IBM BladeCenter Virtual Fabric Solutions

Highlights IBM Virtual Fabric Mode and Switch Independent Mode

Defines the concepts and benefits of virtual NICs (vNICs)

Outlines steps to create various vNIC solutions

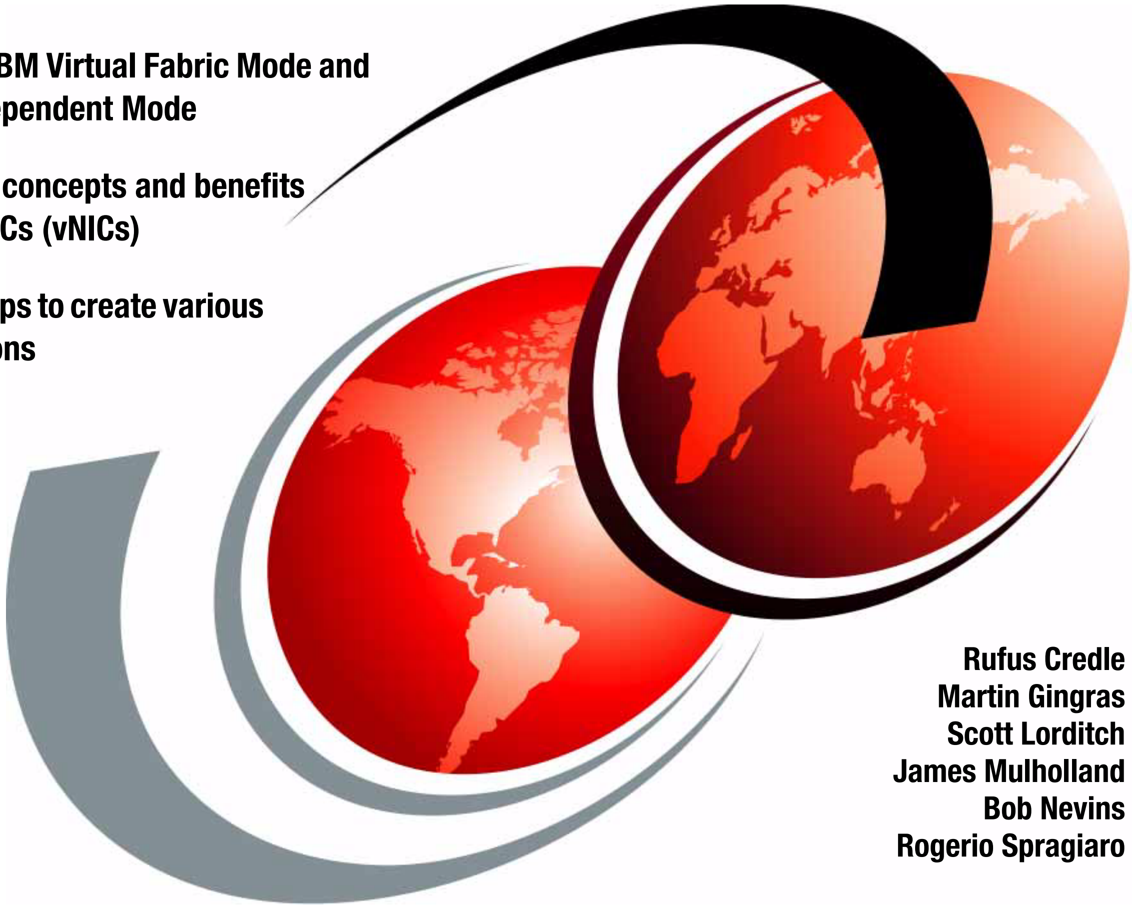

Rufus Credle
Martin Gingras
Scott Lorditch
James Mulholland
Bob Nevins
Rogerio Spragiaro

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM BladeCenter Virtual Fabric Solutions

August 2012

SG24-7966-01

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

Second Edition (August 2012)

This edition applies to the IBM BladeCenter I/O options Emulex 10GbE Virtual Fabric Adapter, Emulex 10GbE Virtual Fabric Adapter II, Broadcom 2-port 10Gb Virtual Fabric Adapter, BNT Virtual Fabric 10Gb Switch Module, and Cisco Nexus 4001I Switch Module.

Contents

Notices . . . vii
Trademarks . . . viii

Preface . . . ix
The team who wrote this book . . . ix
Now you can become a published author, too! . . . xii
Comments welcome . . . xii
Stay connected to IBM Redbooks . . . xiii

Part 1. Introduction to Virtual Fabric solutions . . . 1

Chapter 1. Overview of virtual network interface controllers . . . 3
1.1 Overview of virtualization technologies . . . 4
1.2 Overview of virtual NICs . . . 5
1.3 vNIC modes . . . 6
1.4 Mode comparison . . . 8
1.5 Mode selection . . . 9
1.6 BladeCenter Open Fabric Manager Advanced . . . 10

Chapter 2. Converged fabrics: FCoE and iSCSI capabilities . . . 11
2.1 Fibre Channel over Ethernet protocol stack . . . 12
2.2 iSCSI . . . 14
2.3 iSCSI versus FCoE . . . 15
2.3.1 Key similarities . . . 15
2.3.2 Key differences . . . 16

Part 2. IBM Virtual Fabric Mode . . . 17

Chapter 3. Solution architecture . . . 19
3.1 Overview . . . 20
3.2 Architecture . . . 22
3.2.1 Fibre Channel over Ethernet architecture . . . 25
3.2.2 Internet Small Computer System Interface architecture . . . 27
3.3 Planning considerations . . . 29

Chapter 4. Supported products . . . 31
4.1 Overview of supported components . . . 32
4.2 Emulex 10Gb Virtual Fabric Adapters . . . 33
4.2.1 Emulex 10GbE Virtual Fabric Adapter . . . 33
4.2.2 Emulex 10GbE Virtual Fabric Adapter II . . . 37

4.3 BNT Virtual Fabric 10Gb Switch Module 39
4.4 QLogic Virtual Fabric Extension Module 47

Chapter 5. Configuring the components 51
5.1 Configuring the Emulex 10GbE Virtual Fabric Adapter 52
5.1.1 Upgrading firmware 52
5.1.2 Configuring the adapter in the UEFI 54
5.2 Configuring the BNT Virtual Fabric 10Gb Switch Module 59
5.2.1 IBM System Networking OS CLI 59
5.2.2 isCLI 65
5.2.3 Browser-based interface 66
5.3 Creating vNIC in IBM Virtual Fabric Mode 69
5.3.1 vNIC features and considerations 70
5.3.2 Enabling and disabling vNIC functions on the server 71
5.3.3 Creating the vNICs 75
5.4 Configuring the vNIC group 79
5.5 High availability and vNICs 81
5.6 Configuring vNIC VLANs 85
5.7 Configuring the operating system 88
5.7.1 Windows OneCommand Manager 88
5.7.2 Configuring teams and VLANs 93
5.7.3 Configuring a new team 94
5.7.4 Configuring a VLAN for a physical or team adapter 96
5.7.5 Configuring FCoE and the personality for iSCSI Emulex Virtual Fabric Adapter II 98
5.8 vNIC and VMready 101

Chapter 6. Usage scenarios 103
6.1 Scenario 1: vNICs only with no uplinks 104
6.2 Scenario 2: vNICs only with one uplink 109
6.3 Scenario 3: vNICs only with one uplink trunk 114
6.4 Scenario 4: vNIC FCoE on Emulex II, BNT, and QLogic with IBM DS5300 120
6.4.1 Configuring the QLogic Virtual Fabric Extension Module 121
6.4.2 Implementing the FCoE-enabled BladeCenter with Virtual Fabric Switch and Extension Module 121
6.4.3 Defining the FCoE and FC fabric topology 125
6.4.4 Configuring the BNT Virtual Fabric 10Gb Switch Modules 126
6.4.5 Configuring the QLogic Virtual Extension Modules 132
6.4.6 Switching the Virtual Fabric Extension Module to N-Port Virtualization mode if connected to an existing FC fabric 137
6.4.7 Configuring the FCoE VLAN ID on the CNA 139
6.4.8 Configuring FCoE for IBM Virtual Fabric VNIC 141

6.4.9 Summary 145
6.5 Scenario 5: vNIC iSCSI on Emulex II, BNT, and QLogic with IBM DS5300 145
6.5.1 Key results 146
6.5.2 Configuration details for vNIC mode 146

Part 3. Switch Independent Mode 153

Chapter 7. Solution architecture 155
7.1 Mode comparison 156
7.1.1 IBM Virtual Fabric Mode 156
7.1.2 Switch Independent Mode 156
7.2 Solution overview 158
7.3 Features 160

Chapter 8. Supported products 165
8.1 Overview of Switch Independent Mode 166
8.2 Virtual Fabric adapters 167
8.2.1 Emulex 10GbE Virtual Fabric Adapter II 167
8.2.2 Broadcom 2-port 10Gb Virtual Fabric Adapter 171
8.3 Virtual Fabric Switch Modules 172
8.3.1 BNT Virtual Fabric 10Gb Switch Module 172
8.3.2 Cisco Nexus 4001I Switch Module 180
8.3.3 Brocade Converged 10GbE Switch Module for IBM BladeCenter . 188
8.4 10Gb Ethernet Pass-Thru Module for IBM BladeCenter 192

Chapter 9. Configuring the components 195
9.1 Configuring Virtual Fabric adapters 196
9.1.1 BladeCenter Open Fabric Manager Advanced 196
9.1.2 Upgrading the firmware for Emulex 10GbE Virtual Fabric Adapter II 196
9.1.3 Configuring the Emulex 10GbE Virtual Fabric Adapter II in UEFI . . 196
9.1.4 Configuring the Broadcom 2-port 10Gb Virtual Fabric Adapter in UEFI 197
9.2 Configuring the BNT Virtual Fabric 10Gb Switch Module 206
9.3 Upgrading the firmware for the Cisco Nexus 4001I Switch Module 206
9.4 Creating vNICs with the Broadcom 2-port 10Gb Virtual Fabric Adapter . 209
9.5 Configuring vNIC VLANs 214
9.6 vNIC and VMready 215
9.7 Understanding vNIC representation in an operating system 216
9.7.1 vNICs on the Emulex 10GbE Virtual Fabric Adapter II 216
9.7.2 vNICs on the Broadcom 2-port 10Gb Virtual Fabric Adapter 218
9.7.3 Configuring teams with Emulex OneCommand Manager 225
9.7.4 Configuring teams with Broadcom Advanced Control Suite 229

Chapter 10. Usage scenarios . . . 237
10.1 Overview of the configurations . . . 238
10.1.1 VLAN definitions . . . 239
10.1.2 Switch and adapter vNIC correlation. . . . 241
10.2 Scenario 1: Cisco Nexus 4001I Switch Module configurations . . . 242
10.2.1 vNICs summary. . . . 242
10.2.2 Configuring the Emulex adapter . . . 243
10.2.3 Configuring the Broadcom adapter . . . 255
10.2.4 Configuring the Cisco Nexus 4001I Switch Module . . . 272
10.3 Scenario 2: BNT Virtual Fabric 10Gb Switch Module configurations . . 274
10.3.1 Configuring the Emulex adapter . . . 275
10.3.2 Configuring the Broadcom adapter . . . 287
10.3.3 Configuring the BNT Virtual Fabric 10Gb Switch Module . . . 304
10.3.4 Configuring Switch Independent Mode for Emulex II, BNT, QLogic, and IBM Storwize V7000 . . . 305
10.3.5 Configuring Switch Independent Mode: Emulex Virtual Fabric Adapter II, BNT, and IBM Storwize V7000 or iSCSI. . . . 337
10.4 Scenario 3: Brocade Converged Switch configuration . . . 367
10.4.1 Enabling the Virtual Fabric Adapter Switch Independent Mode function on UEFI. . . . 367
10.4.2 Installing the Windows 2008 drivers . . . 367
10.4.3 Installing Emulex OneCommand Manager . . . 368
10.4.4 Updating the Virtual Fabric Adapter firmware . . . 370
10.4.5 Configuring the ports in operation mode . . . 375
10.4.6 Virtual Fabric Adapter II: Configuring vNICs with FCoE and Brocade converged switch to IBM Storwize V7000 . . . 377

Part 4. Abbreviations, related publications, and index . . . 383

Abbreviations and acronyms . . . 385

Related publications . . . 389
IBM Redbooks . . . 389
Other publications . . . 389
Online resources . . . 390
Help from IBM . . . 390

Index . . . 391

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

Preface

The deployment of server virtualization technologies in data centers requires significant efforts in providing sufficient network I/O bandwidth to satisfy the demand of virtualized applications and services. For example, every virtualized system can host several dozen network applications and services. Each of these services requires certain bandwidth (or speed) to function properly. Furthermore, because of different network traffic patterns that are relevant to different service types, these traffic flows can interfere with each other. They can lead to serious network problems, including the inability of the service to perform its functions.

The IBM® Virtual Fabric solution for IBM BladeCenter® addresses these issues. The solution is based on the IBM BladeCenter H chassis with a 10 Gbps Converged Enhanced Ethernet infrastructure. This infrastructure is built on 10 Gbps Ethernet switch modules in the chassis and the Emulex or Broadcom Virtual Fabric Adapters in each blade server.

This IBM Redbooks® publication provides configuration scenarios that use technology from five global leaders in network switch and adapter technology: Broadcom, Brocade, Cisco, Emulex, and IBM. These scenarios demonstrate the usage of IBM System Networking technology and products from Broadcom, Brocade, Cisco, and Emulex. This book is for clients who want to learn how to implement an IBM Virtual Fabric solution by using the IBM Virtual Fabric Mode offering or the Switch Independent Mode offering. This book explains step-by-step how to configure the adapters and switches.

The team who wrote this book

This second edition was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Raleigh Center.

Rufus Credle is a Certified Consulting IT Specialist and Project Leader at the ITSO, Raleigh Center. He conducts residencies and develops IBM Redbooks and Redpapers™ about various topics, including network operating systems, high availability, IBM industry technology, IBM System x®, and IBM BladeCenter. Rufus has worked for IBM for 31 years in various positions, including assignments in administration and asset management, systems engineering, sales and marketing, and IT services. He has a Bachelor of Science (BS) degree in Business Management from Saint Augustine’s College.

Martin Gingras is a senior IT Specialist working with IBM Global Technology Services® in Canada. He has 13 years experience within the IT industry. His areas of expertise include IBM BladeCenter, System x, Microsoft Windows, VMware, and Linux. Martin has worked the last nine years on the RTSC platform, helping customers resolve difficult cases and working on improving System x and BladeCenter brand from a service point of view. He is involved in many issues involving servers, Ethernet switches and FC switches, and SAN issues. Martin also has a focus in IBM tools such as Dynamic System Analysis (DSA), IBM ToolsCenter Bootable Media Creator, and IBM UpdateXpress System Pack Installer. Martin holds a computer management certification (DEC Informatique de gestion) from Cecep François-Xavier Garneau.

Scott Lorditch is a Consulting Systems Engineer for IBM Systems & Technology Group, Storage Platform. He performs network architecture assessments, and develops designs and proposals for implementing GbE Switch Module products for the IBM BladeCenter. He also developed several training and lab sessions for IBM technical and sales personnel. Previously, Scott spent almost 20 years working on networking in various industries, working as a senior network architect, a product manager for managed hosting services, and manager of electronic securities transfer projects. Scott holds a BS degree in Operations Research with a specialization in computer science from Cornell University.

James Mulholland is an IT Specialist working with IBM Global Technology Services in Portugal and previously in Ireland. He has 15 years of experience in the IT industry. His areas of expertise include Microsoft Windows, Linux, System x, IBM BladeCenter, and various SAN and storage products. During the last 9 years, he focused mainly on SAN, storage implementation, and management. James supports many customers and has been responsible for many implementations and projects throughout Portugal. He holds a masters degree in IT, with a specialty in information security.

Bob Nevins is an IBM and Open Group Master Certified IT Specialist. He has 29 years of experience in the IT industry, with 27 of those years working with IBM networking and IBM cross-brand technical sales. Bob has worked in roles ranging from consultant, software designer, and network architect to teacher and technical specialist. He holds a BS degree in Mathematics and a Master of Business Administration (MBA) degree, both from Duke University.

Rogerio Spragiaro is a Field Technical Sales Specialist for IBM in Brazil. He has hands-on experience with IBM System x that use Windows, Linux, and VMware and spent five years providing level 2 support for x86 servers. As a technical consultant to IBM customers for over ten years, he has a vast knowledge of business and technical environments, such as high performance computing, database, virtualization, and networking solutions that help the sale of IBM smart solutions. Rogerio is certified on IBM System x and BladeCenter products, in addition to Red Hat Virtualization (RHEV).

Thanks to the authors of the previous editions of this book:

Jure Arzensek
Valeria Lima
Norm Stole
David Watts

The authors thank the following people for their contributions to this project:

Kevin Liebl
Maziar Tamadon
Broadcom

Tom Boucher
Leanne Greene
Geoff Gundlach
Jeff Lin
Howard Lo
Brent Mosbrook
Hemal Purohit
Emulex

Manzar Alam
Barry Barnett
Ralph Geiger
Mike Nolterieke
Marc Stracuzza
Tony Tran
IBM Development

Karen Lawrence
Linda Robinson
Alfred Schwab
IBM ITSO

Richard Mancini
Shekhar Mishra
Michael Talplacido
IBM Marketing

Glenn Sullivan
IBM Storage

Additional thanks goes to the following authors who wrote *IBM BladeCenter Virtual Fabric Solutions*, REDP-4673, on which this book is partially based:

Tom Boucher
Scott Lorditch
Mark McCool
David Watts

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Part 1

Introduction to Virtual Fabric solutions

This part provides an introduction to the concepts in this book. It includes the following chapters:

- Chapter 1, “Overview of virtual network interface controllers” on page 3
- Chapter 2, “Converged fabrics: FCoE and iSCSI capabilities” on page 11

1

Overview of virtual network interface controllers

This chapter provides an overview of virtualization technologies, virtual network interface controllers (vNICs), and their modes: IBM Virtual Fabric Mode and Switch Independent Mode. It also highlights the support of BladeCenter Open Fabric Manager Advanced as it relates to deployment of the Emulex and Broadcom 2-port 10Gb Virtual Fabric Adapters when configured for Switch Independent Mode of operation.

This chapter includes the following sections:

- Overview of virtualization technologies
- Overview of virtual NICs
- vNIC modes
- Mode comparison
- Mode selection
- BladeCenter Open Fabric Manager Advanced

1.1 Overview of virtualization technologies

The deployment of server virtualization technologies in data centers requires significant efforts to provide sufficient network I/O bandwidth (or speed) to satisfy the demand of virtualized applications and services. For example, every virtualized system can host several dozen network applications and services, and each of these services requires a certain bandwidth to function properly. Furthermore, because of different network traffic patterns relevant to different service types, these traffic flows might interfere with each other. This interference can lead to serious network problems, including the inability of the service to perform its functions.

Providing sufficient bandwidth and isolation to virtualized applications in a 1 Gbps network infrastructure might be challenging for blade-based deployments where the number of physical I/O ports per blade is limited. For example, of the maximum of eight physical ports per blade, at least two ports are usually used to connect to external storage. Therefore, only six 1-Gbps Ethernet ports are available, with a total network bandwidth of 6 Gb per blade for Gigabit Ethernet infrastructures, leaving no room for future growth.

In addition, traffic flows are isolated on a physical port basis. Also the bandwidth per interface is static with a maximum bandwidth of 1 Gb per flow, thus limiting the flexibility of bandwidth usage. IBM BladeCenter Virtual Fabric solutions address these issues by increasing the number of available Ethernet ports and providing more flexibility in allocating the available bandwidth to meet specific application requirements.

The BladeCenter Virtual Fabric solutions implement vNICs. These solutions are based on the IBM BladeCenter H chassis with a 10 Gbps Converged Enhanced Ethernet (CEE) switch infrastructure and 10 Gbps Virtual Fabric Adapters installed in each blade server.

The basic concept of the BladeCenter Virtual Fabric solution is to take the two 10-Gbps ports that are on a 10 Gbps Virtual Fabric Adapter and split them into eight vNICs. With this configuration, each vNIC or virtual channel can be between 100 MB and 10 Gb in increments of 100 MB. This way, the total of all eight vNICs does not exceed the 20 Gb total that the two 10-Gbps ports offer. The separation of the virtual channels continues into the switch module by using virtual local area network (VLAN) tagging to identify the various data streams.

1.2 Overview of virtual NICs

A *network interface controller* (NIC) is a component in a blade server that allows the server to be connected to a network. The NIC provides the physical point of connection and the internal software for encoding and decoding network packets.

With the introduction of Virtual Fabric in 2009, IBM, Emulex, and BNT® (now known as IBM System Networking) established the new vNIC feature. vNIC enables virtualization of the physical interface. Rather than eight physical connections between the adapter and the switch, there are two 10-Gbps interfaces to two separate switch modules. Also, each 10 Gbps interface can appear to the operating system as four separate physical interfaces.

Tip: The vNIC concept described here is associated with the physical port. Do not confuse it with the concept of virtualization within a hypervisor-based operating system that provides virtual I/O interfaces to the virtual machines.

Splitting up a single interface into four interfaces is achieved by following the Peripheral Component Interconnect Express (PCIe) standard of the PCI device ID. Although not a new standard, each device can have one ID, and that device can have eight functions. Many dual- or quad-port Ethernet or Fibre Channel Adapters also have a single application-specific integrated circuit (ASIC) chip, but the physical interfaces are iterated as a function. Therefore, in a sense, the Virtual Fabric Adapter has virtualized Layer 1 of the open systems interconnection (OSI) model, by virtualizing the physical interface.

Virtualizing the NIC helps to resolve issues that are caused by limited NIC slot availability. By virtualizing a 10 Gbps NIC, its resources can be divided into multiple logical instances or vNICs. Each vNIC appears as a regular, independent NIC to the server operating system or hypervisor, with each vNIC using a portion of the overall bandwidth of the physical NIC. For example, a NIC partition with a maximum bandwidth of 4 Gbps appears to the host applications as a physically distinct 4 Gbps Ethernet adapter.

The vNIC partitions provide traffic forwarding and port isolation based on the destination MAC address, the VLAN tag, or both. It does not provide bridging functions, either within a partition or across partitions.

1.3 vNIC modes

Two operational modes are available. They depend on the type of 10 Gbps expansion card installed in the server and the type of 10 Gbps switch installed in the BladeCenter chassis:

- IBM Virtual Fabric Mode

  IBM Virtual Fabric Mode can be enabled on the combinations listed in Table 1-1.

Table 1-1 Supported combinations using IBM Virtual Fabric Mode (vNIC or vNIC 1 mode)

Supported adapters	**Supported switch: BNT Virtual Fabric 10Gb Switch Module**
Emulex 10GbE Virtual Fabric Adapter, 49Y4235q	Supported
Emulex 10GbE Virtual Fabric Adapter Advanced, 49Y4275	Supported
Emulex 10GbE Virtual Fabric Adapter II, 90Y3550	Supported
Emulex 10GbE Virtual Fabric Adapter II Advanced, 90Y3566	Supported

  In this configuration, each blade server has up to 20 Gb of available network bandwidth (two 10 Gbps ports). In addition, each blade server can split this bandwidth among up to eight vNICs, starting at either 100 MB or 1 Gb (adapter-dependent), and up to a maximum of 10 Gb per one vNIC channel with 100 Mbps increments.

  Channel bandwidth is configured in the IBM System Networking BNT Virtual Fabric Switch for each port and passed to each of the Virtual Fabric Adapters. Dynamic bandwidth support allows the channel bandwidth to be increased or decreased under the control of the Virtual Fabric Switch without resetting the Virtual Fabric Adapter. Also, this mode supports bidirectional bandwidth control. Both the server outbound virtual channel bandwidth limits are enforced by the Virtual Fabric Adapter, and the inbound virtual channel bandwidth limits are enforced by Virtual Fabric Switch. By using NIC teaming at the Virtual Fabric Adapter, an individual vNIC channel can fail over without impacting other vNIC channels on the same physical link.

- Switch Independent Mode

  As the name suggests, Switch Independent Mode is more independent of the switch module used and can be enabled on the combinations listed in Table 1-2 on page 7.

Table 1-2 Supported combinations using Switch Independent Mode (vNIC 2 mode)

Supported adapters	Supported switches		
	BNT Virtual Fabric 10Gb Switch Module	**Cisco Nexus 4001I Switch Module**	**Brocade Converged 10GbE Switch Module**
Emulex 10GbE Virtual Fabric Adapter, 49Y4235	Not supported	Not supported	Not supported
Emulex 10GbE Virtual Fabric Adapter Advanced, 49Y4275	Not supported	Not supported	Not supported
Emulex 10GbE Virtual Fabric Adapter II, 90Y3550	Supported	Supported	Supported
Emulex 10GbE Virtual Fabric Adapter II Advanced, 90Y3566	Supported	Supported	Supported
Broadcom 2-port 10Gb Virtual Fabric Adapter, 81Y3133	Supported	Supported	Supported

Switch Independent Mode extends the IEEE 802.1q VLAN domains in the network infrastructure to specific virtual channels at each server. Up to 8 channels can be configured for each Virtual Fabric Adapter, with outbound bandwidths per channel from 100 Mb to 10 Gb in 100 Mb increments. In the release of an upcoming version of BladeCenter Open Fabric Manager, the channel bandwidth can be configured by using the adapter utility. Inbound channel bandwidth constraints (for example, from the switch to the Virtual Fabric Adapter) are not required.

Combining the Emulex 10GbE Virtual Fabric Adapter II and the BNT Virtual Fabric 10Gb Switch Module supports both modes. The actual mode that is used is specified in the Unified Extensible Firmware Interface (UEFI) menu of the adapter.

With either mode, the integrated Gigabit Ethernet ports on the blade servers are still available, connecting to 1 Gbps Ethernet switches installed in bays 1 and 2 of the chassis. A CIOv expansion card in the blade servers and I/O modules in bays 3 and 4 of the BladeCenter H chassis are now added. Therefore, additional network ports or storage ports can be made available to applications for more expansion capabilities.

The combination of the integrated Gigabit Ethernet ports, the CIOv expansion card, and the eight vNICs of the Virtual Fabric Adapter means that up to 12 I/O ports are available to each blade in the chassis.

For more information about supported configurations, see the BladeCenter Interoperability Guide (BIG) at:

http://www.ibm.com/support/entry/portal/docdisplay?brand=5000020&lndocid=MIGR-5073016

1.4 Mode comparison

Table 1-3 compares IBM Virtual Fabric Mode and Switch Independent Mode.

Table 1-3 Comparison of modes

	IBM Virtual Fabric Mode	**Switch Independent Mode**
Number of virtual channels per physical 10 Gbps port	4	4
Total number of virtual channels per adapter	8	8
Minimum virtual channel bandwidth (can also be set to 0 to disable the channel)	► Emulex 10 Gbps Virtual Fabric Adapter: 100 Mb[a] ► Emulex 10 Gbps Virtual Fabric Adapter II: 100 Mb	100 MB
Server-to-switch bandwidth limit per virtual channel	Yes	Yes
Switch-to-server bandwidth limit per virtual channel	Yes	No
Supported expansion cards	► Emulex 10GbE Virtual Fabric Adapter ► Emulex 10GbE Virtual Fabric Adapter II	► Emulex 10GbE Virtual Fabric Adapter II ► Broadcom 2-port 10Gb Virtual Fabric Adapter
Supported switches	► BNT Virtual Fabric 10Gb Switch Module	► BNT Virtual Fabric 10Gb Switch Module ► Cisco Nexus 4001I Switch Module ► Brocade Converged 10GbE Switch Module
Number of unique customer VLANs per virtual channel	Maximum allowed by applications	Varies by adapter ► Emulex Virtual Fabric Adapter: 128 VLANs (4 channels) ► Broadcom Virtual Fabric Adapter: Maximum allowed by applications

	IBM Virtual Fabric Mode	Switch Independent Mode
IEEE 802.1q VLAN tagging	Optional	Required
VLAN trunking on switch uplinks	Not supported	Supported
Isolated NIC teaming failover per vNIC channel	Supported	Not supported
iSCSI support: Hardware initiator	► Emulex Virtual Fabric Adapter I: Yes ► Emulex Virtual Fabric Adapter II: Yes	► Broadcom Virtual Fabric Adapter: No ► Emulex Virtual Fabric Adapter II: Yes
iSCSI support: Software initiator	Yes	Yes
Fibre Channel over Ethernet (FCoE) channel option	Yes	Yes
Switch stacking	Yes	Yes (BNT switch only)
Maximum number of virtual groups per switch	32 (1 per virtual channel)[b]	► BNT: 1024 ► Cisco: 512

a. Support for a minimum bandwidth allocation of 100 Mbps requires BNT Virtual Fabric 10Gb Switch Module firmware 6.5.3.0 or later and Emulex 10GbE Virtual Fabric Adapter firmware 2.103.269.3x or later. Otherwise, the minimum allocation is 1 Gbps.

b. One vNIC port per physical port per group.

