configuring LPVID.

     NOTE An LPVID is optional for Virtual Fabric mode, but it is required for every function if Switch Independent mode is selected.
   - View the SR-IOV status.

     NOTE SR-IOV is automatically disabled if Multichannel is enabled.

3.6 PXE Boot Parameters Default Values

The default settings for the PXE Boot parameters are listed in the following table.

Table 5 PXE Boot Parameter Default Values

Parameter	Default Value	Valid Values
UMC or Multichannel Mode (Lenovo System x)	The default for this parameter varies depending on the vendor configuration.	For Emulex adapters: ■ Enabled ■ Disabled For Lenovo System x adapters: ■ Virtual Fabric Mode ■ Unified Fabric Protocol Mode ■ Switch Independent Mode ■ Disabled
Personality	The default for this parameter varies depending on the vendor configuration.	If UMC is enabled: ■ NIC ■ iSCSI ■ FCoE ■ Custom If UMC and SR-IOV are disabled: ■ NIC ■ iSCSI ■ FCoE ■ NIC+RoCE **NOTE** RoCE is supported on Windows and Linux only. If UMC is disabled and SR-IOV is enabled: ■ NIC ■ iSCSI ■ FCoE
SR-IOV	Disabled	Enabled Disabled
RoCE Profile **NOTE** This setting is only available if the NIC+RoCE personality is selected. **NOTE** This setting is only available if the Virtualization mode is set to None.	RoCE-2	RoCE-2
Minimum Bandwidth	0%	Must have a total of 100% across all enabled function on that port
Maximum Bandwidth	0%	From the minimum bandwidth value up to 100% (inclusive)
LPVID	0	2 to 4094
PXE Boot Support	The default for this parameter varies depending on the vendor configuration.	Enabled Disabled

Table 5 PXE Boot Parameter Default Values (Continued)

Parameter	Default Value	Valid Values
PXE VLAN	Disabled	Enabled Disabled
PXE VLAN ID	0	2 to 4094
PXE VLAN Priority	0	0 to 7

Chapter 4: Configuring Boot from SAN for the FCoE Protocol

When booting from SAN, the storage device is typically identified by its WWPN and a LUN. By extending the server system BIOS, boot from SAN capability is provided by the boot BIOS contained on an Emulex adapter in the server. If properly configured, the adapter then permanently directs the server to boot from a logical unit (disk) on the SAN as if it were a local disk.

NOTE Not all procedures are required. Emulex adapters usually ship from the factory with the latest version of boot code enabled, so you do not need to enable boot code in those cases. However, if boot code is not enabled, you must enable it. You may want to update the boot code if a more current version is available at http://www.broadcom.com.

This section describes how to set up a system to boot from SAN. The specific procedure to follow is determined by the system architecture and the operating system.

Table 6 Boot from SAN Procedures

Architecture	Operating System	Procedure
x86 and x64	Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2 Windows Server 2016	Section 4.2, Installing Windows Server on a Boot Disk (x64)
	Linux	Section 4.3.1, Configuring Boot from SAN on Linux, Citrix, or VMware (x86 and x64)
	VMware	
	Solaris (SFS driver)	Section 4.4.1, Configuring Boot from SAN on Solaris (x86 and x64)
PowerPC	Solaris (SFS driver)	Section 4.4.1, Configuring Boot from SAN on Solaris (x86 and x64)
UEFI-based systems	Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2 Windows Server 2016	Section 4.1.2, Configuring Boot from SAN on Windows (UEFI) Section 4.3.2, Configuring Boot from SAN on Linux or VMware (UEFI)

4.1 Windows Server

This section explains how to configure boot from SAN on Windows Server operating systems.

4.1.1 Configuring Boot from SAN on Windows (x64)

To configure boot from SAN:

1. If necessary, update the boot code on the adapter. See Chapter 7, Updating and Enabling Boot Code, for more information.
2. If necessary, enable the boot code on the adapter. See Chapter 7, Updating and Enabling Boot Code, for more information.
3. Enable the adapter to boot from SAN. See Section 6.3, Enabling an Adapter to Boot from SAN, for more information.
4. Configure boot devices. See Section 6.5, Configuring Boot Devices, for more information.

5. If desired, configure the boot options on the adapter. See Section 6.7, Configuring Advanced Adapter Parameters, for more information.
6. Install the operating system on a boot disk. For more information, see Section 4.2, Installing Windows Server on a Boot Disk (x64).

4.1.2 Configuring Boot from SAN on Windows (UEFI)

To configure boot from SAN:

1. If necessary, update the boot code on the adapter. See Chapter 7, Updating and Enabling Boot Code for more information.
2. If necessary, enable the boot code on the adapter. See Chapter 7, Updating and Enabling Boot Code for more information.
3. Configure boot devices. See Section 6.5, Configuring Boot Devices, for more information.
4. If desired, configure the boot options on the adapter. See Section 6.7, Configuring Advanced Adapter Parameters, for more information.
5. Install the operating system on a boot disk:
   - For a new installation of a UEFI-aware operating system on a UEFI-based x64 server, see Section 4.1.3, Installing a New Windows UEFI-Aware Operating System on a UEFI-Based x64 Server.
   - To direct a UEFI-based server to a Windows operating system image that is already installed on the SAN, see Section 4.1.4, Directing a UEFI-Based Server to a Windows Server Operating System Image Already Installed on the SAN.
   - To install Windows Server on an unformatted boot disk and configure the system to boot from the SAN disk drive, see Section 4.2, Installing Windows Server on a Boot Disk (x64).

4.1.3 Installing a New Windows UEFI-Aware Operating System on a UEFI-Based x64 Server

This installation procedure assumes LUNs are created in the SAN storage device and zoned appropriately to the host adapter's WWN.

To install a new Windows Server UEFI-aware operating system:

1. From the server system UEFI setup, verify that CD/DVD is the first device in the boot order list.
2. Enable the adapter BIOS setting to allow SAN boot in the Emulex UEFI configuration utility.
3. Configure the boot target and LUN in the Emulex UEFI configuration utility to select the desired target.
4. Boot the host server with the Windows Server DVD inserted. Follow the on-screen prompts to install the appropriate version of Windows Server.

   The Windows installation exposes all available and visible LUNs as disks and partitions numbered 0 to N, where N is the highest number available. These numbers typically are the LUN numbers assigned by the array.
5. Select the disk on which you want to install the operating system.
6. Follow system prompts in the Windows installation.

NOTE By default, Microsoft Windows Servers install with a GPT-formatted disk on an UEFI-aware server.

After the installation is complete, Windows Boot Manager, which is available from the Start Options menu of the host server UEFI, is populated with a media device path pointing to the Windows boot loader utility.

The Windows Boot Manager option is inserted as the first boot option in the boot order list of the host server UEFI. The CD/DVD is the second device in the boot order list.

Upon reboot, the system boots from the LUN set up on the SAN.

4.1.4 Directing a UEFI-Based Server to a Windows Server Operating System Image Already Installed on the SAN

This installation procedure assumes a LUN exists in the SAN storage device, is zoned appropriately to the host adapter's WWN, and a UEFI-aware operating system resides on the target LUN.

To direct an UEFI-based server to a Windows Server operating system Image (installed as UEFI-aware) already installed on the SAN:

1. Enable network boot in the Emulex UEFI configuration utility. See Section 13.4, Enabling an Adapter to Boot from SAN, for more information.
2. Configure the boot target and LUN in the Emulex UEFI configuration utility to point to the desired target. See Section 13.7, Adding Boot Devices, for more information.
3. Select Boot Manager from the system UEFI configuration manager.
4. Select Add Boot Option.
5. Identify the desired target in the list, and continue down the explorer path until you locate the `bootmgfw.efi` file. This file is the boot loader utility for the Windows Server UEFI-aware operating system installation.
6. Input a boot device description (for example, Win2K8_UEFI_SAN) and optional data (if desired) for this device and select Commit Changes.
7. From the Windows Boot Manager, select Change Boot Order.
8. Move the previous input description name (Win2K8_UEFI_SAN) to the desired position in the boot order.
9. Select Commit Changes. The Start Options list now reflects the boot order changes.

   Upon reboot, the server is able to boot from this target LUN on the SAN.

4.2 Installing Windows Server on a Boot Disk (x64)

This procedure installs Windows Server onto an unformatted disk drive and configures the system to boot from the SAN disk drive.

NOTE The computer's system BIOS may require that another adapter take precedence over the Emulex adapter during boot. If this occurs, you must disconnect or disable the other adapter. This lets you configure and build the operating system on the drive connected to the Emulex adapter.

To install Windows Server on a boot disk:

1. Install the Emulex driver for Windows as described in the *Emulex Drivers for Windows User Guide for OneConnect Adapters*.
2. In Windows Explorer, navigate to the folder you specified during driver installation.
3. In the `\AutoPilot Installer\Drivers\`*`drivername`* folder, open the folder that corresponds to your computer type, such as x64. *`drivername`* is the type of driver you downloaded.
4. Copy all the files in this folder onto a formatted floppy disk or a USB device.
5. Boot the target system with the Windows Server setup media. The system displays the Install Windows splash screen.
6. Configure the values for the language, the time and date, and the keyboard. Click Next. Another splash screen appears.
7. Click Install Now. The Where do you want to install Windows? pop up appears.

8. Click Load Driver. Browse to the floppy disk or USB device that contains the driver to load the driver for the appropriate operating system. After selecting the driver, the correct driver location and driver are displayed under the Select driver to be installed dialog.
9. Select Next. After the driver is loaded, the Where do you want to install Windows? pop up appears.
10. Select the drive you configured as the boot device. See Section 6.5, Configuring Boot Devices.

4.3 Linux, Citrix, and VMware Operating Systems

This section explains how to configure boot from SAN on Linux, Citrix, and VMware operating systems.

4.3.1 Configuring Boot from SAN on Linux, Citrix, or VMware (x86 and x64)

To configure boot from SAN:

1. If necessary, update the boot code on the adapter. See Chapter 7, Updating and Enabling Boot Code, for more information.
2. If necessary, enable the boot code on the adapter. See Chapter 7, Updating and Enabling Boot Code, for more information.
3. Enable the adapter to boot from SAN. See Section 6.3, Enabling an Adapter to Boot from SAN, for more information.
4. If necessary, enable spinup delay. See Section 6.7.2, Enabling or Disabling the Spinup Delay, for more information.
5. Configure boot devices. See Section 6.5, Configuring Boot Devices, for more information.
6. If desired, configure the boot options on the adapter. See Section 6.7, Configuring Advanced Adapter Parameters, for more information.
7. Use the driver on the operating system distribution disk to boot the system. If necessary, you can update the driver to the latest version.

4.3.2 Configuring Boot from SAN on Linux or VMware (UEFI)

To configure boot from SAN:

1. If necessary, update the latest boot code on the adapter. See Chapter 7, Updating and Enabling Boot Code, for more information.
2. If necessary, enable the boot code on the adapter. See Chapter 7, Updating and Enabling Boot Code, for more information.
3. Configure boot devices. See Section 6.3, Enabling an Adapter to Boot from SAN, for more information.
4. If desired, configure the boot options on the adapter. See Section 6.7, Configuring Advanced Adapter Parameters, for more information.
5. Use the driver on the operating system distribution disk to boot the system. If necessary, you can update the driver to the latest version.

4.4 Solaris Operating System

This section explains how to configure boot from SAN on the Solaris operating system.

4.4.1 Configuring Boot from SAN on Solaris (x86 and x64)

To configure boot from SAN:

1. If necessary, update the boot code on the adapter. See Chapter 7, Updating and Enabling Boot Code for more information.
2. If necessary, enable the boot code on the adapter. See Chapter 7, Updating and Enabling Boot Code for more information.
3. Enable the adapter to boot from SAN. See Section 6.3, Enabling an Adapter to Boot from SAN, for more information.
4. If necessary, enable spinup delay. See Section 6.7.2, Enabling or Disabling the Spinup Delay, for more information.
5. Configure boot devices. See Section 6.5, Configuring Boot Devices, for more information.
6. If desired, configure the boot options on the adapter. See Section 6.7, Configuring Advanced Adapter Parameters, for more information.
7. Boot the Solaris installation CD and follow the prompts.

NOTE If you need help determining the LUNs to select for boot from SAN, see Section 4.4.2, Determining LUNs to Select for Boot from SAN.

4.4.2 Determining LUNs to Select for Boot from SAN

To determine which LUNs to select:

1. Open a terminal window and leave it open.
2. In the terminal window, select the LUN you are going to use as the SAN boot disk (not the local drive) using the `luxadm probe` command. This shows all the available LUNs. Record this LUN information, which is used throughout this procedure. LUN 0 is used in the example:

```
luxadm probe

Found Fibre Channel device(s):
      Node WWN:50060e8003823800 Device Type:Disk device
      Logical Path:/dev/rdsk/c5t226000C0FF9833AFd6s2
      Node WWN:50060e8003823800 Device Type:Disk device
      Logical Path:/dev/rdsk/c5t226000C0FF9833AFd6s2
Node WWN:50060e8003823800 Device Type:Disk device
```

3. Copy the /dev/rdsk/*nnn* part of the path statement for a drive.
4. In the terminal window, use the `luxadm display` command to show the WWPN or the LUN for which you selected the path in the prior step:

```
luxadm display </dev/rdsk/nnn>
```

5. Record this LUN or WWPN information for use in the procedure.

Chapter 5: Configuring x86/x64 Platforms for the iSCSI Protocol

The iSCSI protocol is an IP-based standard for linking data-storage devices over a network and transferring data by carrying SCSI commands over IP networks. An iSCSI network consists of one or more iSCSI storage units (targets) connected through a copper or optical networking cable to 10GbE network switches and IP routers. One or more servers are connected to this network, which are responsible for transferring data to or from the storage units.

When an operating system receives a request, it generates the SCSI command and then sends an IP packet over an Ethernet connection. At the receiving end, the SCSI commands are separated from the request, and the SCSI commands and data are sent to the SCSI controller and then to the SCSI storage device. iSCSI also returns a response to the request using the same protocol.

5.1 Constructing a Basic iSCSI SAN

Three main components make up an iSCSI SAN:

- **iSCSI initiator** – The initiator allows a given machine to access the storage available in the iSCSI SAN. It requests all SCSI operations like read or write. An initiator is usually located on the host (server) side, either as hardware (iSCSI adapter) or software (iSCSI software initiator). To transport block (SCSI) commands over the IP network, an iSCSI driver must be installed on the iSCSI host. An iSCSI driver is included with the Emulex adapter. For more information on iSCSI initiators, see Chapter 9, Configuring and Managing the iSCSI Initiator with the iSCSISelect Utility.
- **iSCSI target** – An iSCSI SAN has one or more iSCSI targets, which house and make available the storage used within the SAN. The iSCSI target is the storage device itself or an appliance that controls and serves volumes or virtual volumes. The target is the device that performs the SCSI command or bridges it to an attached storage device. iSCSI targets can be disks, RAID arrays, or FC fabrics. For more information, see Chapter 10, Configuring and Managing iSCSI Targets with the iSCSISelect Utility.
- **Networking infrastructure** – The networking infrastructure in an iSCSI SAN uses Ethernet transport. The configuration and complexity of the storage network depend on its intended function and the required capabilities.

5.2 Managing an iSCSI Session

To transmit information from an iSCSI initiator to an iSCSI target, the initiator must first establish a session with the target through an iSCSI login process. The login process:

- Starts a TCP/IP connection
- Verifies that the iSCSI initiator has access to the iSCSI target (authentication)
- Allows negotiation of various parameters

5.2.1 Logging in to an iSCSI Session

An iSCSI session has two phases:

- **Login phase** – iSCSI parameters are negotiated using login requests and responses.
- **Full featured phase** – After authentication has occurred and operational parameters have been set, the initiator begins to perform SCSI I/O functions.

5.2.2 Security

Because the iSCSI protocol operates in the internet environment, security is critical. The iSCSI SAN uses the CHAP security method.

5.2.2.1 CHAP Authentication

CHAP authentication is used to periodically verify the identity of the initiator by the target using a challenge and response mechanism. The challenge and response is established on the initial link and may be repeated at any time afterward. For CHAP to work, the target must know the initiator's secret key, and the initiator must correctly respond to the challenge.

Although the authentication is only one-way, you can negotiate CHAP in both directions for mutual authentication with the help of the same secret set.

5.3 Configuring Boot from SAN for iSCSI

This section provides instructions for configuring boot from SAN for iSCSI on various operating systems using the iSCSISelect utility. It also provides information on using the iSCSISelect utility to perform an MPIO boot configuration.

5.3.1 Setting Up Boot from SAN for iSCSI

In the iSCSI target configuration, you have the option of setting dual network paths to a single boot LUN.

NOTE iSCSI must be enabled for the port before configuring a boot target.

5.3.1.1 Windows Server

To set up boot from SAN for iSCSI on Windows Server:

1. Log in to the iSCSISelect utility by pressing the Ctrl-S keys when prompted.
2. Use the iSCSISelect utility to log in to a target from one adapter and configure it for boot. For more information, see Chapter 9, Configuring and Managing the iSCSI Initiator with the iSCSISelect Utility and Chapter 10, Configuring and Managing iSCSI Targets with the iSCSISelect Utility.

   If the iSCSISelect utility is configured properly, a message during boot time indicates that the iSCSI disk was successfully connected.
3. Install a Windows Server operating system over the iSCSI LUN.

5.3.1.2 Linux and Citrix Servers

To set up boot from SAN for iSCSI on Linux or Citrix:

1. Log in to the iSCSISelect utility by pressing the Ctrl-S keys when prompted.
2. Use the iSCSISelect utility to log in to a target from one adapter and configure it for boot. For more information, see Chapter 9, Configuring and Managing the iSCSI Initiator with the iSCSISelect Utility and Chapter 10, Configuring and Managing iSCSI Targets with the iSCSISelect Utility.

   If the iSCSISelect utility is configured properly, a message during boot time indicates that the iSCSI disk was successfully connected.
3. Install a Linux operating system over the iSCSI LUN. For more information on installing the driver, refer to the *Emulex Drivers for Linux User Guide for OneConnect Adapters*, in the Documents and Downloads area of http://www.broadcom.com. For iSCSI boot, you will need the following information:

   For RHEL operating systems:

a. When running the `vmlinuz` command, include the dd parameter to indicate that a driver update will be done.
b. When prompted for the type of device your installation involves, choose **Specialized Storage Devices**.
c. When prompted to select the drives on which you will install the operating system or any drives you want to automatically mount to the system, select the iSCSI target device on the **Other SAN Devices** tab.

For SLES operating systems:

a. Set `withiscsi=1` to indicate that a driver update will be done.
b. When prompted, click Configure iSCSI Disks to perform target login.

NOTE You may need to run certain `iscsiadm` commands to load the disk after the iSCSI driver is installed. Refer to the iscsiadm documentation for more information.

5.3.1.3 ESXi Server

To install and configure boot from SAN on an ESXi server:

1. Log in to the iSCSISelect utility pressing the Ctrl-S keys when prompted.
2. Use the iSCSISelect utility to log in to a target from one adapter and configure it for boot. For more information, see Chapter 9, Configuring and Managing the iSCSI Initiator with the iSCSISelect Utility and Chapter 10, Configuring and Managing iSCSI Targets with the iSCSISelect Utility.

   If the iSCSISelect utility is configured properly, a message during boot time indicates that the iSCSI disk was successfully connected.
3. Install an ESXi Server operating system over the iSCSI LUN. For more information on installing the driver, refer to the *Emulex Drivers for VMware User Guide for OneConnect Adapters* in the Documents and Downloads area of http://www.broadcom.com.

5.3.2 Booting from SAN for iSCSI MPIO

MPIO support allows the initiator to log in to dual sessions to the same target. In this way I/O can be sent over either TCP/IP connection to the target. If one session fails another session can continue processing I/O without interrupting the application. In iSCSI target configuration, you have the option of setting dual network paths to a single boot LUN.

NOTE Although MPIO boot support allows the initiator to log in to multiple sessions, the iSCSI BIOS currently limits the number of sessions to two iSCSI sessions for a single boot LUN.

5.3.2.1 Windows Server

To install and configure MPIO on Windows Server:

1. Log in to the iSCSISelect utility by pressing the Ctrl-S keys when prompted.
2. Use the iSCSISelect utility to log in to a target from one adapter and configure it for boot by following the steps in Section 8.2, Setting Up a Basic iSCSI Boot Configuration.

   If the iSCSISelect utility is configured properly, a message during boot time indicates that the iSCSI disk was successfully connected.

NOTE Only one instance of the boot LUN must be visible to the server during the installation. The installation might fail if multiple instances of the boot LUN are available to the server. Disable the Spanning Tree Protocol on any ports connected to Windows Server hosts booting using iSCSI. The Spanning Tree Protocol is used to calculate the best

path between switches if there are multiple switches and multiple paths through the network.

3. Install a Windows Server operating system over the iSCSI LUN.
4. After Windows Server is installed, install the MPIO software (such as Microsoft iSCSI Initiator), and reboot the system.
5. After rebooting, verify that the boot LUN is detected by the MPIO driver. This can be done by either of the following two methods:
   — Look at the driver name for the disk device from Device Manager,

   Or

   — Use the MPIO GUI to check for device ID **MS FT2005iSCSIBusType_0x9** under the MPIO Devices tab. The Disk Manager does not show duplicate disks.
6. Reboot the system and log in to the iSCSISelect utility by pressing the Ctrl-S keys when prompted.
7. You can now add an additional boot path with the iSCSISelect utility.

   To add an additional boot path to the same iSCSI LUN:

   a. Follow the steps from Section 5.3.1.1, Windows Server, to set up the second path.
   b. At the Add/Ping iSCSI Target menu, set a valid ISID Qualifier before adding the target. The ISID Qualifier can be any number up to 65535. For more information on when an ISID value must be changed, see Section 5.3.2, Booting from SAN for iSCSI MPIO.
   c. After logging in to the target, from the iSCSI Target Configuration menu, set the Boot Target option for the second target to Yes.
   d. Press the F7 key to show the LUNs behind the target.
   e. Select the same LUN you chose for the initial boot LUN, then press the F3 key to set it to bootable.
8. After Windows Server boots, the MPIO installs drivers on the second path and prompts for reboot. Reboot the server.

5.3.2.2 Linux and Citrix Servers

To install and configure MPIO on Linux or Citrix:

1. Log in to the iSCSISelect utility by pressing the Ctrl-S keys when prompted.
2. Use the iSCSISelect utility to log in to a target from one adapter and configure it for boot by following the steps from Section 5.3.1.2, Linux and Citrix Servers.
3. After the operating system is installed, install the MPIO software and reboot the system.
4. After rebooting, verify that the boot LUN is detected by the MPIO driver. Perform either of the following steps:
   — To see if the multipath is running, type:
   ```
   # /sbin/multipath
   ```
   Or
   — To see if the multipath daemon is running, type:
   ```
   # ps -aelf | grep multipathd
   ```
5. If multipath is running, you can view the current multipath topology by typing
   ```
   # /sbin/multipath -l
   ```
6. Reboot the system and log in to the iSCSISelect utility by pressing the Ctrl-S keys when prompted.
7. Use the iSCSISelect utility to add an additional boot path to the same iSCSI LUN. To do this, follow these steps:

a. Follow the steps from Section 5.3.1.2, Linux and Citrix Servers, to set up the second path.
b. When you get to the Add/Ping iSCSI Target menu, make sure you set a valid ISID Qualifier before adding the target. The ISID Qualifier can be any number up to 65535. For more information on when an ISID value must be changed, see Section 5.3.2, Booting from SAN for iSCSI MPIO.
c. After logging in to the target, from the iSCSI Target Configuration menu, set the Boot Target option for the second target to Yes.
d. Press the F7 key to show the LUNs behind the target.
e. Select the same LUN you chose for the initial boot LUN, then press the F3 key to set it to bootable.
f. Follow the instructions on the bottom of the menus until you are prompted to exit.
g. Save your changes before you exit.
h. Press Y to exit the iSCSISelect utility.

8. After Linux boots, the MPIO installs drivers on the second path and prompts you to reboot. Reboot the server.

5.3.2.3 ESXi Server

To install and configure MPIO on an ESXi server:

1. Log in to the iSCSISelect utility by pressing the Ctrl-S keys when prompted.
2. Use the iSCSISelect utility to log in to a target from one adapter and configure it for boot by following the steps from Section 5.3.1.3, ESXi Server.
3. After the ESXi server is installed, install the MPIO software and reboot the system. The ESXi Server is MPIO by default.
4. Use the iSCSISelect utility to add an additional boot path to the same iSCSI LUN. To do this, follow these steps:
   a. Follow the steps from Section 5.3.1.3, ESXi Server, to set up the second path.
   b. When you access the Add/Ping iSCSI Target menu, set a valid ISID Qualifier before adding the target. The ISID Qualifier can be any number up to 65535. For more information on when an ISID value must be changed, see Section 5.3.2, Booting from SAN for iSCSI MPIO.
   c. After you log in to the target, from the iSCSI Target Configuration menu, set the Boot Target option for the second target to Yes.
   d. Press the F7 key to show the LUNs behind the target.
   e. Select the same LUN you chose for the initial boot LUN, then press the F3 key to set it to bootable.
   f. Follow the instructions on the bottom of the menus until you are prompted to exit.
   g. Save your changes before you exit.
   h. Press Y to exit the iSCSISelect utility.

   After the ESXi Server boots, the MPIO installs drivers on the second path and prompts for a reboot.
5. Reboot the server.

Chapter 6: Using the FCoE Boot BIOS Utility for x86 and x64 Architectures

Before using the FCoE BIOS utility, verify that the boot code is enabled on the adapter as described in Chapter 7, Updating and Enabling Boot Code.

NOTE This section reflects the most recent release of the FCoE BIOS utility. Some selections may not be available if you are using an older version of the utility.

After exiting the FCoE BIOS configuration utility, the system will reboot whether or not any changes have been made.

6.1 Navigating the FCoE BIOS Utility

Use the following methods to navigate the FCoE BIOS utility:

- Press the up and down arrows on the keyboard to move through and select menu options or configuration fields. If multiple adapters are listed, use the up and down arrows to scroll to the additional adapters.
- Press the Page Up key to scroll to the previous page.
- Press the Page Down key to scroll to the next page.
- Press Enter to select a menu option, to select a changed value, to select a row in a configuration dialog, or to change a configuration default.
- Press Esc to go back to the previous menu.

6.2 Starting the FCoE BIOS Utility

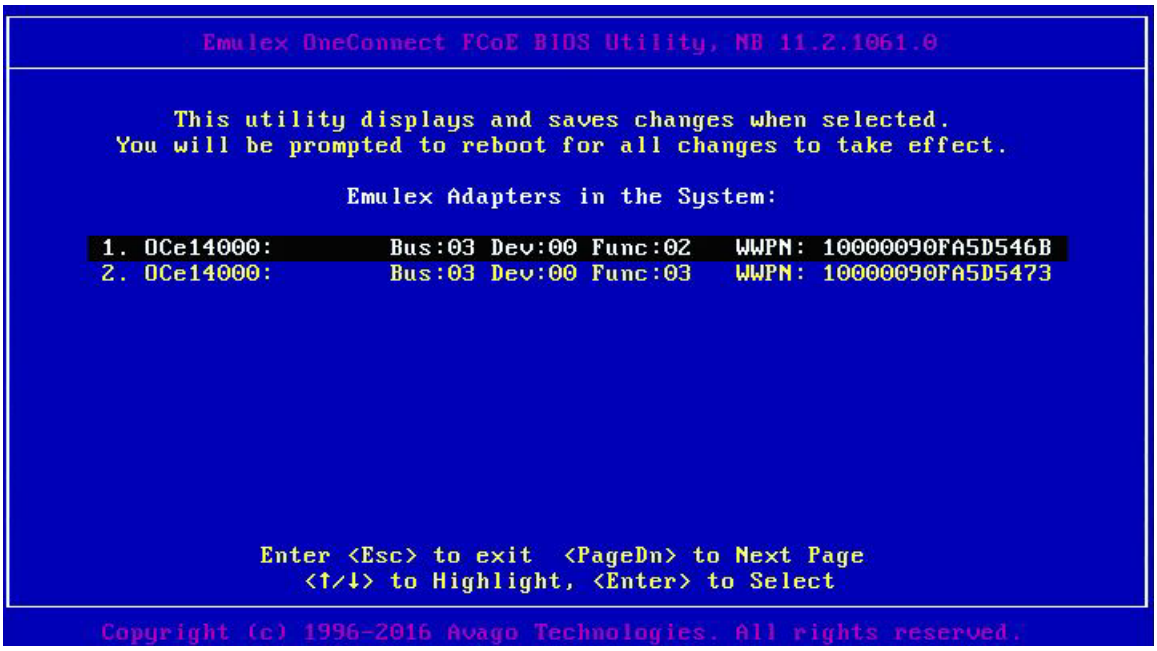
NOTE Links must be established before starting the FCoE BIOS utility; otherwise, you will receive an error message.

To start the FCoE BIOS utility:

1. Turn on the computer and press Alt-E or Ctrl-E immediately (within five seconds) when the Emulex bootup message to start the FCoE BIOS utility appears. An adapter listing (Figure 28) appears.

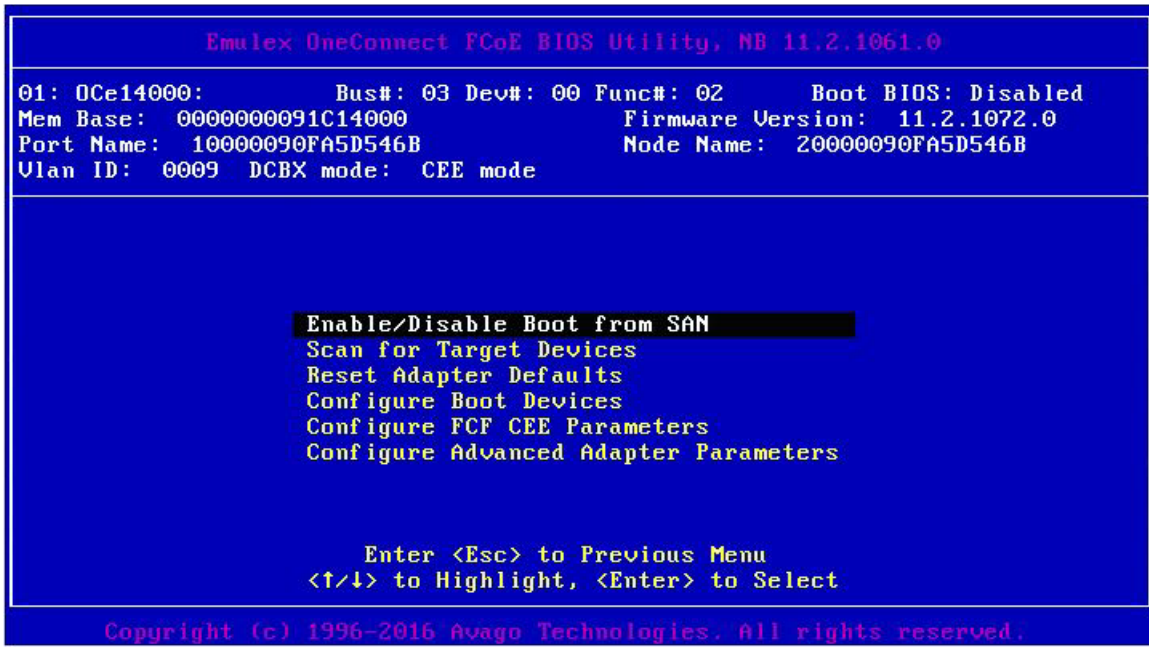
   NOTE If the bootup message does not appear, you must enable x86 BootBIOS. For more information, see Chapter 7, Updating and Enabling Boot Code.

Figure 28 Emulex Adapter Listing



2. Select the adapter to configure and press Enter. The following main configuration menu appears.

Figure 29 Main Configuration Menu



Under normal circumstances, you would first configure boot devices using the FCoE BIOS utility. However, if the adapter is not enabled to boot from SAN, you must first enable the adapter to do so.

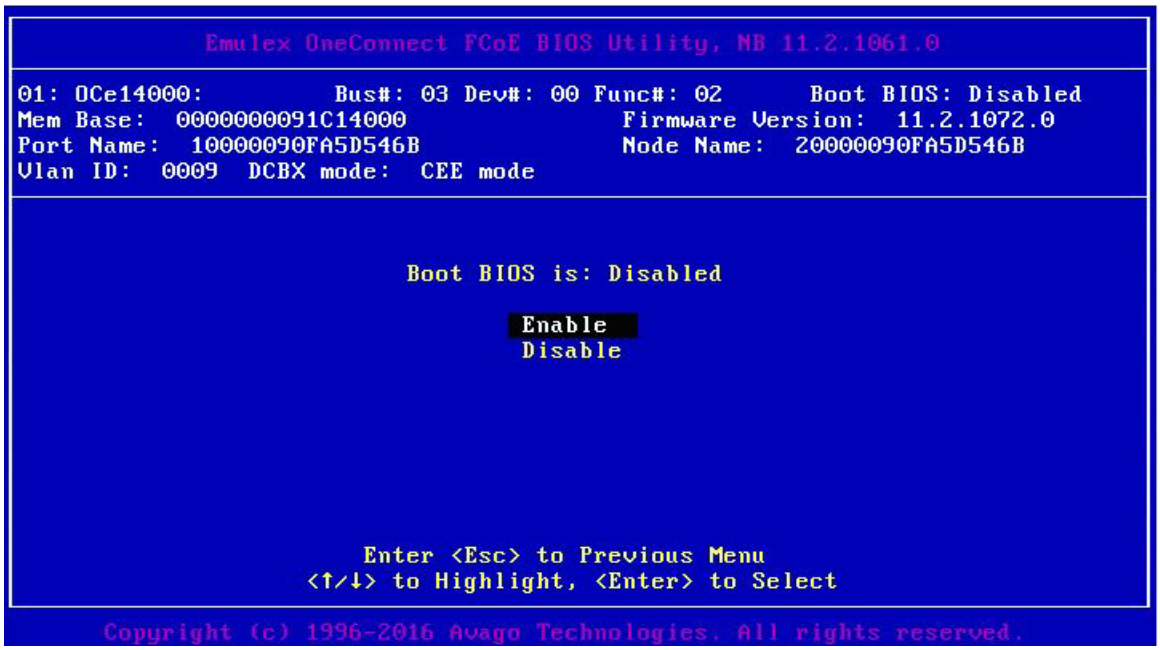
6.3 Enabling an Adapter to Boot from SAN

To enable an adapter to boot from SAN, from the main configuration menu (Figure 29), select **Enable/Disable Boot from SAN** and press Enter.

NOTE Adapters are disabled by default.

Boot from SAN must be enabled on one adapter to use remote boot functionality. After you enable an adapter, its boot status appears.

Figure 30 BIOS Status

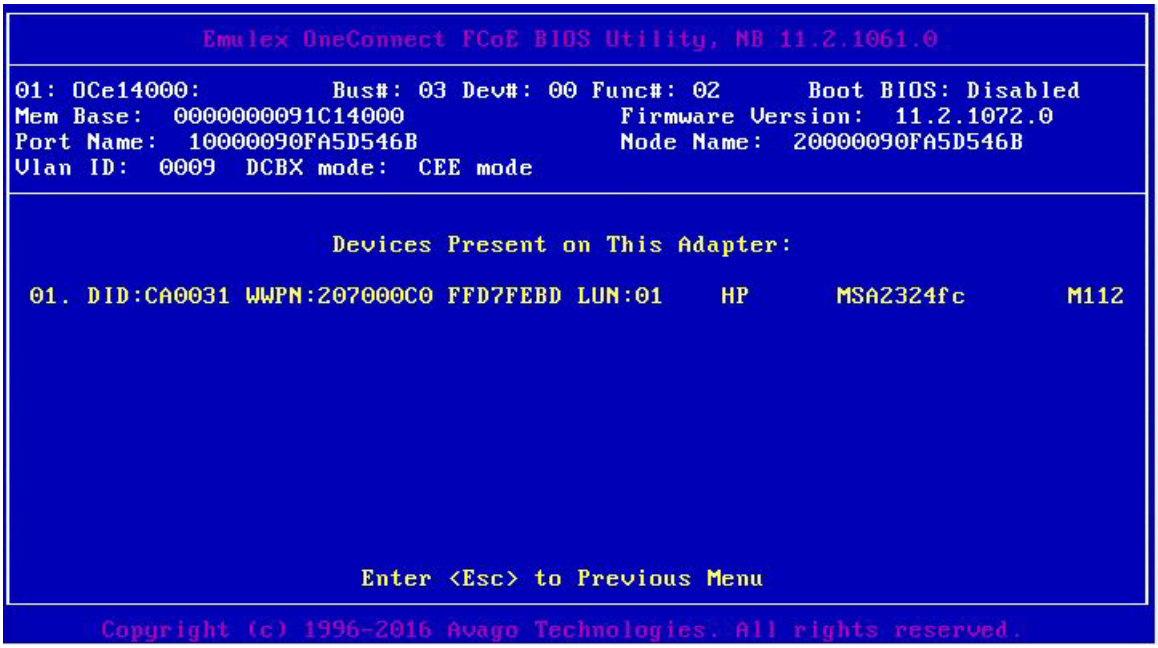


6.4 Scanning for Target Devices

To scan for target devices:

1. From the main configuration menu (Figure 29), select Scan for Target Devices. This option displays a list of discovered target devices (Figure 31). It allows you to quickly verify zoning and SAN connectivity.

Figure 31 Devices Present on the Adapter



2. Press Esc to return to the main configuration menu.

6.5 Configuring Boot Devices

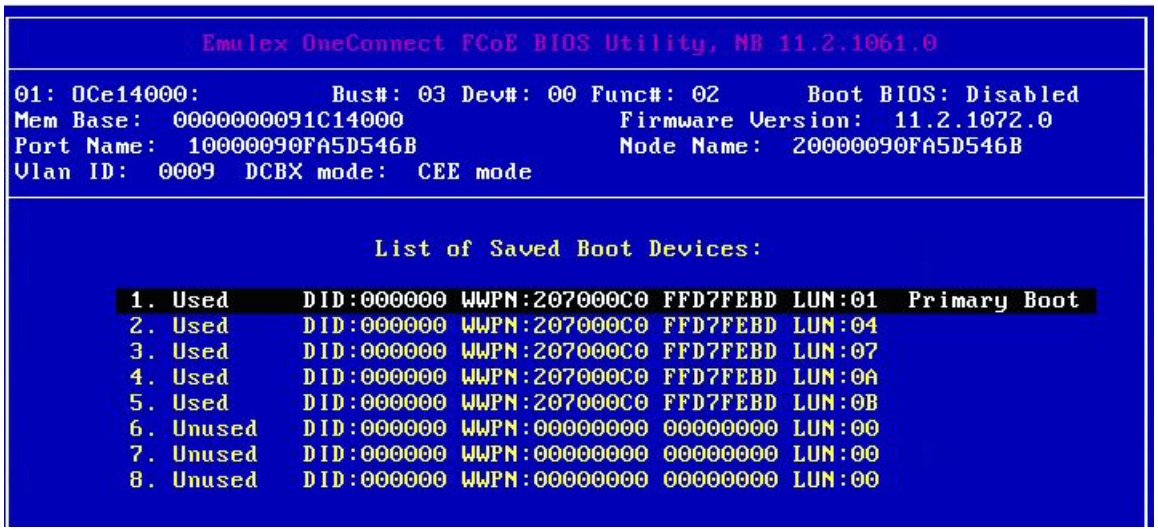
To configure boot devices:

1. On the main configuration menu (Figure 29), select Configure Boot Devices and press Enter.

   A list of eight boot devices appear (Figure 32). The primary boot device is the first entry shown. It is the first bootable device. If the first boot entry fails, the system attempts to boot from the second bootable entry. If the second boot entry fails, the system attempts to boot from the third bootable entry and so on.

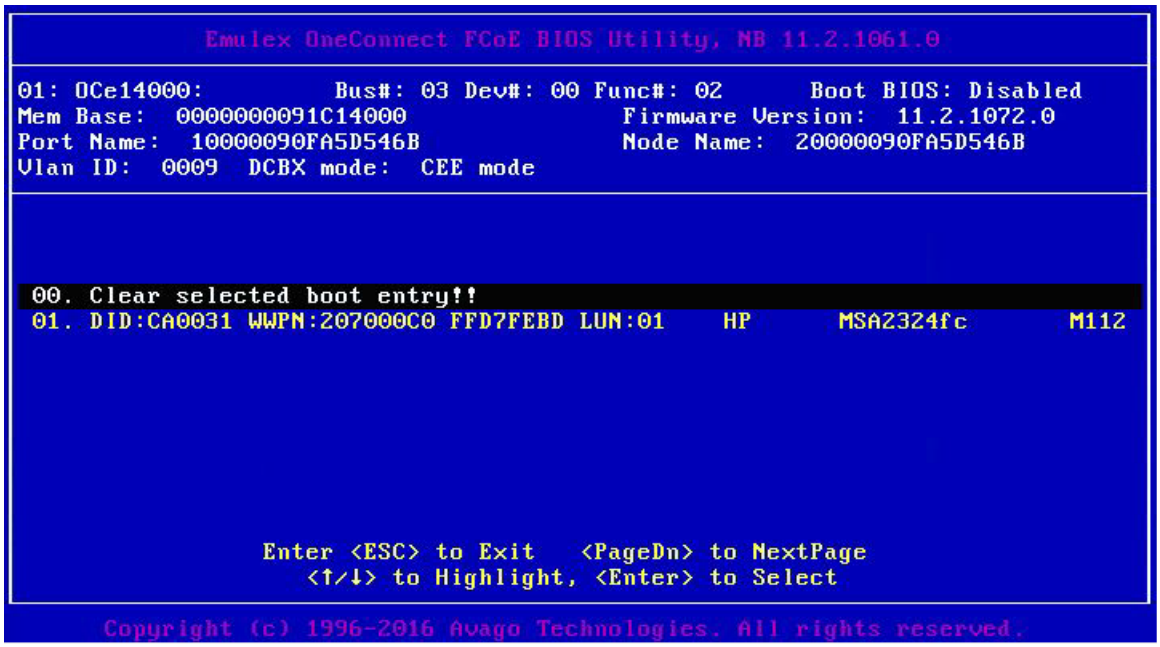
   NOTE This assumes the system BIOS supports BBS and the boot order in the system BIOS is set using the mapped drives from the list of saved boot devices.

Figure 32 List of Saved Boot Devices



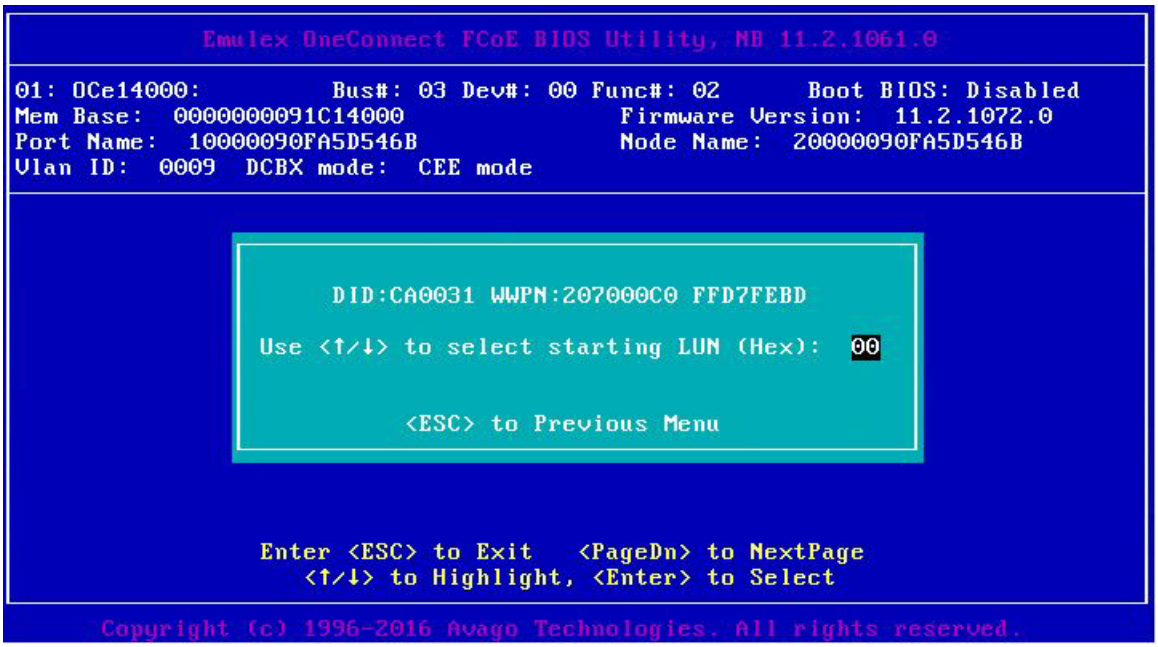
2. Select a boot entry and press Enter. A list similar to the following one appears.

Figure 33 Device Selection List Example



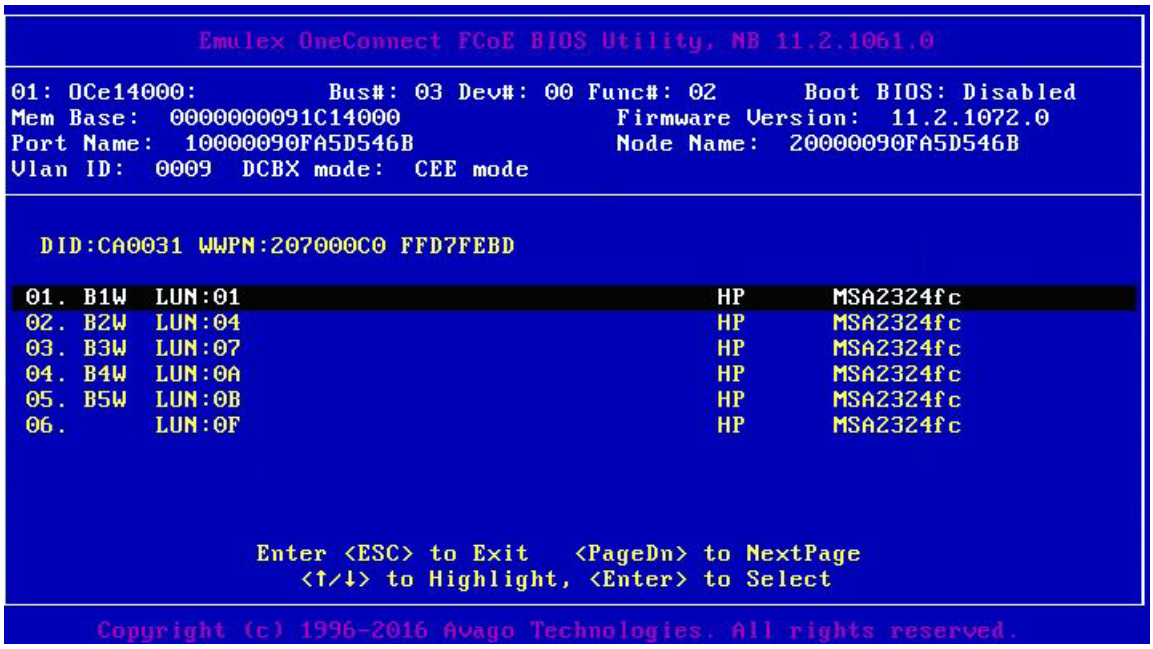
3. Select `00` and press Enter to clear the selected boot entry, or select the WWPN or DID of a device to configure and press Enter.
4. If you select a device, you are asked for the starting LUN (Figure 34). Use the up and down arrows to enter the starting LUN in hexadecimal format and press Enter. The starting LUN can be any number from 0 to 255.

Figure 34 LUN Listing Dialog



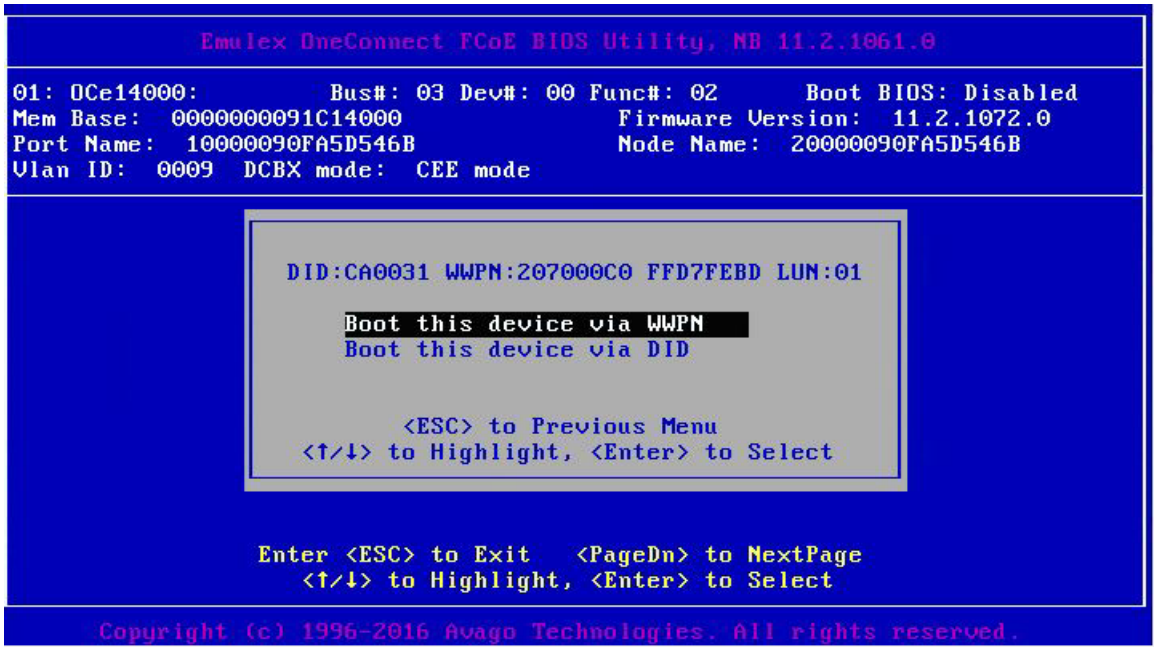
5. A dialog similar to the following one appears. Press Enter.

Figure 35 Selected Boot Device Dialog



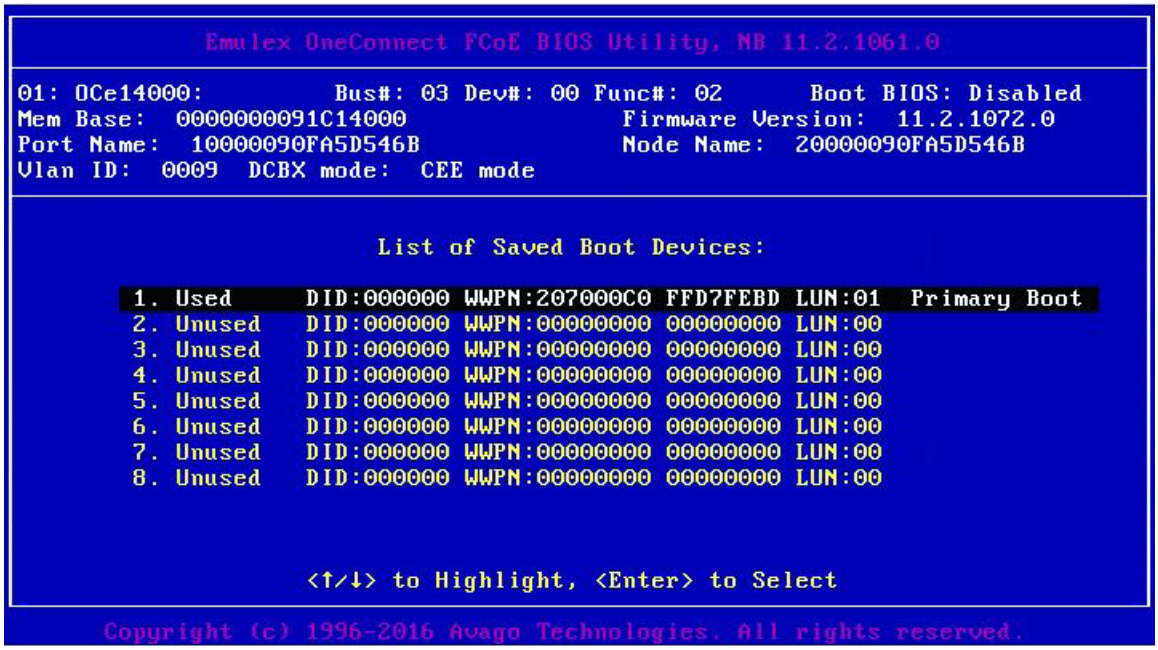
The following Boot Devices dialog (Figure 36) appears.

Figure 36 Boot Devices Dialog



6. Use the up and down arrows to select the boot method and press Enter. In the list that appears you can select to scan and boot from the device by its unique WWPN or by its FC DID.

Figure 37 Primary Boot Device Setup



7. Press Enter to select the change.
8. Press Esc to return to the main configuration menu.
9. Reboot the system for the new boot path to take effect.

6.6 Configuring FCoE FCF CEE Parameters

NOTE For DCBX configuration, the adapter is always in CEE mode.

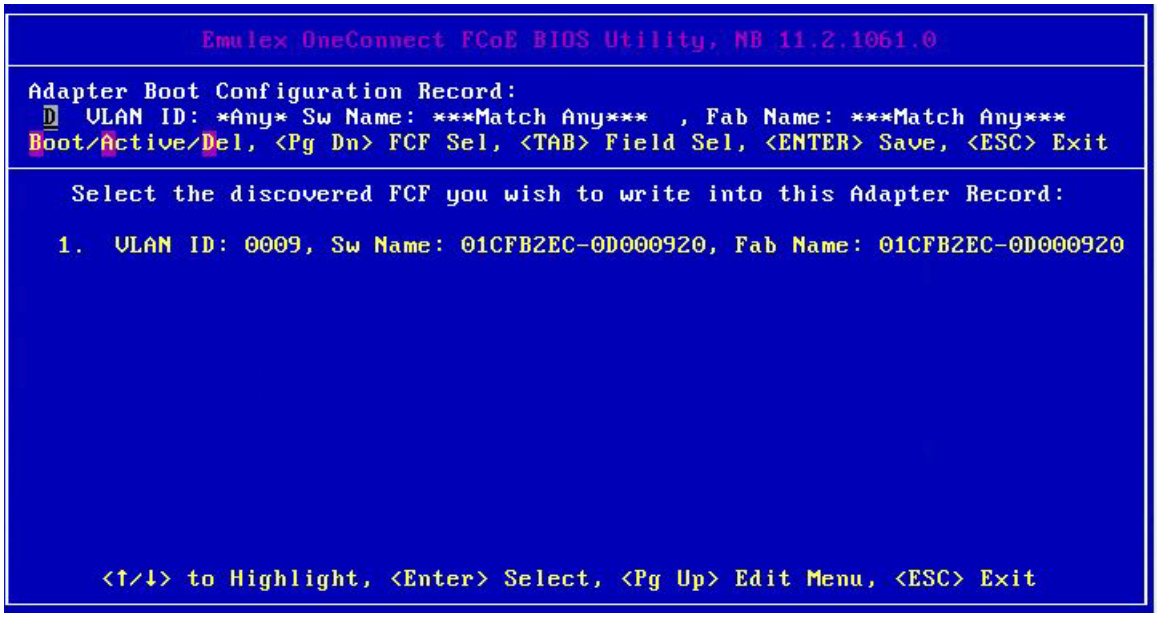
An FCF is the bridging layer to connect FCoE to an FC network. The FCF is usually built into the FCoE switch.

The FCF contains routing information including the switch name, fabric name, and VLAN information used to send the FCoE data to the correct FC fabric.

To configure CEE parameters:

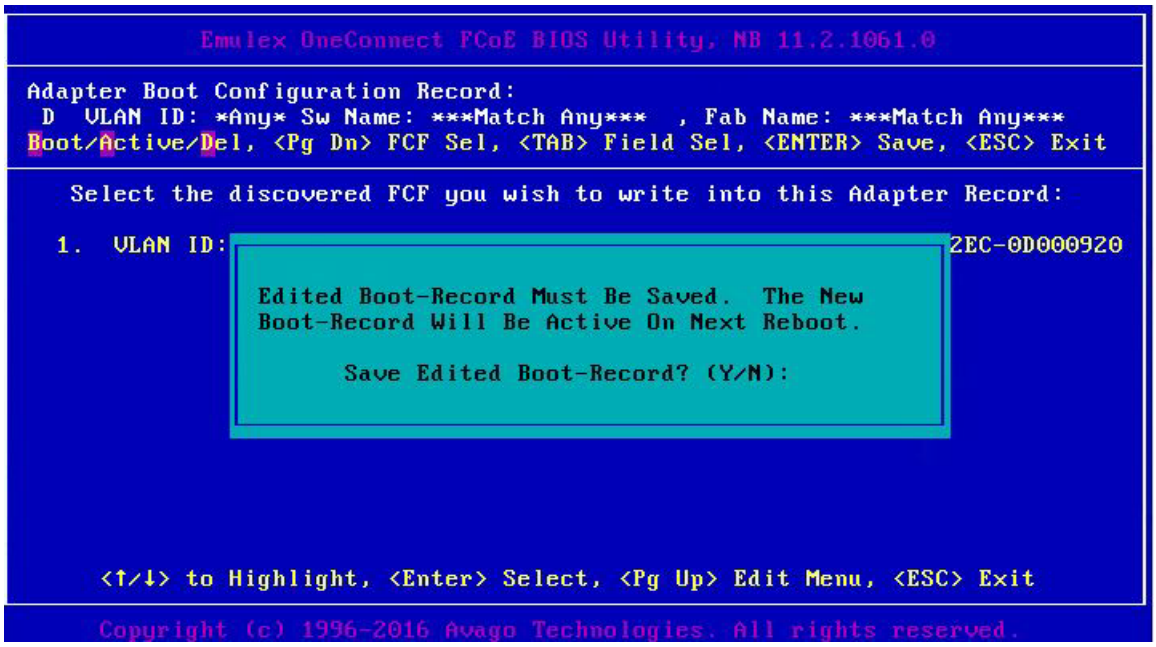
1. If multiple VLANs or FCoE switches are set up and you want to select a particular VLAN to boot from every time, select Configure FCF CEE Parameters from the main configuration menu (Figure 29) and press Enter. A list of discovered FCFs appears.

Figure 38 FCF Listing



2. Select the FCF you want to boot from every time and save it to the adapter record by pressing Enter. The system displays the following message (Figure 39).

Figure 39 Save Edited Boot Record Dialog Box



3. Press Y to save the boot record. The main configuration menu (Figure 29) appears.

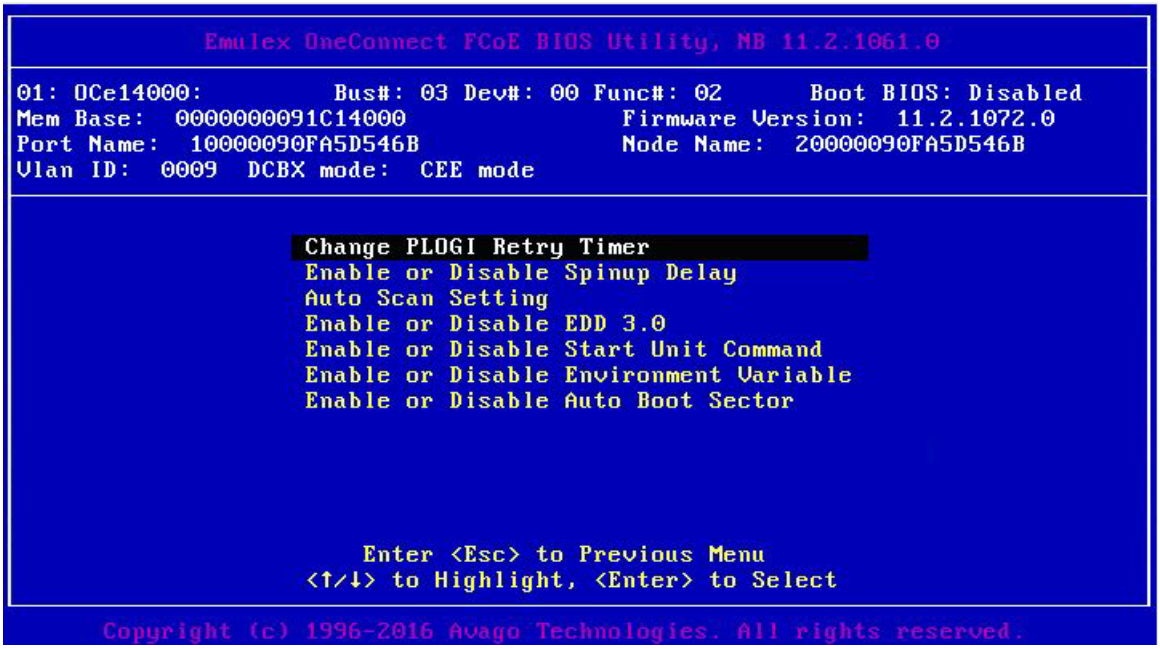
6.7 Configuring Advanced Adapter Parameters

The BIOS utility provides various options that can be modified to provide different behavior. Use the BIOS utility to perform the following tasks:

- Change the port login retry timer.
- Enable or disable spinup delay.
- Set autoscan.
- Enable or disable EDD 3.0.
- Enable or disable the start unit command.
- Enable or disable the environment variable.
- Enable or disable the auto boot sector.

To access the adapter configuration menu, from the main configuration menu (Figure 29), select Configure Advanced Adapter Parameters and press Enter. The following adapter configuration menu (Figure 40) appears.

Figure 40 Advanced Adapter Configuration Menu



The default settings are acceptable for most installations.

To reset all values to their defaults, from the main configuration menu (Figure 29), select Reset Adapter Defaults and press Enter.

6.7.1 Changing the Port Login Retry Timer

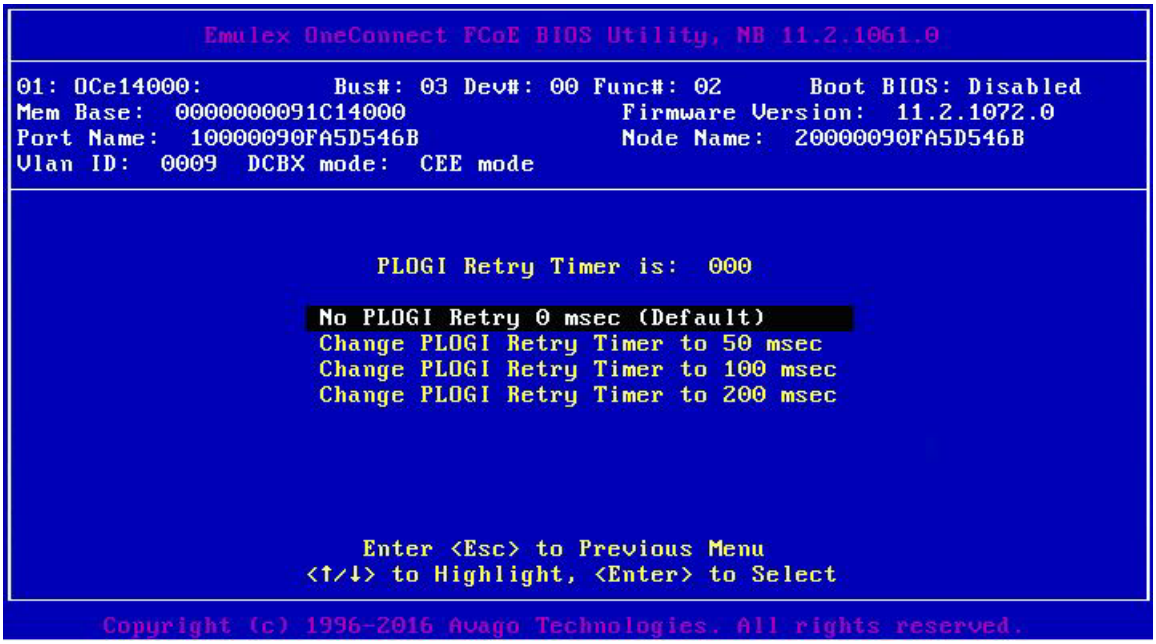
This option is especially useful for Tachyon-based RAID arrays. Rarely, a Tachyon-based RAID array resets itself and the port goes offline temporarily. When the port returns to operation, the port login scans the loop to discover the device. The port login retry interval is the time it takes for one port login to scan the whole loop. You can choose:

- No Port Login Retry: 0 msec – Default
- 50 msec takes 5 to 6 seconds per device
- 100 msec takes 12 seconds per device
- 200 msec takes 22 seconds per device

To set the interval for the port login retry timer:

1. On the main configuration menu (Figure 29), select Configure Advanced Adapter Parameters and press Enter. The adapter configuration menu (Figure 40) appears.
2. Select Change PLOGI Retry Timer and press Enter. Information appears similar to the following menu.

Figure 41 Change the Port Login Retry Timer Menu



3. Select the retry timer interval.
4. Press Enter to accept the new interval.
5. Press Esc to return to the advanced adapter configuration menu.
6. For changes to take effect, reboot the system.

6.7.2 Enabling or Disabling the Spinup Delay

This option lets you enable or disable the disk spinup delay. The factory default setting is disabled.

If at least one boot device has been defined, and the spinup delay is enabled, the boot BIOS searches for the first available boot device.

- If a boot device is present, the boot BIOS boots from it immediately.
- If a boot device is not ready, the boot BIOS waits for the spinup delay and, for up to three additional minutes, continues the boot scanning algorithm to find another multiboot device.

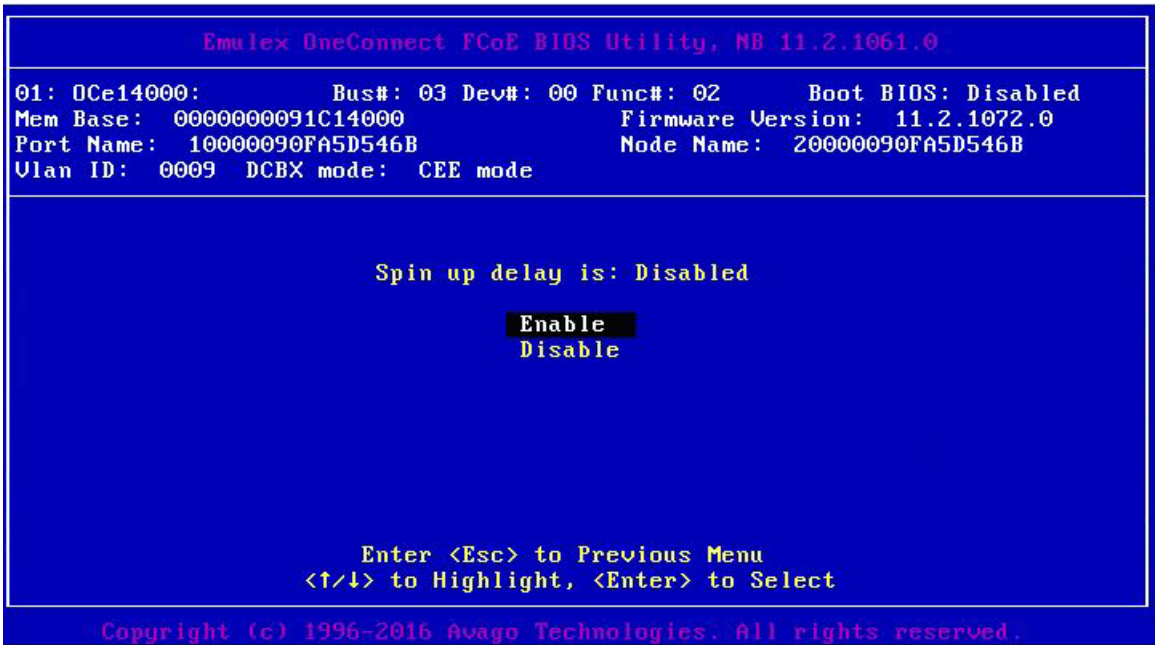
If boot devices have not been defined and auto scan is enabled, the boot BIOS waits for five minutes before scanning for devices.

- In an attached fabric, the boot BIOS attempts to boot from the first target found in the NameServer data.

To enable or disable the spinup delay:

1. On the main configuration menu (Figure 29), select Configure Advanced Adapter Parameters and press Enter. The adapter configuration menu (Figure 40) appears.
2. Select **Enable or Disable Spinup Delay** and press Enter. Information appears similar to the following dialog.

Figure 42 Enable or Disable Spinup Delay Dialog



3. Select whether to enable or disable spinup delay.
4. Press Enter to accept the new value.
5. Press Esc to return to the advanced adapter configuration menu.
6. For changes to take effect, reboot the system.

6.7.3 Setting Auto Scan

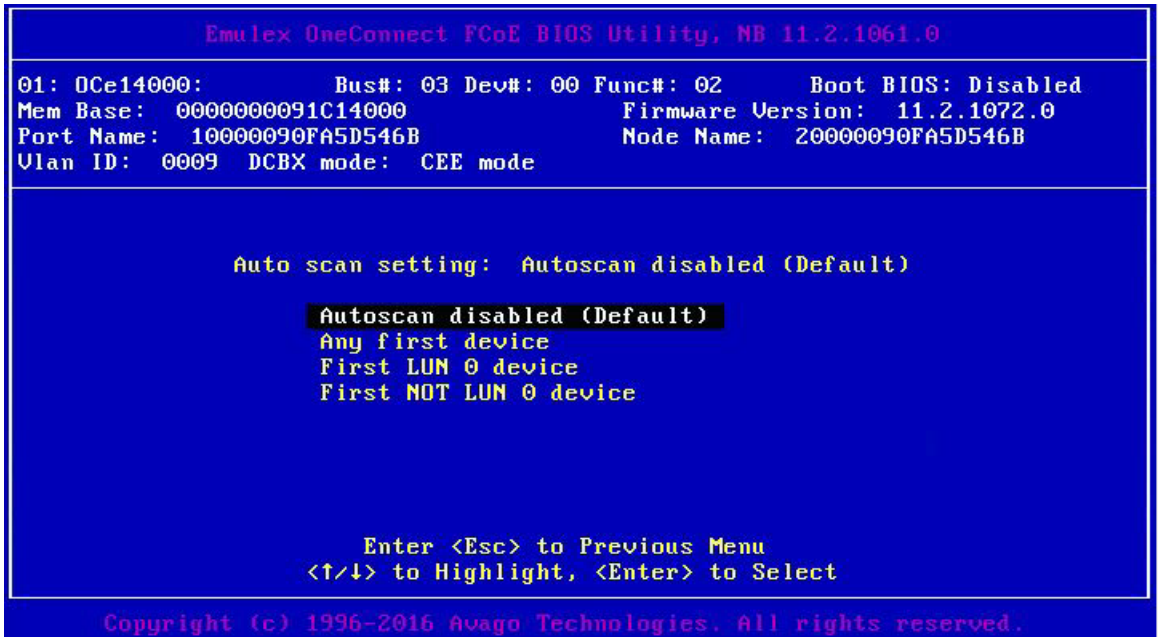
This option allows you to enable or disable the auto scan feature. Auto scan is available only when all eight boot entries are unused (empty). The factory default for auto scan is disabled. If there is more than one adapter with the same PCI Bus number in the system, and each has a boot drive attached, the first PCI-scanned adapter is the boot adapter. Auto scan allows you to boot from the first boot device found in the FC fabric with the following options:

- **Autoscan disabled** – Default
- **Any first device** – The first adapter scans the fabric and the first disk discovered becomes the attempted boot device
- **First LUN 0 device** – The first device discovered with LUN 0
- **First NOT LUN 0 device** – The first device discovered with a LUN other than 0

To set auto scan:

1. From the main configuration menu (Figure 29), select Configure Advanced Adapter Parameters and press Enter. The adapter configuration menu (Figure 40) appears.
2. Select Auto Scan Setting and press Enter. The system displays the following Set Auto Scan menu.

Figure 43 Set Auto Scan Menu



3. Select the appropriate auto scan option.
4. Press Enter to accept the new scan option.
5. Press Esc to return to the advanced adapter configuration menu.
6. For changes to take effect, reboot the system.

6.7.4 Enabling or Disabling EDD 3.0

EDD 3.0 provides additional data to the operating system boot loader during an INT-13h function 48h call (get device parameters). This information includes the path to the boot device and the disk size. The default setting for EDD 3.0 is enabled. If EDD 3.0 is disabled, the setting reverts to EDD 2.1.

To enable or disable EDD 3.0:

1. From the main configuration menu (Figure 29), select Configure Advanced Adapter Parameters and press Enter. The adapter configuration menu (Figure 40) appears.
2. Select **Enable or Disable EDD 3.0** and press Enter. The EDD 3.0 configuration dialog (Figure 44) appears.

Figure 44 Enable or Disable EDD 3.0 Dialog



3. Select whether to enable or disable EDD 3.0.
4. Press Enter to accept the change.
5. Press Esc to return to the advanced adapter configuration menu.
6. For changes to take effect, reboot the system.

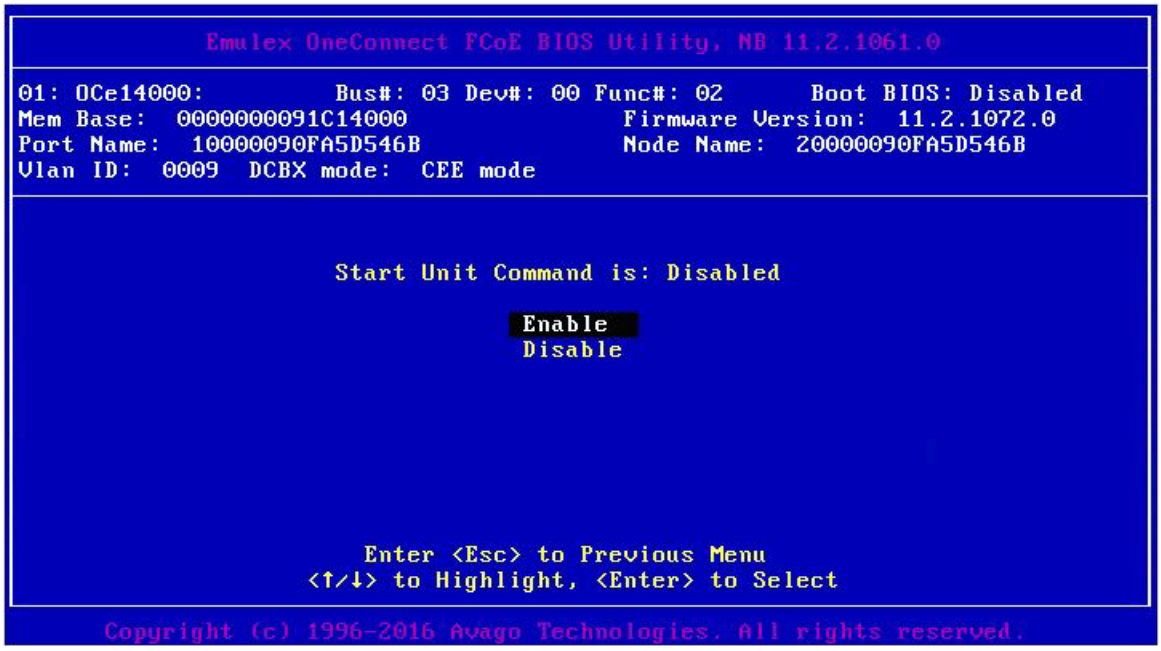
6.7.5 Enabling or Disabling the Start Unit Command

You must know the specific LUN to issue the SCSI start unit command. The default setting is disabled.

To enable or disable the start unit command:

1. From the main configuration menu (Figure 29), select Configure Advanced Adapter Parameters and press Enter. The system displays the adapter configuration menu (Figure 40).
2. Select **Enable or Disable Start Unit Command** and press Enter. The Start Unit Command configuration dialog (Figure 45) appears.

Figure 45 Enable or Disable Start Unit Command Dialog



3. Select whether to enable or disable the start unit command.
4. Press Enter to select the change.
5. Press Esc to return to the advanced adapter configuration menu.
6. For changes to take effect, reboot the system.

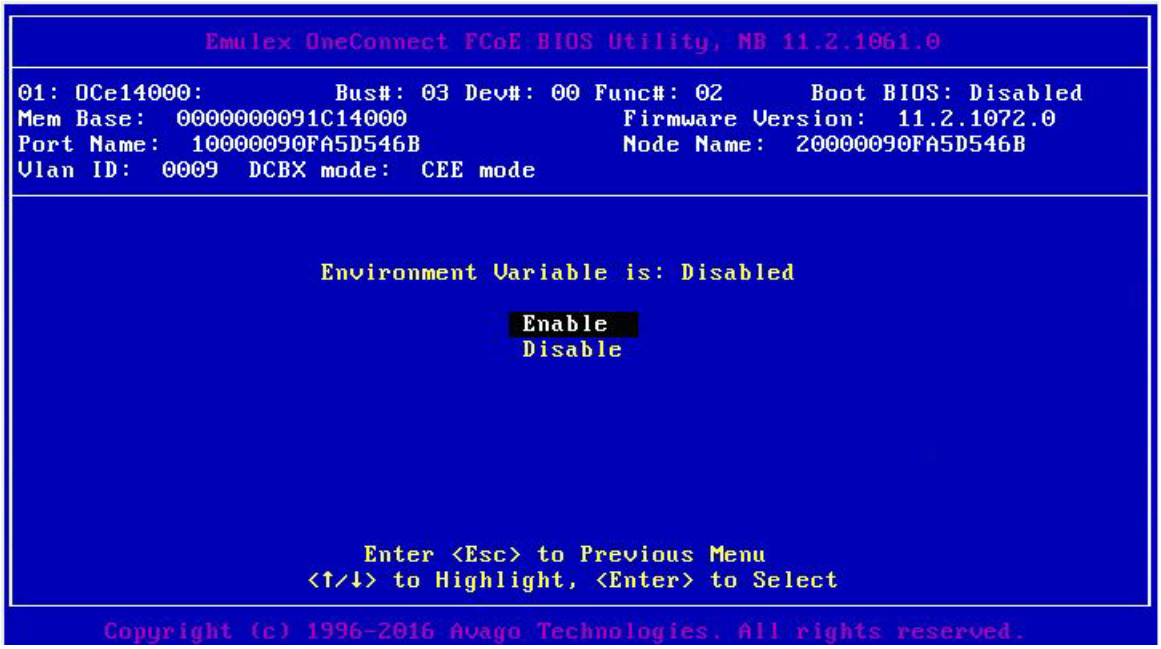
6.7.6 Enabling or Disabling the Environment Variable

Sets the boot adapter order if the system supports the environment variable. The default setting is disabled.