1.5 Mode selection

Use IBM Virtual Fabric Mode if your network environment has the following characteristics:

- You have a BNT switch module.
- You do not already implement IEEE 802.1q VLAN tags in your network infrastructure. For example, you might have implemented multiple IP subnets with Layer 3 separation among the Layer 2 broadcast domains that are provided by switches and routers. Servers might have multiple NIC adapters with attachments to physically separate switches or port-based VLAN groups to provide Layer 2 separation.
- You implemented VLAN tagging only in the network infrastructure so that the tagging is transparent to the applications. This implementation requires the

switches to add and remove VLAN tags, rather than the server applications or NIC drivers adding and removing them.

- You require overlapping IEEE 802.1q VLANs per NIC interface. Two (or more) individual vNICs on a given Virtual Fabric Adapter must not share the VLAN in either mode.
- You want to implement a solution that takes advantage of Single Root IO Virtualization (SR-IOV).

Use Switch Independent Mode if your network environment has the following characteristics:

- You chose to use the Cisco Nexus 4001I Switch Module in your BladeCenter H or HT chassis.
- You want to implement some of the advanced Layer 2 or Layer 3 IBM System Networking BNT switch features that are not available in IBM Virtual Fabric Mode.
- You want to implement the Broadcom 2-port 10Gb Virtual Fabric Adapter to maintain driver compatibility with the base 1 Gbps NIC cards.
- You want to trunk more than 1 VLAN per external port.

1.6 BladeCenter Open Fabric Manager Advanced

IBM plans to update BladeCenter Open Fabric Manager Advanced to support the deployment of Emulex and Broadcom 2-port 10Gb Virtual Fabric Adapters when configured for Switch Independent Mode of operation. After this update, you can configure the Virtual Fabric Adapters across multiple chassis from a central control point. Advanced BladeCenter Open Fabric Manager pushes the parameters to a BladeCenter Open Fabric Manager instance that is within the BladeCenter advanced management module (AMM) in each chassis. In turn, the AMM pushes the parameters to the Virtual Fabric Adapter on each blade. You can configure the channel bandwidth from the Templates panel for the vNIC ports.

You must still perform the VLAN configuration in the blade application drivers and the corresponding switch configuration as described in Part 3, “Switch Independent Mode” on page 153.

2

Converged fabrics: FCoE and iSCSI capabilities

Fibre Channel over Ethernet (FCoE) is a method of sending Fibre Channel (FC) frames directly over an Ethernet network. It relies on a new Ethernet transport with extensions that provide lossless transmission of storage data. Switches that support FCoE are required. Some switches require a license update to support FCoE.

With an Internet Small Computer System Interface (iSCSI) network, you can use an existing traditional Ethernet networking infrastructure, reducing the costs for specialized storage area networking devices, software, and licenses. iSCSI has a head start for many businesses because it already has a stable network infrastructure in place. iSCSI uses the reliable TCP protocol to transport SCSI I/O commands over a network, providing block-level data access without the necessity for specialized hardware requirements. It can also operate with various peripheral devices.

This chapter includes the following sections:

- Fibre Channel over Ethernet protocol stack
- iSCSI
- iSCSI versus FCoE

2.1 Fibre Channel over Ethernet protocol stack

FCoE assumes the existence of a lossless Ethernet, such as one that implements the Data Center Bridging (DCB) extensions to Ethernet. This section highlights, at a high level, the concepts of FCoE as defined in FC-BB-5. The BNT supports FCoE with updated firmware. The Cisco 4001 supports FCoE with a license upgrade.

The basic notion of FCoE is that the upper layers of FC are mapped onto Ethernet, as shown in Figure 2-1. The upper layer protocols and services of FC remain the same in an FCoE deployment. Zoning, fabric services, and similar services still exist with FCoE.

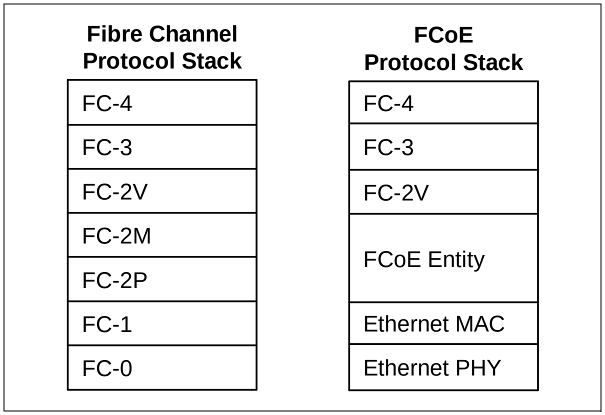

Figure 2-1 FCoE protocol mapping

The difference is that the lower layers of FC are replaced by lossless Ethernet, which also implies that FC concepts, such as port types and lower-layer initialization protocols, must be replaced by new constructs in FCoE. Such mappings are defined by the FC-BB-5 standard and are briefly addressed here.

Figure 2-2 shows another perspective on FCoE layering compared to other storage networking technologies. In this figure, FC and FCoE layers are shown with other storage networking protocols, including iSCSI.

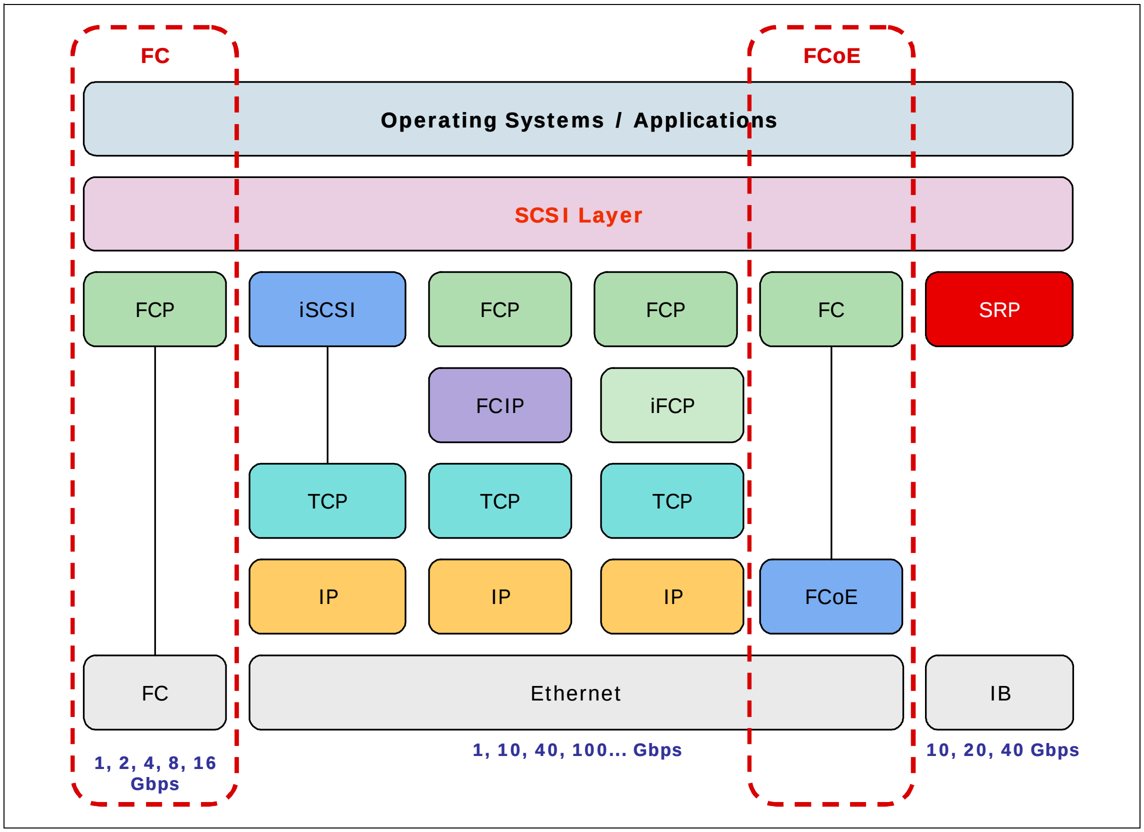

Figure 2-2 Storage Network Protocol Layering

Based on this protocol structure, Figure 2-3 shows a conceptual view of an FCoE frame.

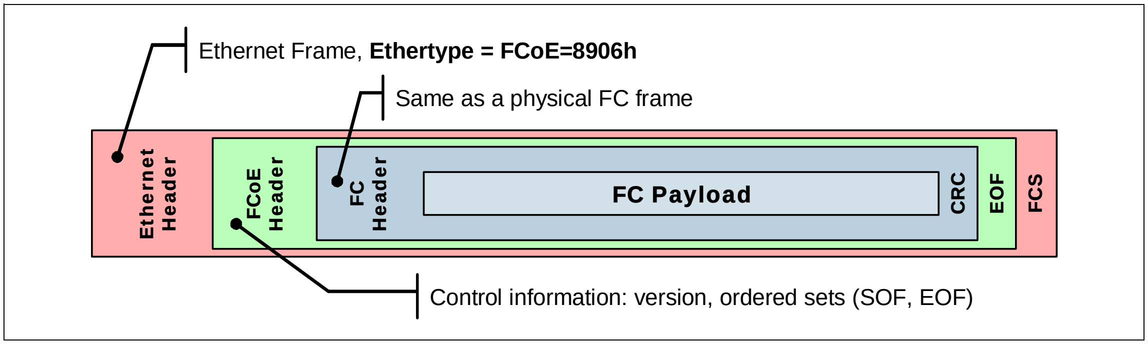

Figure 2-3 Conceptual view of an FCoE frame

2.2 iSCSI

The iSCSI protocol allows for longer distances between a server and its storage when compared to the traditionally restrictive parallel SCSI solutions or the newer serial-attached SCSI (SAS). iSCSI technology can use a hardware initiator, such as a host bus adapter (HBA), or a software initiator to issue requests to target devices. Within iSCSI storage terminology, the initiator is typically known as a *client,* and the target is the storage device. The iSCSI protocol encapsulates SCSI commands into protocol data units (PDUs) within the TCP/IP protocol and then transports them over the network to the target device. The disk is presented locally to the client as shown in Figure 2-4.

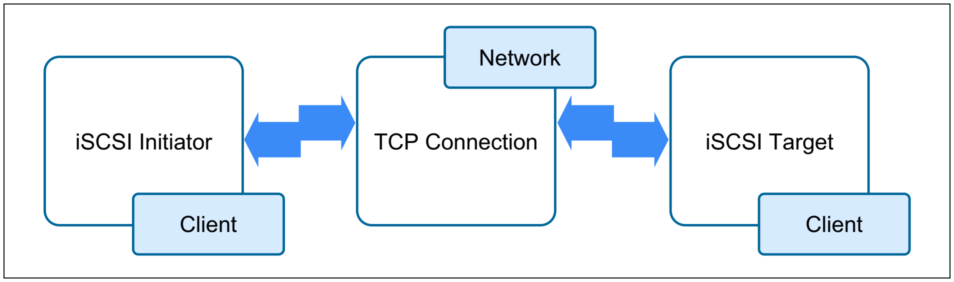

Figure 2-4 iSCSI architecture overview

The iSCSI protocol is a transport for SCSI over TCP/IP. Figure 2-2 on page 13 illustrates a protocol stack comparison between Fibre Channel and iSCSI. iSCSI provides block-level access to storage, as does Fibre Channel, but uses TCP/IP over Ethernet instead of Fibre Channel protocol. iSCSI is defined in RFC 3720, which you can find at:

http://www.ietf.org/rfc/rfc3720.txt

iSCSI uses Ethernet-based TCP/IP rather than a dedicated (and different) storage area network (SAN) technology. Therefore, it is attractive for its relative simplicity and usage of widely available Ethernet skills. Its chief limitations are the relatively lower speeds of Ethernet compared to Fibre Channel and the extra TCP/IP encapsulation required. With lossless 10 Gbps Ethernet now becoming available, the attractiveness of iSCSI is expected to grow rapidly. TCP/IP encapsulation will still be used, but 10 Gbps Ethernet speeds will dramatically increase the appeal of iSCSI.

2.3 iSCSI versus FCoE

The section highlights the similarities and differences between iSCSI and FCoE. However, in most cases, considerations other than purely technical ones will influence your decision in choosing one over the other.

2.3.1 Key similarities

iSCSI and FCoE have the following similarities:

- Both protocols are block-oriented storage protocols. That is, the file system logic for accessing storage with either of them is on the computer where the initiator is, not on the storage hardware. Therefore, they are both different from typical network-attached storage (NAS) technologies, which are file oriented.
- Both protocols implement Ethernet-attached storage.
- Both protocols can be implemented in hardware, which is detected by the operating system of the host as an HBA.
- Both protocols can use the Converged Enhanced Ethernet (CEE), also referred to as *Data Center Bridging*), standards to deliver “lossless” traffic over Ethernet.
- Both protocols are alternatives to traditional FC storage and FC SANs.

2.3.2 Key differences

iSCSI and FCoE have the following differences:

- iSCSI uses TCP/IP as its transport, and FCoE uses Ethernet. iSCSI can use media other than Ethernet, such as InfiniBand, and iSCSI can use Layer 3 routing in an IP network.
- Numerous vendors provide local iSCSI storage targets, some of which also support Fibre Channel and other storage technologies. Relatively few native FCoE targets are available at this time, which might allow iSCSI to be implemented at a lower overall capital cost.
- FCoE requires a gateway function, usually called a *Fibre Channel Forwarder* (FCF), which allows FCoE access to traditional FC-attached storage. This approach allows FCoE and traditional FC storage access to coexist either as a long-term approach or as part of a migration.
- iSCSI-to-FC gateways exist but are not required when a storage device is used that can accept iSCSI traffic directly. The DS5300 that we used in our tests for this book is one such device.
- Except in the case of a local FCoE storage target, the last leg of the connection uses FC to reach the storage. FC uses 8b/10b encoding, which means that, sending 8 bits of data requires a transmission of 10 bits over the wire or 25% payload that is transmitted over the network to prevent corruption of the data. The 10 Gbps Ethernet uses 64b/66b encoding, which has a far smaller payload.
- iSCSI includes IP headers and Ethernet (or other media) headers with every frame, which adds payload.
- The largest payload that can be sent in an FCoE frame is 2112. iSCSI can use jumbo frame support on Ethernet and send 9K or more in a single frame.
- iSCSI has been on the market for several years longer than FCoE. Therefore, the iSCSI standards are more mature than FCoE.
- Perhaps because of the relative immaturity of FCoE, in our testing, we had more difficulty with FCOE. The implementation times were longer. Troubleshooting FCoE end-to-end requires Ethernet networking skills and FC SAN skills.

Part 2

IBM Virtual Fabric Mode

This part provides product and implementation details for the IBM Virtual Fabric Mode. This part includes the following chapters:

- Chapter 3, “Solution architecture” on page 19
- Chapter 4, “Supported products” on page 31
- Chapter 5, “Configuring the components” on page 51
- Chapter 6, “Usage scenarios” on page 103

3

Solution architecture

The Virtual Fabric architecture is based on a virtual channel concept. This architecture depends upon both the Ethernet adapter and the Ethernet switch that are used to provide separation of the virtual channels on the shared 10 Gbps link. This chapter introduces the technology and architecture of IBM Virtual Fabric Mode.

This chapter includes the following sections:

- Overview
- Architecture
- Planning considerations

3.1 Overview

As described in 1.2, “Overview of virtual NICs” on page 5, a single 10 Gbps link can be partitioned into up to four separate virtual channels. Each channel can be assigned a specified bandwidth. Traffic separation is maintained among the channels.

The IBM Virtual Fabric Mode, with the Emulex 10GbE Virtual Fabric Adapter and the BNT Virtual Fabric 10Gb Switch Module, uses a double virtual local area network (VLAN) tagging scheme. By using this scheme, IBM Virtual Fabric Mode maintains separation of the virtual channels on the shared 10 Gbps link. Figure 3-1 shows an example of an implementation of IBM Virtual Fabric Mode and the VLAN tagging that is employed.

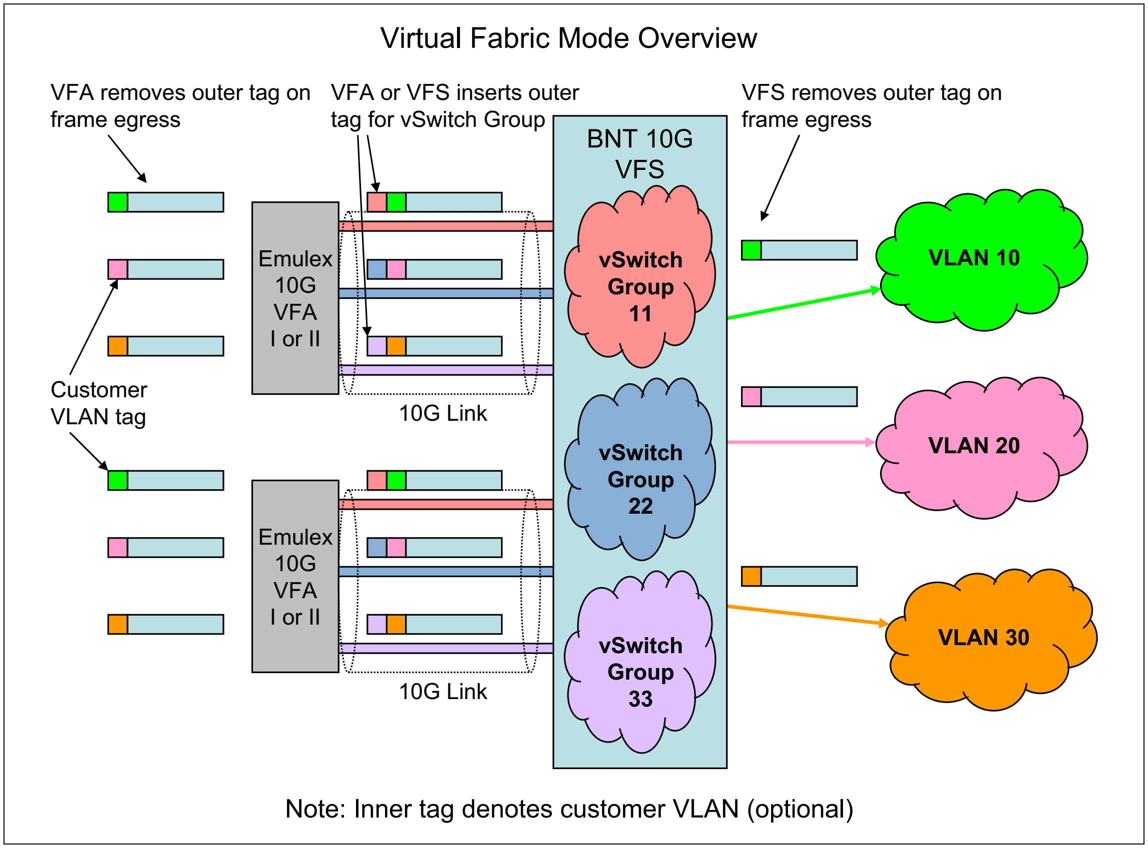

Figure 3-1 Overview of IBM Virtual Fabric Mode

The Virtual Fabric Adapter applies a unique outer tag to the packets. By using this tag, the switch can correctly handle the inbound frames and forward them to the appropriate subset of ports. Likewise, the switch applies a corresponding outer tag on all packets that are received from the network infrastructure by using the uplink ports. The switch also ensures that these packets are sent only to Virtual Fabric Adapters that are members of the corresponding virtual network interface controller (vNIC) group. By using this double tag scheme, the VLANs can be tunneled through the Virtual Fabric Adapter and the switch. Clients do not need to make any concessions in regard to their existing network infrastructure Layer 2 networking scheme.

Configuration is provided through the switch user interfaces. Parameters are passed to the adapter by the switch. Channel bandwidth is dynamic in that it can be changed (from the switch) without resetting the Virtual Fabric Adapter.

The Virtual Fabric Adapter enforces bandwidth limits for each virtual channel. The Virtual Fabric Switch enforces the same virtual channel bandwidth limits on traffic that is flowing to the server.

Bandwidth allocation and metering is controlled by the Virtual Fabric Switch. In cases when outgoing traffic flow for a particular vNIC exceeds established bandwidth limits, the switch might drop packets to maintain the traffic below the configured thresholds.

To isolate vNIC traffic flows from each other, vNICs can be assigned to vNIC logical groups (or *vSwitches*) that behave as independent switches inside the Virtual Fabric Switch. That is, vNICs in the same group can communicate with each other, but vNICs in different groups cannot communicate with each other. This vSwitch group ID is appended as an outer VLAN tag by the Virtual Fabric Switch or the Virtual Fabric Adapter. In addition, to external or internal switch ports, external trunks can also be assigned to a particular vNIC group to isolate upstream connections.

The Virtual Fabric Adapter must be in vNIC mode to support the IBM Virtual Fabric Mode solution. vNIC mode is the default mode of operations for the Virtual Fabric Adapter. In such a configuration, each physical port on the Virtual Fabric Adapter is divided into four virtual vNICs that are visible by the operating system as independent Ethernet controllers. The Virtual Fabric Switch recognizes the vNICs and manages and isolates network traffic flows between them.

The BNT Virtual Fabric 10Gb Switch Module disables several Layer 2 and Layer 3 switch functions when supporting the IBM Virtual Fabric Mode. For example, ACL filters, IBM VMready®, Layer 2 spanning tree, and Layer 3 routing are some functions that are not supported with Virtual Fabric Mode operation. These functions are disabled at the port level, but not for the entire switch. For example, you cannot configure vNIC and VMready on the same switch port (blade server), but you can configure them on different ports on the same switch module.

> **Switch Independent Mode:** These restrictions do not apply to Switch Independent Mode, as described in Chapter 7, “Solution architecture” on page 155.

The IBM System Networking switch supports stacking for both IBM Virtual Fabric Mode and Switch Independent Mode.

3.2 Architecture

Figure 3-2 provides a general view of the vNIC solution architecture and offers an example of configuring the vNICs to have different bandwidth values.

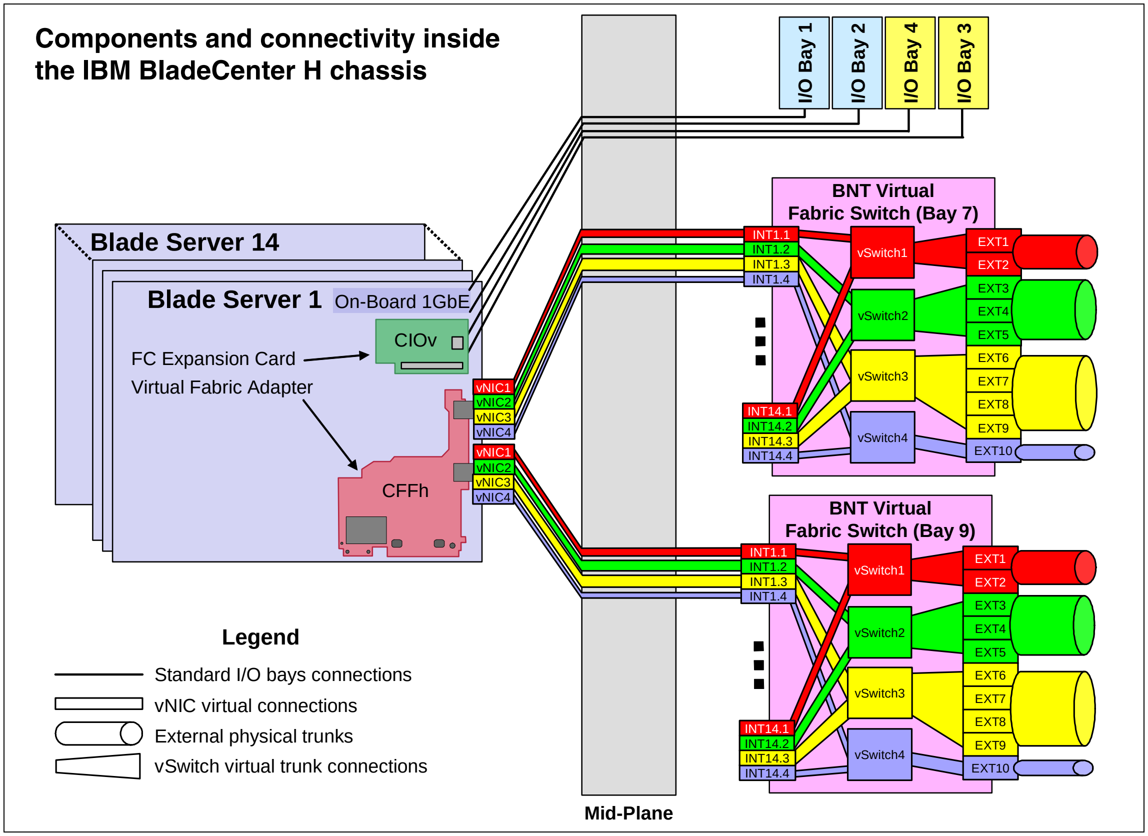

Figure 3-2 BladeCenter Virtual Fabric vNIC solution

The colors represent different virtual port groups to isolate traffic flows for different services and applications. The thickness of each link represents the proportional amount of assigned bandwidth for the particular traffic flow.

The Virtual Fabric architecture provides the capability to subdivide a single 10 Gbps link into up to four virtual channels.

Trunk failover is supported on a per-vNIC basis, so that you can set up NIC teaming in a failover mode for vNICs. If an external trunk or uplink failure occurs, the switch shuts down only those vNICs that belong to the same vNIC group as the failed external trunk. However, other vNICs on the same physical NIC (pNIC) port are untouched and continue to perform their functions.

vNICs are configured on the BNT Virtual Fabric 10Gb Switch Module by dividing internal Virtual Fabric Switch ports into subinterfaces (or logical ports). Each internal interface of the Virtual Fabric Switch can be divided into up to four vNICs. They maintain strict one-to-one mapping to vNICs on the Virtual Fabric Adapter to a logical vSwitch group within the switch and the corresponding uplink ports.

Table 3-1 is a vNIC mapping for Virtual Fabric Adapter and Virtual Fabric Switch, in addition to vNIC identification data, to differentiate vNICs on the host operating system (vNIC Peripheral Component Interconnect Express (PCIe) function ID).

Table 3-1 vNIC mappings and identification

vNIC PCIe function ID	Virtual Fabric Adapter port number	Virtual Fabric Switch bay number	vNIC number	vNIC ID[a]
0	0	7	1	INTx.1
2	0	7	2	INTx.2
4	0	7	3	INTx.3
6	0	7	4	INTx.4
1	1	9	1	INTx.1
3	1	9	2	INTx.2
5	1	9	3	INTx.3
7	1	9	4	INTx.4

a. A vNIC ID of *x,* where *x* is 1 - 14, represents the Virtual Fabric Switch internal port number to which the Virtual Fabric Adapter physical port is connected.