To enable or disable the environment variable:

1. From the main configuration menu (Figure 29), select Configure Advanced Adapter Parameters and press Enter. The adapter configuration menu (Figure 40) appears.
2. Select **Enable or Disable Environment** Variable and press Enter. The Environment Variable configuration dialog (Figure 46) appears.

Figure 46 Enable or Disable Environment Variable Dialog



3. Select whether to enable or disable the environment variable.
4. Press Enter to select the change.
5. Press Esc to return to the advanced adapter configuration menu.
6. For changes to take effect, reboot the system.

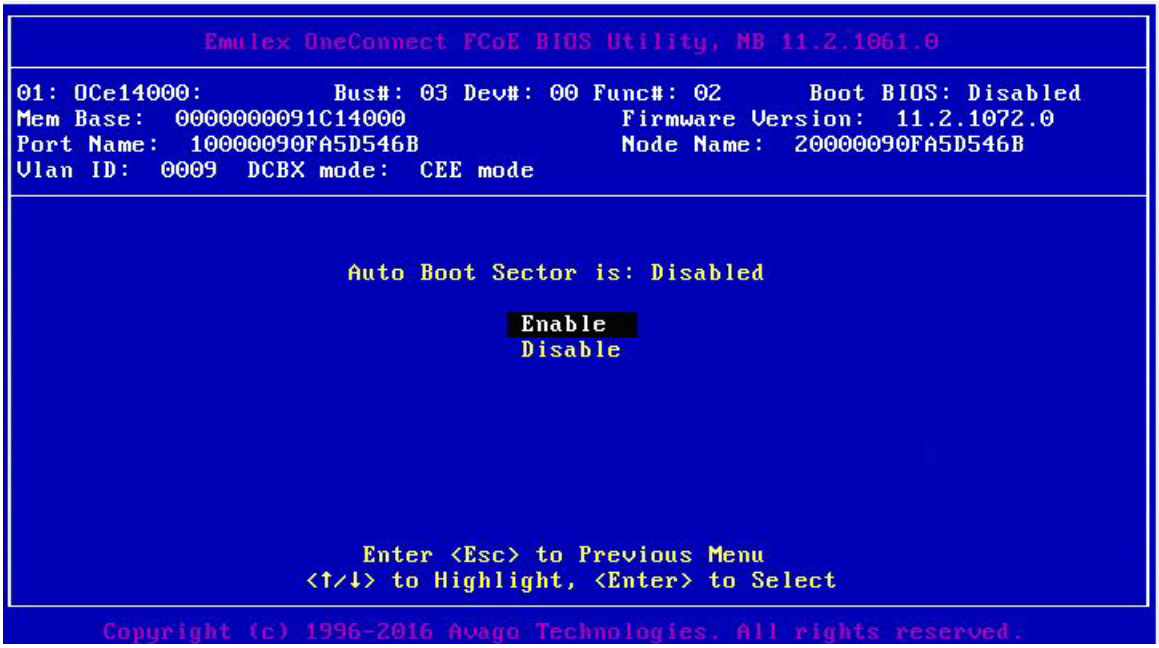
6.7.7 Enabling or Disabling Auto Boot Sector

This option automatically defines the boot sector of the target disk for the migration boot process, which applies to HP MSA1000 arrays only. If there is no partition on the target, the default boot sector format is 63 sectors. The default setting is disabled.

To enable or disable the auto boot sector format:

1. From the main configuration menu (Figure 29), select Configure Advanced Adapter Parameters and press Enter. The adapter configuration menu (Figure 40) appears.
2. Select **Enable or Disable Auto Boot Sector** and press Enter. The Auto Boot Sector configuration dialog (Figure 47) appears.

Figure 47 Enable or Disable Auto Boot Sector Format Dialog



3. Select whether to enable or disable the auto boot sector.
4. Press Enter to accept the new value.
5. Press Esc to return to the advanced adapter configuration menu.
6. For changes to take effect, reboot the system.

6.8 Using Multipath Boot from SAN

Multiboot BIOS is in compliance with BBS. The system must have a multiboot system BIOS to take advantage of this option. The multiboot BIOS allows you to select any boot disk in the system BIOS setup menu. The boot disk can be an FC drive, a SCSI drive, an IDE drive, a USB device, or a floppy drive. The Emulex BIOS supplies the first eight drives to the system BIOS menu. The multiboot BIOS can override the FC drive selected in the BIOS utility.

For example, the system has eight FC disks only. The boot disk has DID 0203E4. However, you can select DID 1A14B5 in the system BIOS setup menu. The boot device is the FC disk with DID 1A14B5 instead of DID 0203E4, as is set in the BIOS utility.

If the system supports multiboot BBS, the local boot disk (drive C) is the first entry in multiboot on the system BIOS setup menu. The list of entries is determined by the list of configured boot entries in the BIOS utility. For example:

```
Adapter 1: boot_entry0, boot_entry1
Adapter 2: boot_entry2, boot_entry3
```

The order of boot entries exported to multiboot (BBS) is

```
boot_entry0, boot_entry1, boot_entry2, and boot_entry3
```

Multiboot allows changing the boot order in the system BIOS, which allows any disk to become the primary boot (C) drive.

6.9 Resetting to Default Values

The BIOS utility enables you to reset BIOS boot parameters to their factory default settings. These defaults are listed in the following table.

NOTE The FCoE CEE parameters (FCF information) are not affected. Only boot adapter specific information is cleared.

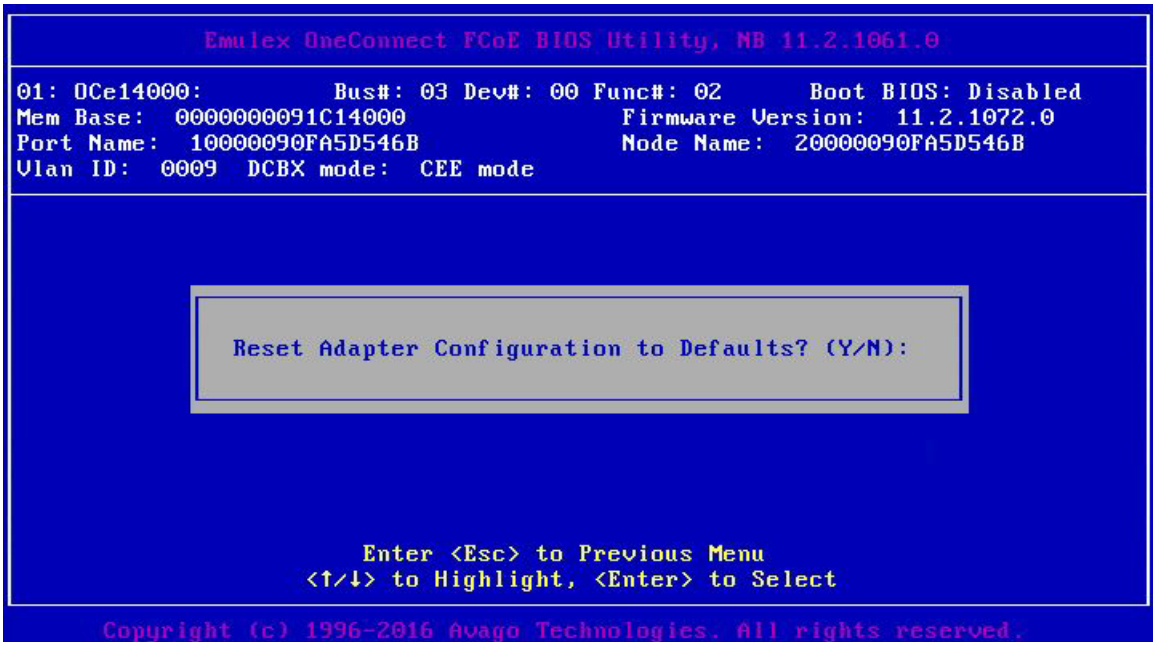
Table 7 Default Adapter Boot Parameter Values

Parameter	Default Value	Valid Values
Boot from SAN	Disabled	Enabled Disabled
EDD 3.0	Enabled (EDD 3.0)	Enabled (EDD 3.0) Disabled (EDD 2.1)
Port Login Retry Timer	Disabled	Disabled 50 msec 100 msec 200 msec
Spinup Delay	Disabled	Enabled Disabled
Auto Scan	Disabled	Enabled Disabled
Start Unit	Disabled	Enabled Disabled
Environmental Variable	Disabled	Enabled Disabled
Auto Boot Sector	Disabled	Enabled Disabled

To reset parameters to their factory default settings:

1. On the main configuration menu (Figure 29), select Reset Adapter Defaults and press Enter. The following message pop up (Figure 48) appears.

Figure 48 Reset Adapter Configuration to Defaults Pop Up



2. Press Y. All settings revert to their factory default values.
3. Press Esc to return to the previous menu.

Chapter 7: Updating and Enabling Boot Code

Broadcom provides utilities to update and enable boot code. The utility you use depends on the operating system and, in some cases, the driver type or system architecture. The following table indicates the utilities currently available for each operating system.

Table 8 Utilities that Update and Enable Boot Code

Utilities	Windows	Linux[1]	VMware[2]	Solaris
OneCommand CNA Manager GUI	X	X	X	X
OneCommand CNA Manager CLI	X	X	X	X
Offline Utilities	X	X		
UEFI Utilities	X	X	X	
OneCommand CNA Manager for VMware vCenter			X	
OneConnect Flash Utility[3]	X	X	X	X
CIM Provider			X	
ethtool		X		
FCA brcmadm Utility				X

1. For the Citrix operating systems, the OneCommand CNA Manager application CLI is not supported locally, but can be managed using a remote system. The OneCommand CNA Manager GUI application is not supported.
2. For VMware operating systems, the OneCommand CNA Manager application GUI and CLI are not supported locally, but can be run remotely using a Windows system.
3. Not supported on IA64–based servers.

After you decide which utility to use, see the appropriate documentation for the proper procedures:

- **OneCommand CNA Manager application** – Refer to the *OneCommand CNA Manager Application User Guide for OneConnect Adapters* or the *OneCommand CNA Manager Command Line Interface User Guide for OneConnect Adapters*.
- **Offline utilities** – Refer to the *Brcmflash and Lpcfg Management Utilities User Guide for OneConnect Adapters*.
- **UEFI utilities**
  - See Chapter 11, Configuring UEFI for Ethernet.
  - See Chapter 12, Configuring UEFI for iSCSI.
  - See Chapter 13, Configuring UEFI for FCoE.
- **OneCommand CNA Manager for VMware vCenter** – Refer to the *OneCommand CNA Manager for VMware vCenter User Guide for OneConnect Adapters.*
- **OneConnect Flash Utility** – Refer to the *Using the OneConnect Flash Utility to Update OneConnect Adapter Firmware* guide.
- **CIM Provider** – Refer to the *CIM Provider Package Installation Guide for OneConnect Adapters*.
- **ethtool** – Refer to the *Emulex Drivers for Linux User Guide for OneConnect Adapters*.
- **FCA emlxadm utility** – Refer to the *FCA Utilities for Solaris User Guide for OneConnect Adapters*.

Chapter 8: Configuring iSCSI Boot Support with the iSCSISelect Utility

This section provides quick installation instructions for setting up a basic boot configuration. For more information on setting up an iSCSI initiator and adding an iSCSI target, see Chapter 9, Configuring and Managing the iSCSI Initiator with the iSCSISelect Utility and Chapter 10, Configuring and Managing iSCSI Targets with the iSCSISelect Utility.

8.1 Navigating the iSCSISelect Utility

Use the following methods to navigate the iSCSISelect utility:

- Press the up and down arrows on the keyboard to move through and select menu options or configuration fields. If multiple adapters are listed, use the up and down arrows to scroll to the additional adapters.
- Press the Tab key to move to the next field, and press shiftbar+Tab key to move to the previous field.
- Press Enter to accept a selection, select a menu option, to select a row in a configuration dialog, or to change a configuration default.
- Press Esc to return to the previous menu or page, cancel a selection or dialog box, or exit the utility.

8.2 Setting Up a Basic iSCSI Boot Configuration

By setting up a basic iSCSI boot configuration, you can configure the initiator, contact network devices, and set up an iSCSI boot target. With iSCSI boot capability, the iSCSISelect utility can quickly and easily provide the target LUN disk to the operating system through the BIOS. After setting up a basic iSCSI boot configuration, you can continue to use the iSCSISelect utility to manage the Emulex OneConnect server iSCSI SAN environment.

When setting up a basic iSCSI boot configuration do the following:

- Configure the iSCSI initiator name and enable boot support
- Configure the network properties
- Add an iSCSI target and enable it as a boot target
- Set the boot LUN
- Reboot the system

To set up a basic iSCSI boot configuration:

1. Log in to the iSCSISelect utility by pressing the Ctrl-S keys when prompted.
2. From the following iSCSI Initiator Configuration dialog, set the initiator name.

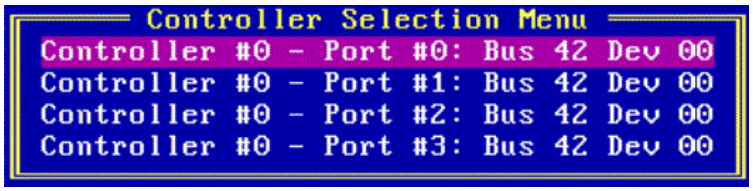
Figure 49 iSCSI Initiator Configuration Dialog



3. Select Save and press Enter to save the initiator name.
4. Select Controller Configuration and press Enter.
   — If you are running a single adapter, the Controller Configuration menu (Figure 51) appears.

— If you are running multiple adapters, the following Controller Selection menu appears.

Figure 50 Controller Selection Menu



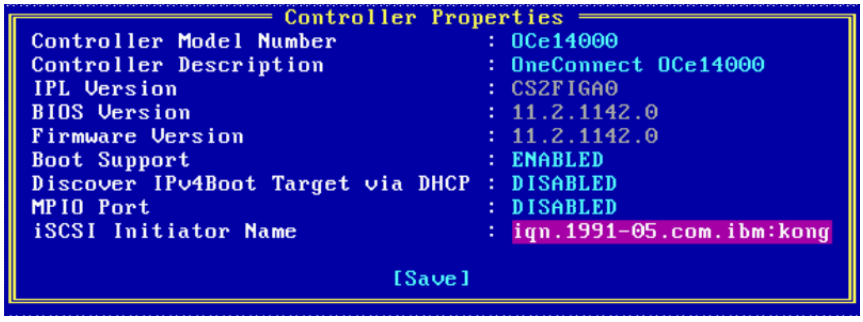
5. From the Controller Selection menu (Figure 50), select the adapter and port then press Enter. For more information, see Section 9.1, Selecting an Adapter.
6. After selecting the adapter, the following Controller Configuration menu appears.

Figure 51 Controller Configuration Menu



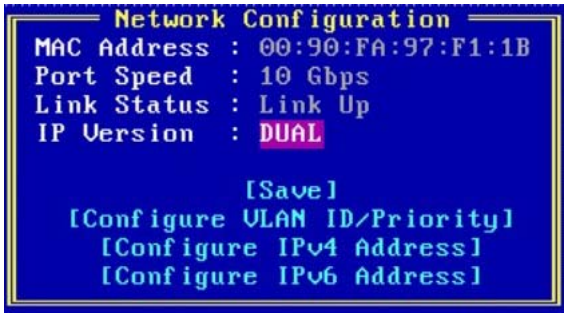
7. From this menu, select Controller Properties and press Enter. The following Controller Properties dialog (Figure 52) appears. For more information on the Controller Properties dialog, see Section 9.2, Viewing the Controller Properties.

Figure 52 Controller Properties Dialog



8. Verify that Boot Support is set to Enabled.
9. Select Save and press Enter.
10. From the Controller Configuration menu (Figure 51), select Network Configuration and press Enter. The following Network Configuration dialog (Figure 53) appears.

Figure 53 Network Configuration Dialog



11. Configure a network IP address. From the Network Configuration dialog, you can do the following:
    - Select the IP version (IPv4, IPv6, or Dual). For more information, see Section 9.3.1, Selecting the IP Version.
    - Set up an IP address. For more information, see Section 9.3.3, Configuring an IP Address.
12. If desired, configure a VLAN ID. For more information, see Section 9.3.2, Configuring VLAN ID and VLAN Priority.
13. After configuring the network, you can add a SCSI target. From the Controller Configuration menu (Figure 51), select iSCSI Target Configuration, then select Add a New iSCSI Target. For more information, see Section 10.1, Adding iSCSI Targets.
14. After configuring the iSCSI initiator and target, reboot the system.

8.3 Booting the System

After setting up the basic boot configuration, exit the iSCSISelect utility and the system automatically reboots. During the next boot, the iSCSI BIOS logs into the configured iSCSI boot target and displays its target and LUN information in the BIOS sign-on banner. For example:

```
<<< Press <Ctrl><S> for iSCSISelect(TM) Utility >>>

iSCSI#0 Port#1 Base 0x00000380002A0000 @ Bus:1B Dev:00 Fun:02
iSCSI#0 Port#2 Base 0x00000380002E0000 @ Bus:1B Dev:00 Fun:03

Initiator iSCSI Name:  iqn.2014-05.com.emulex:iscsi-dell-init
Controller #0:0 Initiator IPV6 Address: FE80:0:0:0:290:FAFF:FE6B:BA2D
Controller #0:0 Initiator IPV4 Address: 0.0.0.0
Controller #0:1 Initiator IPV6 Address: FE80:0:0:0:290:FAFF:FE6B:BA35
Controller #0:1 Initiator IPV4 Address: 192.168.10.56

Drive #0 MSFT   VirtualHD   0   30720 MB
Drive Geometry    3FF    3F    FF
BIOS Installed Successfully!
```

Chapter 9: Configuring and Managing the iSCSI Initiator with the iSCSISelect Utility

The iSCSI initiator or host (also known as the iSCSI client) is a system that attaches to an IP network and initiates requests and receives responses from an iSCSI target. Each iSCSI host is identified by a unique IQN.

After the initiator host system is running, you must configure the initiator to allow access to the iSCSI SAN. To do this, you must:

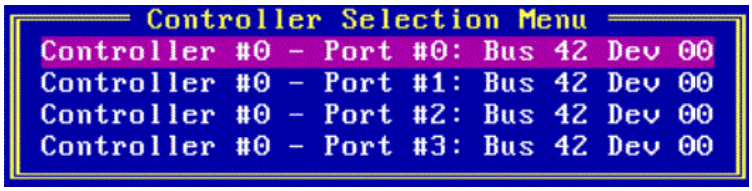
1. Select an adapter. For more information, see Section 9.1, Selecting an Adapter.
2. Configure network properties. For more information, see Section 9.3, Configuring Network Properties.

9.1 Selecting an Adapter

To select an adapter:

1. Log in to the iSCSISelect utility by pressing the Ctrl-S keys when prompted. The system displays the Controller Selection Menu.

Figure 54 Controller Selection Menu



2. Select an adapter to configure the iSCSI target. The following Controller Configuration menu (Figure 55) appears. From this menu, you can configure and manage the iSCSI initiator and targets.

Figure 55 Controller Configuration Menu



To verify that you are configuring the correct port, either:

- Perform port identification; for more information, see Section 9.4, Identifying a Port.
- Check the link status by connecting the port to a 10Gb/s switch port and checking the link status in the Network Configuration dialog (Figure 57); if the status is Link Up, it is the correct port.

NOTE Verify that you back out of the Network Configuration dialog to the Controller Configuration menu before checking the Network Configuration dialog. This allows the Link Status field to refresh when you insert or remove the cable from the adapter.

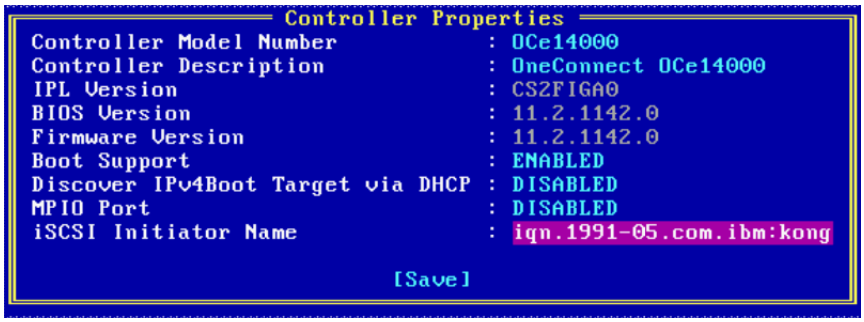
9.2 Viewing the Controller Properties

The Controller Properties dialog displays adapter information and allows you to configure adapter options such as boot support, IPv4 boot target discovery using DHCP, MPIO port configuration, and the iSCSI initiator name.

To view the adapter properties:

1. From the Controller Configuration menu (Figure 55), select Controller Properties and press Enter. The Controller Properties dialog appears.

Figure 56 Controller Properties Dialog



9.2.1 Enabling Boot Support

Boot support is provided for each port or function. If boot support is enabled, you can boot from the specified function. If boot support is disabled, you will not be able to boot from the function.

To enable boot support:

1. From the Controller Properties dialog (Figure 56), select Boot Support and press Enter. A pop up appears.
2. Select **Enabled** or **Disabled** and press Enter.
3. To save your changes, select Save and press Enter.

9.2.2 Enabling IPv4 Boot Target Discovery using DHCP

To enable IPv4 boot target discovery using DHCP:

1. From the Controller Properties dialog (Figure 56), select Discover IPv4 Boot Target via DHCP and press Enter. A pop up appears.
2. Select **Enabled** or **Disabled** and press Enter.
3. To save your changes, select Save and press Enter.

9.2.3 Enabling the MPIO Port

MPIO support allows the initiator to log in dual sessions to the same target. In this way I/O can be sent over either TCP/IP connection to the target. If one session fails, another session can continue processing I/O without interrupting the application. In iSCSI target configuration, you have the option of setting dual network paths to a single boot LUN.

NOTE Although MPIO boot support allows the initiator to log in to multiple sessions, the iSCSI BIOS limits the number of sessions to two iSCSI sessions for a single boot LUN.

Follow these steps to configure MPIO boot support for each operating system.

1. Use the Emulex iSCSI configuration utility to configure the first path to a boot target.
2. Complete normal operating system installation.
3. Install third-party MPIO software for your operating system.
4. Configure the second path to a single boot LUN through the Emulex iSCSI configuration utility.

To enable MPIO support:

1. From the Controller Properties dialog (Figure 56), select MPIO Port and press Enter. A pop up appears.
2. Select **Enabled** or **Disabled** and press Enter.
3. To save your changes, select Save and press Enter.

9.2.4 Configuring the iSCSI Initiator Name

The iSCSI initiator name is a unique identifier for the initiator on the network and configures target access. It must be unique to that system on the network and is identified by an IQN. The iSCSI initiator name is global to the entire system. The iSCSI initiator name must also match the target's ACL. After you log in to the iSCSISelect utility, you can configure the iSCSI initiator name from the iSCSI Initiator Configuration dialog.

NOTE When installing Microsoft software iSCSI initiator service, the iSCSI initiator name for the Emulex OneConnect adapter will change and any IQN name previously configured through the iSCSISelect utility will be overridden. Though this will not affect existing boot sessions and persistent sessions, new target logins can fail because the new IQN name may not match the incoming initiator IQN names configured on the target. To avoid this situation, after installing Microsoft software, you must rename the IQN name to the IQN name you had configured previously.

To configure the iSCSI initiator name:

1. From the Controller Properties dialog (Figure 56), select iSCSI Initiator Name and press Enter.
2. Set or change the iSCSI initiator name.

NOTE If there are multiple adapters in the system, the iSCSI initiator name may be different for each adapter. If this happens, an error message indicates an iSCSI initiator name mismatch. You must save a new initiator name so the iSCSI initiator names match on all of the adapters. If there were pre-existing sessions before the iSCSI initiator name change, the pre-existing sessions use the original name that was used to log in the first time. The new name is used for any sessions added after the name change.

For Emulex adapters, the iSCSI Initiator Name setting is global across all ports on the same adapter. Changing one port to a new initiator name causes the other ports to reflect the new initiator name.

Some vendor adapters allow you to assign each port a different initiator name, or all ports can have the same initiator name.

3. Select Save and press Enter to save the iSCSI initiator name.

9.3 Configuring Network Properties

The Network Configuration dialog displays the MAC address, port speed, and link status for the adapter. From the Network Configuration dialog, you can configure the following items:

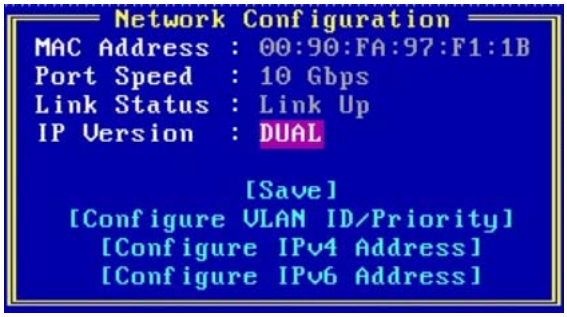
- IP version
- VLAN ID/priority
- IPv4 or IPv6 address

NOTE Certain menu items may not be available depending on the selected IP Version.

To view the network configuration:

From the Controller Configuration menu (Figure 55), select Network Configuration and press Enter. The Network Configuration dialog appears.

Figure 57 Network Configuration Dialog



9.3.1 Selecting the IP Version

To select the IP version:

1. From the Network Configuration dialog (Figure 57), select IP Version and press Enter. A pop up appears.
2. Select IPv4, IPv6, or DUAL and press Enter.
3. Select Save and press Enter to save the IP version.

9.3.2 Configuring VLAN ID and VLAN Priority

A VLAN is a way to partition the network. If the LAN is made up of all devices within a broadcast domain, a VLAN is a broadcast domain made up of switches. You first create a VLAN and then assign ports to it. All ports in a single VLAN are in a single broadcast domain.

You do not have to configure VLANs unless the network is already using them. Some reasons to use a VLAN include:

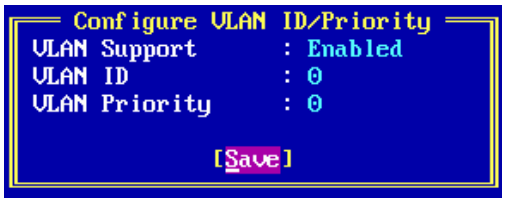
- A LAN increases in size with several devices.
- A LAN has increased broadcast traffic on it.
- Groups of users on a LAN need more security.

A VLAN ID, like an IP address or initiator name, is assigned a value to uniquely identify it on a network. A VLAN priority is set to determine priority order for packets within a VLAN.

To configure a VLAN ID and VLAN priority:

1. On the Network Configuration dialog (Figure 57), select Configure VLAN ID/Priority and press Enter. The Configure VLAN ID/Priority dialog appears.

Figure 58 Configure VLAN ID/Priority Dialog



2. Navigate to VLAN Support and select **Enabled**.
3. Set a VLAN ID number. This is a unique value you assign to each VLAN on a single device. Valid values range from 0 to 4094.
4. Set a VLAN priority, if necessary. This unique value assigns a priority to outbound packets containing a specified VLAN ID. Valid values range from 0 to 7, with 0 the highest priority level.
5. Select Save and press Enter.
6. Press Esc to return to the Network Configuration dialog.

9.3.3 Configuring an IP Address

An IP address is a logical address for the adapter and uniquely identifies the system on a TCP/IP network. Depending on the IP version you selected, you can either configure an IPv4 address or an IPv6 address.

9.3.3.1 Configuring an IPv4 Address

The IP address can be determined statically (manually assigned) or dynamically (the DHCP server obtains an IP address). The method you choose depends on your network configuration. If your network uses a DHCP configuration, you can enable DHCP and obtain an IP address from the DHCP server. If a DHCP server is not available, you must configure a static IP address.

NOTE

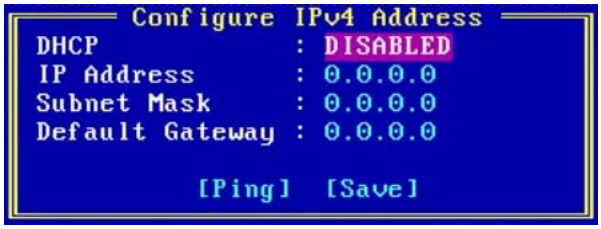
For all Linux-based operating systems, enable DHCP when attempting to assign an IP address to an iSCSI port, because the iSCSI firmware will not detect it as a duplicate IP address if one already exists on a Linux server. Otherwise, you must manually verify that the IP address you assign is not already used.

If you are using target discovery through DHCP, you must add the root path to the DHCP server and enable DHCP discovery through the iSCSISelect utility.

To configure an IPv4 address:

1. On the Network Configuration dialog (Figure 57), verify that the IP Version is set to IPv4.
2. Select Configure IPv4 Address and press Enter. The Configure IPv4 Address dialog appears.

Figure 59 Configure IPv4 Address Dialog



3. From the Configure IPv4 Address dialog, you can do the following:
   - Enable DHCP for automatic assignment of the IP address through a DHCP server. See Section 9.3.3.1.1, Enabling DHCP, for more information.
   - Configure a static IP address (if DHCP is disabled). See Section 9.3.3.1.2, Configuring a Static IPv4 Address, for more information.
   - Ping the iSCSI initiator. See Section 9.3.3.2, Pinging the iSCSI Initiator, for more information.

9.3.3.1.1 Enabling DHCP

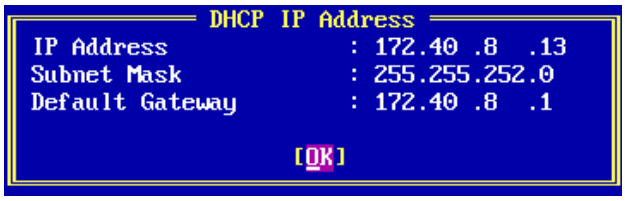
Enabling DHCP provides the initiator with an IP address.

NOTE If you are using the DHCP server to obtain an IP address for the iSCSI initiator, set up a reservation. A reservation assigns a specific IP address based on the MAC address of the iSCSI function. If you do not reserve an IP address through DHCP, you must set the lease length for the iSCSI initiator IP address to unlimited. This prevents the IP address lease from expiring.

To enable DHCP:

1. From the Configure IPv4 Address dialog (Figure 59), select **ENABLED** from the DHCP pop up and press Enter.
2. Select Save and press Enter. The DHCP IP Address dialog box appears.

Figure 60 DHCP IP Address Dialog Box



NOTE If you set a static IP address, a message warns that the static IP address will be lost. Press Y to continue with enabling DHCP.

3. Click OK.

9.3.3.1.2 Configuring a Static IPv4 Address

If a DHCP server is not available, you must manually configure a static IPv4 address.

To configure a static IP address:

1. From the Configure IPv4 Address dialog (Figure 59), verify that DHCP is disabled.
2. Enter the IP address. This is a unique 32-bit address that specifies the location of each device or workstation in the network. This address is required for communication to the iSCSI SAN. For an initiator with its own network interface, you must assign an IP address in the same range as the iSCSI SAN.

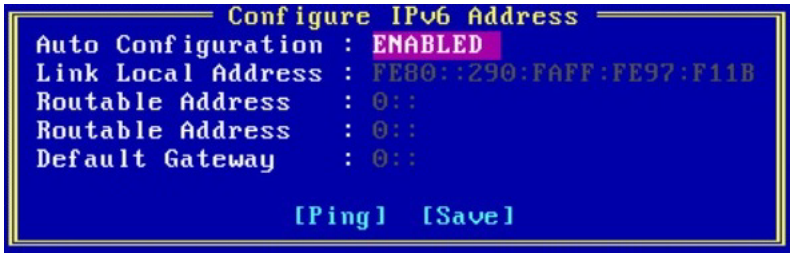
3. Enter the subnet mask. The subnet mask provides a way to segment networks. All hosts (iSCSI initiators or targets) on the same physical network usually have the same subnet mask. For example, if the initiator is in the subnet 255.255.xxx.xxx, all the hosts (initiators or targets) in a sub-network will have the same subnet mask.
4. Enter the default gateway information, if necessary. A gateway is a router on a computer network that serves as an access point to another network and that an initiator uses by default. Any data to addresses not on the initiator's subnet are sent through the default gateway.
5. Select Save and press Enter to save the settings.
6. When you are prompted to save the changes, type Y.

9.3.3.1.3 Configuring an IPv6 Address

To configure an IPv6 address:

1. On the Network Configuration dialog (Figure 57), verify that the IP Version is set to IPv6.
2. Select Configure IPv6 Address and press Enter. The Configure IPv6 Address dialog appears.

Figure 61 Configure IPv6 Address Dialog



3. From the Configure IPv6 Address dialog, you can do the following:
   - Enable auto-configuration for automatic assignment of the IP address. See Section 9.3.3.1.1, Enabling DHCP, for more information.
   - Configure a static IP address (if Auto Configuration is disabled). See Section 9.3.3.1.5, Configuring a Static IPv6 Address, for more information.
   - Ping the iSCSI initiator. See Section 9.3.3.2, Pinging the iSCSI Initiator, for more information.

9.3.3.1.4 Enabling Auto Configuration

To automatically configure the IPv6 address, from the Configure IPv6 Address dialog (Figure 61), select **ENABLED** from the Auto Configuration pop up and press Enter.

9.3.3.1.5 Configuring a Static IPv6 Address

If Auto Configuration is disabled, you must manually configure a static IPv6 address.

To configure a static IP address:

1. From the Configure IPv6 Address dialog (Figure 61), enter the following information in the corresponding fields:
   - **Link Local Address** – The unique address assigned to the port 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 port.
   - **Routable Address 2** – A second routable address assigned to the iSCSI port.
   - **Default Gateway** – A gateway is a router on a computer network that serves as an access point to another network and that an initiator uses by default.
2. Select Save and press Enter to save the settings.
3. You are prompted to save the changes, type Y.

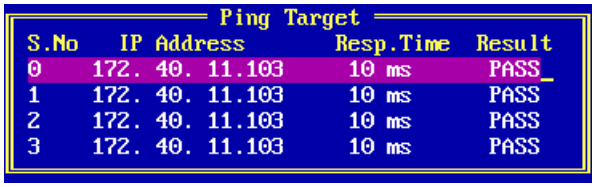
9.3.3.2 Pinging the iSCSI Initiator

After you configure the IPv4 or IPv6 address, you can confirm your network settings by pinging the network. The ping option verifies that the IP address is on the network. If there is another IP entity on that network when you ping, you receive a ping reply back to confirm the network settings. Ping can be a diagnostic tool but it is also a validation that your network is set up properly prior to target login.

To verify that a target is accessible and that you are connected to the network:

1. From the Configure IPv4 Address dialog (Figure 59) or Configure IPv6 Address dialog (Figure 61), select Ping and press Enter.
2. In the Ping Target dialog box, enter the IP address of the iSCSI device to ping. You are notified that the ping is successful. If the ping is unsuccessful, you will receive a failed message. For more information, see Section 14.2, Troubleshooting for the iSCSI Protocol. The following screen is an example of a successful ping.

Figure 62 Successful Ping



NOTE The ping works for any IP address that supports ICMP (Echo).

If you cannot verify the network interface, see Section 14.2, Troubleshooting for the iSCSI Protocol.

9.4 Identifying a Port

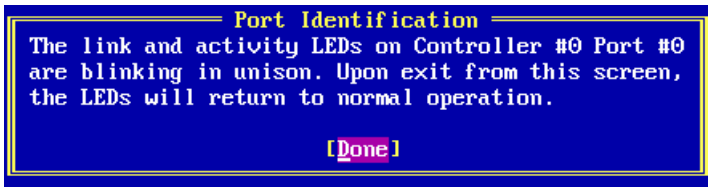
Port identification, or beaconing, helps you determine which port you are configuring by blinking both the link and the activity LEDS of that port.

LEDs blink on the back of the server so there is no confusion as to which physical port you are configuring with the iSCSISelect utility. Port identification lets you correlate the iSCSI software configuration with the hardware port.

NOTE Not all adapters have LEDs that are externally visible. If you are using an add-in card in a blade server environment, port identification is not supported.

To identify a port, from the Controller Configuration menu (Figure 55), select Port Identification and press Enter. The LED status indicators for the selected port blink on the adapter until you select Done and press Enter on the Port Identification pop up (Figure 63).

Figure 63 Port Identification Pop Up



9.5 Erasing the Current Configuration

Erase Configuration erases the configuration of a single adapter. Configuration data is erased for both ports on the selected adapter. The initiator name is global for all Emulex OneConnect adapters in the system. If you have more than one adapter and you erase the configuration on the first one, the Erase Configuration option resets the initiator names back to their default values. If you erase the configuration on the second adapter, the default values are reset only for the second adapter and are not reset globally on all adapters.

NOTE You must select Erase Configuration to clear out existing IQN data if you purchase a different or subsequent license for the adapter.

To erase an adapter configuration:

1. From the Controller Configuration menu (Figure 55), select Erase Configuration and press Enter.
2. A message appears asking if you want to erase the current configuration. Type Y.
3. You are cautioned that the operation removes any existing configuration permanently. Type Y.

   After the adapter configuration is erased, the system displays the Controller Configuration menu (Figure 55).

Chapter 10: Configuring and Managing iSCSI Targets with the iSCSISelect Utility

An iSCSI target is any device that receives iSCSI commands. The device can be an end node, such as a storage device, or it can be an intermediate device, such as a bridge between IP and FC devices. Each iSCSI target is identified by a unique IQN, and each port on the adapter (or on a bridge) is identified by one or more IP addresses.

As the client system, the initiator initiates disk I/O. It communicates with the target, and the target serves the initiator by providing resources to the initiator.

After you configure the iSCSI initiator and have verified that the network interface is working, you can log in to the target to establish the connection. To do this, you must:

- Perform a login to the target (adding a target).
- Ping a target and verify that you can connect to it.
- Reboot the system.

10.1 Adding iSCSI Targets

After the initiator has been configured, you must make an iSCSI target available to the initiator host. Targets can be discovered automatically using the SendTargets discovery mechanism, or you can manually configure targets.

Before you begin the login session, gather the following information:

- iSCSI target name (manual configuration only) – The target name you are adding. This name should be known to you based on how you configured the iSCSI target. For more information, see Section 10.1.1, Using the SendTargets Discovery to Add an iSCSI Target.
- iSCSI target IP address – The IP address of the target into which you are logging.
- TCP port number – The TCP port number of the target portal. Typically this is 3260, which is the well-known port number defined for use by iSCSI.

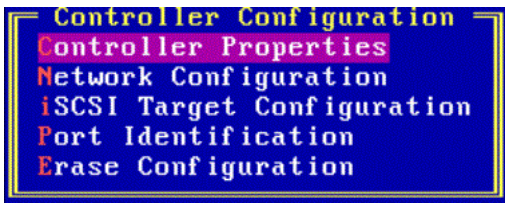
10.1.1 Using the SendTargets Discovery to Add an iSCSI Target

The iSCSI initiator uses the SendTargets discovery to locate names of available targets. Use this method if an iSCSI node has a large number of targets.

To configure an iSCSI target using the SendTargets discovery:

1. Log in to the iSCSISelect utility by pressing the Ctrl-S keys when prompted.
2. Select Controller Configuration and press Enter. The system displays a list of adapters.
3. Select an adapter and press Enter. The following Controller Configuration menu appears.

Figure 64 Controller Configuration Menu



4. Select iSCSI Target Configuration and press Enter. The following iSCSI Target Configuration dialog appears.

Figure 65 iSCSI Target Configuration Dialog



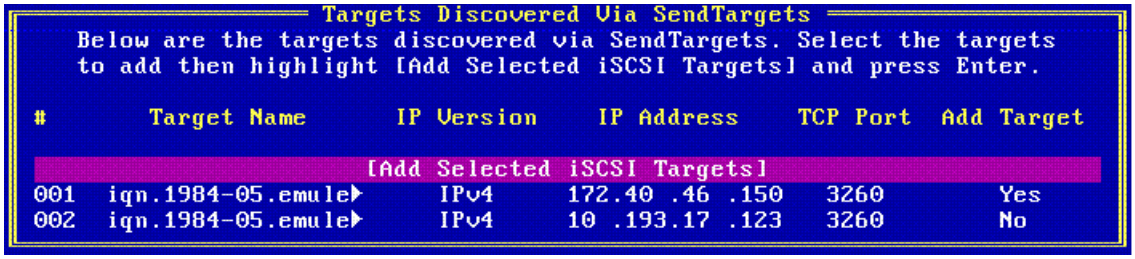
5. Select Add New iSCSI IPv4 Target or Add New iSCSI IPv6 Target and press Enter. The following Add iSCSI Target dialog appears.

Figure 66 Add/Ping iSCSI Target Dialog



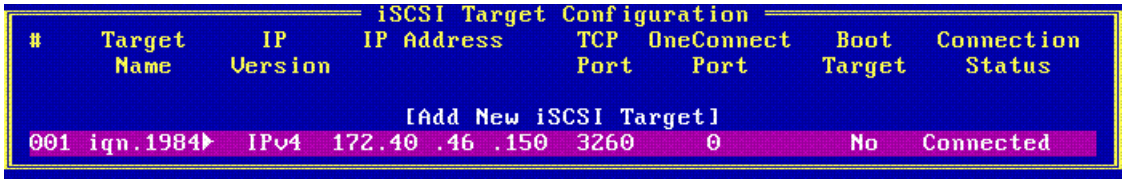
6. Leave the iSCSI Target Name blank.
7. Enter the iSCSI target IP address.
8. Change the TCP port number value, if necessary. The default target port number is 3260.
9. Leave the Boot Target field set to *No* for now; you will return to this field later to enable the boot target.
10. If you are connecting dual sessions to the same target portal group, enter a unique ISID Qualifier, which can be any number up to 65535. For more information on when an ISID value must be changed, see Section 5.3.2, Booting from SAN for iSCSI MPIO.
11. Set the Header Digest to Yes if you want to enable the header digest. If the Header Digest is set to Yes, and the iSCSI initiator is set accordingly, the integrity of an iSCSI protocol data unit's header segment is protected by CRC32C checksum. The default is disabled.
12. Set the Data Digest to Yes if you want to enable the data digest. If the Data Digest is set to Yes, and the iSCSI initiator is set accordingly, the integrity of an iSCSI protocol data unit's data segment is protected by CRC32C checksum. The default is disabled.
13. Select an authentication method (optional). If you are enabling an authentication method, you are prompted to enter the CHAP configuration. For more information, see Section 10.2.2.1, Enabling CHAP Authentication.
14. Select Save/Login. A message reminds you that you have left the iSCSI Target name blank and that the SendTargets mechanism will be used. To continue, press Y.
15. After you send the SendTargets request, the firmware attempts to log in to the target. If it is successful, the Targets Discovered Via SendTargets dialog (Figure 67) displays a list of targets.

Figure 67 Targets Discovered Via SendTargets Dialog



16. Select the targets you want to add from the list and press the F3 key.
17. Add these targets to the list of iSCSI targets available for the initiator to log in to by selecting Add Selected iSCSI Targets, and press Enter. If you enabled CHAP Authentication, you are prompted to enter the CHAP configuration for each target, one at a time. For more information, see Section 10.2.2.1, Enabling CHAP Authentication.
18. From the Controller Configuration menu (Figure 64), select iSCSI Target Configuration and press Enter to view the added target information.

Figure 68 iSCSI Target Configuration Dialog



NOTE If you set the Boot Target option in step 5 before adding the target, the Boot Target appears as No. To enable Boot Target, continue with step 19. Otherwise, you have completed the procedure.

19. To enable Boot Target or make any other changes to the targets, select the target and press Enter. The Edit/Ping iSCSI Target dialog (Figure 70) appears. From this dialog, you can edit the target.

10.1.2 Manually Adding an iSCSI Target

Use this method if an iSCSI node has a small number of targets, or if you want to restrict the targets that an initiator can access. To manually configure a target, you must provide the iSCSI target name.

To configure an iSCSI target manually:

1. Log in to the iSCSISelect utility by pressing the Ctrl-S keys when prompted.
2. Select Controller Configuration and press Enter. The system displays a list of adapters.
3. Select an adapter and press Enter. The Controller Configuration menu (Figure 64) appears.
4. Select iSCSI Target Configuration and press Enter. The iSCSI Target Configuration dialog (Figure 68) appears.
5. Select Add New iSCSI Target and press Enter. The Add/Ping iSCSI Target dialog (Figure 66) appears.
6. Enter the iSCSI target name, iSCSI target IP address, and TCP port number (the default target port number is 3260).
7. If you want to enable the target as a boot target, select Yes.
8. If you are connecting dual sessions to the same target portal group, enter a unique ISID Qualifier, which can be any number up to 65535. For more information on when an ISID value must be changed, see Section 5.3.2, Booting from SAN for iSCSI MPIO.

9. Set the Header Digest to Yes if you want to enable the header digest. If the Header Digest is set to Yes, and the iSCSI initiator is set accordingly, the integrity of an iSCSI protocol data unit's header segment is protected by CRC32C checksum. The default is disabled.
10. Set the Data Digest to Yes if you want to enable the data digest. If the Data Digest is set to Yes, and the iSCSI initiator is set accordingly, the integrity of an iSCSI protocol data unit's data segment is protected by CRC32C checksum. The default is disabled.
11. Select the Authentication Method. For more information, see Section 10.2.2.1, Enabling CHAP Authentication.
12. Select Save/Login and press Enter. If the firmware logs into the target successfully, you receive a successful login message. Press Esc to go to the iSCSI Target Configuration dialog (Figure 68), where you can view information about the added target.
13. To enable Boot Target or make any other changes to the targets, select the target and press Enter.

10.2 Managing an iSCSI Target

With the iSCSISelect utility, you can manage a target by viewing target details or by editing the target configuration.

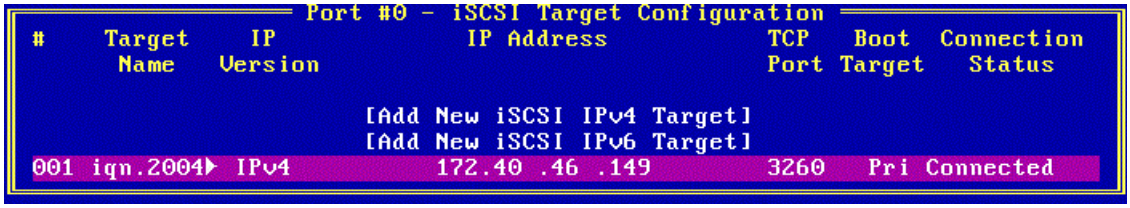
10.2.1 Viewing Target Information

After successfully logging in to the target, you can view the details of the newly added target or manage the target. From the Controller Configuration menu (Figure 64), select iSCSI Target Configuration and press Enter to view the target information.

NOTE The iSCSISelect utility shows LUN information for the first 128 LUNs only.

The following screen is an example of a target detail.

Figure 69 iSCSI Target Configuration Information



In the iSCSI Target Configuration dialog, the function keys located at the bottom of the menu help you manage the target configuration.

- Press Enter to select the target configuration.
- Press the Delete key to delete the target.
- Press the F5 key to log in to the target.
- Press the F6 key to log out of the target.
- Press the F7 key to configure the LUN.
- Press Esc to return to the previous menu.

10.2.2 Editing a Target Configuration

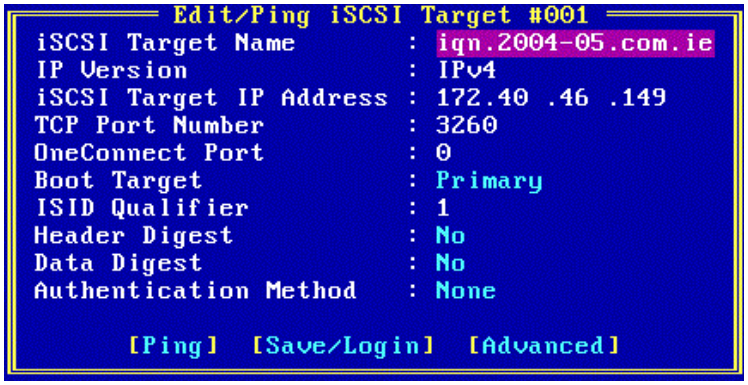
After adding a target, you can edit the iSCSI target configuration or apply other management options to the target.

NOTE If you want to change the target name, you must delete the existing target and add it again with the new target name.

To edit a target configuration:

1. From the iSCSI Target Configuration dialog (Figure 68), select the target and press Enter. The following Edit/Ping iSCSI Target dialog appears.

Figure 70 Edit/Ping iSCSI Target Dialog



2. If you want to enable the target as a boot target, select Boot Target and select either Primary or Secondary in the Boot Target pop up.
3. If you are connecting dual sessions to the same target portal group, enter a unique ISID Qualifier, which can be any number up to 65535. For more information on when an ISID value must be changed, see Section 5.3.2, Booting from SAN for iSCSI MPIO.
4. Set the Header Digest to Yes if you want to enable the header digest. If the Header Digest is set to Yes, and the iSCSI initiator is set accordingly, the integrity of an iSCSI protocol data unit's header segment is protected by CRC32C checksum. The default is disabled.
5. Set the Data Digest to Yes if you want to enable the data digest. If the Data Digest is set to Yes, and the iSCSI initiator is set accordingly, the integrity of an iSCSI protocol data unit's data segment is protected by CRC32C checksum. The default is disabled.
6. To select an authentication method to use when logging in to the target, select an authentication method from the Authentication Method pop up. For more information about authentication methods, see Section 10.2.2.1, Enabling CHAP Authentication.
7. To verify the target connection, select Ping and press Enter. For more information about ping, see Section 10.1, Adding iSCSI Targets.
8. To view more target properties, select Advanced and press Enter. For more information on these properties, see Section 10.2.3, Viewing Advanced Target Properties.
9. To save your changes, select Save/Login and press Enter. The system displays the iSCSI Target Configuration dialog with the revised information.

10.2.2.1 Enabling CHAP Authentication

The iSCSISelect utility uses CHAP to authenticate initiators and targets for added network security. By using a challenge and response security mechanism, CHAP periodically verifies the initiator's identity. This authentication method depends on a secret key known to the initiator and the target only. Although the authentication can be one-way, you can negotiate CHAP in both directions with the help of the same secret set for mutual authentication. You must make sure, however, that what you configure on the target side matches the initiator side. The iSCSISelect utility supports both one-way and mutual authentication.

10.2.2.1.1 Authenticating One-Way CHAP

With one-way CHAP authentication, the target authenticates the initiator. Use one-way CHAP authentication for a one-way challenge and response security method.

To enable one-way CHAP authentication:

1. Log in to the iSCSISelect utility by pressing the Ctrl-S keys when prompted.
2. Select Controller Configuration. The system displays a list of adapters.
3. Select an adapter and press Enter. The Controller Configuration menu (Figure 64) appears.
4. Select iSCSI Target Configuration and press Enter. The iSCSI Target Configuration dialog (Figure 68) appears.
5. Select Add New iSCSI Target and press Enter.
6. Follow steps 5 to 11 in Section 10.1.1, Using the SendTargets Discovery to Add an iSCSI Target.
7. On the Add/Ping iSCSI Target dialog (Figure 66), select Authentication Method, select One-Way Chap, and press Enter. The CHAP Configuration dialog box (Figure 71) appears.

Figure 71 One-Way Chap Configuration Dialog Box



8. Enter the target CHAP name and target secret, then select OK and press Enter.

   The target CHAP name can be any name or sequence of numbers consisting of 12 to 16 characters. However, the name and secret on the target side must match the name and target on the initiator side.

   The target secret can be in string (for example, abcdefg23456) or hexadecimal (for example, 0x01234567890123456789ABCD) format.

   — For string format, the secret can be any name or sequence of numbers with a minimum of 12 and a maximum of 16 bytes, where each character equals one byte.

   — For hexadecimal format, the secret must be a minimum of 12 and a maximum of 16 bytes long, where two characters equal one byte and the character length must be even (24, 26, 28, 30, or 32).

   — The following restrictions apply to hexadecimal formatting:

   - The 0x representation requires an even number of hexadecimal values excluding the 0x (0-9, A-F). For example, the secret 0x123456789012345678 9ABCD is not allowed because the character length is odd, with 23 hexadecimal values excluding the 0x. The secret 0x01234567890123456789ABCD is allowed, with 24 hexadecimal values excluding the 0x.
   - The hexadecimal representation of CHAP secrets must contain hexadecimal values only. The 0x representation may include only the following characters: 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, a, b, c, d, e, f, A, B, C, D, E, and F. No other characters are allowed.

   After enabling the CHAP authentication, the Add/Ping iSCSI Target dialog (Figure 66) appears. For more information, see Section 10.2.2.2, Pinging a Target.

10.2.2.1.2 Authenticating Mutual CHAP

With mutual CHAP authentication, the target authenticates the initiator and the initiator authenticates the target. Use mutual CHAP authentication for a two-way challenge and response security method.

To enable mutual CHAP authentication:

1. Log in to the iSCSISelect utility by pressing the Ctrl-S keys when prompted.
2. Select Controller Configuration and press Enter. The system displays a list of adapters.

3. Select an adapter and press Enter. The Controller Configuration menu (Figure 64) appears.
4. Select iSCSI Target Configuration and press Enter. The iSCSI Target Configuration dialog (Figure 68) appears.
5. Select Add New iSCSI Target and press Enter. The Add/Ping iSCSI Target dialog (Figure 66) appears.
6. Select Authentication Method, select Mutual Chap, and press Enter. The following CHAP Configuration dialog box (Figure 72) appears.

Figure 72 Mutual CHAP Configuration Dialog Box



7. Enter the target CHAP name, target secret, initiator CHAP name, and initiator secret. Select OK and press Enter.

   The target CHAP name can be any name or sequence of numbers with a minimum of 12 and a maximum of 16 characters. However, the name and secret on the target side must match the name and target on the initiator side.

   The target secret can be in string (for example, abcdefg23456) or hexadecimal (for example, 0x01234567890123456789ABCD) format.

   — For string format, the secret can be any name or sequence of numbers with a minimum of 12 and a maximum of 16 bytes, where each character equals one byte.

   — For hexadecimal format, the secret must be a minimum of 12 and a maximum of 16 bytes long, where two characters equal one byte and the character length must be even (24, 26, 28, 30, or 32).

   — The following restrictions apply to hexadecimal formatting:

   - The 0x representation requires an even number of hexadecimal values excluding the 0x (0-9, A-F). For example, the secret 0x123456789012345678 9ABCD is not allowed because the character length is odd, with 23 hexadecimal values excluding the 0x. The secret 0x01234567890123456789ABCD is allowed, with 24 hexadecimal values excluding the 0x.
   - The hexadecimal representation of CHAP secrets must contain hexadecimal values only. The 0x representation may include only the following characters: 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, a, b, c, d, e, f, A, B, C, D, E, and F. No other characters are allowed.

   After enabling the CHAP authentication, the Add/Ping iSCSI Target dialog (Figure 66) appears.

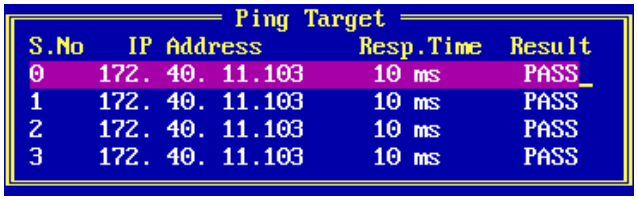
10.2.2.2 Pinging a Target

To ping a target:

1. From the Controller Configuration menu (Figure 64), select iSCSI Target Configuration and press Enter. The iSCSI Target Configuration dialog (Figure 68) appears.
2. Select Add/Ping iSCSI Target or Edit/Ping iSCSI Target and press Enter. The iSCSI Target IP Address is the address that will be pinged.
3. From the Add/Ping iSCSI Target dialog (Figure 66) or the Edit/Ping iSCSI Target dialog (Figure 70), select Ping and press Enter.

   The following Ping Target information displays the results of the ping.

Figure 73 Ping Target Results

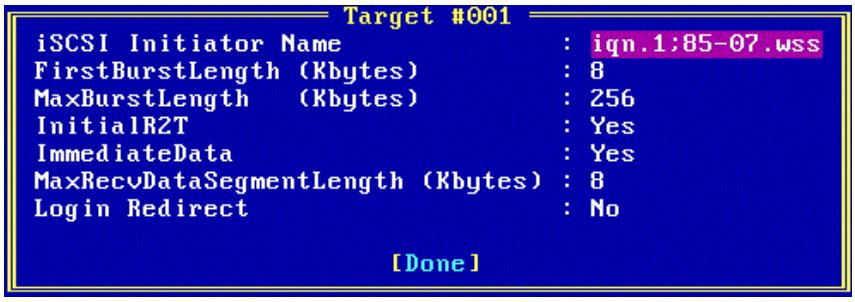


A ping can be unsuccessful for several reasons; for more information, see Section 14.2, Troubleshooting for the iSCSI Protocol.

10.2.3 Viewing Advanced Target Properties

After logging in to a target, you can view more information about it by navigating to the Edit/Ping iSCSI Target dialog (Figure 70), selecting the Advanced option, and pressing Enter. The Advanced Target Properties dialog appears.

Figure 74 Advanced Target Properties Dialog



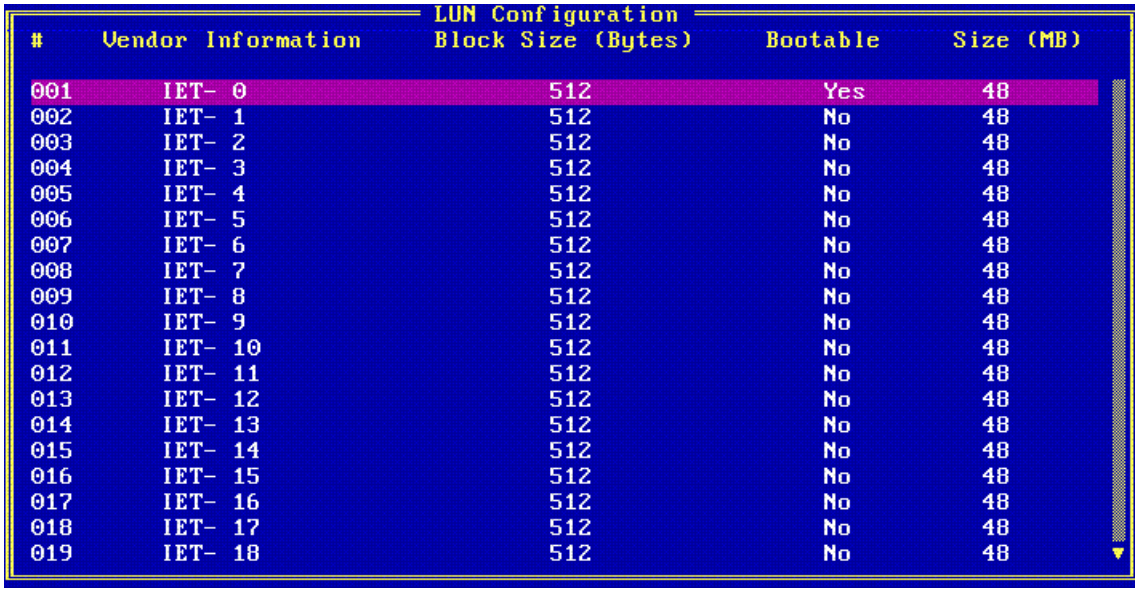
10.2.4 Configuring LUN Settings

A LUN represents an individually addressable logical device that is part of a target. An initiator negotiates with a target to establish connectivity to a LUN.

To configure a LUN:

1. From the iSCSI Target Configuration dialog (Figure 68), press the F7 key to show the LUNs associated with the iSCSI node.

Figure 75 LUN Configuration Menu



2. If the target is a boot target, press the F3 key to set the boot LUN. By setting the LUN to bootable, you can confirm that the boot target is configured properly and you can see the LUN that is booting. By looking at the LUN size, you can also identify each target along with its LUN order.

NOTE The LUN order is determined by the target and is in the order listed.

If there are no LUNs available, the system displays the following message:

```
No LUN available, please check your configuration on the Target.
```

For more information, see Section 14.2, Troubleshooting for the iSCSI Protocol.

10.2.5 Logging In to and Out of a Configured Target

To log in to a target, from the iSCSI Target Configuration dialog (Figure 68), select the disconnected target, and then press the F5key to establish the target session.

If you log out of a target, it is still included in the target list, with the status of disconnected. You are deleting the target session only; you can log back in to it.

To log out of a target, from the iSCSI Target Configuration dialog (Figure 68), select the target and press the F6 key.

10.2.6 Removing a Configured Target

To fully remove a target, from the iSCSI Target Configuration dialog (Figure 68), select the disconnected target and press the Delete key.

10.2.7 Booting the System

After configuring the iSCSI initiator and target, you must reboot the system for the configurations to take effect. When you exit the iSCSISelect utility, the system automatically reboots, and during system startup the Emulex adapter BIOS information appears. For example:

```
<<< Press <Ctrl><S> for iSCSISelect(TM) Utility >>>

iSCSI#0 Port#1 Base 0x00000380002A0000 @ Bus:1B Dev:00 Fun:02
iSCSI#0 Port#2 Base 0x00000380002E0000 @ Bus:1B Dev:00 Fun:03

Initiator iSCSI Name:  iqn.1990-07.com.emulex:00-90-fa-6b-ba-2c
Controller #0:0 Initiator IPV6 Address: FE80:0:0:0:290:FAFF:FE6B:BA2D
Controller #0:0 Initiator IPV4 Address: 0.0.0.0
Controller #0:1 Initiator IPV6 Address: FE80:0:0:0:290:FAFF:FE6B:BA35
Controller #0:1 Initiator IPV4 Address: 0.0.0.0

<Ctrl><S> Pressed-Utility will be invoked after BIOS initialization.
```

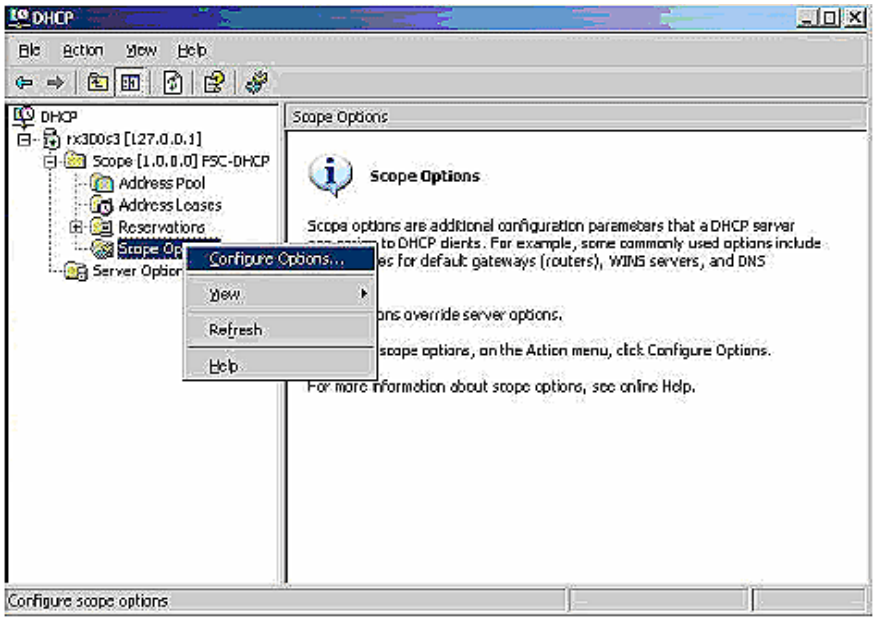
NOTE The iSCSI BIOS logs into the configured iSCSI boot target and shows its target and LUN information in the BIOS sign-on banner only if you have configured an iSCSI boot target. For more information, see Section 8.2, Setting Up a Basic iSCSI Boot Configuration.

10.3 Discovering Targets through DHCP for iSCSI

For the target to be discovered by the initiator through DHCP, you must add the root path to the DHCP server:

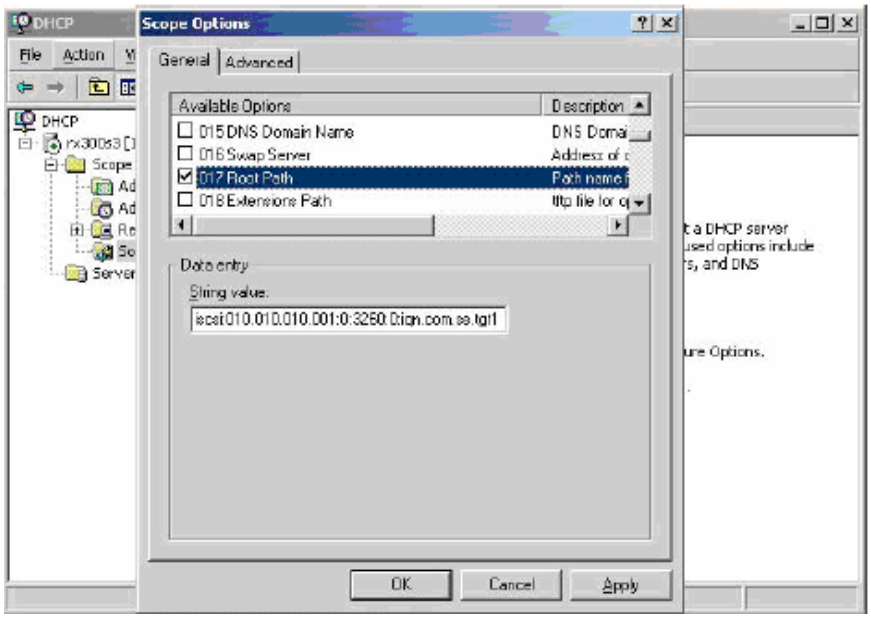
1. From the following DHCP server window (Figure 76), select Scope Options, then right-click and select Configure Options.

Figure 76 DHCP Server Window



2. From the General tab shown in the following dialog (Figure 77), select 017 Root Path.

Figure 77 DHCP Server Scope Options



3. In the String Value field, add the location of the iSCSI target you want to use for iSCSI boot. Enter the string in the following format:

   ```
   iscsi: <ipaddress>:<protocol>:<iscsi port number>:<luns><target name>
   ```

   For example:

   ```
   iscsi:010.010.010.001:0:3260:0:iqn.1992-08.com.netap:sn.15729740
   ```

4. Click Apply. The DHCP server window is ready to discover boot targets.

10.4 Enabling DHCP Discovery through the iSCSISelect Utility

If the DHCP server supports DHCP discovery and you added the root path to the DHCP server, you can enable boot target discovery using DHCP with the iSCSISelect utility. By enabling Discover IPv4 Boot Target via DHCP, you can use DHCP discovery to provide the boot path from the DHCP server.

NOTE If you leave Discover IPv4 Boot Target via DHCP disabled (default), you must manually configure iSCSI boot.

To enable DHCP discovery through iSCSISelect:

1. After configuring the iSCSI Initiator Name, select Controller Properties from the Controller Configuration menu (Figure 64) and press Enter.
2. From the Controller Properties menu, select Discover IPv4 Boot Target via DHCP and press Enter. A pop up appears.
3. To enable the function, select **Enabled** and press Enter.
4. Select Save and press Enter.

Chapter 11: Configuring UEFI for Ethernet

The UEFIBoot supports:

- UEFI version – 2.3.1 compatible drivers (backward compatible with 2.1)
- EDK2 compatible drivers
- The following UEFI protocols:
  - NII protocol
  - Configuration Access protocol
  - Component Name2 protocol
  - Driver Diagnostics protocol
  - Platform to Driver Configuration protocol
  - Firmware Management Protocol
- Switching of personalities from the HII menus using the multimode feature
- Enabling, disabling, and configuring UMC from HII
- Lenovo System x Blade Open Firmware Management Protocol and Universal Configuration Manager

  NOTE Applies to Lenovo System x products only.
- Up to 96 physical adapters or 192 ports
- Single driver – support for all adapters in the system

This version of UEFIBoot is loaded from flash into system memory.

For more information on loading and unloading UEFIBoot, see Appendix E, Loading and Unloading UEFIBoot from the UEFI Shell.

If UEFIBoot is loaded in a UEFI Shell, type `drivers` and press Enter to see if the driver is loaded.

NOTE If you are using a Dell adapter, see Appendix F, Dell UEFI for information on using the Dell UEFI utility.

11.1 Navigating the Emulex NIC Configuration Utility

Use the following methods to navigate the Emulex NIC configuration utility:

- Press the up/down arrows on the keyboard to move through menu options or fields. If multiple adapters are listed, use the up/down arrows to scroll to the additional adapters.
- Press the +, -, or Enter keys to change numeric values.
- Press Enter to select a menu option, to select a row in a configuration dialog, or to change a configuration default.
- Use the navigation entries on the page to move about the utility.
- Press Esc to exit the current screen and show the previous screen.

NOTE HP adapters in HP systems use the F10 key to save the current configuration settings.

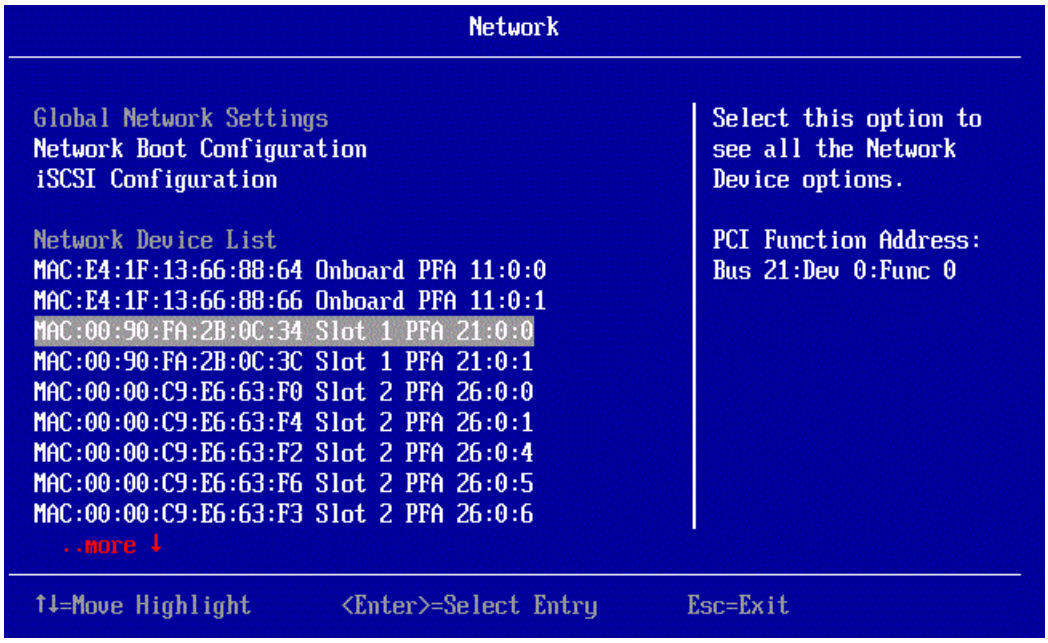
11.2 Starting the Emulex NIC Configuration Utility

Depending on the vendor UEFI configuration, the Emulex NIC configuration utility may appear under different setup menus in the vendor system firmware or BIOS (such as **System Settings > Network**).

To start the Emulex NIC configuration utility using the HII:

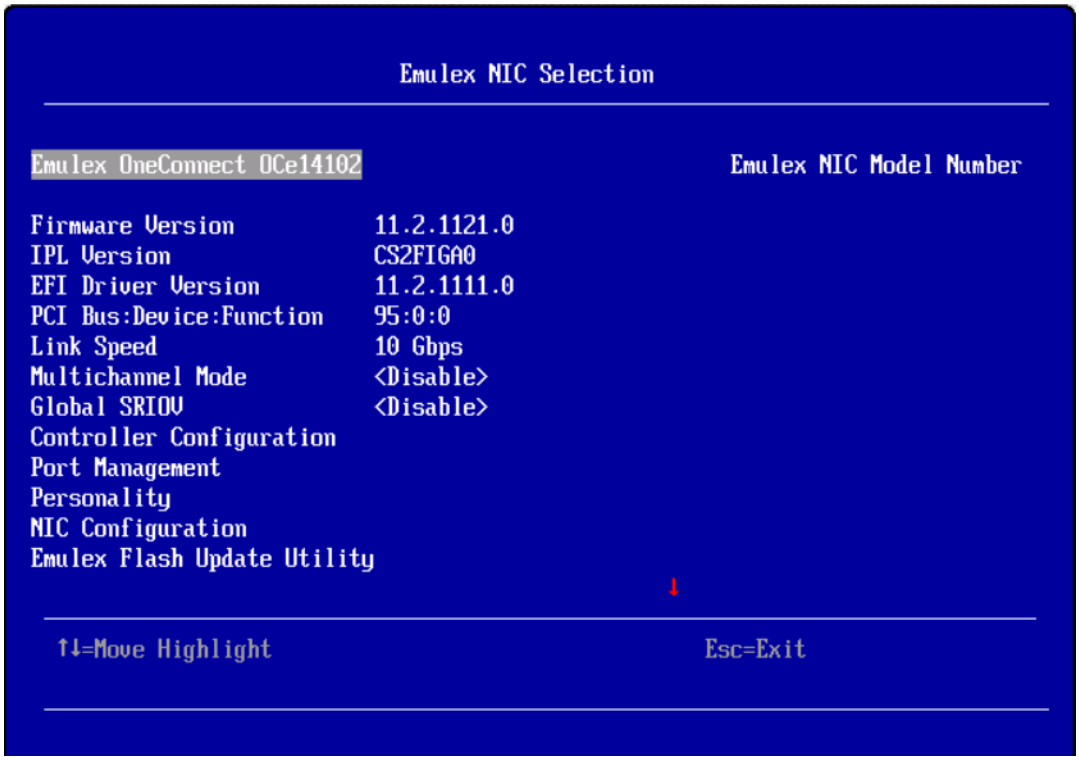
1. The Network dialog (Figure 78) shows a list of the adapters in the system. Your list may vary depending on the installed adapters.

Figure 78 Network Dialog



2. Select an adapter and press Enter. The Emulex NIC Selection dialog (Figure 79) appears.

Figure 79 Emulex NIC Selection Dialog



The Emulex NIC Selection dialog shows the following adapter information and configuration options:

NOTE Some of the following options may or may not appear depending on your adapter and the system in which it is installed.

- **Firmware Version** – The current firmware version installed on the adapter. The firmware version is the same as the boot version.
- **IPL Version** – The current version installed on the adapter.
- **EFI Driver Version** – The current version of the EFI driver installed on the adapter.
- **PCI Bus:Device:Function** – The PCI bus, device, and function for the adapter. This information describes the location of the adapter in the PCI configuration space.
- **Link Speed** – The current maximum port speed on the adapter.
- **Advanced Mode (not shown)** – See Section 11.2.1, Advanced Mode Support, for more information.
- **Multichannel Mode** – See Section 11.2.8, Configuring Universal Multichannel, for more information.
- **Global SRIOV** – The SRIOV setting on an adapter level. The SRIOV setting can be enabled on a port basis, only if this parameter is set to Enabled.
- **Controller Configuration** – The available menu options depend on whether UMC is enabled or disabled. See Section 11.2.7, Configuring the Adapter If UMC Is Disabled or Not Supported, or Section 11.2.8, Configuring Universal Multichannel, for more information.
- **Boot Mode (not shown)** – See Section 11.2.2, Configuring Boot Mode, for more information.

NOTE If the system BIOS supports boot mode configuration, this feature is not displayed in the menu.

- **iBFT Configuration (not shown)**- This menu is available only if Boot mode is disabled. See Section 11.2.3, Configuring iBFT, for more information.

NOTE If the system BIOS supports iBFT configuration, this feature is not displayed in the menu.

- **Port Management** – See Section 11.2.10, Port Management, for more information.
- **Personality** – See Section 11.2.4, Configuring Personality, for more information.
- **Link Reconfiguration (not shown)** – See Section 11.2.5, Link Reconfiguration, for more information.
- **NIC Configuration** – See Section 11.2.6, Configuring Virtual LAN Settings, for more information.

NOTE This menu item is available only when Multichannel mode is disabled.

- **Emulex Flash Update Utility** – See Section 11.3, Downloading Firmware and Boot Code, for more information.
- **FoD (Lenovo only; not shown)** – See Section 11.2.11, Feature on Demand, for more information.
- **Port Identification (not shown)** – See Section 11.4, Identifying a Port, for more information.
- **Erase Configuration (not shown)**– See Section 11.5, Erasing Ports and Adapter Configurations, for more information.

NOTE The Erase Configuration option is not available on HP adapters in HP systems. Instead, use the F7 key to reset the configuration settings to their factory default settings.

11.2.1 Advanced Mode Support

Advanced mode is a driver compatibility option. With Advanced mode enabled, you can run Advanced mode-aware drivers that provide advanced capabilities as listed in the following table. With Advanced mode disabled, you can run older legacy inbox drivers that are not Advanced mode-aware with the latest firmware versions.

NOTE Advanced mode support is enabled by default on Emulex OCe11100-series two-port and four-port adapters. On four-port adapters, the Advanced mode setting is not provided in the HII utilities. The Advanced mode setting on these platforms is implicitly enabled and Advanced mode-aware drivers must be installed. Compatibility with legacy drivers requires that Advanced mode support be disabled on two-port adapters.

Advanced mode is automatically enabled on Emulex OCe14000-series adapters and it cannot be disabled.

Table 9 Advanced Mode Capabilities (by Operating System)

Operating System	Advanced Mode Enabled	Advanced Mode Disabled (Legacy Mode)
Windows	16 RSS queues **NOTE** Supported on Windows Server 2008 R2 and later versions only.	4 RSS queues
	VMQ lookahead split **NOTE** VMQs are supported on Windows Server 2008 R2 and later versions only.	Lookahead split is silently ignored. A small performance penalty for VMQs may occur.

Table 9 Advanced Mode Capabilities (by Operating System) (Continued)

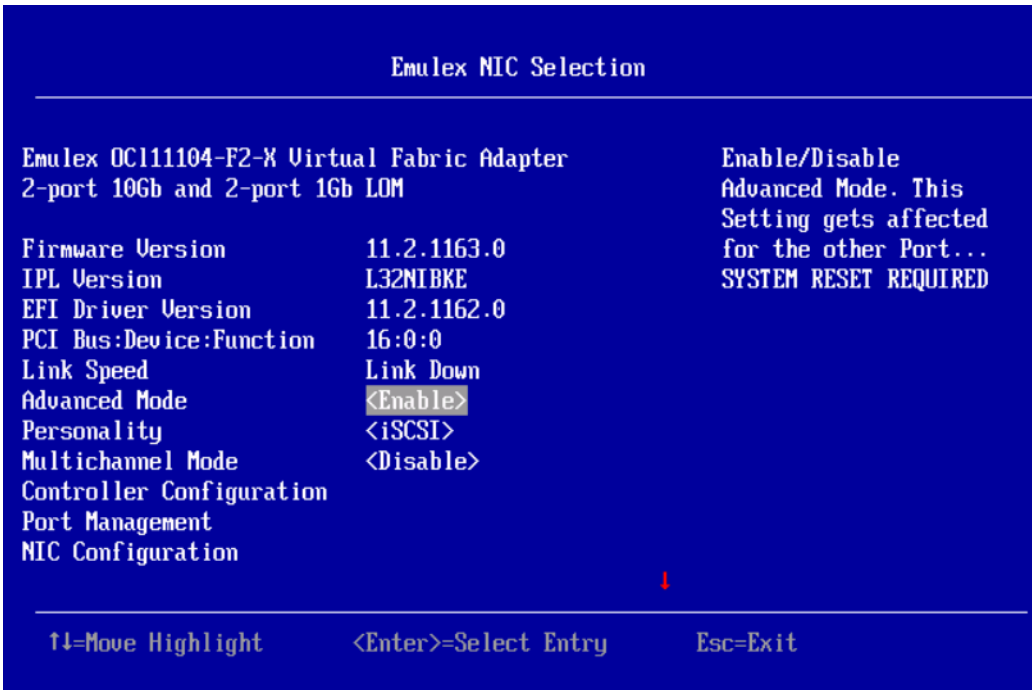
Operating System	Advanced Mode Enabled	Advanced Mode Disabled (Legacy Mode)
Linux and Citrix	16 RSS queues	4 RSS queues
	VFs and PFs can be increased up to 30.	
VMware ESXi	For both 1500 and 9000 MTU: 16 NetQueues per PF in non-VFA 4 NetQueues per PF in VFA	1500 MTU – 8 NetQueues per PF in non-VFA and 4 NetQueues per PF in VFA. 9000 MTU – 4 NetQueues per PF in both VFA and non-VFA.

To configure Advanced mode support:

NOTE Changing the Advanced mode setting requires a reboot.

1. On the Emulex NIC Selection dialog (Figure 80), select Advanced mode and press Enter. A pop up appears.

Figure 80 Emulex NIC Selection Dialog with Advanced Mode Displayed



2. Select **Enable** or **Disable** and press Enter.
3. Reboot the system.

11.2.2 Configuring Boot Mode

This setting configures the boot mode.

- Enable – iSCSI hardware initiator boot mode
- Disable – iSCSI software initiator boot mode

NOTE If boot mode is disabled, the iBFT Configuration menu is available.

To configure boot mode:

1. On the Emulex NIC Selection dialog (Figure 79), select Boot Mode and press Enter. A pop up appears.
2. Select **Enable** or **Disable** and press Enter.

11.2.3 Configuring iBFT

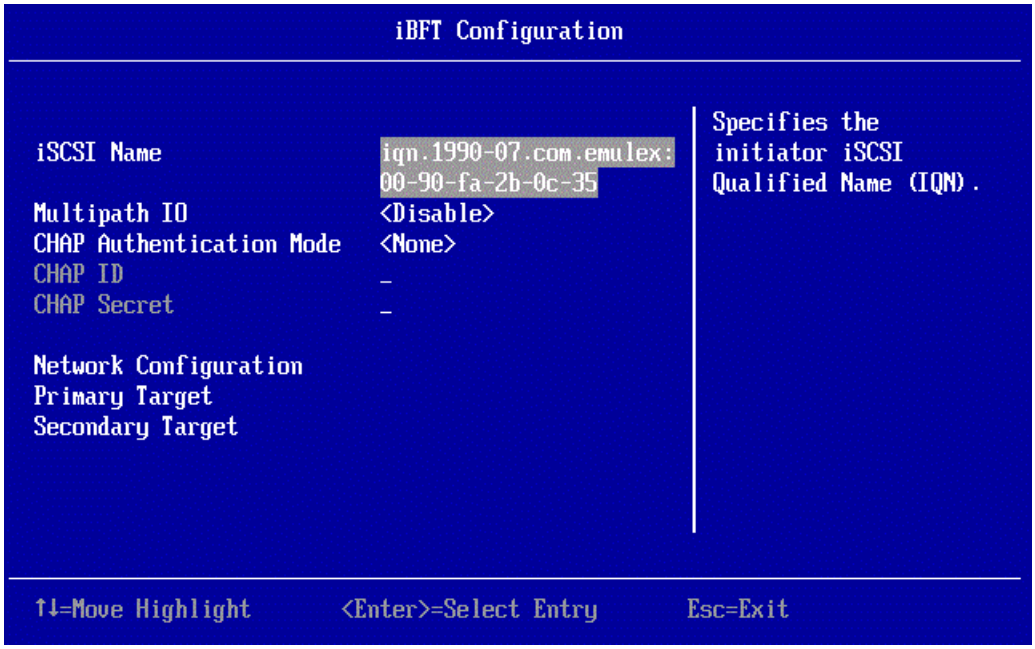
NOTE Boot mode must be set to Disable to access the iBFT Configuration menu.

Configuring iSCSI through the NIC interface (iBFT) is exactly the same as normal iSCSI configuration except it is performed through the NIC interface. For this to occur, you must use the operating system's NIC driver and an iBFT-aware operating system.

To configure iBFT:

1. On the Emulex NIC Selection dialog (Figure 79), select iBFT Configuration and press Enter.
   The iBFT Configuration dialog appears.

Figure 81 iBFT Configuration Dialog



The iBFT Configuration dialog enables you to perform the following tasks:

- Edit the iSCSI initiator name
- Enable multipath I/O
- Configure the CHAP authentication mode
- Configure the network settings
- Configure the primary or secondary targets

After completing the iBFT configuration, save your changes.

To save changes:

1. Press Esc to return to the Emulex NIC Selection dialog (Figure 79).
2. Select Save Changes and press Enter.

11.2.3.1 Changing the iSCSI Name

To edit the iSCSI initiator name:

1. On the iBFT Configuration dialog (Figure 81), select iSCSI Name and press Enter.

   This is a global setting. The iSCSI initiator name configured on one adapter is reflected on all adapters in the system.

   NOTE HP systems are an exception to the previous statement. Each port on an HP system can be configured with a unique initiator name.

11.2.3.2 Selecting MPIO

MPIO support allows the initiator to log in dual sessions to the same target. In this way I/O can be sent over either TCP/IP connection to the target. If one session fails, another session can continue processing I/O without interruption to the application. In iSCSI target configuration, you have the option of setting dual network paths to a single boot LUN.

NOTE Although MPIO boot support allows the initiator to log in to multiple sessions, the iSCSI BIOS currently limits the number of sessions to two iSCSI sessions for a single boot LUN.

You must follow these steps to configure MPIO boot support for each operating system.

1. Use the Emulex iSCSI configuration utility to configure the first path to a boot target.
2. Complete normal operating system installation.
3. Install third-party MPIO software for your operating system.
4. Configure the second path to a single boot LUN through the Emulex iSCSI configuration utility.

To configure MPIO:

1. On the iBFT Configuration dialog (Figure 81), select Multipath I/O and press Enter. A pop up appears.
2. Select **Enable** or **Disable** and press Enter.

11.2.3.3 Configuring CHAP Authentication Mode

Broadcom uses the CHAP authentication method to authenticate initiators and targets for added network security. By using a challenge and response security mechanism, CHAP periodically verifies the initiator's identity. This authentication method depends on a secret key known to the initiator and the target only. Although the authentication can be one-way, you can negotiate CHAP in both directions with the help of the same secret set for mutual authentication. You must make sure however, that what you configure on the target side matches the initiator side. Both one-way and mutual authentication are supported.

To configure the CHAP authentication method:

1. On the iBFT Configuration dialog (Figure 81), select CHAP Authentication Mode and press Enter. The system displays a pop up.
2. Select None, One-way CHAP, or Mutual CHAP and press Enter.
   - **One-way CHAP** – With one-way CHAP authentication, the target authenticates the initiator. Use one-way CHAP authentication for a one-way challenge and response security method – you must configure the CHAP ID and CHAP Secret, which is authenticated by the target.
   - **Mutual CHAP** – With mutual CHAP authentication, the target authenticates the initiator and the initiator authenticates the target. Use mutual CHAP authentication for a two-way challenge and response security method – you must configure different CHAP and mutual CHAP IDs and CHAP secrets.

   NOTE If you configure the CHAP parameters, verify that those parameters match the parameters on the storage side.

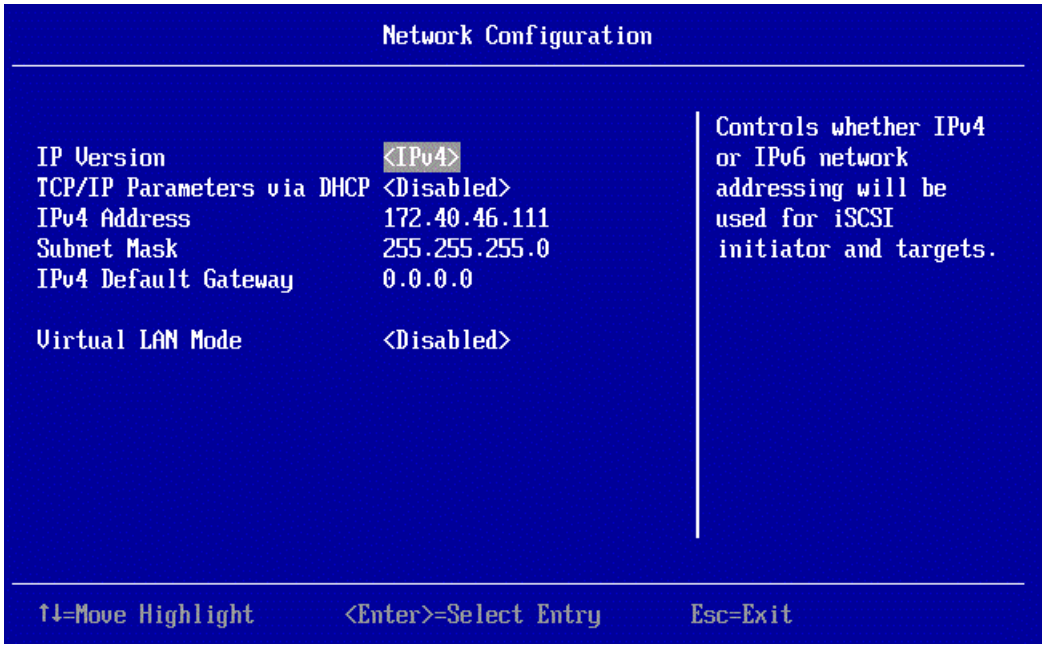
3. Enter the appropriate CHAP ID and CHAP Secret and press Enter.

11.2.3.4 Configuring Network Settings

To configure the network settings:

1. On the iBFT Configuration dialog (Figure 81), select Network Configuration and press Enter. The Network Configuration dialog (Figure 82) appears.

Figure 82 Network Configuration Dialog



The Network Configuration dialog enables you to perform the following tasks:

- Configure the initiator's IP version
- Configure TCP/IP parameters using DHCP

  NOTE This option is available only if IPv4 is selected as the IP version.
- Configure iSCSI parameters using DHCP

  NOTE This option is available only if TCP/IP Parameters via DHCP is enabled.
- Select the virtual LAN mode

After completing the network configuration, save your changes.

To save changes:

1. Press Esc to return to the iBFT Configuration dialog (Figure 81).
2. Press Esc again to return to the Emulex NIC Selection dialog (Figure 79).
3. Select Save Changes and press Enter.

11.2.3.4.1 Configuring the Initiator IP Version

To configure the initiator IP version:

1. On the Network Configuration dialog (Figure 82), select IP version and press Enter. The system displays a pop up.
2. Select IPv4, IPv6, or None and press Enter.

11.2.3.4.2 Configuring TCP/IP Parameters Using DHCP

NOTE This option is available only if IPv4 is selected as the IP Version.

This option controls the source of the initiator's IP address through DHCP or static assignment.

To configure TCP/IP parameters using DHCP:

1. On the Network Configuration dialog (Figure 82), select TCP/IP Parameters via DHCP and press Enter. A pop up appears.
2. Select **Enabled** or **Disabled** and press Enter.
3. If this setting is disabled, you must enter the following information for the iSCSI initiator:
   - IPv4 address
   - Subnet mask
   - IPv4 default gateway

11.2.3.4.3 Configuring iSCSI Parameters Using DHCP

NOTE This option is available only if TCP/IP Parameters via DHCP is enabled.

This option enables you to retrieve the iSCSI target parameters using DHCP.

To configure iSCSI parameters using DHCP:

1. On the Network Configuration dialog (Figure 82), select iSCSI Parameters via DHCP and press Enter. The system displays a pop up.
2. Select **Enabled** or **Disabled** and press Enter.

11.2.3.4.4 Selecting Virtual LAN Mode

If virtual LAN mode is enabled, a VLAN tag can be used by iSCSI boot.

To enable virtual LAN mode:

1. On the Network Configuration dialog (Figure 82), select Virtual LAN Mode and press Enter. The system displays a pop up.
2. Select **Enabled** and press Enter.

11.2.3.5 Configuring the Primary and Secondary Targets

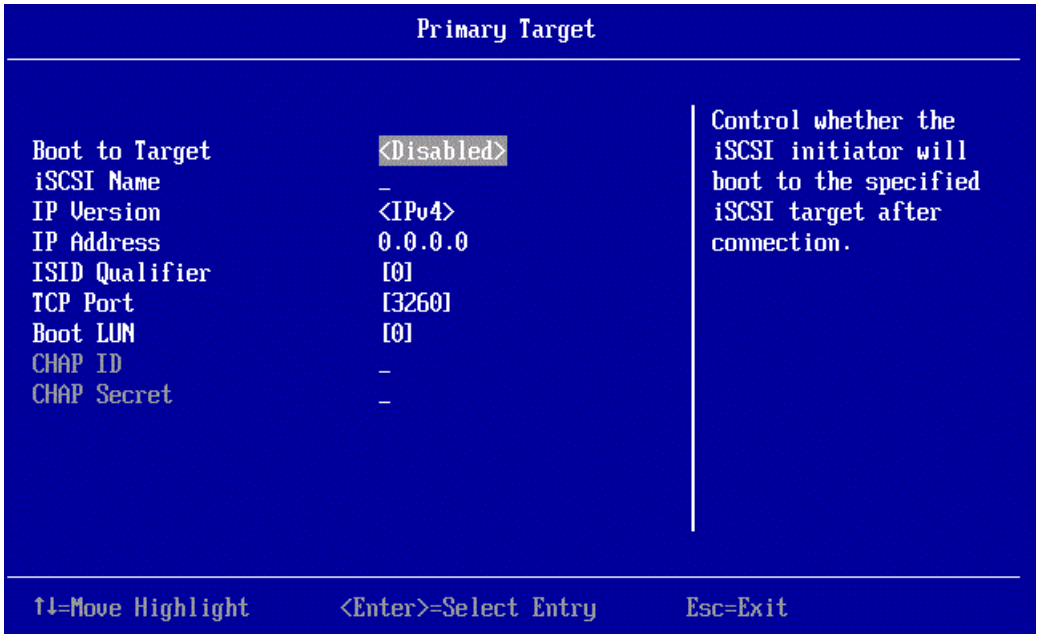
To configure the iSCSI target parameters:

1. On the iBFT Configuration dialog (Figure 81), select Primary Target or Secondary Target and press Enter.

   NOTE Configuration of the primary and secondary targets is identical. For this reason, this section covers the configuration of the primary target only.

   The Primary Target dialog (Figure 83) appears.

Figure 83 Primary Target Dialog



The Primary Target dialog enables you to configure the following items:

- Boot target
- iSCSI target name
- Target IP version
- Target IP address
- Target ISID qualifier
- Target TCP port
- Target boot LUN
- Target CHAP ID and secret

NOTE These options are available only if the initiator's CHAP authentication mode is properly set.

After completing the primary or secondary target configuration, you must save your changes.

To save changes:

1. Press Esc to return to the Emulex NIC Selection dialog (Figure 79).
2. Select Save Changes and press Enter.

11.2.3.5.1 Configuring the Boot Target

NOTE Only one boot target can be enabled at a time. If the primary target has the boot target setting enabled, the secondary target must have the boot target setting disabled.

To configure the boot target:

1. On the Primary Target dialog (Figure 83), select Boot to Target and press Enter. The system displays a pop up.
2. Select **Enabled** or **Disabled** and press Enter.

11.2.3.5.2 Changing the Target iSCSI Name

This setting specifies the IQN of the specified target.

To edit the iSCSI target name:

1. On the Primary Target dialog (Figure 83), select iSCSI Name and press Enter.
2. Edit the target name and press Enter.

11.2.3.5.3 Configuring the Target IP Version

To configure the target IP version:

1. On the Primary Target dialog (Figure 83), select IP version and press Enter. A pop up appears.
2. Select IPv4 or IPv6 and press Enter.

11.2.3.5.4 Configuring the Target IP Address

To configure the target IP address:

1. On the Primary Target dialog (Figure 83), select IP Address and press Enter.
2. Enter the target IP address and press Enter.

11.2.3.5.5 Configuring the Target ISID Qualifier

If you are connecting dual sessions to the same target portal group, enter a unique ISID Qualifier, which can be any number up to 65535. For more information on when an ISID value must be changed, see Section 5.3.2, Booting from SAN for iSCSI MPIO.

To configure the target ISID qualifier:

1. On the Primary Target dialog (Figure 83), select ISID Qualifier and press Enter.
2. Enter the target ISID qualifier value and press Enter.

11.2.3.5.6 Configuring the Target TCP Port

To configure the target TCP port:

1. On the Primary Target dialog (Figure 83), select TCP Port and press Enter.
2. Enter the target TCP port value and press Enter.

11.2.3.5.7 Configuring the Target Boot LUN

To configure the target boot LUN:

1. On the Primary Target dialog (Figure 83), select Boot LUN and press Enter.
2. Enter the target boot LUN value and press Enter.

11.2.3.5.8 Configuring the Target CHAP ID and Secret

NOTE These options are available only if the initiator's CHAP authentication mode has been properly set.

To configure the target CHAP ID and secret:

1. Ensure that the initiator's CHAP authentication mode is properly set.
2. On the Primary Target dialog (Figure 83), select CHAP ID and press Enter.
3. Enter the target CHAP ID value and press Enter.
4. Select CHAP Secret and press Enter.
5. Enter the target CHAP secret value and press Enter.

11.2.4 Configuring Personality

The *personality* reflects the protocol or protocols of the adapter. This option specifies a list of available protocols that can be configured on an adapter. Depending on the personalities for which the adapter is licensed, one or more of the following selections appears:

- **NIC** – The NIC personality implies that all the enabled functions provide NIC/TOE functionality.

  NOTE TOE functionality is not supported on Emulex OCe14000-series adapters.

- **iSCSI or FCoE** – These personalities are enabled on one function per adapter port and include NIC functionality on the other enabled functions. Only one of each storage protocol is allowed on each port.

  NOTE Changing the personality to iSCSI or FCoE on one physical port also changes the personality on all other ports from the same adapter.

- **NIC+RoCE** – For more information on the NIC+RoCE personality, see Section 11.2.4.1, Selecting the NIC+RoCE Personality.

  NOTE The NIC+RoCE personality is available only on Emulex OCe14000-series adapters running Windows or Linux with both UMC and SR-IOV disabled.

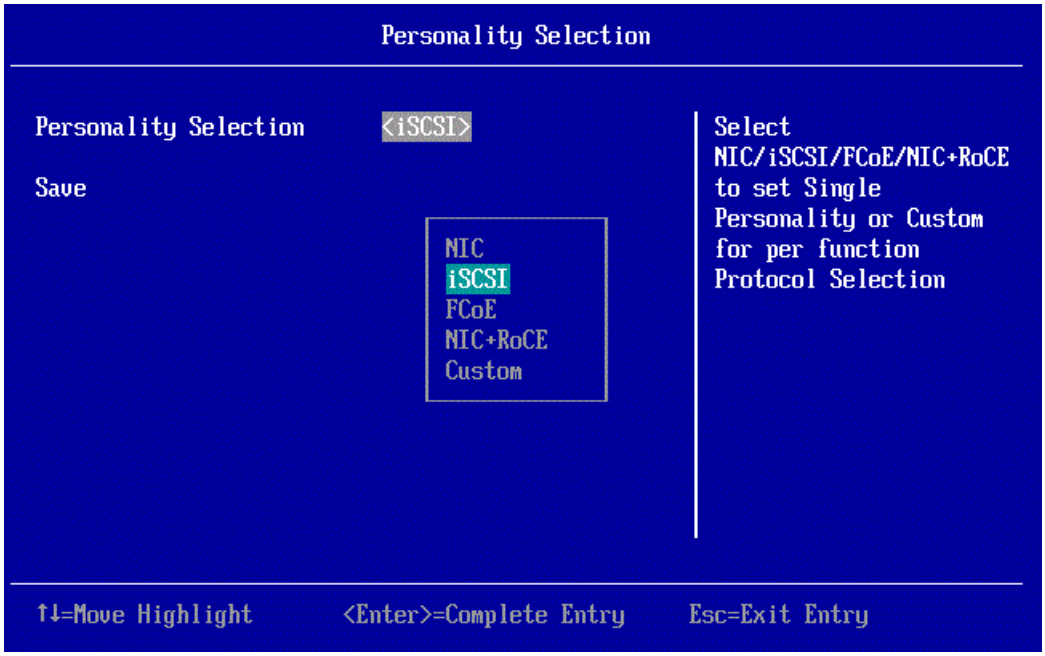
- **Custom (OCe14000-series adapters only)**– For more information on the Custom personality, see Section 11.2.4.2, Selecting the Custom Personality.

The menu displays the available personalities only, including both free and licensed personalities.

To view the personalities supported on the adapter:

1. Select Personality on the Emulex NIC Selection dialog (Figure 79), and then press Enter. The Personality Selection dialog appears.
2. Select the current personality setting and press Enter. A pop up appears with a list of available personalities.

Figure 84 Personality Selection Dialog with Pop Up



To change the personality of the adapter:

1. On the Personality Selection pop up (Figure 84) select the desired personality and press Enter.
2. Select Save and press Enter.

NOTE Additional configuration changes can be made after changing the Personality setting; however, a reboot is required before the new Personality setting will take effect.

11.2.4.1 Selecting the NIC+RoCE Personality

NOTE The NIC+RoCE personality is available only on Emulex OCe14000-series adapters running Windows or Linux with both UMC and SR-IOV disabled.

NOTE The RDMA application profile setting is available only when the NIC+RoCE personality is selected.

NOTE The RDMA application profile setting is available only when the Virtualization mode is set to None.

The NIC+RoCE personality implies that all the enabled functions provide RoCE functionality.

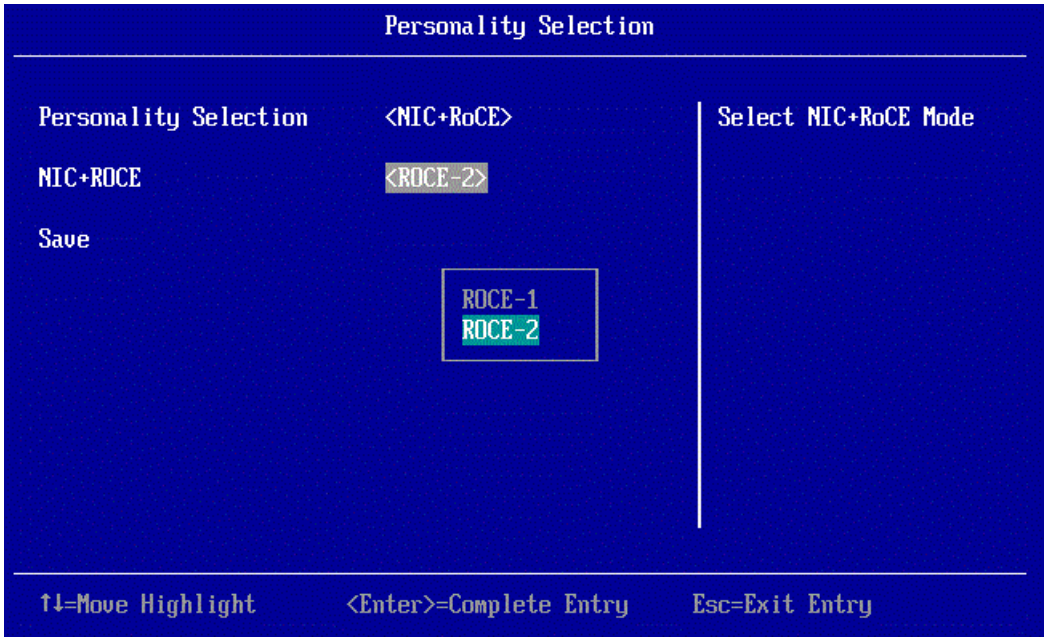
If NIC+RoCE is enabled, you must select a NIC+RoCE profile. The NIC+RoCE profile settings allows you to select the appropriate configuration for your particular system setup. The RDMA application profile choices include:

- **RoCE-1** – Not supported.
- **RoCE-2** – The RoCE-2 profile supports Windows SMB Direct, Linux iSER, NFS-RDMA, and VM Migration.

To select the NIC+RoCE personality:

1. From the Personality Selection dialog (Figure 84), select NIC+RoCE and press Enter. A pop up appears.

Figure 85 NIC+RoCE Profile Dialog with Pop Up



2. Select the desired NIC+RoCE profile and press Enter.
3. Select Save and press Enter.

NOTE Additional configuration changes can be made after changing the Personality setting; however, a reboot is required before the new Personality setting will take effect.

11.2.4.2 Selecting the Custom Personality

The Custom personality allows you to select the protocol type for each function. The iSCSI and FCoE personalities can be enabled on only one function per adapter port.

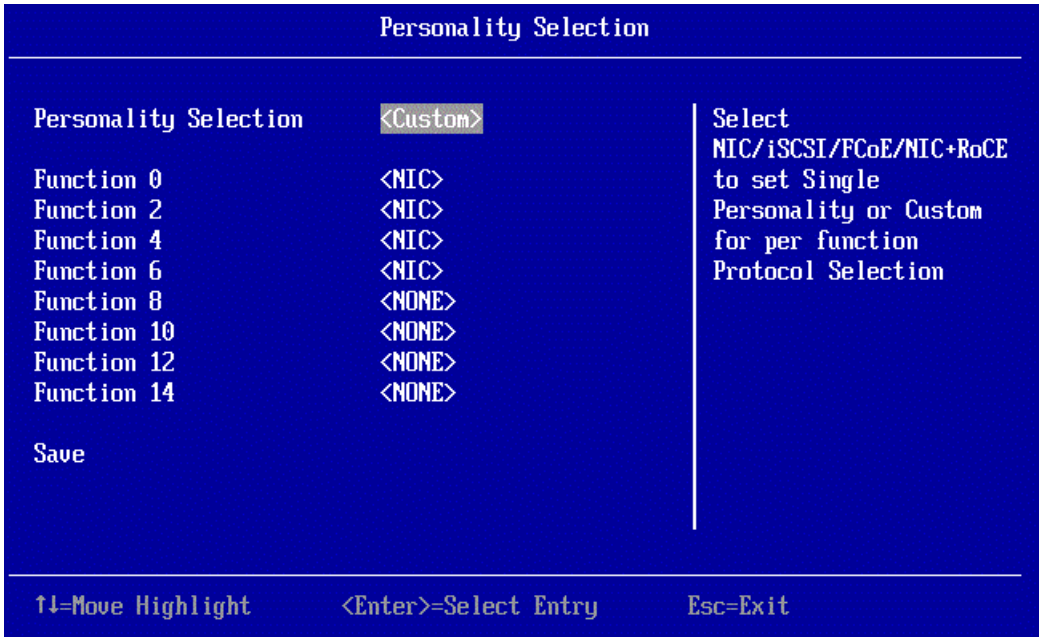
NOTE Two iSCSI functions and two FCoE functions are not allowed on a single port.

NOTE If higher order functions have been enabled and you want to disable them, you must first enable the Custom personality before the option to disable the functions appears.

To select the Custom personality:

1. From the Personality Selection dialog (Figure 84), select Custom and press Enter. The Custom Personality Selection dialog appears.

Figure 86 Custom Personality Selection Dialog



2. Select the desired personality for the appropriate function from the pop up and press Enter.
3. Select Save and press Enter.

NOTE Additional configuration changes can be made after changing the Personality setting; however, a reboot is required before the new Personality setting will take effect.

11.2.5 Link Reconfiguration

NOTE This option is available only if the adapter supports link reconfiguration.

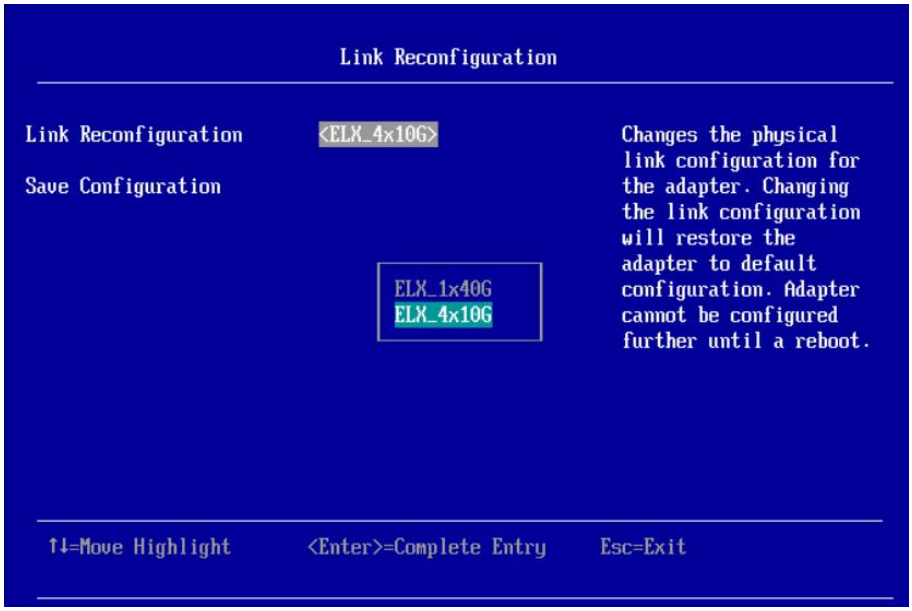
NOTE This feature is not supported on all Emulex OCe14000-series adapters.

The Link Reconfiguration option allows you to modify the physical link configuration for the adapter from a list of supported configurations.

To reconfigure the physical link for an adapter:

1. On the Emulex NIC Selection dialog (Figure 79), select Link Reconfiguration and press Enter.
   The Link Reconfiguration dialog appears.
2. Select the Link Reconfiguration setting and press Enter. A pop up appears with a list of available configurations.

Figure 87 Link Reconfiguration Dialog with Pop Up



3. Select the desired configuration and press Enter.
4. Select Save Configuration and press Enter.

Changing the link configuration restores the adapter to default configuration. The adapter cannot be configured further until a reboot.

11.2.6 Configuring Virtual LAN Settings

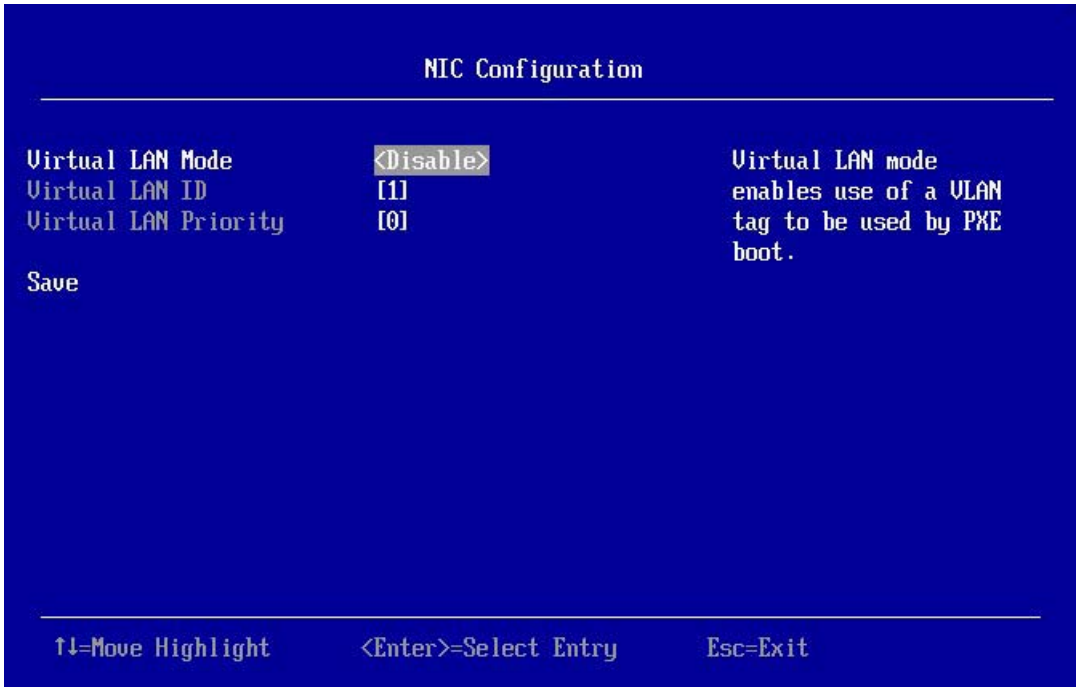
NOTE This feature is available only when Multichannel mode is disabled.

NOTE VLAN configuration is available only if the server's system BIOS does not provide the option to configure VLAN.

To view the VLAN settings for the adapter:

1. Verify that Multichannel mode is set to Disabled on the Emulex NIC Selection dialog (Figure 79).
2. Select NIC Configuration on the Emulex NIC Selection dialog (Figure 79), and then press Enter. The NIC Configuration dialog appears.

Figure 88 NIC Configuration Dialog



The NIC Configuration dialog enables you to perform the following tasks:

- Enable or disable VLAN mode
- Configure the VLAN ID
- Configure the VLAN priority

NOTE For FTS adapters, an additional menu option is available called PXE Boot. See Section 11.2.6.4, Enabling or Disabling PXE Boot, for more information.

For HP adapters in HP systems, additional menu options are available. These options include Legacy PXE Boot and HP Shared Memory Feature. For more information, see and Section 11.2.6.6, Configuring the HP Shared Memory Feature.

11.2.6.1 Enabling or Disabling VLAN Mode

To configure the VLAN mode:

1. On the NIC Configuration dialog (Figure 88), highlight the current Virtual LAN Mode setting.
2. Enable or disable the Virtual LAN Mode setting.
3. Select Save and press Enter.

11.2.6.2 Configuring the VLAN ID

NOTE Virtual LAN mode must be enabled to configure this feature.

To configure a VLAN ID:

1. On the NIC Configuration dialog (Figure 88), highlight the current Virtual LAN ID setting.
2. Select a number from 1 to 4094.
3. Select Save and press Enter.

11.2.6.3 Configuring the VLAN Priority

NOTE Virtual LAN mode must be enabled to configure this feature.

To set the VLAN priority level:

1. On the NIC Configuration dialog (Figure 88), highlight the current Virtual LAN Priority setting.
2. Select a number from 0 to 7.

   This unique value assigns a priority to outbound packets containing a specified VLAN ID. Valid values range from 0 to 7, with 0 being the highest priority level.
3. Select Save and press Enter.

11.2.6.4 Enabling or Disabling PXE Boot

NOTE This feature is available for FTS adapters only.

To configure PXE boot support:

1. On the NIC Configuration dialog (Figure 88), highlight the current PXE Boot setting.
2. Enable or disable the PXE Boot setting.
3. Select Save and press Enter.

NOTE If PXE boot is enabled, during system startup PXE contacts the DHCP server for an IP address to boot from the network.

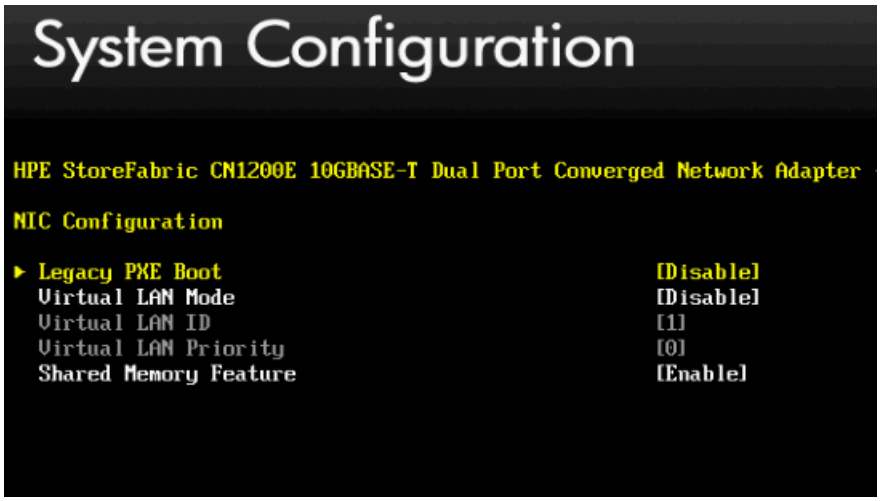
11.2.6.5 Configuring Legacy PXE Boot

NOTE This feature is available for HP adapters in HP systems only.

To configure legacy PXE boot support:

1. On the following NIC Configuration dialog, highlight the current Legacy PXE Boot setting.

Figure 89 NIC Configuration Dialog (HP Adapters Only)



2. Enable or disable the Legacy PXE Boot setting.
3. Press the F10 key to save the changes.

NOTE If Legacy PXE Boot is enabled, during system startup PXE contacts the DHCP server for an IP address to boot from the network.

11.2.6.6 Configuring the HP Shared Memory Feature

NOTE This feature is available for HP adapters in HP systems only.

When this setting is enabled, the system no longer uses reserved memory regions so devices can be directly assigned to guest virtual machines.

To set the HP Shared Memory Feature:

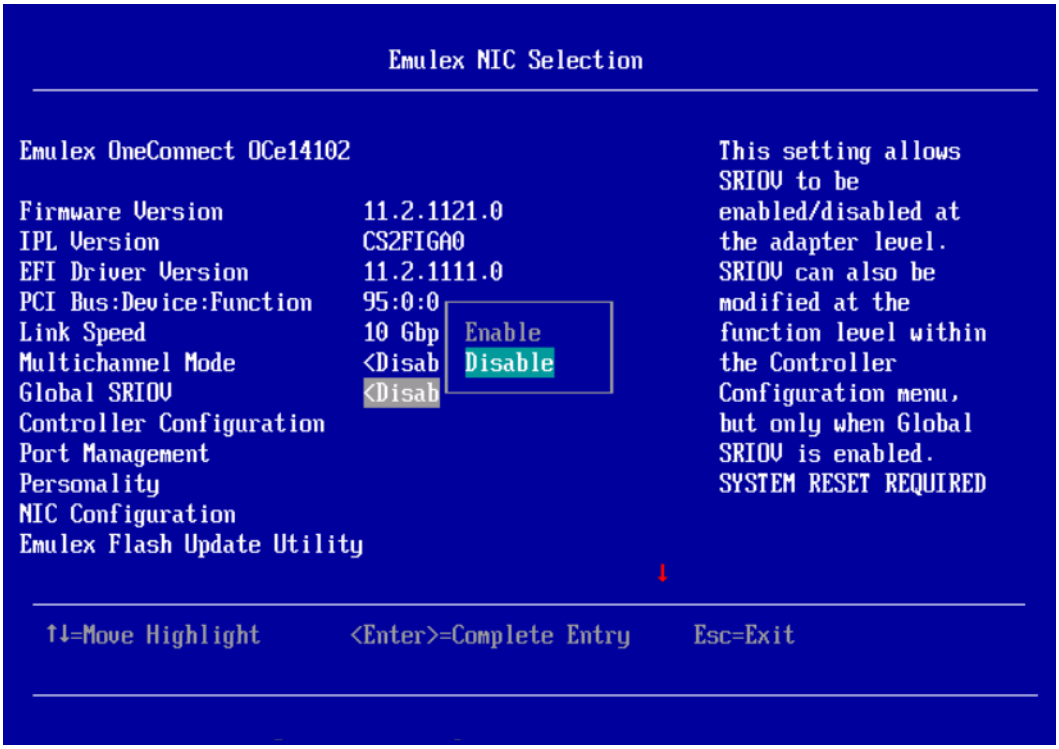
1. On the NIC Configuration dialog (Figure 89), highlight the current HP Shared Memory Feature setting.
2. Enable or disable the HP Shared Memory Feature.
3. Press the F10 key to save the changes.

11.2.7 Configuring the Adapter If UMC Is Disabled or Not Supported

NOTE If UMC is enabled on Emulex or HP adapters, see Section 11.2.8, Configuring Universal Multichannel. If you are using a Lenovo System x adapter, see Section 11.2.9, Configuring Multichannel for Lenovo System x Adapters.

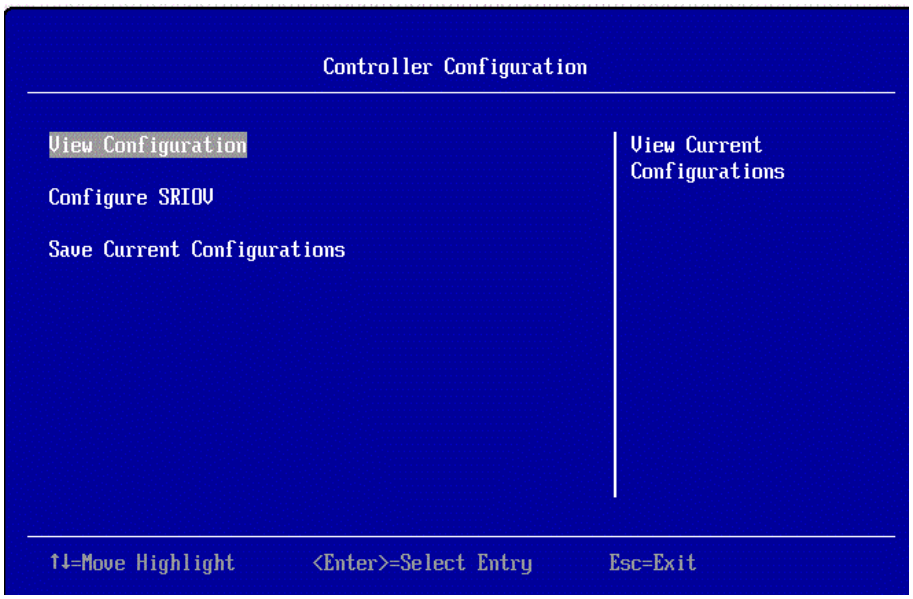
The Global SRIOV setting enables and disables SR-IOV on an adapter. When SR-IOV is enabled at the adapter level, it can be turned off individually for each function if required. Disabling SR-IOV on all functions disables Global SRIOV. The value of this setting is determined by the collective SR-IOV states of all functions on the adapter. If SR-IOV is enabled on at least one function, then Global SRIOV will be set to enabled.

Figure 90 Global SRIOV Dialog



- If UMC is disabled or is not supported, and Global SRIOV is enabled, the following Controller Configuration menu (Figure 91) appears when you select Controller Configuration on the Emulex NIC Selection dialog (Figure 79).

Figure 91 Controller Configuration Menu (UMC Disabled)



The Controller Configuration menu enables you to perform the following tasks:

- View the current configuration
- Configure SR-IOV
- Configure NIC mode

NOTE This setting is available only if SR-IOV is enabled, the Personality Selection is set to NIC, and a two-port Emulex OCe14000-series adapter is in use.

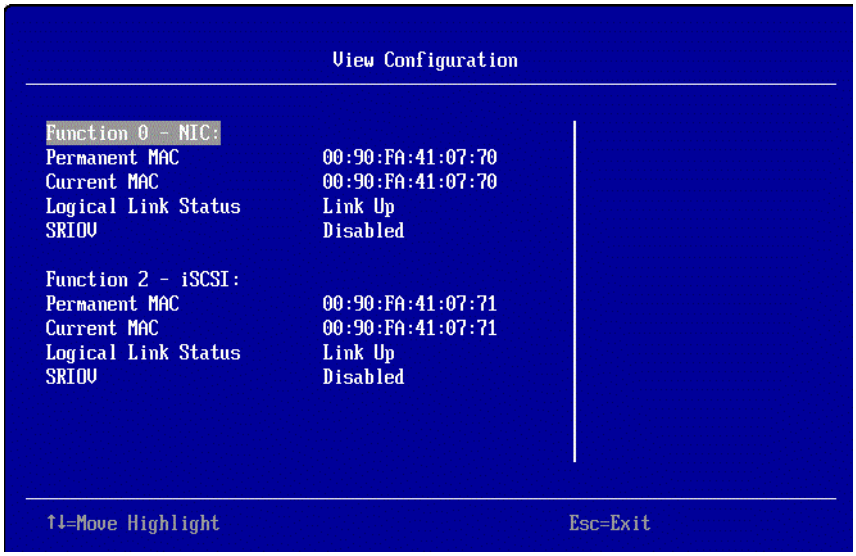
- Save the current configuration

11.2.7.1 Viewing the Configuration

To view the current configuration:

1. On the Controller Configuration menu (Figure 91), select View Configuration and press Enter. The View Configuration information appears.

Figure 92 View Configuration



From the View Configuration screen, you can view information for that specific function, including the permanent and current MAC addresses, the logical link status, and the SR-IOV status for NIC functions.

2. When you are finished viewing the configurations, press Esc to return to the Controller Configuration menu.

11.2.7.2 Configuring SR-IOV

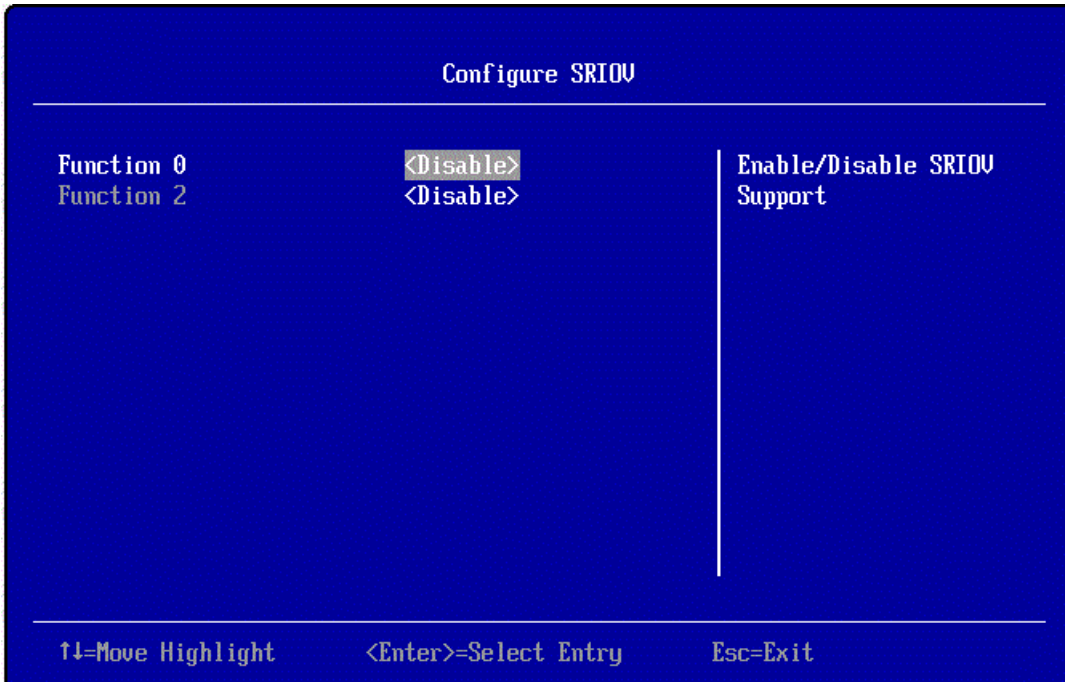
If the system BIOS supports SR-IOV, you can enable it. SR-IOV support can only be enabled if UMC is disabled or it is not supported.

NOTE SR-IOV is not supported with the NIC+RoCE personality.
SR-IOV is only available for NIC functions and not for storage (iSCSI or FCoE) functions.

To configure SR-IOV:

1. On the NIC Configuration dialog (Figure 88), set the Global SRIOV parameter to Enabled.
2. On the Controller Configuration menu (Figure 91), select Configure SRIOV and press Enter. The Configure SRIOV dialog appears.

Figure 93 Configure SRIOV Dialog



3. Select the current setting for a particular NIC function and a pop up appears.
4. Select **Enable** or **Disable** and press Enter.
5. When you are finished, press Esc to return to the Controller Configuration menu.

For more information on SR-IOV configuration, see the appropriate Emulex driver's user guide.

11.2.7.3 Configuring NIC Mode

NOTE This setting is available only if SR-IOV is enabled, the Personality Selection is set to NIC, and a two-port Emulex OCe14000-series adapter is in use.

This setting allows you to control the VF count on a two-port Emulex OCe14000-series adapter. The choices include:

- **NIC Only Profile** – The VF count is restricted to 31 VFs per port and QoS is supported for those VFs.

  NOTE For two-port Emulex OCe14000B-series adapters only, this setting supports up to 63 VFs per port.

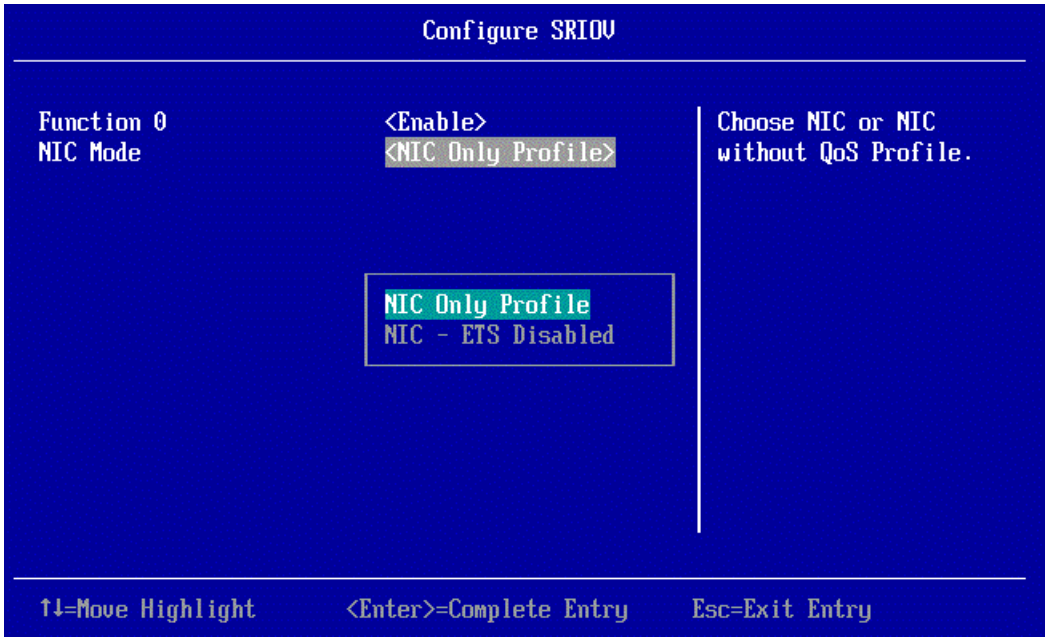
- **NIC-ETS Disabled** – The VF count is set to 63 VFs per port but QoS is not supported for those VFs.

  NOTE This setting is not supported on Emulex OCe14000B-series adapters.

To select the NIC mode:

1. From the Configure SRIOV dialog (Figure 93), select the NIC Mode setting for a particular NIC function and a pop up appears (Figure 94).

Figure 94 Configure NIC Mode Pop Up



2. Select NIC Only Profile or NIC-ETS Disabled and press Enter.

NOTE For two-port Emulex OCe14000B-series adapters only, the NIC Only Profile setting supports up to 63 VFs per port.

The NIC-ETS Disabled setting is not supported on Emulex OCe14000B-series adapters.

3. When you are finished, press Esc to return to the Controller Configuration menu.

11.2.7.4 Saving the Current Configuration

To save the configuration changes, on the Controller Configuration menu (Figure 91), select Save Current Configurations and press Enter.

11.2.8 Configuring Universal Multichannel

UMC provides the ability to configure multiple PCI functions or I/O channels for each adapter port. For more information on UMC support, refer to the *Emulex Universal Multichannel Reference Manual*.

NOTE HP adapters support UMC, FLEX-10, and FLEX-QBG mode.

If UMC is enabled and the adapter also supports Virtual Fabric mode or Unified Fabric Protocol mode, see Section 11.2.9, Configuring Multichannel for Lenovo System x Adapters, for more information. If UMC is disabled or not supported on the adapter, see Section 11.2.7, Configuring the Adapter If UMC Is Disabled or Not Supported.

If UMC is enabled, you must configure the minimum and maximum bandwidths settings for the iSCSI and FCoE storage functions in the NIC BIOS before they can be configured through their respective utilities (iSCSI or FCoE BIOS). Otherwise, the Logical Link for that function will still show as down, and you will be unable to log in to targets or find LUNs behind those targets.

UMC functionality is supported on Emulex OneConnect adapters running in 10GbE mode only. UMC is not supported in 1GbE mode.

If UMC functionality is supported on the system, the Emulex UEFI NIC utility enables you to perform the following tasks:

- Enable or disable UMC functionality (Multichannel mode)
- View the current adapter configuration
- Configure the minimum and maximum bandwidth for each channel
- Configure the LPVID for each channel

NOTE Your system may not support all UMC options.

For Emulex OCe14000-series adapters, the Emulex NIC configuration utility displays all of the functions on a port, regardless of whether the system supports ARI or not. For more information on ARI, see Section 11.2.8.1, ARI Support.

To enable UMC support:

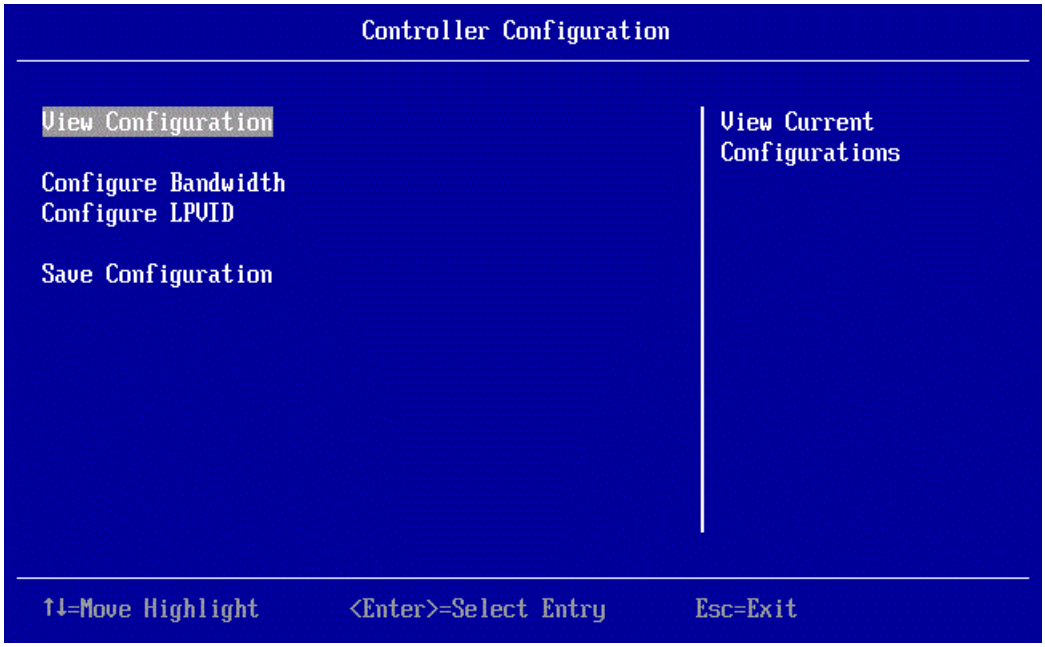
1. On the Emulex NIC Selection dialog (Figure 79), select Multichannel Mode and press Enter. A pop up appears.
2. Select **Enable** to activate UMC support and press Enter.

NOTE HP FLEX mode cannot be enabled in this menu. If an adapter is in FLEX mode, the Multichannel mode is displayed as either FLEX-10 or FLEX-QBG, as appropriate. Some settings, such as Multichannel and Personality, and the Multichannel parameters, cannot be configured from the UEFI BIOS utility in FLEX mode.

To configure the UMC configuration:

1. On the Emulex NIC Selection dialog (Figure 79), select Controller Configuration and press Enter. The following Controller Configuration menu appears.

Figure 95 Controller Configuration Menu (UMC Enabled)



The Controller Configuration menu enables you to perform the following tasks:

- View the current configuration of each function.
- Configure the minimum and maximum bandwidth percentage.
- Configure the LPVID per function.
- Save the current configuration.

11.2.8.1 ARI Support

The Emulex OCe14000-series adapters support ARI, which provides the ability to configure additional PCI functions or I/O channels for each physical adapter port if UMC is enabled.

You can configure up to 16 functions on a one-port Emulex OCe14400-series adapter, up to 8 functions per port on a one or two-port Emulex OCe14100-series adapter, and up to 4 functions per port on a four-port Emulex OCe14100-series adapter.

The maximum number of functions allowed on an adapter is controlled by the adapter's IPL file and the system's support for ARI.

The following requirements must be met to support more than eight functions on an adapter.

- The system hardware (the motherboard and BIOS) must support ARI.
- ARI must be enabled in the system BIOS.
- The host operating system must support ARI:
  - Windows Server 2012 and newer versions
  - RHEL 5.9 and newer versions
  - RHEL 6.4 and newer versions
  - SLES 11 SP2 and newer versions
  - ESXi 5.1 and newer versions
- The application management tools, including the Emulex OneCommand CNA Manager application must support ARI.
- ARI must be enabled in the firmware using the Emulex OneCommand CNA Manager application.

If these conditions are not met, you may be able to configure more than eight functions, but only up to eight functions will be running and discovered after a reboot.

11.2.8.2 Configuring More than Four Functions per Port on Emulex OCe14000-Series Adapters

By default on 1-port or 2-port Emulex OCe14000-series adapters, only four functions are enabled when UMC is enabled. To enable the other four functions, you must manually enable those functions for them to be operational.

To configure more than four functions per port on 1-port or 2-port Emulex OCe14000-series adapters, you must perform the following steps:

1. Ensure that UMC is enabled. By default, UMC enables four functions per port. See Section 11.2.8, Configuring Universal Multichannel, for information on configuring UMC.
2. Select the Custom mode in the Personality Selection dialog (Figure 84). This enables you to view all eight functions per port. See Section 11.2.4.2, Selecting the Custom Personality, for information on selecting the Custom mode.

   NOTE The options to enable and configure more than four functions per physical port are not available unless Custom mode is selected.

3. For each function to be enabled, change the function's personality type to the desired protocol. See Section 11.2.4.2, Selecting the Custom Personality, for information on selecting the personality.

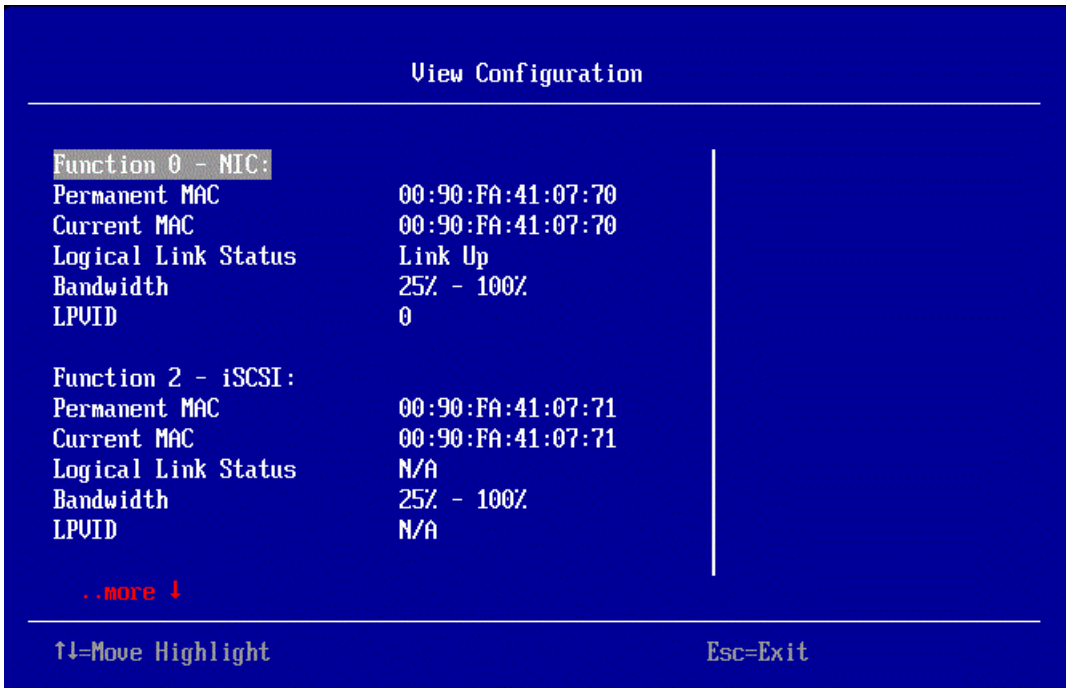
4. Assign the appropriate bandwidth and LPVID. See Section 11.2.8.4, Configuring Minimum and Maximum Bandwidths, for information on configuring the bandwidth and Section 11.2.8.5, Configuring the LPVID, for information on configuring the LPVID.

11.2.8.3 Viewing the Configuration

To view the current UMC configuration:

1. On the Controller Configuration menu (Figure 95), select View Configuration and press Enter. The View Configuration information appears.

Figure 96 UMC Configuration for Function 0



From the View Configuration display, you can view information for that specific function, including the permanent and current MAC addresses, the logical link status, the minimum and maximum bandwidth settings, and LPVID of all of the NIC functions.

NOTE Logical Link Status and LPVID for storage functions (iSCSI and FCoE) appears as N/A.

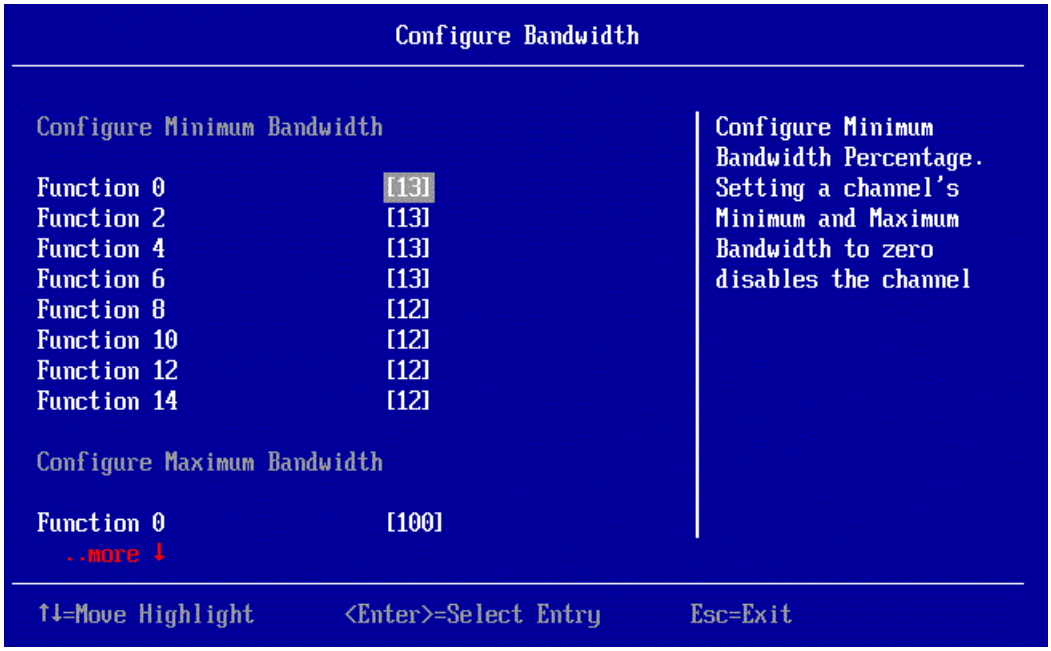
2. When you are finished viewing the configurations, press Esc to return to the Controller Configuration menu.

11.2.8.4 Configuring Minimum and Maximum Bandwidths

To configure bandwidth:

1. On the Controller Configuration menu (Figure 95), select Configure Bandwidth and press Enter. The Configure Bandwidth dialog (Figure 97) appears.

Figure 97 Configure Bandwidth Dialog



2. Select a specific function under Configure Minimum Bandwidth or Configure Maximum Bandwidth and press Enter.
   - The Minimum Bandwidth value is the least amount of bandwidth the function can provide. It is represented as a percentage. The Minimum Bandwidth value must be less than or equal to the Maximum Bandwidth value. The total of the Minimum Bandwidth values for all enabled functions on that port must be equal to 100.

     NOTE A Minimum Bandwidth value of 0 is a valid value. If all of the partitions' Minimum Bandwidth values are zero, the bandwidth is distributed equally among the current active partitions. If a specific partition's Minimum Bandwidth and Maximum Bandwidth values are 0, that partition's logical link will be brought down.
   - The Maximum Bandwidth value is the greatest amount of bandwidth the function can provide. It is represented as a percentage.
3. Enter the value for the specified function and press Enter.

   NOTE To disable a function, set the Minimum Bandwidth and Maximum Bandwidth values to 0.
4. When you are finished configuring bandwidths, press Esc to return to the Controller Configuration menu.

11.2.8.5 Configuring the LPVID

The LPVID is used to enforce a VLAN ID on all traffic originating from an IP address, channel, or PCI function. If the operating system for that PCI function has set up a VLAN ID, the operating system-configured VLAN ID takes precedence over the LPVID for transmit packets while the operating system-configured VLAN ID and LPVID-tagged packets will both be received. If the operating system has not set up any VLAN IDs, then the LPVID is used for tagging.

NOTE LPVID and user-configured VLAN IDs from the operating system must be different.

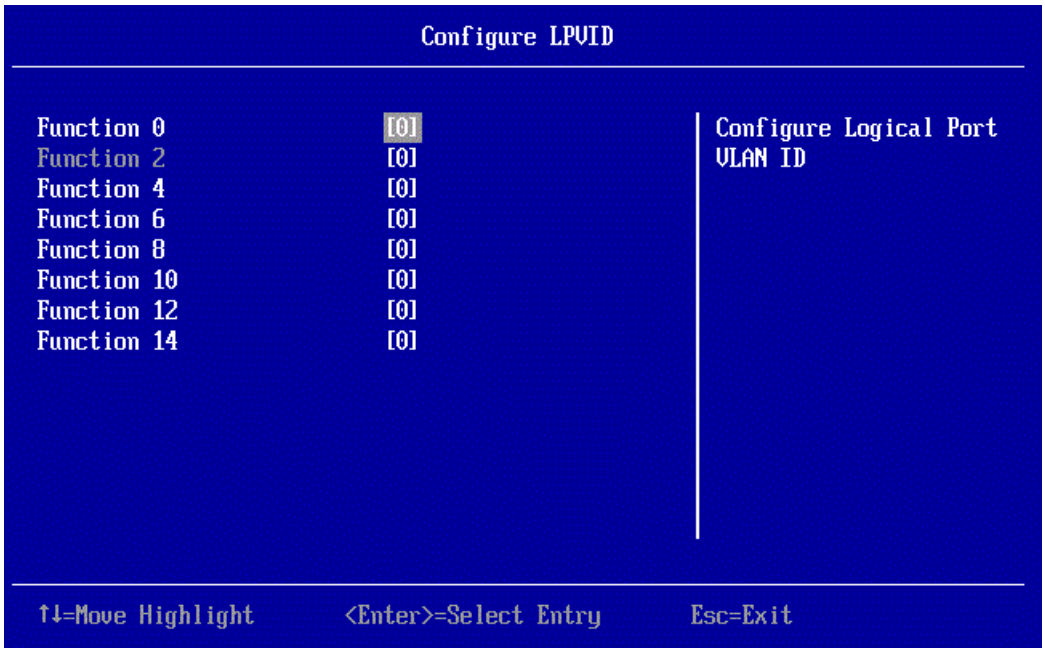
LPVIDs also need to be configured on the switch port.

Each LPVID must be unique and is relevant for NIC traffic only. The LPVID is not supported for storage functions. For iSCSI storage functions, you must configure a VLAN ID through iSCSISelect or through the host. For more information, see Section 9.3.2, Configuring VLAN ID and VLAN Priority.

To configure an LPVID:

1. On the Controller Configuration menu (Figure 95), select Configure LPVID and press Enter. The Configure LPVID dialog appears.

Figure 98 Configure LPVID Dialog



2. Select the LPVID value for a specific function and press Enter.
3. Enter a value and press Enter. The LPVID range is 2 to 4094. A value of 0 disables the LPVID.

NOTE LPVID values 1 and 4095 are currently reserved and cannot be used or configured.

4. When you are finished configuring LPVIDs, press Esc to return to the Controller Configuration menu.

11.2.8.6 Saving the Configuration

To save the configuration changes, on the Controller Configuration menu (Figure 95), select Save Configuration and press Enter.

11.2.9 Configuring Multichannel for Lenovo System x Adapters

NOTE This section only pertains to certain Emulex OneConnect adapters that support Multichannel for Lenovo System x Virtual Fabric or UFP.

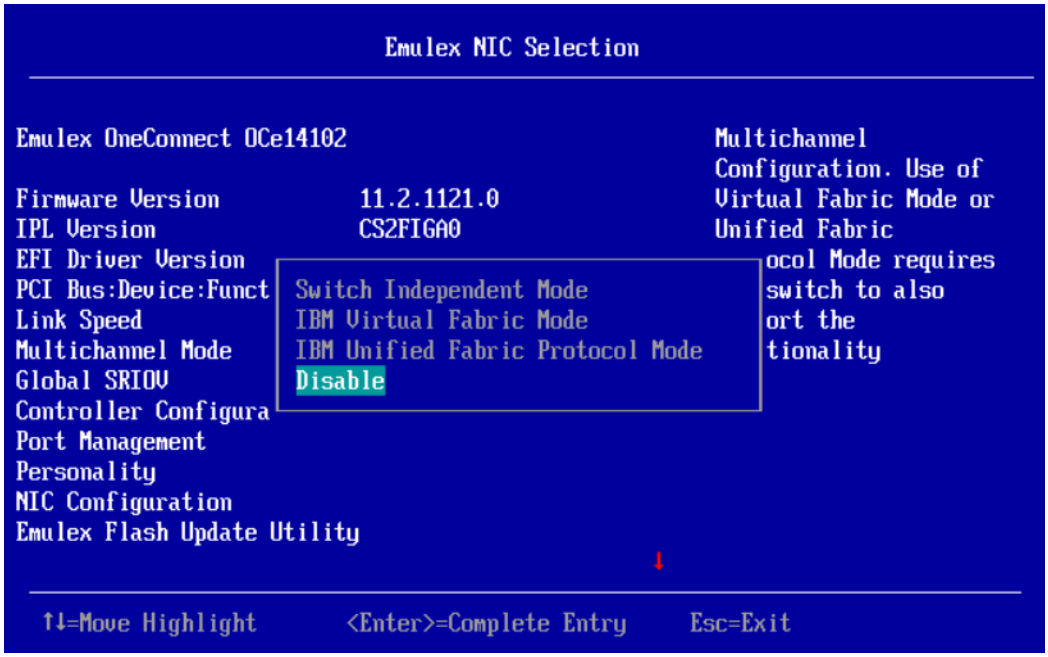
Multichannel for Lenovo System x provides the ability to configure multiple PCI functions or I/O channels for each adapter port.

NOTE Setting up Multichannel for Lenovo System x Virtual Fabric or UFP depends on cooperation with adjacent switches.

To enable Multichannel for Lenovo System x support:

1. On the Emulex NIC Selection dialog (Figure 99), select Multichannel Mode and press Enter. A pop up appears.

Figure 99 Lenovo System x Multichannel Mode Selection Pop Up



2. Select one of the following settings and press Enter:
   - Switch Independent Mode
   - Virtual Fabric Mode (referred to as IBM Virtual Fabric Mode on some Emulex OCe14000-series adapters and all Emulex OCe11100-series adapters)
   - Unified Fabric Protocol Mode (referred to as IBM Unified Fabric Protocol Mode on some Emulex OCe14000-series adapters and all Emulex OCe11100-series adapters)
   - Disabled

NOTE For more information on the available Multichannel modes, see Section 11.2.9.2, Multichannel Modes.

With the Multichannel mode selected, you can now configure Multichannel for Lenovo System x support. See Section 11.2.9.3, Multichannel for Lenovo System x Configuration, for more information.

11.2.9.1 Configuring More than Four vNICs per Port on Emulex OCe14000-Series Adapters

By default on 1-port or 2-port Emulex OCe14000-series adapters, only four vNICs are enabled when Multichannel is enabled. To enable the other four vNICs, you must manually enable those vNICs for them to be functional.

To configure more than four vNICs per port, you must perform the following steps:

1. Verify that Multichannel is enabled. By default, Multichannel enables four vNICs per port. See Section 11.2.9, Configuring Multichannel for Lenovo System x Adapters, for information on configuring Multichannel.
2. Select the Custom mode in the Personality Selection dialog (Figure 84). This enables you to view all eight vNICs per port. See Section 11.2.4.2, Selecting the Custom Personality, for information on selecting the Custom mode.
3. For each vNIC to be enabled, change the personality type to the desired protocol. See Section 11.2.4.2, Selecting the Custom Personality, for information on selecting the personality.

4. Assign the appropriate bandwidth and LPVID. See Section 11.2.8.4, Configuring Minimum and Maximum Bandwidths, for information on configuring the bandwidth and Section 11.2.8.5, Configuring the LPVID, for information on configuring the LPVID.

11.2.9.2 Multichannel Modes

NOTE The following modes are available only on certain adapters and systems that support Virtual Fabric mode and Unified Fabric Protocol mode.

Additional PFs are supported when ARI is enabled on the system. For more information on ARI, see Section 11.2.8.1, ARI Support.

A Lenovo System x Virtual Fabric-enabled switch provides the ability to configure an LPVID for a virtual channel or I/O channel on an adapter port. If Multichannel for Lenovo System x is supported on the system, you can select one of the following modes:

- **Virtual Fabric Mode** – Select this mode if an Emulex OneConnect adapter is attached to a Lenovo System x Virtual Fabric-enabled switch. This mode supports up to four PFs per port.

  NOTE This mode is referred to as IBM Virtual Fabric mode on some Emulex OCe14000-series adapters and all Emulex OCe11100-series adapters.

  Virtual Fabric mode is not supported on 40GbE switches.

- **Unified Fabric Protocol Mode** – Select this mode if an Emulex OneConnect adapter is attached to a Lenovo System x UFP-enabled switch. In this mode, the maximum PFs per port for each adapter configuration is listed below:
  - **2-port 10GbE adapters** – 8 PFs per port
  - **4-port 10GbE adapters** – 4 PFs per port
  - **1-port 40GbE adapters** – 16 PFs per port

  NOTE This mode is referred to as IBM Unified Fabric Protocol mode on some Emulex OCe14000-series adapters and all Emulex OCe11100-series adapters.

  Some Lenovo System x switches support both Unified Fabric Protocol mode and Virtual Fabric mode.

- **Switch Independent Mode** – Select this mode if you are using a switch other than a Lenovo System x Virtual Fabric or UFP-enabled switch. In this mode, the maximum PFs per port for each adapter configuration is listed below:
  - **2-port 10GbE adapters** – 8 PFs per port
  - **4-port 10GbE adapters** – 4 PFs per port
  - **1-port 40GbE adapters** – 16 PFs per port

  NOTE Multichannel for Lenovo System x functionality is supported on Emulex OneConnect adapters running in 10GbE mode only. Multichannel for Lenovo System x is not supported in 1GbE mode.

  An LPVID is optional for Virtual Fabric mode, but it is required for every function if using Switch Independent mode.

11.2.9.3 Multichannel for Lenovo System x Configuration

Multichannel for Lenovo System x provides the ability to configure multiple PCI functions or I/O channels for each physical adapter port.

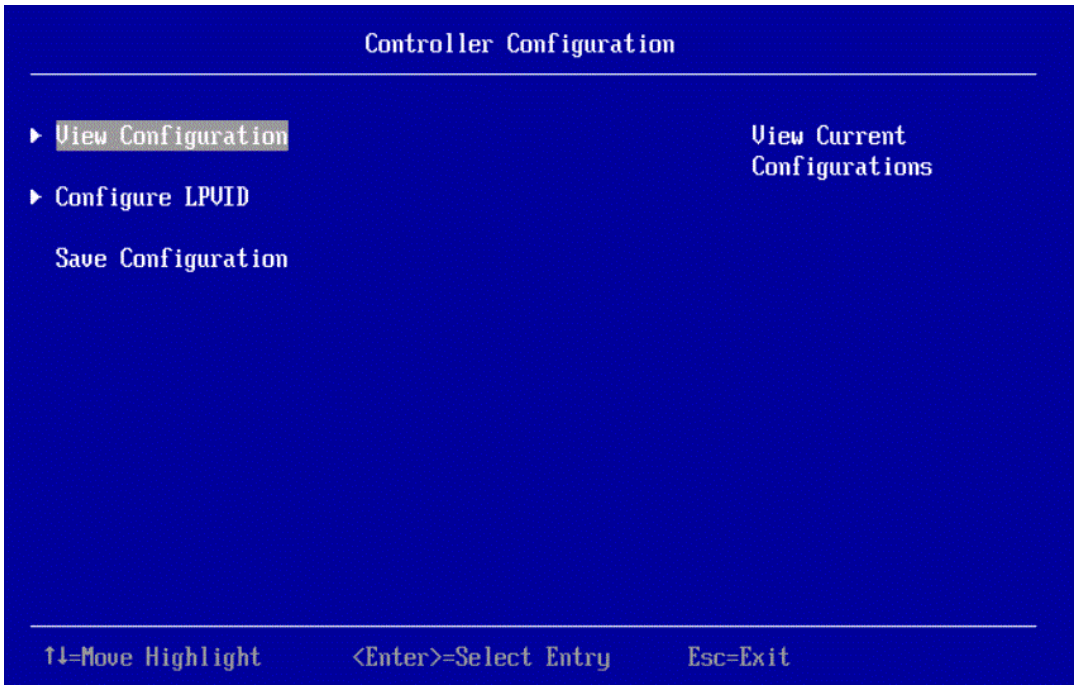
See the appropriate section below based on the mode you have selected.

11.2.9.3.1 Virtual Fabric Mode Configuration

To configure Multichannel for Lenovo System x if the Multichannel mode is set to Virtual Fabric mode:

1. On the Emulex NIC Selection dialog (Figure 99), select Controller Configuration and press Enter. The following Controller Configuration menu appears.