This function has the following general characteristics:

- Each server pNIC port is divided into up to four vNICs.
- Operating system configurations detect eight unique NICs (2 ports x 4 vNICs).
- All vNIC parameters are configured from the Virtual Fabric Switch user interface.

  The Data Center Bridging Exchange (DCBX) protocol is used between the switch and NIC to convey configuration information.
- The user enables vNICs and allocates the bandwidth.
  - The allowable vNIC bandwidth range is 100 Mbps or 1 Gbps to 10 Gbps (Virtual Fabric Adapter model dependent):
    - The increments are 100 Mbps.
    - The default bandwidth setting is 2.5 Gbps.
    - The sum of all four vNICs cannot exceed 10 Gbps.
  - Virtual Fabric Switch-to-server bandwidth does metering on a per vNIC per port basis.
  - Server-to-Virtual Fabric Switch bandwidth does metering on a per-vNIC basis.
- The user assigns vNICs and, optionally, uplinks to vNIC groups:
  - Groups serve to isolate virtual NIC traffic that flows on the same physical port:
    - Existing VLANs within the client network are not affected.
    - No forwarding occurs between uplinks that are assigned to vNIC groups.
    - Up to 32 vNIC groups are supported per Virtual Fabric Switch.
    - An uplink (port or trunk) can belong to only one vNIC group.
    - A server port (pNIC or vNIC) can belong to only one vNIC group.
  - The failover mechanism is virtual port aware.

Concurrency has the following key rules:

- vNIC and VMready

  VMready is a unique function that enables the network to be virtual machine aware. The network can be configured and managed at the virtual port (vport) level, rather than just the physical port level.

  Some ancillary functions of VMready that come with vCenter integration work with the Virtual Fabric Adapter in either pNIC or vNIC mode. Running the

Virtual Fabric Adapter in pNIC mode and running VMready are supported concurrently.

For the latest information, check the latest release notes of the BNT Virtual Fabric 10Gb Switch Module firmware at:

http://www.ibm.com/support/fixcentral

- vNIC and switch stacking

  A *stack* is a group of BNT Virtual Fabric 10Gb Switch Modules that work together as a unified system and, therefore, is managed as a single entity.

  The network views the stack as a single switch entity, and the stack can be accessed and managed as a whole by using standard switch IP interfaces. After the stacking links are established, the number of ports that are available in a stack is equal to the total number of remaining ports of all the switches that are part of the stack. The numbers of available IP interfaces, VLANs, trunks, trunk links, and other switch attributes are not aggregated among the switches in a stack. The totals for the stack as a whole are the same as for any single switch that is configured in stand-alone mode.

  The BNT Virtual Fabric Switch supports stacking and Virtual Fabric Mode at the same time with the most recent firmware.

3.2.1 Fibre Channel over Ethernet architecture

A Fibre Channel over Ethernet (FCoE) network, in general, contains servers, lossless Ethernet (DCB-capable) switches, Fibre Channel Forwarders (FCFs) that provide Fibre Channel (FC) fabric services, and storage devices. A traditional FC network might or might not be present. However, in practice, in deployments for the first several years of FCoE, a traditional FC is likely to almost always be present.

Converged Network Adapters

Beginning with the server, Figure 3-3 illustrates a converged network adapter (CNA) in a server.

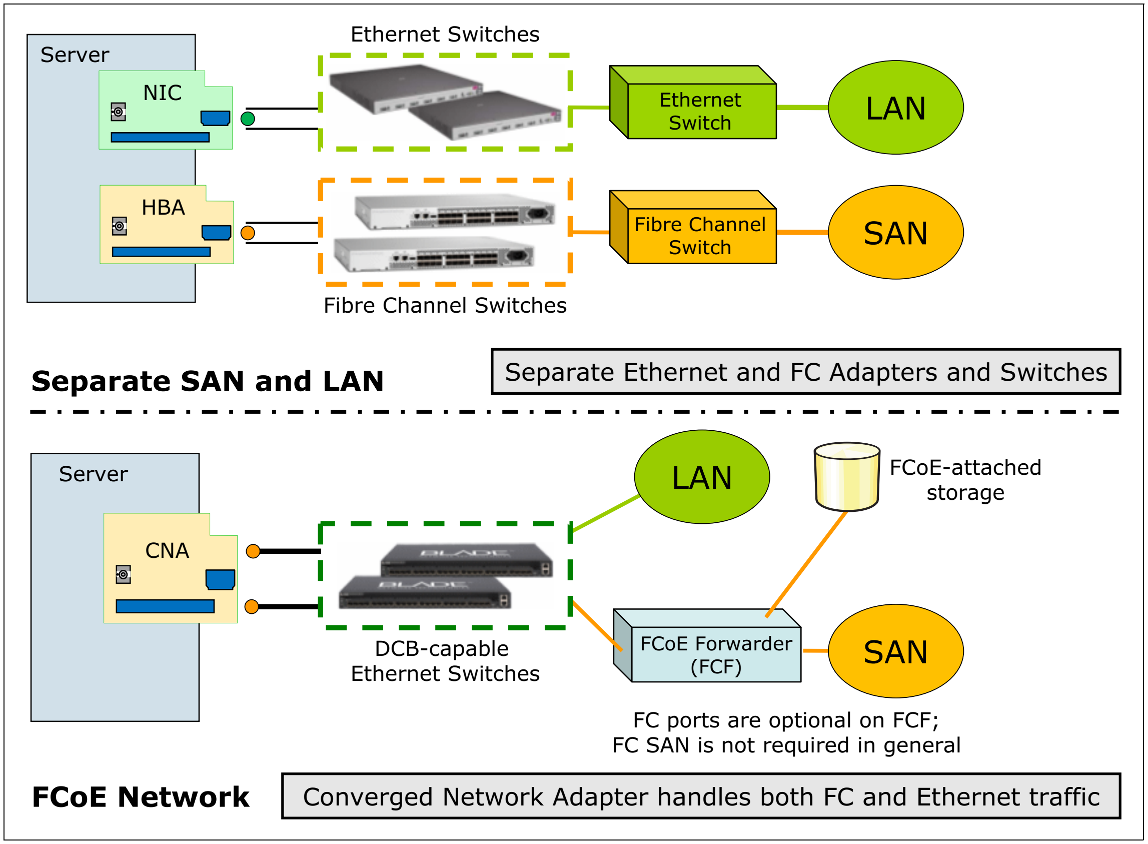

Figure 3-3 Converged Network Adapter (CNA) in a server

In a traditional SAN or LAN environment, the server might have an Ethernet adapter for data traffic and an FC HBA for storage traffic. With FCoE, those two adapters are replaced with a CNA that services both protocol stacks. A single cable from the CNA connects to a lossless Ethernet switch, which (eventually) provides connectivity to an FCF and storage devices.

This consolidation of network adapters, cables, and intermediate switches (DCB-capable Ethernet switches that replace at least some of the intermediate FC switches) provides much of the motivation for FCoE. The reduction in equipment, power, and maintenance costs is anticipated (or at least hoped) to be significant over some time.

Storage devices in an FCoE network might also have CNAs. Few storage devices are implementing them now. Initial deployments typically assume that the storage might remain attached to the traditional FC network in the beginning. Over time as FCoE matures, more storage devices might be expected to support FCoE directly.

3.2.2 Internet Small Computer System Interface architecture

Figure 3-4 gives a conceptual overview of the Internet Small Computer System Interface (iSCSI) protocol layers. As the diagram illustrates, the iSCSI solution requires an initiator (host), a target (generally a storage device), and a carrier network.

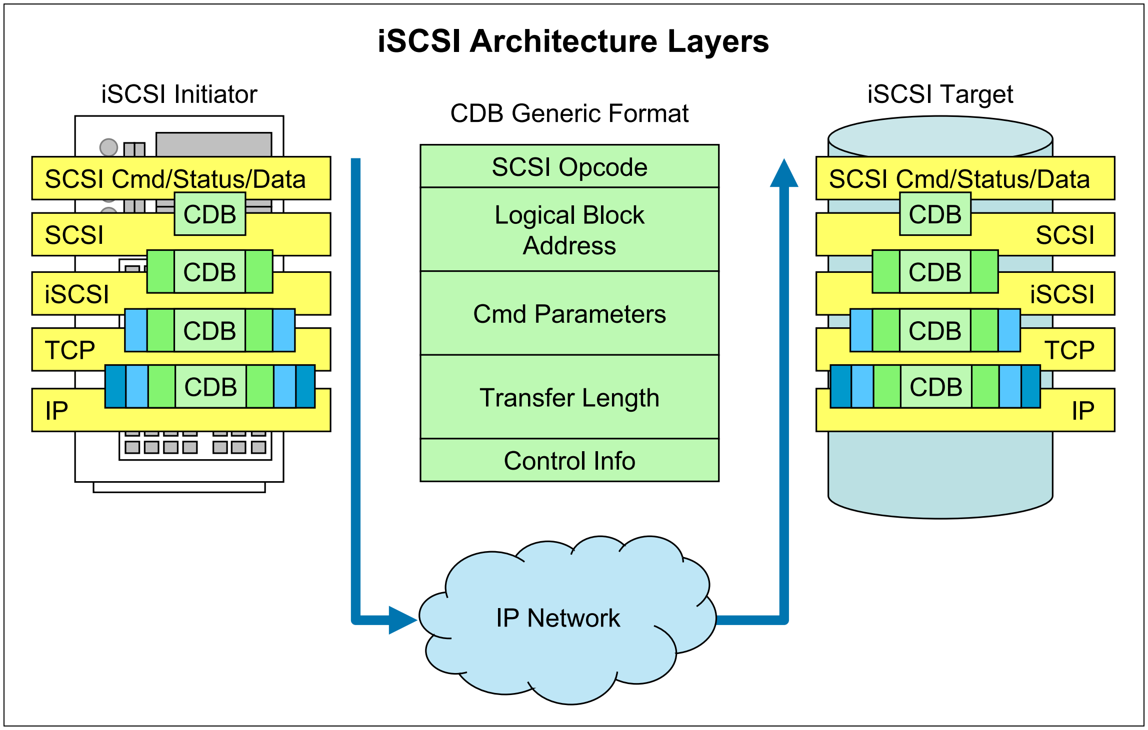

Figure 3-4   iSCSI protocol layers

As mentioned previously, the iSCSI protocol use the TCP/IP protocol to transport iSCSI protocol data units (PDUs), which are the most basic forms of message exchange between the host and the storage controller. The PDU transports both information and SCSI command descriptor blocks (CDBs) between the initiator and target. There, they receive the required data and response, which in turn might require a reply. The PDU also provides ordering and control information.

TCP/IP was chosen for the following reasons:

- It is field proven.
- It can reliably traverse almost any physical network media.
- It can deliver almost error free data that is in order.
- It provides congestion control.
- It acknowledges packets that are received and resends unacknowledged packets.
- The benefits outweighed the use of alternative protocols.
- iSCSI supports SCSI-3 command sets.

Figure 3-5 shows the makeup of an iSCSI PDU and its place within the Ethernet frame.

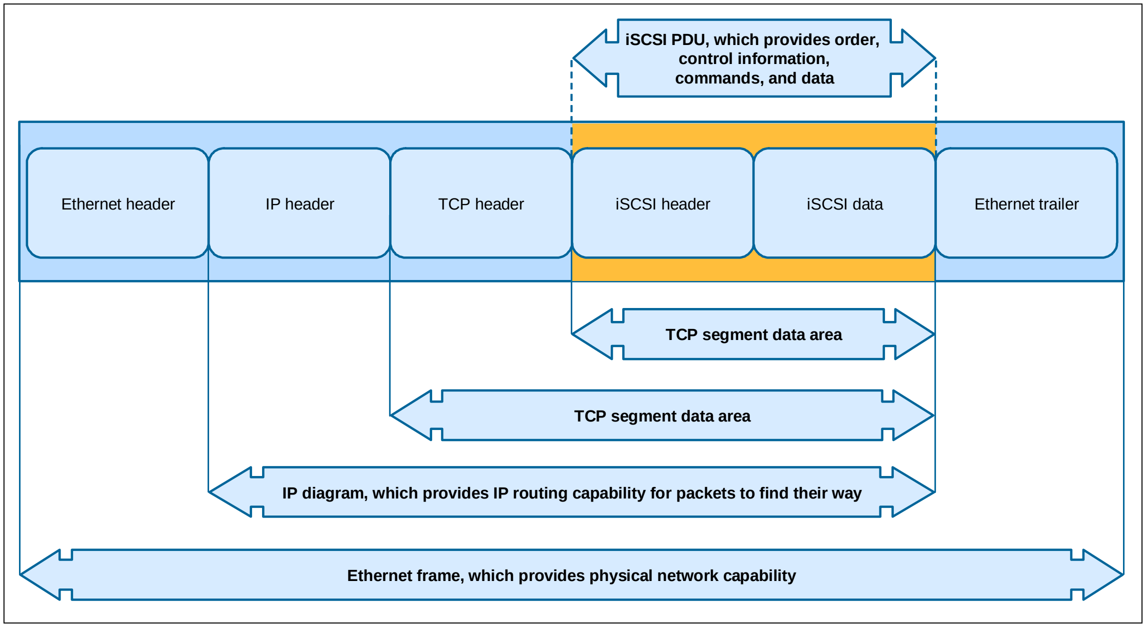

Figure 3-5 iSCSI encapsulation overview

The PDU consists of an iSCSI header where the data length is specified and iSCSI data is encapsulated and transported within the TCP/IP packet. A PDU is not restricted to one TCP segment and can span over more than one TCP segment. Alternatively, it can have more than one iSCSI PDU in a single TCP segment data area. Each TCP segment is encapsulated within an IP datagram. TCPI/IP is responsible for reassembling the TCP segment in the correct order within the datagram on the target side and for delivering the encapsulation to the iSCSI layer in the same order in which it was sent. After arriving at the iSCSI target or initiator, it is opened, and actual iSCSI data is revealed for storage or processing.

The iSCSI protocol works effectively over IP networks, without needing to change the TCP/IP protocol. Figure 3-5 on page 28 shows an overview of iSCSI encapsulation.

3.3 Planning considerations

When planning a Virtual Fabric solution that uses IBM Virtual Fabric Mode, take the following considerations into account:

- The adapter must operate in vNIC mode.
- Up to four vNICs can be configured on each pNIC on the Virtual Fabric Adapter (a total of eight vNICs on a 2-port Virtual Fabric Adapter).
- The default bandwidth to be assigned to a vNIC is 2.5 Gbps, which can be changed only by configuring the Virtual Fabric Switch.
- The maximum bandwidth that can be assigned to a vNIC interface is 10 Gbps.
- The minimum bandwidth that can be assigned to a vNIC depends upon the adapter used:
  - Emulex 10GbE Virtual Fabric Adapter: 100 Mb minimum

    Support for a minimum bandwidth allocation of 100 Mbps has the following requirements:

    - BNT Virtual Fabric 10-Gb Switch Module firmware bnt_fw_bcsw_24-10g-6.8.8.0_anyos_noarch or later

      For more information, see Fix Central at:

      `http://www.ibm.com/support/fixcentral`

    - Emulex 10GbE Virtual Adapters for IBM BladeCenter 4.1.334.47 (a single flash image that contains firmware and boot code)

      For more information, see the Emulex site at:

      `http://www.emulex.com`

    Otherwise, the minimum allocation is 1 Gbps.

  - Emulex 10GbE Virtual Fabric Adapter II: 100 Mb minimum
- The sum of the bandwidth allocated for all four vNICs on the same pNIC cannot exceed 10 Gb.

- Up to 32 vNIC groups are supported on a Virtual Fabric Switch.
- vNICs on the same Virtual Fabric Adapter must be assigned to separate vNIC groups at the Virtual Fabric Switch.
- Only one external port or one trunk made up of external ports is allowed per vNIC group.
- External trunks and ports can belong to, at most, one vNIC group at a time Trunks and ports are not required to belong to a vNIC group.
- Internal ports not connected to a Virtual Fabric Adapter in vNIC mode can belong to, at most, one vNIC group at a time.
- Virtual machines (VMs) and other virtual entities that are associated with vNICs are automatically detected by Virtual Fabric Switch when VMready is enabled. However, these virtual entities cannot be assigned to any of the VMready VM groups, because vNIC groups used by vNICs are fully isolated from other switch entities.
- Currently, BladeCenter Open Fabric Manager is not supported for IBM Virtual Fabric Mode.
- Each vNIC can accommodate one of the following traffic types: regular Ethernet, iSCSI, or FCoE.
- For vNICs on the Emulex Virtual Fabric Adapter, iSCSI and FCoE are mutually exclusive. iSCSI and FCoE *cannot* be used at the same time.
- In terms of vNIC Bandwidth Metering, FCoE vNIC does not use egress metering.

  Enhanced Transmission Selection (ETS) and priority flow control (PFC) must be enabled to ensure lossless transmission for FCoE traffic. ETS does traffic shaping. You can configure a minimum bandwidth for each traffic class, for example, 40% for FCoE priority 3 and 60% for the Ethernet traffic. FCoE traffic gets a 40% minimum guaranteed bandwidth. If the Ethernet traffic uses only 30% bandwidth, FCoE traffic can use 70%. If there is no other Ethernet traffic, FCoE traffic can use 100%. The FCoE vNIC can use up to 100% of the bandwidth, with a minimum guaranteed bandwidth of 40%.
- Each vNIC group can contain traffic of one type only: regular Ethernet, iSCSI, or FCoE. Traffic of different types cannot be mixed within any vNIC group.
- vNICs are not supported simultaneously on the same switch ports as VMready, nor on the same switch ports as DCBX or FCoE.

> **Tip:** All vNIC parameters are configured on the BNT Virtual Fabric 10Gb Switch Module.

4

Supported products

This chapter describes the products that are used in IBM Virtual Fabric Mode. This offering consists of the following supported adapters and switch modules:

- Emulex 10GbE Virtual Fabric Adapter and Virtual Fabric Adapter Advanced for IBM BladeCenter

  These adapters are the first generation Emulex Virtual Fabric Adapter cards for BladeCenter. They are supported only in the IBM Virtual Fabric Mode offering.

- Emulex 10GbE Virtual Fabric Adapter II and Virtual Fabric Adapter II Advanced for IBM BladeCenter

  These expansion cards are the second generation Virtual Fabric Adapters from Emulex. They can be used in IBM Virtual Fabric Mode and in Switch Independent Mode.

- BNT Virtual Fabric 10Gb Switch Module

This chapter includes the following sections:

- Overview of supported components
- Emulex 10Gb Virtual Fabric Adapters
- BNT Virtual Fabric 10Gb Switch Module
- QLogic Virtual Fabric Extension Module

4.1 Overview of supported components

Table 4-1 list the components that are supported by the IBM Virtual Fabric Mode.

Table 4-1 Solution support matrix for IBM Virtual Fabric Mode in BladeCenter

Description	Machine type or part number	Quantity
Chassis		
IBM BladeCenter H	8852	Varies
IBM BladeCenter HT	8740/8750	Varies
Blade servers		
HS22	7870	Varies
HS22V	7871	Varies
HS23[a]	7875	Varies
HX5	7872	Varies
Virtual Fabric Adapters		
Emulex 10GbE Virtual Fabric Adapter	49Y4235	1 per CFFh slot
Emulex 10GbE Virtual Fabric Adapter II	90Y3550	1 per CFFh slot
Virtual Fabric Switches		
BNT Virtual Fabric 10Gb Switch Module (see 4.3, "BNT Virtual Fabric 10Gb Switch Module" on page 39)	46C7191	1 or 2 per chassis
QLogic Virtual Fabric Extension Module (see 4.4, "QLogic Virtual Fabric Extension Module" on page 47)	46M6172	1 or 2 per chassis
External Virtual Fabric Switch connections		
Copper-based		
Small form-factor pluggable plus (SFP+) Copper Direct Attach Cable - 1 M	59Y1936	1 per external port
SFP+ Copper Direct Attach Cable - 3 M	59Y1940	1 per external port
SFP+ Copper Direct Attach Cable - 7 M	59Y1944	1 per external port
Fiber Optics-based		
IBM 10GBase-SR 10GbE 850 nm Fiber SFP+ Transceiver	44W4408	1 per external port

Description	Machine type or part number	Quantity
BLADE 1000Base-SX small form-factor pluggable (SFP)	81Y1622	1 per external port
IBM BNT 1000BASE-T SFP Transceiver	81Y1618	1 per external port

a. An HS23 can have two distinct Emulex cards, each with two ports, and their configurations are independent of each other. One is on the system board, and the other attaches as a mezzanine card.

4.2 Emulex 10Gb Virtual Fabric Adapters

You can use two generations of Emulex Virtual Fabric Adapters in IBM Virtual Fabric Mode:

- First-generation Emulex 10GbE Virtual Fabric Adapter (part number 49Y4235) and Emulex 10GbE Virtual Fabric Adapter Advanced (part number 49Y4275)

  These adapters support IBM Virtual Fabric Mode only. They do not support Switch Independent Mode.

- The Emulex 10GbE Virtual Fabric Adapter II (part number 90Y3550) and Emulex 10GbE Virtual Fabric Adapter II Advanced (part number 90Y3566)

  These adapters support both IBM Virtual Fabric Mode and Switch Independent Mode. These adapters also support Serial over LAN by using 10 Gbps switches and support the *technology preview* of SR-IOV on Red Hat Enterprise Linux 6.

The following sections describe the details of these expansion cards.

4.2.1 Emulex 10GbE Virtual Fabric Adapter

The Emulex 10GbE Virtual Fabric Adapter for IBM BladeCenter is a dual-port 10 Gbps Ethernet card. It supports 1 Gbps or 10 Gbps traffic, or up to eight virtual network interface controller (vNIC) devices. The vNICs are configured to meet a mix of network connectivity and throughput demands for today's complex server application environments. Each physical 10 Gbps port can be divided into four virtual ports with bandwidth allocation in 100 Mbps increments to a maximum of 10 Gbps per physical port. The adapter is a CFFh expansion card. Therefore, it requires switch modules to be installed in I/O bays 7 and 9.

The Emulex 10GbE Virtual Fabric Adapter Advanced supports Fibre Channel over Ethernet (FCoE) and Internet Small Computer System Interface (iSCSI) hardware initiator functions. It also supports the features that are available on the standard card.

Table 4-2 lists the Emulex 10GbE Virtual Fabric Adapter ordering part numbers and feature codes.

Table 4-2 Ordering part numbers and feature codes

Description	Part number	Feature code
Emulex 10GbE Virtual Fabric Adapter	49Y4235	5755
Emulex 10GbE Virtual Fabric Adapter Advanced	49Y4275	2435
Emulex 10GbE Virtual Fabric Adapter Advanced Upgrade	49Y4265	2436

The Emulex 10GbE Virtual Fabric Adapter has the following features:

- Connection to either 1 Gbps or 10 Gbps data center infrastructure (1 Gbps and 10 Gbps auto-negotiation)
- PCI Express 2.0 x8 host interface
- IBM BladeCenter CFFh form factor, which can be combined with a CIOv expansion card on the same blade server
- Operation as either as an 8-port vNIC or as a 2-port 1/10Gb Ethernet adapter:
  - Virtual port bandwidth allocation in 100 Mbps increments
  - Support for up to eight virtual ports
- Wake on LAN support
- Full-duplex (FDX) capability
- Bus-mastering support
- Direct memory access (DMA) support
- Preboot Execution Environment (PXE) support
- IPv4/IPv6 TCP, User Datagram Protocol (UDP) checksum offload
  - Large send offload (LSO)
  - Large receive offload
  - Receive-side scaling (RSS)
  - IPv4 TCP Chimney Offload
- Virtual local area network (VLAN) insertion and extraction
- Jumbo frames up to 9000 bytes
- Load balancing and failover support

  This support includes adapter fault tolerance (AFT), switch fault tolerance (SFT), adaptive load balancing (ALB), teaming support, and IEEE 802.3ad.

- Enhanced Ethernet (draft)
  - Enhanced Transmission Selection (ETS) (P802.1Qaz)
  - Priority-based flow control (PFC) (P802.1Qbb)
  - Data Center Bridging Capability Exchange (DCBX) Protocol, CIN-DCBX and Converged Enhanced Ethernet (CEE) DCBX (P802.1Qaz)
- FCoE and iSCSI support, which is provided by Emulex Virtual Fabric Adapter Advanced
- Configurable adapter from the BNT Virtual Fabric 10Gb Switch Module, which is the basis for the IBM Virtual Fabric Mode offering.

The expansion card has two modes of operation:

- Virtual Fabric vNIC (or vNIC) Mode

  This mode is the basis for Virtual Fabric functions. In vNIC mode, each physical port is visible to the blade server as four vNICs. The default bandwidth for each vNIC is 2.5 Gbps. Bandwidth for each vNIC can be configured from 1 Gbps to 10 Gbps, in 100 Mbps increments. Total bandwidth of all vNICs combined cannot exceed 10 Gbps per physical port. vNICs can also be configured to have zero bandwidth if you must allocate the available bandwidth to fewer than eight vNICs.

  In IBM Virtual Fabric Mode, the expansion card communicates with the BNT Virtual Fabric 10Gb Switch Module, which provides independent control for each vNIC.

- Physical NIC (pNIC) Mode

  In pNIC mode, the expansion card can operate as a standard 10 Gbps or 1 Gbps, 2-port Ethernet expansion card. In pNIC mode, the expansion card functions with any supported 10 Gbps switch installed in I/O module bays 7 and 9.

  In addition, the Emulex Virtual Fabric Adapter Advanced supports iSCSI hardware initiator functions on both physical ports. When the ports operate in pNIC mode, you can use up to two iSCSI hardware initiators. When the ports are configured in vNIC mode, you can use two iSCSI initiators and six vNICs.

  The standard Emulex 10GbE Virtual Fabric Adapter can be upgraded to the same features as the Emulex 10GbE Virtual Fabric Adapter Advanced with the addition of the Advanced Upgrade option, part number 49Y4265.

Figure 4-1 shows the Emulex 10GbE Virtual Fabric Adapter.

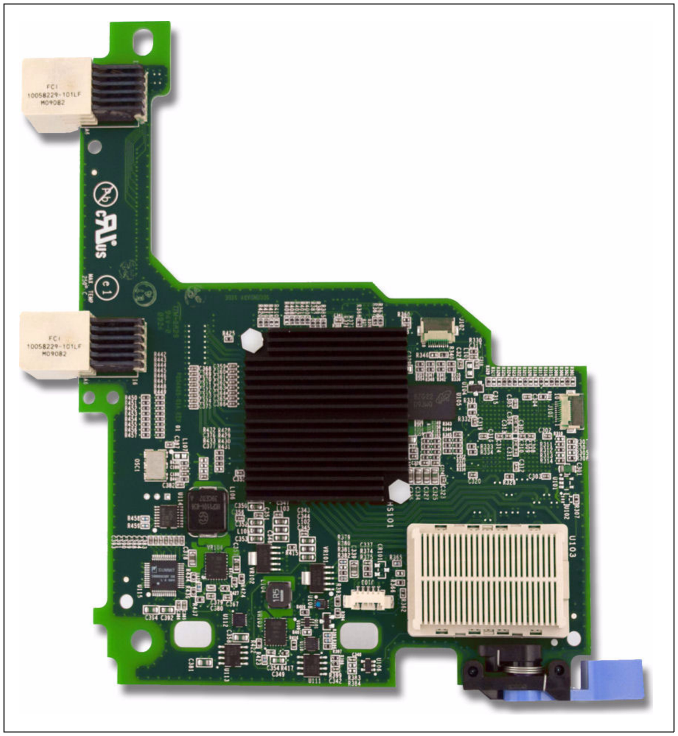

Figure 4-1 Emulex 10GbE Virtual Fabric Adapter

Figure 4-2 shows the placement of the adapter in an HS22 blade server. The adapter is installed in the CFFh expansion slot.