Figure 100 Virtual Fabric Mode Controller Configuration Menu



The Controller Configuration menu enables you to perform the following tasks:

- **View the current configuration of each function** – See Section 11.2.8.3, Viewing the Configuration.
- Configure the LPVID per function – See Section 11.2.8.5, Configuring the LPVID.
- Save the current configuration – See Section 11.2.8.6, Saving the Configuration.

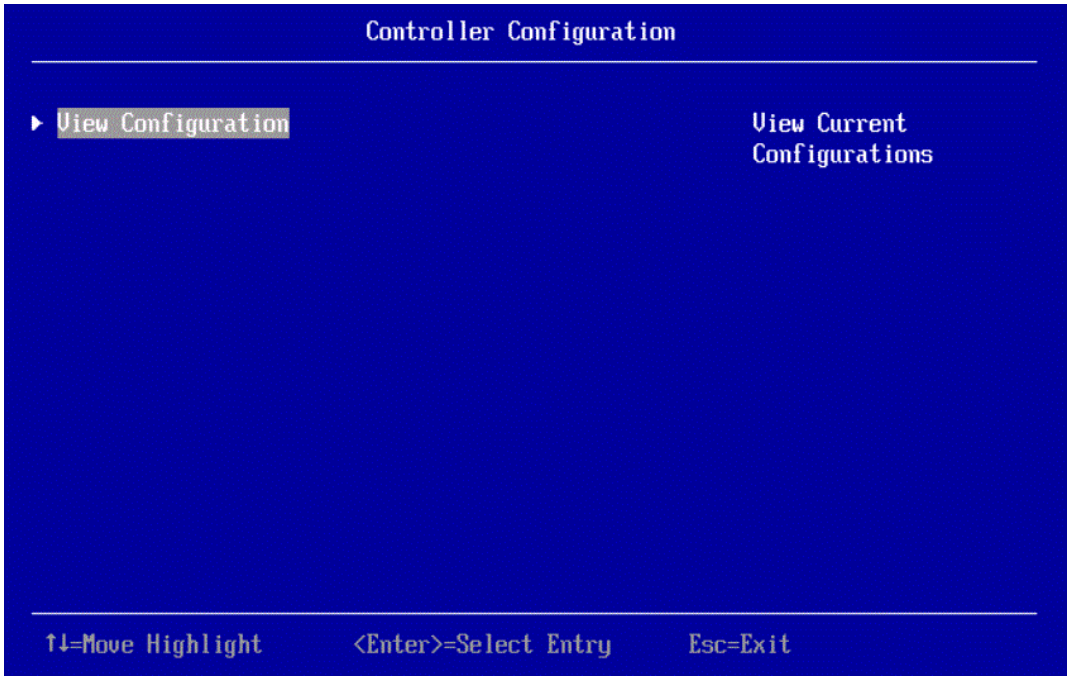
11.2.9.3.2 Unified Fabric Protocol Mode Configuration

NOTE By default on 1-port or 2-port Emulex OCe14000-series adapters, only four vNICs are enabled when Multichannel is enabled. To enable the other four vNICs, you must manually enable those vNICs using the steps provided in Section 11.2.9.1, Configuring More than Four vNICs per Port on Emulex OCe14000-Series Adapters.

To configure Multichannel for Lenovo System x if the Multichannel mode is set to Unified Fabric Protocol mode:

1. On the Emulex NIC Selection dialog (Figure 99), select Controller Configuration and press Enter. The Controller Configuration menu (Figure 101) appears.

Figure 101 UFP Controller Configuration Menu



From the Controller Configuration menu, you can view the current configuration of each function. See Section 11.2.8.3, Viewing the Configuration, for more information.

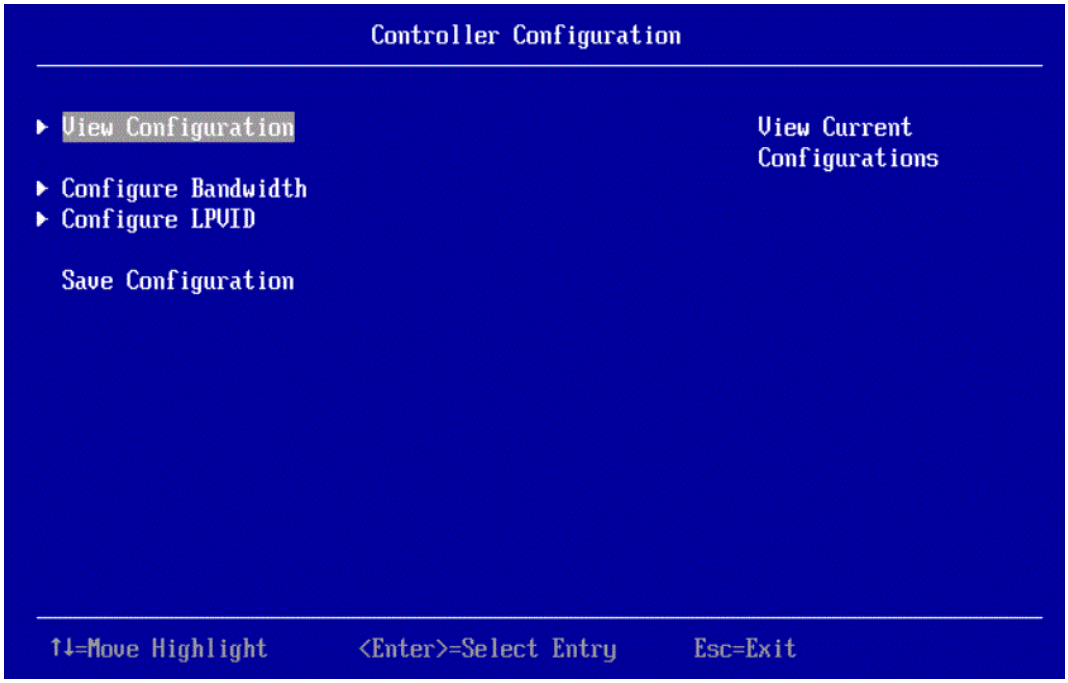
11.2.9.3.3 Switch Independent Mode Configuration

NOTE By default on 1-port or 2-port Emulex OCe14000-series adapters, only four vNICs are enabled when Multichannel is enabled. To enable the other four vNICs, you must manually enable those vNICs using the steps provided in Section 11.2.9.1, Configuring More than Four vNICs per Port on Emulex OCe14000-Series Adapters.

To configure Multichannel for Lenovo System x if the Multichannel mode is set to Switch Independent mode:

1. On the Emulex NIC Selection dialog (Figure 99), select Controller Configuration and press Enter. The following Controller Configuration menu (Figure 102) appears.

Figure 102 Switch Independent Mode Controller Configuration Menu



From the Controller Configuration menu, you can perform the following tasks:

- View the current configuration of each function – See Section 11.2.8.3, Viewing the Configuration.
- Configure the minimum and maximum bandwidth percentage – See Section 11.2.8.4, Configuring Minimum and Maximum Bandwidths.
- Configure the LPVID per function – See Section 11.2.8.5, Configuring the LPVID.

  NOTE An LPVID is required for every function if using Switch Independent mode, but it is optional for Virtual Fabric mode and Unified Fabric Protocol mode.
- Save the current configuration – See Section 11.2.8.6, Saving the Configuration.

11.2.10 Port Management

The Port Management option is used to enable or disable the physical ports.

If port 0 is disabled, it brings down the physical link for port 0 and the power consumption of the adapter is lowered; however, the PCI functions associated with that port must remain enabled.

If any other port is disabled, all of the corresponding functions associated with that port are disabled and removed from the PCI configuration space. Disabling these ports also lowers the power consumption of the adapter.

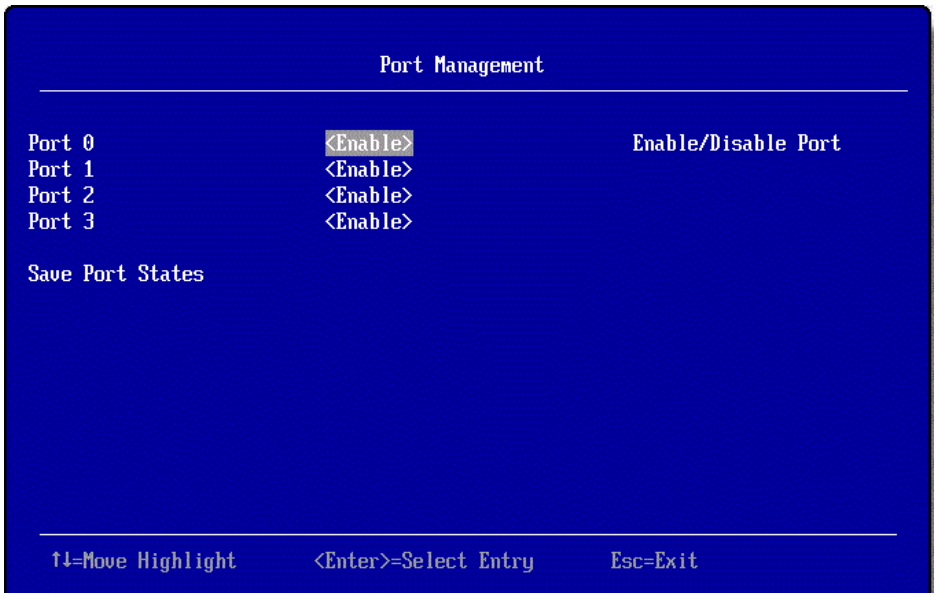
Enabling any of the ports restores the previously removed PCI functions and restores the power consumption to its normal state.

NOTE Enabling and disabling ports is applicable in both Multichannel and non-Multichannel configurations. If Multichannel mode is enabled, all virtual functions associated with the disabled port are also disabled.

To configure port management:

1. On the Emulex NIC Selection dialog (Figure 99), select Port Management and press Enter. The following Port Management dialog appears.

Figure 103 Port Management Dialog



2. Select the current setting for a particular port and a pop up appears.
3. Select **Enable** or **Disable** and press Enter.
4. When you are finished, select Save Port States to save the current settings.
5. Press Esc to return to the Controller Configuration menu.

NOTE A reboot is required for this setting to take effect.

11.2.11 Feature on Demand

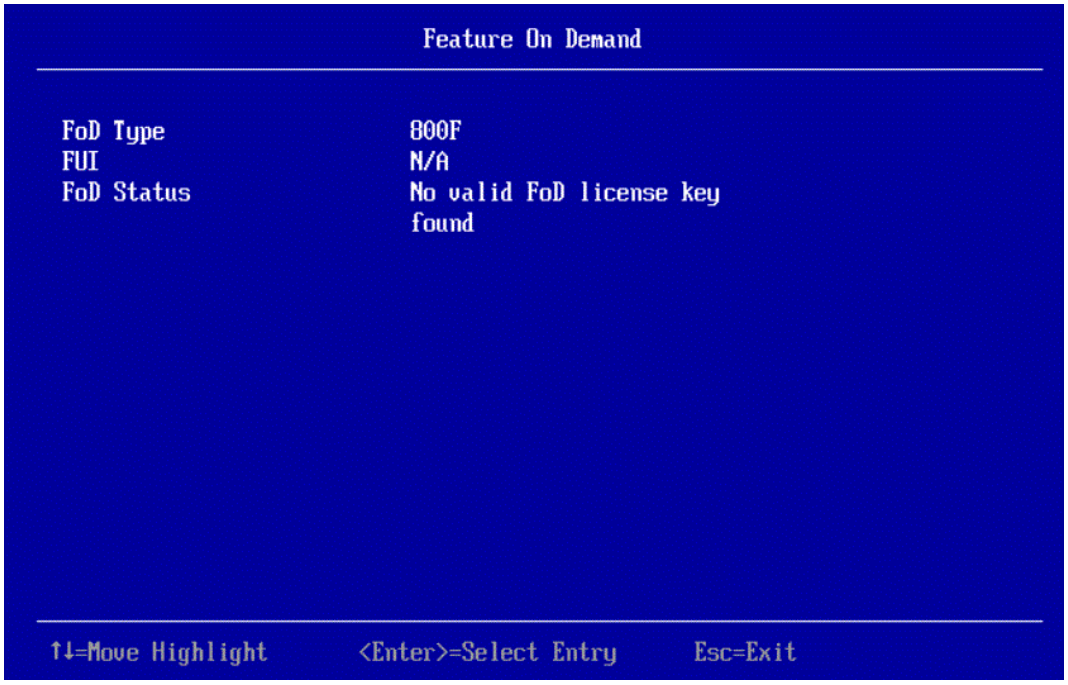
NOTE The Feature on Demand function is available only on some Lenovo System x adapters.

FoD is a Lenovo System x proprietary protocol that enables storage functions on Emulex adapters. This feature requires a unique license key to be applied from Lenovo System x's Integrated Management Module. For more information on obtaining a license key, contact your Lenovo System x representative.

The FoD Type and FUI fields are required by Lenovo System x to generate a unique license key for the specific adapter for which you have requested a license. Each adapter in the system must obtain its own license key.

To view Feature on Demand information, from the Emulex NIC Selection dialog (Figure 99), select Feature On Demand and press Enter. The Feature On Demand information appears.

Figure 104 Feature On Demand



From the Feature On Demand screen, you can view information on the feature type, the FoD Unique Identifier, and its current status.

- **FoD Type** – The FoD Type field is unique for each platform. Different products have different FoD types. For example, LOM products have a different FoD type than mezzanine cards.
- **FUI** – The FUI field is internally generated by the Emulex firmware using the FoD type and the adapter's serial number. The FUI is unique to each adapter because each card has its own unique serial number.
- **FoD Status** – The FoD Status field indicates whether a license key has been applied properly or not.
  - If the FoD status field is set to Enabled, it indicates that the FoD license key has been successfully applied and FoD is enabled. If FoD is enabled, you will see all of the storage personalities under the Personality menu option, including NIC, iSCSI, and FCoE.
  - In all other cases, the FoD Status field indicates that a failure has occurred or that an FoD license has not been applied.

11.3 Downloading Firmware and Boot Code

A single firmware image contains the latest version of the firmware and boot code.

The firmware and UEFI NIC, iSCSI, and FCoE drivers are contained in one image file with the .UFI extension. This file must be flashed through the NIC interface exposed by HII.

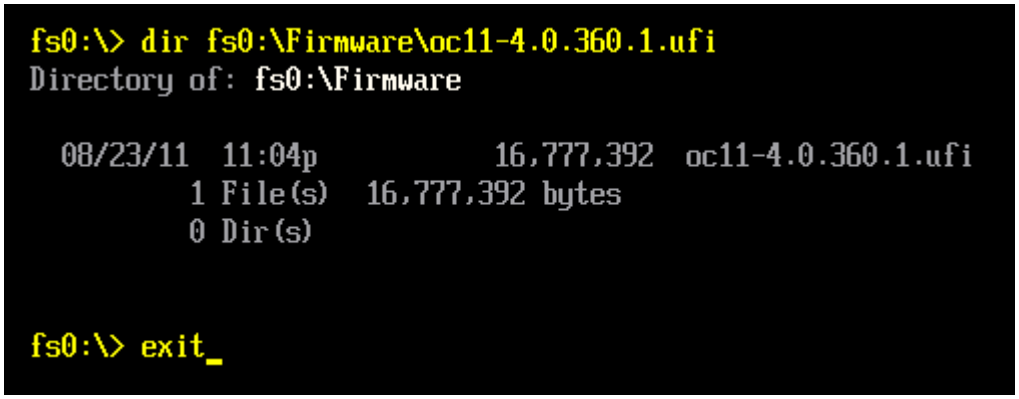
To view the adapter's current firmware and boot code version, follow the steps in Section 11.2, Starting the Emulex NIC Configuration Utility.

The Emulex NIC Selection dialog displays the current firmware version for the selected adapter and a link to the Emulex Flash Update Utility, which you can use to download firmware and boot code.

To download firmware and boot code:

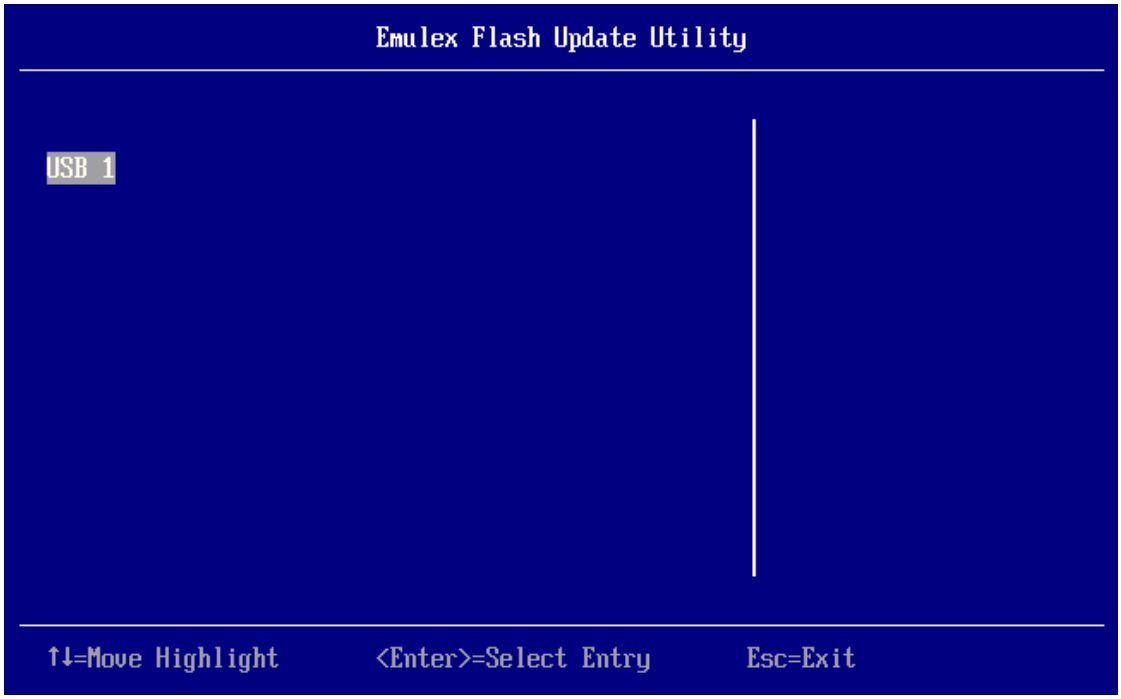
1. In an open UEFI shell, copy the firmware and boot code file into a directory on the UEFI partition media.

Figure 105 UEFI Shell with Firmware and Boot Code File



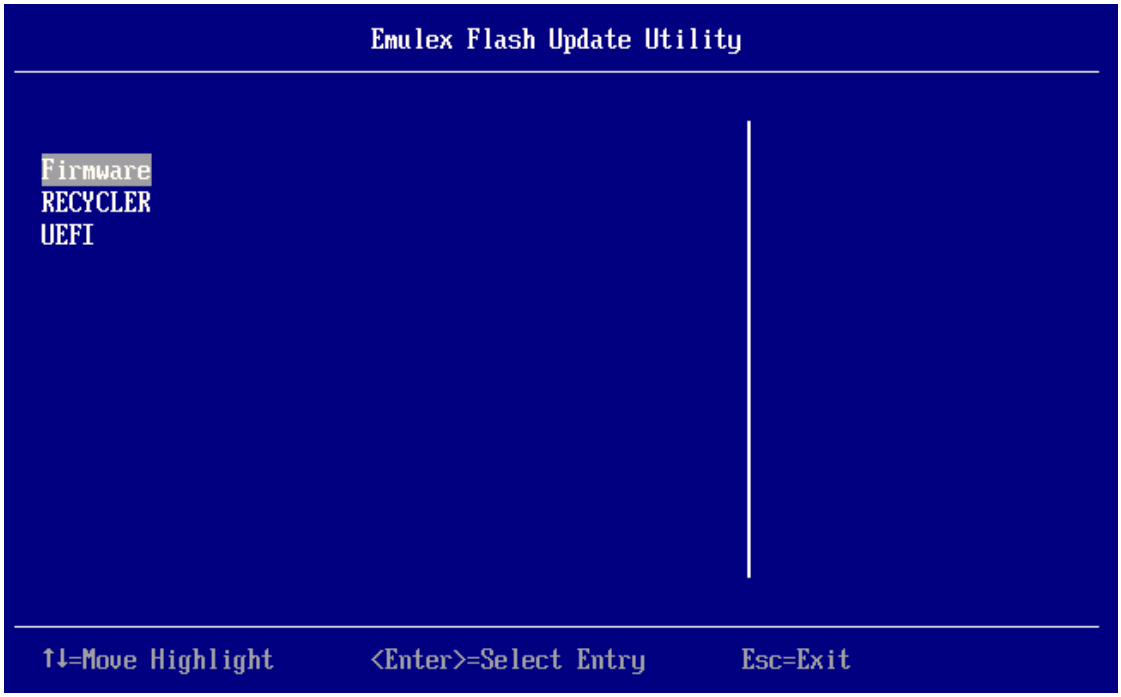
2. Exit the UEFI shell and launch the Emulex NIC configuration utility.
3. From the Network device list, select the NIC adapter you want to modify and press Enter.
4. The Emulex NIC Selection dialog (Figure 79) shows information for the selected adapter. Select Emulex Flash Update Utility and press Enter. The Emulex Flash Update Utility dialog (Figure 99) appears. This utility displays all available media and installs the flash file on the adapter.

Figure 106 Emulex Flash Update Utility Dialog



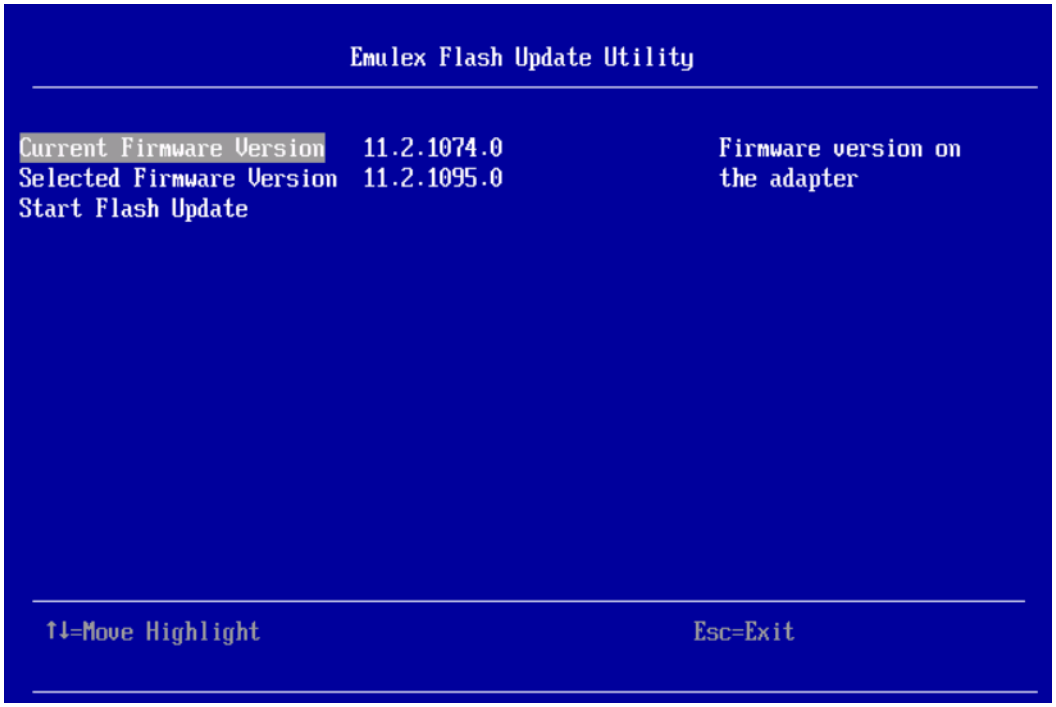
5. Press Enter to select the media containing the flash file.
6. Navigate to the directory containing the flash file and press Enter (Figure 107).

Figure 107 Emulex Flash Update Utility with Directory Name



7. Use the arrow keys to select the flash file and press Enter.
8. Use the arrow keys to select Start Flash Update, and press Enter to begin the update process.

Figure 108 Emulex Flash Update Utility, Start Flash Update

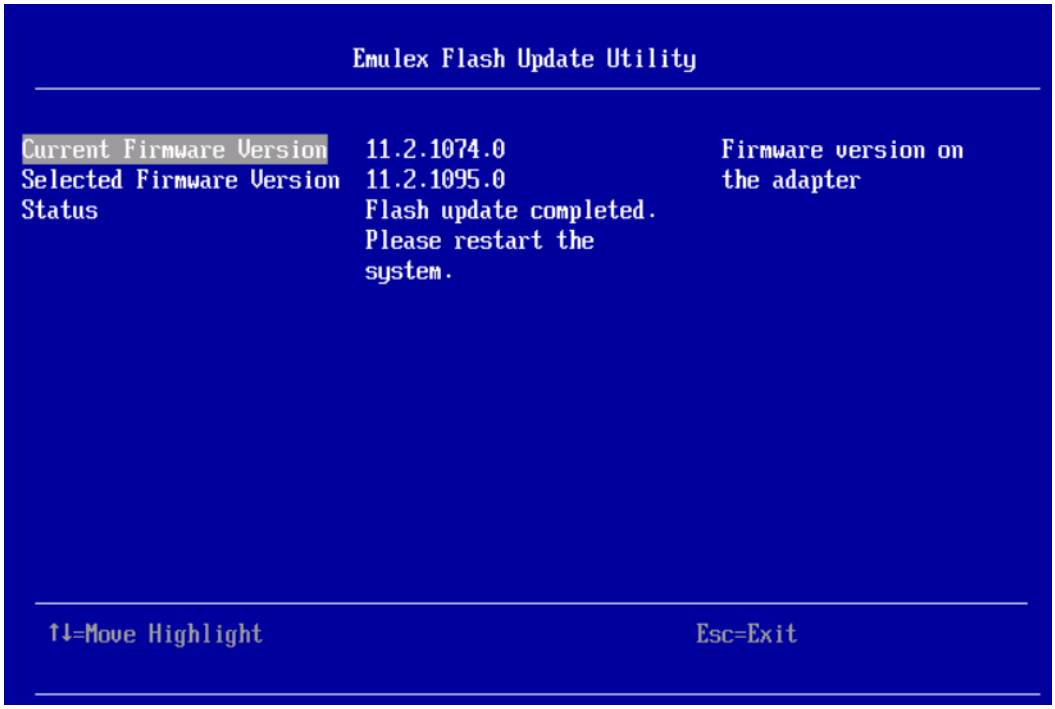


NOTE This firmware flash procedure only flashes the physical interface that was selected in the HII menu. All Emulex adapters must contain the same firmware version to function properly. The above procedure must be repeated for each physical interface installed in the system before proceeding to the next step.

NOTE If a jumper is plugged into the ASIC GPIO pin on the adapter, it indicates that only secure firmware versions are allowed to be flashed on the adapter. Upgrade from earlier unsecure firmware versions does not involve a digital signature check since the earlier code does not support this feature. Therefore, upgrade to secure firmware versions will be successful. If a secure version of the firmware (version 11.0 or later) is installed on an Emulex OCe14000B-series adapter and you want to update to an earlier unsecure version of firmware, you must remove the secure firmware jumper block before performing the update. See the adapter installation guide for more information. Firmware upgrade from a secure firmware version to an unsecure firmware version results in an error when the jumper is in place. If this happens, an error message indicating digital signature failure appears.

9. The flash update begins. It takes about two minutes to complete.

Figure 109 Emulex Flash Update Utility, Completed



10. Reset the system. The latest version information for the adapter appears in the Emulex NIC Selection dialog (Figure 79).

11.4 Identifying a Port

To physically identify a port on the Emulex NIC Selection dialog (Figure 79), select Port Identification and press Enter. The LEDs on the adapter start blinking indicating the selected port.

NOTE Not all adapters have externally–visible LEDs. If you are using an add-in card in a blade server environment, port identification is not supported.

If port identification is not supported on the system, you will receive a message stating that it is not supported.

11.5 Erasing Ports and Adapter Configurations

NOTE On OCe11100-series adapters, if you select the Erasing Ports and Adapter Configurations setting, all previous configuration settings are returned to their factory default settings except for the current personality selection. You must perform this action to provide a clean environment for new configuration settings to take effect.

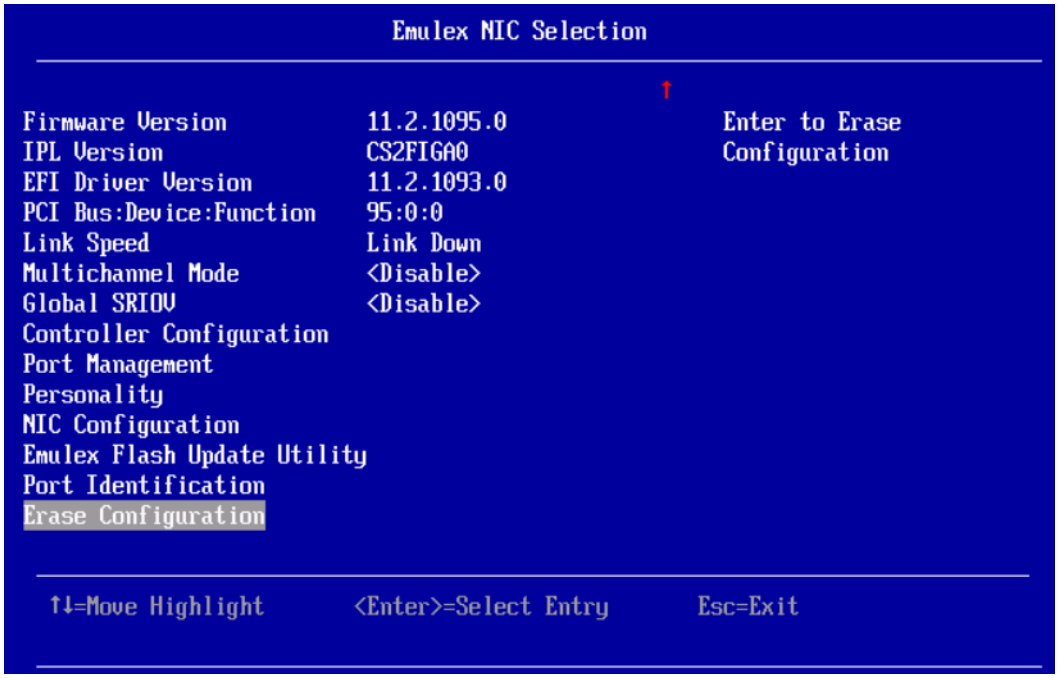
On OCe14000 series adapter, if you select the Erasing Ports and Adapter Configurations setting, all previous configuration settings are returned to their factory default settings including the current personality selection. Performing this action ensures a clean environment for new configuration settings to take effect.

NOTE The Erase Configuration option is not available on HP adapters in HP systems. Instead, use the F7 key to reset the configuration settings to their factory default settings.

To erase ports and the adapter configuration:

1. On the Emulex NIC Selection dialog (Figure 79), select Erase Configuration and press Enter (Figure 110) .

Figure 110 Erase Configuration



2. Press Enter to delete the adapter configuration, or press Esc to stop the operation.

11.6 UEFI NIC Diagnostics

The UEFI NIC driver diagnostics protocol can be used to run diagnostic tests on each NIC function of the card. This protocol is used by a platform management utility to let you run driver specific diagnostics on an adapter.

11.6.1 EFI_DRIVER_DIAGNOSTICS_PROTOCOL

Syntax:

```
drvdiag [-c] [-l XXX] [-s] [-e] [-m] [driverhandle [devicehandle [childhandle]]]
```

Description:

The UEFI NIC driver diagnostics protocol can be used to run diagnostic tests on each NIC function of the card. Diagnostics can be run in standard mode, extended mode, and manufacturing mode.

Parameters:

`-c`	Diagnose all child devices.
`-l XXX`	Diagnose using the ISO 639-2 language specified by XXX.
`-s`	Run diagnostics in standard mode. Diagnostics in standard mode runs the LED test, Link test, Get MAC test, and the DMA test (Read, Write, and Read and Write)
`-e`	Run diagnostics in extended mode. Diagnostics in extended mode runs the ARM Timer test, the MAC Loopback test, and the Physical Loopback test.

`-m`	Run diagnostics in manufacturing mode. Diagnostics in manufacturing mode run the Low-level Subsystem Network Loopback test.
`driverhandle`	Handle of the driver being configured.
`devicehandle`	Handle of a device the driverhandle is managing.
`childhandle`	Handle of a device that is a child of the devicehandle.

Examples:

The following examples show you a way of using the EFI_DRIVER_DIAGNOSTICS_PROTOCOL.

The `driver` command identifies the handle of the driver:

```
Shell> drivers

      122 Emulex 10G NIC
```

The `drvdiag` command lists all the devices available for diagnostics. Each Ctrl [XXX] corresponds to a NIC function, which can be physical or virtual. The command returns handles for NIC functions on both ports.

```
Shell> driverdiag

      Drv[122] Ctrl[121]
      Drv[122] Ctrl[123]
```

To run the standard diagnostic tests on function 0:

```
Shell> drvdiag -s 122 121
```

To run the standard diagnostic tests on all NIC functions:

```
Shell> drvdiag -s 122
```

Chapter 12: Configuring UEFI for iSCSI

NOTE For systems with multiple adapters, the UEFI system firmware or boot code uses the highest version driver installed on any of the adapters. Adapters with older versions of EFIBoot are managed by the more recent version, but only as long as the adapter with the most recent version is in the system. The latest firmware and boot code must be installed on each adapter in the system to ensure that each adapter runs the latest firmware and boot code.

UEFIBoot supports:

- Configuration Access and Component Name2 UEFI protocols.
- Selectable SCSI device path through the driver configuration protocol
- Up to 96 adapters in a system
- Up to 256 targets (selectable through the driver configuration protocol)
- Up to 256 LUNs
- Multiutility setup

This version of UEFIBoot is loaded from flash into system memory.

If UEFIBoot is loaded in a UEFI Shell, type `drivers` and press Enter to see if the driver is loaded.

NOTE If you are using a Dell adapter, see Appendix F, Dell UEFI for information on using the Dell UEFI utility.

12.1 Navigating the Emulex iSCSI Configuration Utility

The Emulex iSCSI configuration utility has menus and configuration dialogs. Use the following methods to navigate them:

- Press the up and down arrows on the keyboard to navigate menu options or configuration fields. If multiple adapters are listed, use the up and down arrows to scroll to the additional adapters.
- Press the +, -, or Enter keys to change numeric values.
- Press Enter to select an option.
- Press Esc to exit the current screen and show the previous screen.

NOTE HP adapters in HP systems use the F10 key to save the current configuration settings.

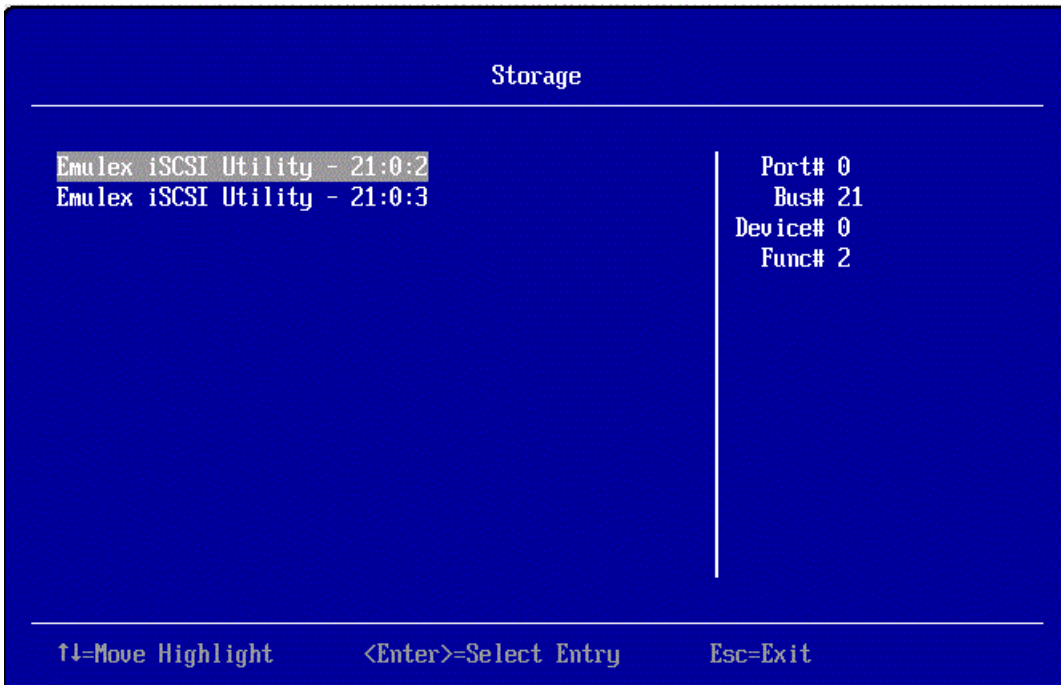
12.2 Starting the Emulex iSCSI Configuration Utility

Depending on the vendor UEFI configuration, the Emulex iSCSI configuration utility may appear under different setup menus in the vendor system firmware or BIOS (such as Storage).

To start the Emulex iSCSI configuration utility:

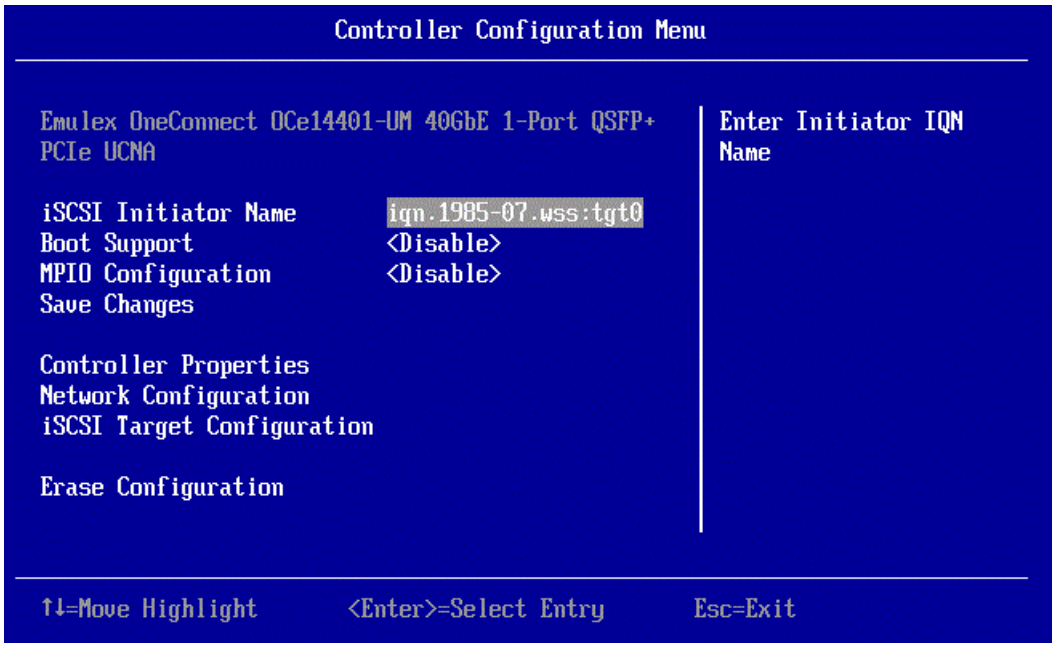
1. Select the Emulex iSCSI Utility for a particular function in the Storage dialog and press Enter.

Figure 111 Storage Dialog



2. The Controller Configuration Menu appears.

Figure 112 Controller Configuration Menu



3. Highlight iSCSI Initiator Name and press Enter to edit the initiator name.
   This is a global setting. The initiator name configured on one adapter is reflected on all adapters in the system.

NOTE HP systems are an exception. Each port on an HP system can be configured with a unique initiator name.

4. To save changes, select Save Changes and press Enter.

12.3 Configuring MPIO

MPIO support allows the initiator to log in dual sessions to the same target. In this way I/O can be sent over either TCP/IP connection to the target. If one session fails another session can continue processing I/O without interruption to the application. In iSCSI target configuration, you have the option of setting dual network paths to a single boot LUN.

NOTE Although MPIO boot support allows the initiator to log in to multiple sessions, the iSCSI BIOS currently limits the number of sessions to two iSCSI sessions for a single boot LUN.

You must follow these steps to configure MPIO boot support for each operating system.

1. Use the Emulex iSCSI configuration utility to configure the first path to a boot target.
2. Complete normal operating system installation.
3. Install third-party MPIO software for your operating system.
4. Configure the second path to a single boot LUN through the Emulex iSCSI configuration utility.

To configure MPIO:

1. From the Controller Configuration Menu (Figure 112), select MPIO Configuration and press Enter. The MPIO Configuration pop up appears.
2. Select **Enable** or **Disable** and press Enter.
3. To save your changes, select Save Changes and press Enter.

12.4 Configuring Boot Support

Boot support is provided for each port or function. If boot support is enabled, you can boot from the specified function. If boot support is disabled, you will not be able to boot from the function.

To configure boot support:

1. From the Controller Configuration Menu (Figure 112), select Boot Support and press Enter. The Boot Support pop up appears.
2. Select **Enable** or **Disable** and press Enter.
3. To save your changes, select Save Changes and press Enter.

12.5 Configuring TCP ACK Delay

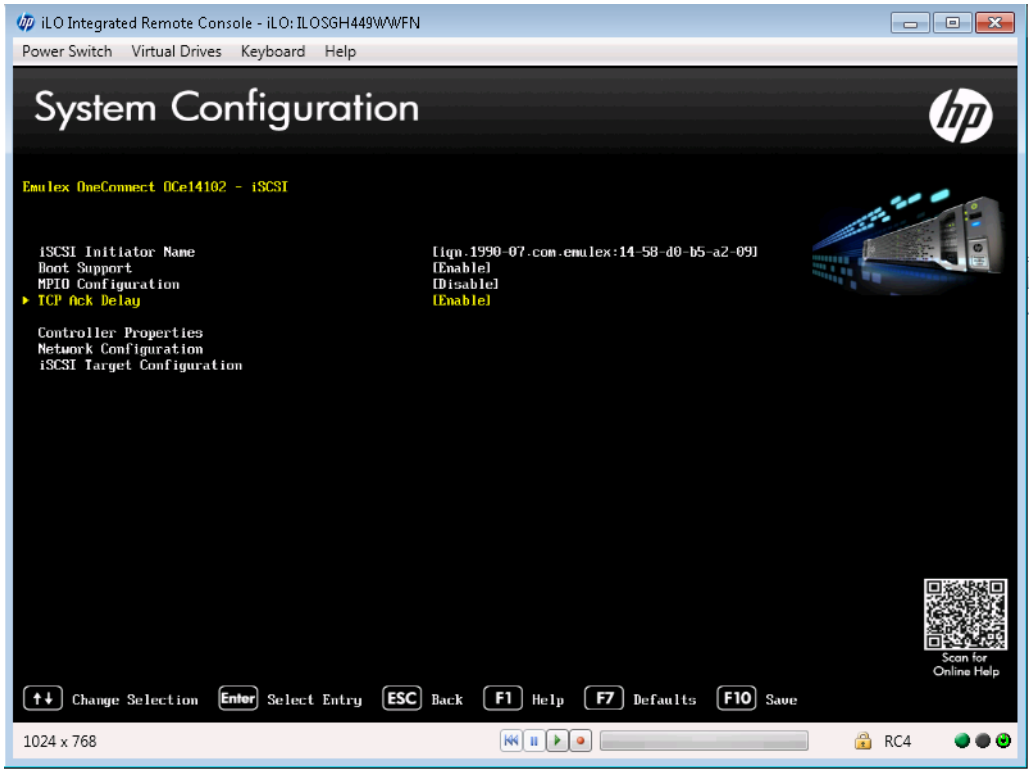
NOTE This functionality is available only on HP adapters that support UEFI.

TCP ACK delay is provided for each port or function. If TCP ACK delay is enabled, several ACK responses may be combined into a single response. If TCP ACK delay is disabled, an ACK is sent for each TCP packet received.

To configure TCP ACK delay:

1. From the HP Controller Configuration Menu (Figure 113), select TCP Ack Delay and press Enter.

Figure 113 Configuring TCP ACK Delay



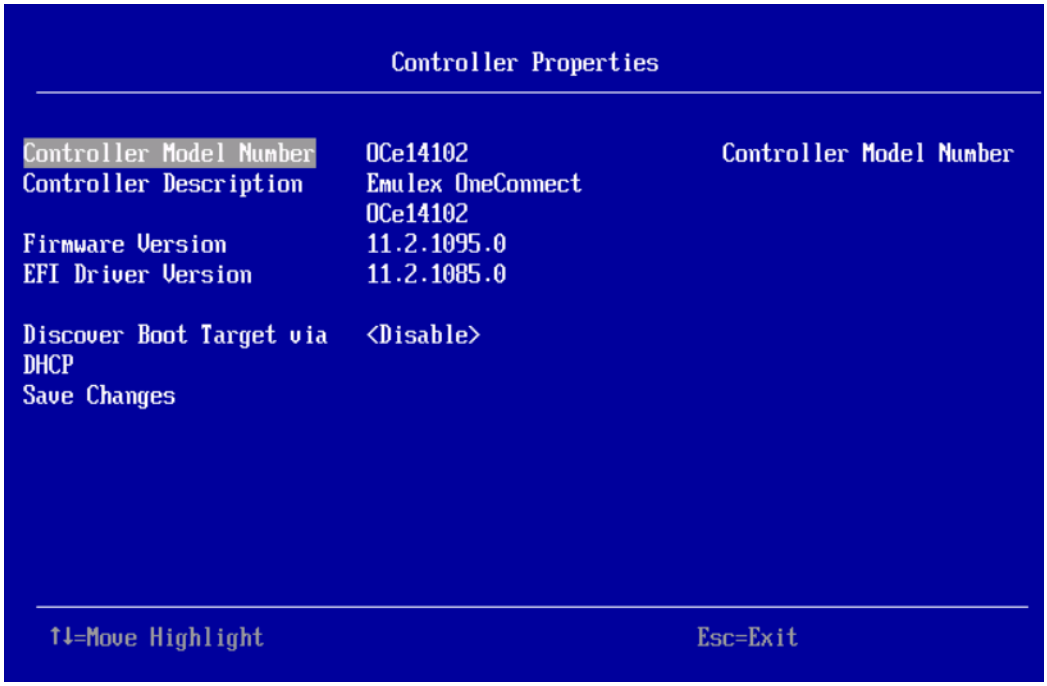
2. Select **Enable** or **Disable** and press Enter.
3. To save your changes, select Save.

12.6 Viewing the Controller Properties

To view the adapter properties:

1. From the Controller Configuration Menu (Figure 112), select Controller Properties and press Enter. The Controller Properties menu (Figure 114) appears.

Figure 114 Controller Properties Menu



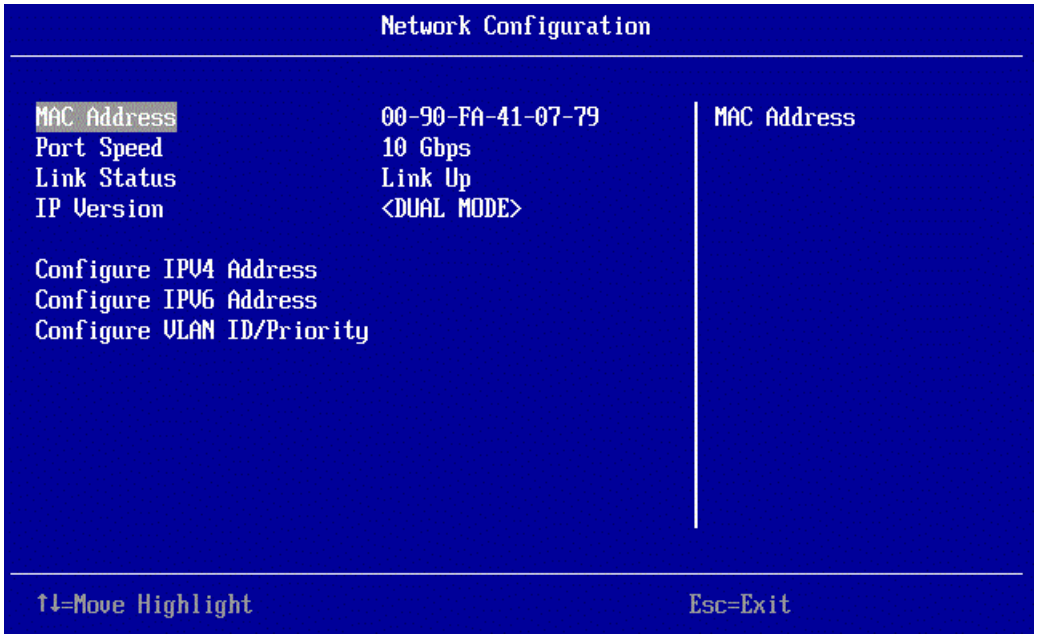
The Controller Properties menu displays the model number, description, BIOS version, and firmware version for the adapter. From the Controller Properties menu, you can configure the discovery of the boot target using DHCP. See Section 12.9.1, Discovering and Adding Boot Targets through DHCP, for more information.

12.7 Configuring the Network

To configure the network:

1. From the Controller Configuration menu (Figure 112), select Network Configuration and press Enter. The Network Configuration dialog (Figure 115) appears.

Figure 115 Network Configuration Dialog



The Network Configuration dialog displays the MAC address, port speed, and link status for the adapter. From the Network Configuration dialog, you can configure the following items:

- IP version
- IPv4 address
- IPv6 address
- VLAN ID and priority

12.7.1 Configuring the IP Version

To configure the IP version:

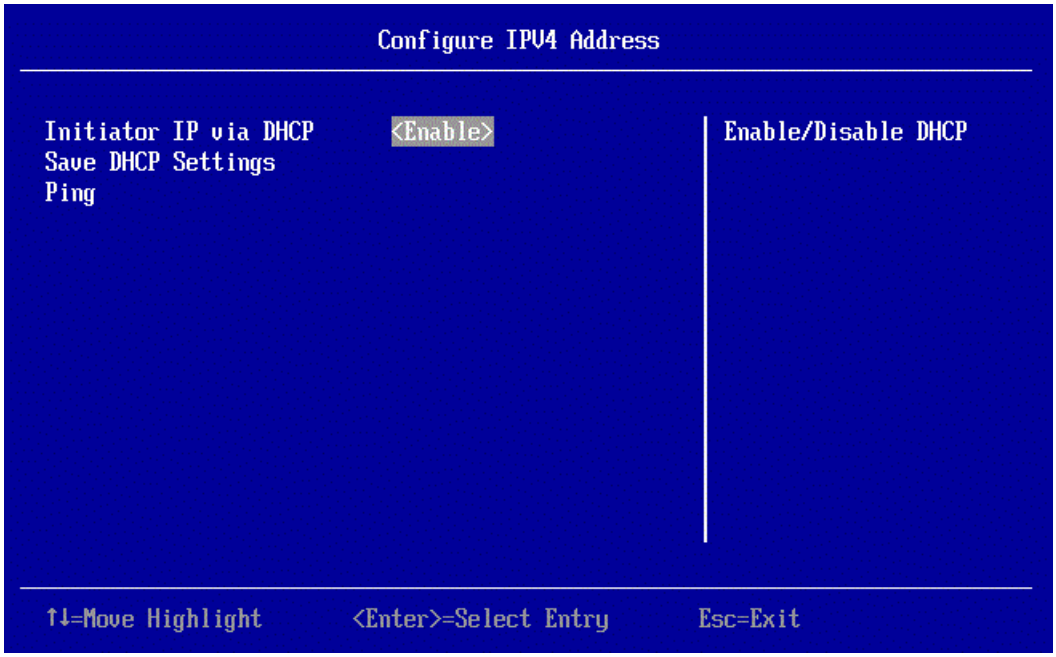
1. On the Network Configuration dialog (Figure 115), select IP version and press Enter. The IP Version pop up appears.
2. Select IPV4, IPV6, or DUAL MODE and press Enter.

12.7.2 Configuring an IPv4 Address

To configure an IPv4 address:

1. On the Network Configuration dialog (Figure 115), verify that the IP version is set to IPV4 or DUAL MODE.
2. Select Configure IPV4 Address and press Enter. The Configure IPV4 Address dialog (Figure 116) appears.

Figure 116 Configure IPV4 Address Dialog



From the Configure IPV4 Address dialog, you can do the following:

- Enable DHCP for automatic assignment of the IP address through a DHCP server. See Section 12.7.2.1, Automatically Assigning an IP Address through a DHCP Server, for more information.
- Configure a static IP address (if Initiator IP via DHCP is disabled). See Section 12.7.2.2, Manually Assigning an IP Address, for more information.
- Ping the iSCSI initiator. See Section 12.7.2.3, Pinging the iSCSI Initiator, for more information.

12.7.2.1 Automatically Assigning an IP Address through a DHCP Server

To enable DHCP for automatic assignment of the IP address through a DHCP server:

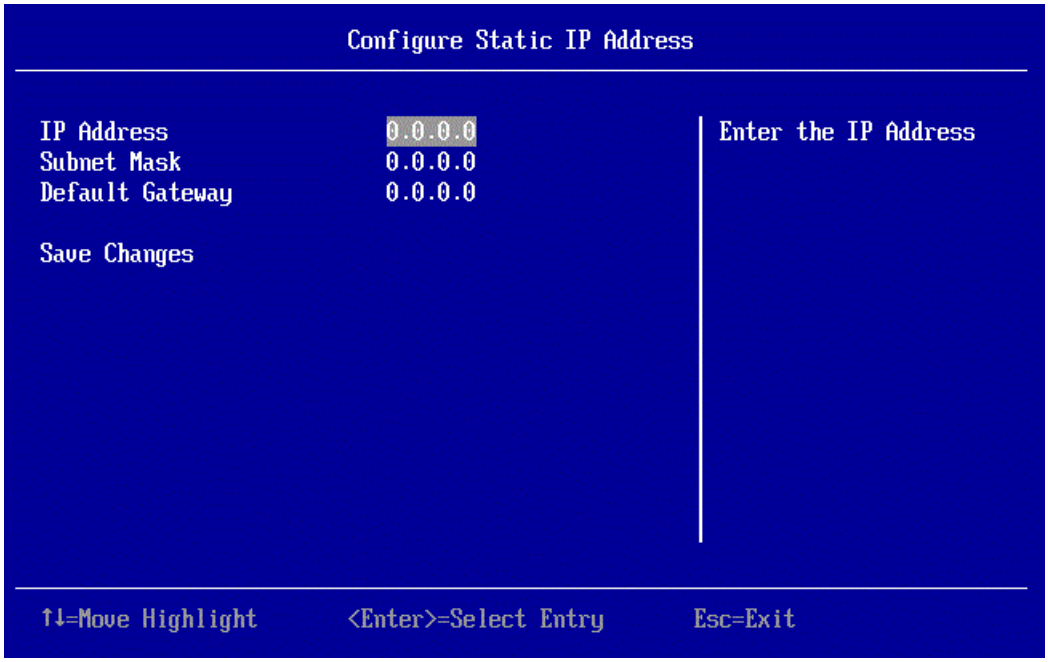
1. On the Configure IPV4 Address dialog (Figure 116), select Initiator IP via DHCP and press Enter. The Initiator IP via DHCP pop up appears.
2. Select **Enable** and press Enter.
3. Select Save DHCP Settings and press Enter.

12.7.2.2 Manually Assigning an IP Address

To manually assign an IP address:

1. On the Configure IPV4 Address dialog (Figure 116), verify that the Initiator IP via DHCP is set to Disable. If you change this setting from `Enable` to `Disable`, select Save DHCP Settings and press Enter.
2. Select Configure Static IP Address and press Enter. The Configure Static IP Address dialog (Figure 117) appears.

Figure 117 Configure Static IP Address Dialog



3. Enter the IP address, subnet mask, and default gateway in the corresponding fields.
4. Select Save Changes and press Enter.

12.7.2.3 Pinging the iSCSI Initiator

To ping the iSCSI initiator:

1. On the Configure IPV4 Address dialog (Figure 116), select Ping and press Enter. A pop up text box appears.
2. Enter the IP address of the iSCSI initiator in the text box and press Enter.
   - If the ping is successful, a reply message appears with the iSCSI initiator IP address.
   - If the ping is unsuccessful, a failure message appears.

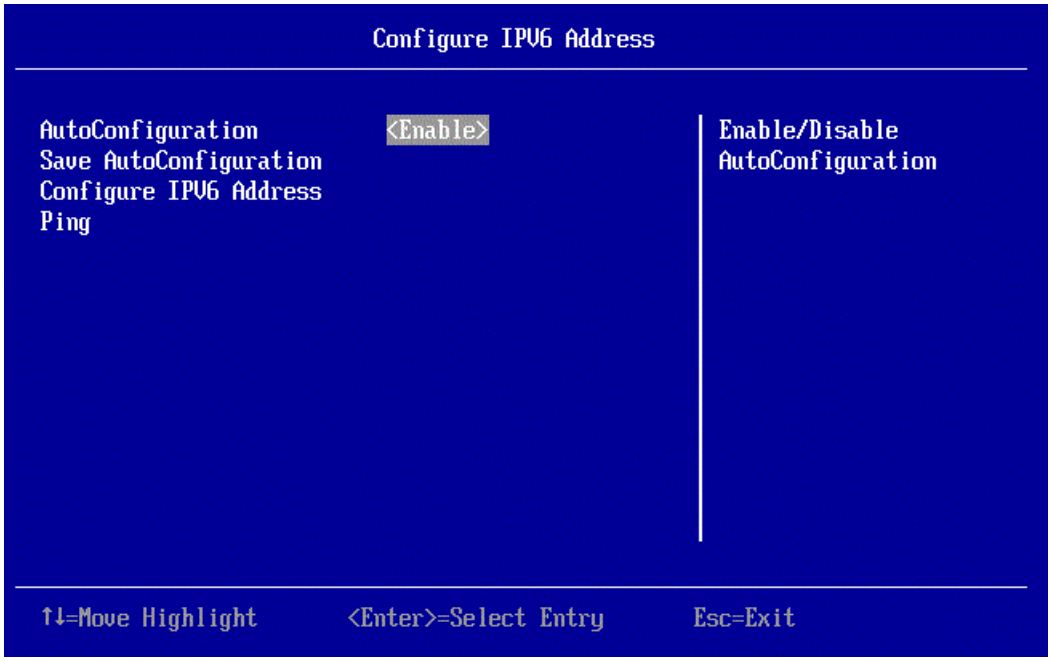
12.7.3 Configuring an IPv6 Address

NOTE For HP adapters in HP systems, all of the following IPv6 configuration options are available from the Configure IPV6 Address dialog. The instructions in the following sections still apply to HP adapters; however, use the F10 key to save the configuration settings.

To configure an IPv6 address:

1. On the Network Configuration dialog (Figure 115), verify that the IP version is set to IPV6 or DUAL MODE.
2. Select Configure IPV6 Address and press Enter. The Configure IPV6 Address dialog (Figure 118) appears.

Figure 118 Configure IPV6 Address Dialog



From the Configure IPV6 Address menu, you can do the following:

- Enable automatic configuration of the IP address. See Section 12.7.3.1, Automatically Configure an IP Address, for more information.
- Configure the IPv6 address. See Section 12.7.3.2, Configuring the IPv6 Address, for more information.
- Ping the iSCSI initiator. See Section 12.7.3.3, Pinging the iSCSI Initiator, for more information.

12.7.3.1 Automatically Configure an IP Address

To enable automatic IP address configuration:

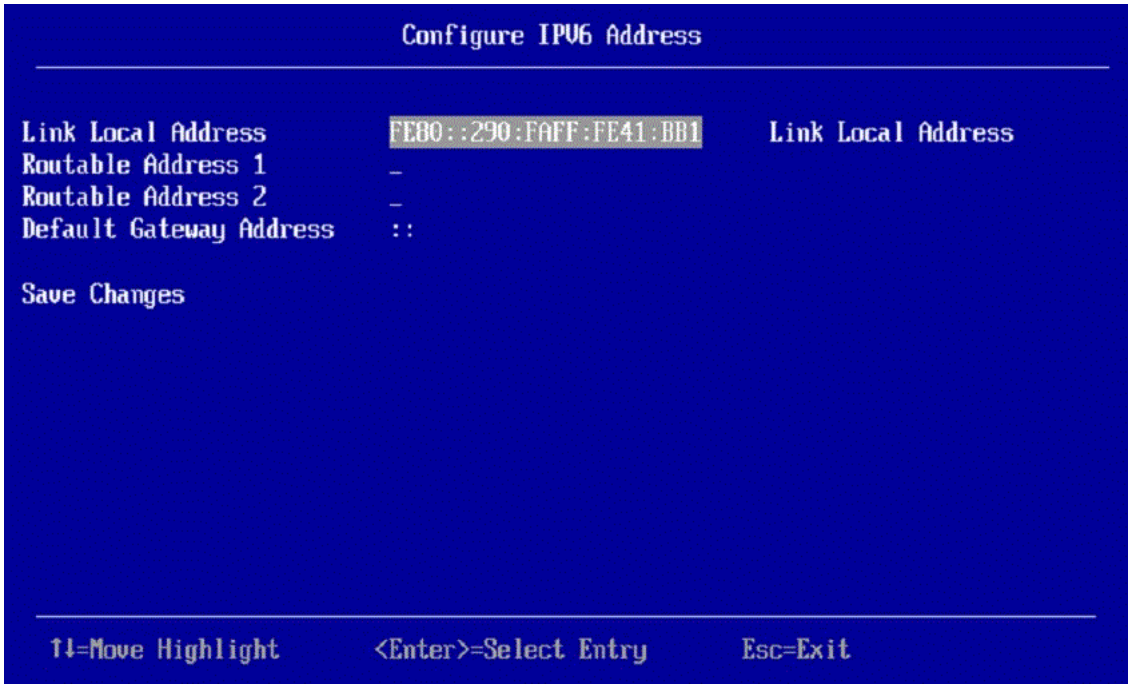
1. On the Configure IPV6 Address dialog (Figure 118), select AutoConfiguration and press Enter. The AutoConfiguration pop up appears.
2. Select **Enable** and press Enter.
3. Select Save AutoConfiguration and press Enter.

12.7.3.2 Configuring the IPv6 Address

To configure the IPv6 address:

1. On the Configure IPV6 Address dialog (Figure 118), select Configure IPV6 Address and press Enter. The Configure IPV6 Address dialog (Figure 119) appears.

Figure 119 Configure IPV6 Address Settings Dialog



2. Enter the following information in the corresponding fields:
   - **Link Local Address** – The unique address assigned to the port 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 port.
   - **Routable Address 2** – A second routable address assigned to the iSCSI port.
   - **Default Gateway** – A gateway is a router on a computer network that serves as an access point to another network and that an initiator uses by default.
3. Select Save Changes and press Enter.
4. Press Esc to exit the menu.

12.7.3.3 Pinging the iSCSI Initiator

To ping the iSCSI initiator:

1. On the Configure IPV6 Address dialog (Figure 118), select Ping and press Enter. A pop up text box appears.
2. Enter the IP address of the iSCSI initiator in the text box and press Enter.
   - If the ping is successful, a reply message appears with the iSCSI initiator IP address.
   - If the ping is not successful, a failure message appears.

12.7.4 Configuring VLAN ID and Priority

A VLAN is a way of partitioning the network. If the LAN is made up of all devices within a broadcast domain, a VLAN is a broadcast domain made up of switches. You first create a VLAN and then assign ports to a VLAN. All ports in a single VLAN are in a single broadcast domain.

You do not have to configure VLANs unless the network is already using them. Some reasons why VLANs are used include:

- A LAN increases in size with numerous devices.

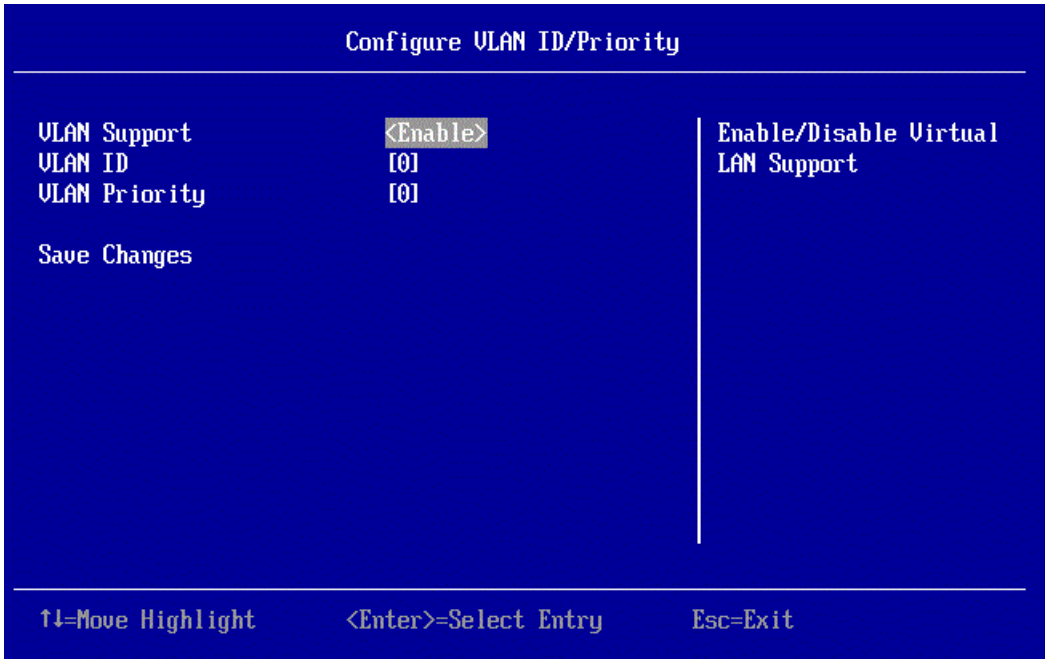
- A LAN has extensive broadcast traffic on it.
- Groups of users on a LAN need more security.

A VLAN ID, like an IP address or initiator name, is assigned a value to uniquely identify it on a network. A VLAN priority is set to determine what packet gets priority order within a VLAN.

To configure a VLAN ID and priority:

1. On the Network Configuration dialog (Figure 115), select Configure VLAN ID/Priority and press Enter. The Configure VLAN ID/Priority dialog appears.

Figure 120 Configure VLAN/ID Priority Dialog



2. To enable VLAN support:
   a. Select VLAN Support and press Enter. The VLAN Support pop up appears.
   b. Select **Enable** and press Enter.
3. To assign a VLAN ID number:
   a. Select VLAN ID and press Enter. This is a unique value you assign to each VLAN on a single device. Valid values range from 2 to 4094.
   b. Enter a VLAN ID value and press Enter.
4. To set a VLAN priority, if necessary:
   a. Select VLAN PRIORITY and press Enter. This unique value assigns a priority to outbound packets containing a specified VLAN ID. Valid values range from 0 to 7, with 0 being the highest priority level.
   b. Enter a VLAN priority value and press Enter.
5. Select Save Changes and press Enter.
6. Press Esc to return to the Network Configuration dialog.

12.8 Updating Firmware

To update firmware, you must use the NIC firmware update utility, which revises the iSCSI function with a single firmware download image. For more information on downloading firmware, see Section 11.3, Downloading Firmware and Boot Code.

12.9 Adding and Configuring Targets

After the initiator has been configured, you must make an iSCSI target available to the initiator host. Targets can be discovered automatically using the DHCP discover mechanism, the SendTargets discovery mechanism, or you can manually configure targets.

Before you begin the login session, gather the following information:

- **iSCSI target name (manual configuration only)** – The target name you are adding. The iSCSI target name is not required if you are using DHCP or SendTargets discovery. It is required for manually configured targets only. This name should be known to you based on how you configured the iSCSI target.
- **iSCSI target IP address** – The IP address of the target into which you are logging.
- **TCP port number** – The TCP port number of the target portal. Typically this is 3260, which is the well-known port number defined for use by iSCSI.

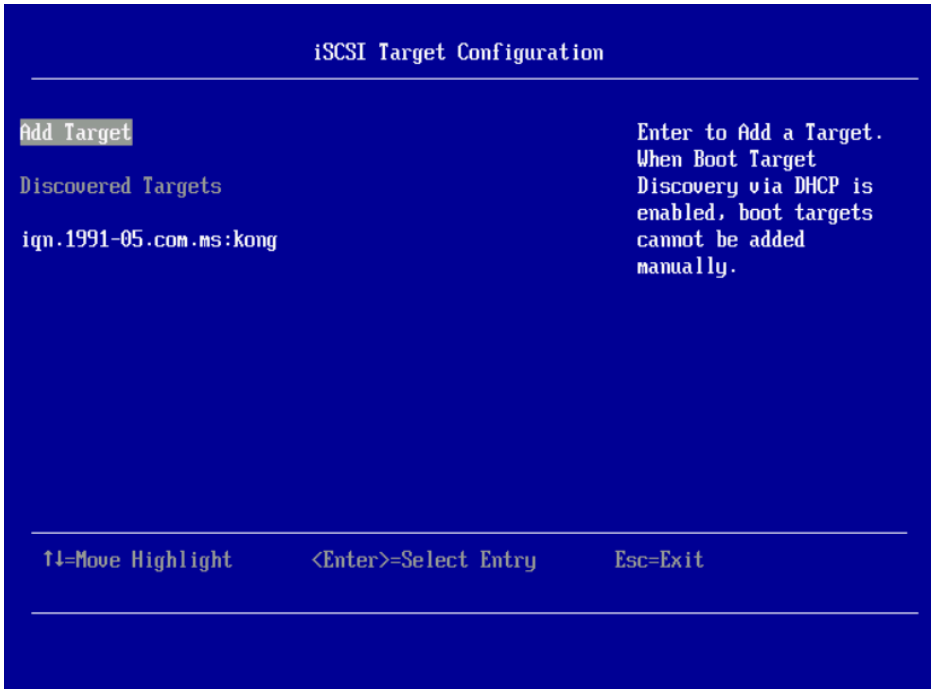
12.9.1 Discovering and Adding Boot Targets through DHCP

To automatically discover and add boot targets through DHCP:

1. On the Controller Properties dialog (Figure 114), set Discover Boot Target via DHCP to `Enable` and press Enter.
2. Select Save Changes and press Enter.
3. Reboot the system.
4. On the Controller Configuration Menu (Figure 112), select iSCSI Target Configuration and press Enter. The iSCSI Target Configuration dialog (Figure 121) shows the discovered targets.

For more information on configuring a DHCP boot target, see Section 10.3, Discovering Targets through DHCP for iSCSI.

Figure 121 iSCSI Target Configuration Dialog



12.9.2 Using the SendTargets Discovery to Add an iSCSI Target

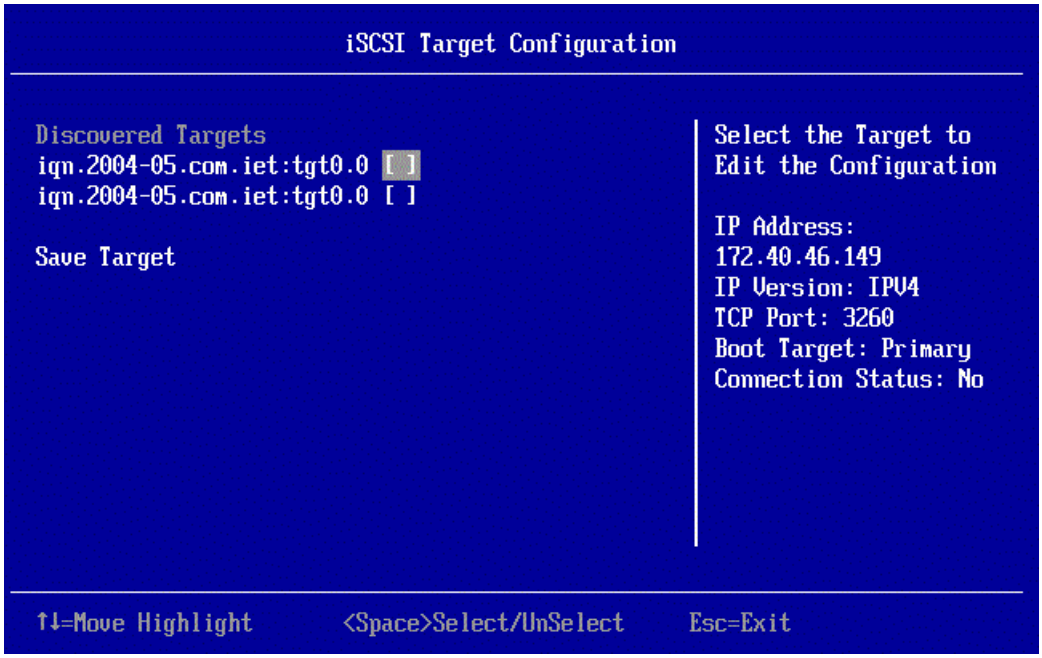
The iSCSI initiator uses the SendTargets discovery mechanism to locate names of available targets. Use this method if an iSCSI node has a large number of targets.

To configure an iSCSI target using the SendTargets discovery:

1. On the Controller Properties dialog (Figure 114), set Discover Boot Target via DHCP to `Disable` and press Enter.
2. Select Save Changes and press Enter.
3. Reboot the system.
4. From the Controller Configuration Menu (Figure 112), select iSCSI Target Configuration and press Enter. The iSCSI Target Configuration dialog (Figure 121) appears.
5. Select Add Target and press Enter. The Add/Ping iSCSI Target dialog (Figure 123) appears.
6. Leave the iSCSI Target Name blank.
7. Enter the iSCSI target IP address.
8. Change the TCP port number value, if necessary. The default number is 3260.
9. If you are connecting dual sessions to the same target portal group, enter a unique ISID Qualifier, which can be any number up to 65535. For more information on when an ISID value must be changed, see Section 5.3.2, Booting from SAN for iSCSI MPIO.
10. Leave the Boot Target field set to No for now; you will return to this field later to enable the boot target.
11. Set the Header Digest to Yes if you want to enable the header digest. If the Header Digest is set to Yes, and the iSCSI initiator is set accordingly, the integrity of an iSCSI protocol data unit's header segment is protected by CRC32C checksum. The default is disabled.
12. Set the Data Digest to Yes if you want to enable the data digest. If the Data Digest is set to Yes, and the iSCSI initiator is set accordingly, the integrity of an iSCSI protocol data unit's data segment is protected by CRC32C checksum. The default is disabled.

13. Select an authentication method (optional). If you are enabling an authentication method, you are prompted to enter CHAP configuration.
14. Select Save/Login. A message reminds you that you have left the iSCSI Target name blank and that the SendTargets mechanism will be used. If you want to continue, press Y.
15. After you send the SendTargets request, the firmware attempts to log in to the target. If it is successful, the Targets Discovered Via SendTargets dialog displays a list of targets.

Figure 122 Discovered Targets Dialog



16. Select the targets you want to add from the list and press the space bar.
17. Add these targets to the list of iSCSI targets available for the initiator to log in to by selecting Save Target and press Enter. If you enabled CHAP Authentication, you are prompted to enter CHAP configuration for each target, one at a time.
18. From the Controller Configuration Menu (Figure 112), select iSCSI Target Configuration and press Enter to view the added target information.

NOTE If you set the Boot Target option in step 5 before adding the target, the Boot Target appears as No. To enable Boot Target, continue with step 19. Otherwise, you have completed the procedure.

19. To enable Boot Target or make any other changes to the targets, select the target and press Enter. The Edit/Ping iSCSI Target dialog (Figure 124) appears. From this dialog, you can edit the target.

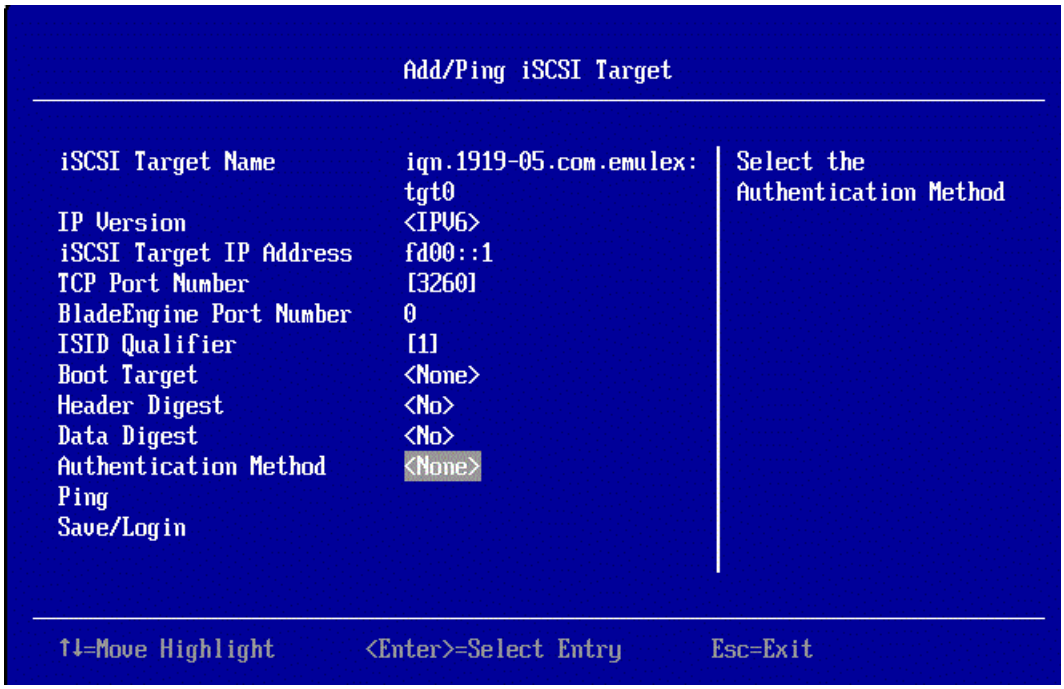
12.9.3 Manually Adding, Discovering, and Managing Boot Targets

To manually add and discover boot targets:

1. On the Controller Properties dialog (Figure 114), set Discover Boot Target via DHCP to `Disable` and press Enter.
2. Select Save Changes and press Enter.
3. Reboot the system.

4. On the iSCSI Target Configuration dialog (Figure 121), select Add Targets and press Enter. The Add/Ping iSCSI Target dialog appears.