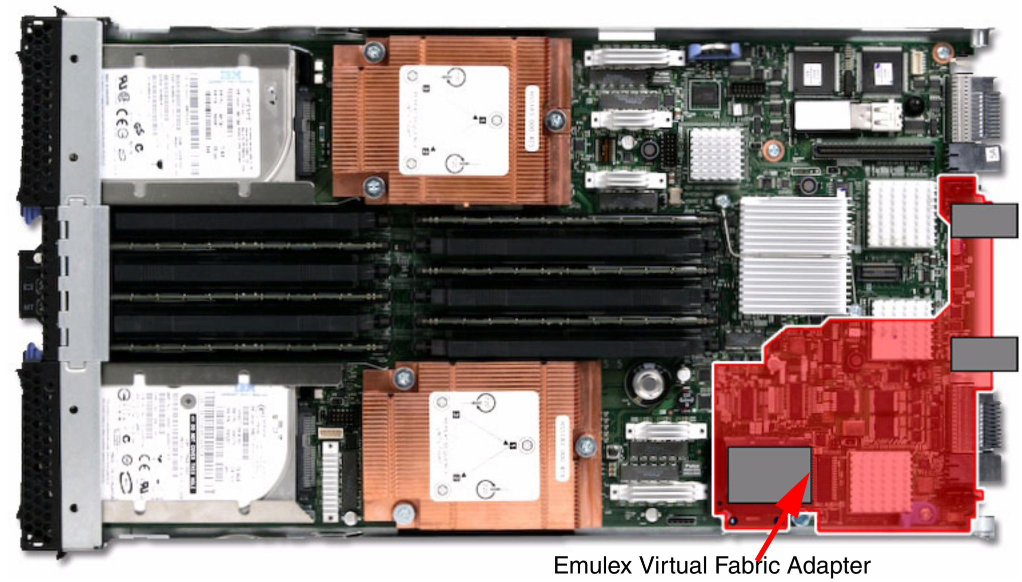

Figure 4-2 Emulex Virtual Fabric Adapter placement in HS22 blade server

For more information, see the *Emulex 10GbE Virtual Fabric Adapter and Virtual Fabric Adapter Advanced for IBM BladeCenter*, TIPS0748.

4.2.2 Emulex 10GbE Virtual Fabric Adapter II

Emulex 10GbE Virtual Fabric Adapter II enhances the existing BladeCenter Virtual Fabric portfolio. Although its basic features are similar to the original Emulex 10GbE Virtual Fabric Adapter, Emulex 10GbE Virtual Fabric Adapter II provides more flexible vNIC support. In addition to IBM Virtual Fabric Mode, it supports Switch Independent Mode. Emulex Virtual Fabric Adapter II Advanced offers FCoE and iSCSI hardware initiator functions, in addition to the features that are available on the standard card.

Figure 4-3 shows Emulex 10GbE Virtual Fabric Adapter II.

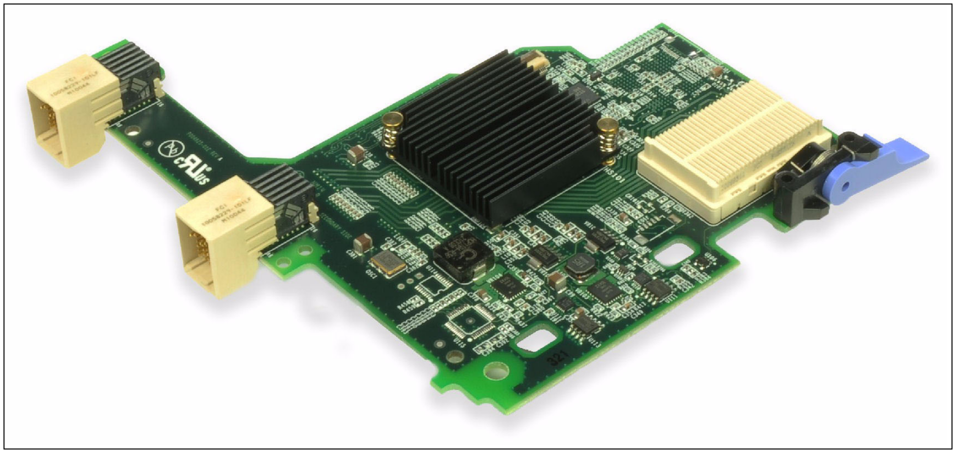

Figure 4-3 Emulex 10GbE Virtual Fabric Adapter II

Table 4-3 lists the part numbers for ordering these expansion cards.

Table 4-3 Part numbers and feature codes for ordering Emulex adapters

Description	Part number	Feature code
Emulex Virtual Fabric Adapter II (CFFh)	90Y3550	A1XG
Emulex Virtual Fabric Adapter II Advanced (CFFh)	90Y3566	A1XH
Emulex Virtual Fabric Adapter Advanced Upgrade	49Y4265	2436

Tip: Emulex Virtual Fabric Adapter Advanced Upgrade is the same part number (49Y4265) for both the Emulex 10GbE Virtual Fabric Adapter and the Emulex 10GbE Virtual Fabric Adapter II.

The original Emulex 10GbE Virtual Fabric Adapter can be configured in pNIC or vNIC mode. The new Emulex 10GbE Virtual Fabric Adapter II now offers two types of vNIC:

- IBM Virtual Fabric Mode, which is the same as vNIC mode on the original Emulex Virtual Fabric Adapters

  In this mode, the Emulex Virtual Fabric Adapter card obtains the vNIC parameters from the switch by using DCBX. This mode works only with the BNT Virtual Fabric 10Gb Switch Module.

- Switch Independent Mode, where the adapter can work with BNT Virtual Fabric 10Gb Switch Module or the Cisco Nexus 4001I Switch Module

Table 4-4 compares the Emulex 10GbE Virtual Fabric Adapter with the Emulex 10GbE Virtual Fabric Adapter II.

Table 4-4 Feature comparison of Emulex Virtual Fabric Adapters

Function	Emulex Virtual Fabric Adapter	Emulex Virtual Fabric Adapter II
pNIC mode	Yes	Yes
IBM Virtual Fabric Mode	Yes	Yes
Switch Independent Mode	No	Yes
iSCSI hardware initiator	Yes (Advanced[a])	Yes (Advanced[a])
FCoE	Yes (Advanced[a])	Yes (Advanced[a])
Serial over LAN	No	Yes
High energy efficiency	No	Yes[b]

a. Available on Emulex Virtual Fabric Adapter II Advanced or with Advanced Upgrade.
b. Emulex Virtual Fabric Adapter II is 20% more energy efficient.

4.3 BNT Virtual Fabric 10Gb Switch Module

The BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter offers the most bandwidth of any blade switch. This switch module also represents the perfect migration platform for clients who are still at 1 Gb outside the chassis by seamlessly integrating in the existing 1 Gbps infrastructure. This switch is the first 10 Gbps switch for IBM BladeCenter that supports converged networking. That is, it supports the ability to transmit CEE to an FCoE-capable, top-of-rack switch. This feature is available with firmware release 6.1 and later.

With the CEE and FCoE functions, you can transfer storage, network, Voice over IP (VoIP), video, and other data over the common Ethernet infrastructure. By using the QLogic Virtual Fabric Extension Module, clients can achieve FCoE gateway functionality inside the BladeCenter chassis.

This Redbooks publication concentrates on virtual fabric. The BNT Virtual Fabric 10Gb Switch Module can be used in IBM Virtual Fabric Mode and Switch Independent Mode. In IBM Virtual Fabric Mode, all vNIC parameter configuration is performed on the BNT Virtual Fabric 10Gb Switch Module. Therefore, the switch module is a key part of the offering. With this offering, clients can form 8 vNICs from one physical NIC and manage them in virtual groups. The switch module can be managed by using a command-line interface (CLI) or web browser interface that provides all of the benefits of I/O virtualization at a 10 Gbps speed.

You might have a chassis with multiple servers of which several operate at 1 Gbps, several operate at 10 Gbps, and several transmit converged packets. In this situation, this single switch can handle all of these workloads and connect to a 1 Gbps infrastructure, to a 10 Gbps infrastructure, or both.

With the extreme flexibility of the IBM System Networking switch, you can take advantage of the technologies that are required for multiple environments. For 1 Gbps uplinks, you can benefit from SFP transceivers. For 10 Gbps uplinks, you can choose SFP+ transceivers (short range or long range for longer distances) or direct-attached copper (DAC) cables (also known as *twinax active cables*) for shorter distances. DAC cables are more cost-effective, use less power, and can be up to 7 meters in length. They are ideal for connecting chassis, connecting to a top-of-rack switch, or connecting to an adjacent rack.

Figure 4-4 shows the switch module.

Figure 4-4 BNT Virtual Fabric 10Gb Switch Module

Table 4-5 lists the part number and feature code to use to order the module.

Table 4-5 BNT Virtual Fabric 10Gb Switch Module part number and feature code

Description	Part number	Feature code
BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter	46C7191	1639

The part number includes the following items:

- One BNT Virtual Fabric 10Gb Switch Module
- Three-meter, mini-USB-to-DB9 serial console cable
- One filler module
- BNT Virtual Fabric 10Gb Switch Module Installation Guide
- BNT user license agreement
- Important Notices document
- Documentation CD-ROM

Tip: SFP+ transceivers are not included. You must purchase them separately.

To communicate outside of the chassis, you must have SFP+ transceivers or SFP+ DAC cables connected. DAC cables have SFP+ transceivers on both ends. You have the flexibility to expand your bandwidth as preferred, by using anywhere from one connection up to ten connections per switch.

Table 4-6 lists the part numbers for ordering the SFP+ transceivers, FC cables, and DAC cables.

Tip: Both 10 Gbps and 1 Gbps SFP+ modules are available for uplink.

Table 4-6 IBM part numbers for ordering SFP+ transceivers, FC cables, and DAC cables

Description	Part number	Feature code
10Gb SFP+		
IBM 10GBase-SR 10GbE 850 nm Fiber SFP+ Transceiver	44W4408	4942
BNT SFP+ Transceiver	46C3447	5053
1Gb SFP+		
BLADE 1000BASE-T (RJ45) SFP Transceiver	81Y1618	3268
BLADE 1000BASE-SX SFP Transceiver	81Y1622	3269
DAC cables		
0.5 m Molex Direct Attach Copper SFP+ Cable	59Y1932	3735
1 m Molex Direct Attach Copper SFP+ Cable	59Y1936	3736
3 m Molex Direct Attach Copper SFP+ Cable	59Y1940	3737
7 m Molex Direct Attach Copper SFP+ Cable	59Y1944	3738

Description	Part number	Feature code
FC cables		
3 m Intel Connects Optical Cable	46D0153	3852
10 m Intel Connects Optical Cable	46D0156	3853
30 m Intel Connects Optical Cable	46D0159	3854

Table 4-7 lists additional transceivers and DAC cable options that are available directly from IBM System Networking.

Table 4-7 BNT part numbers for ordering SFP-based transceivers and cables

Description	Part number
BLADE 1000Base-T SFP	BN-CKM-S-T
BLADE 1000Base-SX SFP	BN-CKM-S-SX
BLADE 10GBase-LR SFP+	BN-CKM-SP-LR
SFP+ Copper Direct Attach Cable - 1 M	BN-SP-CBL-1M
SFP+ Copper Direct Attach Cable - 3 M	BN-SP-CBL-3M
SFP+ Copper Direct Attach Cable - 7 M	BN-SP-CBL-7M

The BNT Virtual Fabric 10Gb Switch Module includes the following features and functions:

- Form-factor
  - Single-height, high-speed switch module
- Internal ports
  - Fourteen internal auto-negotiating ports: 1 Gbps or 10 Gbps to the server blades
  - Two internal, full-duplex 100 Mbps ports connected to the management module
- External ports
  - Up to ten 10 Gbps SFP+ ports (also support 1 Gbps SFP if required, flexibility of mixing 1 Gb/10 Gb)

    Oversubscription ratio (14 internal ports to 10 external ports) is low, which makes the switch module suitable for the most performance-intensive environments.

  - One 10/100/1000-Mb copper RJ45 used for management or data
  - An RS-232 mini-USB connector for a serial port, which provides an additional means to install software and configure the switch module
- Scalability and performance
  - Autosensing 1 Gbps/10 Gbps internal and external Ethernet ports for bandwidth optimization
  - Non-blocking architecture with wire-speed forwarding of traffic and full line rate performance of 480-Gbps full duplex
  - Media access control (MAC) address learning: Automatic update and support for up to 32-Kb MAC addresses
  - Up to 128 IP interfaces per switch
  - Static, EtherChannel, and Link Aggregation Control Protocol (LACP) (IEEE 802.3ad) link aggregation, up to 100 Gb of total bandwidth per switch, up to 18 trunk groups, and up to eight ports per group
  - Support for jumbo frames (up to 12288 bytes)
  - Broadcast and multicast storm control
  - Internet Group Management Protocol (IGMP) snooping for limit flooding of IP multicast traffic (IGMP V1, V2, and V3)
  - Configurable traffic distribution schemes over trunk links, based on source and destination IP addresses, MAC addresses, or both
  - Fast port forwarding and fast uplink convergence for rapid Spanning Tree Protocol (STP) convergence
  - Stacking support (Clients can stack up to 8 BNT Virtual Fabric 10Gb Switch Modules.)
- Availability and redundancy
  - Virtual Router Redundancy Protocol (VRRP) for Layer 3 router redundancy
  - IEEE 802.1D STP for providing Layer 2 redundancy with PVRST+
  - IEEE 802.1s Multiple STP (MSTP) for topology optimization, up to 128 STP instances supported by single switch
  - IEEE 802.1w Rapid STP (RSTP), providing rapid STP convergence for critical delay-sensitive, traffic-like voice, or video
  - Layer 2 trunk failover to support active and standby configurations of network adapter teaming on blades
  - Interchassis redundancy (Layer 2 and Layer 3)

- VLAN support
  - Up to 1024 VLANs supported per switch, VLAN numbers from 1 to 4095, where 4095 is a dedicated VLAN that is used for management module connection only
  - 802.1Q VLAN tagging support on all ports
  - Protocol-based VLANs
- Security
  - VLAN-based, MAC-based, and IP-based access control lists (ACLs)
  - 802.1X port-based authentication
  - Multiple user IDs and passwords
  - User access control
  - Radius, Terminal Access Controller Access-Control System Plus (TACACS+), and Lightweight Directory Access Protocol (LDAP)
- Quality of service (QoS)
  - Up to eight queues per port
  - Support for IEEE 802.1p, IP type of service (ToS)/Differentiated Services Code Point (DSCP), and ACL-based (MAC/IP source and destination addresses, VLANs) traffic classification and processing
  - Traffic shaping and remarking based on defined policies
  - Eight Weighted Round Robin (WRR) priority queues per port for processing qualified traffic
- Layer 3 functions
  - IP forwarding
  - IP filtering with ACLs (up to 4096 ACLs supported)
  - VRRP for router redundancy
  - Support for up to 128 static routes
  - Routing protocol support (Router Information Protocol (RIP) V1, RIP V2, Open Shortest Path First protocol (OSPF)V1, V2, and V3, BGP-4) and up to 1024 entries in a routing table
  - IPv6 routing, which includes static routes and OSPFv3 (requires firmware V6.3 or later)
  - Support for Dynamic Host Configuration Protocol (DHCP) Relay
  - IPv6 host management
  - IPv6 forwarding based on static routes

- Manageability
  - Simple Network Management Protocol (SNMP V1, V2, and V3)
  - HTTP and HTTPS browser-based interface (BBI)
  - Industry standard CLI and IBM Networking OS/AlteonOS CLI
  - Telnet interface for CLI
  - SSH v1/v2
  - Serial interface for CLI
  - Scriptable CLI
  - Firmware image update (Trivial File Transfer Protocol (TFTP) and File Transfer Protocol (FTP))
  - Network Time Protocol (NTP) for switch clock synchronization
  - IBM System Networking BLADEHarmony Manager support
- Monitoring
  - Switch LEDs for external port status and switch module status indication
  - Port mirroring to analyze network traffic that passes through the switch
  - Change tracking and remote logging with syslog feature
  - Power-on self test (POST) diagnostics
- Special functions
  - Serial over LAN
- Virtualization features
  - VMready
  - Virtual Fabric Adapter vNIC support
- Converged Enhanced Ethernet and FCoE features
  - FCoE allows Fibre Channel (FC) traffic to be transported over Ethernet links.
  - FCoE Initialization Protocol snooping enforces point-to-point links for FCoE traffic outside the regular FC topology.
  - PFC (IEEE 802.1Qbb) extends the 802.3x standard flow control to allow the switch to pause traffic, based on the 802.1p priority value in each packet VLAN tag.
  - ETS (IEEE 802.1Qaz) provides a method for allocating link bandwidth, based on the 802.1p priority value in each packet VLAN tag.
  - DCBX (IEEE 802.1AB) allows neighboring network devices to exchange information about their capabilities.

– Supports the QLogic Virtual Fabric Extension Module for IBM BladeCenter, which provides FCoE gateway functions inside the BladeCenter chassis.

VMready is a unique solution that enables the network to be virtual machine (VM)-aware. The network can be configured and managed for virtual ports (vports), rather than just for physical ports. With VMready, as VMs migrate across physical hosts, their network attributes also migrate. VMs can be added, moved, and removed, and yet retain the same ACLs, QoS, and VLAN attributes. VMready allows for a define-once-use-many configuration that evolves as the server and network topologies evolve. VMready works with all virtualization products, including VMware, Hyper-V, Xen, and KVM, without modification of VM hypervisors or guest operating systems. It is available as part of the 6.1 software code (and later).

VMready compatibility with Virtual Fabric solutions is as follows:

- VMready is not supported with IBM Virtual Fabric Mode.
- VMready is supported with Switch Independent Mode.

The switch module supports the following Institute of Electrical and Electronics Engineers (IEEE) standards:

- IEEE 802.1D STP with PVRST+
- IEEE 802.1s MSTP
- IEEE 802.1w RSTP
- IEEE 802.1p Tagged Packets
- IEEE 802.1Q Tagged VLAN (frame tagging on all ports with VLANs enabled)
- IEEE 802.1x port-based authentication
- IEEE 802.2 Logical Link Control
- IEEE 802.3ad Link Aggregation Control Protocol
- IEEE 802.3x Full-duplex Flow Control
- IEEE 802.3ab 1000BASE-T Gigabit Ethernet
- IEEE 802.3ae 10GBASE-SR 10 Gbps Ethernet fiber optics short range
- IEEE 802.3ae 10GBASE-LR 10 Gbps Ethernet fiber optics long range
- IEEE 802.3z 1000BASE-SX Gigabit Ethernet

The following network cables are supported for the BNT Virtual Fabric 10Gb Switch Module:

- 10GBASE-SR for 10 Gbps ports: 850Nm wavelength, multimode fiber, 50μ or 62.5μ (300 meters maximum), with LC duplex connector
- 1000BASE-T for RJ45 port:
  - UTP Category 6 (100 meters maximum)
  - UTP Category 5e (100 meters maximum)
  - UTP Category 5 (100 meters maximum)
  - EIA/TIA-568B 100-ohm STP (100 meters maximum)

For more information, see *BNT Virtual Fabric 10Gb Switch Module At-a-Glance Guide*, TIPS0708.

4.4 QLogic Virtual Fabric Extension Module

The QLogic Virtual Fabric Extension Module for the IBM BladeCenter (Figure 4-5) is another example of how IBM is at the forefront of offering new technology to clients.

Figure 4-5 QLogic Virtual Fabric Extension Module for IBM BladeCenter

IBM was among the first to deliver FCoE across System x and BladeCenter. This new module goes a step further by offering clients I/O convergence inside the chassis. Clients who use the BNT Virtual Fabric 10Gb Switch Module for their LAN traffic can now combine it with this module and a converged network adapter (CNA). For example, they might use the QLogic 2-port 10 Gb CNA. With this combination, they converge their LAN and storage area network (SAN) on a single network. The QLogic Virtual Fabric Extension Module offers six ports of 8 Gb FC connectivity, without the need for separate FC expansion cards in the BladeCenter servers.

Table 4-8 lists the part numbers for ordering the QLogic Virtual Fabric Extension Module.

Table 4-8 Part number information

Description	Part number	Feature code
QLogic Virtual Fabric Extension Module	46M6172	4799

The part number includes the following items:

- One QLogic Virtual Fabric Extension Module for IBM BladeCenter
- Support CD
- The IBM Important Notices document
- Warranty information

The QLogic Virtual Fabric Extension Module comes without SFP+ modules, which you must order additionally. Table 4-9 lists the part number that is supported.

Table 4-9 Supported SFP+ for the QLogic Virtual Fabric Extension Module

Description	Part number	Feature code
IBM 8 Gb SFP+ SW Optical Transceiver	44X1964	5075

The QLogic Virtual Fabric Extension Module for IBM BladeCenter has the following features:

- Standard I/O module form factor
- Six external autosensing FC ports that operate at 8 Gbps, 4 Gbps, or 2 Gbps
- External ports can be configured as Full Fabric (GL, G, F, FL, E) FC ports or Transparent Fabric (TF) FC ports
- Up to 40 Gbps of internal bandwidth to the switch module
- Two internal full-duplex 100 Mbps Ethernet interfaces for management
- POST diagnostics and status reporting
- Support for Non-Disruptive Code Load Activation (NDCLA)
- Registered State Change Notification (RSCN)
- Support for standards-based FC-SW2 interoperability
- Support for transparent mode (N_Port ID Virtualization (NPIV))
- Error detection
  - Cyclic redundancy check (CRC)
  - 8-byte and 10-byte conversion

  - Parity
  - Long frame and short frame
  - D_ID mismatch
  - S_ID mismatch
- Frame bundling
  - No frame bundling

    Frames are intermixed from different source ports.
  - Soft lockdown

    The I/O module waits for the sequence to be completed or a gap in the frame traffic to occur before it services requests from a different port.
- Configurable Fabric Address Notification (FAN)
- Support for up to 239 Fabric Extension Modules depending on the configuration
- 8 Gb switch fabric aggregate bandwidth: 224 Gbps at full duplex
- Maximum frame size: 2148 bytes (2112 byte payload)
- Nonblocking architecture to prevent latency
- Support for the Call Home function
- Support for Domain Name Service (DNS)
- Support for IP Version 6
- Support for Internet Protocol security (IPsec)
- Support for separate trap community strings for each trap address
- Support for SNMP Version 3
- Support for vital product data (VPD)
- Support for optional SFP+ modules

The QuickTools web interface software feature comes with the switch module. The switch supports the following fabric management. All management connections go through the management module:

- Web interface through QuickTools
- CLI through the Telnet program
- Switch SNMP agent

  By using this agent, a network management workstation can receive configuration values, traffic information, and FC failure data through SNMP and the Ethernet interface.

For more information about QLogic configurations and support, see *QLogic Virtual Fabric Extension Module for IBM BladeCenter*, TIPS0717.

5

Configuring the components

This chapter explains how to configure the components to operate in IBM Virtual Fabric Mode. To configure the components to operate in Switch Independent Mode, see Chapter 9, “Configuring the components” on page 195.

When configuring the solution in this mode, you configure the virtual network interface controller (vNIC) parameters on the BNT Virtual Fabric 10Gb Switch Module. The vNIC parameters might include bandwidth allocation, vNIC groups, virtual local area networks (VLANs), and Converged Enhanced Ethernet (CEE). The switch transfers the configuration settings to the Emulex 10GbE Virtual Fabric Adapter by using the Data Center Bridging Exchange (DCBX) Protocol. The adapter requires little configuration. You need only to enable vNIC mode in the Unified Extensible Firmware Interface (UEFI) utility of the blade server and the required personality of the adapter within the management utility. Although most of the configuration is performed on the switch, the adapter also requires some configuration changes.

The BNT Virtual Fabric 10Gb Switch Module can be configured by using any of the following tools:

- IBM System Networking OS command-line interface
- Industry-standard CLI (isCLI)
- Browser-based interface (BBI)

This chapter includes the following sections:

- Configuring the Emulex 10GbE Virtual Fabric Adapter
- Configuring the BNT Virtual Fabric 10Gb Switch Module
- Creating vNIC in IBM Virtual Fabric Mode
- Configuring the vNIC group
- High availability and vNICs
- Configuring vNIC VLANs
- Configuring the operating system
- vNIC and VMready

5.1 Configuring the Emulex 10GbE Virtual Fabric Adapter

This section describes how to upgrade firmware and configure the Emulex 10GbE Virtual Fabric Adapter cards.

Emulex 10GbE Virtual Fabric Adapter II: This section illustrates the process by using the new Emulex 10GbE Virtual Fabric Adapter II. The procedures are similar on the older Emulex 10GbE Virtual Fabric Adapter cards, but with minor differences.

5.1.1 Upgrading firmware

Before you begin the setup, upgrade the firmware on the Emulex 10GbE Virtual Fabric Adapter II to the latest level. You can upgrade firmware in the following ways:

- Windows Server 2008, 2008 R2 and 2003 online update
- Windows Preinstallation Environment (PE) offline update
- Linux online code update
- Linux offline code update
- VMware ESX code update package
- Code update by using a bootable DVD media
- UEFI code update

The upgrade procedure is documented in the readme files for these products.

Tip: Use UpdateXpress Systems Pack Installer (UXSPI) and Bootable Media Creator, which are the standard means to update firmware.

The following procedure uses the Windows online package to upgrade the firmware. The package uses the `elxflash.exe` file to perform the update.

> **Tip:** When using the `elxflash.exe` file for the firmware update, you might need to set some parameters before you do the update. For more information about which parameters require preliminary setup, see the readme file for the executable file.

To upgrade the firmware:

1. Enter the **elxflash /update** command.
2. Verify that no error messages are displayed. Figure 5-1 shows an example where the firmware updated successfully.

```
>elxflash /update

Thu May 19 15:41:27 2011
HBA=OCe11100-NIC, Port Type=NIC, MAC=00-00-C9-B2-57-02,
Update=Firmware, Image=oc11-4.0.200.10.ufi, New=4.0.200.10,
Old=3.103.351.10, Status=Success

Thu May 19 15:41:27 2011
HBA=OCe11100-NIC, Port Type=NIC, MAC=00-00-C9-B2-57-06,
Update=Firmware, Image=oc11-4.0.200.10.ufi, New=4.0.200.10,
Old=3.103.351.10, Status=Success

(lines deleted for clarity)

Thu May 19 15:41:28 2011
HBA=OCe11100-NIC, Port Type=NIC, MAC=00-00-C9-B2-57-09,
Update=Firmware, Image=oc11-4.0.200.10.ufi, New=4.0.200.10,
Old=3.103.351.10, Status=Success

elxflash: All required updates succeeded - Return Code=0
```

Figure 5-1 Windows online code update by using the elxflash executable

The output is truncated here for clarity. The command produces eight sections, three of which are shown in Figure 5-1. Each section represents one vNIC.

5.1.2 Configuring the adapter in the UEFI

> **Important:** The instructions in this section are based on firmware version 4.0.1061.0 for the Emulex 10GbE Virtual Fabric Adapter II. This firmware was chosen to give you a preview of the UEFI layout.

To configure the Emulex 10GbE Virtual Fabric Adapter II by using the blade server UEFI:

1. Power on or restart the blade.
2. When prompted during POST, press F1 to display the System Configuration and Boot Management panel (also known as the *UEFI utility main menu*, see Figure 5-2).