Figure 123 Add/Ping iSCSI Target Dialog



5. Enter the target IP address and TCP port number (the default target port number is 3260).
6. If you are connecting dual sessions to the same target portal group, enter a unique ISID Qualifier, which can be any number up to 65535. For more information on when an ISID value must be changed, see Section 5.3.2, Booting from SAN for iSCSI MPIO.
7. Leave the Boot Target field set to No for now; you will return to this field later to enable the boot target.
8. Set the Header Digest to Yes if you want to enable the header digest. If the Header Digest is set to Yes, and the iSCSI initiator is set accordingly, the integrity of an iSCSI protocol data unit's header segment is protected by CRC32C checksum. The default is disabled.
9. Set the Data Digest to Yes if you want to enable the data digest. If the Data Digest is set to Yes, and the iSCSI initiator is set accordingly, the integrity of an iSCSI protocol data unit's data segment is protected by CRC32C checksum. The default is disabled.
10. Select an authentication method (optional). If you are enabling an authentication method, you are prompted to enter CHAP configuration. For more information about authentication methods, see Section 12.9.5, Selecting an Authentication Method.
11. Select Save/Login and press Enter to discover targets.

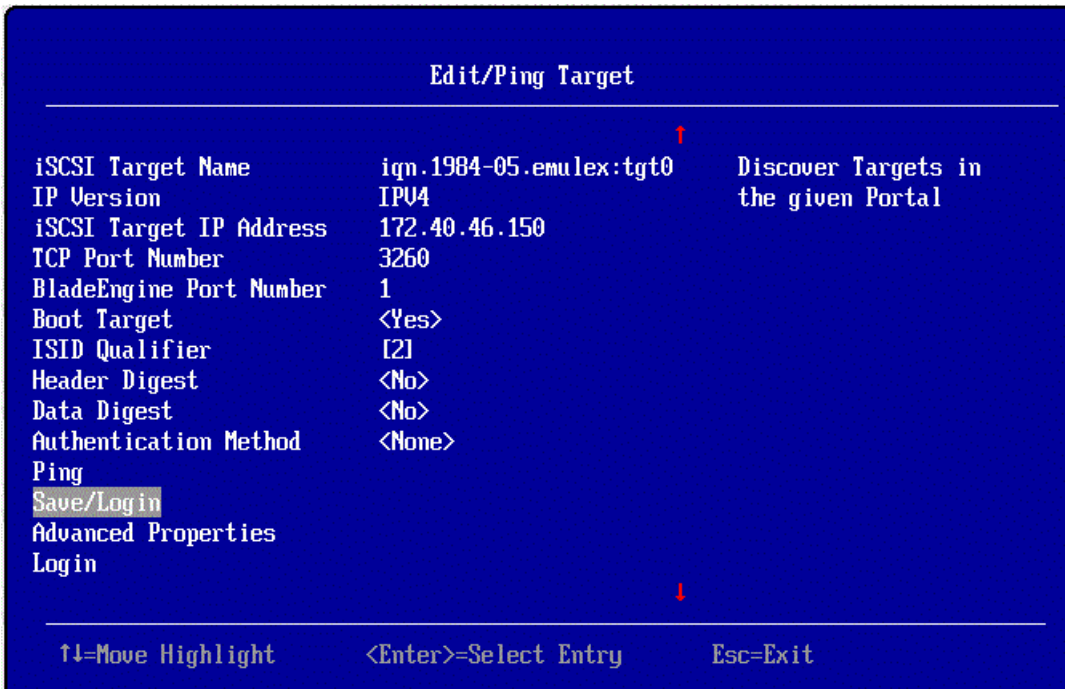
12.9.4 Setting a Boot Target

The discovered target must be set as a boot target to enable iSCSI boot.

To set a boot target:

1. On the iSCSI Target Configuration dialog (Figure 121), select the target name and press Enter. The Edit/Ping Target dialog (Figure 124) appears.

Figure 124 Edit/Ping Target Dialog



2. Select the Boot Target option and press Enter. The Boot Target pop up appears.
3. Select Primary or Secondary and press Enter.
4. Select Save/Login and press Enter.

12.9.5 Selecting an Authentication Method

The Emulex iSCSI configuration utility uses CHAP to authenticate initiators and targets for added network security. By using a challenge and response security mechanism, CHAP periodically verifies the initiator's identity. This authentication method depends on a secret key known to the initiator and the target only. Although the authentication can be one-way, you can negotiate CHAP in both directions with the help of the same secret set for mutual authentication. You must make sure however, that what you configure on the target side matches the initiator side. The Emulex iSCSI configuration utility supports both one-way and mutual authentication.

To configure the CHAP authentication method:

1. On the iSCSI Target Configuration dialog (Figure 121), select the target name and press Enter. The Edit/Ping Target dialog (Figure 124) appears.
2. Select the Authentication Method option and press Enter. The Authentication Method pop up appears.
3. Select None, One-way CHAP, or Mutual CHAP and press Enter.
   - **One-way CHAP** – With one-way CHAP authentication, the target authenticates the initiator. Use one-way CHAP authentication for a one-way challenge and response security method – you must configure the username and password (secret), which is authenticated by the target.
   - **Mutual CHAP** – With mutual CHAP authentication, the target authenticates the initiator and the initiator authenticates the target. Use mutual CHAP authentication for a two-way challenge and response security method – you must configure different CHAP and mutual CHAP usernames and passwords.

NOTE When you configure the CHAP parameters, verify that those parameters match the parameters on the storage side.

4. Select Save/Login and press Enter.

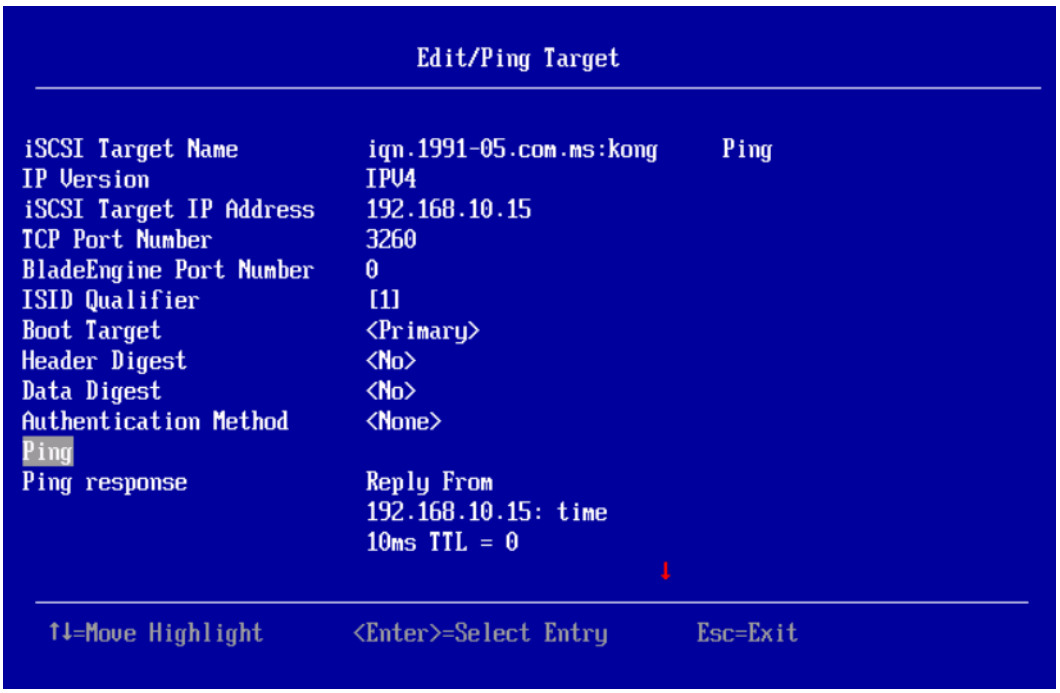
12.9.6 Pinging a Target

NOTE If you want to verify that you can connect to a target before you log in, you must ping the target before you configure the target on the Edit/Ping Target dialog.

To ping a target:

1. On the iSCSI Target Configuration dialog (Figure 121), select the target name and press Enter. The Edit/Ping Target dialog (Figure 124) appears.
2. Select Ping and press Enter. If the ping is successful, the system displays a reply message pop up.

Figure 125 Successful Target Ping Pop Up



If the ping is not successful, the system displays a failure message. A ping can be unsuccessful for several reasons; for more information, see Section 14.2, Troubleshooting for the iSCSI Protocol.

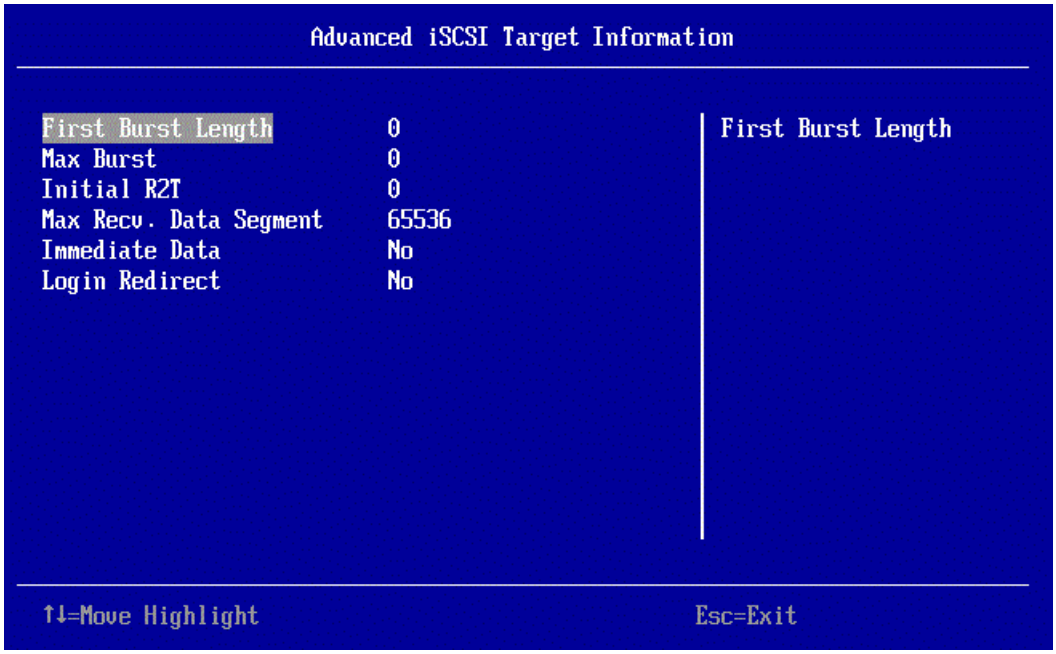
12.9.7 Viewing Advanced Properties

To view advanced properties:

1. On the Edit/Ping Target dialog (Figure 124), select Advanced Properties and press Enter. The Advanced iSCSI Target Information dialog (Figure 126) appears.

NOTE You may need to select ..more to show additional configuration options.

Figure 126 Advanced iSCSI Target Information Dialog



12.9.8 Logging In to or Out of a Target

To log in to or out of a target:

1. On the Edit/Ping Target dialog (Figure 124), select Login or Logout to explicitly log in to or log out of a target, and then press Enter.

NOTE You may need to select ..more to show additional configuration options.

12.9.9 Deleting a Target

NOTE A target can be deleted only if it is not a boot target.

If a target is set as a boot target:

1. On the Edit/Ping Target dialog (Figure 124), select Boot Target and press Enter. The Boot Target pop up opens.
2. Select No and press Enter.
3. Select Save/Login and press Enter.

To delete the target:

1. On the Edit/Ping Target dialog (Figure 124), select ..more to show additional configuration options.
2. Select Delete Target and press Enter.

12.9.10 Configuring LUNs

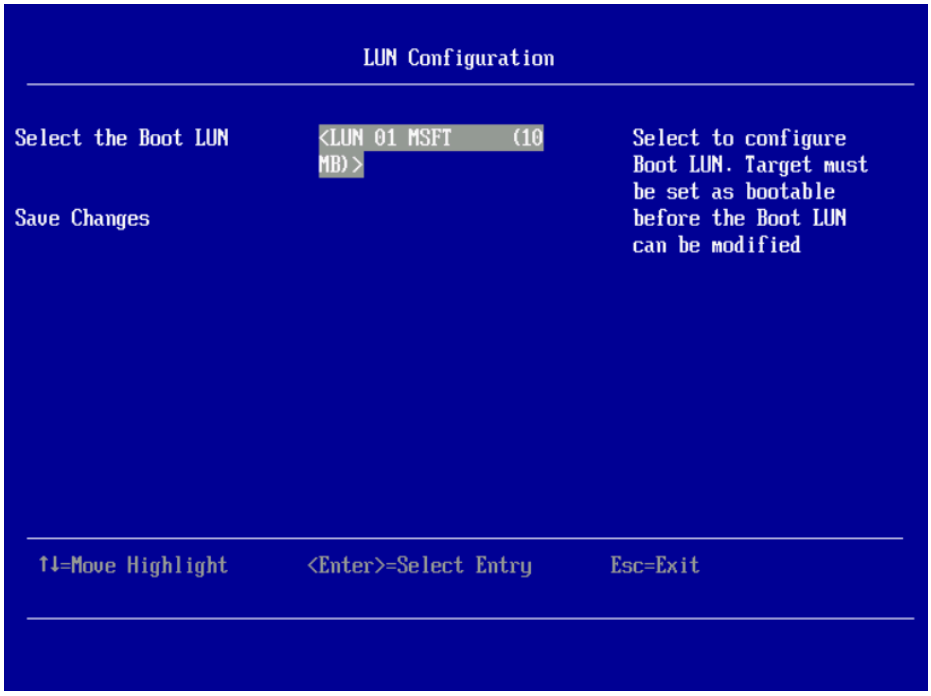
A LUN represents an individually addressable logical device that is part of a target. An initiator negotiates with a target to establish connectivity to a LUN.

To configure a LUN:

1. On the Edit/Ping Target dialog (Figure 124), select LUN Configuration and press Enter. The LUN Configuration dialog (Figure 127) opens.

NOTE You may need to select ..more to show additional configuration options.

Figure 127 LUN Configuration Dialog



2. If the target is a boot target, select any single LUN as a boot LUN from the dropdown list.

NOTE If the target is not a boot target, you cannot select any LUNs.

3. Select Save Changes and press Enter.

12.9.11 iSNS Configuration

NOTE This functionality is available only on UEFI-capable HP adapters.

The iSNS Protocol allows automated discovery, management, and configuration of iSCSI and FC devices (using iFCP gateways) on a TCP/IP network.

To view the iSNS Configuration dialog:

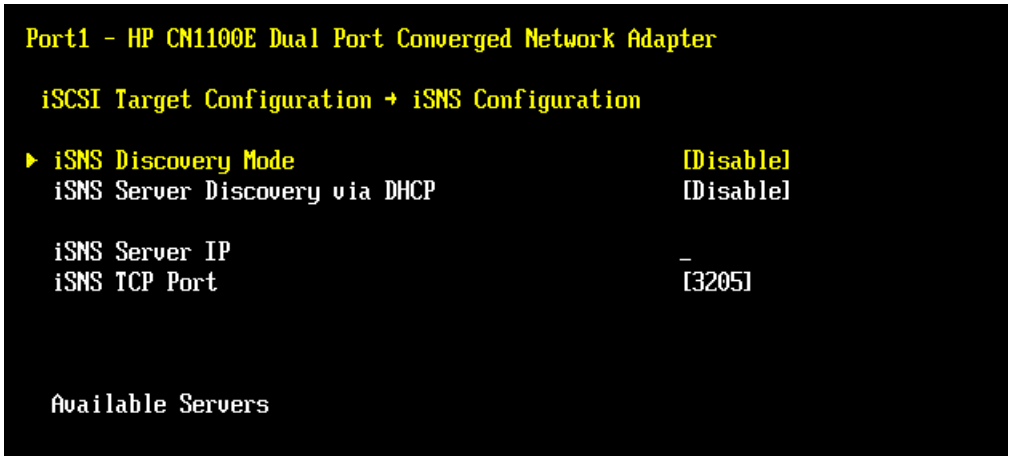
1. On the Controller Configuration Menu (Figure 112), select iSCSI Target Configuration and press Enter. The iSCSI Target Configuration menu (Figure 128) appears.

Figure 128 iSCSI Target Configuration Menu



2. Select iSNS Configuration and press Enter. The iSNS Configuration dialog appears.

Figure 129 iSNS Configuration Dialog



The iSNS Configuration dialog enables you to perform the following tasks:

- Enable or disable iSNS discovery mode
- Enable or disable iSNS server discovery using DHCP
- Configure the iSNS server IP address and TCP port number manually
- View a list of available iSNS servers

12.9.11.1 Enabling iSNS Discovery

This option enables or disables iSNS discovery. iSNS targets are only discovered if this option is enabled.

To enable or disable iSNS discovery:

1. From the iSNS Configuration dialog (Figure 129), select iSNS Discovery Mode and press Enter. A pop up appears.
2. Select either **Enable** or **Disable** and press Enter.
3. Press the F10 key to save the current settings.

12.9.11.2 Enabling iSNS Server Discovery Using DHCP

This option discovers iSNS servers configured on the network using DHCP.

NOTE If you leave iSNS Server Discovery via DHCP disabled (default), you must manually configure the iSNS server.

To enable or disable iSNS server discovery using DHCP:

1. From the iSNS Configuration dialog (Figure 129), select iSNS Server Discovery via DHCP and press Enter. A pop up appears.
2. Select either **Enable** or **Disable** and press Enter.
3. Press <F10> to save the current settings.

12.9.11.3 Configuring the iSNS Server IP Address and TCP Port Number Manually

An iSNS server can also be configured manually by entering the server IP address and the TCP port. Only one iSNS server can be configured at present.

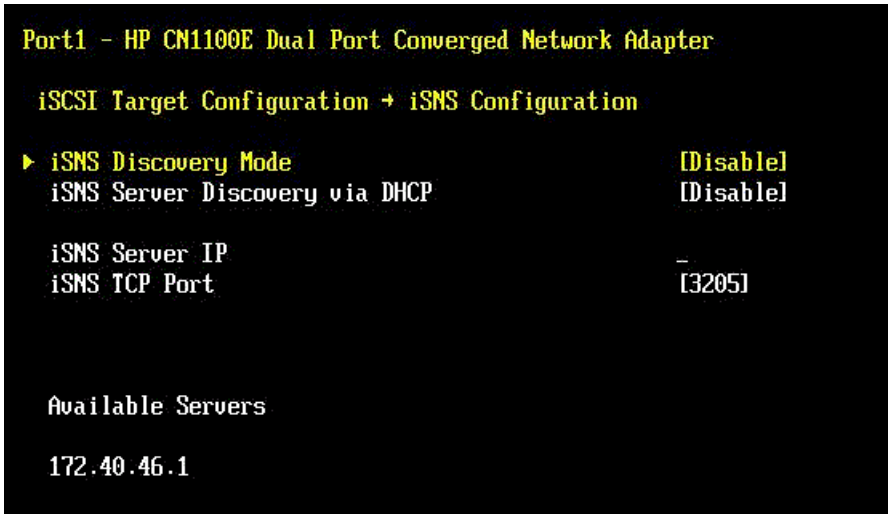
NOTE This option is available only if the iSNS Server Discovery via DHCP option is disabled.

To add an iSNS server manually:

1. From the iSNS Configuration dialog (Figure 129), select iSNS Server IP and press Enter.
2. Enter the iSNS server's IP address and press Enter.
3. Select iSNS TCP Port and press Enter. The default iSNS Port is 3205. The valid range is 1024 to 65535.
4. Press the F10 key to save the current settings.

   The added server appears under the Available Servers heading on the iSNS Configuration dialog as shown here.

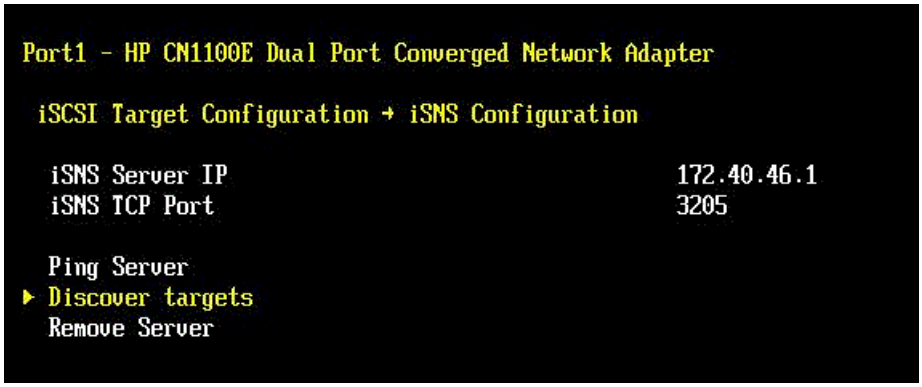
Figure 130 Available Servers Dialog



12.9.11.4 Discovering Targets Using the iSNS Server

To configure targets using the iSNS server, from the iSNS Configuration dialog, select the iSNS server under the Available Servers list (Figure 130) and press Enter. The iSNS server IP and TCP port information dialog appears (Figure 131).

Figure 131 iSNS Server Options Dialog



The following iSNS server options are available on this dialog:

- **Ping Server** – Select this option to check connectivity to the server.
- **Discover Targets** – Select this option to show a list of targets available on the iSNS server.

  One or more targets can be added by logging in to the targets individually. All added targets are listed under Discovered Targets on the iSCSI Target Configuration menu (Figure 128).
- **Remove Server** – Select this option to remove the server from the Available Servers list.

NOTE This removes the server only and not the targets added by this server. To remove the targets, select the target from the Discovered Targets list under the iSCSI Target Configuration menu (Figure 128), select Delete Target, and press Enter.

12.10 Erasing the Configuration

Erase Configuration erases the configuration of a single adapter. Configuration data is erased for both ports on the selected adapter. Erase configuration restores the default settings for that particular adapter only. If an initiator name is global for all adapters, the IQN configured on the first adapter is reflected on all adapters.

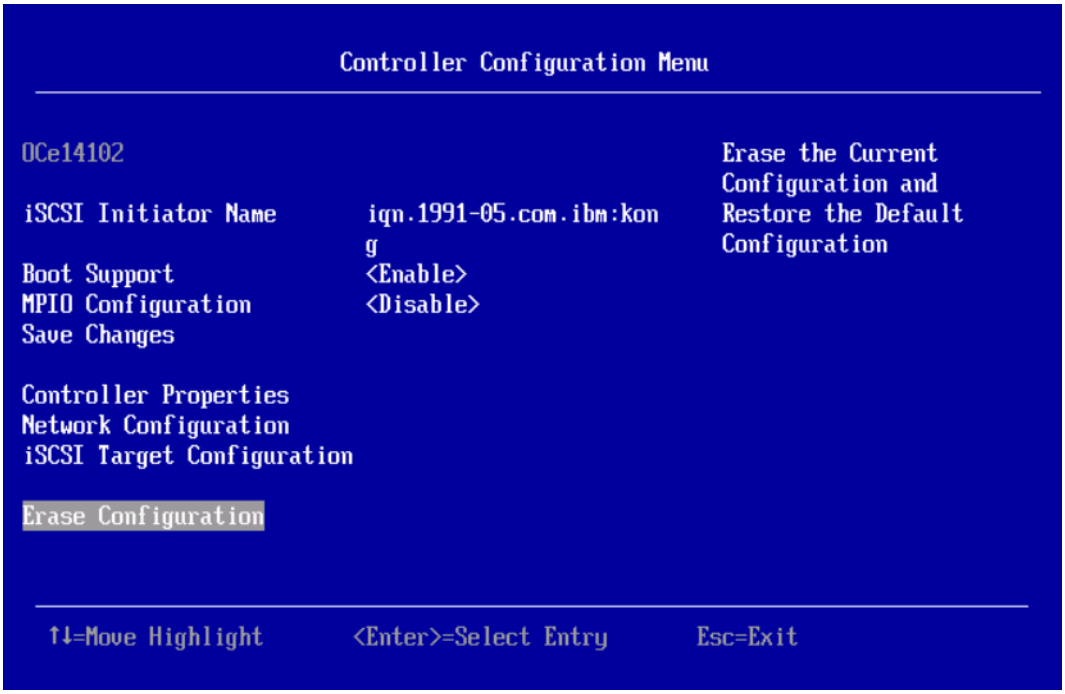
NOTE The Erase Configuration option is not available on HP adapters in HP systems. Instead, use the F7 key to reset the configuration settings to their factory default settings. For HP systems, the initiator name on all ports of the selected adapter are reset to their respective default values.

You must select Erase Configuration to clear out existing IQN data if you purchase a different or subsequent license for the adapter.

To erase an adapter configuration:

1. From the Controller Configuration Menu (Figure 112), select Erase Configuration and press Enter (Figure 132).

Figure 132 Erase Configuration Dialog



2. Press Enter to erase the configuration, or press Esc to stop the operation.
   When the adapter configuration is erased, the Controller Configuration Menu (Figure 112) appears.

Chapter 13: Configuring UEFI for FCoE

NOTE If you have several adapters in your system, the UEFI system firmware or boot code uses the highest version driver installed on any of the adapters. Adapters with older versions of EFIBoot are managed by the more recent version, but only as long as the adapter with the most recent version is in the system. The latest firmware and boot code must be installed on each adapter in the system to ensure that each adapter runs the latest firmware and boot code.

UEFIBoot supports:

- Fabric point to-point topology only
- All applicable UEFI 2.3 driver protocols
- All UEFI-aware operating systems
- Selectable fibre or SCSI device path through the driver configuration application
- Up to 256 ports
- Up to eight boot targets (selectable through the driver configuration application)
- Up to 4096 LUNs
- Multiutility – Setup and firmware updates

The Emulex FCoE UEFIBoot is loaded from the flash in to the system memory.

To verify that the Emulex UEFIBoot driver is loaded, enter the UEFI Shell, type `drivers`, and then press Enter.

NOTE If you are using a Dell adapter, see Appendix F, Dell UEFI for information on using the Dell UEFI utility.

13.1 Navigating the Emulex FCoE Configuration Utility

Use the following methods to navigate the FCoE configuration utility:

Use the following methods to navigate the Emulex NIC configuration utility:

- Press the up/down arrows on the keyboard to move through menu options or fields. If multiple adapters are listed, use the up/down arrows to scroll to the additional adapters.
- Press the +, -, or Enter keys to change numeric values.
- Press Enter to select a menu option, to select a row in a configuration dialog, or to change a configuration default.
- Use the navigation entries on the page to move about the utility.
- Press Esc to exit the current screen and show the previous screen.

NOTE HP adapters in HP systems use the F10 key to save the current configuration settings.

13.2 Starting the Emulex FCoE Configuration Utility

To start the Emulex FCoE configuration utility using HII:

1. Exit the UEFI shell.

Figure 133 Exiting the UEFI Shell

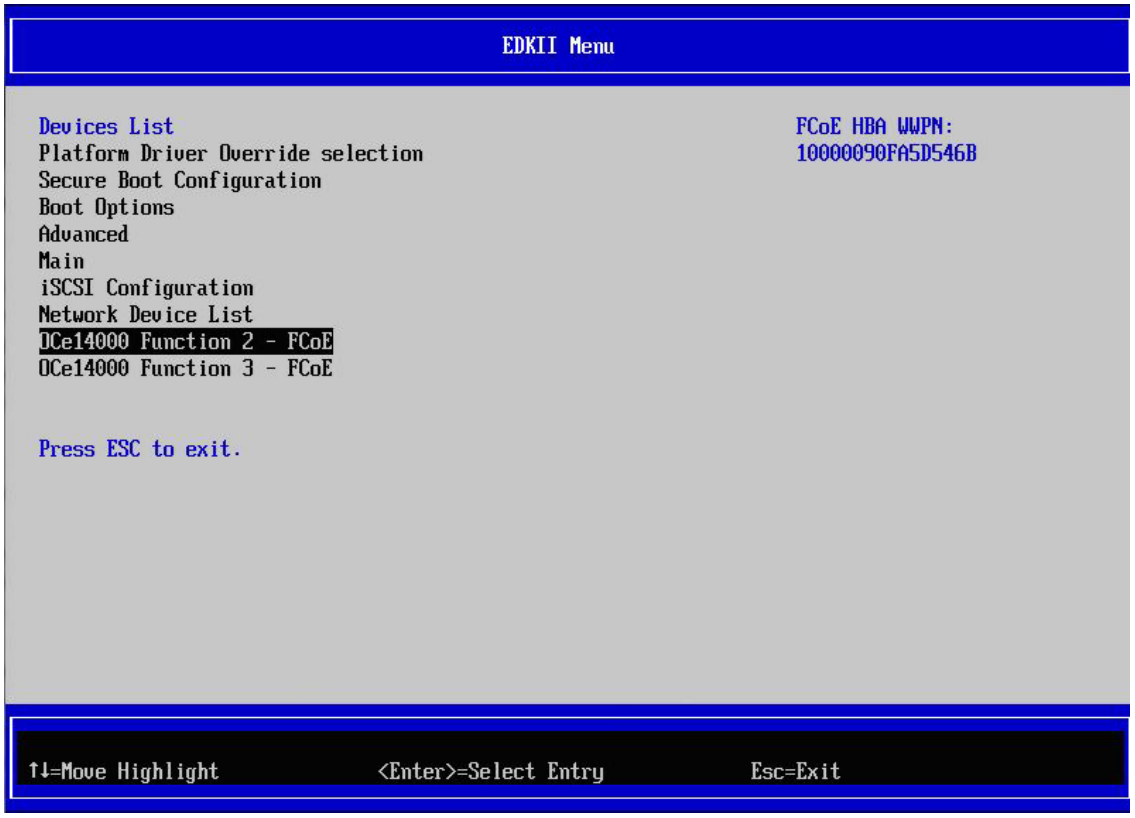


The system displays a front page.

NOTE Depending on the vendor UEFI configuration, the menu item on the front page that allows you to configure devices can be called Storage, Device Manager, Device Settings, or something else; see the documentation that accompanied the server.

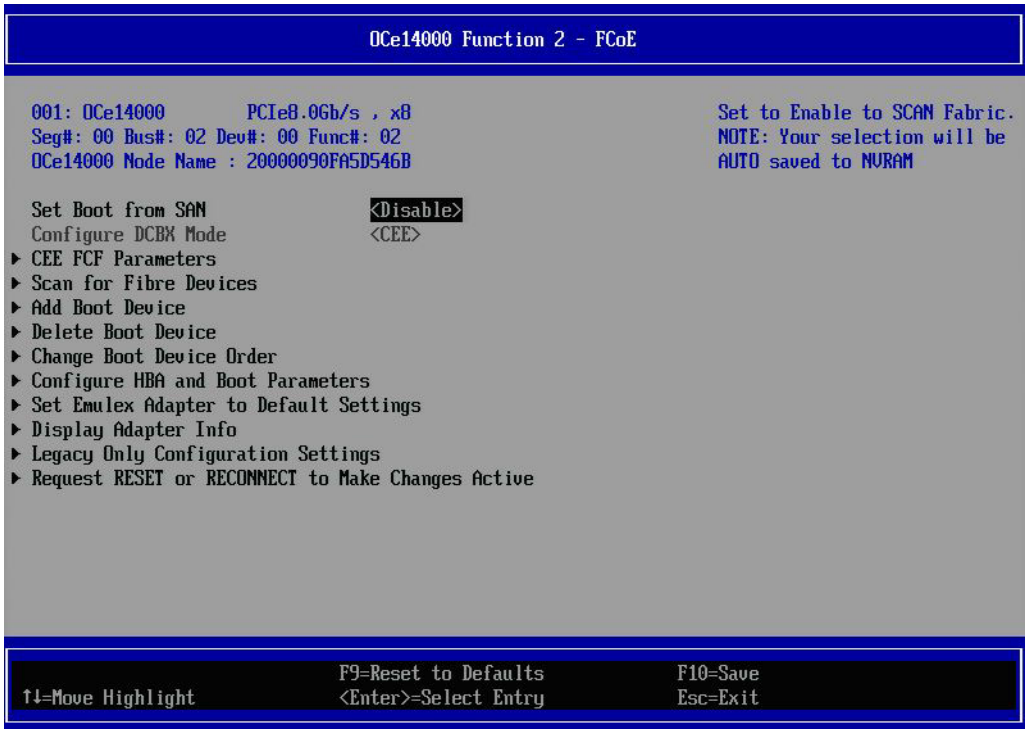
2. Select the menu item that lets you configure devices. In the Device Manager (Figure 134), a list of all the adapters in the system is shown. Your list may vary depending on the installed adapters.

Figure 134 Device Settings List



3. Select the adapter you want to configure and press Enter. The Emulex Adapter Configuration main menu (Figure 135) appears.

Figure 135 Emulex Adapter Configuration Main Menu



NOTE DCBX mode is always set to CEE and cannot be configured.ing

13.3 Updating Firmware and Boot Code

To update the firmware and boot code to the latest versions, you must use the NIC firmware update utility, which revises the FCoE function with a single firmware and boot code image. For more information on revising firmware versions, see Chapter 11, Configuring UEFI for Ethernet.

13.4 Enabling an Adapter to Boot from SAN

To enable an adapter to boot from SAN:

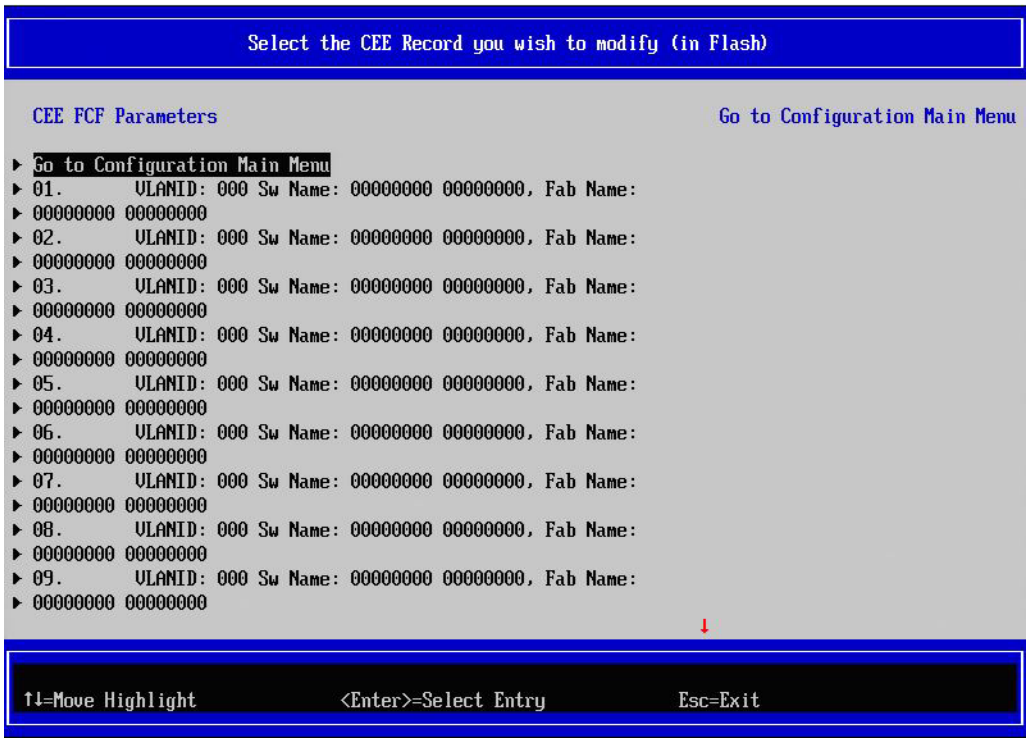
1. From the Device Manager list (Figure 134), select the adapter that you want to change the network boot setting for and press Enter. The Emulex Adapter Configuration main menu (Figure 135) appears.
2. Select Set Boot from SAN and press Enter. A pop up appears.
3. Select **Enable** and press Enter. The selection is changed to NVRAM and the current setting appears.

13.5 Configuring CEE FCF Parameters

To configure CEE FCF parameters:

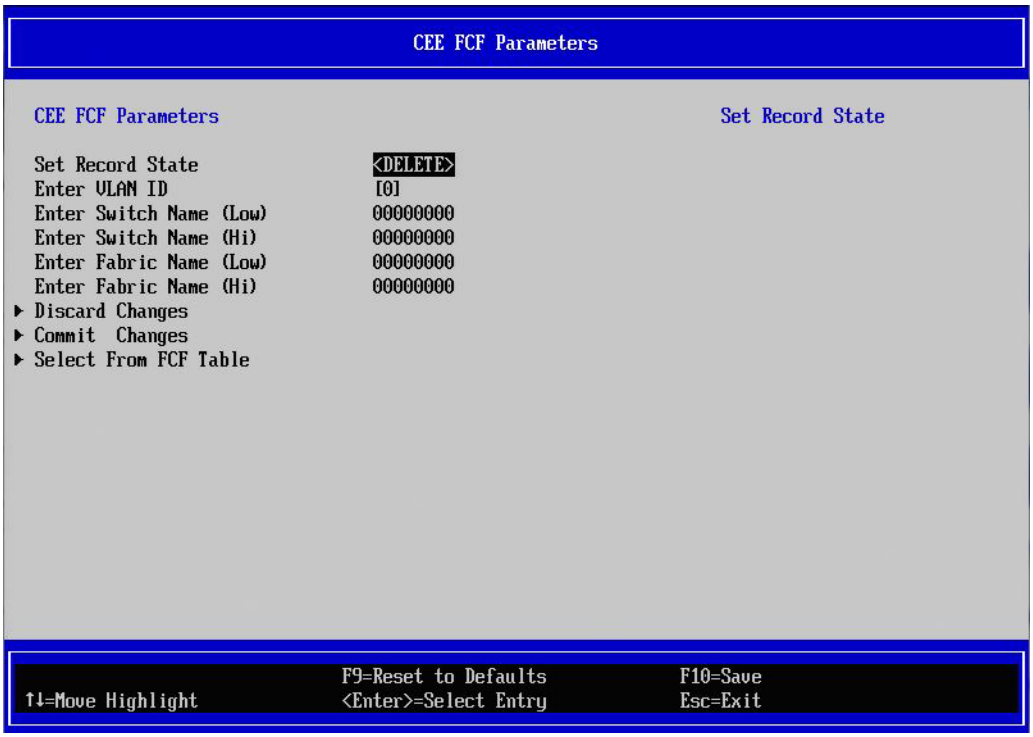
1. From the Device Manager list (Figure 134), select the adapter to configure and press Enter. The Emulex Adapter Configuration main menu (Figure 135) appears.
2. Select Configure CEE FCF Parameters and press Enter. The system displays a listing of the current CEE FCF record information.

Figure 136 CEE Record Selection List



3. Select the CEE FCF record to modify and press Enter. The system displays the current record information.

Figure 137 CEE FCF Record Information



4. Change the record information as needed.
   - Set Record State can be set to unused, active, or boot.
   - VLAN ID must be a three digit hexadecimal number.
   - Switch Name (Low). Enter the low bits of the FC switch's WWN to which to connect. This must be an 8-digit hexadecimal number.
   - Switch Name (Hi). Enter the high bits of the FC switch's WWN to which to connect. This must be an 8-digit hexadecimal number.
   - Fabric Name (Low). Enter the low bits of the FC fabric's WWN to which to connect. This must be an 8-digit hexadecimal number.
   - Fabric Name (Hi). Enter the high bits of the FC fabric's WWN to which to connect. This must be an 8-digit hexadecimal number.
5. Do one of the following:
   - To save your changes, select Commit Changes and press Enter. The changed CEE FCF record information appears.
   - To discard all changes, select Discard Changes and press Enter. The system displays the CEE record selection list (Figure 136).
   - To discard your changes and use the CEE parameters from the FCF table, select Select From FCF Table and press Enter. The system displays the FCF table parameters.
6. Press Enter. The Emulex Adapter Configuration main menu (Figure 135) appears.

13.6 Scanning for Fibre Devices

When LUNs are set up on the SAN before POST has completed (that is, before the boot driver has been started), you can select Scan for Fibre Devices or Add Boot Device to discover all available LUNs. Although this procedure does not perform a complete adapter initialization, it runs faster than selecting Reconnect Devices.

NOTE If you dynamically add LUNs after POST has completed and the driver has been started, you must select Reconnect Devices to perform a complete adapter initialization and discovery. If this step is not performed, all LUNs may not be properly discovered.

Select Reconnect Devices after adding any discovered LUNs to the NVRAM list.

To scan for fibre devices:

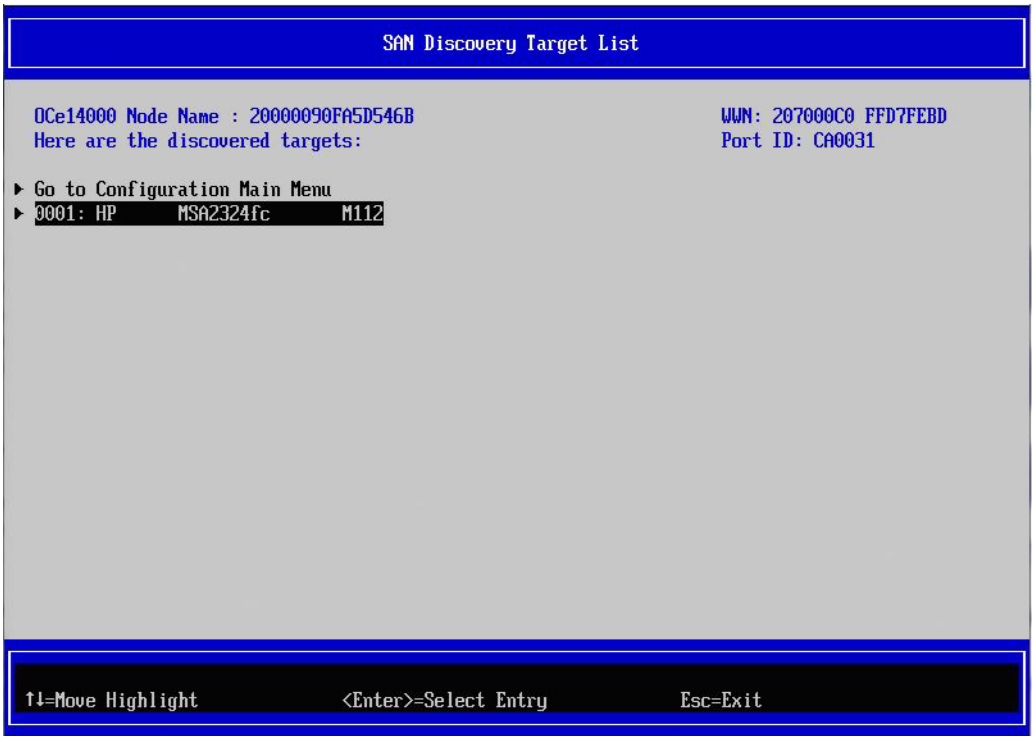
1. From the Device Manager list (Figure 134), select the adapter you want to scan for fibre devices and press Enter. The Emulex Adapter Configuration main menu (Figure 135) appears.
2. Select Scan for Fibre Devices and press Enter. A list of discovered targets appear. This list of discovered target devices lets you quickly determine SAN connectivity only and provides you with a mechanism to have the port logged in for zoning.

13.7 Adding Boot Devices

To add a boot device:

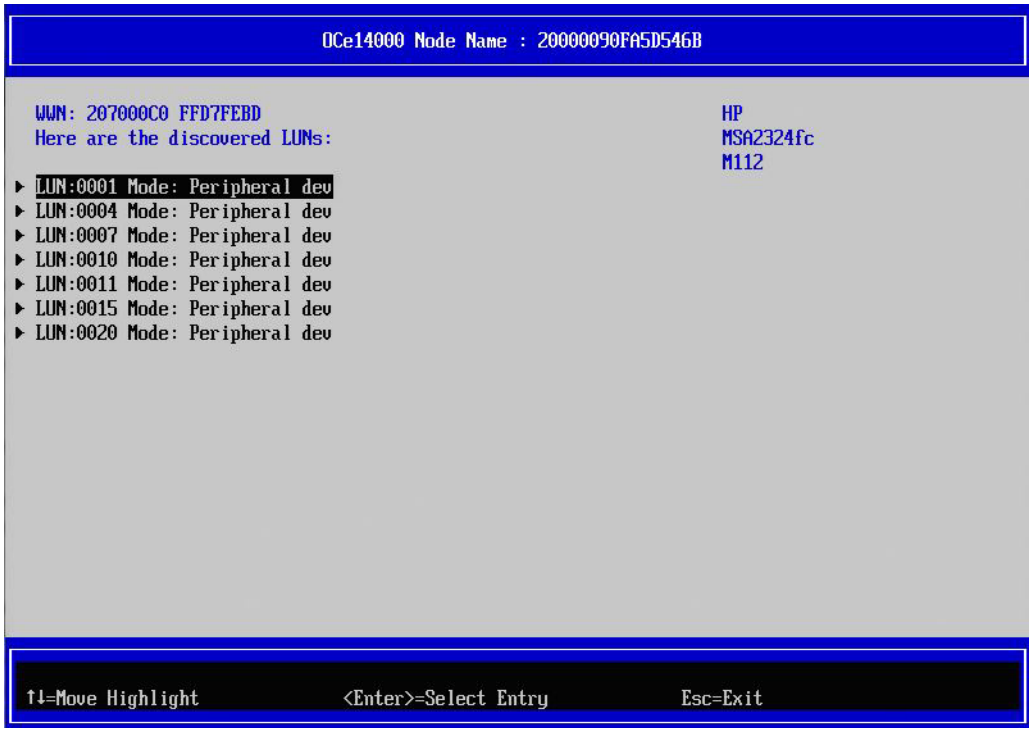
1. From the Device Manager list (Figure 134), select the adapter to which you want to add a boot device and press Enter. The Emulex Adapter Configuration main menu (Figure 135) appears.
2. Select Add Boot Device and press Enter. A list of discovered targets appears (Figure 138).

Figure 138 SAN Discovery Targets List



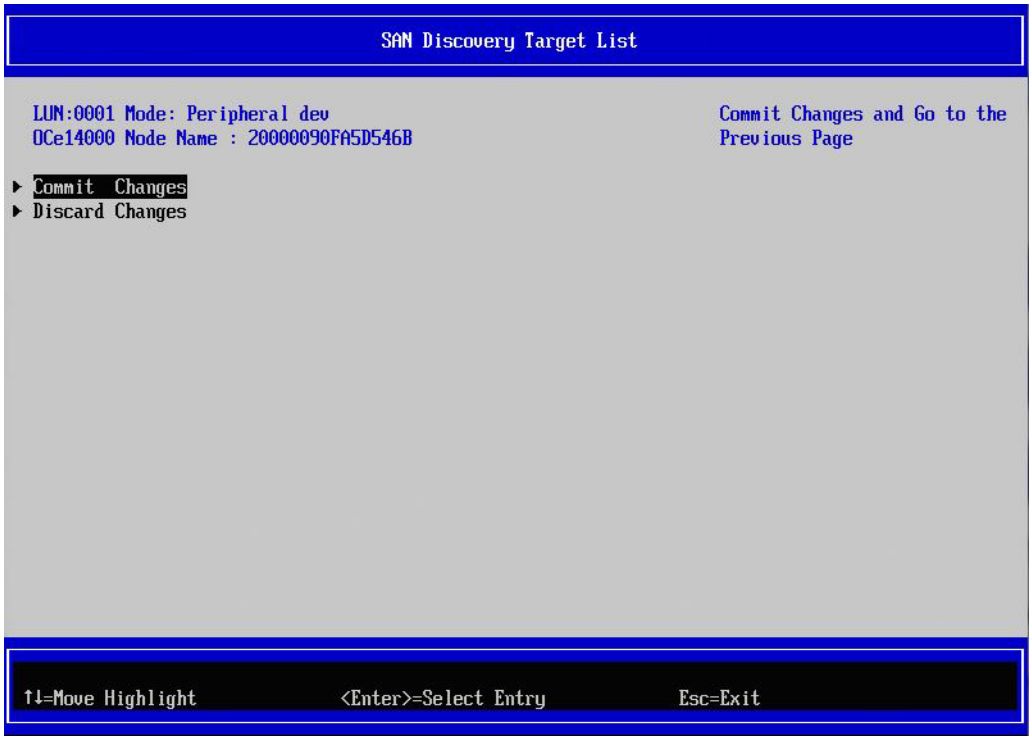
3. Select a target and press Enter. The LUN list appears.

Figure 139 LUN Listing



4. Select a LUN from the list and press Enter. The SAN Discovery Target List (Figure 140) appears.

Figure 140 SAN Discovery Target List



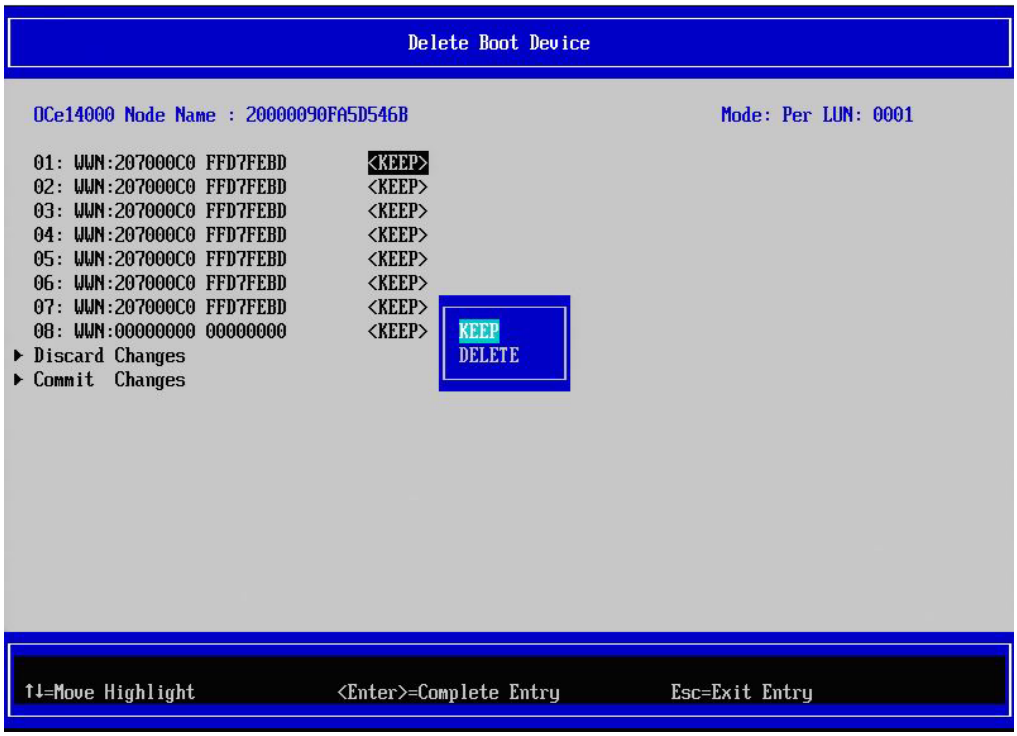
5. Select Commit Changes and press Enter to return to the Emulex Adapter Configuration main menu.

13.8 Deleting Boot Devices

To delete a boot device:

1. From the Device Manager list (Figure 134), select the adapter from which you want to delete a boot device and press Enter. The Emulex Adapter Configuration main menu (Figure 135) appears.
2. Select Delete Boot Device and press Enter. The system displays a list of eight boot devices.
3. Select the device and press Enter. The Delete Boot Device dialog for that device appears (Figure 141).

Figure 141 Delete Boot Device



4. Select DELETE and press Enter.
5. Select Commit Changes and press Enter to return to the Emulex Adapter Configuration main menu.

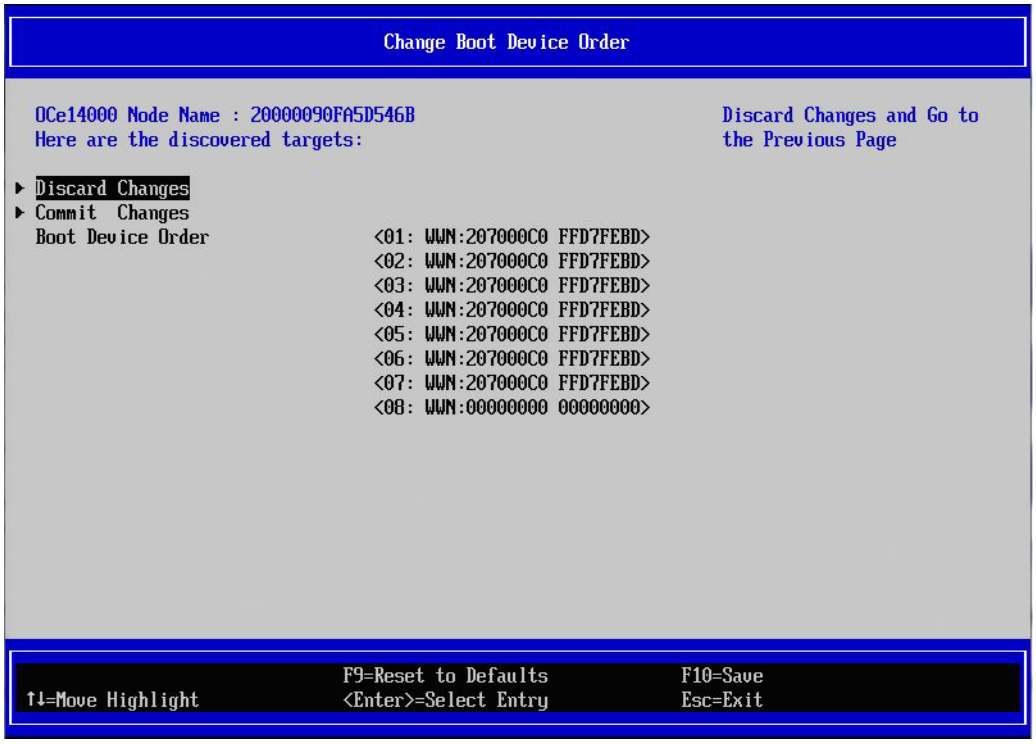
13.9 Changing Boot Device Order

NOTE The boot device order has no relationship to the system BIOS boot device order. Changing the boot device order with this procedure only changes the order in which the devices are discovered by UEFIBoot.

To change the boot device order:

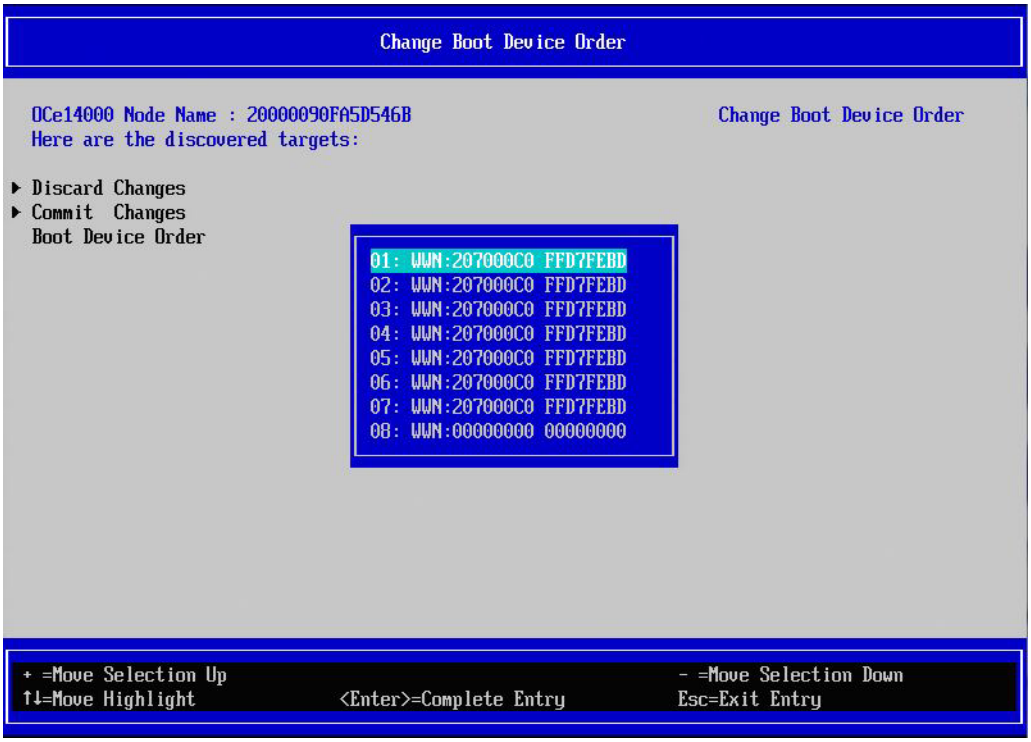
1. From the Device Manager list (Figure 134), select the adapter for the boot device order you want to change and press Enter. The Emulex Adapter Configuration main menu (Figure 135) appears.
2. Select Change Boot Device Order and press Enter. The Change Boot Device Order dialog (Figure 142) appears.

Figure 142 Change Boot Device Order Dialog



3. Select Boot Device Order and press Enter.
4. Select a device from the list of eight boot devices and press Enter.
5. Use the + or - keys to change the order of the selected device in the dialog box and press Enter. The Change Boot Device Order dialog (Figure 143) updates and displays the new boot device order.

Figure 143 Change Boot Device Order with Revised Boot Order



6. Verify that the boot device list order is correct or make any additional changes you want and press Enter.

13.10 Configuring Adapter Parameters

This section explains how to configure the boot parameters for an adapter.

13.10.1 Changing the Port Login Retry Timer

This option lets you set the interval for the port login 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 port login retry interval scans the loop to discover this device.

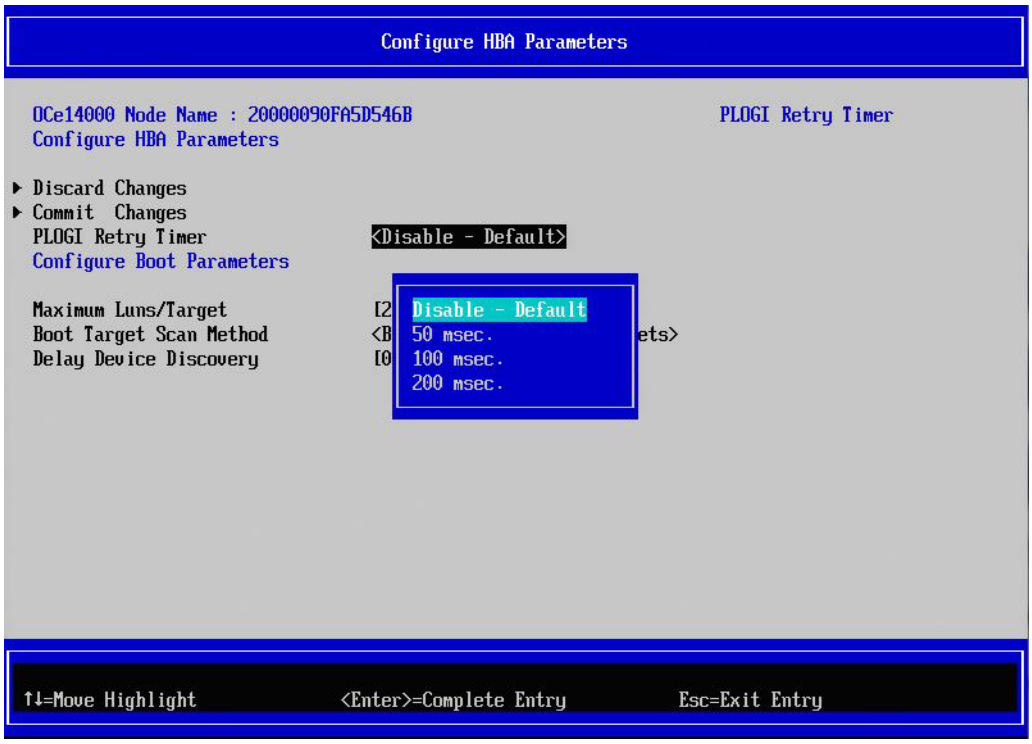
You can choose:

- Disable (default setting)
- 50 Msec
- 100 Msec
- 200 Msec

To change timer values:

1. From the Device Manager list (Figure 134), select the adapter you want to configure and press Enter. The Emulex Adapter Configuration main menu (Figure 135) appears.
2. Select Configure HBA Parameters and press Enter. The Configure HBA Parameters dialog appears (Figure 144).

Figure 144 Configure HBA Parameters Dialog – Port Login Retry Timer Pop Up



3. Select PLOGI Retry Timer and press Enter. The Port Login Retry Timer pop up appears.
4. Select a retry timer option and press Enter. The screen is refreshed with the modified value.

NOTE Press Esc to return to the UEFI utility menu.

5. Select Commit Changes and press Enter to save the changes and return to the Emulex Adapter Configuration main menu.

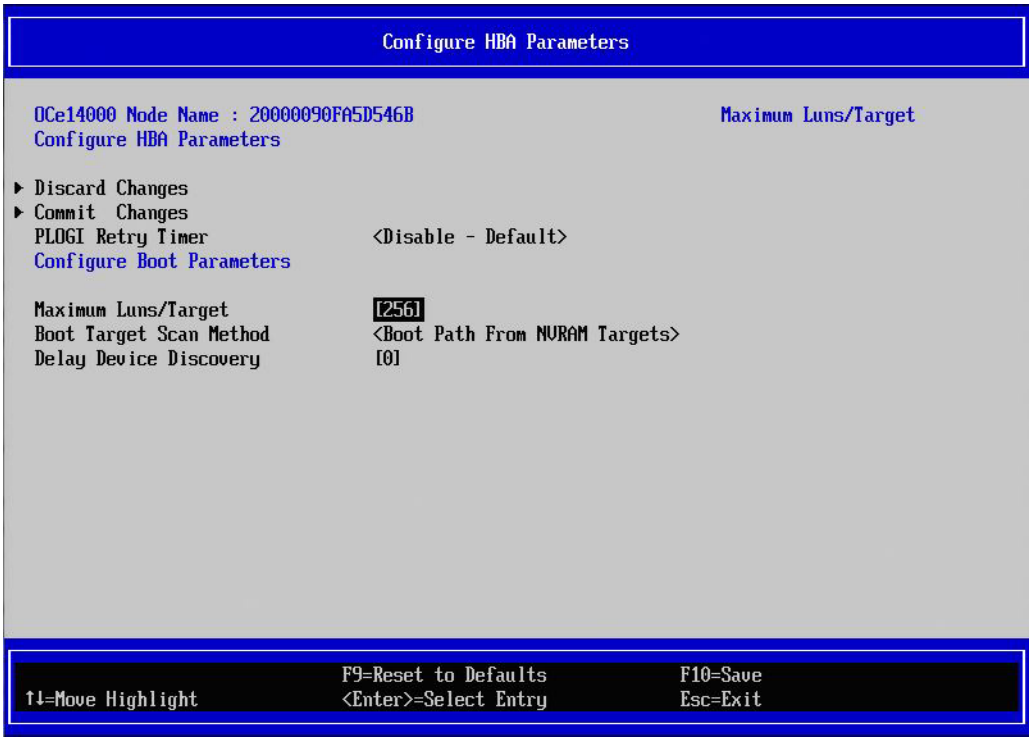
13.10.2 Changing the Maximum LUNs per Target

The maximum number of LUNs represents the maximum number of LUNs that are polled during device discovery. The minimum value is 1 and the maximum value is 4096. The default is 256.

To change the maximum number of LUNs:

1. From the Device Manager list (Figure 134), select the adapter for the LUNs per Target information you want to change and press Enter. The Emulex Adapter Configuration main menu (Figure 135) appears.
2. Select Configure HBA Parameters and press Enter. The Configure HBA Parameters dialog (Figure 145) appears.

Figure 145 Configure HBA Parameters Dialog



3. Select Maximum LUNs/Target and press Enter.
4. Enter a number from 1 to 4096 and press Enter. The screen is refreshed with the modified value.

   NOTE The default and typical maximum number of LUNs in a target device is 256. A higher number of maximum LUNs causes the discovery to take more time.

5. Select Commit Changes and press Enter to save the changes and return to the Emulex Adapter Configuration main menu.

13.10.3 Changing Boot Target Scan Method

This option is available only if none of the eight boot entries is configured to boot from the DID or WWPN. The Emulex Adapter Configuration main menu is used to configure up to eight boot entries. With boot scan enabled, the first device issues a Name Server Inquiry.

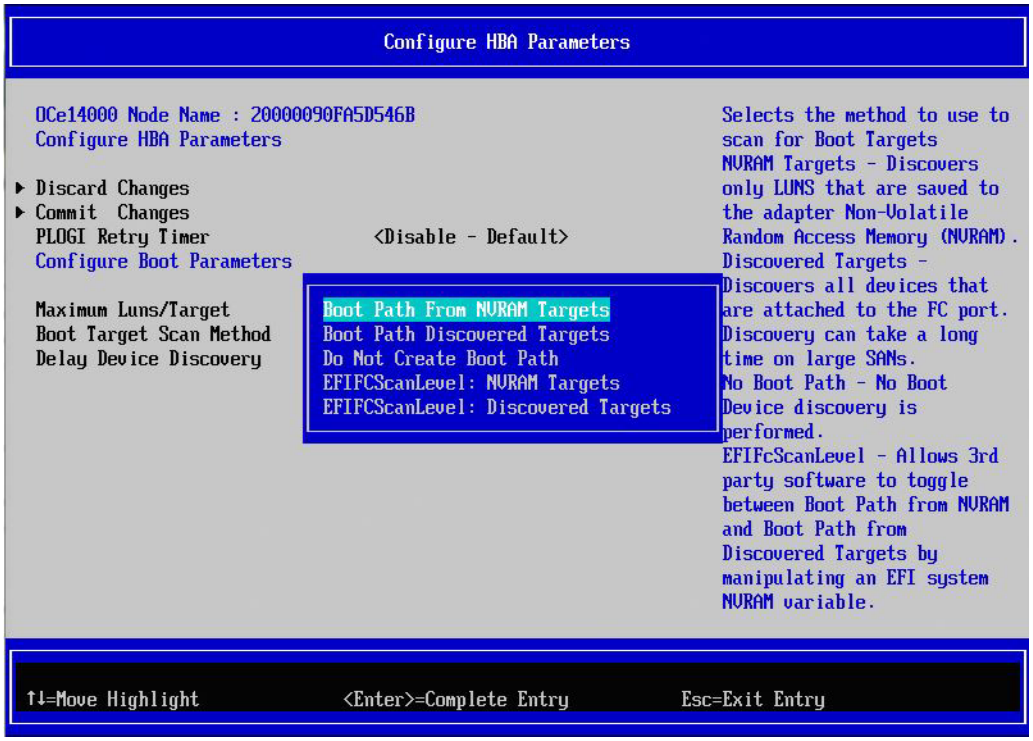
The boot scan options are:

- **Boot Path from NVRAM Targets** – Boot scan discovers LUNs saved to the adapter's NVRAM only. Select up to eight attached devices to use as potential boot devices. Limiting discovery to a set of eight selected targets can greatly reduce the time it takes for the EFIBoot driver to complete discovery.
- **Boot Path from Discovered Targets** – Boot scan discovers all devices attached to the FC port. Discovery can take a long time on large SANs if this option is used.
- **Do not create boot path**.
- **EFIFCScanLevel: NVRAM Targets** – Boot scan sets the EFIFCScanLevel environment variable to inform the driver to configure targets in the NVRAM boot table only.
- **EFIFCScanLevel - Discovered Targets** – Boot scan sets the EFIFCScanLevel environment variable to inform the driver to configure all available targets on the SAN.

To change the boot target scan method:

1. From the Device Manager list (Figure 134), select the adapter for the boot target scan method you want to change and press Enter. The Emulex Adapter Configuration main menu (Figure 135) appears.
2. Select Configure HBA Parameters and press Enter. The Configure HBA Parameters dialog (Figure 145) appears.
3. Select Boot Target Scan Method and press Enter. The Boot Target Scan Method pop up appears.

Figure 146 Configure HBA Parameters Dialog – Boot Target Scan Method Pop Up



4. Select a boot scan option and press Enter. The screen is refreshed with the modified value.

NOTE If you have a large SAN and set the boot path to Boot Path Discovered Targets, discovery takes a long time. Press Esc to return to the UEFI utility menu.

5. Select Commit Changes and press Enter to save the changes and return to the Emulex Adapter Configuration main menu.

13.10.4 Changing Device Discovery Delay

This parameter sets a delay to occur after a loop initialization and before a scan is initiated to discover the target. The default is off or 0 seconds. Change the default if you have an HP MSA1000 or HP MSA1500 RAID array, and if both of the following conditions exist:

- The MSA array is direct connected or part of an arbitrated loop (for example, daisy chained with a JBOD).
- The boot LUN is not reliably discovered. In this case, a delay may be necessary to let the array complete a reset.

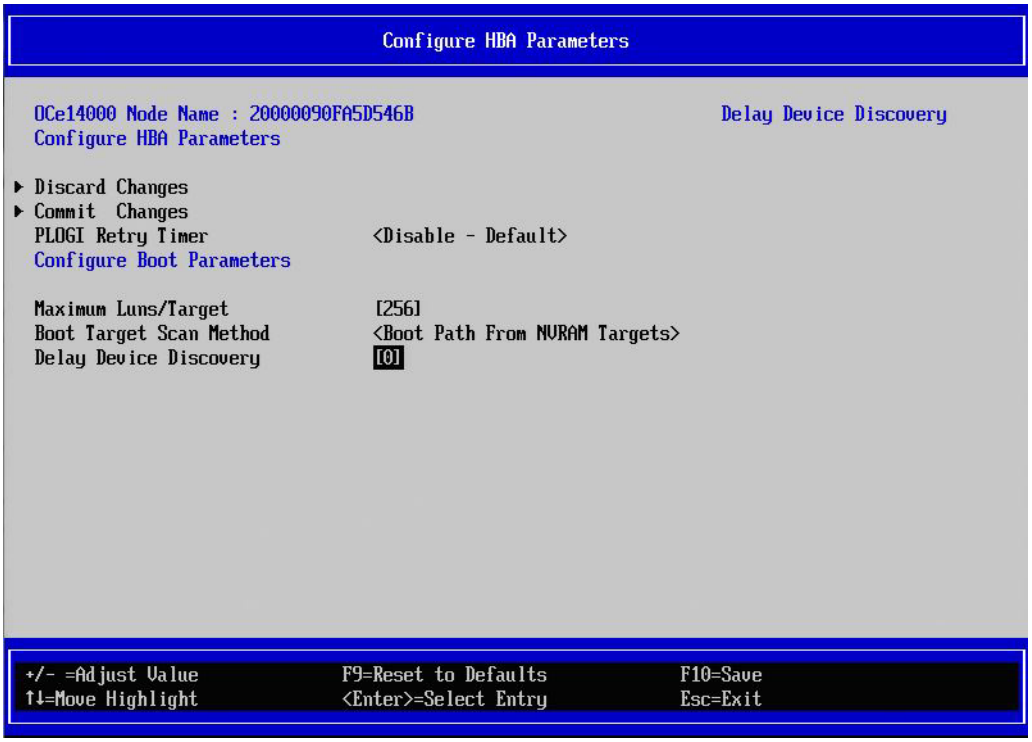
CAUTION Do not change the delay device discovery time if the MSA array is connected to a fabric switch. Setting it to any other time guarantees that the maximum delay time is seen on every loop initialization.

If both of the above conditions exist, typically set this parameter to 20 seconds. However, the delay should be only long enough for the array to be reliably discovered after a reset. Your value may be different.

To change the delay device discovery value:

1. From the Device Manager list (Figure 134), select the adapter for the device discovery delay setting you want to change and press Enter. The Emulex Adapter Configuration main menu (Figure 135) appears.
2. Select Configure HBA Parameters and press Enter. The Configure HBA Parameters dialog appears.

Figure 147 Configure HBA Parameters Dialog – Delay Device Discovery



3. Select Delay Device Discovery and press Enter. The Delay Device Discovery dialog appears.
4. Use the +/- keys to change the delay device discovery value in increments of 10 seconds and press Enter. The screen is refreshed with the modified value.
5. Select Commit Changes and press Enter to save the changes and return to the Emulex Adapter Configuration main menu.

13.11 Resetting Emulex Adapters to Default Values

The UEFI utility enables you to clear the NVRAM target list and set all boot device WWNNs back to 0, along with setting the adapter back to the default values. These default values are listed in the following table.

Table 10 Adapter Default Values

Parameter	Default Value	Valid Values
Enable/Disable BIOS	Disabled	Enabled Disabled
Port Login Retry Timer	Disabled	Disabled 50 msec 100 msec 200 msec
Boot Target Scan	Boot path from NVRAM targets	Boot path from NVRAM targets Boot path discovered targets Do not create boot path
Maximum LUNs Setting	0256	0–4096
Delay Device Discovery	0000	0000–0255
Advanced Mode	Enabled	Enabled Disabled
PXE Boot Support	The default for this parameter varies depending on the vendor configuration.	Enabled Disabled
SR-IOV	Disabled	Enabled Disabled
VLAN Support	Disabled	Enabled Disabled
VLAN ID	0	0–4094
VLAN Priority	0	0–7
UMC Support	The default for this parameter varies depending on the vendor configuration.	Enabled Disabled
Function En/Dis	Disabled	Enabled Disabled
Bandwidth	0%	Must have a total of 100% across all ports.
LPVID	0	2–4094
Switch Option (Lenovo System x Virtual Fabric-capable configuration, if available)	Virtual Fabric Mode	Virtual Fabric Mode Unified Fabric Protocol Mode Switch Independent Mode

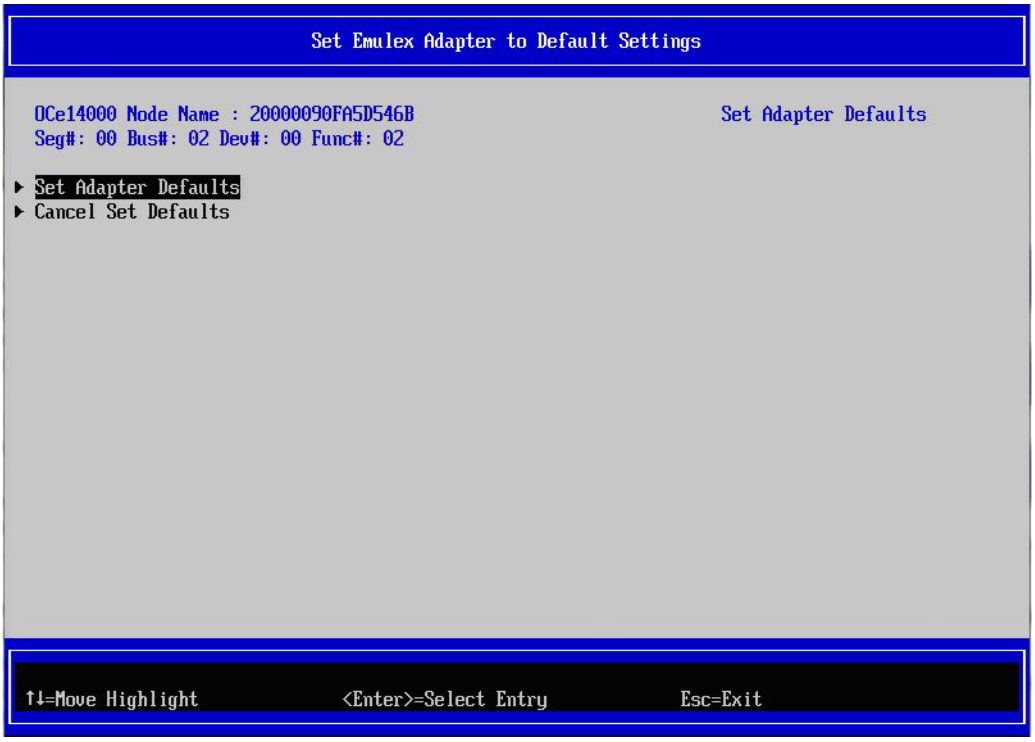
NOTE The following example sets the default values for FC/FCoE adapters only. To set other adapters to their default settings, you must use the utility for that specific protocol.

To set Emulex adapters to their default settings:

1. From the Device Manager list (Figure 134), select the adapter for the default settings you want to change and press Enter. The Emulex Adapter Configuration main menu (Figure 135) appears.

2. Select Set Emulex Adapter to Default Settings and press Enter. The Set Emulex Adapter to Default Settings dialog (Figure 148) appears.

Figure 148 Set Adapter Defaults Dialog



NOTE This sets the adapter to the FCoE driver default settings only.

3. Select Set Adapter Defaults and press Enter to set the adapter back to its default values. The Emulex Adapter Configuration main menu (Figure 135) appears.

13.12 Displaying Adapter Information

The adapter information shown contains the following information about the selected adapter:

- Adapter status
- Network boot status
- Link speed
- Topology
- Firmware version

NOTE Refer to this firmware version if reporting an issue to Broadcom Technical Support.

- Universal boot version
- EFI Boot version

To show adapter information:

1. From the Device Manager list (Figure 134), select the adapter for the information you want to view and press Enter. The Emulex Adapter Configuration main menu (Figure 135) appears.

2. Select Display Adapter Info and press Enter. The system displays information about the selected adapter(Figure 149).

Figure 149 Controller Information



13.13 Configuring Legacy Only Settings

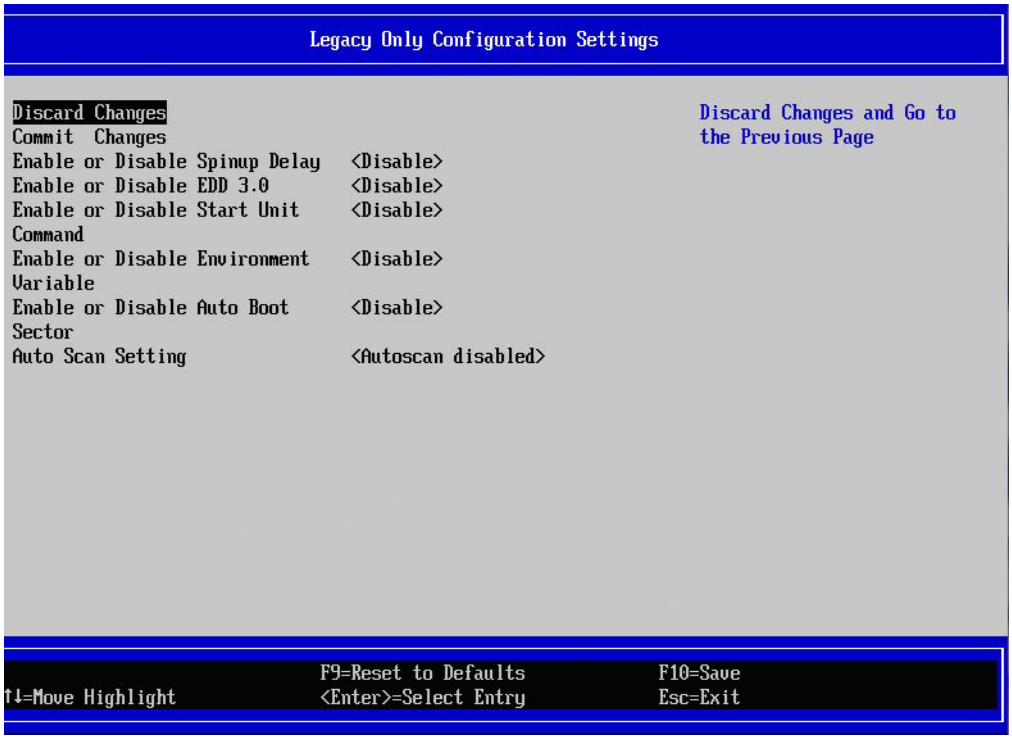
Configuration settings for some adapters are available only through the Legacy Only Configuration Settings dialog. The configuration settings include:

- Enabling or disabling spinup delay
- Enabling or disabling EDD 3.0
- Enabling or disabling start unit command
- Enabling or disabling environment variable
- Enabling or disabling auto boot sector
- Configuring the auto scan setting

To configure any of these settings:

1. From the Emulex Adapter Configuration main menu (Figure 135), select Legacy Only Configuration Settings and press Enter. The Legacy Only Configuration Settings dialog appears (Figure 150).