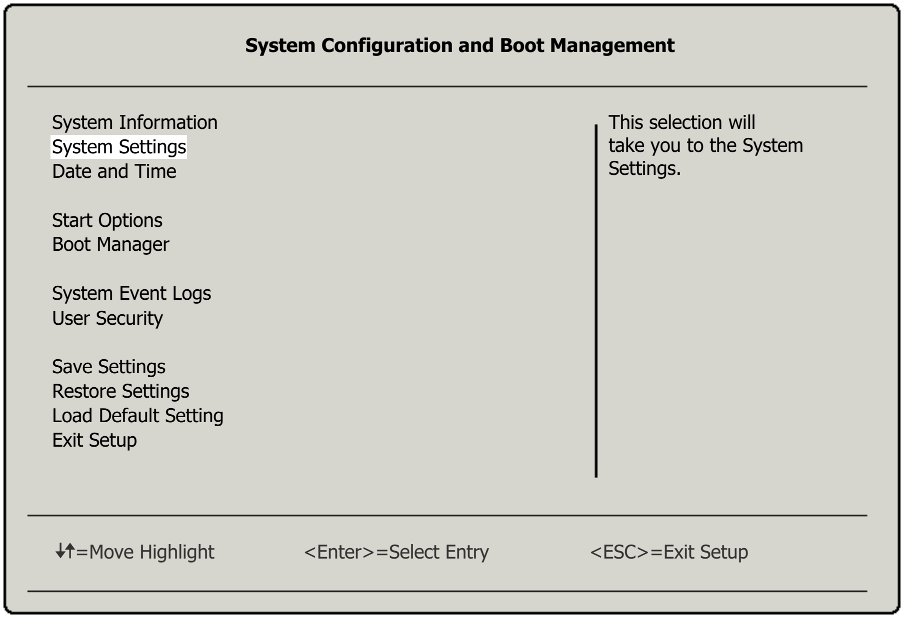

Figure 5-2 Blade UEFI utility main menu

3. Select **System Settings**.
4. In the next panel, select **Network**

5. In the Network configuration panel (Figure 5-3), which shows the Emulex 10GbE Virtual Fabric Adapter II ports as the two physical ports on the adapter, select the first port.

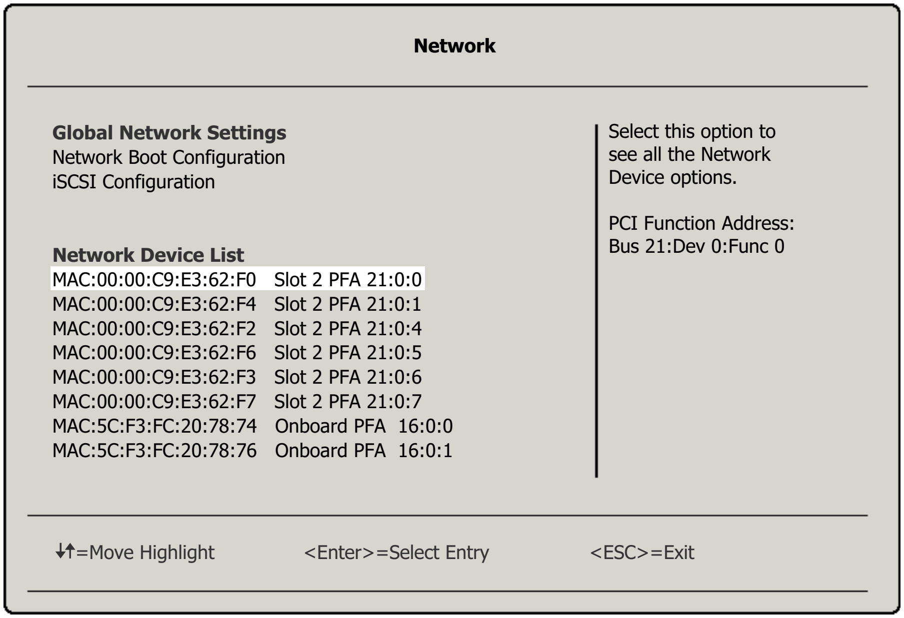

Figure 5-3 Network configuration panel in the blade UEFI utility

6. To enable IBM Virtual Fabric Mode (`vNIC`), enable Multichannel first.

   Unless Multichannel is enabled, the card operates as two physical 10 Gbps Ethernet ports. When you enable Multichannel, it is enabled on both 10 Gbps physical ports. For information about enabling Multichannel, see 10.3.1, “Configuring the Emulex adapter” on page 275.

Tips:

- Switch Independent Mode is not available on the older Emulex 10GbE Virtual Fabric Adapter.
- SR-IOV is disabled when using Switch Independent Mode.
- For more information, see “Switch Independent Mode” on page 153.

7. With Multichannel enabled, select **Switch Configuration**. In this example, of using IBM Virtual Fabric Mode, two options are available: Switch Independent Mode or IBM Virtual Fabric Mode (Figure 5-4).

Hint: When enabling Multichannel configuration, you might be prompted to reboot your system.

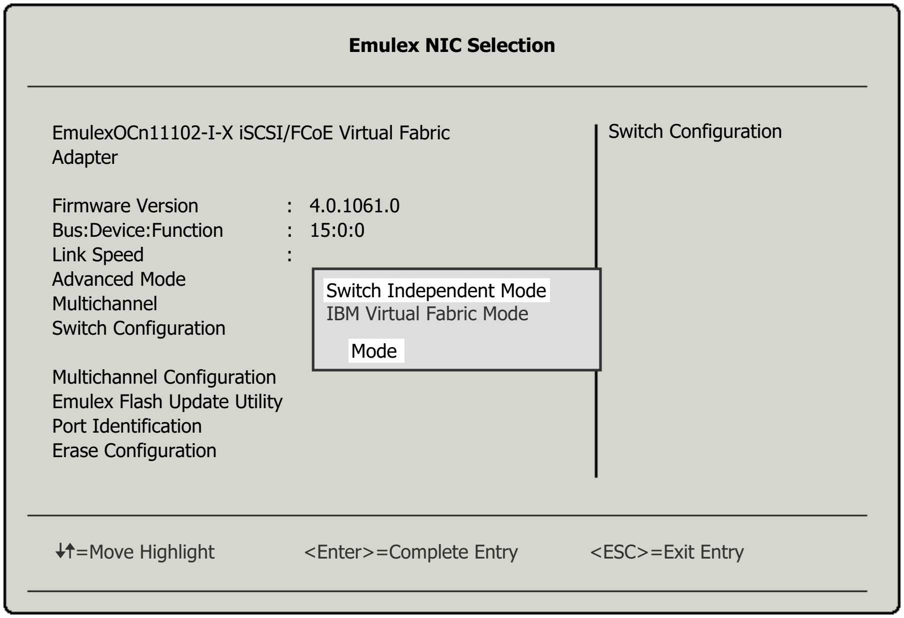

Figure 5-4 Switch Configuration mode selection

8. Select **Multichannel Configuration.**
9. In the Function Configuration panel (Figure 5-5), which lists the vNIC settings, where four functions represent four vNICs, select **Function 0**.

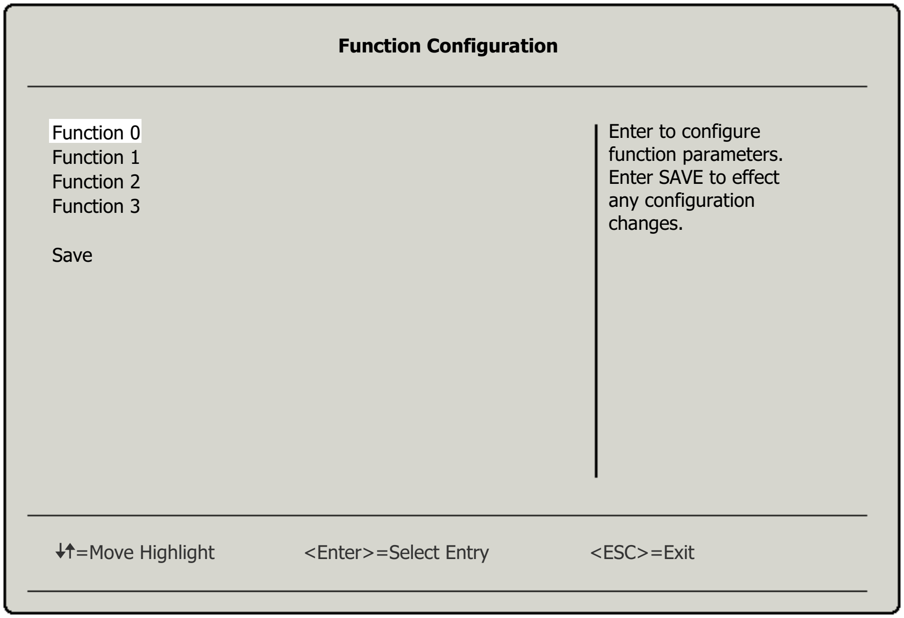

Figure 5-5 Function Configuration panel

10. In the Multichannel Configuration panel, configure the Logical Port VLAN ID and Bandwidth parameters for the vNIC. For more information about these parameters, see “Logical Port VLAN ID parameter” on page 58 and “Bandwidth parameter” on page 58.

    Set the LPVID and Bandwidth values for the remaining vNICs (Functions 1, 2, and 3), and then configure the vNICs on the second physical port. Finally, exit to the main UEFI menu, and save the settings.

11. Save the settings after making changes.

    When Multichannel is enabled on either of the physical ports, it is enabled on both physical ports and can also be disabled on both ports. Review the vNIC settings for the second physical port to ensure that the configuration is correct.

12. At the UEFI main menu (Figure 5-6), select **Save Settings**, and restart the blade server.

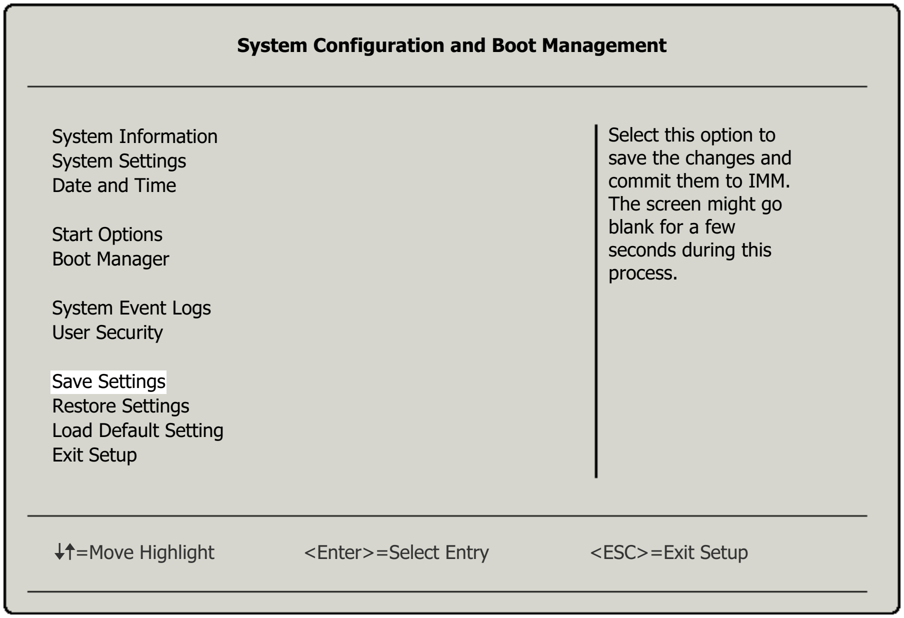

Figure 5-6 Save settings and exit the setup

Logical Port VLAN ID parameter

In IBM Virtual Fabric Mode, the Multichannel Configuration panel shows only one parameter, *Logical Port VLAN ID* (LPVID). Configuration of this parameter is optional because vNIC functions do not depend on it. LPVID is the inner VLAN ID in double-tagged frames. The outer VLAN ID and bandwidth parameters come from the BNT Virtual Fabric 10Gb Switch Module. However, if you are configuring Switch Independent Mode, LPVID is a required parameter.

When configuring for Switch Independent Mode, you must configure the PLVID parameter. The adapter uses LPVID values to isolate traffic within each individual vNIC. For network traffic to pass correctly, on the blade-facing port of the Virtual Fabric Switch Module, allow all of the VLAN IDs that are specified on the four vNIC functions.

Bandwidth parameter

The Bandwidth parameter is another important parameter for the vNIC bandwidth allocation. You can enter values 0 - 100 in this field. The value is the percentage of the total bandwidth that is allocated to a particular vNIC.

5.2 Configuring the BNT Virtual Fabric 10Gb Switch Module

This section provides an overview of the tools available for configuring the BNT Virtual Fabric 10Gb Switch Module. This switch module can be configured by using the following methods:

- A browser-based interface (BBI)
- A CLI (IBM System Networking OS or isCLI)
- Simple Network Management Protocol (SNMP)

Future plans include management of switch modules from BladeCenter Open Fabric Manager.

The examples in this section use BNT Virtual Fabric 10Gb Switch Module firmware version 6.7.5.0, which was the current version at the time this book was written.

5.2.1 IBM System Networking OS CLI

IBM System Networking OS CLI is a simple and intuitive, menu-based user interface that provides complete switch module administration. To use the IBM System Networking OS CLI on the switch module, use the following communication methods:

- Telnet session
- Serial connection
- SSH connection

For more information about IBM System Networking OS CLI commands, see *BLADEOS 6.5 Menu-Based CLI Command Reference BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter*, at:

`http://www.bladenetwork.net/userfiles/file/PDFs/VFSM_CR_6-5.pdf`

Logging in to the IBM System Networking OS CLI

IBM System Networking OS CLI consists of a main menu and hierarchically organized submenus. To log in to the switch module by using **`telnet`**:

1. Use the **`telnet`** command to connect to the switch module with IP address 9.42.171.39:

   ```
   telnet 9.42.171.39
   ```

2. Type the password. The default password is `admin`.

The IBM System Networking OS CLI main menu is displayed and serves as the starting point for all switch module management and configuration tasks (Figure 5-7).

```
[Main Menu]

May 25 14:01:06 9.42.171.39 NOTICE  mgmt: admin(admin) login from host
9.44.16.209
     info     - Information Menu
     stats    - Statistics Menu
     cfg      - Configuration Menu
     oper     - Operations Command Menu
     boot     - Boot Options Menu
     maint    - Maintenance Menu
     diff     - Show pending config changes  [global command]
     apply    - Apply pending config changes [global command]
     save     - Save updated config to FLASH [global command]
     revert   - Revert pending or applied changes [global command]
     exit     - Exit  [global command, always available]

>> Main#
```

Figure 5-7 IBM System Networking OS CLI main menu

Upgrading the firmware

This section guides you through an example of upgrading the BNT Virtual Fabric 10Gb Switch Module firmware by using the IBM System Networking OS CLI. At the time this book was written, firmware 6.7.5.0 was the current level. Therefore, this procedure upgrades the switch module to this firmware version.

Requirement: BNT Virtual Fabric 10Gb Switch Module firmware 6.7.5.0 is now available and is required for the Emulex 10GbE Virtual Fabric Adapter II. We used it in our testing for this book. This firmware level is also required if the Emulex 10GbE Virtual Fabric Adapter is updated to the latest firmware version.

The firmware update package is available from the IBM support website at:

http://www.ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5088963

The firmware code consists of the following parts:

- Boot image
- Operating system (OS) image 1
- OS image 2

When upgrading firmware, upgrade the boot image and one of the two OS images. Consider leaving the other OS image at the existing firmware level, so that you can easily activate it if any problems arise with the new level.

You can use any Trivial File Transfer Protocol (TFTP) or FTP server to transfer the boot and OS images to the BNT Virtual Fabric 10Gb Switch Module. The example uses the TFTP server that is available on the BladeCenter advanced management module (AMM).

From the firmware upgrade package, we moved the following two files into the `tftproot` directory in AMM local storage:

- `GbESM-24-10G-6.7.5.0_Boot.img`, which is the new boot image
- `GbESM-24-10G-6.7.5.0_OS.img`, which is the new OS image

Figure 5-8 shows the contents of the `tftproot` directory in AMM local storage.

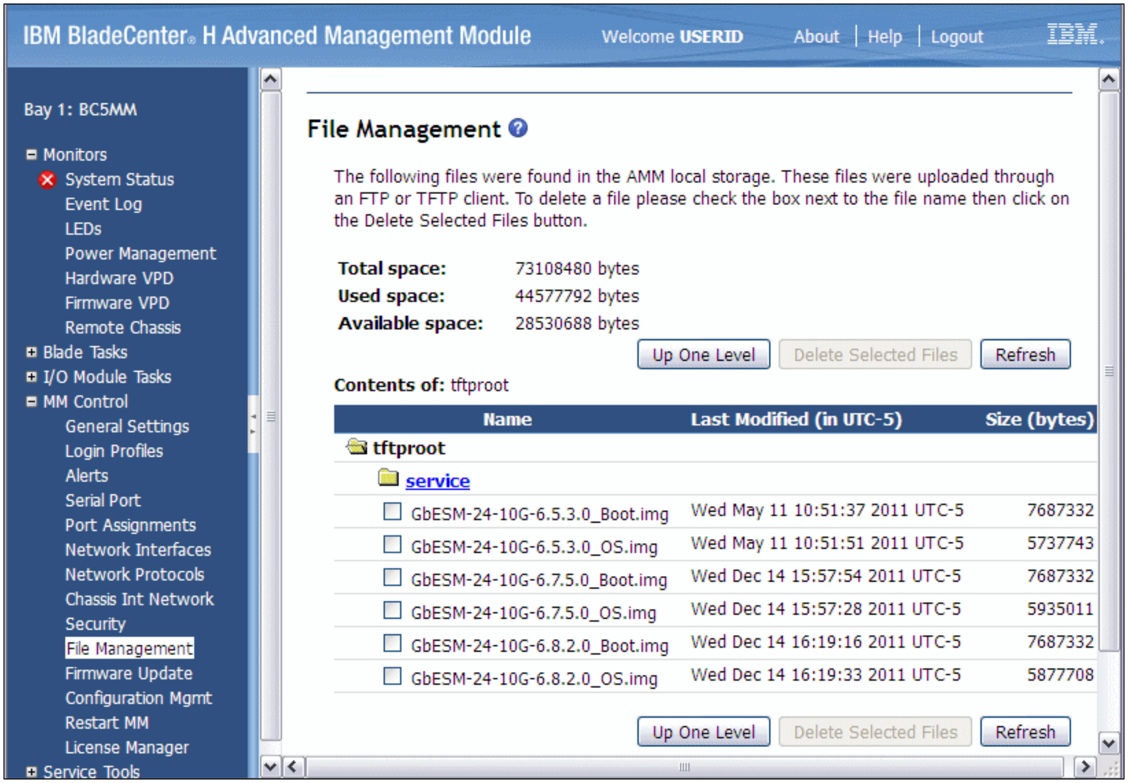

Figure 5-8 Firmware files in the tftproot directory

After ensuring that TFTP server functions are enabled on the AMM, proceed with the firmware upgrade by using the IBM System Networking OS CLI.

Important: The following steps use BNT Virtual Fabric 10Gb Switch Module firmware version 6.8.2.0 to give you a preview of the GUI and the IBM System Networking OS CLI layout.

Upgrading the OS image

To upgrade the OS image:

1. Transfer the new OS image to one of the image banks by using the following command:

   ```
   /boot/gtimg X TADDR GbESM-24-10G-6.7.5.0_OS.img
   ```

   Where *X* represents the OS image bank (1 or 2) and TADDR represents the TFTP server IP address.

   In this example, by using image bank 2, we enter the following command:

   ```
   /boot/gtimg 2 9.42.171.73 GbESM-24-10G-6.7.5.0_OS.img
   ```

2. Verify that no error messages are displayed. The results in Figure 5-9 on page 63 confirm a successful update of OS image 2 to the new firmware version, version 6.7.5.0. Before the upgrade, the switch module was set to boot from image bank 1. Notice in Figure 5-9 that the switch module changed to image bank 2 during the update process.

```
>> Main# /boot/gtimg 2 9.42.171.73 GbESM-24-10G-6.7.5.0_OS.img
Enter username for FTP server or hit return for TFTP server:
image2 currently contains Software Version 6.1.2
 that was downloaded at  6:11:38 Thu Jan  5, 2000.
New download will replace image2 with file "GbESM-24-10G-6.7.5.0_OS.img"
 from FTP/TFTP server 9.42.171.73.
Confirm download operation [y/n]: y
Starting download...
File appears valid
Download in progress
.................................................................
.................................................................
................................................................
Image download complete (5737743 bytes)
Writing to flash...This takes about 90 seconds. Please wait
Write complete (5737743 bytes), now verifying FLASH...
Verification of new image2 in FLASH successful.
image2 now contains Software Version 6.7.5
AMM CodeLevel 3 : 0605WMJ03000   IBMNOS Im2 11/17/2010
Switch is currently set to boot software image1.
Do you want to change that to the new image2? [y/n]
May 25 14:57:28 9.42.171.39 INFO    mgmt: image2 downloaded from host
9.42.171.73, file 'GbESM-24-10G-6.7.5.0_OS.img', software version 6.7.5
y
Next boot will use new software image2.
>>
May 25 14:57:46 9.42.171.39 NOTICE  mgmt: boot image changed
Boot Options#
May 25 14:57:46 9.42.171.39 INFO    mgmt: Firmware downloaded to image2
```

Figure 5-9 Upgrade of OS image 2

Upgrading the boot image

To begin the transfer:

1. Transfer the new boot image by using the following command:

   `/boot/gtimg boot TADDR GbESM-24-10G-6.7.5.0_Boot.img`

   Where *TADDR* represents the IP address of the TFTP server.

   In this example, we enter the following command:

   `/boot/gtimg boot 9.42.171.73 GbESM-24-10G-6.7.5.0_Boot.img`

2. Verify that the new boot manager image was installed without any errors. Also verify that `GbESM-24-10G-6.7.5.0 Boot.img` software version 6.7.5 was downloaded to the kernel (Figure 5-10).

```
>> Boot Options# /boot/gtimg boot 9.42.171.73 GbESM-24-10G-
6.7.5.0_Boot.img
Enter username for FTP server or hit return for TFTP server:
boot kernel currently contains Software Version 6.1.2
New download will replace boot kernel with file "GbESM-24-10G-
6.7.5.0_Boot.img" from FTP/TFTP server 9.42.171.73.
Confirm download operation [y/n]: y
Starting download...
File appears valid
Download in progress
.................................................................
.................................................................
.................................................................
Boot image (FS, 7687332 bytes) download complete.
Writing to flash...This can take up to 90 seconds. Please wait
FS Sector now contains Software Version 6.7.5

Boot image (Kernel, 7687332 bytes) download complete.
Writing to flash...This can take up to 90 seconds. Please wait
Kernel Sector now contains Software Version 6.7.5

Boot image (Boot, 7687332 bytes) download complete.
Writing to flash...This can take up to 90 seconds. Please wait
Boot Sector now contains Software Version 6.7.5

>>May 25 15:13:17 9.42.171.39 INFO    mgmt: boot kernel downloaded from
host 9.42.171.73, file 'GbESM-24-10G-6.7.5.0_Boot.img', software version
6.7.5

Boot Options#
May 25 15:13:17 9.42.171.39 INFO    mgmt: Firmware downloaded to boot
kernel
```

Figure 5-10 Confirmation of the boot image upgrade

Resetting the switch module

To reset the switch module, enter the following command:

```
/boot/reset
```

5.2.2 isCLI

The IBM System Networking OS CLI is a proprietary CLI environment that is developed by IBM to configure and manage its switch modules. System administrators who are more familiar with the Cisco CLI have the option of using the isCLI. isCLI is a Cisco-like user interface that allows complete management and configuration support of BNT switch modules. The IBM System Networking OS CLI is the default CLI mode, but you can easily change between the IBM System Networking OS CLI and isCLI.

For more information about isCLI commands, see *BLADEOS 6.5 ISCLI–Industry Standard CLI Command Reference BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter* at:

http://www.bladenetwork.net/userfiles/file/IBM%20BladeCenter/VFSM_IS_6-5.pdf

isCLI privilege levels

The isCLI mode has three privilege levels:

- User EXEC mode

  This initial mode has the most limited privileges. You can perform basic tasks, such as listing system information, but you cannot set or change operating parameters. With this privilege level set, the isCLI prompt ends with the greater than symbol (>), as in the following example:

  ```
  Switch>
  ```

- Privileged EXEC mode

  This mode has a higher privilege level. By using this mode, you can set or change operating parameters of the switch module. With this privilege level set, the isCLI prompt ends with the number sign (#), as in the following example:

  ```
  Switch#
  ```

  Privileged EXEC mode is accessed from User EXEC mode by using the following command:

  ```
  enable
  ```

- Global Configuration Mode

  This mode is the most privileged mode. When in this mode, you can change the running configuration, and set or change global parameters.

  Use the following command to enter this mode from Privileged EXEC mode:

  ```
  configure terminal
  ```

To change the privilege level, enter the commands shown in Figure 5-11. This example shows how to enable Privileged EXEC and then Global Configuration Mode.

```
Router>enable

Enable privilege granted.

Router#configure terminal

Enter configuration commands, one per line.  End with Ctrl/Z.

Router(config)#
```

Figure 5-11   Enabling EXEC privileges in isCLI

To enable a different privilege level, see “isCLI privilege levels” on page 65.

Changing between the IBM System Networking OS CLI and isCLI

To change from the IBM System Networking OS CLI to isCLI, enter the following command:

```
boot/mode iscli
```

Restart the switch module after this command is run.

To revert from isCLI to the IBM System Networking OS CLI, enter the following command:

```
boot cli-mode IBMNOS-cli
```

Restart the switch module after the command is run.

The selected CLI mode is preserved across reboots, even when you restore the switch module to factory defaults.

5.2.3 Browser-based interface

You can use the BBI to perform most management, administration, and configuration tasks on the BNT Virtual Fabric 10Gb Switch Module. BBI is an easy-to-use and intuitive web-based management interface. It is enabled by default, but can be disabled by using the CLI.

For more information about BBI commands, see the following web pages:

- Publications for the BNT Virtual Fabric 10Gb Switch Module - IBM BladeCenter H, HT

  http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5080917

- IBM BladeCenter support

  http://www.bladenetwork.net/IBM-BladeCenter-support.html

- *BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter Application Guide*

  http://www.bladenetwork.net/userfiles/file/VFSM_AG_6-8.pdf

Starting the BBI

To log in to the switch module and start BBI:

1. Open a web browser, and connect to the switch module IP address. Alternatively, you can access BBI from the AMM web interface:

   a. In the left panel of the AMM web interface, select **I/O Module Tasks** → **Configuration**.

   b. Select the appropriate slot number of the I/O device that you want to access, and then select **Advanced Options** → **Start CLI/Web Session**.

c. In the Start CLI/Web Session (Figure 5-12), verify that the Protocol field is set to **Web** (not Telnet), and then click **Start Session**.

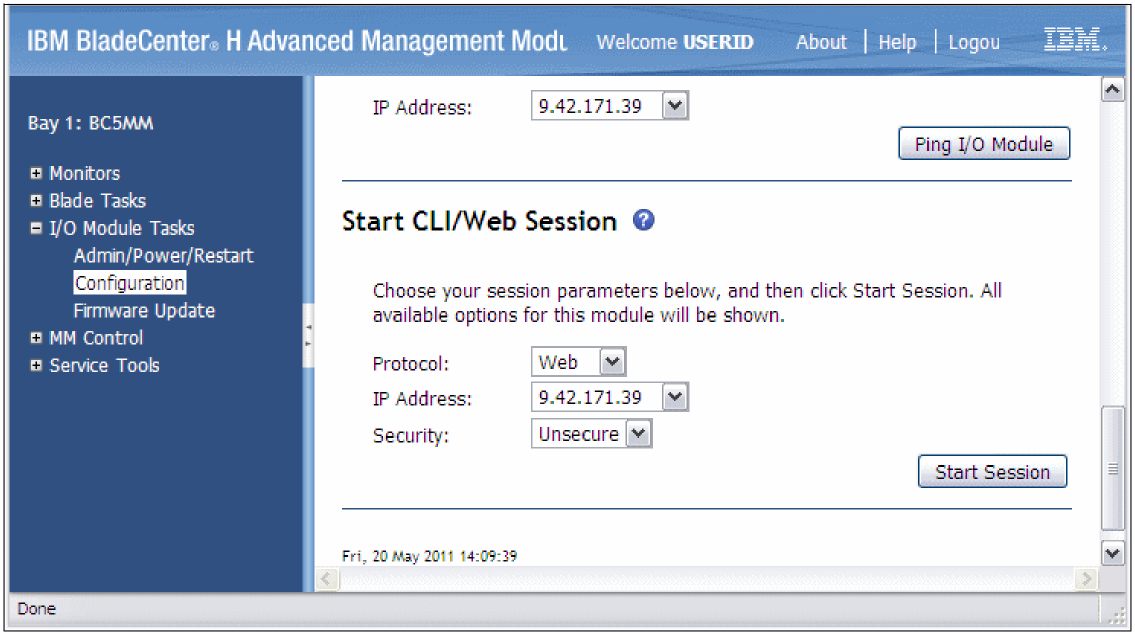

Figure 5-12 AMM web interface: Starting the CLI web session

2. Log in to the switch module, which has the following default credentials:

   – Username: `admin`
   – Password: `admin`

   > **Security:** For added security, change the password from the default.