Figure 150 Legacy Only Configuration Settings Dialog



13.13.1 Enabling or Disabling the Spinup Delay

This option lets you enable or disable the disk spinup delay. The factory default setting is disabled.

If at least one boot device has been defined, and the spinup delay is enabled, the boot BIOS searches for the first available boot device.

- If a boot device is present, the boot BIOS boots from it immediately.
- If a boot device is not ready, the boot BIOS waits for the spinup delay and, for up to three additional minutes, continues the boot scanning algorithm to find another multiboot device.

If boot devices have not been defined and auto scan is enabled, then the boot BIOS waits for five minutes before scanning for devices.

- In an attached fabric, the boot BIOS attempts to boot from the first target found in the NameServer data.

To enable or disable the spinup delay:

1. On the Legacy Only Configuration Settings dialog (Figure 150), select **Enable or Disable Spinup Delay** and press Enter. A pop up appears.
2. Select **Enable** or **Disable** and press Enter.
3. Select Commit Changes and press Enter to save the changes and return to the Legacy Only Configuration Settings dialog.

13.13.2 Enabling or Disabling EDD 3.0

EDD 3.0 provides additional data to the operating system boot loader during an INT-13h function 48h call (get device parameters). This information includes the path to the boot device and the disk size. The default setting for EDD 3.0 is enabled. If EDD 3.0 is disabled, the setting reverts to EDD 2.1.

To enable or disable EDD 3.0:

1. On the Legacy Only Configuration Settings dialog (Figure 150), select **Enable or Disable EDD 3.0** and press Enter. A pop up appears.
2. Select **Enable** or **Disable** and press Enter.
3. Select Commit Changes and press Enter to save the changes and return to the Legacy Only Configuration Settings dialog.

13.13.3 Enabling or Disabling the Start Unit Command

You must know the specific LUN to issue the SCSI start unit command. The default setting is disabled.

To enable or disable the start unit command:

1. On the Legacy Only Configuration Settings dialog (Figure 150), select **Enable or Disable Start Unit Command** and press Enter. A pop up appears.
2. Select **Enable** or **Disable** and press Enter.
3. Select Commit Changes and press Enter to save the changes and return to the Legacy Only Configuration Settings dialog.

13.13.4 Enabling or Disabling the Environment Variable

Sets the boot adapter order if the system supports the environment variable. The default setting is disabled.

To enable or disable the environment variable:

1. On the Legacy Only Configuration Settings dialog (Figure 150), select **Enable or Disable Environment Variable** and press Enter. A pop up appears.
2. Select **Enable** or **Disable** and press Enter.
3. Select Commit Changes and press Enter to save the changes and return to the Legacy Only Configuration Settings dialog.

13.13.5 Enabling or Disabling Auto Boot Sector

This option automatically defines the boot sector of the target disk for the migration boot process, which applies to HP MSA1000 arrays only. If there is no partition on the target, the default boot sector format is 63 sectors. The default setting is disabled.

To enable or disable the auto boot sector format:

1. On the Legacy Only Configuration Settings dialog (Figure 150), select **Enable or Disable Auto Boot Sector** and press Enter. A pop up appears.
2. Select **Enable** or **Disable** and press Enter.
3. Select Commit Changes and press Enter to save the changes and return to the Legacy Only Configuration Settings dialog.

13.13.6 Configuring the Auto Scan Setting

This option lets you enable or disable the auto scan feature. Auto scan is available only when all eight boot entries are unused (empty). The factory default for auto scan is disabled. Auto scan lets you boot from the first boot device found in the FC fabric with the following options:

- **Autoscan disabled** – Default

- **Any first device** – The first adapter scans the fabric and the first disk discovered becomes the attempted boot device
- **First LUN 0 device** – The first device discovered with LUN 0
- **First NOT LUN 0 device** – The first device discovered with a LUN other than 0

To configure auto scan:

1. On the Legacy Only Configuration Settings dialog (Figure 150), select Auto Scan Setting and press Enter. A pop up appears.
2. Select the appropriate auto scan option and press Enter.
3. Select Commit Changes and press Enter to save the changes and return to the Legacy Only Configuration Settings dialog.

13.14 Requesting a Reset or Reconnect

After you have made certain changes to the UEFI BIOS, you may need to reset or reconnect the adapter for the changes to take effect.

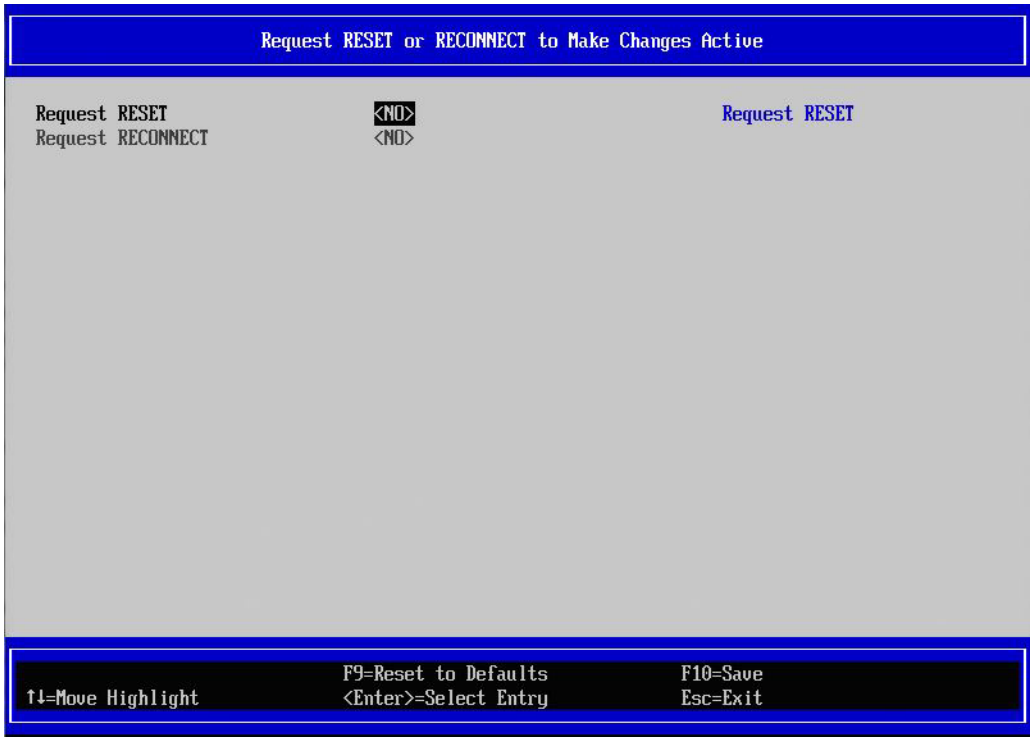
Reset forces the system to reboot after you exit the UEFIBoot (HII) utility.

Reconnect unloads and reloads the driver. This allows any changes to take effect immediately without requiring a reboot.

To request a reset or reconnect, do the following:

1. From the Emulex Adapter Configuration main menu (Figure 135), select Request RESET or RECONNECT to Make Changes Active. The Request RESET or RECONNECT to Make Changes Active dialog appears (Figure 151).

Figure 151 Request RESET or RECONNECT to Make Changes Active Dialog



2. Do one of the following:

- To request a reset:
  a. Set Request RESET to yes to notify the system BIOS that you want to perform a system reset.
  b. Exit the UEFIBoot (HII) utility and return to the system BIOS setup screens.
  c. A pop up window prompts you to reset the system. Any changed Emulex adapter settings become active during the subsequent reboot.
- To request a reconnect:
  a. Set Request RECONNECT to yes to notify the system BIOS that you want to perform a disconnect and reconnect on the adapter port.
  b. The changed Emulex adapter settings are activated without a system reset.

Chapter 14: Troubleshooting

This section provides troubleshooting information for the supported networking protocols.

14.1 Troubleshooting for the NIC Protocol

This section provides troubleshooting information for the PXESelect utility.

14.1.1 Issues During the PXE Boot Process

Situation: If any of the following occur during the boot process:

- The PXE boot ROM sends a DHCP request extended with PXE specific options, but the request is not received.
- The DHCP responses do not contain the DHCP options (DHCPOFFERs) that include the NBP filename and boot server lists.
- The PXE client attempts to download the specified NBP over TFTP from one of the specified boot servers, but fails.
- The PXE client does not execute the downloaded NBP.

Outcome: If any of these issues occur, the boot process continues using the next available device in the boot sequence.

14.1.2 PXE Error Codes

Table 11 PXE Error Codes

Error Codes	Possible Reason	Comment/Solution
Init Error Codes		
PXE-E01: PCI Vendor and Device IDs do not match!	The PCI BIOS passes Bus/Device/Function numbers in AX register when it initializes the OptionROM. If the vendor ID & device ID did not match the vendor ID and the device ID in the PCI configuration space for the device, with the UNDI device PCI structure defined in UNDI driver initialization code, the system displays this error.	This message should never be seen in the production level BIOS.
PXE-E04: Error reading PCI Configuration space	Any of the PCI BIOS INT 1Ah calls to read the PCI configuration space failed.	This should not happen with any production level motherboard BIOS.
PXE-E07: PXE BIOS calls not supported.	The PCI system BIOS does not support INT 1Ah services.	This should not happen with any PCI BIOS as all PCI BIOS must support these services.
ARP/TFTP Error Codes		
PXE-E11:ARP timeout	An ARP request is tried four times and the client did not receive an ARP reply.	Check the DHCP server settings. If you are connected through a switch, also check the switch settings.
PXE-E32: TFTP Open timeout.	A DHCP address and boot file name has been obtained and attempts to download the file name from a remote server.	■ Check that the TFTP is configured properly. ■ Check that the filename is properly placed in the tftpboot directory. ■ Check that the PXE server is configured for TFTP not MTFTP.
PXE-E35: TFTP read timeout.	A TFTP server is not able to receive the next packet.	Check that the hardware and TFTP server are functioning properly.
PXE-E36: Error received from TFTP server.	A TFTP server sends an error packet.	Check the TFTP server settings and boot file settings.
PXE-E38: TFTP cannot open connection.	A TFTP client is not able to open a TFTP connection with the TFTP server.	Check that the correct boot file is copied into the default TFTP boot path.
PXE-E3B: TFTP error-File not found	A requested boot file is not found on TFTP server.	Check that the correct boot file is copied into the default TFTP boot path.
PXE-E3C: TFTP error - Access violation.	The TFTP server does not have enough access rights to open or read the file.	Check the file permissions for the TFTP server.
PXE-E3F: TFTP packet size is invalid	TFTP packet received is larger than 1456 bytes.	Check the MTU on the Server side NIC interface.
Boot/DHCP Error Codes		
PXE-E51: No DHCP or Proxy DHCP offers were received.	The client did not receive any valid DHCP, BOOTP, or Proxy DHCP offers.	Check that the boot file name is configured correctly for PXE Clients in the DHCP setup.

Table 11 PXE Error Codes (Continued)

Error Codes	Possible Reason	Comment/Solution
PXE-E52: Proxy DHCP offers were received. No DHCP offers were received.	The client did receive at least one valid proxyDHCP offer, but did not receive any valid DHCP or BOOTP offers.	Check that the DHCP server and PXE server are configured properly.
PXE-E53: No boot filename received.	The client did receive a valid BOOTP/DHCP offer, but does not have a boot file name to download in the DHCP offer.	Check that the boot file name is configured correctly for PXEClients in the DHCP setup.
PXE-E55: Proxy DHCP service did not reply to request on port 4011.	The client issued a proxy DHCP request to the DHCP server on port 4011 but did not receive a reply.	Check that the DHCP server and TFTP server are properly configured on different servers. This issue may happen only if both are configured on separate servers.
UNDI Error Codes		
PXE-E60: Invalid UNDI API function number	The underlying UNDI drivers do not support the UNDI API sent from BaseCode/NBP.	Check with the hardware vendor.
PXE-E61: Media test failed, check cable.	No physical link exists on the PXE client port on which the boot is attempted.	Check the cables on the NIC interface.
PXE-E63: Error while initializing the NIC.	An issue in initializing the hardware exists as part of the UNDI_Initialize API.	Check with the hardware vendor.
BaseCode/UNDI Loader Error Codes		
PXE-EC1: BaseCode ROM ID structure was not found.	The UNDI boot module can not find the BaseCode ROM ID structure. This can be caused by a corrupted BaseCode ROM image.	
PXE-EC3: BaseCode ROM ID structure is invalid.	The BaseCode ROM ID structure is invalid. This can be caused by a corrupted BaseCode image in ROM.	
PXE-EC4: UNDI ROM ID structure was not found.	The BaseCode loader module could not locate the UNDI ROM ID structure. The UNDI ROM image is probably corrupted.	
PXE-EC5: UNDI ROM ID structure is invalid.	The UNDI ROM image is probably corrupted.	
PXE-EC8:!PXE structure was not found in UNDI driver code segment.	The UNDI ROM image is probably corrupted.	
Bootstrap and Discovery Error Codes		
PXE-E74: Bad or missing PXE menu and /or prompt information	PXE tags were detected but the boot menu and/or boot prompt were not found or were invalid.	
PXE-E77: Bad or missing discovery server list.	Two possible reasons. One is that Multicast and Broadcast discovery are both disabled. The other is that Multicast and Broadcast are enabled, but the server list tag is not found or is invalid.	

Table 11 PXE Error Codes (Continued)

Error Codes	Possible Reason	Comment/Solution
PXE-E78: Could not locate boot server.		
PXE-E79: NBP is too big to fit in free base memory.	The NBP file size is larger than the free base memory.	Check the NBP file size.
Miscellaneous Error Codes		
PXE-EA0: Network boot cancelled by keystroke	Pressing Ctrl+C keys during DHCP discovery.	

14.2 Troubleshooting for the iSCSI Protocol

This section includes troubleshooting information for the iSCSISelect utility.

Table 12 Troubleshooting the iSCSISelect Utility

Issue	Answer/Solution
iSCSI BIOS banner is not displayed during system post.	■ Check the motherboard BIOS configuration and make sure that the Option ROM is enabled and is set for the PCIe slot into which the adapter is inserted. ■ If the Option ROM is enabled and is set for the PCIe slot into which the adapter is inserted and the issue remains, erase the adapter configuration. For more information, Section 9.5, Erasing the Current Configuration.
Unable to ping iSCSI target IP address	■ Check cable connections and make sure they are securely connected at both ends. Make sure the link light is ON at both ends of the cable. ■ Verify that you have assigned a valid IP address with the correct subnet mask to the interfaces connected to the network. ■ Check for duplicate IP addresses. ■ Make sure you are connected to the correct port. ■ If the IP address is coming from DHCP, check that the DHCP server is up and available. ■ Is the target on a different subnet? Check the default gateway and make sure the IP addresses of both the default gateway and the remote host are correct. ■ Check link status under the Network Configuration dialog. If you try changing the initiator link status, you can back out one screen then go back to the Network Configuration dialog to see a view of the current link status.
Login to iSCSI target fails or login succeeds, but no LUNs are available.	■ Check the initiator IQN name and the target IQN name. Verify that they are properly specified. ■ Check the target's ACL settings (if any) to verify that the initiator's IQN name is listed and can be recognized. ■ Make sure that the target login parameters are compatible with the initiator's. ■ If this checklist confirms that initiator and target are configured correctly and the issue remains, erase the adapter configuration. For more information, see Section 9.5, Erasing the Current Configuration.
Login to iSCSI target is successful, but iSCSI I/O causes the system to hang or freeze.	Check the jumbo frame settings on the iSCSI target. If enabled, change the frame size to 1514 and retry. This isolates any issues in the network related to jumbo frames. Jumbo frames, if supported by the entire network, provide increased performance. **NOTE** Emulex OneConnect adapters support jumbo frames.

Table 12 Troubleshooting the iSCSISelect Utility (Continued)

Issue	Answer/Solution
Unable to boot from the iSCSI target.	■ Check the target setup. Check the target ACL to verify that the initiator's iSCSI name is listed and can be recognized. ■ Check the initiator name and make sure it is correct. ■ If applicable, verify that you have selected the iSCSI LUN as the boot LUN in the system BIOS setting. ■ Check the system BIOS for boot device priority order. Make sure that the boot device is at the top. If using DHCP: ■ If Boot Support is not enabled, enable it on the iSCSI Initiator dialog. (For more information, see Chapter 9, Configuring and Managing the iSCSI Initiator with the iSCSISelect Utility.) ■ If DHCP boot support is enabled, check the DHCP setup and also verify that you have added the root path to the DHCP server. For more information, see Section 8.2, Setting Up a Basic iSCSI Boot Configuration. ■ If this checklist confirms that the initiator and target are configured correctly and the issue remains, erase the adapter configuration. For more information, see Section 9.5, Erasing the Current Configuration.
BIOS post shows `BIOS Not Installed` message.	■ If you *have not* configured an iSCSI boot target, this is proper normal behavior. ■ If you *have* configured an iSCSI boot target and the BIOS cannot find a bootable LUN, then refer to the solutions provided in the issue *Unable to boot from the iSCSI target*.
If solutions for issues 1-3 and 5-6 all fail.	Use the iSCSISelect utility to erase the adapter configuration. To clear the adapter configuration: 1. From the Controller Configuration menu, select Erase Configuration and press Enter. 2. The system displays a message asking if you want to erase the current configuration. Press Y. 3. You are cautioned that the operation removes any existing configuration permanently. Press Y. **NOTE** If you have multiple adapters, you must erase the configuration of each adapter separately. For more information, see Section 9.5, Erasing the Current Configuration. After you erase the adapter configuration, reboot the system and then reconfigure the adapter.
You receive this POST error message: `Initiator iSCSI Name mismatch, Please use iSCSISelect to set a single name for all controllers. Press <Ctrl><S> to enter iSCSISelect. (Reboot required)`	In the iSCSI BIOS, the iSCSI initiator name can be different if there is more than one Emulex adapter in the system. If the iSCSI initiator name is different on multiple adapters, you receive a POST error message indicating an iSCSI name mismatch. You must enter the iSCSISelect utility and save a new initiator name on the first iSCSISelect utility menu so the iSCSI initiator name on all adapters match. All logins from the multiple adapters use the new name.

You receive any of the following POST error messages or warnings:
`Redboot Initialization failed...`
`Firmware Load failed...`
`DDR config failed...`
`DDR callibrate failed...`
`DDR test failed...`
`SEEPROM checksum failed...`
`Secondary firmware image loaded...`
	These POST messages indicate that you must reload the adapter firmware using the Emulex OneCommand CNA Manager application or one of the Emulex online or offline utilities. If the errors remain, contact Broadcom Technical Support.

You receive this POST error message: `Firmware halted. This firmware does not support this controller.`	The firmware loaded on this adapter is not supported. Load the appropriate firmware for the adapter.

14.3 Troubleshooting for the FCoE Protocol

Table 13 Troubleshooting the FCoE Protocol

Issue	Answer/Solution
Cisco Nexus switch configuration situations: ■ Windows creates the NTFS partition properly, but then reports that `The hard disk containing the partition or free space you chose has a LUN greater than 0. Setup cannot continue.` (Dell 1850 server). ■ Windows reboots successfully, but then gets stuck during the GUI portion of the installation right from the beginning. (HP DL385G2 server).	Ensure that the FCoE switch ports are configured correctly for the adapter's FCoE settings.

Appendix A: Configuring iSCSI through a DHCP Server Using Vendor-Specific Option 43

An iSCSI initiator requires the following parameters to boot from a target:

- Its IP address
- Its unique node name
- The boot iSCSI target IP address
- The boot target name. If not available, this may be discovered using the iSCSI discovery protocol.
- The target non-default TCP port number (if applicable)
- Header and data digest settings to be used with the target
- Authentication parameters, if applicable

If your boot BIOS allows it, you can configure all of the above parameters from a centrally configured DHCP server using vendor-specific option 43. This appendix documents this method of configuration. For this method, the initiator must be configured (using non-DHCP means) with the appropriate DHCP vendor ID. The method and format for specifying the vendor ID is outside the scope of this document. The initiator offers this Vendor ID to the DHCP server to retrieve data in the format described in the following sections.

Two other configuration methods are not documented in this appendix:

- Manual configuration using iSCSISelect or SMCLP/MILI
- Pre-boot configuration using CLP /BIOS 3.0

A.1 Format of Vendor-Specific Option 43

The format for vendor-specific option 43 is as follows:

```
iscsi:<ipaddress>:<protocol>:<iscsi port number>:<lun>:<target name>
```

Fields enclosed in angular brackets (including the angular brackets) must be replaced with their corresponding values. All fields are case insensitive.

See the example at the end of this document.

A.1.1 Description of Parameters

A.1.1.1 <ipaddress>

Replace this parameter with a valid IPv4 address in dotted decimal notation. This is a mandatory field.

A.1.1.2 <protocol>

Replace this parameter with a decimal number indicating the TCP port. The default TCP port is 3260.

A.1.1.3 <iscsi port number>

Replace this parameter with a decimal number ranging from 1 to 65535 (inclusive). It is an optional field.

A.1.1.4 <lun>

This parameter is a hexadecimal representation of the logical unit number of the boot device. It is an optional field. If not provided, LUN 0 is assumed to be the boot LUN. It is an 8-byte number which must be specified as a hexadecimal number consisting of 16 digits, with an appropriate number of 0's padded to the left, if required.

A.1.1.5 <target name>

Replace this parameter with a valid iSCSI target 'iqn' name of up to 223 characters. This is a mandatory field.

A.1.2 Example

```
iscsi:010.010.010.001:0:3260:0:iqn.1992-08.com.netap:sn.151729740
```

The above example specifies the following:

- Target IP address: 010.010.010.001
- Target protocol: 0
- Target TCP port: 3260
- Target boot LUN: 0
- Target iqn name: iqn.1992-08.com.netap:sn.151729740

Appendix B: Example for Installing and Configuring Linux or Citrix for PXE Boot and UEFIBoot

This section provides examples for installing and configuring Linux and Citrix for PXE boot and UEFIBoot.

B.1 Linux and Citrix PXE Server Remote Installation Procedure

PXE configuration requires a PXE server and the PXE client.

Setting up a PXE server requires the following configurations:

- NFS server
- TFTP server
- DHCP server
- PXE boot server

The pxelinux functionality occurs in this order:

1. The client machine boots to PXE, which requests a DHCP address.
2. The DHCP server responds with an IP address for the client machine, the address of a TFTP server, and a filename to load (`pxelinux.0`) from that server.
3. The client downloads pxelinux.0 from the specified TFTP server and executes it.
4. The `pxelinux.0` file searches the pxelinux.cfg directory on the server for a configuration file that matches the IP address of the machine. If no matches are found, it attempts to load a file called default.
5. The configuration file loaded by pxelinux.0 has instructions on what to do next. Some of the choices include boot to local hard drive, boot to an image file (floppy image), or load vmlinuz and initrd.img.
6. The client searches for a configuration file with the IP address converted to hexadecimal (for example, 192.168.1.60 becomes C0A8013C) or the MAC address of your PXE boot client's Ethernet card with a prefix of *01*. The MAC address must be separated with dashes instead of colons.

   In this example, the client looks for the following configuration file names and uses the first one it finds.

   ```
   01-00-00-C9-5B-75-A8
   C0A8013C
   C0A8013
   C0A801
   C0A80
   C0A8
   C0A
   C0
   C
   default
   ```

7. The default file's contents should look similar to the following:

   ```
   prompt 1
   default linux
   timeout 100

   label linux
   kernel vmlinuz
   append initrd=initrd.img ramdisk_size=9216 noapic acpi=off
   ```

B.1.1 PXE Server

This section provides configuration scripts for the PXE server.

B.1.1.1 NFS Server Configuration Script

```
[root@bglinux156 ~]# mkdir /work
[root@bglinux156 ~]# mount 10.192.194.110:/work /mnt
[root@bglinux156 ~]# cd /mnt/
[root@bglinux156 ~]# cp -r rhel5564/ /work/ #copy the entire
 directory to “/work”
[root@bglinux156 ~]# cd /work/rhel5564
[root@bglinux156 ~]# cp -r images/ /tftpboot/linux-install/
[root@bglinux156 ~]# vim /etc/exports###(Add “/work *(rw,sync)”)
[root@bglinux156 ~]# exportfs -a (reflect changes in NFS Server)
```

1. Restart NFS Services.

```
[root@bglinux156 ~]# /etc/rc.d/init.d/portmap restart
Stopping portmap: [ OK ]
Starting portmap: [ OK ]
[root@bglinux156 ~]# /etc/rc.d/init.d/nfslock restart
Stopping NFS locking: [ OK ]
Stopping NFS statd: [ OK ]
Starting NFS statd: [ OK ]
[root@bglinux156 ~]# /etc/rc.d/init.d/nfs restart
Shutting down NFS mountd: [ OK ]
Shutting down NFS daemon: [ OK ]
Shutting down NFS quotas: [ OK ]
Shutting down NFS services: [ OK ]
Starting NFS services: [ OK ]
Starting NFS quotas: [ OK ]
Starting NFS daemon: [ OK ]
Starting NFS mountd: [ OK ]
```

2. Check if the NFS services are running.

```
[root@bglinux156 ~]# rpcinfo -p
 program vers proto port
 100000 2 tcp 111 portmapper
 100000 2 udp 111 portmapper
 100021 1 udp 56782 nlockmgr
 100021 3 udp 56782 nlockmgr
 100021 4 udp 56782 nlockmgr
 100021 1 tcp 44855 nlockmgr
 100021 3 tcp 44855 nlockmgr
 100021 4 tcp 44855 nlockmgr
 100024 1 udp 766 status
 100024 1 tcp 769 status
 100011 1 udp 815 rquotad
 100011 2 udp 815 rquotad
 100011 1 tcp 818 rquotad
 100011 2 tcp 818 rquotad
 100003 2 udp 2049 nfs
 100003 3 udp 2049 nfs
 100003 4 udp 2049 nfs
 100003 2 tcp 2049 nfs
```

```
100003 3 tcp 2049 nfs
100003 4 tcp 2049 nfs
100005 1 udp 828 mountd
100005 1 tcp 831 mountd
100005 2 udp 828 mountd
100005 2 tcp 831 mountd
100005 3 udp 828 mountd
```

B.1.1.2 TFTP Server Setup

```
[root@bglinux156 ~]# vi /etc/xinetd.d/tftp

service tftp
{
 socket_type = dgram
 protocol = udp
 wait = yes
 user = root
 server = /usr/sbin/in.tftpd
 server_args = -s /tftpboot/linux-install #This line specifies
 #path where the pxe boot searches for
 #the images
 disable = no #convert this line from yes to no
 per_source = 11
 cps = 100 2
 flags = IPv4
}

[root@bglinux156 ~]# /sbin/chkconfig --level 345 xinetd on
[root@bglinux156 ~]# /sbin/chkconfig --level 345 tftp on
[root@bglinux156 ~]# service xinetd restart
```

B.1.1.3 DHCP Server Configuration

1. Install the following RPMs in the server if not installed:

```
[root@bglinux156 ~]# rpm -qa | grep dhcp
dhcpv6-client-1.0.10-18.el5
dhcp-devel-3.0.5-23.el5
dhcp-3.0.5-23.el5
```

2. Make the following entry in `/etc/dhcpd.conf` file:

```
#
# DHCP Server Configuration file.
# see /usr/share/doc/dhcp*/dhcpd.conf.sample
#
ddns-update-style interim;
 subnet 192.168.1.0 netmask 255.255.255.0 {
 range 192.168.1.50 192.168.1.90;
 default-lease-time 3600;
 max-lease-time 4800;
 option subnet-mask 255.255.255.0;
 option domain-name “pxe_text”;
 option time-offset -8;
 }
```

```
host bglinux45{
 next-server 192.168.1.1;
 hardware ethernet 00:00:C9:5B:75:A8;
 fixed-address 192.168.1.60;
 option host-name “linux-test”;

  filename “pxelinux.0”;
 }
###This sets up a DNS server that will assign IP Address 192.168.1.60
###to the client machine that has MAC Address “00:00:C9:5B:75:A8”
###assigned to a PXE capable NIC
###The only thing that needs to be changed in the above, is the MAC
###Address to match that of the NIC in the client box
###The IP address that follows the next-server option should be the IP
###address of the tftp server

[root@bglinux156 ~]# service dhcpd restart
Shutting down dhcpd: [ OK ]
Starting dhcpd: [ OK ]
```

B.1.1.4 PXE Boot Server Configuration

```
[root@bglinux156 ~]# rpm -qa | grep system-config-netboot
system-config-netboot-cmd-0.1.45.1-1.el5
system-config-netboot-0.1.45.1-1.el5
```

B.1.2 Copying Files to the TFTP Server with the Graphical Version of the Network Booting Tool

Copy the necessary files to start the installation to the TFTP server so they can be found when the client requests them. Run the Network Booting Tool on the NFS, FTP, or HTTP server. A separate PXE server is not necessary.

To use the graphical version of the Network Booting Tool, you must be running the X Window System, have root privileges, and have the redhat-config-netboot.rpm package installed.

To start the Network Booting Tool from the desktop, do one of the following:

- Go to the Main Menu Button (on the Panel) and select **System Settings >Server Settings >Network Booting Service**.

  Or

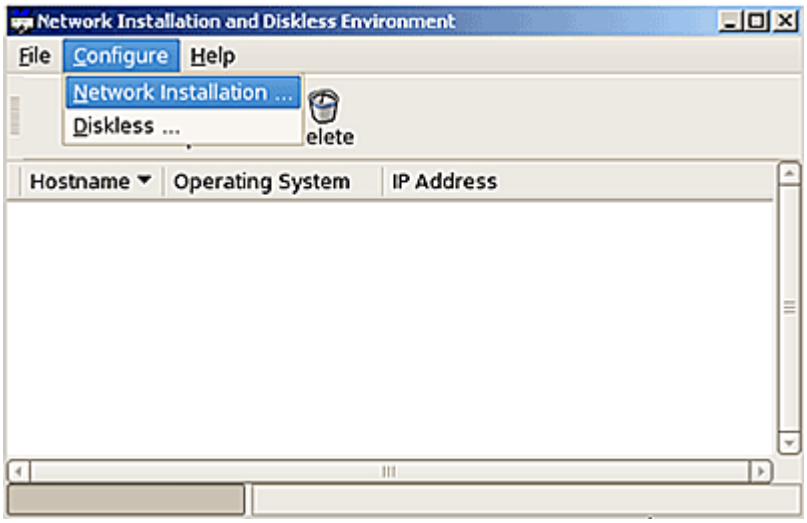
- Type the command `system-config-netboot` at a shell prompt. (For example, in an XTerm or a GNOME terminal.)

NOTE If you are starting the Network Booting Tool for the first time, select **Network Install from the First Time Druid**.

Copy the files necessary to start the installation to the TFTP server.

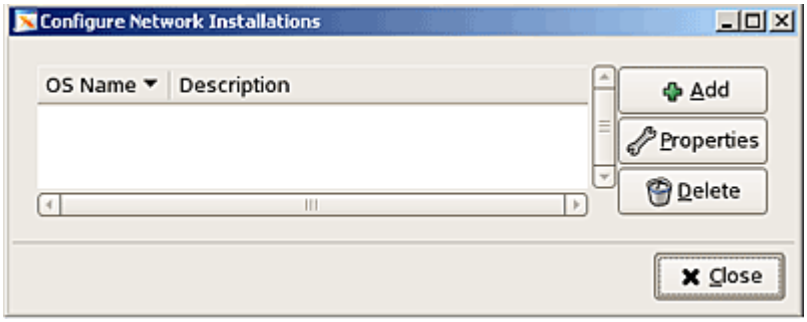
1. Select **Configure > Network Installation** from the pull-down menu (Figure 152).

Figure 152 Network Installation and Diskless Environment Window



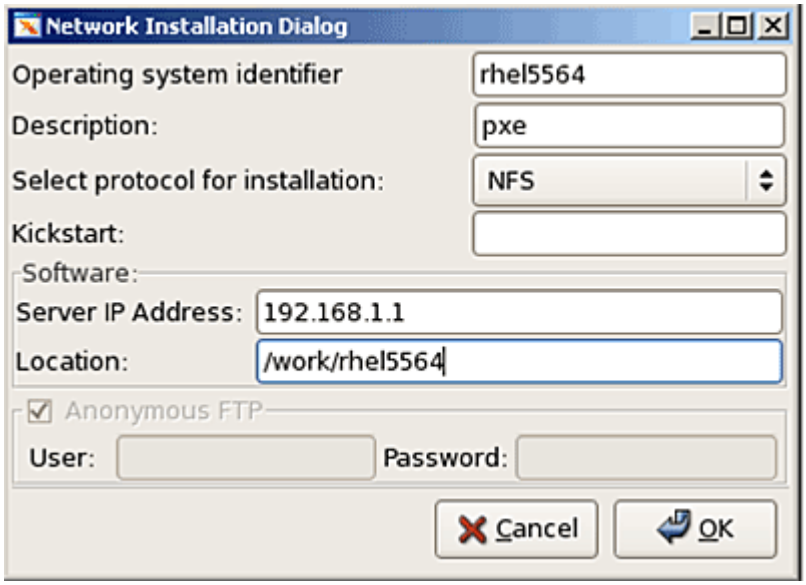
The following Configure Network Installations window appears.

Figure 153 Configure Network Installations Window



2. Click Add.
   The following Network Installation Dialog window (Figure 154) appears.

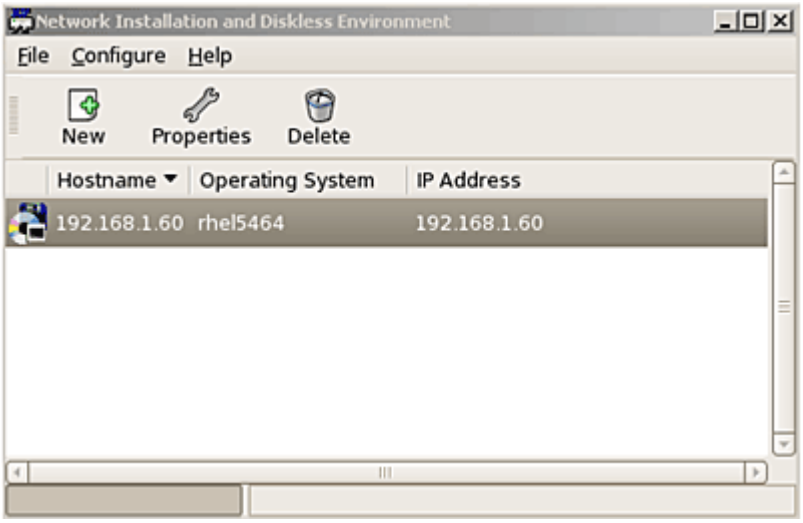
Figure 154 Network Installation Dialog Window



3. Enter the operating system identifier and description.
4. Select pxe as the protocol for installation.
5. Enter the IP address of the TFTP server (as set up in the DHCP server configuration) in the server IP address field.
6. Enter a directory location for the server.

   The following Network Installation and Diskless Environment window shows the IP address of the TFTP server.

Figure 155 Network Installation and Diskless Environment Window with IP Address

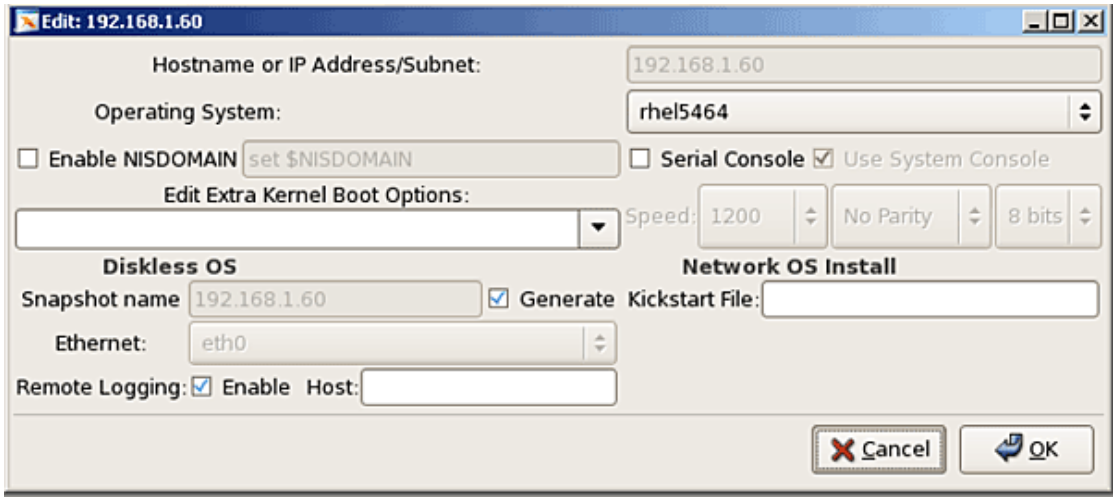


B.1.2.1 Configuring the Host

To configure the host:

1. Double-click the IP address row in the Network Installation and Diskless Environment window (Figure 155). The following Edit window (Figure 156) appears.

Figure 156 Edit Window



2. In the Hostname or IP Address/Subnet field, enter the IP address, fully qualified hostname, or a subnet of systems that should be allowed to connect to the PXE server for installations. In Figure 156, the Hostname or IP Address/Subnet field is the IP address of the client machine.
3. In the case of a multiboot environment, select an operating system preference.
4. Select the operating system identifier to install on this client. The list is populated from the network install instances created from the Network Installation dialog box.
5. If you use a serial console, check the Serial Console check box.
6. Specify the location of a kickstart file, such as `http://server.example.com/kickstart/ks.cfg`. This file can be created with the Kickstart Configuration.

NOTE Ignore the Snapshot name and Ethernet options. They are used for diskless environments only.

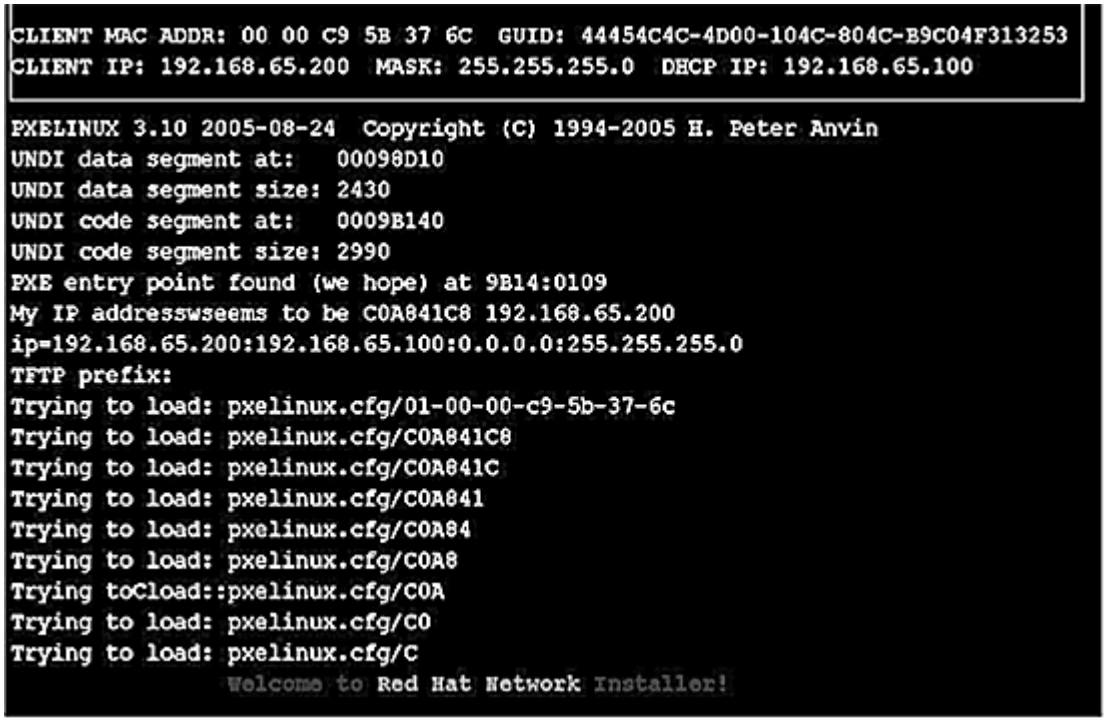
7. Click OK. This closes the window and creates configuration files within:
   `/tftpboot/linux-install and /tftpboot/linux-install/pxelinux.cfg`
8. Go to `/tftpboot/pxelinux.cfg` and open the file `C0A8013C`. The `C0A8013C` file is created when you assign the client IP using `system-config- netboot`. This is the hexadecimal address of the IP assigned to the client NIC. It is 192.168.1.60 in the example in Figure 155.

B.1.2.2 Enabling the PXE Client with a PXE-Enabled Adapter

The PXE client must have a PXE-enabled adapter.

1. Power on the PXE client.
2. In the Enabling Network Boot window, press the F12 key to enable network boot (Figure 157). Pressing the F12 key accesses the DHCP server IP and the Boot Linux image.

Figure 157 Enabling Network Boot Window



Appendix C: Examples for Configuring and Booting UEFI NIC

This section provides examples for configuring and booting UEFI NIC.

C.1 UEFI NIC Server Configuration Script for SLES11 SPx

```
[root@orleansrhel5564 ~]# ls /root/sles11sp164/
SLES-11-SP1-DVD-x86_64-GM-DVD1.iso
[root@orleansrhel5564 ~]# mkdir test1
[root@orleansrhel5564 ~]# mkdir test2
[root@orleansrhel5564 ~]# mount -o loop
sles11sp164/SLES-11-SP1-DVD-x86_64-GM-DVD1.iso test1
[root@orleansrhel5564 ~]# ls test1
ARCHIVES.gz COPYING.degpg-pubkey-3d25d3d9-36e12d04.asc media.1
boot  COPYRIGHT gpg-pubkey-7e2e3b05-4be037ca.asc NEWS
ChangeLog COPYRIGHT.degpg-pubkey-9c800aca-4be01999.asc pubring.gpg
content directory.yast gpg-pubkey-a1912208-446a0899.asc README
content.ascdocu gpg-pubkey-b37b98a9-4be01a1a.asc suse
content.keygpg-pubkey-0dfb3188-41ed929b.ascINDEX.gz
control.xmgpg-pubkey-1d061a62-4bd70bfa.asclicense.tar.gz
COPYING gpg-pubkey-307e3d54-4be01a65.ascls-lR.gz
[root@orleansrhel5564 ~]# mount -o loop test1/boot/x86_64/efi test2
[root@orleansrhel5564 ~]# ls -lar test2
total 28
-rwxr-xr-x 1 root root 48 May 20 2010 .packages.boot-efi
drwxr-xr-x 3 root root 2048 May 20 2010 efi
drwxr-x--- 20 root root 4096 Jan 5 14:04 ..
drwxr-xr-x 3 root root 16384 Jan 1 1970 .
[root@orleansrhel5564 ~]# ls -lar test2/efi/boot/
bootx64.efi elilo.conf initrd linux
[root@orleansrhel5564 ~]# ls -lar test2/efi/boot/
total 26838
-rwxr-xr-x 1 root root 3231872 May 20 2010 linux
-rwxr-xr-x 1 root root 23999623 May 20 2010 initrd
-rwxr-xr-x 1 root root 512 May 20 2010 elilo.conf
-rwxr-xr-x 1 root root 241318 May 20 2010 bootx64.efi
drwxr-xr-x 3 root root 2048 May 20 2010 ..
drwxr-xr-x 2 root root 2048 May 20 2010 .
[root@orleansrhel5564 ~]# cp test2/boot/efi/* /tftpboot/linux-install/
```

C.1.1 NFS Server Configuration Script

```
[root@bglinux156 ~]# /etc/rc.d/init.d/portmap restart
Stopping portmap: [ OK ]
Starting portmap: [ OK ]
[root@bglinux156 ~]# /etc/rc.d/init.d/nfslock restart
Stopping NFS locking: [ OK ]
Stopping NFS statd: [ OK ]
Starting NFS statd: [ OK ]
[root@bglinux156 ~]# /etc/rc.d/init.d/nfs restart
```

```
Shutting down NFS mountd: [ OK ]
Shutting down NFS daemon: [ OK ]
Shutting down NFS quotas: [ OK ]
Shutting down NFS services: [ OK ]
Starting NFS services:[ OK ]
Starting NFS quotas: [ OK ]
Starting NFS daemon: [ OK ]
Starting NFS mountd: [ OK ]
#### Check if the NFS services are running:
[root@bglinux156 ~]# rpcinfo -p
 program vers proto port
 100000 2 tcp 111 portmapper
 100000 2 udp 111 portmapper
 100021 1 udp 56782 nlockmgr
 100021 3 udp 56782 nlockmgr
 100021 4 udp 56782 nlockmgr
 100021 1 tcp 44855 nlockmgr
 100021 3 tcp 44855 nlockmgr
 100021 4 tcp 44855 nlockmgr
 100024 1 udp 766 status
 100024 1 tcp 769 status
 100011 1 udp 815 rquotad
 100011 2 udp 815 rquotad
 100011 1 tcp 818 rquotad
 100011 2 tcp 818 rquotad
 100003 2 udp 2049 nfs
 100003 3 udp 2049 nfs
 100003 4 udp 2049 nfs
 100003 2 tcp 2049 nfs
 100003 3 tcp 2049 nfs
 100003 4 tcp 2049 nfs
 100005 1 udp 828 mountd
 100005 1 tcp 831 mountd
 100005 2 udp 828 mountd
 100005 2 tcp 831 mountd
 100005 3 udp 828 mountd
```

C.1.2 DHCP Server Configuration

```
[root@orleansrhel5564 ~]# cat /etc/dhcpd.conf
```

1. Install the following RPMs in the server if not installed:

```
[root@bglinux156 ~]# rpm -qa | grep dhcp
dhcpv6-client-1.0.10-18.el5
dhcp-devel-3.0.5-23.el5
dhcp-3.0.5-23.el5
```

2. Make the following entry in `/etc/dhcpd.conf` file:

```
#
# DHCP Server Configuration file.
#      see /usr/share/doc/dhcp*/dhcpd.conf.sample
#

ddns-update-style interim;
```

```
#ignore client-updates;
subnet 192.168.47.0 netmask 255.255.255.0 {
range 192.168.47.50 192.168.47.100;
default-lease-time 3600;
max-lease-time-4800;
option subnet-mask 255.255.255.0;
option domain-name "pxe_test";
option time-offset -8; # Eastern Standard Time
}

# we want the nameserver to appear at a fixed address
      host orleansrhel5564 {
      next-server 192.168.47.1;#IP of the NFS Server
      hardware ethernet 00:00:c9:5b:a5:26;#mac address of
            #the client
      #machine

      fixed-address 192.168.47.60; #IP assigned to the
      #client machine
      option host-name "linux-test";
      filename "bootx64.efi"; #Used for UEFI boot
      }

###This sets up a DNS server that will assign IP Address 192.168.47.60
###to the client machine that has MAC Address "00:00:c9:5b:a5:26"
###assigned to a PXE(UEFI) capable NIC
###The only thing that needs to be changed in the above, is the MAC
###Address to match that of the NIC in the client box
###The IP address that follows the next-server option should be the IP
###address of the tftp server
[root@bglinux156 ~]# service dhcpd restart
Shutting down dhcpd: [ OK ]
Starting dhcpd: [ OK ]
```

C.1.3 TFTP Server Setup

```
[root@orleansrhel5564 ~]# cat /etc/xinetd.d/tftp
service tftp
{
 disable = no
 socket_type = dgram
 protocol    = udp
 wait        = yes
 user        = root
 server      = /usr/sbin/in.tftpd
 server_args = -s /tftpboot/linux-install
 per_source = 11
 cps         = 100 2
 flags       = IPv4
```

C.2 UEFI NIC Server Configuration Script for RHEL 6.x

This section provides UEFI NIC server configuration scripts to use on RHEL 6.x operating systems.

C.2.1 Setup Information

PXE client server: Dell R710 with Emulex adapter

PXE server: HP running RHEL 5.5 64 bit

C.2.2 NFS Configuration

```
[root@orleans ~]# mkdir /work
[root@orleans ~]# mount 10.192.194.110:/work /mnt   #### Mounting network drive
[root@orleans ~]# cd /mnt/
[root@orleans ~]# cp -r rhel6364/ /work/  #### Copying rhel6.1-64 bit os to PXE
server mechine
[root@orleans ~]# vim /etc/exports       ####(Add "/work    *(rw,sync)"
[root@orleans ~]#  exportfs -a           ####(reflect changes in NFS Server)


#### Restart NFS services
[root@orleans ~]#/etc/rc.d/init.d/portmap restart
[root@orleans ~]# /etc/rc.d/init.d/nfslock restart
[root@orleans ~]# /etc/rc.d/init.d/nfs restart


#### Check if NFS is running fine

[root@orleans ~]# rpcinfo -p

   program vers proto   port
    100000    2   tcp    111  portmapper
    100000    2   udp    111  portmapper
    100021    1   udp  56782  nlockmgr
    100021    3   udp  56782  nlockmgr
    100021    4   udp  56782  nlockmgr
    100021    1   tcp  44855  nlockmgr
    100021    3   tcp  44855  nlockmgr
    100021    4   tcp  44855  nlockmgr
    100024    1   udp    766  status
    100024    1   tcp    769  status
    100011    1   udp    815  rquotad
    100011    2   udp    815  rquotad
    100011    1   tcp    818  rquotad
    100011    2   tcp    818  rquotad
    100003    2   udp   2049  nfs
    100003    3   udp   2049  nfs
    100003    4   udp   2049  nfs
    100003    2   tcp   2049  nfs
    100003    3   tcp   2049  nfs
    100003    4   tcp   2049  nfs
```

```
    100005    1   udp    828  mountd
    100005    1   tcp    831  mountd
    100005    2   udp    828  mountd
    100005    2   tcp    831  mountd
    100005    3   udp    828  mountd
```

C.2.3 TFTP Configuration

```
[root@orleans ~]# vi /etc/xinted.d/tftp

# default: off
# description: The tftp server serves files using the trivial file transfer \
#       protocol.  The tftp protocol is often used to boot diskless \
#       workstations, download configuration files to network-aware printers, \
#       and to start the installation process for some operating systems.
service tftp
{
        disable                 = no
        socket_type             = dgram
        protocol                = udp
        wait                    = yes
        user                    = root
        server                  = /usr/sbin/in.tftpd
        server_args             = -s /tftpboot/linux-install
        per_source              = 11
        cps                     = 100 2
        flags                   = IPv4
}


##### Restart TFTP service

[root@orleans ~]# service xinetd restart
Stopping xinetd:                                           [  OK  ]
Starting xinetd:                                           [  OK  ]
[root@orleans ~]# chkconfig tftp on
```

C.2.4 PXE Configuration

```
[root@orleans ~]# cd /work/rhel6364/

[root@orleans RHEL6364]# cd images/pxeboot/

[root@orleans pxeboot]# cp -r * /tftpboot/linux-install/
#### Copy vmlinuz and initrd to TFTP root directory

[root@orleans pxeboot]# mkdir -p /root/test1

[root@orleans images]# cd /work/rhel6364/images/

[root@orleans images]# mount -o loop efiboot.img /root/test1      ###### Extract
efiboot.img to get bootx64.efi and BOOTX64.conf files
```

```
[root@orleans images]# cd /root/test1

[root@orleans test1]# cd efi/boot/

[root@orleans boot]# cp -r * /tftpboot/linux-install/    ###### Copy bootx64.efi
and BOOTX64.conf to TFTP root directory

#### Editing efidefault to the following content

[root@orleans boot]# cat /tftpboot/linux-install/efideault
default=0
timeout 10
splashimage=(nd)/splash.xpm.gz
title RHEL6364
        root (nd)
        splashimage /splash.xpm.gz
        kernel /vmlinuz keymap=us lang=en_US  method=nfs:192.168.1.1:/RHEL6364
ip=dhcp noipv6
        initrd /initrd.im
```

C.2.5 Configuring DHCP

```
[root@orleans ~]# cat  /etc/dhcpd.conf
#
# DHCP Server Configuration file.
#   see /usr/share/doc/dhcp*/dhcpd.conf.sample
ddns-update-style interim;
ignore client-updates;
#allow booting;
#allow bootp;
subnet 192.168.1.0 netmask 255.255.255.0 {
 range 192.168.1.50 192.168.1.90;
 default-lease-time 3600;
 max-lease-time 4800;
 option subnet-mask 255.255.255.0;
# option router 192.168.1.1;
 option domain-name "pxe_text";
# option name-server 192.168.1.20;
 option time-offset -8;
 }

host bglinux45{
  next-server 192.168.1.1;    #### IP address of PXE server interface connected
back to back to PXE client interface
  hardware ethernet 00:00:C9:BB:C7:8F;  #### MAC address of PXE client interface
  fixed-address 192.168.1.60;  #### IP address to be assigned to PXE client
interface
  option host-name "linux-test";
  filename "bootx64.efi";
 }

#### Restart DHCP service
```

```
[root@orleans ~]# service dhcpd restart
Shutting down dhcpd:                                       [  OK  ]
Starting dhcpd:                                            [  OK  ]
```

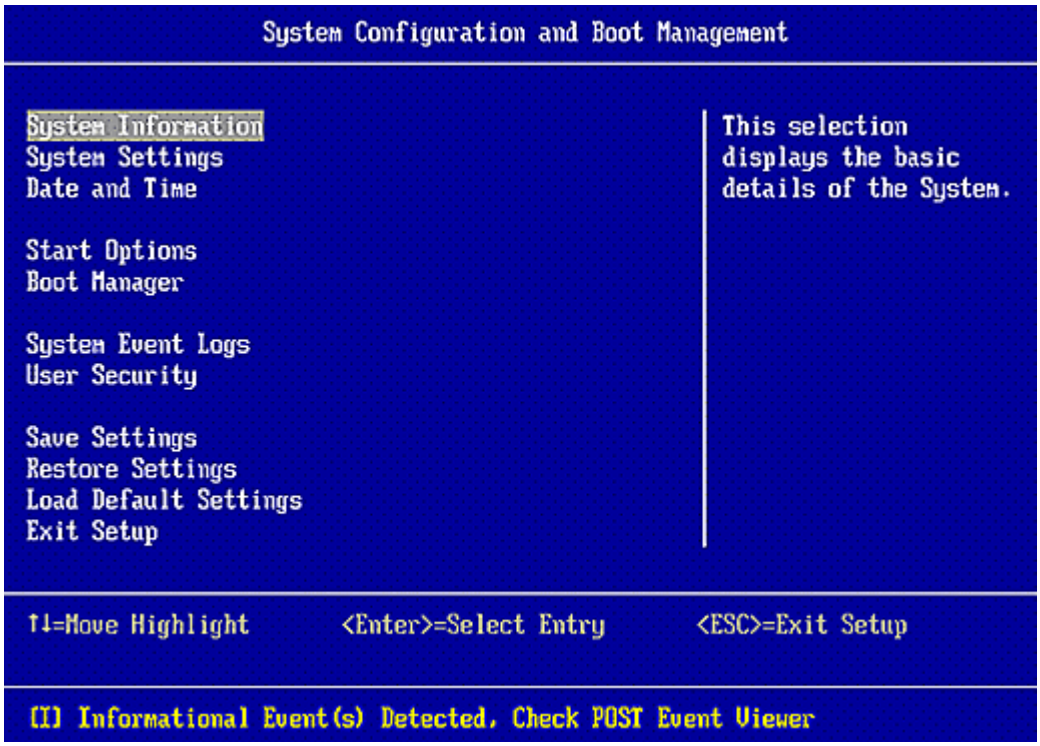
NOTE Reboot the system under test and boot from the PXE client interface.

C.3 Configuring the Lenovo System x HS22 Blade with e-Raptor MEZZ

To install and update the driver, follow these steps:

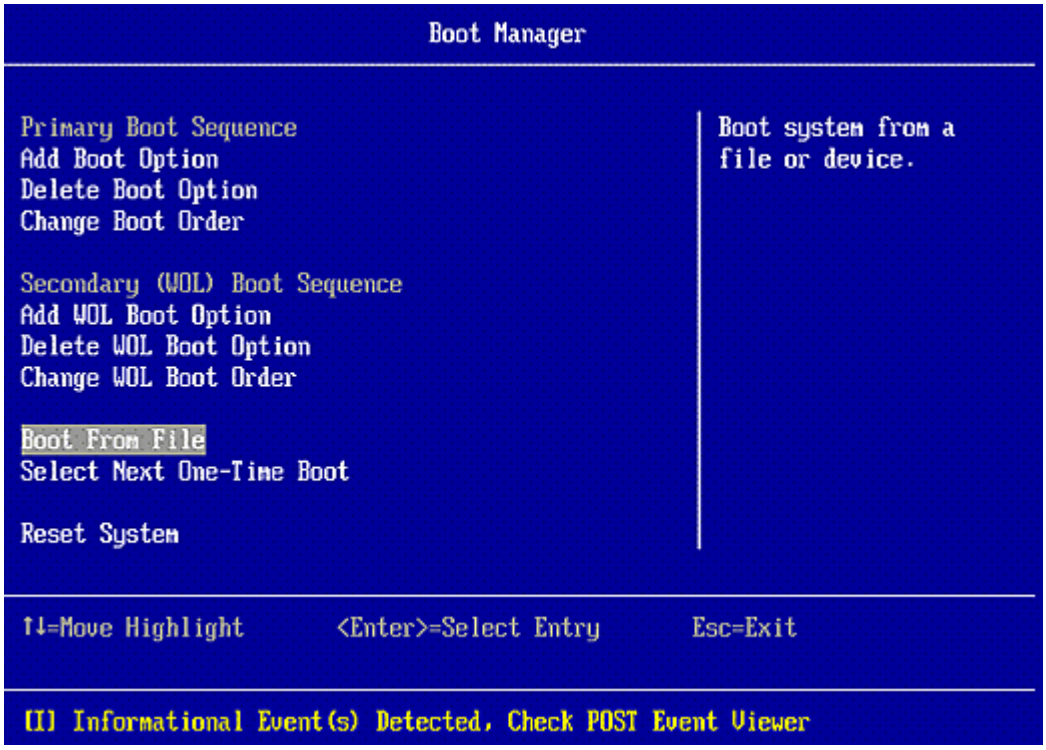
1. Press the F1 key during the system boot and log in to the following System Configuration and Boot Management dialog.

Figure 158 System Configuration and Boot Management Dialog



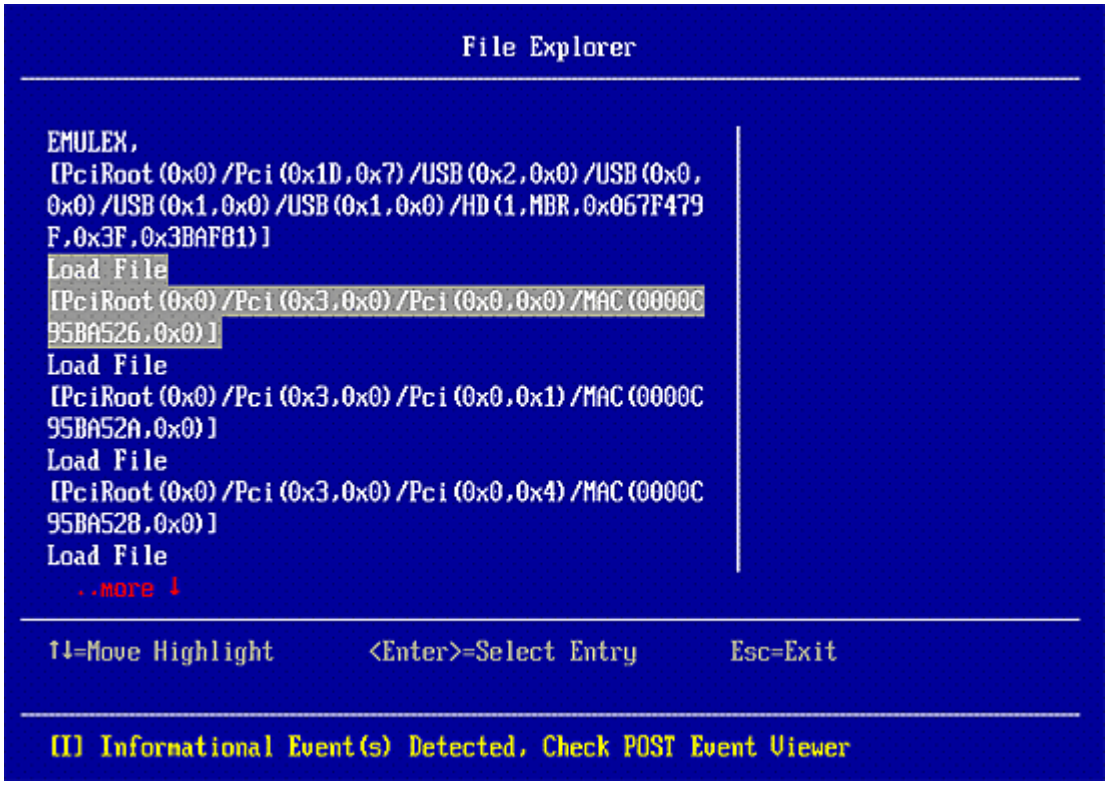
2. Select Boot Manager and press Enter. The following Boot Manager dialog (Figure 159) appears.

Figure 159 Boot Manager Dialog



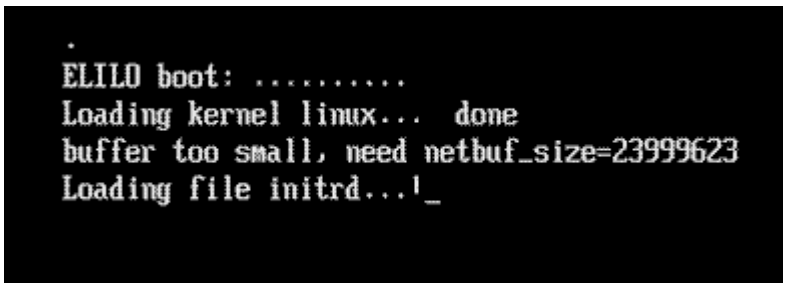
3. Select Boot from File and press Enter. The following File Explorer dialog (Figure 160) appears.

Figure 160 File Explorer Dialog



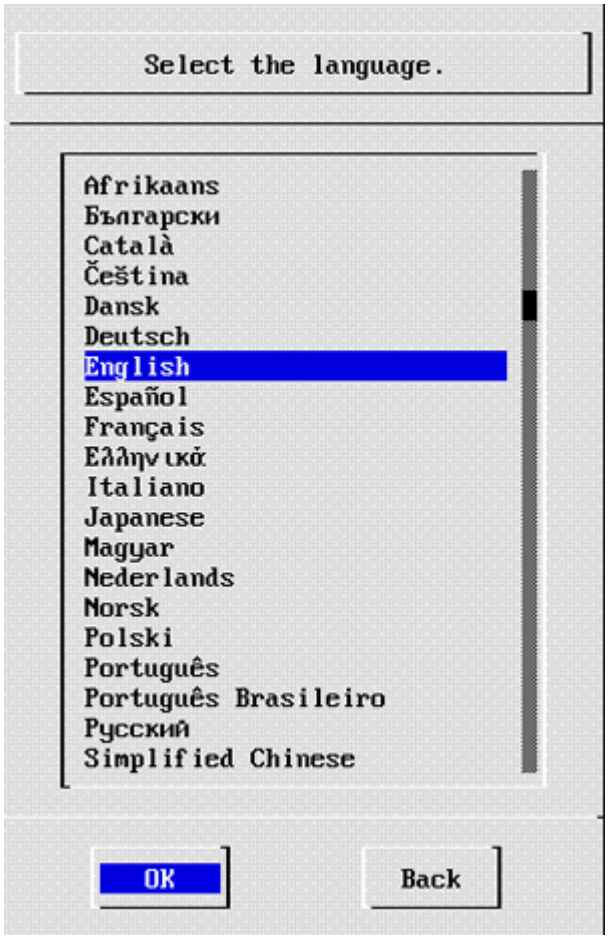
4. Select the card from which the UEFI NIC boot is to be performed. Press Enter.
   A message similar to the following pop up appears.

Figure 161 ELILO Boot Message



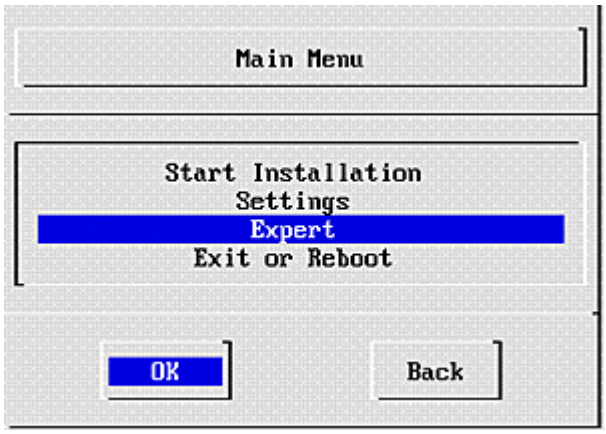
The following Language Selection menu (Figure 162) appears.

Figure 162 Language Selection Menu



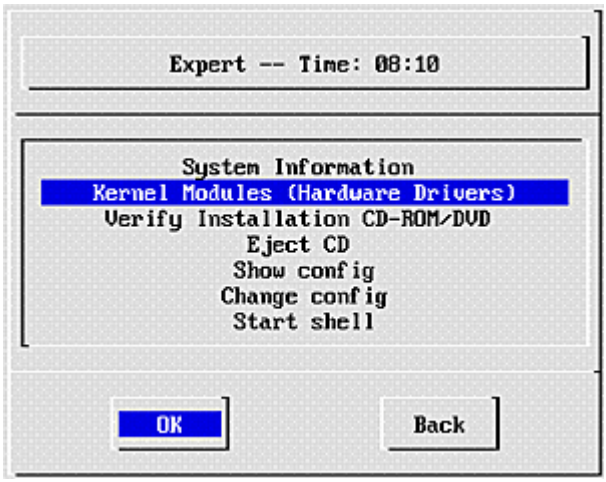
5. Select a language and click OK. The following Main Menu appears.

Figure 163 Main Menu



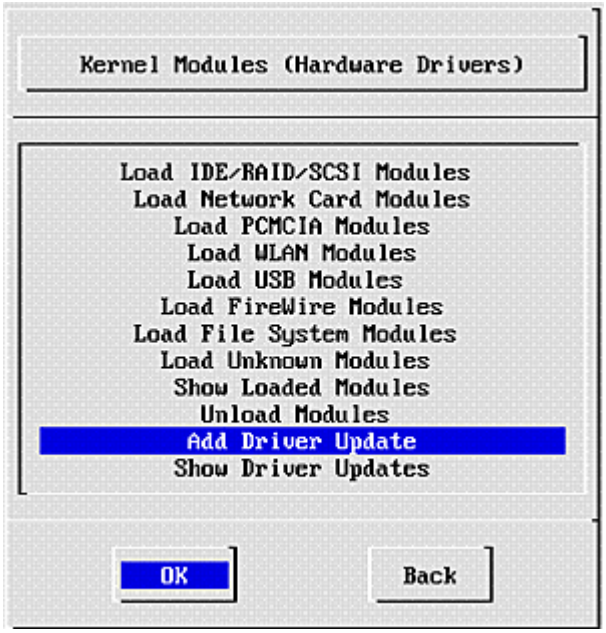
6. Select Expert and click OK. The following Expert menu (Figure 164) appears.

Figure 164 Expert Menu



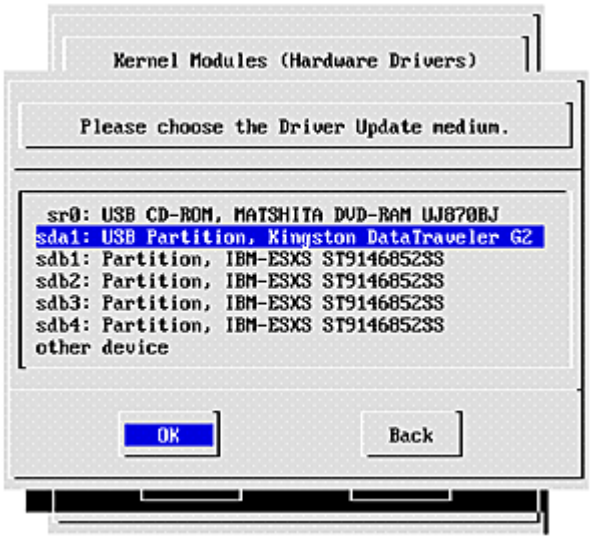
7. Select Kernel Modules (Hardware Drivers) and click OK. The following Kernel Modules (Hardware Drivers) menu appears.

Figure 165 Kernel Modules (Hardware Drivers) Menu



8. Select Add Driver Update and click OK. The following Driver Update Medium Selection listing (Figure 166) appears.

Figure 166 Driver Update Medium Selection Listing



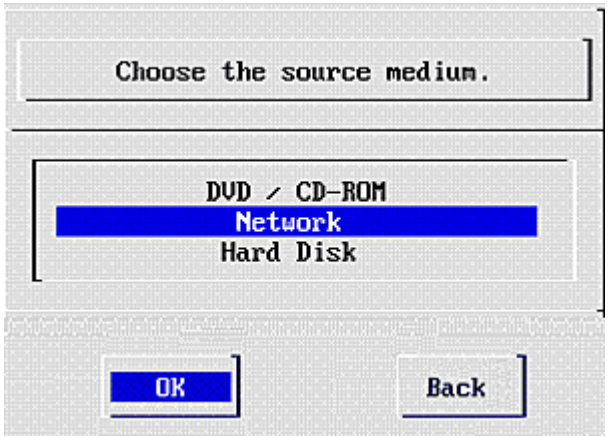
9. Select a medium and click OK. The following Driver Update List confirmation dialog box appears.

Figure 167 Driver Update List Dialog Box



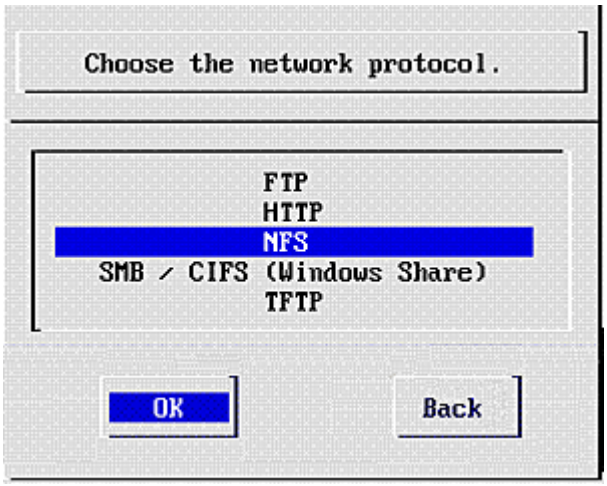
10. Click OK. The Expert menu (Figure 164) appears.
11. Click Back. The Main Menu (Figure 163) appears.
12. Select Start Installation and click OK. The following Source Medium Selection menu (Figure 168) appears.

Figure 168 Source Medium Selection Menu



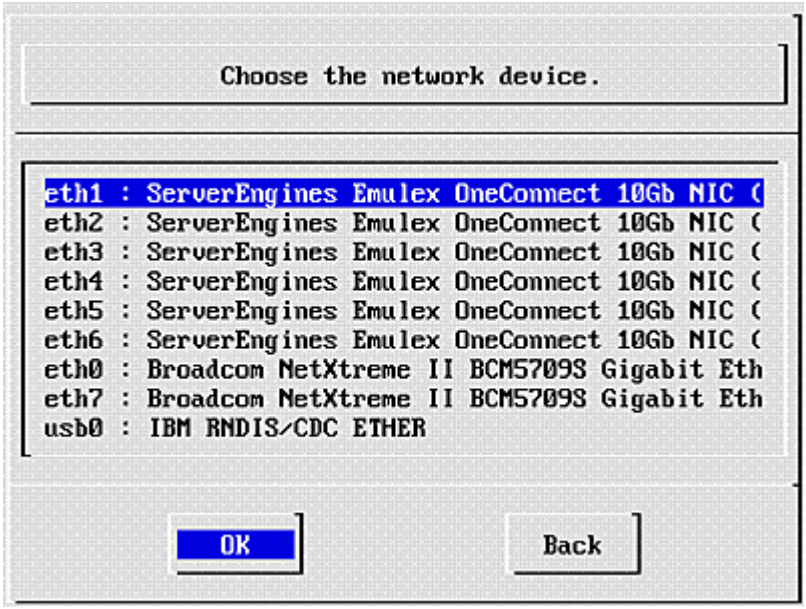
13. Select Network and click OK. The following Network Protocol Selection menu appears.

Figure 169 Network Protocol Selection Menu



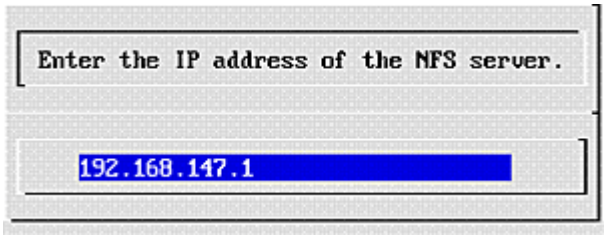
14. Select NFS and click OK. The following Network Device Selection listing (Figure 170) appears.

Figure 170 Network Device Selection Listing



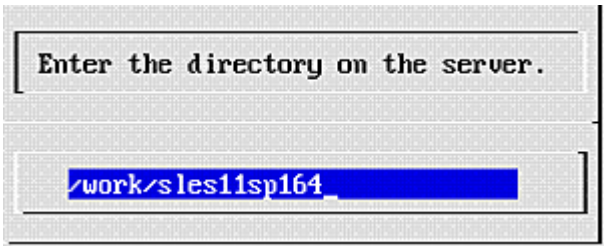
15. Select the device and click OK. The DHCP request is sent. A following dialog box prompts you to enter the IP address of the NFS server.

Figure 171 NFS Server IP Address Dialog Box



16. Enter the IP address and press Enter. A following dialog box prompts you to enter the directory path of the server.

Figure 172 Server Directory Dialog Box



17. Enter the directory path and press Enter. The system displays, information similar to that shown in Figure 173 and Figure 174.

Figure 173 Starting Script

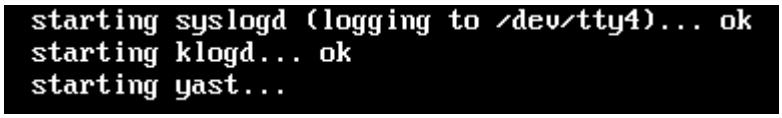
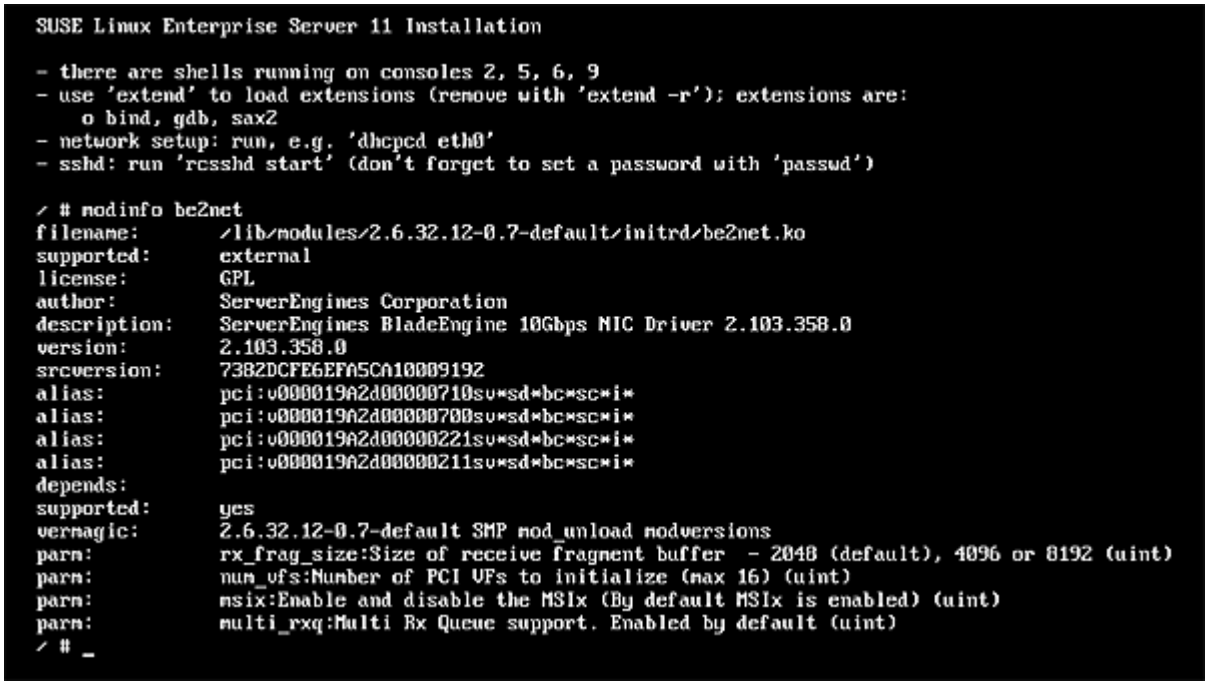


Figure 174 Driver Update Confirmation



C.4 UEFI NIC IPv6 PXE Setup Procedure for SLES11 SP2

This section provides UEFI NIC PXE setup instructions for configuring IPv6 on the SLES11 SP2 operating system.

C.4.1 Setup Information

- PXE server: CISCO UCS C240 with SLES11 SP2 operating system
- PXE client server: DELL R620 Server F6PCP adapter installed

C.4.2 PXE Server Setup

Launch `yast2`

C.4.2.1 Configure Network Adapter

1. In the Network Devices section, select Network Settings.
2. Select the adapter port you wish the PXE server to use, then click Edit.
3. Select Statically assigned IP Address and enter the following settings:
   - IP Address: 192.168.1.1
   - Subnet Mask: /24
4. Under the Additional Addresses section, click Add to add an alias.
   - Alias Name: ipv6static
   - IP Address: fc00:ba49:1625:fb0f::137
   - Netmask: /64
5. Click OK.

C.4.2.2 Disable Firewall

1. In the Security and Users section, select Firewall.
2. Select Disable Firewall Automatic Starting.
3. Click Stop Firewall Now.

C.4.2.3 Install Services

1. In the Network Services section, select DHCP Server.
2. Choose to install and accept all defaults.
3. Install NFS server, TFTP server, and FTP server one by one by accepting all defaults.

C.4.3 Configure DHCPD6

```
Brussels-sles11sp2:~ # vim /etc/dhcpd6.conf

default-lease-time 2592000;
preferred-lifetime 604800;
option dhcp-renewal-time 3600;
option dhcp-rebinding-time 7200;
allow leasequery;
option dhcp6.name-servers fc00:ba49:1625:fb0f::1;
option dhcp6.domain-search "pxetest.org";
option dhcp6.info-refresh-time 21600;
option dhcp6.bootfile-url code 59 = string;
# The subnet where the server is attached
# (i.e., the server has an address in this subnet)
subnet6 fc00:ba49:1625:fb0f::/64 {
option dhcp6.bootfile-url "tftp://[fc00:ba49:1625:fb0f::137]/bootx64.efi";
# Nine addresses available to clients
# (the tenth client should get NoAddrsAvail)
range6 fc00:ba49:1625:fb0f::42:10 fc00:ba49:1625:fb0f::42:99;
# Use the whole /64 prefix for temporary addresses
# (i.e., direct application of RFC 4941)
range6 fc00:ba49:1625:fb0f:: temporary;
}
```

Edit `/etc/sysconfig/dhcpd` and verify `DHCPD6_INTERFACE` is set to the network adapter port on which you want the DHCP server to listen.

```
DHCPD6_INTERFACE="ethx"
```

C.4.4 Configure TFTP

```
Brussels-sles11sp2:~ # vim  /etc/xinetd.d/tftp
# default: off
# description: tftp service is provided primarily for booting or when a \
#       router need an upgrade. Most sites run this only on machines acting as
#       "boot servers".
service tftp
{
        socket_type     = dgram
        protocol        = udp
        wait            = yes
```

```
        flags           = IPv6 IPv4
        user            = root
        server          = /usr/sbin/in.tftpd
        server_args     =  -s /tftpboot
        disable         = no
        per_source      = 11
        cps             = 100 2
}
```

C.4.4.1 Configure ELILO Boot Loader

Retrieve ELILO from the installation DVD and create a configuration file:

1. Place the first SUSE Linux Enterprise Server install DVD in the DVD drive of the PXE Server.
2. Open a root terminal.
3. Mount the boot image and copy the contents.

```
Brussels-sles11sp2:~ # mkdir /tmp/efi-img
Brussels-sles11sp2:~ #  mount  -o loop  /media/<name of
disc>/boot/x86_64/efi     /tmp/efi-img
Brussels-sles11sp2:~ #  cp /tmp/efi-img/efi/boot/*   /tftpboot
```

C.4.4.2 Create /tftpboot/elilo.conf with the Following Contents

```
Brussels-sles11sp2:~ # vim  /tftpboot/elilo.conf
prompt
timeout=50
delay=100
image=/linux
label=suse
initrd=/initrd
append="install=ftp://[fc00:ba49:1625:fb0f::137]/sles  ipv6only=1  sysrq=1  dd
vga=791 manual=1 textmode=0 showopts"
```

C.4.4.3 Copy Installation Files for FTP Transfer

```
Brussels-sles11sp2:~ #  cd   /media/<name of disc>
Brussels-sles11sp2:~ #  mkdir   /srv/ftp/sles
Brussels-sles11sp2:~ #  cp -r *  /srv/ftp/sles
```

C.4.4.4 Configure FTP to Enable Support for IPv6

```
Brussels-sles11sp2:~ # vim     /etc/vsftpd.conf
# Activate directory messages - messages given to remote users when they
# go into a certain directory.
#
dirmessage_enable=YES
#
# It is recommended that you define on your system a unique user which the
# ftp server can use as a totally isolated and unprivileged user.
#
nopriv_user=ftpsecure
#
# You may fully customise the login banner string:
#
#ftpd_banner="Welcome to FOOBAR FTP service."
#
```

```
# You may activate the "-R" option to the builtin ls. This is disabled by
# default to avoid remote users being able to cause excessive I/O on large
# sites. However, some broken FTP clients such as "ncftp" and "mirror" assume
# the presence of the "-R" option, so there is a strong case for enabling it.
#
#ls_recurse_enable=YES
#
# You may specify a file of disallowed anonymous e-mail addresses. Apparently
# useful for combatting certain DoS attacks.
#
#deny_email_enable=YES
#
# (default follows)
#
#banned_email_file=/etc/vsftpd.banned_emails
#
# If  enabled,  all  user  and  group  information in
# directory listings will be displayed as "ftp".
#
#hide_ids=YES

# Local FTP user Settings
#
# Uncomment this to allow local users to log in.
#
#local_enable=YES
#
# Default umask for local users is 077. You may wish to change this to 022,
# if your users expect that (022 is used by most other ftpd's)
#
#local_umask=022
#
# Uncomment to put local users in a chroot() jail in their home directory
# after login.
#
#chroot_local_user=YES
#
# You may specify an explicit list of local users to chroot() to their home
# directory. If chroot_local_user is YES, then this list becomes a list of
# users to NOT chroot().
#
#chroot_list_enable=YES
#
# (default follows)
#
#chroot_list_file=/etc/vsftpd.chroot_list
#
# The maximum data transfer rate permitted, in bytes per second, for
# local authenticated users. The default is 0 (unlimited).
#
#local_max_rate=7200
# Anonymus FTP user Settings
#
# Allow anonymous FTP?
```

```
#
anonymous_enable=YES
#
# Anonymous users will only be allowed to download files which are
# world readable.
#
anon_world_readable_only=YES
#
# Uncomment this to allow the anonymous FTP user to upload files. This only
# has an effect if the above global write enable is activated. Also, you will
# obviously need to create a directory writable by the FTP user.
#
anon_upload_enable=YES
#
# Default umask for anonymus users is 077. You may wish to change this to 022,
# if your users expect that (022 is used by most other ftpd's)
#
#anon_umask=022
#
# Uncomment this if you want the anonymous FTP user to be able to create
# new directories.
#
anon_mkdir_write_enable=YES
#
# Uncomment this to enable anonymus FTP users to perform other write operations
# like deletion and renaming.
#
anon_other_write_enable=YES
#
# If you want, you can arrange for uploaded anonymous files to be owned by
# a different user. Note! Using "root" for uploaded files is not
# recommended!
#
#chown_uploads=YES
#chown_username=whoever
#
# The maximum data transfer rate permitted, in bytes per second, for anonymous
# authenticated users. The default is 0 (unlimited).
#
#anon_max_rate=7200
# Log Settings
#
# Log to the syslog daemon instead of using an logfile.
#
syslog_enable=YES
#
# Uncomment this to log all FTP requests and responses.
#
#log_ftp_protocol=YES
#
# Activate logging of uploads/downloads.
#
#xferlog_enable=YES
#
```

```
# You may override where the log file goes if you like. The default is shown
# below.
#
#vsftpd_log_file=/var/log/vsftpd.log
#
# If you want, you can have your log file in standard ftpd xferlog format.
# Note: This disables the normal logging unless you enable dual_log_enable below.
#
#xferlog_std_format=YES
#
# You may override where the log file goes if you like. The default is shown
# below.
#
#xferlog_file=/var/log/xferlog
#
# Enable this to have booth logfiles. Standard xferlog and vsftpd's own style log.
#
#dual_log_enable=YES
#
# Uncomment this to enable session status information in the system process
listing.
#
#setproctitle_enable=YES
# Transfer Settings
#
# Make sure PORT transfer connections originate from port 20 (ftp-data).
#
connect_from_port_20=YES
#
# You may change the default value for timing out an idle session.
#
#idle_session_timeout=600
#
# You may change the default value for timing out a data connection.
#
#data_connection_timeout=120
#
# Enable this and the server will recognise asynchronous ABOR requests. Not
# recommended for security (the code is non-trivial). Not enabling it,
# however, may confuse older FTP clients.
#
#async_abor_enable=YES
#
# By default the server will pretend to allow ASCII mode but in fact ignore
# the request. Turn on the below options to have the server actually do ASCII
# mangling on files when in ASCII mode.
# Beware that turning on ascii_download_enable enables malicious remote parties
# to consume your I/O resources, by issuing the command "SIZE /big/file" in
# ASCII mode.
# These ASCII options are split into upload and download because you may wish
# to enable ASCII uploads (to prevent uploaded scripts etc. from breaking),
# without the DoS risk of SIZE and ASCII downloads. ASCII mangling should be
# on the client anyway..
#
```

```
#ascii_upload_enable=YES
#ascii_download_enable=YES
#
# Set to NO if you want to disallow the  PASV  method of obtaining a data
# connection.
#
#pasv_enable=NO
# PAM setting. Do NOT change this unless you know what you do!
#
pam_service_name=vsftpd
# Set to ssl_enable=YES if you want to enable SSL
ssl_enable=NO
# Limit passive ports to this range to assis firewalling
pasv_min_port=30000
pasv_max_port=30100
anon_mkdir_write_enable=YES
anon_root=/srv/ftp
anon_upload_enable=YES
chroot_local_user=NO
ftpd_banner=Welcome message
idle_session_timeout=900
local_enable=YES
log_ftp_protocol=YES
max_clients=10
max_per_ip=3
pasv_enable=YES
ssl_sslv2=NO
ssl_sslv3=NO
ssl_tlsv1=YES
write_enable=YES
listen=NO
listen_ipv6=YES
```

C.4.4.5 Restart the Necessary Services

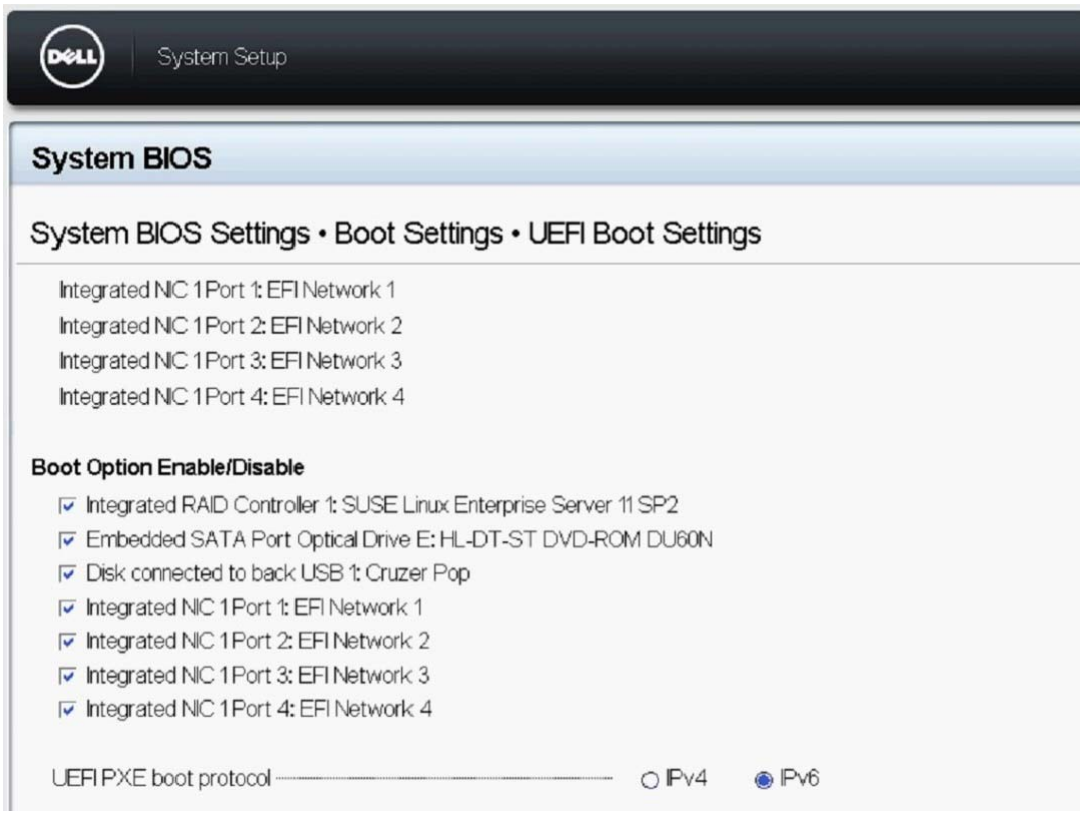
```
Service nfsserver restart
Service xinetd restart
Service vsftpd restart
Service dhcpd6 restart
```

C.4.5 PXE Client Setup

To set up the PXE client:

1. Boot the server and press the F2 key to enter system setup.
2. Select **System BIOS Settings > Boot Settings** and change the boot mode to UEFI.
3. Reboot the server.
4. Press the F2 key to enter system setup.
5. Select **System BIOS Settings > Boot Settings > UEFI Boot Settings** and set the UEFI PXE boot protocol to IPv6 in the following dialog.

Figure 175 UEFI Boot Settings Dialog



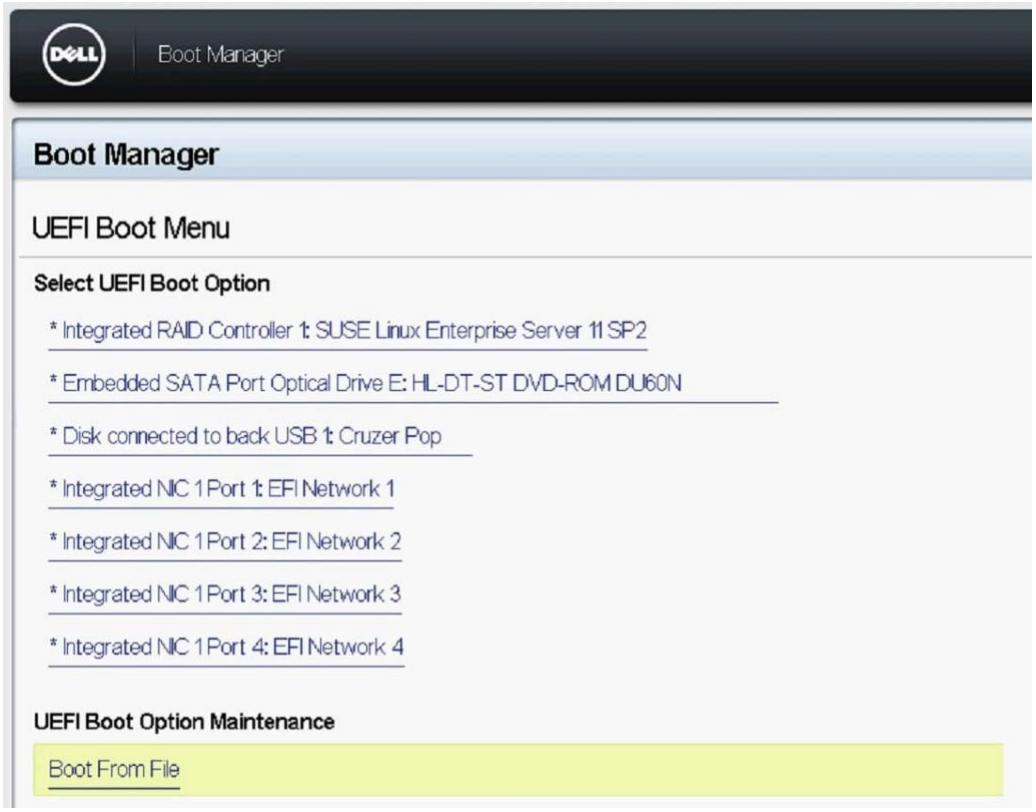
6. Reboot the server.
7. Press the F11 key and enter the UEFI boot manager. The following Boot Manager Main Menu (Figure 176) appears.

Figure 176 Boot Manager Main Menu



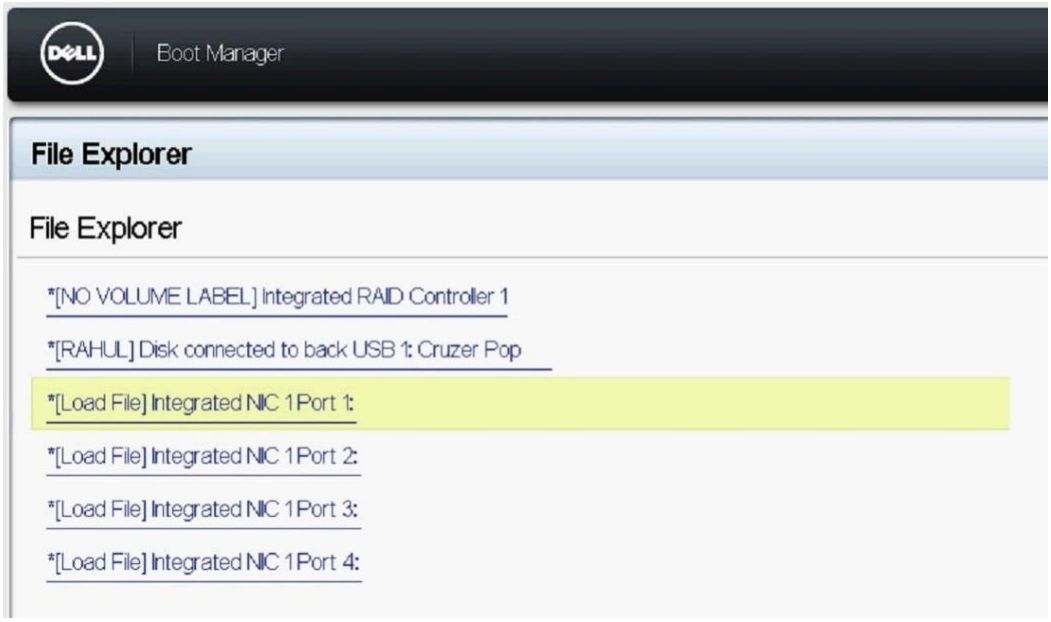
8. Select UEFI Boot Menu. The following UEFI Boot Menu (Figure 177) appears.

Figure 177 UEFI Boot Menu



9. Select Boot From File. The following File Explorer dialog appears.

Figure 178 File Explorer Dialog



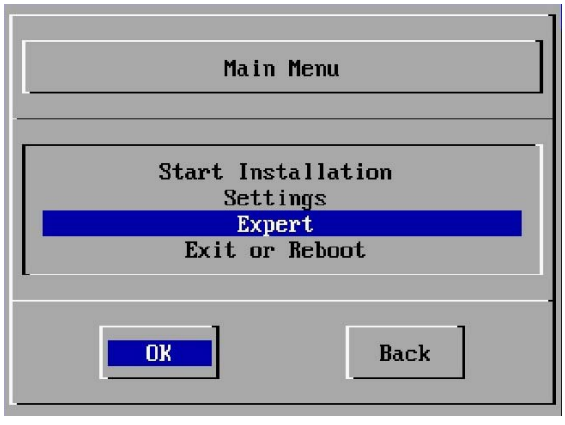
10. Select the port on the client through which PXE boot must be performed and press Enter. The system displays a status message and then displays the following Language Selection menu.

Figure 179 Language Selection Menu



11. Select a language and click OK. The following Main Menu appears.

Figure 180 Main Menu



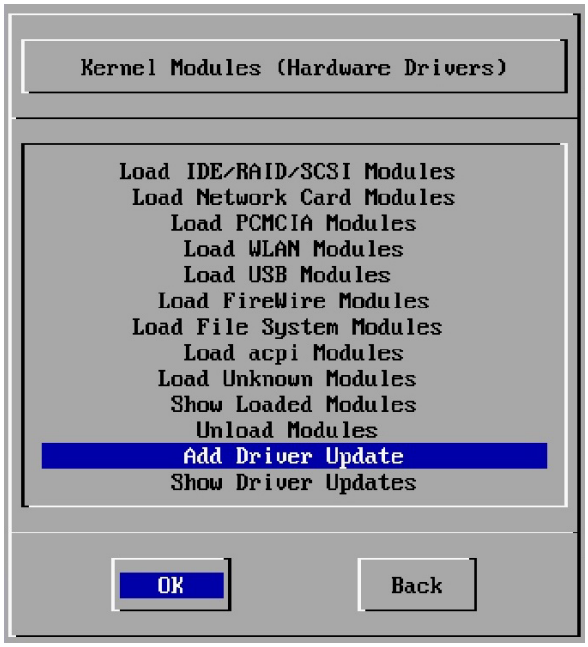
12. Select Expert and click OK. The following Expert menu (Figure 181) appears.

Figure 181 Expert Menu



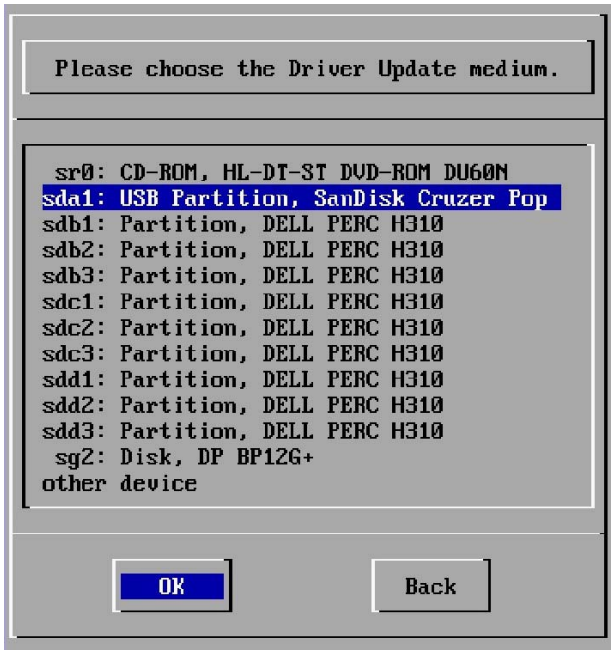
13. Select Kernel Modules (Hardware Drivers) and click OK. The following Kernel Modules (Hardware Drivers) menu appears.

Figure 182 Kernel Modules (Hardware Drivers) Menu



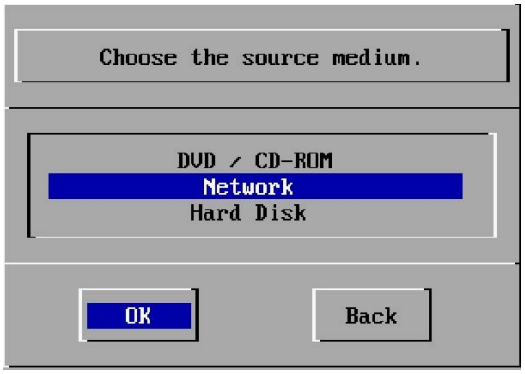
14. Select Add Driver Update and click OK. The following Driver Update Medium Selection list (Figure 183) appears.

Figure 183 Driver Update Medium Selection List



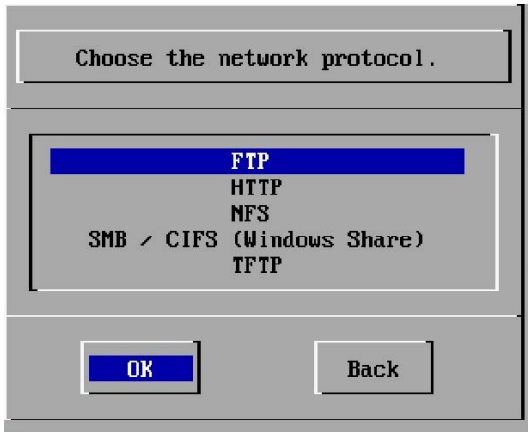
15. Select a medium from the list and click OK. The Expert menu (Figure 181) appears.
16. Click Back. The Main Menu (Figure 180) appears.
17. Select Start Installation and click OK. The following Source Medium list appears.

Figure 184 Source Medium List



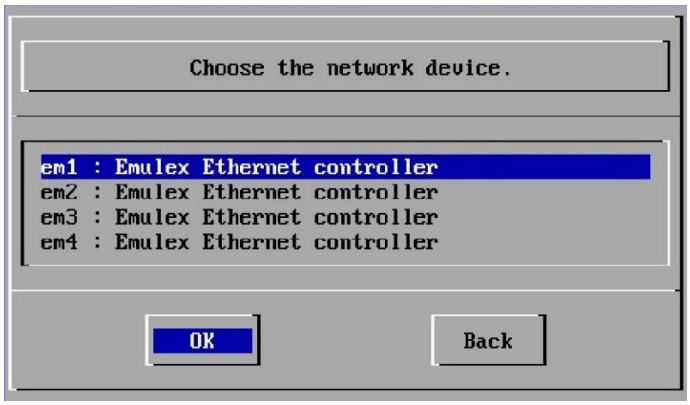
18. Select Network and click OK. The following Network Protocol list (Figure 185) appears.

Figure 185 Network Protocol List



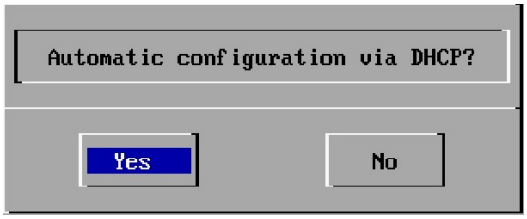
19. Select FTP and click OK. The following Network Device list appears.

Figure 186 Network Device List



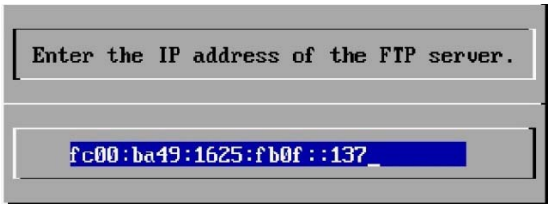
20. Select the device and click OK. The following dialog box prompts you for automatic configuration using DHCP.

Figure 187 Automatic Configuration via DHCP Dialog Box



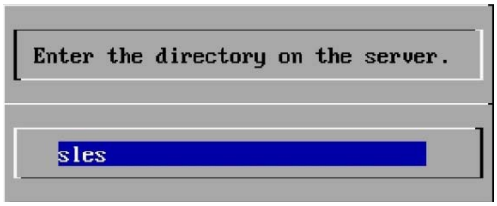
21. Click Yes. The DHCP6 request is sent. The following dialog box (Figure 188) prompts you to enter the IPv6 address of the FTP server.

Figure 188 FTP Server IPv6 Address Dialog Box



22. Enter the IPv6 address and press Enter. The following dialog box prompts you to enter the directory path of the server.

Figure 189 Server Directory Path Dialog Box



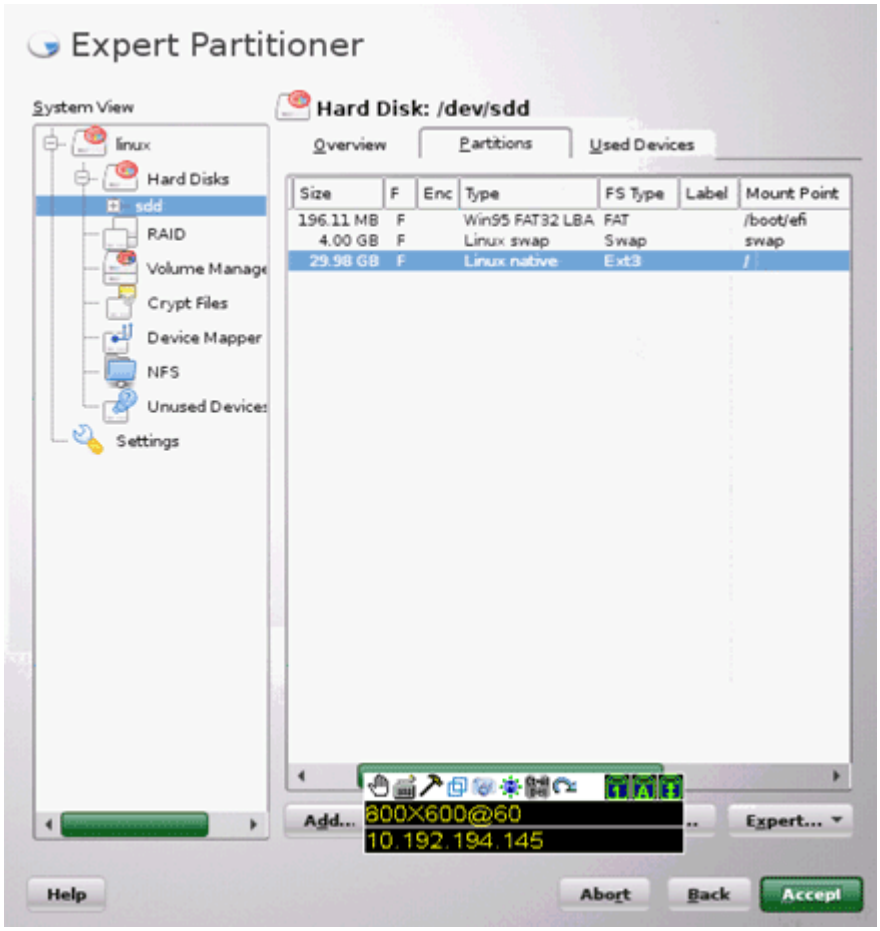
23. Enter the server's directory path and press Enter. A pop up will open if you must enter a username and password for the FTP server.
24. Select NO for this selection and select NO for using HTTP Proxy. The installation begins and the system displays the License Agreement page.
25. After you accept the license agreement, all of the operating system packages are downloaded using FTP over IPv6.

Appendix D: Example for Installing and Booting UEFI FCoE

The following steps use the Linux SLES operating system as an example.

1. Boot to UEFI Linux from DVD.
2. Select UEFI boot from the BIOS boot manager or boot to shell and execute /efi/boot/elilo.efi from removable media.
3. Follow normal boot instructions from YaST.
4. During the Installation Settings, select Partitioning, then select Custom partitioning (for experts).
5. In SystemViews, select a hard disk.
6. Select the Partitions tab as shown in the following window.

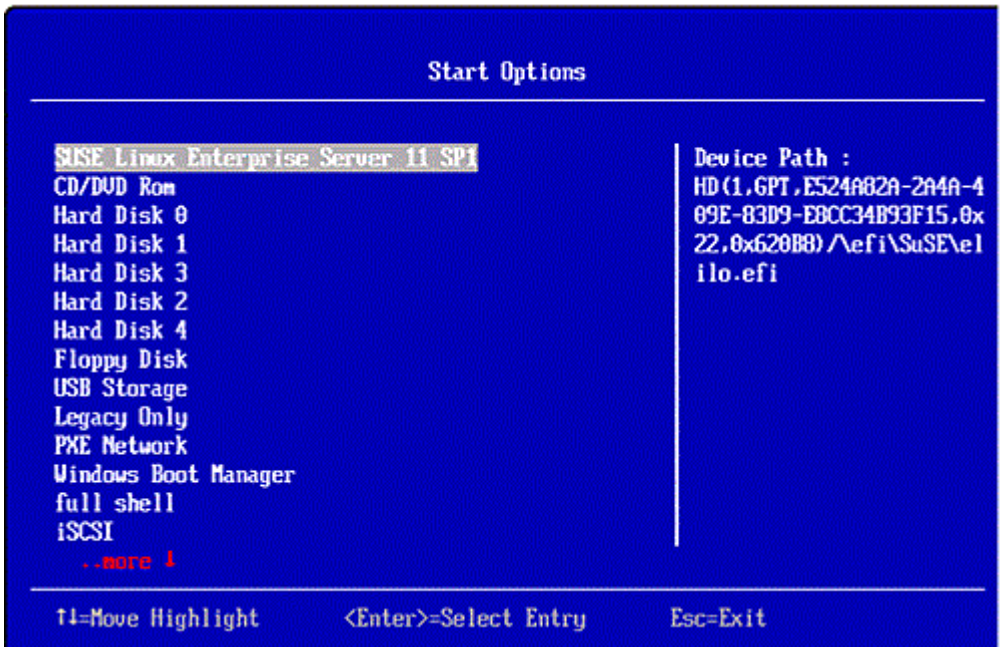
Figure 190 Partitions Tab



7. Delete any old partitions on the disk.
8. To create a GPT disklabel (UEFI does not support booting from disks with an MS-DOS partition table), click Expert (above the Accept button).
9. Create a partition table and check GPT (required for UEFI boot).
10. To add an UEFI system partition (VFAT or FAT), click Add.

11. Select a custom size (100 to 400MB) depending on what UEFI driver and utilities you want to store on the UEFI system partition. In general, using 200MB is sufficient if the plug-in card detail for UEFI system partition is unknown.
12. In the Formatting options dialog box, select **Format Partition > File system > FAT**.
13. Under Mounting options, select /boot/efi and click Finish. The boot drive now has a GPT disk label that contains a FAT EFI boot partition.
14. Perform the remainder of the partitioning as with any Linux installation. The remainder of the installation is not unique to UEFI, for example, add a swap partition and an EXT3 partition for root.
15. After the installation is completed, view the partitions using the `parted -l` command.
16. When the system boots, SUSE Linux Enterprise Server 11 SP1 now shows in the system Start Options.

Figure 191 UEFI FCoE Start Options



Appendix E: Loading and Unloading UEFIBoot from the UEFI Shell

This section explains how to load and unload Offbeat from the UEFI shell.

E.1 Loading UEFIBoot from the UEFI Shell

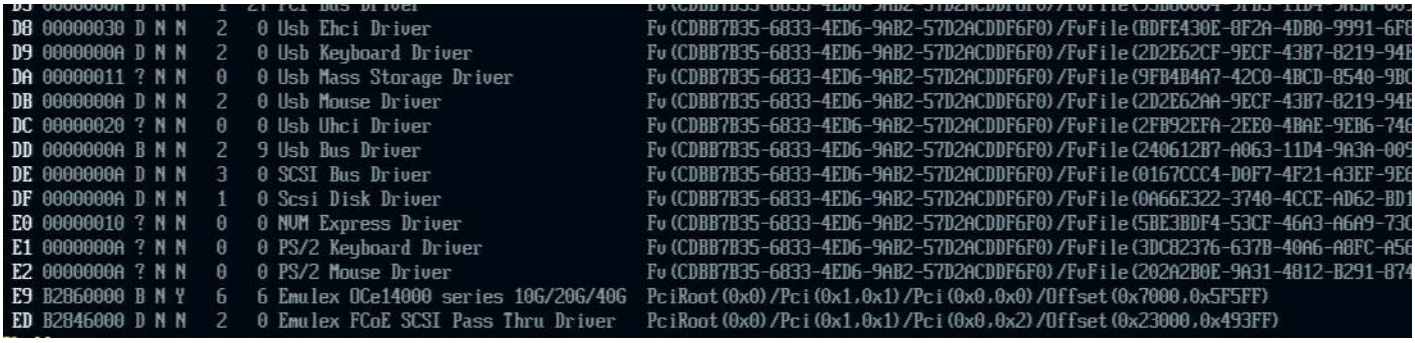
The UEFI boot code is distributed with the firmware in a.UFI file. This file can be downloaded to a NIC, iSCSI, or FCoE adapter through the HII interface in the NIC UEFI code.

E.2 Unloading UEFIBoot from the UEFI Shell

To unload UEFIBoot:

1. View Emulex driver handle information by typing `drivers` and pressing Enter at the shell prompt. The system displays a list of drivers as shown in the following window.

Figure 192 Driver Listing



2. Write down the Emulex NIC, FCoE, or iSCSI driver handles. For the above case, the Emulex NIC driver handle is 122 and the Emulex iSCSI driver handle is 126.
3. To unload the Emulex NIC driver in the previous example, type `unload 122` and press Enter.
4. To unload the Emulex iSCSI driver in the previous example, type `unload 126` and press Enter.

Appendix F: Dell UEFI

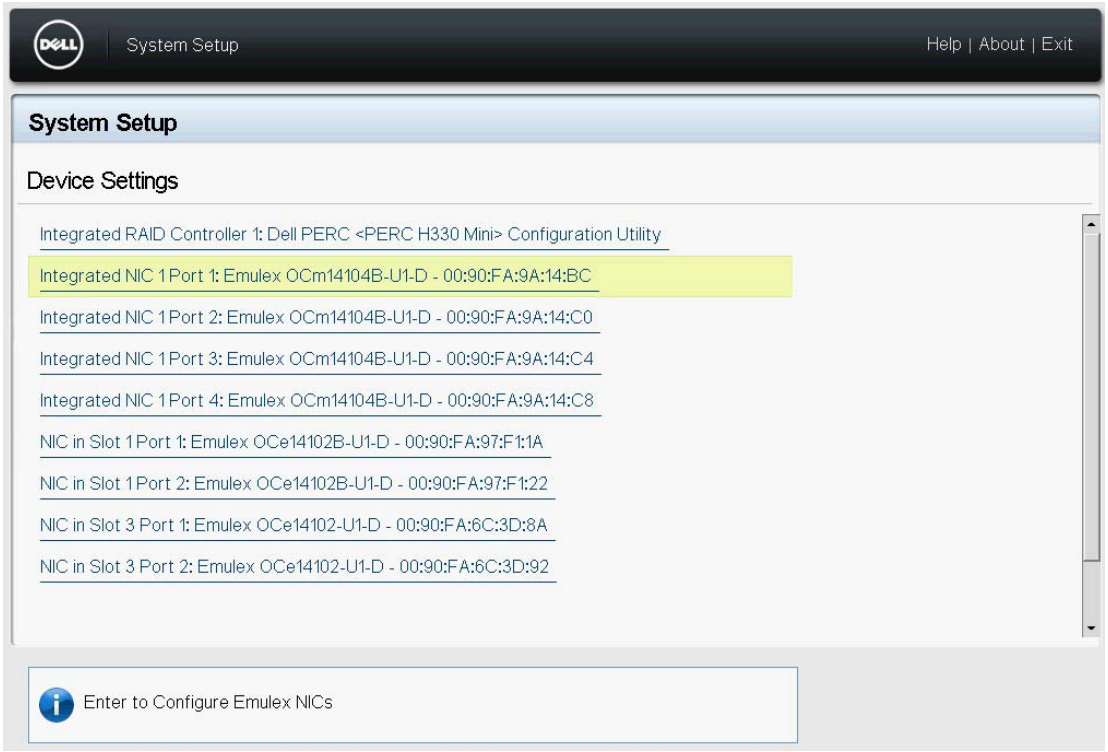
The following sections define and describe the fields and configuration settings for Dell adapters using the Dell UEFI utility.

The Dell UEFI adapter configuration utility is a unified configuration utility that lets you configure all of the available protocols through a single GUI interface, unlike individual configuration utilities. The iSCSI and FCoE configuration can be configured using these menus even if those functions are not physically present. The configuration will be applied after those respective functions are enabled.

F.1 Accessing the Main Configuration Page

From the **System Setup** > **Device Settings** page, select the adapter you want to view.

Figure 193 Device Settings Page



The Main Configuration page (Figure 194) for that adapter appears.

F.2 Main Configuration Page

This page displays device information and lets you configure device parameters.

Figure 194 Main Configuration Page

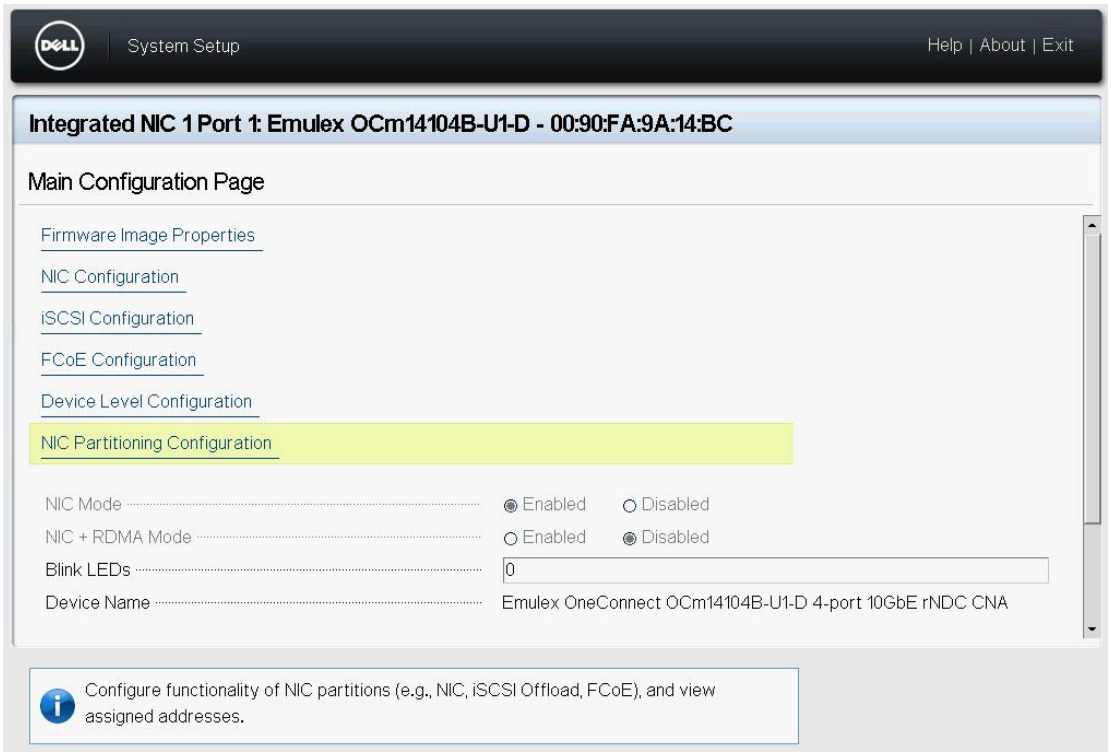


Table 14 Main Configuration Page Menu Options

Menu Option	Description and Available Settings
Firmware Image Properties	View adapter firmware version information. See Section F.2.1, Firmware Image Properties Page, for more information.
NIC Configuration	Configure boot protocol, Wake on LAN, link speed, and VLAN. See Section F.2.2, NIC Configuration Page, for more information.
iSCSI Configuration	Configure general, initiator, and target parameters for iSCSI boot. See Section F.2.3, iSCSI Configuration Page, for more information. **NOTE** This menu appears only when NIC + RDMA mode is disabled.
FCoE Configuration	Configure FCoE boot parameters. See Section F.2.4, FCoE Configuration Page, for more information. **NOTE** This setting appears only when NIC + RDMA mode is disabled.
Device Level Configuration	View and configure global device level parameters. See Section F.2.5, Device Level Configuration Page, for more information.
NIC Partitioning Configuration	Configure functionality of NIC partitions and view assigned addresses. See Section F.2.6, NIC Partitioning Configuration Page, for more information. **NOTE** This menu appears only when NIC + RDMA mode is disabled.
NIC Mode	Always enabled.
NIC + RDMA Mode	Configure port usage for both Ethernet and RDMA traffic. **NOTE** When this setting is enabled, Par cannot be enabled. **NOTE** This setting is at the port level. Therefore, one port can have NIC + RDMA mode enabled while another port can have it disabled.

Table 14 Main Configuration Page Menu Options (Continued)

Menu Option	Description and Available Settings
Blink LEDs	Identify the physical port by blinking the port LED for up to 15 seconds. Available settings are 0 to 15 seconds (default).
Device Name	The name of the adapter.
Chip Type	The name and revision of the chip.
PCI Device ID	The PCI device ID of the port.
PCI Address	The PCI address assigned by the BIOS.
Link Status	The network link status reported by the adapter.

F.2.1 Firmware Image Properties Page

The Firmware Image Properties page displays adapter firmware version information.

Figure 195 Firmware Image Properties Page



Table 15 Firmware Image Properties Page Menu Options

Menu Option	Description and Available Settings
Family Firmware Version	Firmware version information for the installed adapter.
Controller IPL Version	IPL version information for the installed adapter.
Controller BIOS Version	BIOS version information for the installed adapter.
EFI Version	UEFI driver version information for the installed adapter.

F.2.2 NIC Configuration Page

The NIC Configuration page allows you to configure the boot protocol, Wake on LAN, link speed, and VLAN.

Figure 196 NIC Configuration Page

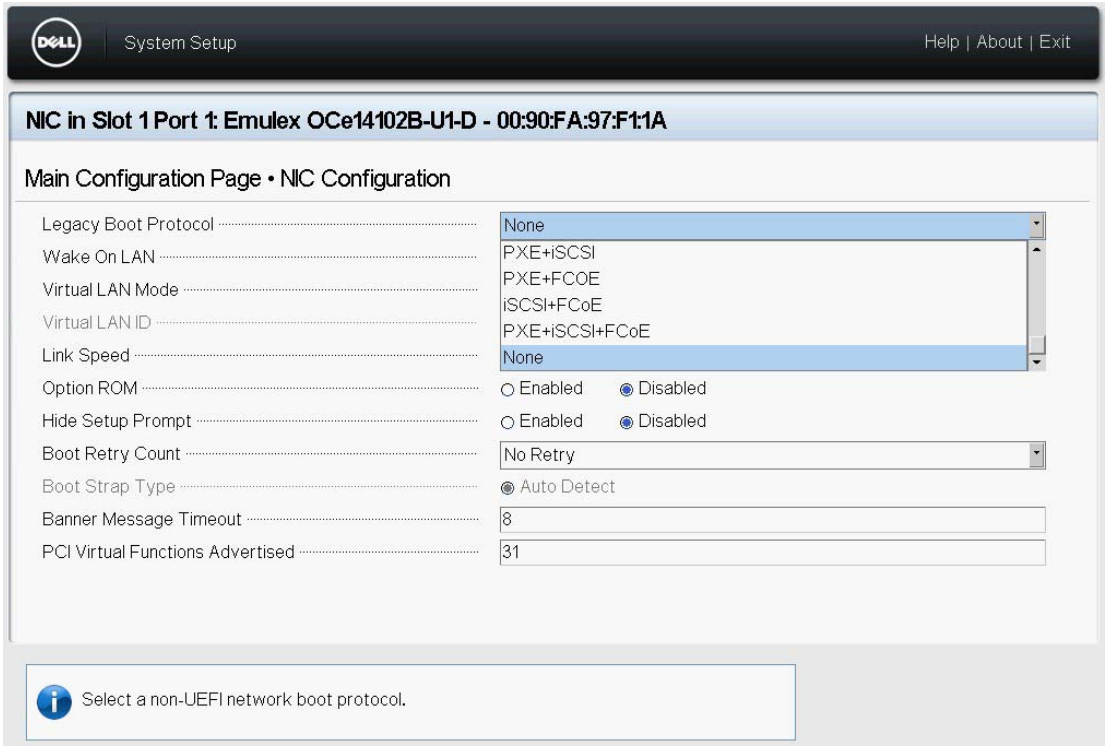


Table 16 NIC Configuration Page Menu Options

Menu Option	Description and Available Settings
Legacy Boot Protocol	Select a non-UEFI network boot protocol. Available settings may include: ■ PXE (default) ■ iSCSI ■ FCoE ■ PXE+iSCSI ■ PXE+FCoE ■ iSCSI+FCoE ■ PXE+iSCSI+FCoE ■ None
Wake On LAN	Enable or disable Wake On LAN. Available settings include: ■ Enabled (default) ■ Disabled
Virtual LAN Mode	**NOTE** This setting is for legacy PXE boot only. For UEFI PXE VLAN configuration through the system BIOS, refer to the Dell system documentation. Enable or disable virtual LAN mode. Available settings include: ■ Enabled ■ Disabled (default)
Virtual LAN ID	**NOTE** This setting is for legacy PXE boot only. For UEFI PXE VLAN configuration through the system BIOS, refer to the Dell system documentation. Specify the ID for VLAN mode. The VLAN ID valid range is 0 to 4095.
VIrtual LAN Priority	Specify the Priority for VLAN mode. The VLAN Priority valid range is 0 to 7.
Physical Link Speed	The current maximum port speed on the adapter.
Link Speed	Specify the port speed used for the selected protocol. Available settings include: ■ Auto Negotiated (default) ■ 1Gb/s ■ 10Gb/s
Option ROM	Enable or disable legacy boot protocols in the Option ROM. Available settings include: ■ Enabled ■ Disabled (default)
Hide Setup Prompt	Enable or disable the option ROM setup prompt (banner) during POST. Available settings include: ■ Enabled ■ Disabled (default)

Table 16 NIC Configuration Page Menu Options (Continued)

Menu Option	Description and Available Settings
Boot Retry Count	Control the number of retries in case of boot failure. Available settings include: ■ No Retry (default) ■ 1 Retry ■ 2 Retries ■ 3 Retries ■ 4 Retries ■ 5 Retries ■ 6 Retries ■ Indefinite Retries
Boot Strap Type	Control the boot strap method used to boot to the operating system. Available settings include: ■ Auto Detect (default)
Banner Message Timeout	Control the number of seconds that the Option ROM banner appears during POST. The valid range is 0 to 14. The default value is 8.
PCI Virtual Functions Advertised	The number of PCI VFs to be advertised on this port in non-NPar mode. The number of VFs advertised on this port must be less than or equal to the number of VFs supported on this port. This value varies based on the adapter in use. For dual port Emulex OCe14000-series adapters, the maximum number of supported VFs per port is 24. For dual port Emulex OCe14000B-series adapters, the maximum number of supported VFs per port is 56. For quad-port adapters, the maximum number of supported VFs per port is 28.

NOTE

For Vindicator 2, Vindicator 3, Pavelow 2, Pavelow 3 adapters, the requirement that both the legacy and the UEFI BIOS functionality be the same mandates that the Legacy Boot Protocol be set to PXE on the first port of the adapter when defaults are loaded.
PXE boot is enabled only when loading defaults on the first port due to the architecture of the UEFI BIOS. The first port will have PXE Boot enabled and the rest of the ports will be set to disabled.
When defaults are loaded on any port other than the first port, the Legacy Boot Protocol will be set to None only on that specific port.
When defaults are loaded, the Virtualization mode is set to None which in turn affects the Legacy Boot Protocol on all the other ports. Hence on the ports other than the one on which the default is loaded, the Legacy Boot Protocol will be set to None if a storage protocol was selected earlier and PXE if PXE alone or PXE in combination with storage protocol was selected earlier.

F.2.3 iSCSI Configuration Page

The iSCSI Configuration page (Figure 197) lets you configure general, initiator, and target parameters for iSCSI boot.

Figure 197 iSCSI Configuration Page

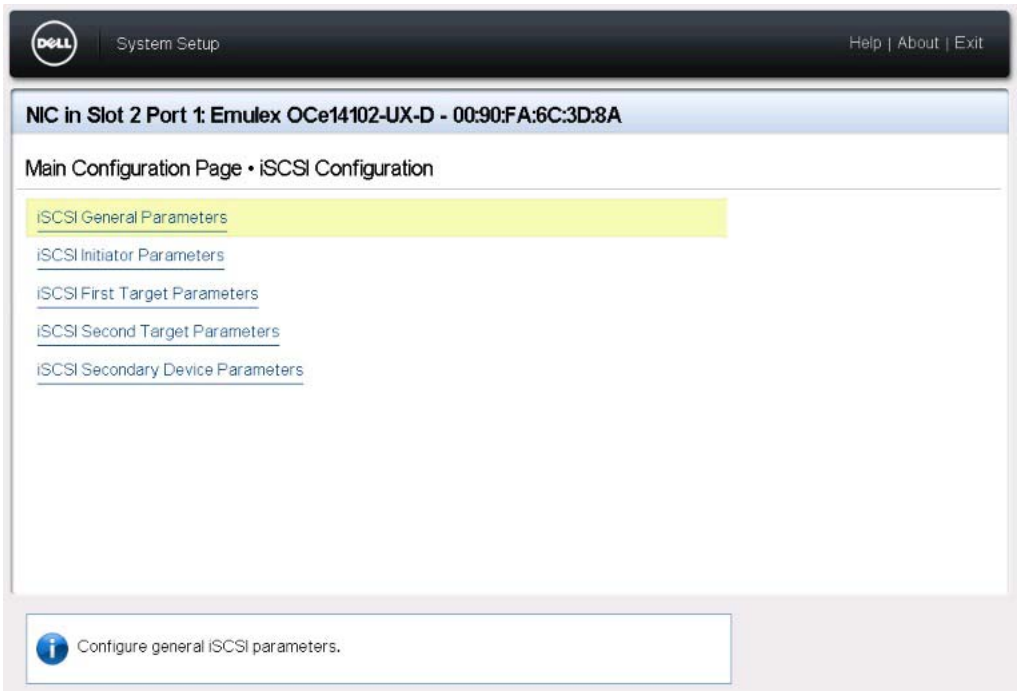


Table 17 iSCSI Configuration Page Menu Options

Menu Option	Description and Available Settings
iSCSI General Parameters	Configure general iSCSI parameters. See Section F.2.3.1, iSCSI General Parameters Page, for more information.
iSCSI Initiator Parameters	Configures iSCSI initiator parameters. See Section F.2.3.2, iSCSI Initiator Parameters Page, for more information.
iSCSI First Target Parameters	Enable connection and configure parameters for the first iSCSI target. See Section F.2.3.3, iSCSI First Target Parameters Page, for more information.
iSCSI Second Target Parameters	Enable connection and configure parameters for the second iSCSI target. See Section F.2.3.4, iSCSI Second Target Parameters Page, for more information.
iSCSI Secondary Device Parameters	Configures the iSCSI secondary boot device parameters. See Section F.2.3.5, iSCSI Secondary Device Parameters Page, for more information.
iSCSI Target Dynamic Discovery Parameters	Configures the iSCSI target dynamic discovery parameters. See Section F.2.3.6, iSCSI Target Dynamic Discovery Parameters Page, for more information.

F.2.3.1 iSCSI General Parameters Page

The iSCSI General Parameters page (Figure 198) lets you configure the general iSCSI parameters.

Figure 198 iSCSI General Parameters Page

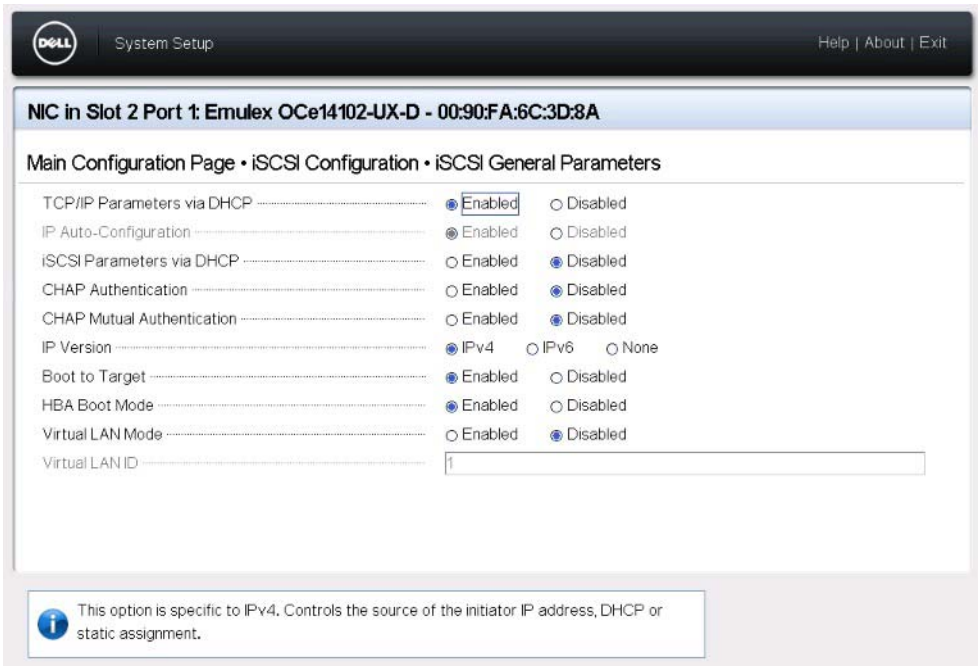


Table 18 iSCSI General Parameters Page Menu Options

Menu Option	Description and Available Settings
TCP/IP Parameters via DHCP	Control the source of the initiator IP address: DHCP or static assignment. Available settings include: ■ Enabled ■ Disabled (default) **NOTE** This option is specific to IPv4.
IP Auto-Configuration	Control the source of the initiator IP address: DHCP or static assignment. Available settings include: ■ Enabled ■ Disabled (default) **NOTE** This option is specific to IPv6.
iSCSI Parameters via DHCP	Enable the acquisition of iSCSI target parameters from DHCP. Available settings include: ■ Enabled ■ Disabled (default)
CHAP Authentication	Enable the ability of the initiator to use CHAP authentication when connecting to the iSCSI target. Available settings include: ■ Enabled ■ Disabled (default)

Table 18 iSCSI General Parameters Page Menu Options (Continued)

Menu Option	Description and Available Settings
CHAP Mutual Authentication	Setting to enable mutual CHAP authentication between the iSCSI initiator and target. To use mutual CHAP authentication, specify an initiator secret on the Initiator Parameters page and configure that secret on the target. Available settings include: ■ Enabled ■ Disabled (default)
IP Version	Control whether IPv4 or IPv6 network addressing is used for the iSCSI initiator. Available settings include: ■ IPv4 ■ IPv6 ■ None (default) – Dual mode
Boot to Target	Control whether the iSCSI initiator boots to the specified iSCSI target after connection. **NOTE** This is only applicable to UEFI iSCSI boot. Available settings include: ■ Enabled ■ Disabled (default)
Virtual LAN Mode	Enable or disable iSCSI Virtual LAN mode, which allows a VLAN tag to be used by iSCSI boot. Available settings include: ■ Enabled ■ Disabled (default)
Virtual LAN ID	Specify the ID (tag) to be used for iSCSI VLAN mode. The valid range for the VLAN ID is from 0 to 4095.
VIrtual LAN Priority	Specify the Priority for VLAN mode. The VLAN Priority valid range is 0 to 7.

F.2.3.2 iSCSI Initiator Parameters Page

The iSCSI Initiator Parameters page lets you configure the iSCSI initiator parameters.

Figure 199 iSCSI Initiator Parameters Page

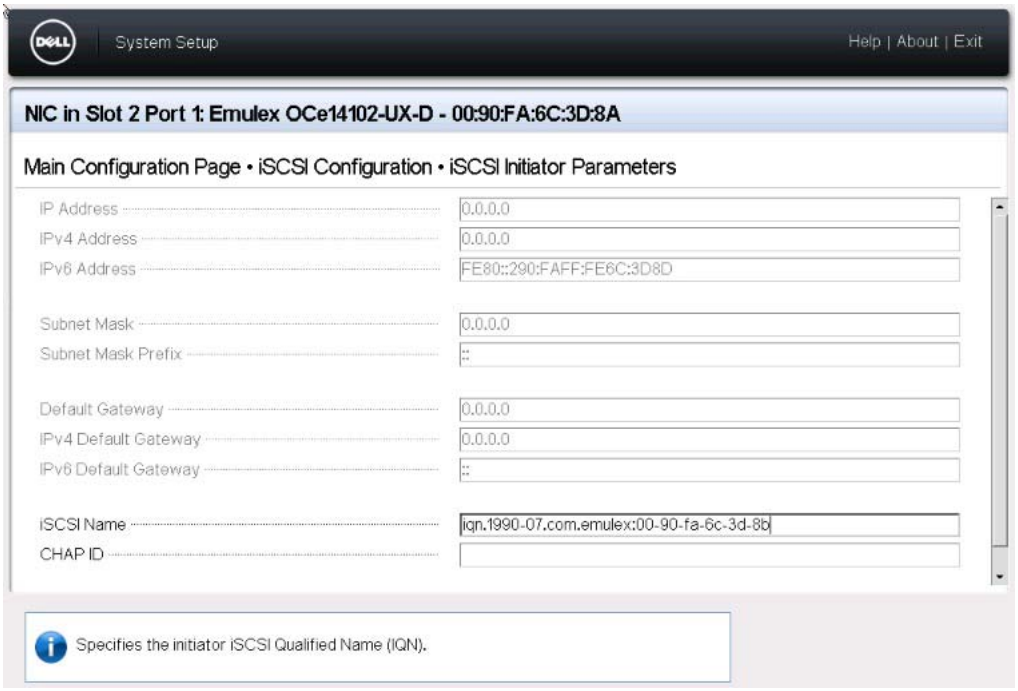


Table 19 iSCSI Initiator Parameters Page Menu Options

Menu Option	Description and Available Settings
IP Address	Specify the iSCSI Initiator IP address. The value should be in either IPv4 or IPv6 format and can be 2 to 39 characters long.
IPv4 Address	Specify the iSCSI Initiator IPv4 address. The value must be in IPv4 format and can be 7 to 15 characters long.
IPv6 Address	Specify the iSCSI Initiator IPv6 address. The value must be in IPv6 format and can be 2 to 39 characters long.
Subnet Mask	Specify the iSCSI Initiator IPv4 subnet mask. The value must be in IPv4 format and can be 7 to 15 characters long.
Subnet Mask Prefix	Specify the iSCSI Initiator IPv6 subnet mask. The value must be in IPv6 format and can be 2 to 39 characters long.
Default Gateway	Specify the iSCSI Initiator default gateway. The value should be in either IPv4 or IPv6 format and can be 2 to 39 characters long.
IPv4 Default Gateway	Specify the iSCSI Initiator IPv4 default gateway. The value must be in IPv4 format and can be 7 to 15 characters long.
IPv6 Default Gateway	Specify the iSCSI Initiator IPv6 default gateway The value must be in IPv6 format and can be 2 to 39 characters long.
iSCSI Name	Specify the initiator iSCSI Qualified Name (IQN). The valid range is 0 to 223 characters.
CHAP ID	Specify the initiator CHAP ID. The valid range is 0 to 128 characters. If this can be set on a per target basis, then the first target value must be stored here.
CHAP Secret	Specify the ISCSI Initiator CHAP Secret. The valid range is 0 or 12 to 16 characters. If this can be set on a per target basis, the value for the first target must be stored here.

F.2.3.3 iSCSI First Target Parameters Page

The iSCSI First Target Parameters page lets you enable a connection for the first iSCSI target and configure parameters.

Figure 200 iSCSI First Target Parameters Page



Table 20 iSCSI First Target Parameters Page Menu Options

Menu Option	Description and Available Settings
Connect	Enable connecting to the first iSCSI target. Available settings include: ■ Enabled ■ Disabled (default)
IP Address	Specify the first target IP address. The value should be in either IPv4 or IPv6 format and can be 2 to 39 characters long.
TCP Port	Specify the first target TCP port number. The valid range is 1 to 65535.
Boot LUN	Specify the first iSCSI storage target boot LUN. The valid range is 0 to 18446744073709551615.
iSCSI Name	Specify the first iSCSI storage target initiator IQN. The valid range is 0 to 223 characters.
CHAP ID	Specify the first ISCSI storage target CHAP ID. The valid range is 0 to 128 characters. If only a single value of the CHAP ID is supported, then it is stored here.
CHAP Secret	Specify the first ISCSI storage target CHAP Secret. The valid range is 0 or 12 to 16 characters. If only a single value of the CHAP Secret is supported, then it is stored here.
IP Version	Control whether IPv4 or IPv6 network addressing is used for the first iSCSI target. Available settings include: ■ IPv4 (default) ■ IPv6

Table 20 iSCSI First Target Parameters Page Menu Options (Continued)

Menu Option	Description and Available Settings
Header Digest	Enable the header digest for connecting to the iSCSI target. Available settings include: ■ Enabled ■ Disabled (default) The integrity of an iSCSI PDU header segment is protected by CRC32C checksum.
Data Digest	Enable the data digest for connecting to the iSCSI target. Available settings include: ■ Enabled ■ Disabled (default) The integrity of an iSCSI PDU header segment is protected by CRC32C checksum.
ISD Qualifier	Initiator session identifier. Value should be in the range of 1 to 65535.

F.2.3.4 iSCSI Second Target Parameters Page

The iSCSI Second Target Parameters page lets you enable a connection for the second iSCSI target and configure parameters.

Figure 201 iSCSI Second Target Parameters Page

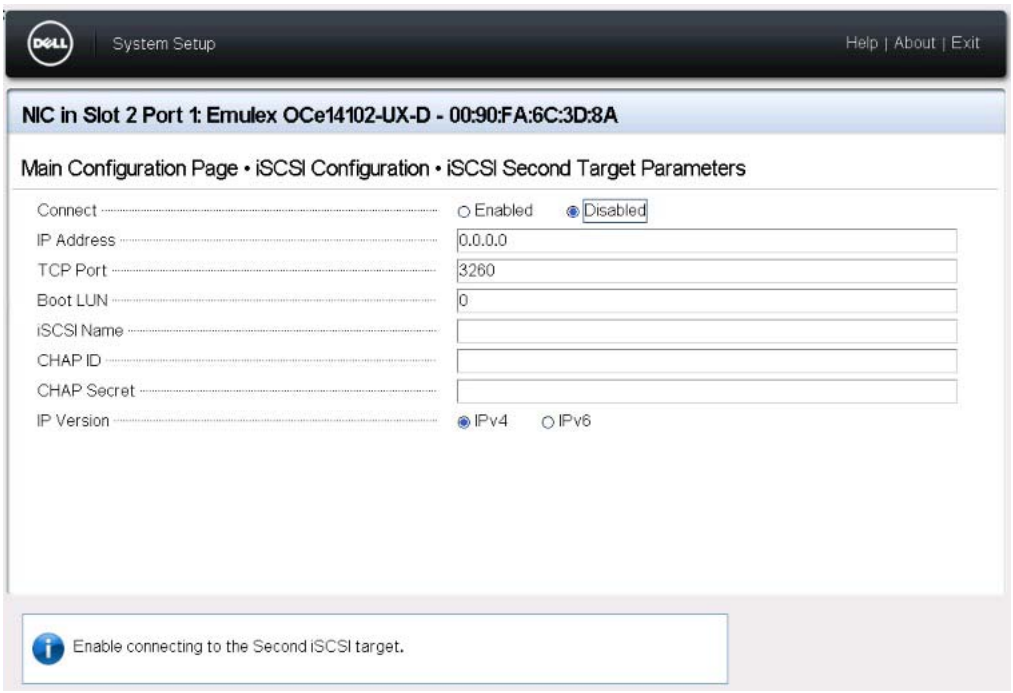


Table 21 iSCSI Second Target Parameters Page Menu Options

Menu Option	Description and Available Settings
Connect	Enable connecting to the second iSCSI target. Available settings include: ■ Enabled ■ Disabled (default)
IP Address	Specify the second target IP address. The value should be in either IPv4 or IPv6 format and can be 2 to 39 characters long.
TCP Port	Specify the second target TCP port number. The valid range is 1 to 65535.
Boot LUN	Specify the second iSCSI storage target boot LUN. The valid range is 0 to 18446744073709551615.
iSCSI Name	Specify the second iSCSI storage target initiator IQN. The valid range is 0 to 223 characters.
CHAP ID	Specify the second ISCSI storage target CHAP ID. The valid range is 0 to 128 characters. If only a single value of the CHAP ID is supported for all targets, then it is stored here.
CHAP Secret	Specify the second ISCSI storage target CHAP Secret. The valid range is 0 or 12 to 16 characters. If only a single value of the CHAP Secret is supported, then it is stored here.
IP Version	Control whether IPv4 or IPv6 network addressing is used for the second iSCSI target. Available settings include: ■ IPv4 (default) ■ IPv6

F.2.3.5 iSCSI Secondary Device Parameters Page

The iSCSI Secondary Device Parameters page lets you configure the MAC address for a secondary iSCSI boot adapter.

Figure 202 iSCSI Secondary Device Parameters Page

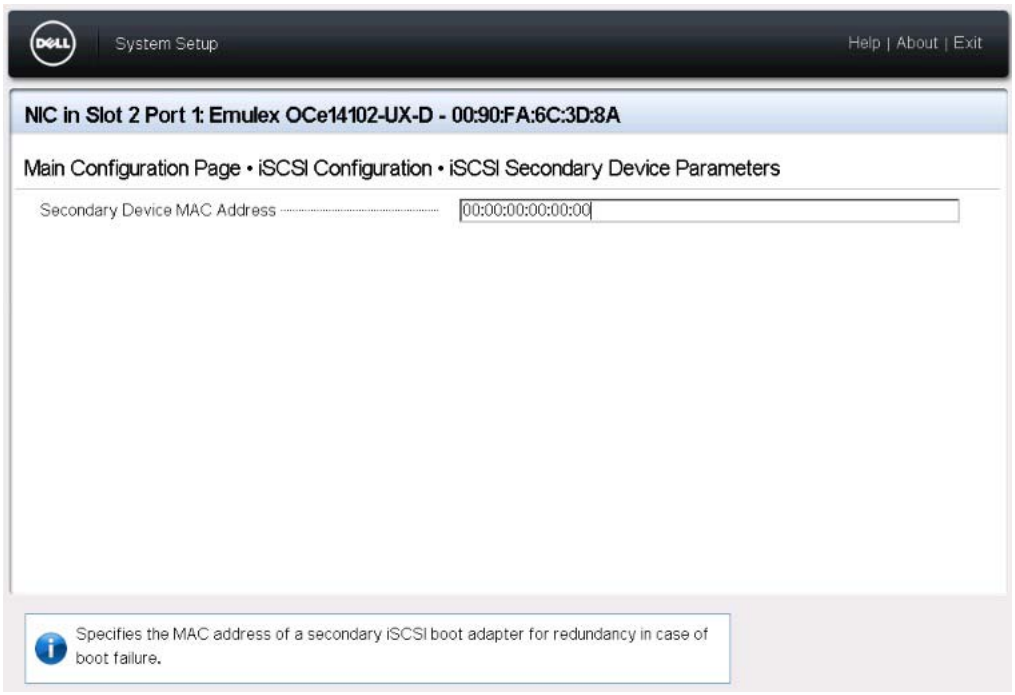


Table 22 iSCSI Secondary Device Parameters Page Menu Options

Menu Option	Description and Available Settings
Secondary Device MAC Address	Configure the MAC address for a secondary iSCSI boot adapter in case of boot failure. The MAC address can have a maximum of 17 characters with a colon (:) separator between every two characters.

F.2.3.6 iSCSI Target Dynamic Discovery Parameters Page

The iSCSI Target Dynamic Discovery Parameters page lets you to discover the iSCSI targets.

Figure 203 iSCSI Target Dynamic Discovery Parameters Page

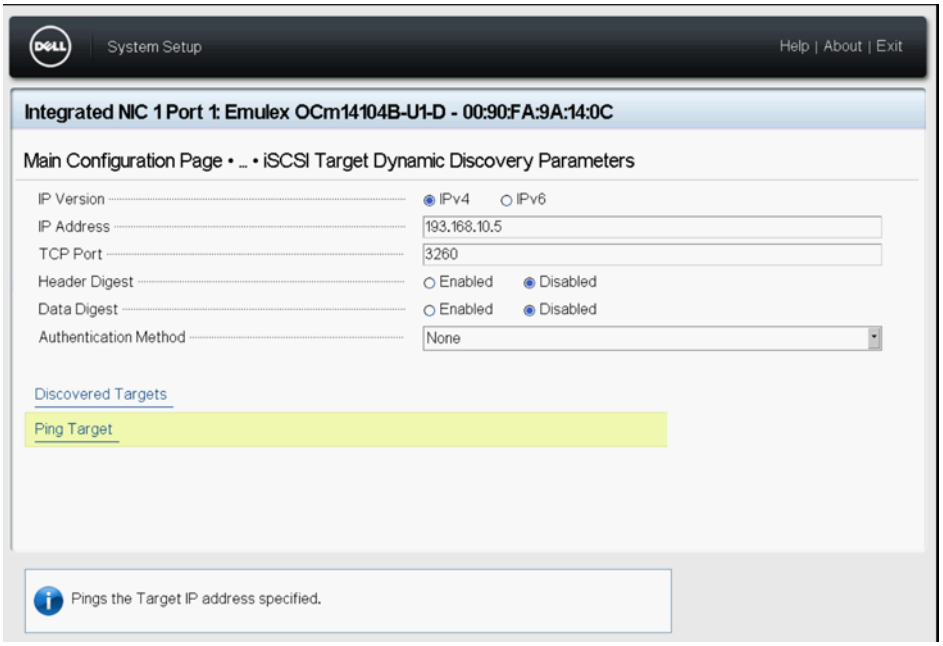


Table 23 iSCSI Target Dynamic Discovery Parameters Page Menu Options

Menu Option	Description and Available Settings
IP Version	Control whether IPv4 or IPv6 network addressing is used for the iSCSI target. Available settings include: ■ IPv4 ■ IPv6
IP Address	Specify the target IP address. The value should be in either IPv4 or IPv6 format and can be 2 to 39 characters long.
TCP Port	Specify the target TCP port number. The valid range is 1 to 65535. The default value is 3260.
Header Digest	Enable the header digest for connecting to the iSCSI target. Available settings include: ■ Enabled ■ Disabled (default) The integrity of an iSCSI PDU header segment is protected by CRC32C checksum.
Data Digest	Enable the data digest for connecting to the iSCSI target. Available settings include: ■ Enabled ■ Disabled (default) The integrity of an iSCSI PDU header segment is protected by CRC32C checksum.

Table 23 iSCSI Target Dynamic Discovery Parameters Page Menu Options (Continued)

Menu Option	Description and Available Settings
Authentication Method	Select an authentication method (optional). If you are enabling an authentication method, you are prompted to enter CHAP configuration. Available settings include: ■ None (default) ■ CHAP Authentication - Configure target side CHAP details. See Section F.2.3.6.1, Configuring CHAP Authentication. ■ CHAP Mutual Authentication - Configure target and initiator side CHAP details. See Section F.2.3.6.2, CHAP Mutual Authentication.
Discovered Targets	Check the connectivity and discover the listed targets. Click Discovered Targets, the iSCSI Target Configuration dialog appears. See Section F.2.3.6.3, Configuring Discovered Targets, for information on Configuring the target.
Ping Target	Check the connectivity to the target portal IP defined in the IP Address tab and displays success or failure messages.

F.2.3.6.1 Configuring CHAP Authentication

In the iSCSI Target Dynamic Discovery Parameters page (Figure 202), if you choose the Authentication Method as CHAP Authentication. You must configure the target side CHAP details. With CHAP Authentication, the target authenticates the initiator.

Figure 204 iSCSI Target Dynamic Discovery Parameters Page

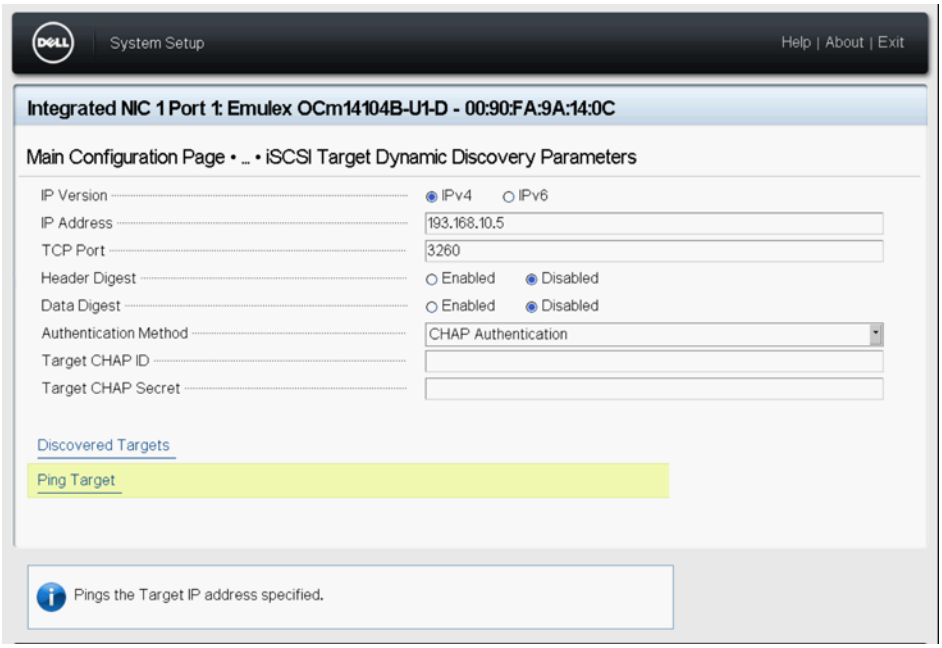


Table 24 CHAP Authentication Options

Menu Option	Description and Available Settings
Target CHAP ID	Specify the target CHAP ID. The valid range is 0 to 128 characters.
Target CHAP Secret	Specify the target CHAP secret. The valid range is 0 or 12 to 16 characters.

F.2.3.6.2 CHAP Mutual Authentication

In the iSCSI Target Dynamic Discovery Parameters page (Figure 202), if you choose the Authentication Method as CHAP Mutual Authentication. You must configure the target and initiator side CHAP details. With CHAP Mutual Authentication, the target authenticates the initiator and the initiator authenticates the target.