   Then click **Submit**.

After a successful login, the BBI Switch Dashboard is displayed (Figure 5-13).

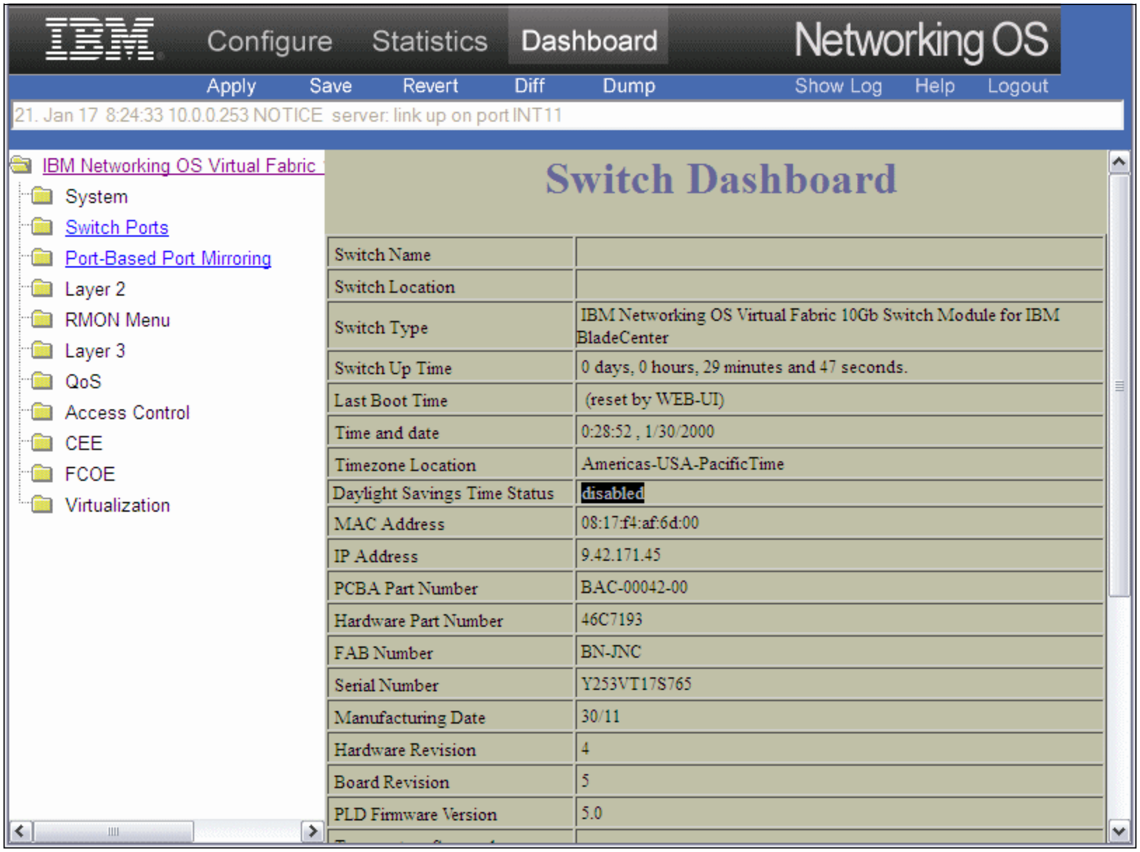

Figure 5-13 BNT Virtual Fabric 10Gb Switch Module dashboard

The switch dashboard is your starting point for switch module management.

5.3 Creating vNIC in IBM Virtual Fabric Mode

This section explains how to create vNICs in a BNT Virtual Fabric 10Gb Switch Module by using the following methods:

- BBI
- isCLI
- IBM System Networking OS CLI (the BNT CLI)

Each vNIC is created and configured on the switch. The properties of each vNIC are communicated to the Emulex adapter by using the DCBX Protocol.

5.3.1 vNIC features and considerations

The Virtual Fabric Adapter for IBM BladeCenter provides the following vNIC features:

- Up to four vNICs are supported on each internal switch port.
- Each vNIC can accommodate one of the following traffic types:
  - Regular Ethernet
  - Internet Small Computer System Interface (iSCSI)
  - Fibre Channel over Ethernet (FCoE)
- vNICs with traffic of the same type can be grouped together, in addition to regular internal ports, external uplink ports, and trunk groups, to define vNIC groups for enforcing communication boundaries.
- If a failure occurs on the external uplink ports that are associated with a vNIC group, the switch can signal affected vNICs for failover and allow other vNICs to continue operation.
- Each vNIC can be allocated a percentage of the 10 Gbps bandwidth on the link (from NIC to switch and from switch to NIC).

You can use the BNT Virtual Fabric 10Gb Switch Module as the single point of vNIC configuration.

When creating vNICs, keep in mind the following points:

- You can create up to four vNICs on any blade.

  Only 10 Gb of total bandwidth can be allocated to the vNICs on one port on a server blade. Oversubscription is not supported. If less than 10 GB of total bandwidth are allocated to all of the defined vNICs on a port, the deallocated bandwidth is unusable.
- Bandwidth allocations are policed on both inbound and outbound traffic.

  Unlike configurations that use quality of service and the vNIC functions that are built into the IBM AIX® operating system, no provision allows a vNIC to exceed its allocated bandwidth at any time.

Bandwidth allocations can be changed from the `switch` command line or browser interface. These changes are propagated to the servers immediately. They do not require bringing down the network or rebooting the server. For example, you can have one bandwidth allocation scheme in use during business hours and have a different one during off hours when batch processing is in progress.

5.3.2 Enabling and disabling vNIC functions on the server

To configure vNICs on one blade, begin by enabling vNIC functions. This process uses a master toggle switch that you can use to temporarily disable vNIC without removing all of the configuration details. You run this command once per switch, regardless of how many blades will use the vNIC.

Configuration by using the BBI method

To enable vNIC mode by using the BBI method, set the Global vNIC On/Off parameter to **On** as shown in Figure 5-14. Then click **Submit**.

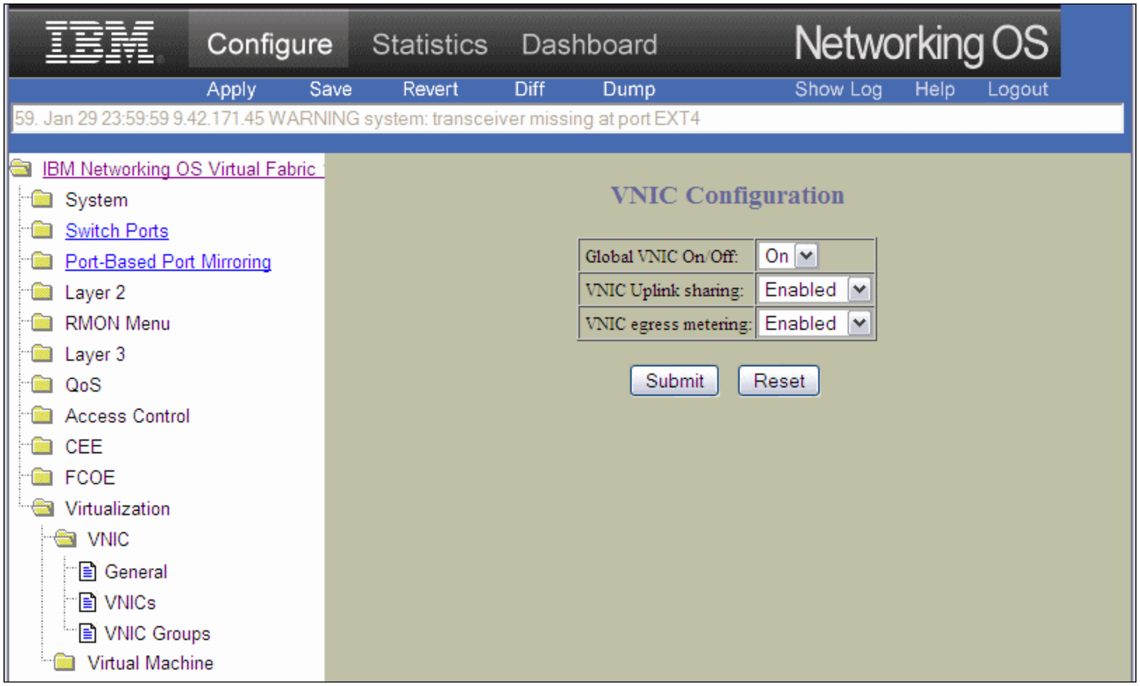

Figure 5-14 Enabling vNIC mode

To disable this mode, setting the Global vNIC On/Off parameter to **Off**.

Configuration by using the isCLI method

To enable vNIC mode by using the isCLI method, enter the **vnic enable** command as shown in Example 5-1.

Example 5-1 Enabling vNIC mode by using the isCLI method

```
BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter.
Enter password:
System Information at 19:12:16 Sun Jan  8, 2000

Router>
Jan  8 19:12:16 9.42.171.39 NOTICE  mgmt: admin(admin) login from host
9.44.168.83
en

Enable privilege granted.
Router#configure terminal
Enter configuration commands, one per line.  End with Ctrl/Z.
Router(config)#vnic enable
Router(config)#
```

To disable vNIC mode by using isCLI, enter the **no vnic enable** command as shown in Example 5-2.

Example 5-2 Disabling vNIC mode

```
Router(config)#no vnic enable
Router(config)#
```

Successful configuration: When you successfully enable or disable vNIC mode by using the isCLI method, no message is displayed.

Configuration by using the IBM System Networking OS CLI method

To enable vNIC mode by using the IBM System Networking OS CLI method:

1. Go to the `/cfg/virt/vnic` directory.
2. In the Main Menu panel (Figure 5-15), enter **cfg** and press Enter.

```
[Main Menu]
Jan 15 16:56:14 9.42.171.39 NOTICE  mgmt: admin(admin) login from host
9.44.168.155
     info     - Information Menu
     stats    - Statistics Menu
      cfg      - Configuration Menu
      oper     - Operations Command Menu
     boot     - Boot Options Menu
     maint    - Maintenance Menu
     diff     - Show pending config changes  [global command]
     apply    - Apply pending config changes [global command]
     save     - Save updated config to FLASH [global command]
     revert   - Revert pending or applied changes [global command]
     exit     - Exit  [global command, always available]
>> Main#
Jan 15 16:56:33 9.42.171.39 WARNING ntp: cannot contact primary NTP server
9.42.
171.253
cfg
```

Figure 5-15 The Main Menu

3. In the Configuration Menu panel (Figure 5-16), enter **virt** and press Enter.

```
[Configuration Menu]
     sys      - System-wide Parameter Menu
     port     - Port Menu
     qos      - QOS Menu
     acl      - Access Control List Menu
     pmirr    - Port Mirroring Menu
     l2       - Layer 2 Menu
     l3       - Layer 3 Menu
     cee      - CEE Configuration Menu
     fcoe     - Fiber Channel Over Ethernet Configuration Menu
     rmon     - RMON Menu
     virt     - Virtualization Menu
     setup    - Step by step configuration set up
     dump     - Dump current configuration to script file
     ptcfg    - Backup current configuration to FTP/TFTP server
     gtcfg    - Restore current configuration from FTP/TFTP server
     cur      - Display current configuration
```

Figure 5-16 The Configuration Menu

4. On the Virtualization Menu, enter **vmpolicy**, and press Enter (Figure 5-17).
5. On the VNIC Global Configuration Menu, enter **on** to enable vNIC functions.

```
>> Configuration# virt
------------------------------------------------------------
[Virtualization Menu]
     vmpolicy - Virtual Machines Policy Configuration Menu
     vnic     - vNIC Configuration Menu
     vmgroup  - Virtual Machines Groups Menu
     vmprof   - Virtual Machine Profiles Menu
     vmware   - VMware-specific Settings Menu
     vmrmisc  - Miscellaneous VMready Configuration Menu
     enavmr   - Enable VMready
     disvmr   - Disable VMready
     cur      - Display all current virtualization settings

>> Virtualization# vnic
------------------------------------------------------------
[VNIC Global Configuration Menu]
     port     - Port vNIC Configuration Menu
     vnicgrp  - VNIC Group Configuration Menu
     emeter   - Globally enable/disable vNIC egress metering
     ulshare  - Globally enable/disable vNIC uplink sharing
     on       - Globally turn vNIC feature ON
     off      - Globally turn vNIC feature OFF
     cur      - Display current vNIC configuration

>> VNIC Global Configuration#
```

Figure 5-17 Enabling vNIC by using the IBM System Networking OS CLI

5.3.3 Creating the vNICs

This section explains how to create vNICs. In most installations, you are likely to have two BNT Virtual Fabric 10Gb Switch Modules in your chassis. Therefore, you must complete this process twice, once on each of the two switches.

If you configure the two physical ports identically, especially if trunk failover is to be used for high availability, you use the same commands on both switches.

When issuing the commands, keep in mind the following points:

- Blade numbers (internal port numbers) are 1 through 14, and vNIC numbers are 1 through 4.
- Bandwidth is specified in 100 Mbps increments, which also specifies the percentage of 10 Gbps bandwidth that is allocated to a specific vNIC.

Configuration by using the BBI method

When you use the BBI method, select the **Configure** tab (Figure 5-18), and select the vNIC. In the Configuration panel on the right side, for State, select **enabled**, and then set the Max Bandwidth. In this example, we entered `30` for 3 Gbps.

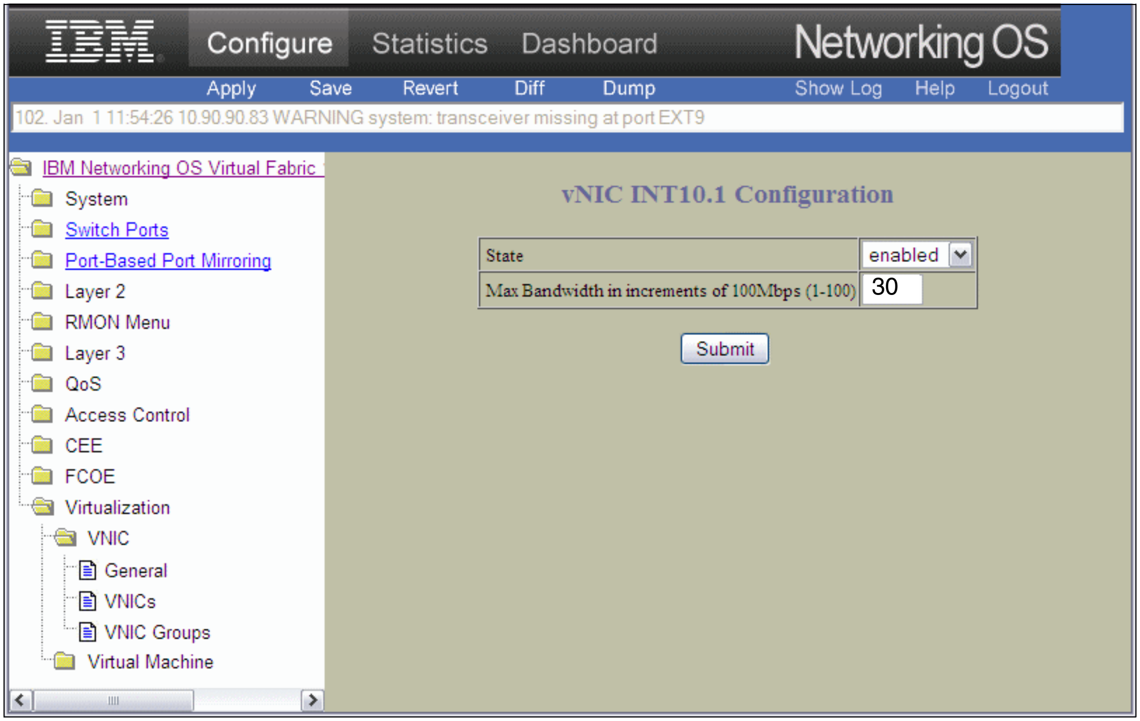

Figure 5-18 Creating a vNIC

Configuration by using the isCLI method:

The isCLI method has the following command syntax:

```
vnic port <port alias or number> index <1-4>
bandwidth <1-100>
enable or no enable
```

For example, you want to complete the following configuration tasks by using the isCLI method:

1. Enable vNIC # 8 on blade 8.
2. Allocate 30% of the 10 Gb to this vNIC (3 Gb).
3. Enable vNIC #2 on blade 8.
4. Allocate 70% of the 10 Gb to this vNIC (7 Gb).

Figure 5-19 shows the commands to do these tasks and the resulting output.

```
Jan  9 17:02:38 9.42.171.39 NOTICE  mgmt: admin(admin) login from host
9.44.168.202
en
Enable privilege granted.
Router#configure terminal
Enter configuration commands, one per line.  End with Ctrl/Z.
Router(config)#vnic port 8 index 1
Router(vnic_config)#band
Router(vnic_config)#bandwidth 30
Router(vnic_config)#enable
Router(config)#vnic port 8 index 2
Router(vnic_config)#bandwidth 70
Router(vnic_config)#enable
Router(vnic_config)#
```

Figure 5-19 Commands for the isCLI method to create a vNIC

Configuration by using the IBM System Networking OS CLI method

For the IBM System Networking OS CLI method, the following commands are relevant:

port <port number>	Displays the port vNIC menu.
vnic <1-4>	Displays the vNIC menu for the selected vNIC.
bw <1-100>	Configures the maximum bandwidth allocated to this vNIC, in increments of 100 Mbps. For example, 1 = 100 Mbps, 10 = 1000 Mbps.
ena	Enables the selected vNIC.
dis	Disables the selected vNIC.
cur	Displays the current vNIC port parameters.

Figure 5-20 shows the commands that were used in the example and the resulting output.

```
> Virtualization# vnic
----------------------------------------------------------
[VNIC Global Configuration Menu]
     port     - Port vNIC Configuration Menu
     vnicgrp  - VNIC Group Configuration Menu
     on       - Globally turn vNIC feature ON
     off      - Globally turn vNIC feature OFF
     cur      - Display current vNIC configuration
>> VNIC Global Configuration# port
Enter port (INT1-INT14):       4
----------------------------------------------------------
[Port INT4 vNICs Menu]
     vnic     - VNIC Configuration Menu
     cur      - Display current port vNIC configuration
>> Port INT4 vNICs# vnic
Enter vNIC (1-4):      2
----------------------------------------------------------
[vNIC INT4.2 Menu]
     bw       - Set maximum bandwidth of the vNIC
     ena      - Enable vNIC
     dis      - Disable vNIC
     cur      - Display current vNIC configuration
>> vNIC INT4.2# ena
Current status: disabled
New status:     enabled
Warning: "Tagging" is enabled on vNIC port INT4
vNIC INT4.2# bw 20
Current vNIC Maximum Bandwidth: 25 (increments of 100Mbps)
Pending vNIC Maximum Bandwidth: 20 (increments of 100Mbps)
```

Figure 5-20 Configuring the vNIC by using the IBM System Networking OS CLI

5.4 Configuring the vNIC group

A vNIC group consists of one or more vNICs, optionally one or more uplink (external) ports, and optionally one or more non-vNIC server blade-facing (internal) ports. For a vNIC to communicate, it must be a member of a single vNIC group, and only one vNIC from a port can be in any specific group.

To configure a vNIC group:

1. Enable or disable a group. You can create up to 32 groups.

   ```
   /cfg/virt/vnic/vnicgrp <number>/{ena or dis}
   ```

2. Add a vNIC to the group:

   ```
   /cfg/virt/vnic/vnicgrp <number>/addvnic <blade#.vnic#>
   ```

3. Add a physical port to the group:

   ```
   /cfg/virt/vnic/vnicgrp <number>/addport <INT or EXT port>
   ```

4. Add an external trunk group to the vNIC group:

   ```
   /cfg/virt/vnic/vnicgrp <number>/addtrnk <trunk #>
   ```

5. Enable or disable trunk failover, as explained in 5.5, “High availability and vNICs” on page 81:

   ```
   /cfg/virt/vnic/vnicgrp <number>/failover {ena or dis}
   ```

Tip: You can use the **`remvnic`**, **`remport`**, and **`remtrnk`** commands to remove the corresponding items from the group.

Figure 5-21 shows the commands to add a vNIC and a trunk to a vNIC group.

```
>> Virtualization# vnic
------------------------------------------------------------
[VNIC Global Configuration Menu]
     port     - Port vNIC Configuration Menu
     vnicgrp  - VNIC Group Configuration Menu
     on       - Globally turn vNIC feature ON
     off      - Globally turn vNIC feature OFF
     cur      - Display current vNIC configuration
>> VNIC Global Configuration# vnicgrp
Enter vNIC Group (1-32):       8
------------------------------------------------------------
[vNIC Group 8 Menu]
     vnicvlan - Set VLAN number to vNIC group
     failover - Enable/disable uplink failover
     addvnic  - Add vNIC to vNIC group
     remvnic  - Remove vNIC from vNIC group
     addport  - Add port to vNIC group
     remport  - Remove port from vNIC group
     addtrnk  - Add trunk to vNIC group
     remtrnk  - Remove trunk from vNIC group
     ena      - Enable vNIC group
     dis      - Disable vNIC group
     del      - Delete vNIC group
     cur      - Display current vNIC group configuration
>> vNIC Group 8# addvnic INT4.2
>> vNIC Group 8# addtrnk 3
>> vNIC Group 8# ena
Current status: disabled
New status:     enabled
>> vNIC Group 8#
```

Figure 5-21 Using the IBM System Networking OS CLI to create a vNIC Group

Figure 5-22 shows the equivalent commands for the isCLI method.

```
vnic vnicgroup <number 1-32>
member INT<port>.<vnic number 1-4>
port INT<port number or alias> or EXT <port number or alias>
vlan <vlan number 1-4094>
trunk <trunk number>
failover
enable
exit
```

Figure 5-22 Commands for the isCLI method

The **exit** command is required to end the configuration for a specific group.

> **Tip:** You can use the **no member**, **no port**, **no trunk**, **no vlan**, and **no failover** commands to remove the corresponding items from the group.

Figure 5-23 shows an example of the commands by using the isCLI method.

```
Router>Jan 10 17:20:24 9.42.171.39 NOTICE  mgmt: admin(admin) login from
host 9.44.168.202
en
Enable privilege granted.
Router#configure terminal
Enter configuration commands, one per line.  End with Ctrl/Z.
Router(config)#vnic vnicgroup 2
Router(vnic group config)#member INT8.1
Router(vnic group config)#vlan 4091
Warning: VLAN 4091 is moved to STG 1.
Router(vnic group config)#failover
Router(vnic group config)#trunk 2
Router(vnic group config)#enable
Warning: STP is turned off for port EXT1 in STG 1
Warning: Changed the pvid of uplink port EXT1 in vNIC group 2 to 4091
Warning: Deleted port EXT1 from VLAN 1
Router(vnic group config)#
```

Figure 5-23 Creating a vNIC group

Currently, only one uplink port *or* one trunk group of uplink ports that are aggregated together can be part of a vNIC group. Any uplink port or trunk can be part of only one vNIC group at a time. Therefore, it is not currently possible for multiple groups to share uplinks. Now, only static trunk groups are supported. Link Aggregation Control Protocol (LACP) trunks are supported in Switch Independent Mode (see Part 3, “Switch Independent Mode” on page 153).

Each vNIC group is associated with a vNIC VLAN, as explained in 5.6, “Configuring vNIC VLANs” on page 85.

5.5 High availability and vNICs

The trunk failover function enables high-availability designs when used with standard ports. This failover function is also available when vNICs are used. With vNICs, the uplink ports that are members of a group (vnicgrp) might fail due to a failure on the upstream device or a cable that is failing or erroneously removed.

In this case, the associated vNIC is brought down by the Virtual Fabric Switch. In Windows, this situation appears as a *network cable unplugged* condition on the server.

To deliver high availability by using the trunk failover function:

1. Configure the other Virtual Fabric Switch in the chassis to use vNICs for the server in question. Optionally, use the same vNIC numbers, associated groups, and bandwidth allocation on both switches.
2. Team together the corresponding vNIC ports, as detected by the operating system, by using the available tools for NIC teaming. In general, perform NIC teaming in active mode or standby mode by using the following specifications:
   - For Windows servers, use the available utility from Emulex.
   - Linux servers must have the local *bonding* driver.
   - ESX has support for multiple NIC teaming modes by using the network configuration GUI.

Multiple vNICs that are part of the same physical 10 Gbps port can support high availability, but each vNIC does so independently. To implement support for high availability:

1. Enable failover for all vNIC groups for which the server ports are members, on both of the Virtual Fabric Switches.
2. Configure NIC teaming (as explained previously) on corresponding pairs of vNICs on the server. Ensure that one member of the team is physically connected to one of the Virtual Fabric Switches and that the other member of the team to the other Virtual Fabric Switch.
3. Configure uplink ports or trunks as members of the corresponding vNIC groups on the two switches. If possible, connect these uplinks to separate upstream physical switches.

Figure 5-24 shows a sample configuration with only one blade. The same configuration commands are required for each blade used for which failover is desired.