Figure 205 iSCSI Target Dynamic Discovery Parameters Page

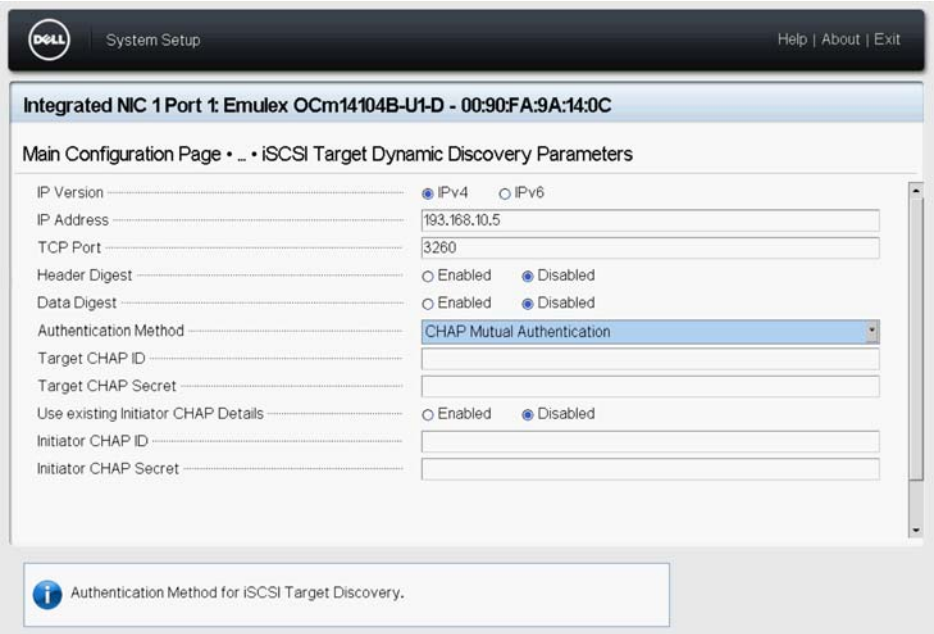


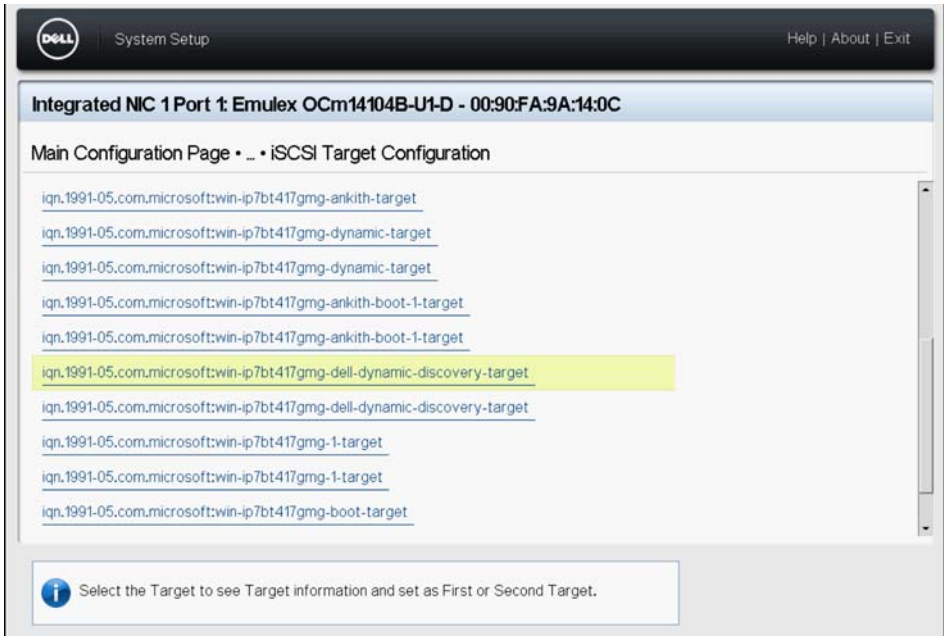
Table 25 CHAP Mutual Authentication Options

Menu Option	Description and Available Settings
Target CHAP ID	Specify the target CHAP ID. The valid range is 0 to 128 characters.
Target CHAP Secret	Specify the target CHAP secret. The valid range is 0 or 12 to 16 characters.
Use Existing Initiator CHAP Details	Enable using existing initiator CHAP details from the configured ISCSI initiator parameters. Available settings include: ■ Disabled (default) ■ Enabled
Initiator CHAP ID	Specify the initiator CHAP ID. The valid range is 0 to 128 characters. This option is available only if the Use Existing Initiator CHAP Details is disabled.
Initiator CHAP Secret	Specify the ISCSI Initiator CHAP Secret. The valid range is 0 or 12 to 16 characters. This option is available only if the Use Existing Initiator CHAP Details is disabled.

F.2.3.6.3 Configuring Discovered Targets

The iSCSI Target Configuration page lets you select the target for which you want to update the priority.

Figure 206 iSCSI Target Configuration Page



1. Select the target for which you want to update the priority. The following Target Priority Selection Page appears.

Figure 207 Target Priority Selection Page

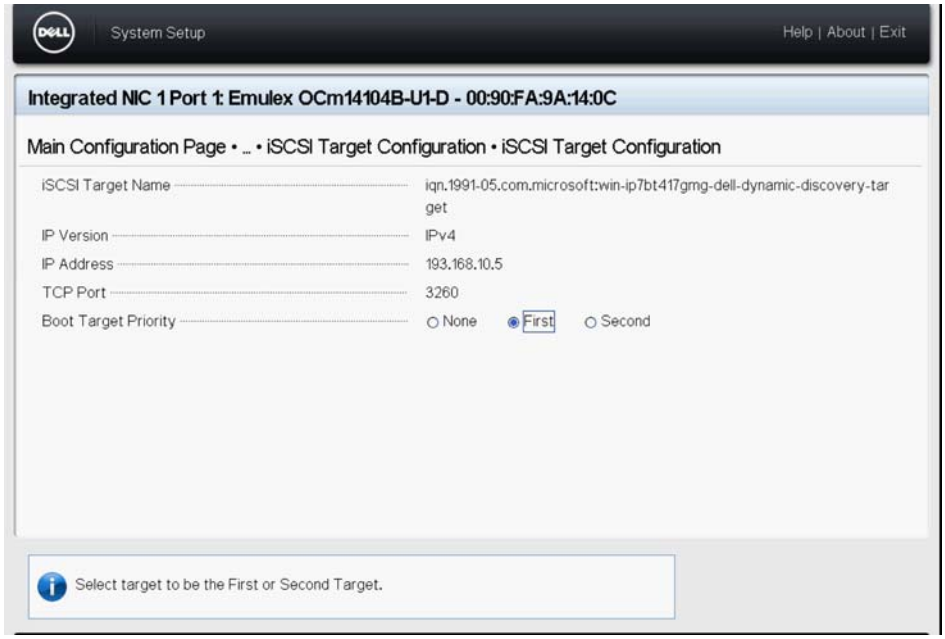


Table 26 iSCSI Target Dynamic Discovery Parameters Page Menu Options

Menu Option	Description and Available Settings
iSCSI Target Name	View the iSCSI Target Name.
IP Version	View network addressing used for the iSCSI target.
IP Address	View the target IP address.
TCP Port	View the target TCP port number.
Boot Target Priority	Select the Boot Target Priority. Available options include: ■ None (default) ■ First ■ Second The priority chosen is reflected in the respective iSCSI Target Parameters Page.

F.2.4 FCoE Configuration Page

The FCoE Configuration page lets you configure FCoE boot parameters.

Figure 208 FCoE Configuration Page

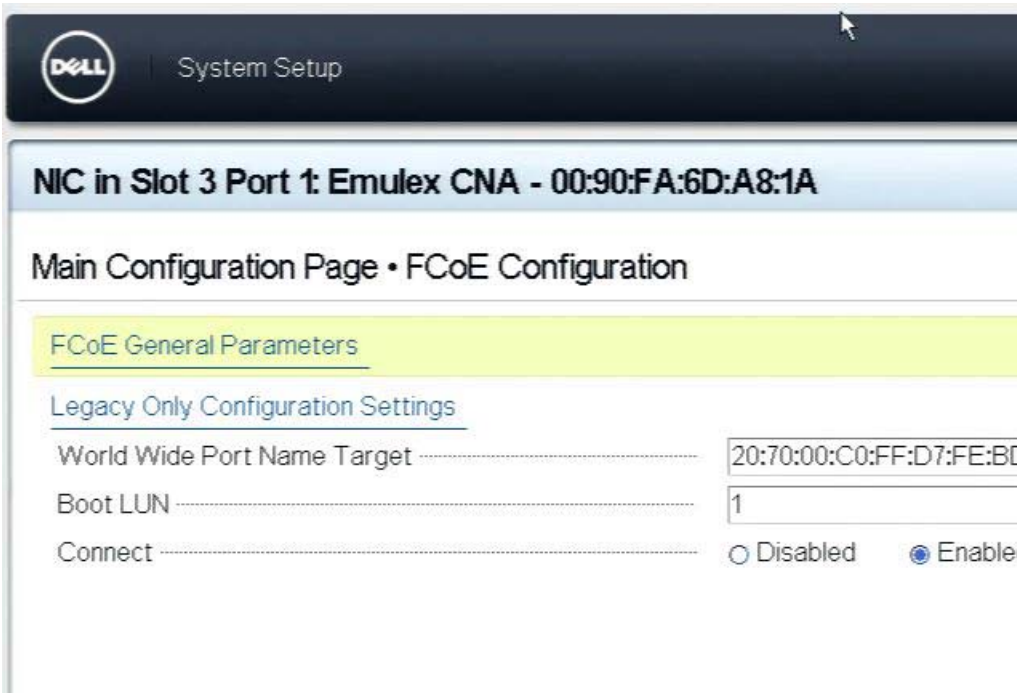


Table 27 FCoE Configuration Page Menu Options

Menu Option	Description and Available Settings
FCoE General Parameters	Configure FCoE general parameters. See Section F.2.4.1, FCoE General Parameters Page, for more information.
Legacy Only Configuration Settings	Configure settings for some adapters. See Section F.2.4.2, Legacy Only Configuration Settings Page, for more information.
Configure FCoE Targets	Configure FCoE targets. See Section F.2.4.3, Configure FCoE Targets Page, for more information.
World Wide Port Name Target	The first FCoE storage target WWPN. **NOTE** The Boot Scan Selection option must be set to Specified LUN before you can update this field.
Boot LUN	The first FCoE storage target LUN that will be used by the FCoE initiator as the system boot device. **NOTE** The Boot Scan Selection option must be set to Specified LUN before you can update this field.
Connect	Specify whether the FCoE initiator should connect to the first FCoE storage target defined. Available settings include: ■ Enabled ■ Disabled (default)

F.2.4.1 FCoE General Parameters Page

The FCoE General Parameters page lets you configure FCoE general parameters.

Figure 209 FCoE General Parameters Page

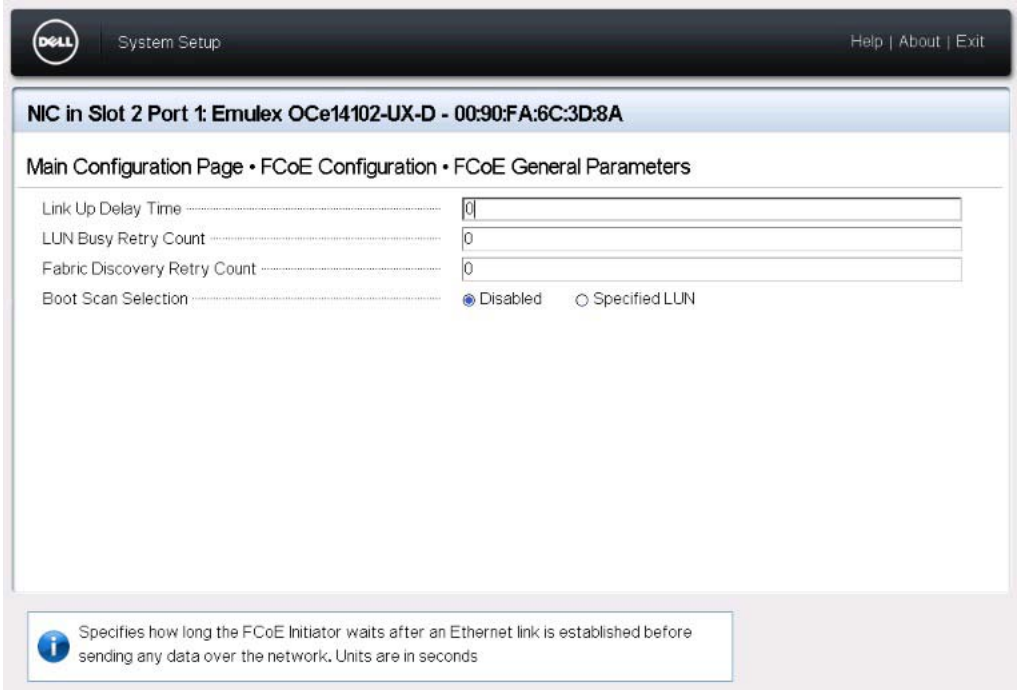


Table 28 FCoE General Parameters Page Menu Options

Menu Option	Description and Available Settings
Link Up Delay Time	Specify how long (in seconds) the FCoE Initiator waits after an Ethernet link is established before sending any data over the network. The valid range is 0 to 255 seconds.
LUN Busy Retry Count	Specify the number of connection retries the FCoE boot initiator attempts if the FCoE target LUN is busy. The valid range is 0 to 60.
Fabric Discovery Retry Count	Specify the retry count for FCoE fabric discovery. The valid range is 0 to 60.
Boot Scan Selection	Specify the adapter initiator behavior for booting the system from specified FC boot targets or fabric discovered targets. Available settings include: ■ **Disabled (default)** – Initiator does not attempt to boot. ■ **First LUN** – Attempt boot from the first bootable LUN of the targets discovered by adapter. ■ **First LUN 0** – Attempt boot from the first bootable LUN 0 of the targets discovered by adapter. ■ **First LUN Not LUN 0** – Attempt boot from the first bootable LUN that is not LUN 0 of the targets discovered by adapter. ■ **Fabric Discovered LUN** – Attempt boot based on the LUN assignment provided by the fabrics management infrastructure. ■ **Specified LUN** – Initiator attempts to boot to the specified World Wide Port Name and LUN for the first target.

F.2.4.2 Legacy Only Configuration Settings Page

The Legacy Only Configuration Settings page (Figure 210) lets you configure FCoE configuration settings for some adapters.

Figure 210 Legacy Only Configuration Settings Page

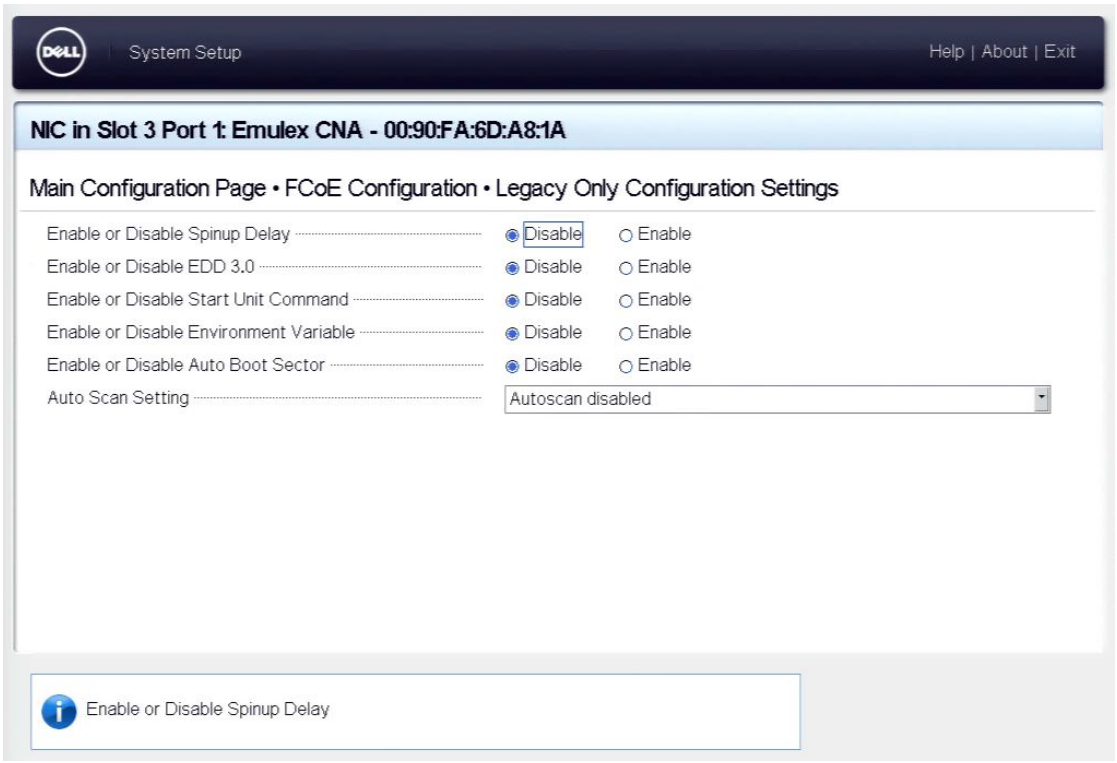


Table 29 Legacy Only Configuration Settings Page Menu Options

Menu Option	Description and Available Settings
Enable or Disable Spinup Delay	This option allows you to enable or disable the disk spinup delay. The factory default setting is disabled. If at least one boot device has been defined, and the spinup delay is enabled, the boot BIOS searches for the first available boot device. ■ If a boot device is present, the boot BIOS boots from it immediately. ■ If a boot device is not ready, the boot BIOS waits for the spinup delay and, for up to three additional minutes, continues the boot scanning algorithm to find another multiboot device. If boot devices have not been defined and auto scan is enabled, then the boot BIOS waits for five minutes before scanning for devices. ■ In an attached fabric, the boot BIOS attempts to boot from the first target found in the NameServer data.
Enable or Disable EDD 3.0	EDD 3.0 provides additional data to the operating system boot loader during an INT-13h function 48h call (get device parameters). This information includes the path to the boot device and the disk size. The default setting for EDD 3.0 is enabled. If EDD 3.0 is disabled, the setting reverts to EDD 2.1.
Enable or Disable Start Unit Command	You must know the specific LUN to issue the SCSI start unit command. The default setting is disabled.

Table 29 Legacy Only Configuration Settings Page Menu Options (Continued)

Menu Option	Description and Available Settings
Enable or Disable Environment Variable	This option sets the boot adapter order if the system supports the environment variable. The default setting is disabled.
Enable or Disable Auto Boot Sector	This option automatically defines the boot sector of the target disk for the migration boot process, which applies to HP MSA1000 arrays only. If there is no partition on the target, the default boot sector format is 63 sectors. The default setting is disabled.
Auto Scan Setting	This option allows you to enable or disable the auto scan feature. Auto scan is available only when all eight boot entries are unused (empty). The factory default for auto scan is disabled. Auto scan allows you to boot from the first boot device found in the FC fabric with the following options: ■ **Autoscan disabled** – Default ■ **Any first device** – The first adapter scans the fabric and the first disk discovered becomes the attempted boot device ■ **First LUN 0 device** – The first device discovered with LUN 0 ■ **First NOT LUN 0 device** – The first device discovered with a LUN other than 0

F.2.4.3 Configure FCoE Targets Page

The Configure FCoE Targets page lets you configure the boot attributes.

Figure 211 Configure FCoE Targets Page

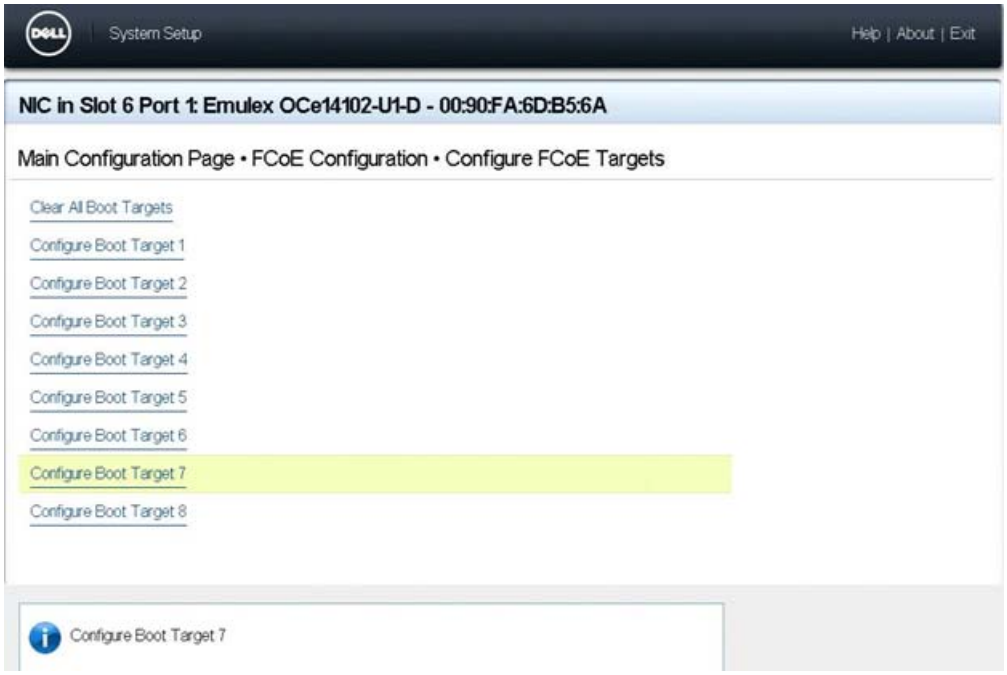


Table 30 Configure FCoE Configuration Page Menu Options

Menu Option	Description and Available Settings
Clear All Boot Targets	This option allows you to erase the boot list contents when you do not want to reset the remaining FCoE parameters
Configure Boot Target	This option allows you to configure boot targets 1 to 8. When you click on any of the Configure Boot Target links an immediate FCoE scan of the SAN occurs and the SAN Discovery Target List page (Figure 212) appears. The FCoE scan occurs only on the first selection of he Configure Boot Target link. On subsequent selections of the Configure Boot Target links the WWPN/LUN Target-List data obtained from the scan is stored internally and re-displayed.

F.2.4.3.1 SAN Discovery Target List Page

The SAN Discovery Target List page lets you select a WWPN/LUN to trigger the dynamic update of the corresponding `FCoEWWPNTarget` and `FCoEBootTargetLUN` xUEFI variable.

Figure 212 SAN Discovery Target List Page

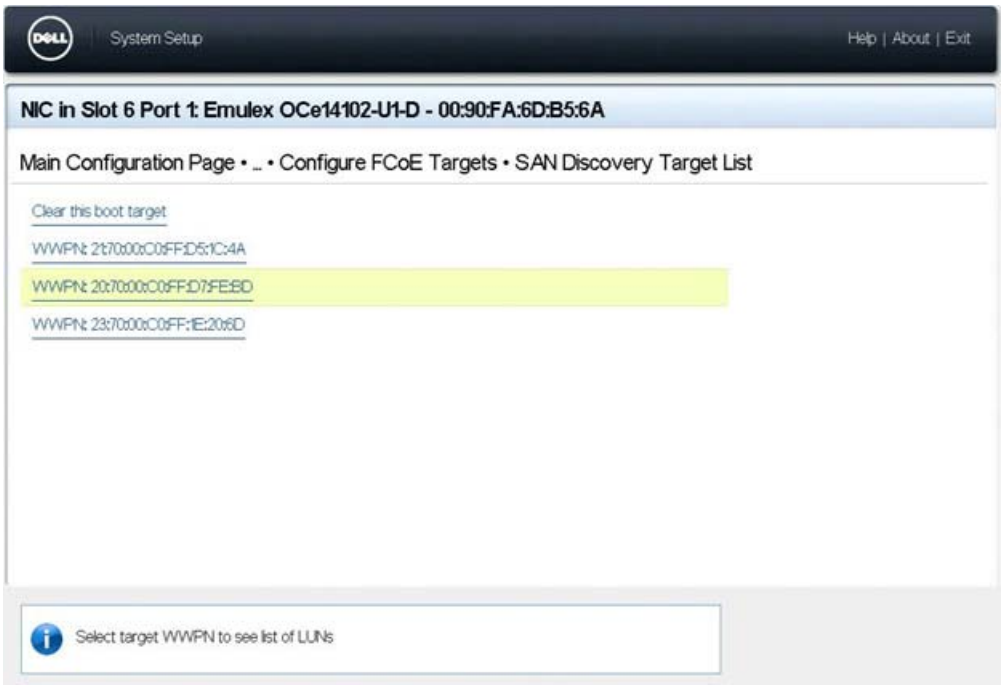
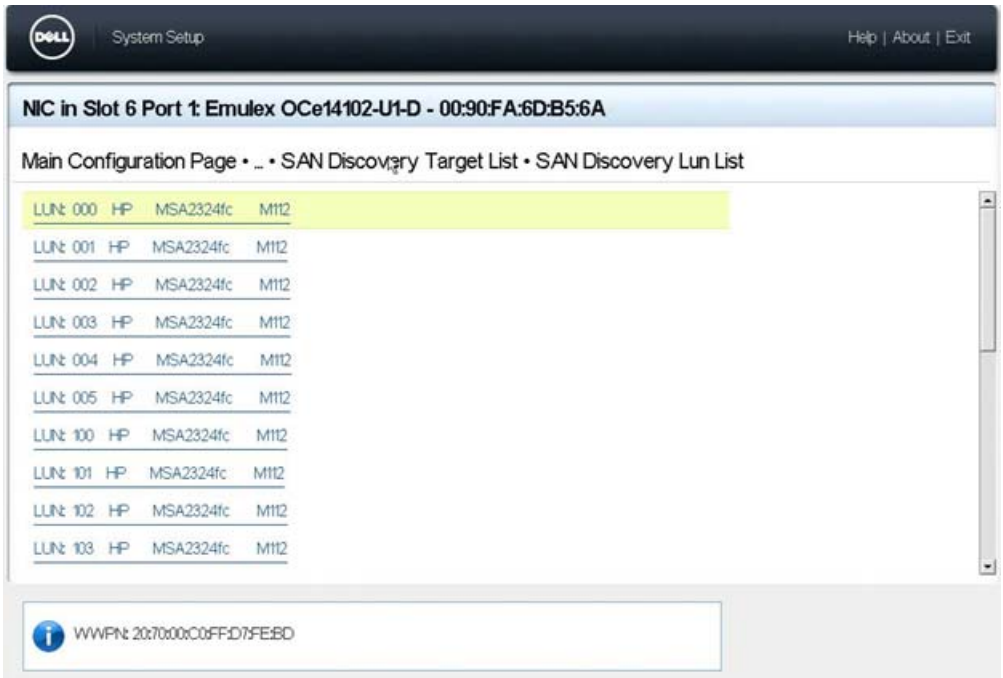


Table 31 SAN Discovery Target List Page Menu Options

Menu Option	Description and Available Settings
Clear this boot target	This option allows you to clear the currently selected boot target. Choose the first item in the list which is ALWAYS a00:00:00:00:00:00:00:00 WWPN and 00 LUN.
WWPN/LUN	This option allows you to select the desired target to be used for FCoE boot configuration for the particular boot target you are configuring. Only one item from the displayed WWPN/LUN may be selected to trigger the dynamic update of the corresponding `FCoEWWPNTarget` and `FCoEBootTarget` LUN xUEFI variable.The FCoE boot BIOS driver target boot-list table location corresponding to the Configure Boot Target previously chosen will also be updated appropriately. When you click on any of the WWPN/LUN the available LUNs for the selected target are displayed in the SAN Discovery LUN List Page (Figure 213).

Figure 213 SAN Discovery LUN List Page



F.2.5 Device Level Configuration Page

The Device Level Configuration page displays the global device level parameters and lets you configure them.

Figure 214 Device Level Configuration Page

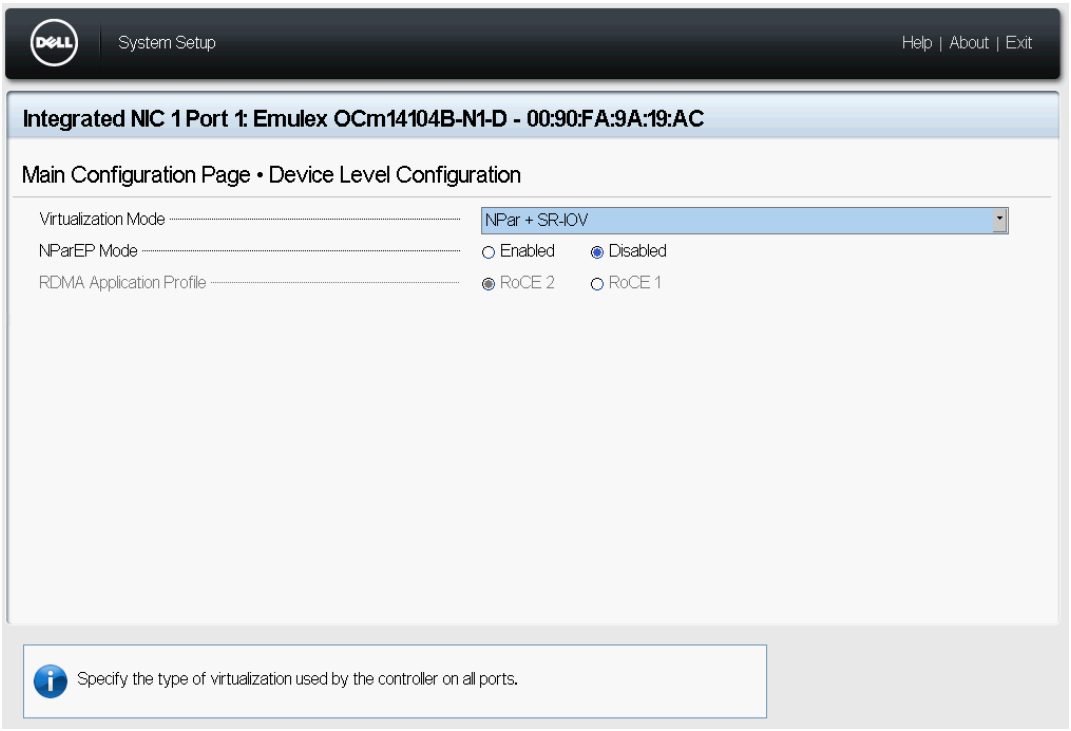


Table 32 Device Level Configuration Page Menu Options

Menu Option	Description and Available Settings
Virtualization Mode	Specify the virtualization mode setting of the adapter. Available settings include: ■ None (default) ■ NPar ■ SR-IOV ■ NPar+SR-IOV **NOTE** In the NPar+SR-IOV mode, SR-IOV is enabled only on the first partition on each port in NPar and NPar-EP modes. The maximum number of VFs supported on each port is 24 for a dual-port P1 adapter, 56 for a dual-port P2 adapter and 28 for a quad-port P1 or P2 adapter. **NOTE** When the Virtualization mode is set to NPAR+SRIOV, the firmware chooses to apply the NIC only profile to the adapter. If NParEP mode support is disabled, the first eight functions of the adapter are by default set to NIC only.
NParEP Mode	Enable more than eight partitions on the device. It must not be enabled if the system and operating system do not support devices with more than eight PCI PFs. **NOTE** This menu option is available only when the Virtualization mode is set to NPar. Available settings include: ■ Enabled ■ Disabled
RDMA Application Profile	The RDMA application profile allows you to select the appropriate configuration for your particular system setup. **NOTE** This menu option is available only when the NIC+RoCE personality is selected. **NOTE** This menu option is available only when the Virtualization mode is set to None. The RDMA application profile choices include: ■ **RoCE-1** – Not supported. ■ **RoCE-2** – The RoCE-2 profile supports Windows SMB Direct, Linux iSER, NFS-RDMA, and VM Migration.

F.2.6 NIC Partitioning Configuration Page

The NIC Partitioning Configuration page (Figure 215) displays NIC partition information and lets you configure their functionality.

Figure 215 NIC Partitioning Configuration Page

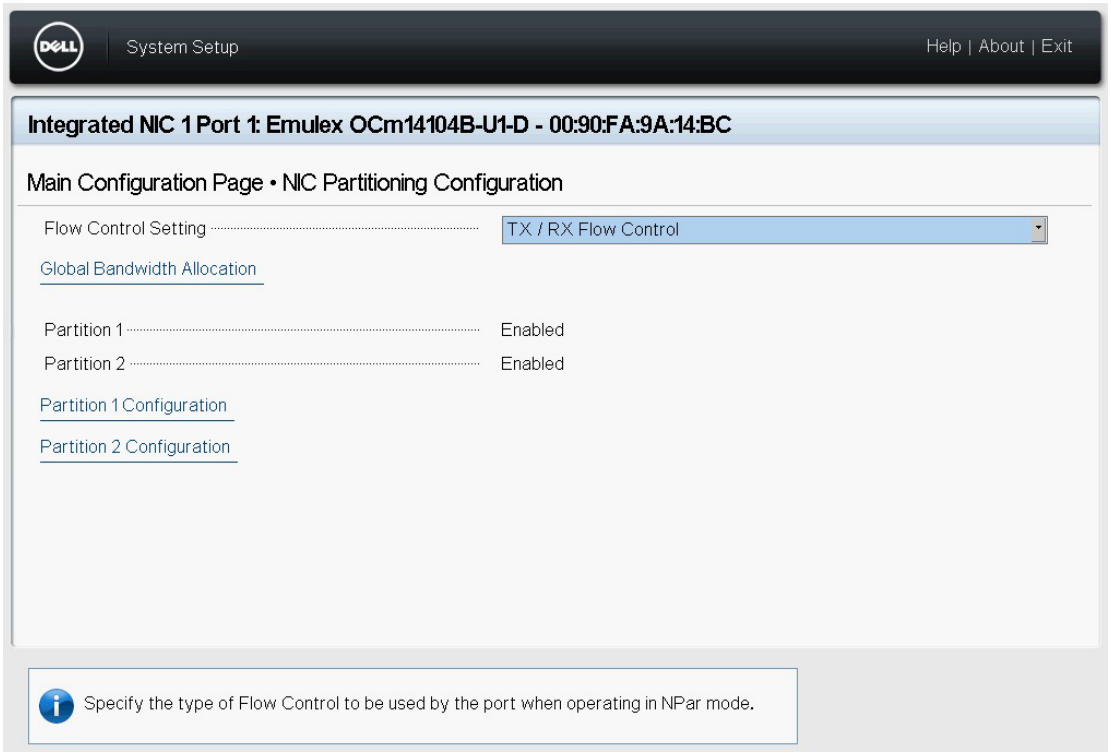


Table 33 NIC Partitioning Configuration Page Menu Options

Menu Option	Description and Available Settings
Flow Control Setting	Configure the type of flow control used by the adapter. **NOTE** Ethernet flow control cannot be changed if priority flow control is enabled on the switch ports and enabled in the firmware (through the Emulex OneCommand CNA Manager application). Available settings include: ■ TX: Send Pause on RX Overflow ■ RX: Throttle TX on Pause Received ■ TX / RX Flow Control
Global Bandwidth Allocation	Assign TX bandwidth allocation across enabled partitions on the port. See Section F.2.6.1, Global Bandwidth Allocation Page, for more information.
Partition n	View the current partition state. One listing for each partition appears with n representing the number of the partition.
Partition n Configuration	View the partition configuration. One listing for each partition appears with n representing the number of the partition. See Section F.2.6.2, Partition Configuration Page, for more information.

F.2.6.1 Global Bandwidth Allocation Page

The Global Bandwidth Allocation page (Figure 216) lets you configure the TX bandwidth allocation across enabled partitions on the port.

Figure 216 Global Bandwidth Allocation Page



Table 34 Global Bandwidth Allocation Page Menu Options

Menu Option	Description and Available Settings
Partition n Minimum TX Bandwidth	The minimum transmit bandwidth value is the least amount of bandwidth that the partition can provide. It is represented as a percentage. One listing for each partition appears with "n" representing the number of the partition. The valid range is 0 to 100 percent for each enabled partition. The minimum transmit bandwidth value must be less than or equal to the maximum transmit bandwidth value. The total of the minimum transmit bandwidth values for all enabled functions on that port must be equal to 100. **NOTE** A minimum transmit bandwidth value of 0 is a valid value. If all of the partitions' minimum transmit bandwidth values are zero, the bandwidth is distributed equally among the current active partitions.
Partition n Maximum TX Bandwidth	The maximum transmit bandwidth value is the greatest amount of bandwidth that the function can provide. It is represented as a percentage. One listing for each partition appears with "n" representing the number of the partition. The valid range is 0 to 100 percent for each enabled partition.

F.2.6.2 Partition Configuration Page

The Partition Configuration page (Figure 217) and (Figure 218) display information for a particular partition and lets you configure some partition parameters. A separate page for each available partition appears.

Only one storage function on each port is allowed. After a storage function is enabled on a partition, the storage protocols appear as unavailable on the other partitions.

Figure 217 Partition 1 Configuration Page

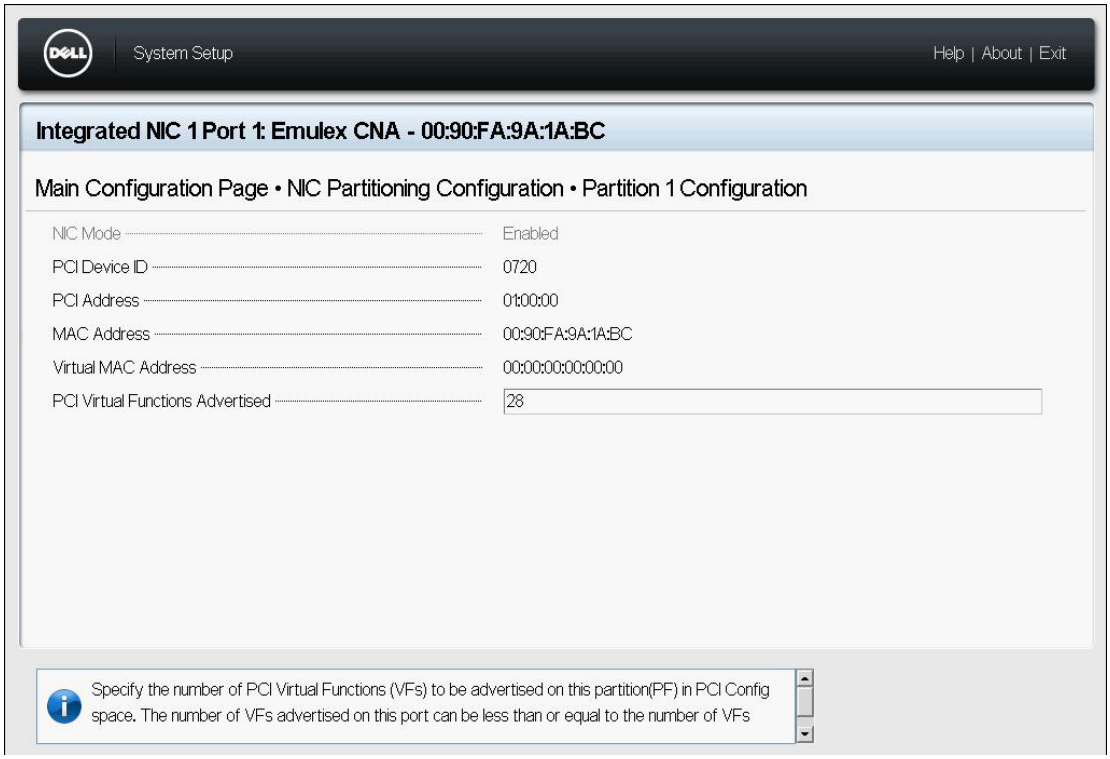


Table 35 Partition 1 Configuration Page Menu Options

Menu Option	Description and Available Settings
NIC Mode	Enable or disable the NIC personality on the partition. Available settings include: ■ Enabled ■ Disabled The default setting is Enabled on all partitions. **NOTE** The NIC mode setting cannot be changed on Partition 1 of each port.
PCI Device ID	The PCI device ID for the partition.
PCI Address	The PCI address for the partition. This value is assigned by the system BIOS.
MAC Address	The permanent MAC address assigned at the factory.
Virtual MAC Address	A programmatically assignable MAC address for the partition.
PCI Virtual Functions Advertised	The number of PCI VFs to be advertised on this partition. The number of VFs advertised on this port must be less than or equal to the number of VFs supported on this port. This value varies based on the adapter in use. For dual port Emulex OCe14000-series adapters, the maximum number of supported VFs per port is 24. For dual port Emulex OCe14000B-series adapters, the maximum number of supported VFs per port is 56. For quad-port adapters, the maximum number of supported VFs per port is 28.

Figure 218 Partition 2 Configuration Page

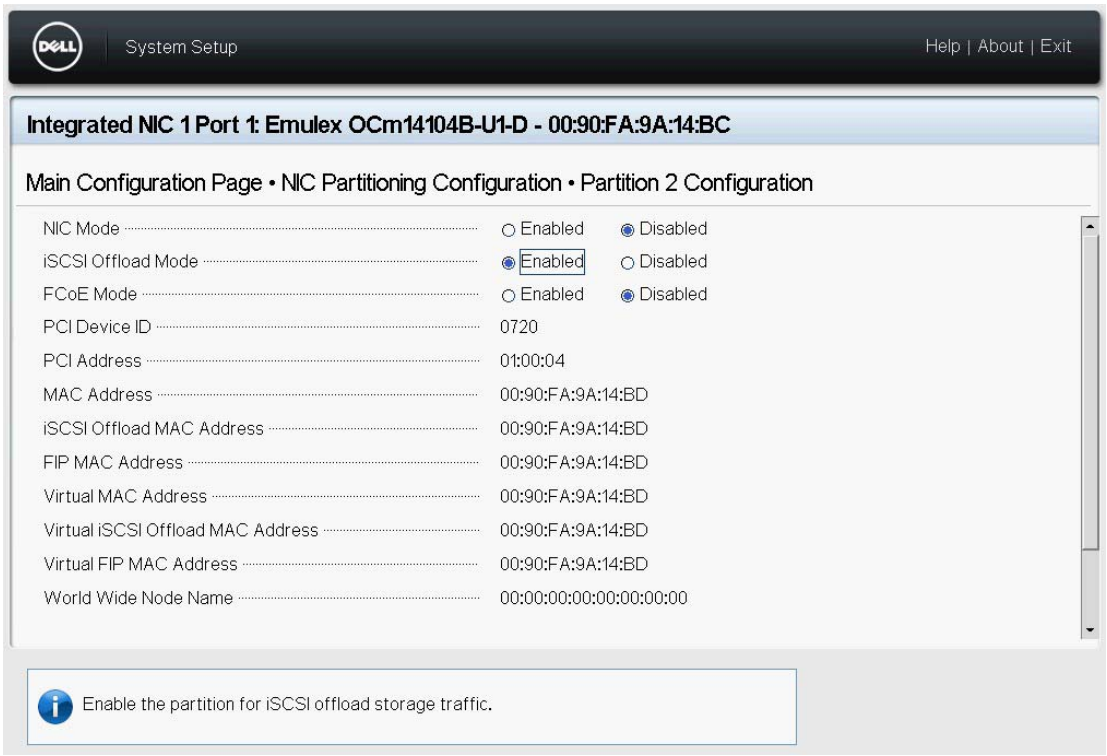


Table 36 Partition 2 Configuration Page Menu Options

Menu Option	Description and Available Settings
NIC Mode	Enable or disable the NIC personality on the partition. Available settings include: ■ Enabled ■ Disabled The default setting is Enabled on all partitions. **NOTE** The NIC mode setting cannot be changed on Partition 1 of each port.
iSCSI Offload Mode	Enable or disable the iSCSI offload personality on the partition. Available settings include: ■ Enabled ■ Disabled (default)
FCoE Mode	Enable or disable the FCoE personality on the partition. Available settings include: ■ Enabled ■ Disabled (default)
PCI Device ID	The PCI device ID for the partition.
PCI Address	The PCI address for the partition. This value is assigned by the system BIOS.
MAC Address	The permanent MAC address assigned at the factory.
iSCSI Offload MAC Address	The permanent MAC address for iSCSI offload assigned at the factory.
FIP MAC Address	The permanent FIP-MAC address for FCoE assigned at the factory.
Virtual MAC Address	A programmatically assignable MAC address for the partition.

Table 36 Partition 2 Configuration Page Menu Options (Continued)

Menu Option	Description and Available Settings
Virtual iSCSI Offload MAC Address	Configure the MAC address for iSCSI offload. The setting can have a maximum of 17 characters with a colon (:) separator between every two characters.
Virtual FIP MAC Address	A programmatically assignable FIP-MAC address for the partition FCoE.
World Wide Node Name	The FC World Wide Node Name ID for FCoE. A maximum of 23 characters is supported. A colon separator (:) must be used between every two characters.
Virtual World Wide Node Name	Specify the FC World Wide Node Name ID for FCoE. A maximum of 23 characters is supported. A colon separator (:) must be used between every two characters.
World Wide Port Name	The FC World Wide Port Name ID for FCoE. A maximum of 23 characters is supported. A colon separator (:) must be used between every two characters.
Virtual World Wide Port Name	Specify the FC World Wide Port Name ID for FCoE. A maximum of 23 characters is supported. A colon separator (:) must be used between every two characters.

F.3 Secure Firmware Download

A single firmware image contains the latest version of the firmware and boot code.

The firmware and UEFI NIC, iSCSI, and FCoE drivers are contained in one image file with the .UFI extension.

If a jumper is plugged into the ASIC GPIO pin on the adapter, it indicates that only secure firmware versions are allowed to be flashed on the adapter. Upgrade from earlier unsecure firmware versions does not involve a digital signature check since the earlier code does not support this feature. Therefore, upgrade to secure firmware versions will be successful. If a secure version of the firmware (version 11.0 or later) is installed on an Emulex OCe14000B-series adapter and you want to update to an earlier unsecure version of firmware, you must remove the secure firmware jumper block before performing the update. See the adapter installation guide for more information. Firmware upgrade from a secure firmware version to an unsecure firmware version results in an error when the jumper is in place. If this happens, an error message indicating digital signature failure appears.

Appendix G: Configuring PXE Boot for NIC on Dell OCe14000-Based Systems

This section describes how to use the PXESelect utility to set up a PXE bootable network by configuring a Dell OCe14000-based adapter.

NOTE For Dell adapter models D6T93, F6PCP, JJPC0, and P3V42, PXE boot can be performed only on the first NIC partition on each port in the legacy x86 BIOS.

For more information on the PXE protocol, the PXE Boot process, and remotely installing with PXE on different operating systems, see Chapter 3, Configuring PXE Boot for NIC on OCe14000-Series Adapters.

G.1 Running the PXESelect Utility

To run the PXESelect utility, start or restart the computer. When prompted, press Ctrl-P keys. If you are running the PXESelect utility with multiple adapters, all of the adapters are displayed when you start the utility. For example:

```
Press <Ctrl><P> for PXESelect(TM)Utility

OCe14102B-UM NIC#0 Port#1 Base 0x00000380001A0000 @ Bus:1B Dev:00 Fun:00
OCe14102B-UM NIC#0 Port#2 Base 0x00000380001E0000 @ Bus:1B Dev:00 Fun:00
- Initializing ...Done.
```

The PXE Configuration menu appears after the boot BIOS initializes to begin the PXE configuration.

NOTE A UEFI-capable system typically does not show the prompt for running the PXESelect utility unless it is configured for legacy booting. See your system configuration manual for information on performing a legacy boot. All configuration that can be performed in the PXESelect utility can instead be performed in the UEFI configuration utility. See Appendix F, Dell UEFI for more information.

G.2 Navigating the PXESelect Utility

Use the following methods to navigate the PXESelect utility:

- Press the up or down arrows on the keyboard to move through and select menu options or configuration fields. If multiple adapters are listed, use the up and down arrows to scroll to the additional adapters.
- Press the left or right arrows on the keyboard to move through and select setting options.
- Press Enter to accept a selection, select a menu option, or to change a configuration default.
- Press the F7 key to save the changes on the screen.
- Press Esc to return to the previous menu or page, cancel a selection or dialog box, or exit the utility.

G.3 Setting Up a PXE Bootable Network

After the PXE boot BIOS is initialized, you can use the PXESelect utility to set up a PXE bootable network by configuring the adapters.

To configure adapters for PXE boot:

1. From the following Controller List (Figure 219), use the up or down arrow keys to select the adapter you want to configure and press Enter.

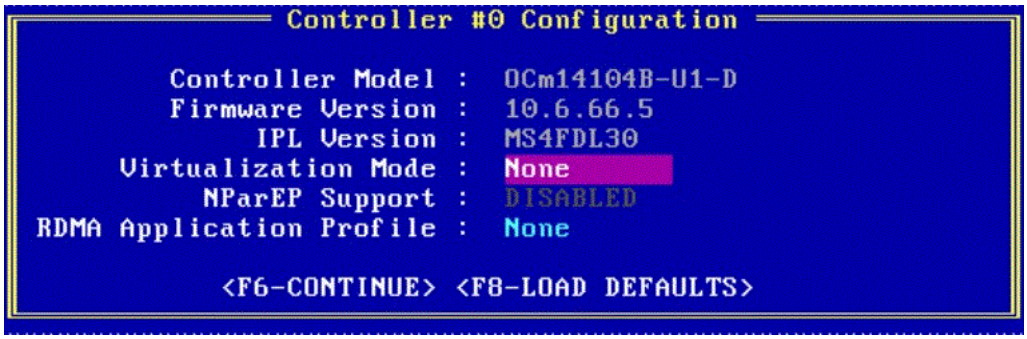
Figure 219 Controller List



NOTE The system displays the Controller List only if there are two or more adapters connected.

The following Controller Configuration dialog appears.

Figure 220 Controller Configuration Dialog



Depending on the adapter, one or more of the following items are displayed:

- **Controller Model** – Displays the adapter's model number.
- **Firmware Version** – Displays the current firmware version.
- **IPL Version** – Displays the IPL version.
- Virtualization Mode – Select one of the following options:
  - **NPar** – If NPar is enabled, you can access up to eight virtual network interfaces (all partitions are set to NIC by default). For more information on NPar support, see Section G.3.1, Configuring NPar.
  - **SRIOV** – If the system BIOS supports SR-IOV, you can enable it. For more information on SR-IOV configuration, refer to the appropriate Emulex driver's user guide.
  - NPar+SRIOV – SRIOV can be enabled over NPar supported adapters.
  - **None** – If there is no virtualization mode selected, there is only one NIC function on each port and there are no storage functions.

  NOTE If the virtualization mode has changed since the last boot, a reboot will occur when you exit the PXESelect utility.
- **NParEP Support** – If NParEP support is enabled, each port can be partitioned into isolated PFs (channels) up to total of 16 functions. See Section G.3.2.2, Enabling NParEP Support, for more information.

NOTE NParEP support is only available on Dell 13G or newer systems.

— **RDMA Application Profile** – This setting allows you to select the appropriate configuration for the particular system setup.

NOTE The RDMA application profile setting is available only when the NIC+RoCE personality is selected.

NOTE The RDMA application profile setting is available only when the Virtualization mode is set to None.

The RDMA application profile choices include:

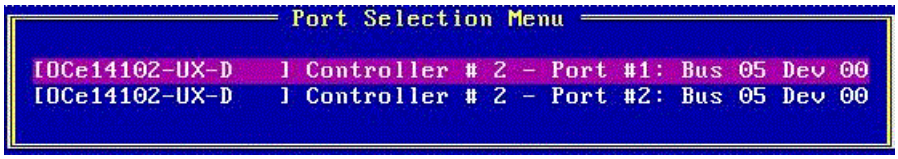
- RoCE-1 – Not supported.
- RoCE-2 – The RoCE-2 profile supports Windows SMB Direct, Linux iSER, NFS-RDMA, and VM Migration.

2. After you set the options in the Controller Configuration dialog (Figure 220), press the F7 key to save the current settings.
3. Press the F6 key to proceed.

NOTE If you want to erase the current configuration, press the F8 key. See Section G.3.3.10, Erasing Ports and Adapter Configurations, for more information.

The following Port Selection Menu appears.

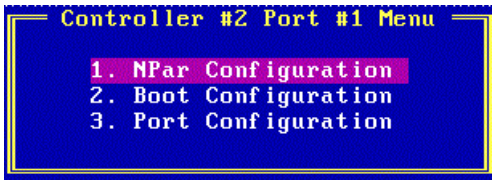
Figure 221 Port Selection Menu



NOTE The Port Selection Menu appears only if there are two or more ports connected.

4. Select the port you want to configure and press Enter. The following Port Menu appears.

Figure 222 Port Menu



5. From the Port Menu, you can select one of the following options:
   — **NPar Configuration** – See Section G.3.1, Configuring NPar, for more information.
   — **Boot Configuration** – See Section G.3.3, Configuring Boot Options, for more information.
   — **Port Configuration** – See Section 7.3.1, Configuring Port Options, for more information.

   Select the desired configuration option and press Enter.

G.3.1 Configuring NPar

NPar support provides the ability to configure multiple physical functions or I/O partitions for each physical adapter port.

NOTE

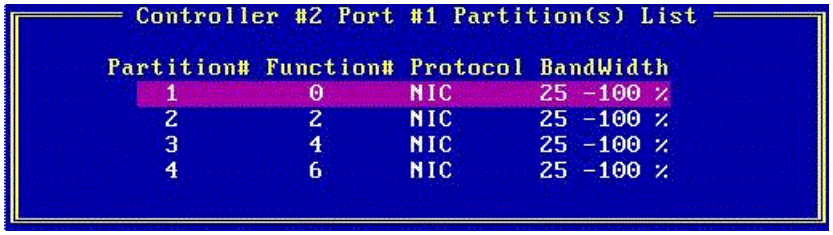
The settings in this menu can be configured only if the virtualization mode is set to NPar. Otherwise, the settings in this menu are for viewing purposes only.

If NPar support is enabled, you must configure the NPar settings (minimum and maximum bandwidths) for iSCSI and FCoE storage functions in the NIC BIOS before they can be configured further from their respective utilities (iSCSI or FCoE BIOS). Otherwise, the Logical Link for that function will still show as down, and you will not be able to log in to targets or find LUNs behind those targets.

To configure NPar:

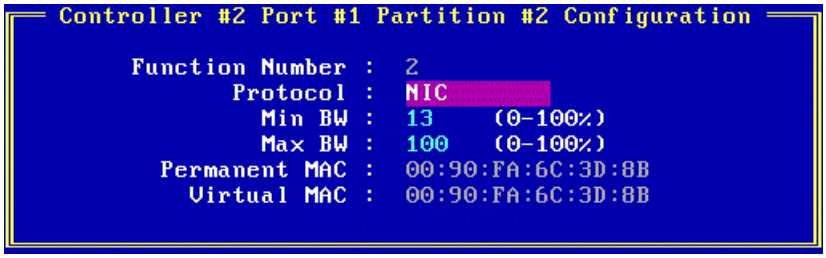
1. From Port Menu (Figure 222), select NPar Configuration and press Enter. The following Partition List appears.

Figure 223 Partition(s) List



2. Select the partition you want to configure and press Enter. The following Configure Partition dialog appears.

Figure 224 Configure Partition Dialog



3. From the Configure Partition dialog, you can perform the following tasks:
   - View the function number.
   - Select the protocol for each function, except for PF #0 and PF #1.
   - Configure the minimum and maximum bandwidth for each channel
   - View the permanent and virtual MAC addresses.

G.3.2 Selecting the Protocol

This option specifies a list of available protocols that can be configured on an adapter. Depending on the protocols for which the adapter is licensed, one or more of the following selections is available:

- **NIC** – Implies that all the enabled functions provide NIC functionality
- **iSCSI or FCoE** – These protocols are enabled on one function per adapter port and include NIC functionality on the other enabled functions

  NOTE Only one of each storage protocol is allowed on each port. Two iSCSI functions and two FCoE functions are not allowed on a single port.
- **RoCE** – The RoCE personality implies that all the enabled functions provide RoCE functionality.

  NOTE The RoCE personality setting is available only when the Virtualization mode is set to None and the RDMA application profile is set to RoCE-2.
- **None** – Allows you to disable that particular function (0 to 7).

The menu displays the available personalities only, including both free and licensed personalities.

To select the protocol:

1. From the Configure Partition dialog (Figure 224), use the left or right arrow keys to select the desired protocol.
   - NIC
   - iSCSI
   - FCoE
   - RoCE
   - None
2. Press the F7 key to save.

   NOTE If the protocol selection has changed since the last boot, a reboot will occur when you exit the PXESelect utility.

G.3.2.1 Configuring Minimum and Maximum Bandwidth

To configure bandwidth:

1. On the Configure Partition dialog (Figure 224), use the use the up or down arrow keys to select Min BW or Max BW.
2. Use the left or right arrow keys to select the desired bandwidth value.
   - The Minimum Bandwidth value is the least amount of bandwidth the function can provide. It is represented as a percentage. The Minimum Bandwidth value must be less than or equal to the Maximum Bandwidth value. The total of the Minimum Bandwidth values for all enabled functions on that port must be equal to 100.

     NOTE A Minimum Bandwidth value of 0 is a valid value. If all of the partitions' Minimum Bandwidth values are zero, the bandwidth is distributed equally among the current active partitions.
   - The Maximum Bandwidth is the greatest amount of bandwidth the function can provide. It is represented as a percentage.
3. Press the F7 key to save.

G.3.2.2 Enabling NParEP Support

NOTE NParEP support is available only on Dell 13G or newer systems.

NParEP support provides the ability to configure additional PCI functions or I/O partitions for each physical adapter port if NPar is enabled.

You can configure up to 16 functions on a one-port Emulex OCe14400-series adapter, up to 8 functions per port on a one or two-port Emulex OCe14100-series adapter, and up to 4 functions per port on a four-port Emulex OCe14100-series adapter.

The maximum number of functions allowed on an adapter is controlled by the adapter's IPL file and the system's support for NParEP.

The following requirements must be met to support more than eight functions on an adapter.

- The Dell system hardware (the motherboard and BIOS) must support NParEP.
- The host operating system must support NParEP:
  - Windows Server 2012 and newer versions
  - RHEL 5.9 and newer versions
  - RHEL 6.4 and newer versions
  - SLES 11 SP2 and newer versions
  - ESXi 5.0 and newer versions
- The application management tools, including the Emulex OneCommand CNA Manager application, must support NParEP.
- NParEP must be enabled in the firmware using the Emulex OneCommand CNA Manager application.

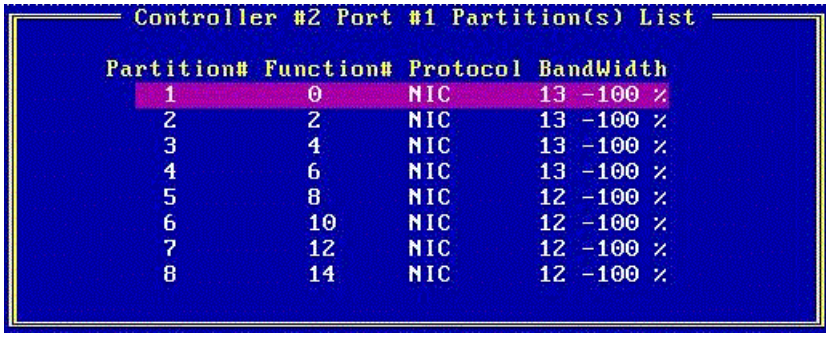
If these conditions are not met, you may be able to configure more than eight functions, but only up to eight functions will be running and discovered after a reboot.

To enable NParEP:

1. From Controller Configuration dialog (Figure 220), select NParEP Support.
2. Select **Enabled** and press the F7 key to save.

If NParEP is enabled, the system displays additional functions on the following Partition List.

Figure 225 Partition(s) List (NParEP Enabled)



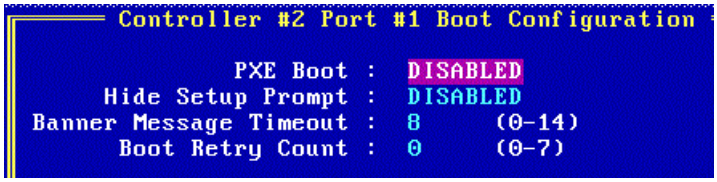
See Section G.3.1, Configuring NPar, for more information on configuring the additional functions.

G.3.3 Configuring Boot Options

To view the Boot Configuration dialog:

1. From the Port Menu (Figure 222), select Boot Configuration and press Enter. The following Boot Configuration dialog appears.

Figure 226 Boot Configuration Dialog



From the Boot Configuration dialog, you can perform the following tasks:

- Configure PXE boot support – See Section G.3.3.1, Configuring PXE Boot Support, for instructions.
- Configure the Hide Setup prompt – See Section G.3.3.2, Configuring the Hide Setup Prompt, for instructions.
- Configure the banner message timeout – See Section G.3.3.3, Configuring the Banner Message Timeout, for instructions.
- Configure the boot retry count – See Section G.3.3.4, Configuring the Boot Retry Count, for instructions.

G.3.3.1 Configuring PXE Boot Support

To configure PXE boot support:

1. On the Boot Configuration dialog (Figure 226), use the left or right arrow keys to enable or disable the PXE Boot setting.
2. Press the F7 key to save.

NOTE If PXE boot is enabled, during system startup PXE contacts the DHCP server for an IP address to boot from the network.

G.3.3.2 Configuring the Hide Setup Prompt

This setting enables or disables the CTRL<P> setup prompt (banner) during POST.

To configure the Hide Setup prompt:

1. On the Boot Configuration dialog (Figure 226), use the left or right arrow keys to enable or disable the Hide Setup Prompt setting.
2. Press the F7 key to save.

G.3.3.3 Configuring the Banner Message Timeout

This setting controls the number of seconds that the PXE banner message appears during POST. The valid range is 0 to 14.

To configure the banner message timeout:

1. On the Boot Configuration dialog (Figure 226), use the left or right arrow keys to select the desired amount of time (in seconds) for the Banner Message Timeout setting.
2. Press the F7 key to save.

G.3.3.4 Configuring the Boot Retry Count

This setting specifies the number of boot retry attempts before control is returned to the system BIOS. Valid values include:

- 0 (no retry)
- 1 to 6
- 7 (indefinite retries)

To configure the boot retry count:

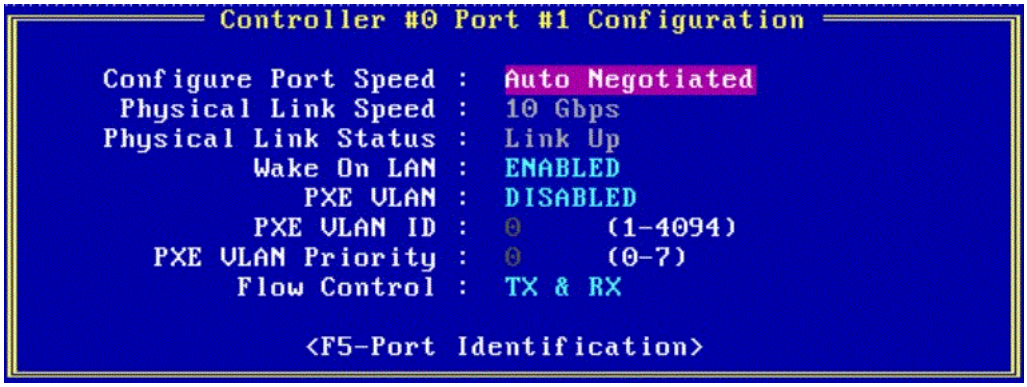
1. On the Boot Configuration dialog (Figure 226), use the left or right arrow keys to select the desired number of retries for the Boot Retry Count setting.
2. Press the F7 key to save.

14.3.1 Configuring Port Options

To view the Port Configuration dialog:

1. From the Port Menu dialog (Figure 222), select Port Configuration and press Enter. The following Port Configuration dialog appears.

Figure 227 Port Configuration Dialog



2. From the Port Configuration dialog, you can perform the following tasks:
   - Set the configured port speed – See Section G.3.3.5, Setting the Configured Port Speed, for instructions.
   - View the physical link speed and link status.
   - Configure WoL – See Section G.3.3.6, Configuring Wake on LAN, for instructions.
   - Configure PXE VLAN, including the PXE VLAN ID and priority – see Section G.3.3.7, Configuring the PXE VLAN ID and Priority, for instructions.
   - Configure flow control – See Section G.3.3.8, Configuring Flow Control, for instructions.

     NOTE Flow Control is only available when the Virtualization mode is set to NPar.
   - Identify the port – See Section G.3.3.9, Physically Identifying the Port, for instructions.

G.3.3.5 Setting the Configured Port Speed

To set the port speed:

1. On the Port Configuration dialog (Figure 227), use the left or right arrow keys to select Auto Negotiated, 10G, or 1G for the Configured Port Speed setting.
2. Press the F7 key to save.

G.3.3.6 Configuring Wake on LAN

Wake on LAN configuration is dependent on the system BIOS and operating system settings. The following table describes the different configurations.

Table 37 Wake on LAN Settings and Results

BIOS WoL	Operating System WoL	Results
Enabled	Enabled	Since the BIOS and operating system have the same setting, after the operating system reboots, WoL is enabled.
Enabled	Disabled	After the operating system reboots, the driver disables WoL.
Enabled	Enabled, but changed to Disabled after reboot	After the operating system reboots, WoL is set to Disabled, so WoL is disabled.
Disabled	Enabled	After the operating system reboots, the driver does not change WoL, so WoL is disabled.
Disabled	Disabled	Since the BIOS and operating system have the same setting, after the operating system reboots, WoL is disabled.
NOTE For situations 2 to 5, in order to re-enable WoL, you must enable WoL in the operating system, reboot the system, and then enable WoL in the system BIOS.		

To configure Wake on LAN:

1. From the Port Configuration dialog (Figure 227), use the left or right arrow keys to enable or disable the Wake on LAN setting.
2. Press the F7 key to save.

G.3.3.7 Configuring the PXE VLAN ID and Priority

To configure a PXE VLAN ID and set the priority level:

1. On the Port Configuration dialog (Figure 227), use the left or right arrow keys to enable the PXE LAN setting.
2. For the PXE VLAN ID, use the left or right arrow keys to enter a number from 1 to 4094.
3. For the PXE VLAN Priority level, use the left or right arrow keys to enter a number from 0 to 7.

   This unique value assigns a priority to outbound packets containing a specified VLAN ID. Valid values range from 0 to 7, with 0 the highest priority level.
4. Press the F7 key to save.

After you exit the PXESelect utility, the system will reboot for the configuration to take effect.

G.3.3.8 Configuring Flow Control

This setting allows you to configure the type of flow control used by the adapter.

NOTE Flow Control is only available when the Virtualization mode is set to NPar.

Ethernet flow control cannot be changed if priority flow control is enabled on the switch ports and enabled in the firmware (through the Emulex OneCommand CNA Manager application).

To configure flow control:

1. From the Port Configuration dialog (Figure 227), use the left or right arrow keys to select one of the following settings for the Flow Control setting.
   - TX & RX
   - TX (Send Pause on RX Overflow)

— RX (Throttle TX on Pause Received)
— Disabled

2. Press the F7 key to save.

G.3.3.9 Physically Identifying the Port

To physically determine which port you are configuring by causing the link and activity LEDs of that port to blink:

1. On the Port Configuration dialog (Figure 227), press the F5 key.
2. The LEDs on the adapter begin blinking. The selected port LED status indicators blink on the adapter for approximately 15 seconds.

NOTE Not all adapters have LEDs that are visible externally. If you are using an add-in card in a blade server environment, the port identification or beaconing capability does not work.

G.3.3.10 Erasing Ports and Adapter Configurations

NOTE If this setting is selected, all previous configuration settings are returned to their factory default settings including the current protocol selection. You must perform this action to provide a clean environment for new configuration settings to take effect.

To erase the ports and adapter configuration:

1. On the Controller Configuration dialog (Figure 220), press the F8 key to restore the factory default settings. A warning appears asking if you want to erase the current configuration for all ports of the adapter.
2. Press Y to delete the configuration. You will receive another warning asking you to confirm the permanent removal of the configuration.
3. Press Y to delete the configuration.

To exit the PXESelect utility after erasing the ports and adapter configuration:

1. Follow the instructions on the bottom of the individual menus until you are prompted to exit.
2. Press Y to exit. Depending on what settings were changed, a reboot may be necessary.

NOTE For older systems, depending on the memory allocation method supported, the PXESelect utility automatically reboots even if no changes were made to the system.

G.4 PXE Boot Parameters Default Values

The default settings for the PXE Boot parameters are listed in the following table.

Table 38 PXE Boot Parameter Default Values

Parameter	Default Value	Valid Values
Hide Setup Prompt	Disabled	Enabled Disabled
Banner Message Timeout	8	0–14
Port Speed	Auto Negotiated	Auto Negotiated 10G 1G

Table 38 PXE Boot Parameter Default Values (Continued)

Parameter	Default Value	Valid Values
PXE Boot	Disabled **NOTE** On some Dell LOMs, this setting is enabled by default on port 1.	Enabled Disabled
Boot Retry Count	0	0 (no retry) 1–7 Retries
Wake on LAN	Enabled	Enabled Disabled
Flow Control	TX & RX	TX & RX TX RX Disabled
PXE VLAN	Disabled	Enabled Disabled
PXE VLAN ID	0	1–4094
PXE VLAN Priority	0	0–7
Virtualization Mode	None	NPar SRIOV None
NParEP Support	Disabled	Enabled Disabled
Protocol	NIC	NIC FCoE iSCSI RoCE None
Min BW	25%	The total bandwidth is distributed evenly across all the enabled partitions.
Max BW	100%	From the minimum bandwidth value up to 100% (inclusive).


BROADCOM

Emulex® OneConnect® Adapter

Installation Guide

Version 11.2
December 30, 2016

Corporate Headquarters
San Jose, CA

Website
www.broadcom.com

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

Copyright © 2003–2016 by Broadcom. All Rights Reserved.

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

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.

Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Chapter 1: Introduction

This manual describes how to install an Emulex® OneConnect ®adapter.

NOTE Illustrations in this manual are only examples. The actual hardware may vary.

Figure 1 Two-Port Adapter

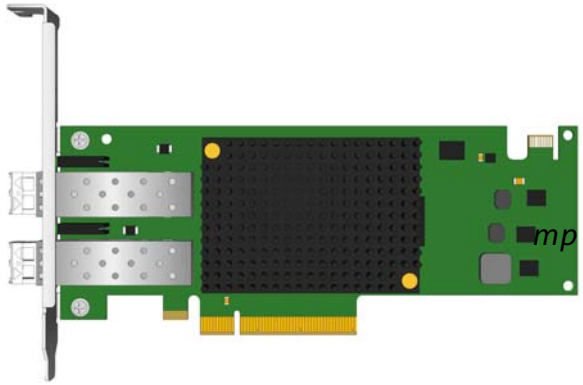
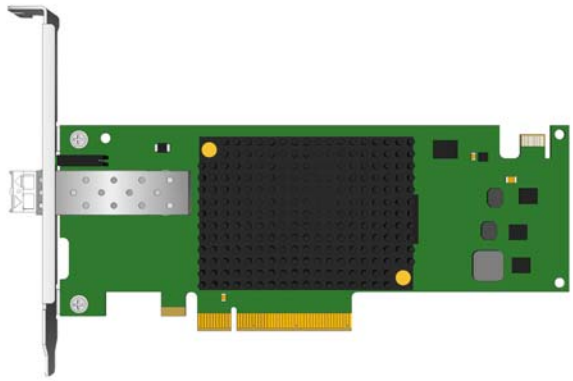


Figure 2 Single-Port Adapter



1.1 Adapter Identification

Each adapter ships with several numbers clearly marked on the board. Record these numbers before installation for later reference.

- **Institute of Electrical and Electronics Engineers, Inc. (IEEE) address** – a unique 64-bit identifier used for system configuration.
- **World Wide Name (WWN)** – derived from the IEEE address; the Fibre Channel (FC) industry uses the WWN for FC connectivity.
- **Serial number** – assigned by Broadcom® and used when communicating with Broadcom.

NOTE If the adapter has two ports, it has two IEEE addresses and two WWNs, one for each port.

Chapter 2: Installation

The Emulex OneConnect adapter can be connected to fiber optic cables (used with embedded optical transceivers) or direct-attach copper (DAC) cables. Additional information for these cable connections and other installation procedures are described in the following sections.

If necessary, before installing the adapter you can change the bracket to a low-profile version, as explained in Section 2.1, Changing the Bracket.

If you do not need to change the bracket, proceed to Section 2.2, Installing the Adapter.

2.1 Changing the Bracket

A standard, full-height Peripheral Component Interconnect Express (PCIe) bracket is factory-installed; however, a low-profile bracket is included in the box with any adapter that supports the low-profile bracket change. The low-profile mounting bracket is shorter than the standard bracket; approximately 3.11 in. (7.9 cm) compared to 4.75 in. (12.06 cm) long.

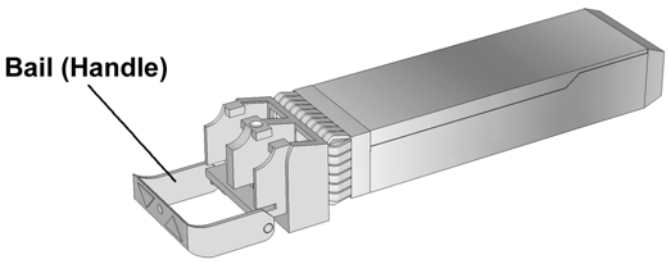
To change the adapter from a full-height to a low-profile bracket:

1. Some adapters come with optical transceivers embedded in their cage assemblies. These optical transceivers must be removed before the bracket can be removed. If the adapter contains optical transceivers, continue with the following steps; otherwise, proceed to step 4.

   CAUTION This is a delicate operation – take care not to damage the optical transceiver.

   The adapter uses different types of optical transceivers. Figure 3 shows an example of one type with the bail (handle) extended.

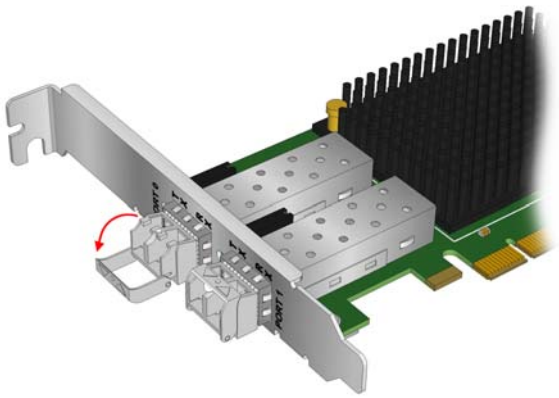
Figure 3 Typical Optical Transceiver



2. To remove a transceiver, pull the bail (handle) out and down to release the latch and gently pull the transceiver out. Do not force it. After the latch is released, the transceiver slides out easily.

   Figure 4 shows a transceiver with the latch released (bail extended) and another transceiver latched in place.

Figure 4 Releasing the Latch on an Optical Transceiver



3. Using electrostatic discharge (ESD) precautions, store the transceiver in an ESD-safe place.
4. Remove the mounting bracket screws from the top of the adapter. Figure 5 shows the screws that are removed from the bracket.

Figure 5 Removing the Bracket



5. Remove the bracket and store it for future use.

6. Align the new mounting bracket tabs with the holes in the adapter.

   NOTE Be careful not to push the bracket past the transceiver or the RJ-45 housing's grounding tabs. Ensure the light emitting diodes (LEDs) are properly aligned with the holes in the bracket.

7. Reinstall the screws that attach the adapter to the bracket.
8. Reinstall the transceiver by sliding it into the housing. When the latch engages, it clicks.
9. Push the bail back into place.

2.2 Installing the Adapter

To install the adapter:

1. Turn off and unplug the computer.
2. Remove the computer case.
3. For a stand-up adapter, remove the blank panel from an empty PCIe bus slot. See Section 2.1, Changing the Bracket to change the brackets if needed.

   NOTE To prevent damage to the adapter, select a slot that will not cause the adapter to interfere with the case.

4. Insert the adapter into the selected slot. Press firmly until the adapter is seated.
5. Secure the adapter's mounting bracket to the case with a panel screw or clip.
6. Replace the computer case and tighten the case screws.

The adapter is now installed in the server and is ready for device attachment.

2.3 Attaching Devices

Two types of cable can be connected to the adapter:

- Fibre optic cable
- DAC cable

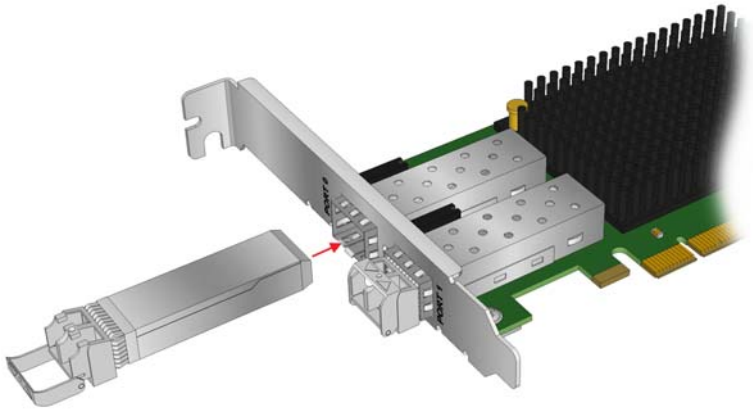
An adapter does not allow normal data transmission on either optical or copper links, unless it is connected to a compatible optical or copper interface connection.

NOTE Use short range (SR), direct-connecting receptacle (DCR) multimode fiber, or 1 to 5 meter small form-factor pluggable plus (SFP+) 10GbE copper cable for interface connections to a network.

To attach devices to the adapter:

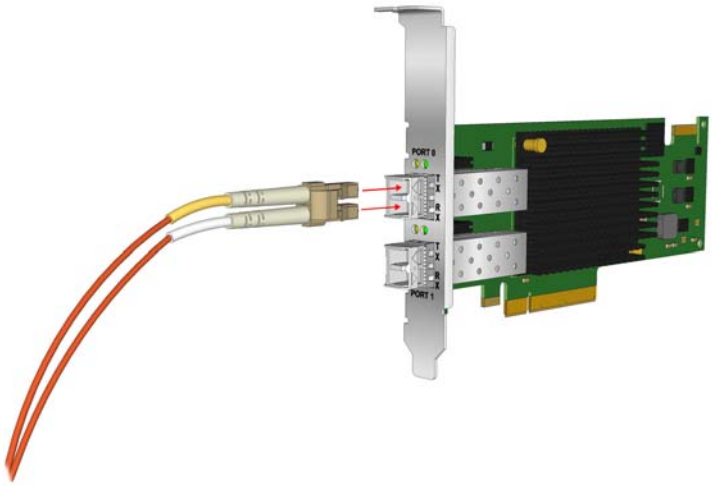
1. Connect the appropriate cable to the adapter.
   When connecting an optical cable, ensure the cages have optical transceivers installed in them (see Figure 6).

Figure 6 Installing an Optical Transceiver



After the optical transceivers are installed, insert the optical cable into the LC connectors on the adapter (see Figure 7).

Figure 7 Connecting a Fiber Optic Cable



When connecting a DAC cable, make sure that the SFP+ cages do not have optical transceivers installed in them. To remove optical transceivers, see step 2 in Section 2.1, Changing the Bracket. After confirming there are no optical transceivers installed, insert the DAC transceiver into the SFP+ cage on the adapter (see Figure 8).

Figure 8 Connecting a DAC Cable



2. After the appropriate cable is connected to the adapter, connect the other end of the cable to a suitable device, such as a Top of Rack (TOR) EE switch.

After the device is connected to the adapter, you are ready to apply power to the computer.

2.4 Applying Power

To apply power:

1. Verify that the adapter is securely installed in the computer.
2. Verify that the correct device is attached.
3. Plug in and turn on the computer.
4. Observe the boot banner for power-on self-test (POST) results.

2.5 Viewing the LEDs

You can view green and yellow LEDs through openings in the adapter's mounting bracket. The green LED indicates firmware operation and the yellow LED indicates port activity or link speed. Each port has a corresponding set of green and yellow LEDs.

Figure 9 LED Indicators



Amber LED (link state):

- **On** (constantly) – Link up
- **Off** – Link down

Green LED (Ethernet activity):

- **Blink** – Activity on the Ethernet link
- **Off** – No activity on the Ethernet link


BROADCOM