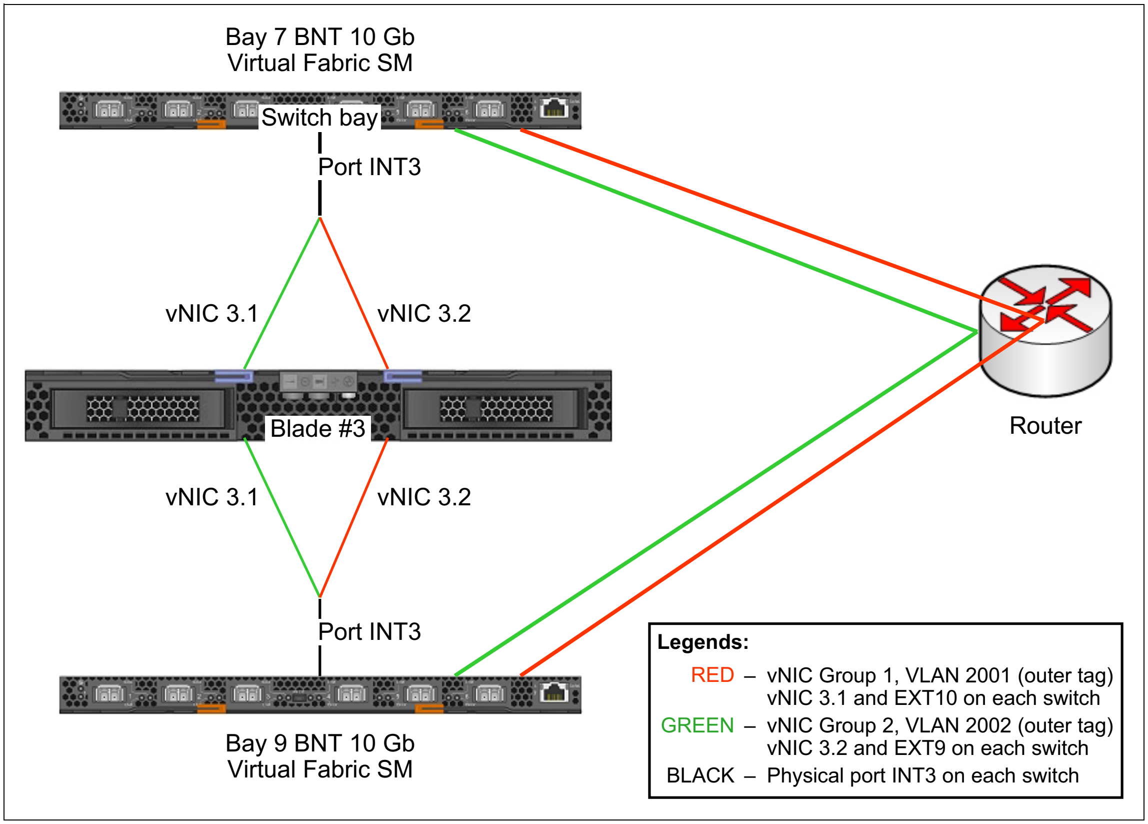

Figure 5-24 High availability with vNIC

Figure 5-25 shows the commands of the isCLI method for the blade configuration.

```
vnic port INT3 index 1
   enable
vnic port INT3 index 2
   enable
vnic vnicgroup 1
   failover
   vlan 2001
   member INT3.1
   port EXT10
   enable
   exit

vnic vnicgroup 2
   failover
   vlan 2002
   member INT3.2
   port EXT9
   enable
   exit
```

Figure 5-25 Commands for the isCLI method

Figure 5-26 shows the IBM System Networking OS CLI configuration, which is identical on each of the two switches.

```
/cfg/virt/vnic /port 3/vnic 1/ena
/cfg/virt/vnic/port 3/vnic 2/ena
/cfg/virt/vnic/vnicgrp 1/ena
   failover ena
   vnicvlan 2001
   addvnic 3.1
   addport EXT10
/cfg/virt/vnic/vnicgrp 2/ena
   failover ena
   vnicvlan 2002
   addvnic 3.2
   addport EXT9
```

Figure 5-26 Commands for the IBM System Networking OS CLI method

For more information, see:

- *BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter Application Guide*:

  http://www.bladenetwork.net/userfiles/file/VFSM_AG_6-8.pdf

- Publications for the BNT Virtual Fabric 10Gb Switch Module - IBM BladeCenter H, HT

  http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5080917

5.6 Configuring vNIC VLANs

Every vNIC channel is associated with a vNIC group with a unique VLAN ID. These group VLANs differ from conventional VLANs that are configured on the Virtual Fabric Switch in significant ways:

- vNIC VLANs are implemented as *outer tags* on the Ethernet frames that pass through the switch. These outer tags are *never* forwarded to a blade or to an upstream switch. They serve solely to isolate traffic in a vNIC group from other vNIC groups and from conventional VLANs.
- Within a group, conventional VLANs configured by an operating system on a server blade (such as with the **vconfig** command in Linux) are implemented as *inner tags*. These inner tags pass through the group and the Virtual Fabric Switch unchanged. The VLAN configuration on upstream switches must match the VLANs on the blades, but the Virtual Fabric Switch passes the inner tags through (Figure 5-27 on page 86).
- The same conventional VLAN can enter the chassis in more than one group, but traffic between groups never occurs within the Virtual Fabric Switch.
- On the Virtual Fabric Switch, a VLAN number can be a conventional VLAN, a vNIC group, or a VMready group, but not more than one type at the same time.

Keep in mind the following additional notes about vNIC VLANs:

- Routing of traffic between vNIC VLANs is not supported, nor is routing between a vNIC VLAN and a conventional VLAN.
- When external (uplink) ports are added to a vNIC group, they are also added to the vNIC VLAN and are always aggregated together. Currently, it is not possible to split these ports and aggregate them in more than one trunk.
- When internal (blade-facing) ports are added to a vNIC group, it is assumed that no vNICs are created on those blades (with the **addport** command, not the **addvnic** command). Traffic can then flow between vNICs and these internal ports. Traffic is forwarded by using the vNIC VLAN as the outer tag.

Any inner tags that are associated with the non-vNIC interface on the ports are unchanged, as with the vNICs.

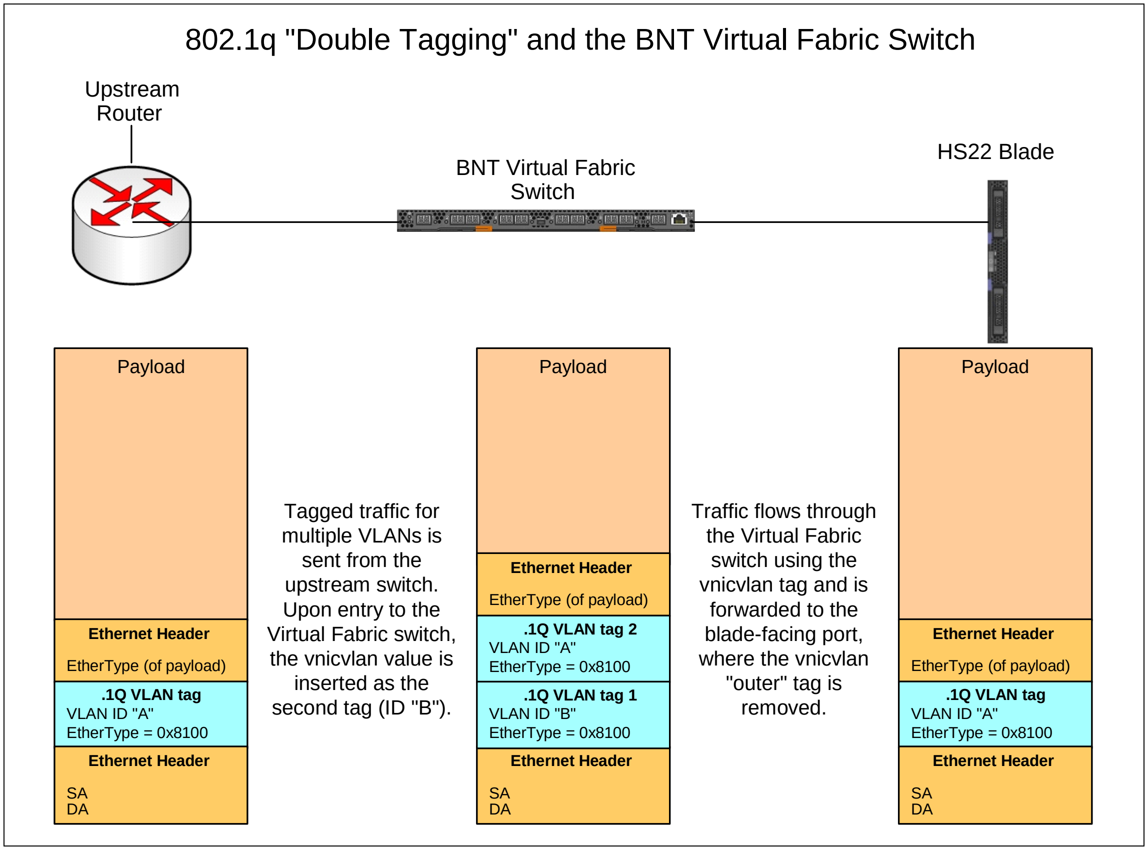

Figure 5-27 Double tagging

vNIC VLANs are configured with the associated vNIC groups.

To configure vNIC VLANS, you can use the following isCLI commands:

```
vnic vnicgroup <number 1-32>
vlan <vlan number 1-4094>
```

Alternatively, you can use the IBM System Networking OS CLI commands:

```
/cfg/virt/vnic/vnicgrp <number>
vnicvlan <1-4094>
```

You cannot use the selected VLAN number within the Virtual Fabric Switch for any other purpose. However, it is of local significance in that it is never detected by the blades (other than by the vNIC driver) and never forwarded out of the Virtual Fabric Switch to an upstream switch. Use high numbers that are not used in the client's network as vNIC VLANs (for example, 4000 and up).

Figure 5-28 shows the BBI window to configure a vNIC group and the associated VLAN.

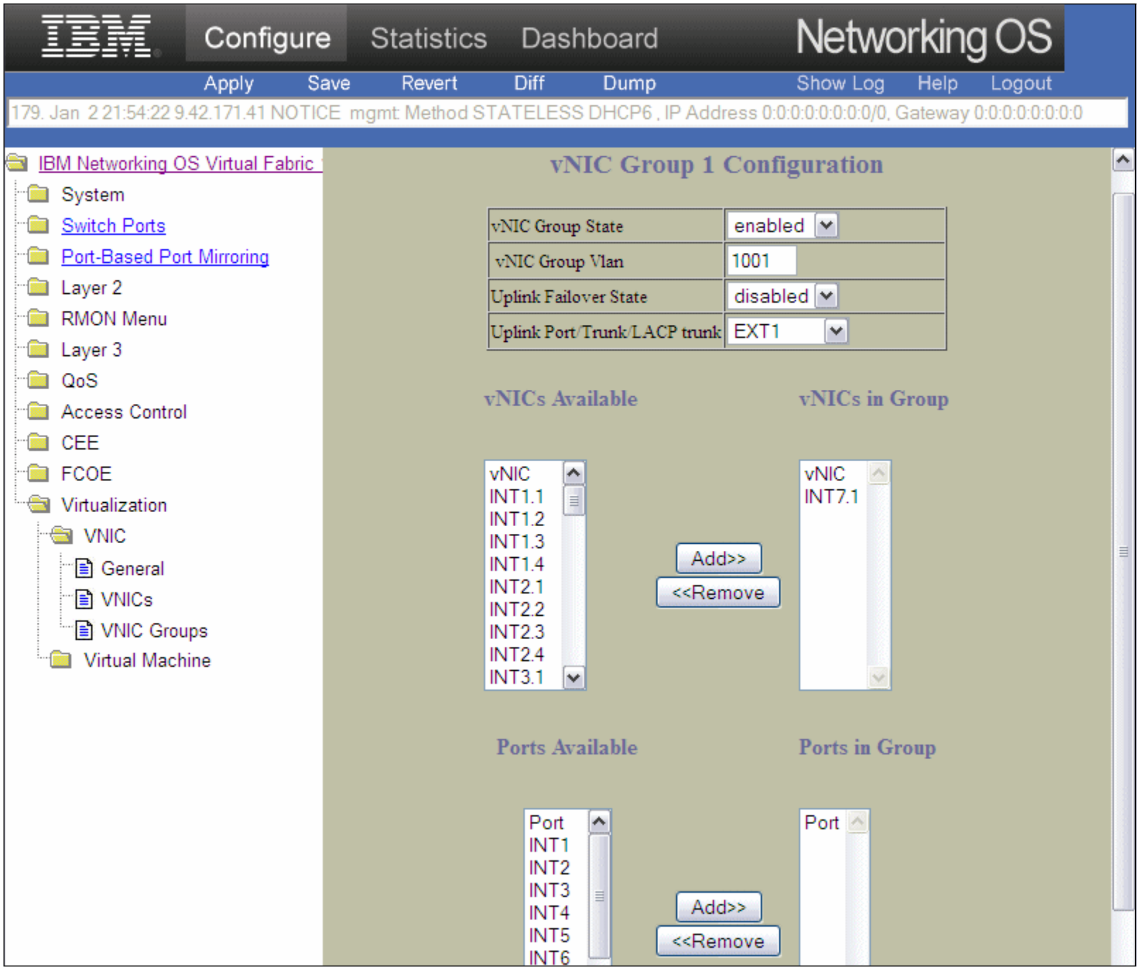

Figure 5-28 Configuring the vNIC group and the associated VLAN

5.7 Configuring the operating system

This section explains how to configure teams and VLANs with the Emulex 10GbE Virtual Fabric Adapter. The following operating systems support the Emulex 10GbE Virtual Fabric Adapter:

- Microsoft Windows Server 2003 (x86 and x64)
- Microsoft Windows Server 2008 (x86 and x64)
- Microsoft Windows Server 2008 Release 2 (x64 only)
- Red Hat Enterprise Linux 5.5 and 5.6
- Red Hat Enterprise Linux 6.0 and 6.1
- SUSE Linux Enterprise Server 10.3 and 10.4
- SUSE Linux Enterprise Server 11.1
- VMware ESX 4.0, 4.1, and 5.0
- Citrix XenServer 5.6
- Citrix XenServer 6.0

> **Important:** Before installing the device drivers and firmware, make sure that the following firmware for other BladeCenter devices is at the latest level.

OneCommand Manager is the Emulex application that enables the centralized management of Emulex OneConnect converged network adapters (CNAs) and host bus adapters (HBAs) from a centralized management console. OneCommand Manager provides a GUI and a scriptable CLI for administration. You can use the OneCommand Manager to configure load balancing and teaming.

5.7.1 Windows OneCommand Manager

You can install the OneCommand Manager application in Windows by using either of the following methods:

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

Attended installation in Windows

To do an attended installation of the OneCommand Manager application in Windows:

1. From the Emulex website, download the x64 or x86 OneCommand Manager Enterprise Kit installation file to your system.

   > **IA64:** For IA64 systems, use the x86 OneCommand Manager Enterprise installation file.

2. Go to the directory for which you downloaded the file.
3. Double-click the **elxocm*version*.exe** file.
4. In the Emulex OCManager Enterprise window (Figure 5-29), click **Next**.

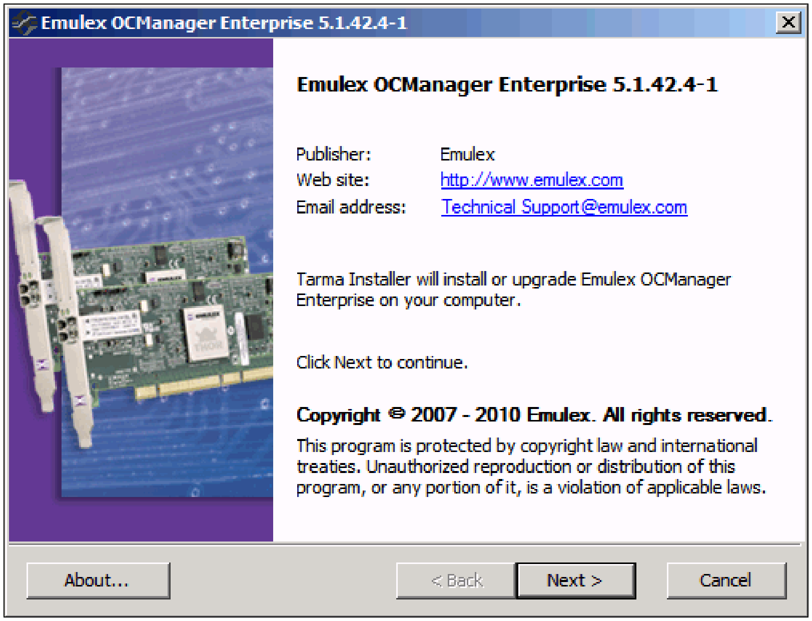

Figure 5-29 Emulex OCManager Enterprise Kit Installer welcome window

5. In the Installation options window (Figure 5-30), click **Install**. The installation process begins.

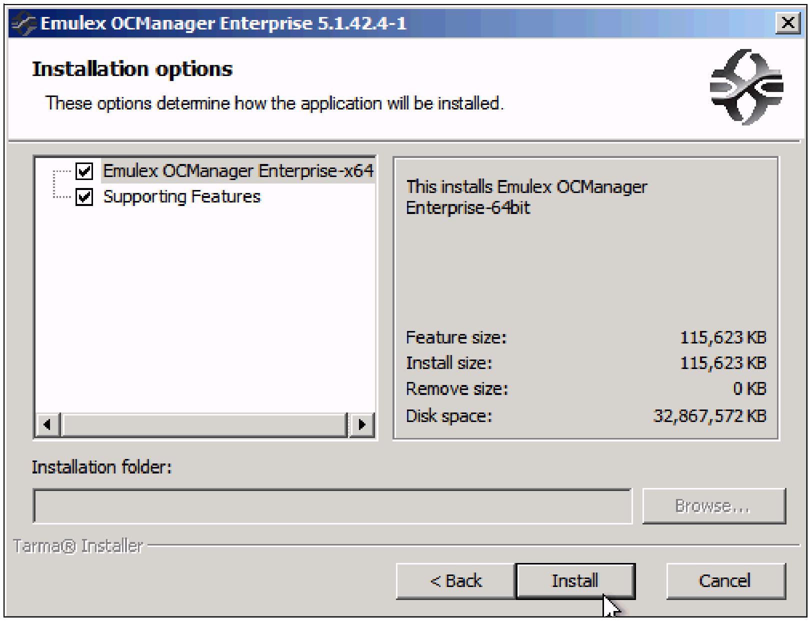

Figure 5-30 Emulex OCManager Installation options window

6. After the installation process completes, in the Management Mode dialog box (Figure 5-31), choose the management mode you want. For this example, we selected **Local Management Plus**. Then click **OK**.

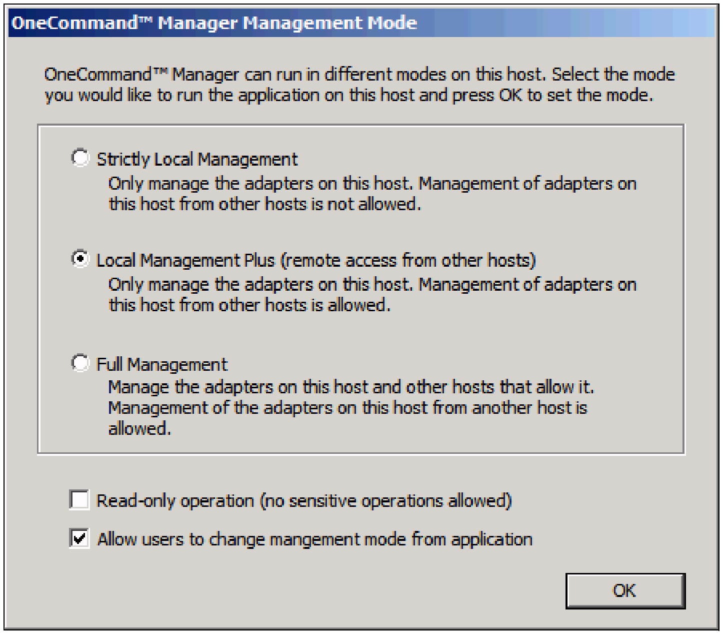

Figure 5-31 OneCommand Manager Management Mode options

Tip: Allow users to change the management mode for the application at this stage.

The Installation Completed window (Figure 5-32) opens when the installation is finished.

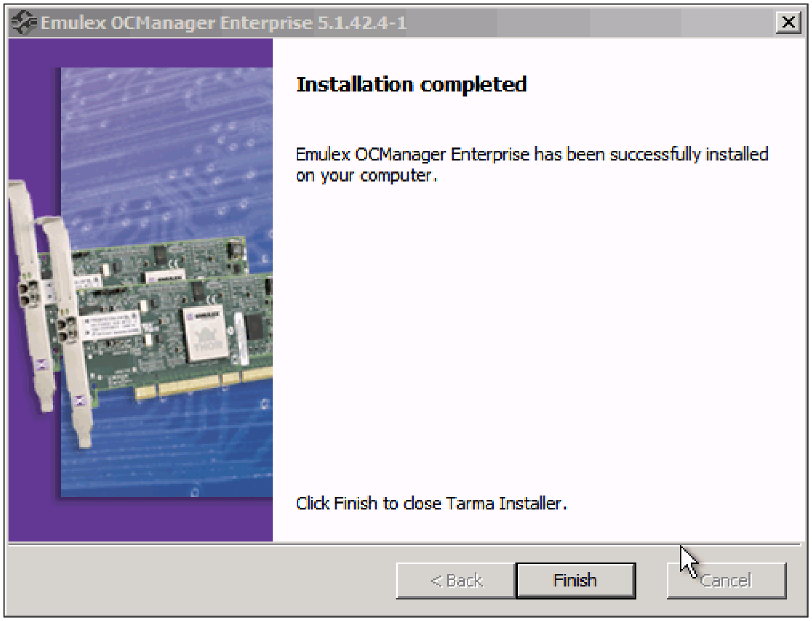

Figure 5-32 Emulex OCManager Installation completion

Unattended installation in Windows

To do an unattended installation of the OneCommand Manager application in Windows:

1. Download the x64 or x86 OneCommand Manager Enterprise Kit installation file to your system from the Emulex web page at:

   http://www.emulex.com/downloads/emulex/windows/windows-7-x64/management.html

2. At a Windows command prompt, run the unattended installation of OneCommand Manager. The kit is activated with the optional switch `/q` or `/q2`. The `/q` switch displays progress reports. The `/q2` switch does not display progress reports. For example, at the command prompt, type either of the following commands:

```
elxocm-windows-x86-5.1.42.4-1.exe /q
elxocm-windows-x64-5.1.42.4-1.exe /q2
```

3. Select a mode:
   a. Add the mode argument. The following mode values are possible:
      - Local Only Management Mode
      - Local Plus Management Mode
      - Full Management Mode
      - Local Plus Management Mode and Read Only
      - Full Management Mode and Read Only
   b. Add the ability to change the mode by adding the change argument with the values selected. For example, at the command prompt, type:

      ```
      elxocm-windows-x86-5.1.42.4-1.exe mmode=3 achange=1 /q2
      ```

      The following change values are possible:
      - Do not allow Management Mode to change
      - Allow Management Mode to change

5.7.2 Configuring teams and VLANs

A team (or bond in Linux) of adapters functions as a single virtual network interface and appears the same as a nonteamed adapter to other network devices. Teaming offers the following benefits:

- Virtual adapters can balance a large network load across several adapters.
- Teams improve system availability by providing a failover path for critical network connections.

A protocol address, such as an IP address, is assigned to the physical adapter. However, when OneCommand NIC Teaming and Multiple VLAN Manager (OneCommand NIC Teaming Manager) are installed, the protocol address is assigned to the team adapter. It is not assigned to the physical adapters that make up the team.

The **`IPCONFIG /all`** command for Windows or the **`IFCONFIG`** command for Linux shows the IP and media access control (MAC) addresses of the virtual adapter, and not of the individual physical adapters.

5.7.3 Configuring a new team

To configure a new team:

1. Start the OneCommand NIC Teaming Manager application.
2. In the OC NIC Teaming and VLAN Manager window (Figure 5-33), click **Create Team** to create a team configuration.

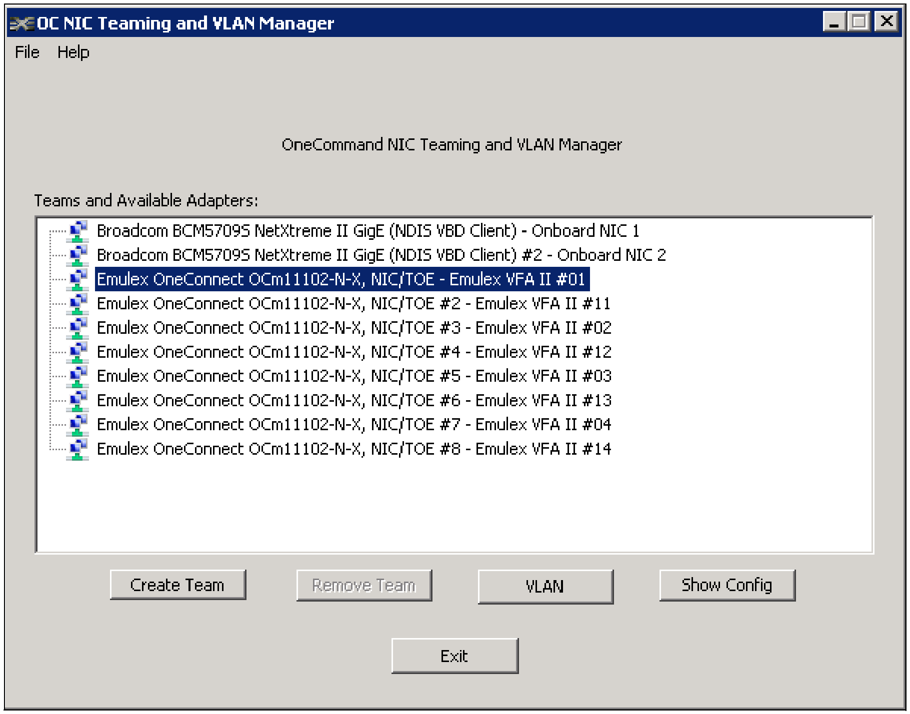

Figure 5-33 Configuring a new team by using OneCommand NIC Teaming Manager

3. In the OC NIC Teaming and VLAN Manager – Create Team window (Figure 5-34 on page 95):
   a. In the Team Name field, enter a name of the new team.
   b. For Team Type, select a type. The default is **Failover**.
   c. If necessary, for Load Distributed By, change the Team Type criteria:

      Default (Failover) Selects the port with the least traffic load for the session.

      Destination MAC Address
      Performs an `XOR` on the destination MAC address to determine which port must carry the load.

Destination IP Address

Performs an `XOR` on the destination and source IP address to determine which port must carry the load.

d. For Auto FailBack, select **Enable** (default) or **Disable**.

e. In the Available Network Adapters area, which lists the adapters that are not members of any team, to add an adapter to the team:

   i. Select the adapter from this list.

   ii. Click **Add** to move the adapter in the Team Member Adapters list, which lists the adapters that are members of the team.

f. Click **OK**.

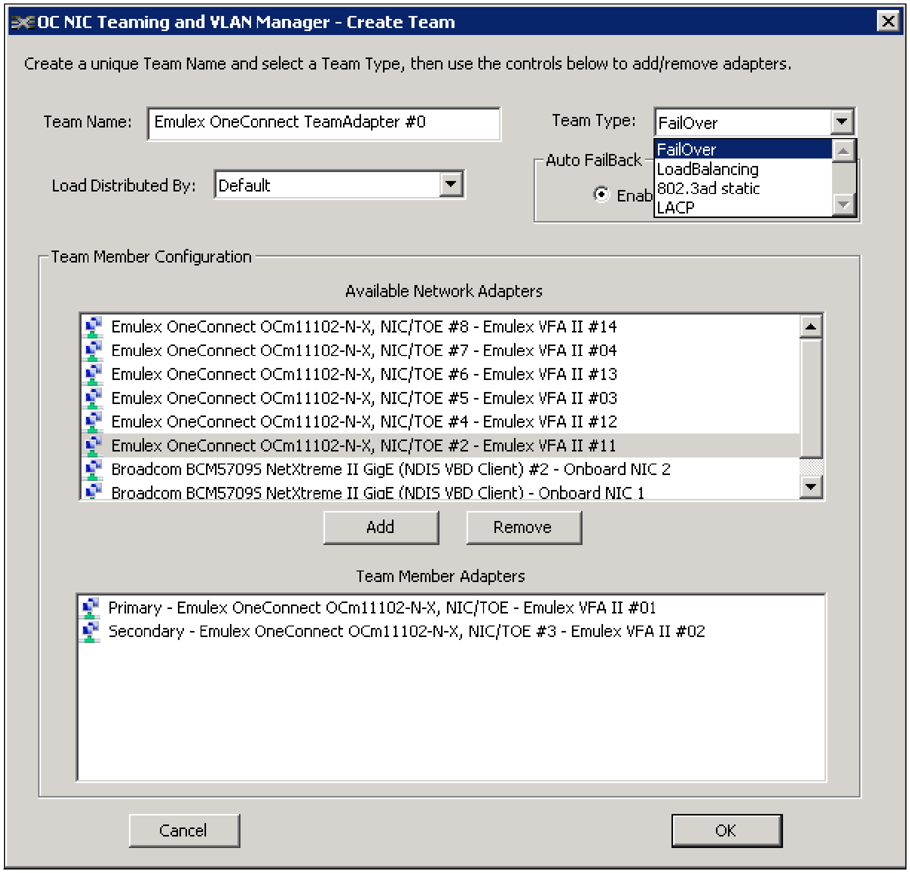

Figure 5-34 Configuring a new team

After the team is created, a window is displayed that shows that the adapter was added to the team.

5.7.4 Configuring a VLAN for a physical or team adapter

You can configure VLANs at the physical adapter level or at the team level. Configuring the VLAN at both the physical and team levels can cause double tagging. Therefore, avoid it. As a result, all virtual adapter members of the team must have the VLAN disabled at the physical or operating system level.

To configure a VLAN for a physical or team adapter, by using the OneConnect software:

1. In the OC NIC Teaming and VLAN Manager window (Figure 5-35), in the Team and Available Adapters area, select the physical or team adapter to which you want to add a VLAN. Then click **VLAN**.

> **Tip:** In the OC NIC Teaming and VLAN Manager window (Figure 5-35), if you select TeamAdapter #0 only, the VLAN is available for you to select it. In this window, you can see and confirm the adapter settings.

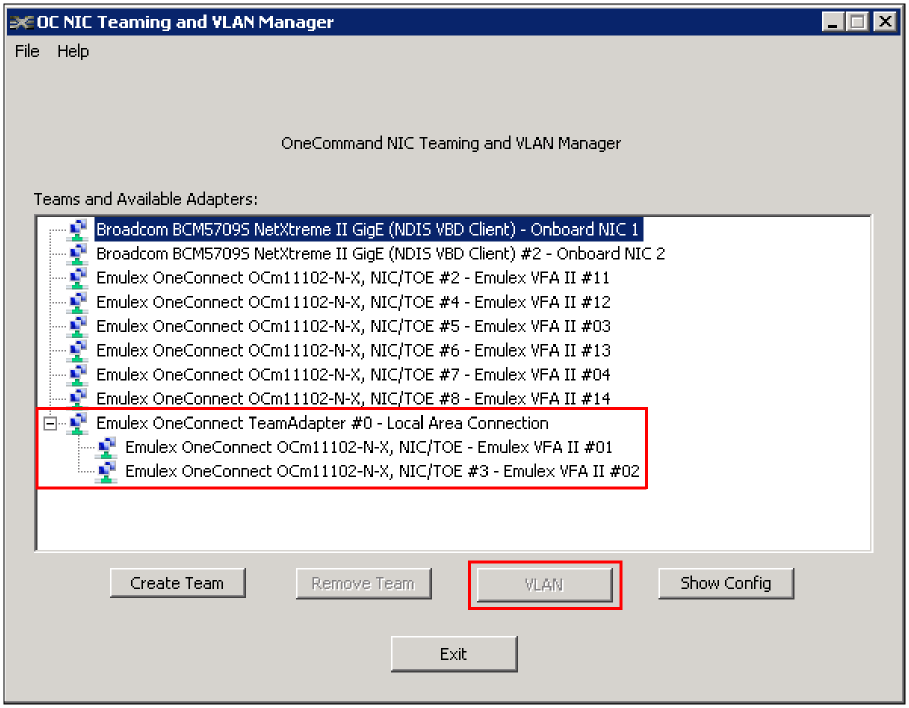

Figure 5-35 Configuring a new team

2. In the OC NIC Teaming and VLAN Manager – Add/Remove VLAN dialog box (Figure 5-36):

   a. Enter a valid VLAN ID in the range 1 - 4094, for example `50`.

   b. Click **Add** to add the VLAN to the adapter. The VLAN Name shows the VLAN Name in the format `Vlan_<VLAN ID>`.

   > **Multiple VLANs:** You can create multiple VLANs for an adapter. The VLANs Configured list shows the list of all VLANs that are configured for the adapter.

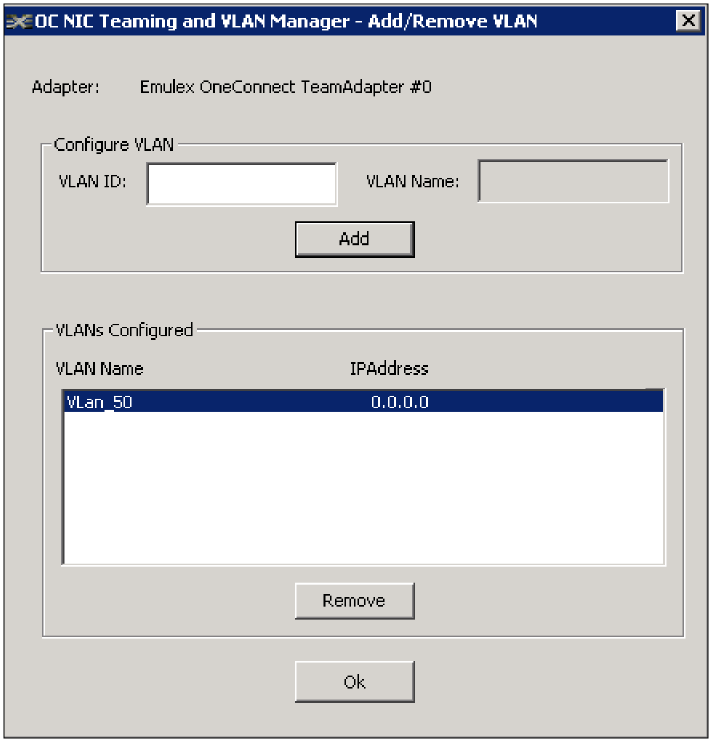

Figure 5-36 Confirmation of configured VLAN

   c. Click **OK** to complete the configuration.

The VLAN is now created.

5.7.5 Configuring FCoE and the personality for iSCSI Emulex Virtual Fabric Adapter II

By using the Emulex OneConnect adapters, you can change the personality of the adapter, can reboot the host, and can have the adapter come back up running the new personality. The OneConnect adapters are currently capable of running NIC-only, NIC and FCoE, and NIC and iSCSI personalities.

In some cases, the adapters are preconfigured to support multiple personalities. In other cases, a feature enablement license must be installed before the adapter can support multiple personalities. Also, the three personalities might not always be available on an adapter. For example, a NIC and FCoE adapter can change to a NIC-only or NIC and iSCSI adapter, but an iSCSI adapter cannot change to a NIC and FCoE adapter. Licensing is required to use the personalities.

To change the personality for the adapter, you need Emulex OneCommand Manager or the OneCommand Manager CLI. You can also use the F1 UEFI utility. In writing this book, we found that installing the Emulex OneCommand Manager was the easier method to perform the change.

> **Tip:** You can now perform the personality change by F1 UEFI utility on the blade server with 4.x firmware.

The OneCommand Manager CLI has the following commands to manage the personality type:

- **`ChangePersonality`**

  This command is supported by Windows, Solaris, Linux, and VMware ESX, and has the following syntax:

  ```
  hbacmd ChangePersonality <WWPN|MAC> <personality_type>
  ```

  `<personality_type>` must reflect the NIC, iSCSI, or FCoE depending on your configuration.

- **`ShowPersonalities`**

  This command shows the personality types and has the following syntax:

  ```
  hbacmd ShowPersonalities <WWPN|MAC>
  ```

Tip: You can install one (or more) driver kits for the current personality, and then change the personality and no longer have the drivers that are necessary to run the adapter. If you change personalities, you must install the appropriate drivers. These drivers are available on the Emulex website at:

http://www.emulex.com/downloads/ibm/vfa-software-kits/ocm5142-bc-windows/drivers.html

For more information about using the OneCommand Manager CLI, see the *Emulex OneCommand Manager Command Line Interface Version 5.1 User Manual* at:

http://www-dl.emulex.com/support/utilities/onecommand/519/corekit_user_manual.pdf

To change the personality, which was required during our testing scenarios, open the OneCommand Manager application and complete these steps:

1. In the OneCommand Manager window (Figure 5-37), in the left pane, select the adapter information.
2. On the **Adapter Information** tab in the right pane, in the Personality area, select the personality. In this example, we select **iSCSI**. Then click **Apply**.

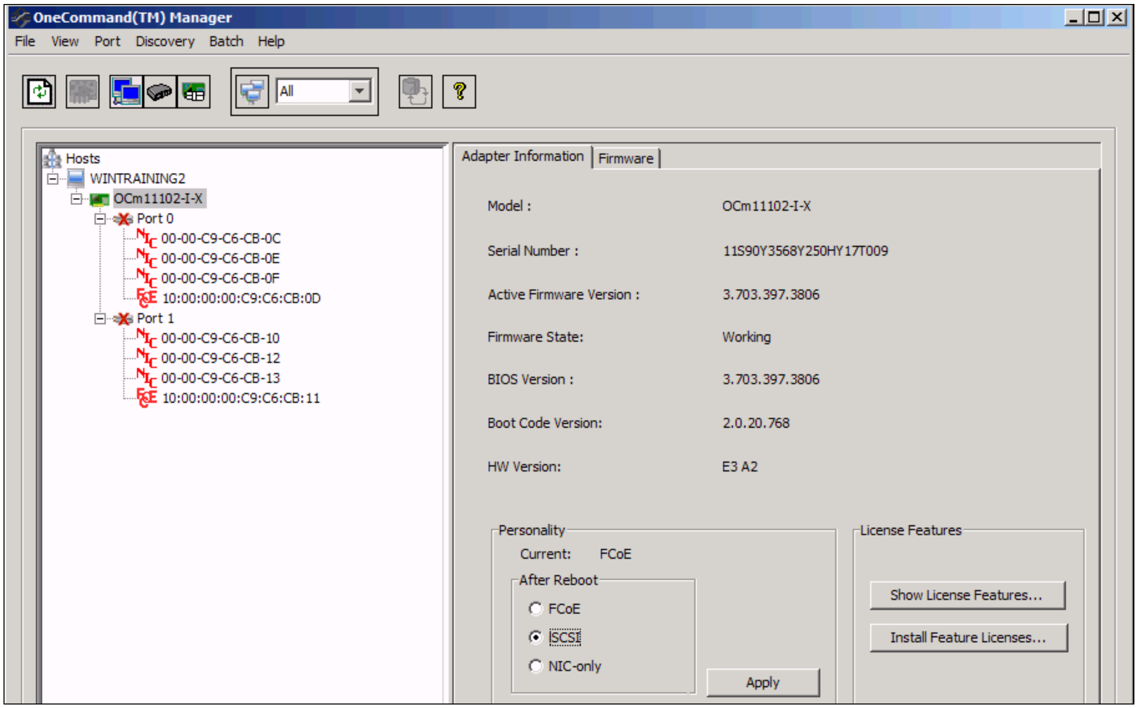

Figure 5-37 OneCommand Manager choosing a particular personality

3. In the Change Adapter Personality dialog box (inset in Figure 5-38), click **OK** to restart your system and to activate the change.

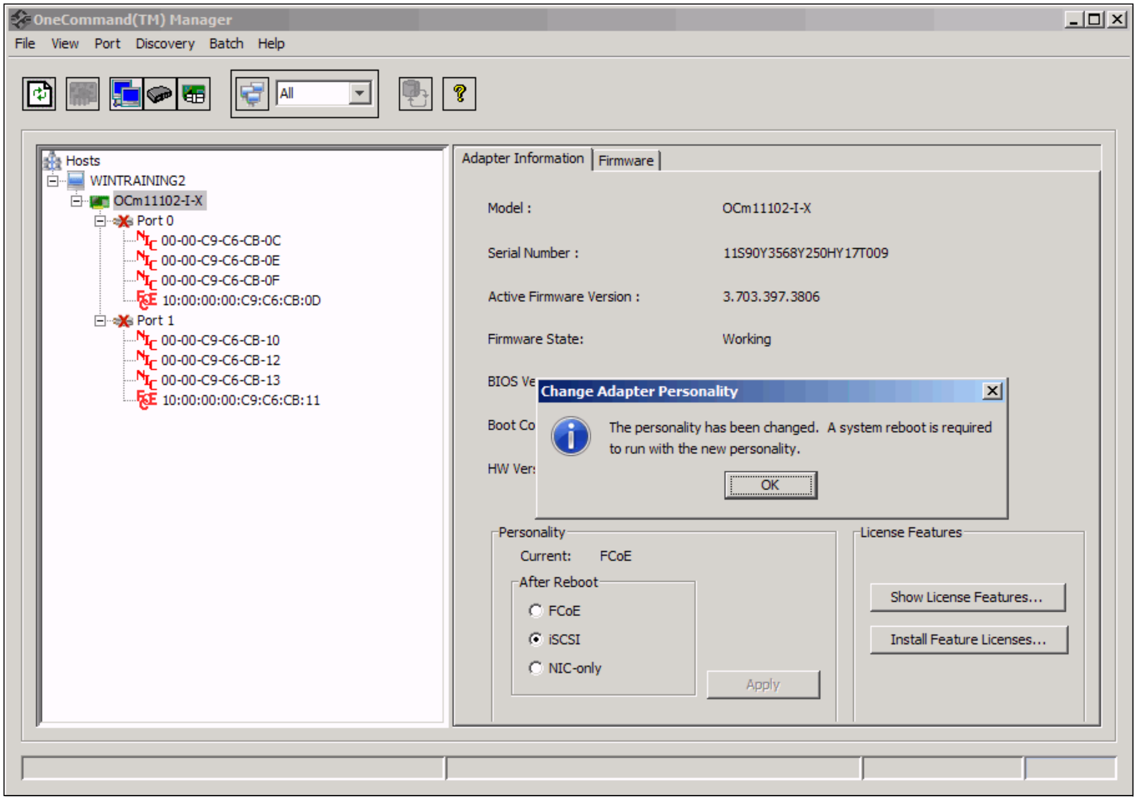

Figure 5-38 Emulex personality change reboot requirement

4. After your system restarts, verify your selection, and proceed with your required configuration. To configure an FCoE and iSCSI solution, see 5.8, “vNIC and VMready” on page 101.

For more information about OneCommand Manager, see the following user manuals:

- *Emulex OneCommand Manager Application Version 5.1 User Manual*

  http://www-dl.emulex.com/support/utilities/onecommand/519/onecommand.pdf

- *Emulex OneCommand Manager Command Line Interface Version 5.1 User Manual*

  http://www-dl.emulex.com/support/utilities/onecommand/519/corekit_user_manual.pdf

See also 6.4, “Scenario 4: vNIC FCoE on Emulex II, BNT, and QLogic with IBM DS5300” on page 120, and 6.5, “Scenario 5: vNIC iSCSI on Emulex II, BNT, and QLogic with IBM DS5300” on page 145.

5.8 vNIC and VMready

You can use VMready only with the BNT Virtual Fabric 10Gb Switch Module. By using VMready, you have much greater control over the allocation of network resources. Also, the server administrator can assign network policies at a group level, as opposed to a physical server level.

Access control list (ACL), quality of service (QoS), and VLAN settings in VMready are set at the group level. When virtual machines (VMs) move because of a VMotion event, the ACLs, QoS, and VLAN assignments move with them. This movement ensures that the appropriate network resources are allocated to the VMs, regardless of where they are in the network.

Versions of BNT Virtual Fabric 10Gb Switch Module firmware before 6.7.2 did not support a blade-facing port in both a vNIC group and a VMready group at the same time. However, with the release of 6.7.2 and later, this support is now possible. With versions before 6.7.2, you can, however, use vNIC to support certain blades and VMready on other blades. This support is possible by using the same switch and is done within the same chassis, but not the same port.

However, if you are running the adapters in Switch Independent Mode, you do not need to use vNIC groups and, therefore, can implement VMready groups.

The VMready software of the switch makes the switch virtualization aware. Servers that run hypervisor software with multiple instances of one or more operating systems can present each instance as an independent VM. With VMready, the switch automatically discovers VMs that are connected to the switch.

Virtualization is used to allocate server resources based on logical needs, rather than on a strict physical structure. With appropriate hardware and software support, servers can be virtualized to host multiple instances of operating systems, or VMs. Each VM has its own presence in the network and runs its own service applications.

A hypervisor manages the various virtual entities, such as VMs and virtual switches, that are on the host server. Depending on the virtualization solution, a virtualization management server can be used to configure and manage multiple hypervisors across the network. With some solutions, VMs can also migrate between host hypervisors, moving to different physical hosts and maintaining their virtual identity and services.

The IBM Networking OS 6.8 VMready feature supports up to 2,048 virtual entities in a virtualized data center environment. The switch automatically discovers the virtual entities that are attached to switch ports. It also distinguishes between regular VMs, Service Console Interfaces, and Kernel or Management Interfaces in a VMware environment.

Virtual entities can be placed into VM groups on the switch to define communication boundaries. Virtual entities in the same VM group can communicate with each other, but not virtual entities in different groups. VM groups also allow for configuring group-level settings, such as virtualization policies and ACLs.

The administrator can also preprovision virtual entities by adding their MAC addresses (or their IPv4 address or VM name in a VMware environment) to a VM group. When a virtual entity with a preprovisioned MAC address is connected to the switch, the switch automatically applies the appropriate group membership configuration.

The BNT Virtual Fabric 10Gb Switch Module with VMready also detects the migration of virtual entities across different hypervisors. As virtual entities move, the BNT Virtual Fabric 10Gb Switch Module with the IBM NMotion® feature automatically moves the appropriate network configuration. NMotion gives the switch the ability to maintain assigned group membership and associated policies, even when a virtual entity moves to a different port on the switch. VMready also works with VMware Virtual Center (vCenter) management software.

By connecting with a vCenter, the BNT Virtual Fabric 10Gb Switch Module can collect information about more distant virtual entities, synchronize switch and virtual entity configuration, and extend migration properties.

For more information, see *Implementing a VM-Aware Network Using VMready*, SG24-7985.

Support: VM groups and policies, virtual entity preprovisioning, and virtual entity migration features are not supported simultaneously on the same ports as vNICs.

6

Usage scenarios

This chapter explains how to create three virtual fabric configurations that have the following combinations:

- Virtual network interface controllers (vNICs)
- Uplink ports
- Uplink trunks

All of the configurations performed are done by using interfaces into the BNT Virtual Fabric 10Gb Switch Module.

> **Tip:** Ensure that vNIC mode is enabled in the Emulex Virtual Fabric Adapters. vNIC mode is the default.

This chapter includes the following sections:

- Scenario 1: vNICs only with no uplinks
- Scenario 2: vNICs only with one uplink
- Scenario 3: vNICs only with one uplink trunk
- Scenario 4: vNIC FCoE on Emulex II, BNT, and QLogic with IBM DS5300
- Scenario 5: vNIC iSCSI on Emulex II, BNT, and QLogic with IBM DS5300

Tip: For the Virtual Fabric Adapter, you can determine the Peripheral Component Interconnect Express (PCIe) function ID. You can determine this function by using either of the following methods:

- Running the **`ethtool -i`** command in Linux
- Viewing the Network Adapter Properties in Windows Device Manager

6.1 Scenario 1: vNICs only with no uplinks

Scenario 1 creates group VGRP 1. This group consists of vNIC ports and has the following configuration:

- vNIC INT1.2 (blade in slot 1, vNIC 2) of the switch in I/O bay 7, set to a maximum bandwidth of 5 Gbps (50% of 10 Gbps)
- vNIC INT2.1 (blade 2, vNIC 1) of the switch in I/O bay 7, set to a maximum bandwidth of 2 Gbps (20% of 10 Gbps)
- One vNIC group that contains the two vNICs (vNIC Group VLAN 127)

You can use the IBM Networking OS command script in Example 6-1 to implement this configuration.

Example 6-1 IBM Networking OS script to create scenario 1

```
/c/virt/vnic
        on
/c/virt/vnic/port INT1/vnic 2
        ena
        bw 50
/c/virt/vnic/port INT2/vnic 1
        ena
        bw 20
/c/virt/vnic/vnicgrp 1
        ena
        vnicvlan 127
        addvnic INT1.2
        addvnic INT2.1
apply
```

You can also use the Industry-standard CLI (isCLI) command script in Example 6-2 to implement this configuration.

Example 6-2 isCLI script to create Scenario 1

```
vnic enable
vnic port INT1 index 2
   enable
   bandwidth 50
vnic port INT2 index 1
   enable
   bandwidth 20
vnic vnicgroup 1
   enable
   vlan 127
   member 1.2
   member 2.1
   exit
```

In addition, you can use the Virtual Fabric Switch browser-based interface (BBI) to implement this configuration:

1. On the **Configure** tab of the BBI, to enable vNIC operations, expand **Virtualization** → **VNIC**, and select **General** (Figure 6-1).

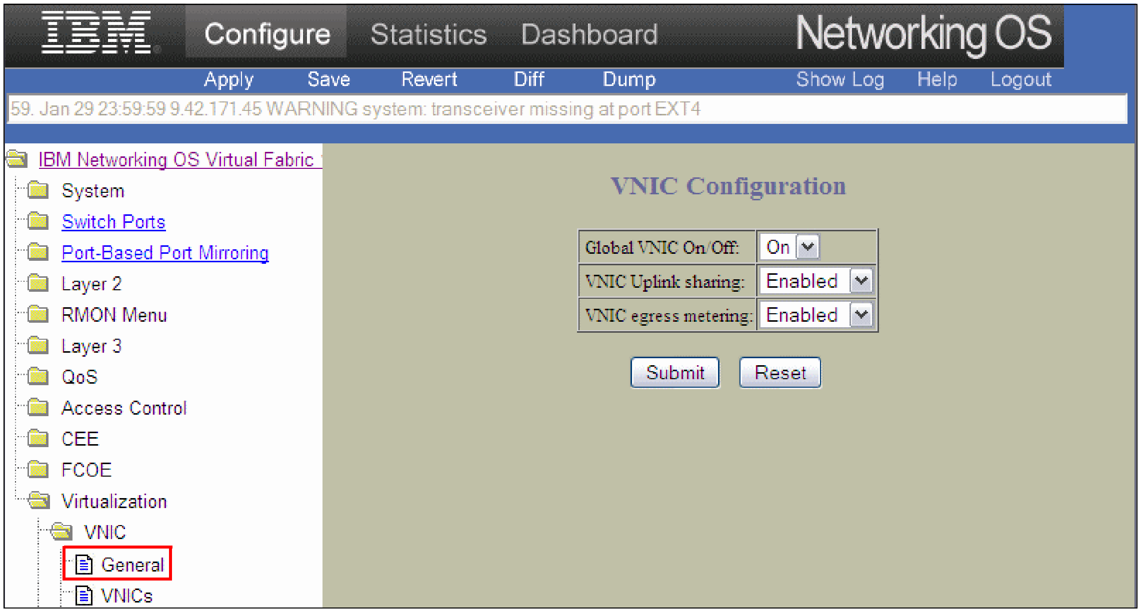

Figure 6-1 Enabling vNIC operations on the vNIC Configuration panel

2. In the VNIC Configuration pane on the right side, for Global VNIC On/Off, select **On**, and then click **Submit**.
3. Select **Apply** → **Save**.

> **Tip:** After *every* change, click **Submit**, and then select **Apply** → **Save** to ensure that the setting persists after any switch reboot.

4. Enable and configure each vNIC. Expand **Virtualization** → **VNIC**, and then select **VNICs** (Figure 6-2).

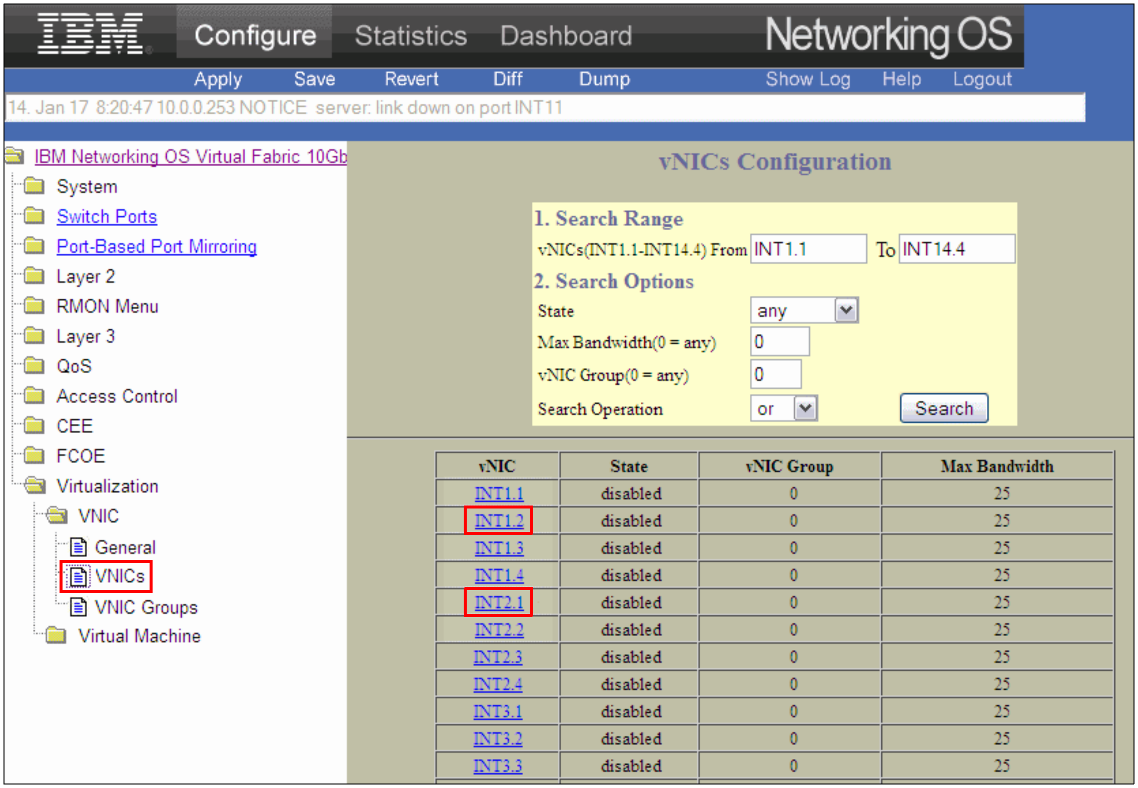

Figure 6-2 vNICs Configuration panel

5. For the first vNIC (Figure 6-3), for State, select **enabled**. Then enter the Max Bandwidth allocation. In this example, we enter 50. Then click **Submit**.

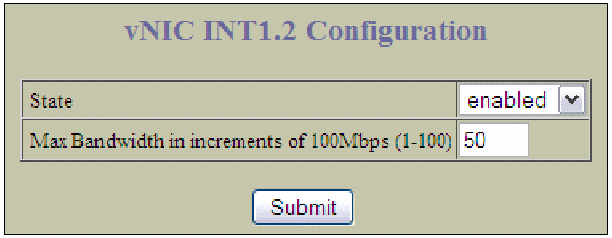

Figure 6-3 Enabling vNIC INT1.2

6. For the second vNIC (Figure 6-4), for State, select **enabled**, and then enter the Max Bandwidth. Again, in this example, we enter 50. Then click **Submit**.

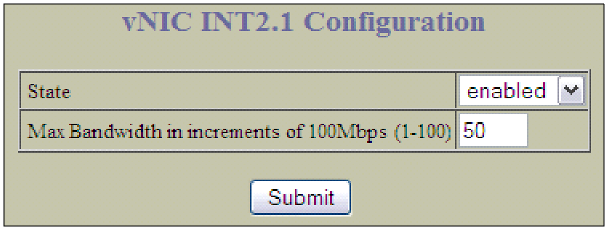

Figure 6-4 Enabling vNIC INT2.1

7. Enable and configure the vNIC group. Expand **Virtualization** → **VNIC**, and then select **VNIC Groups** (Figure 6-5).
8. In the right pane, click **vNIC Group 1**.

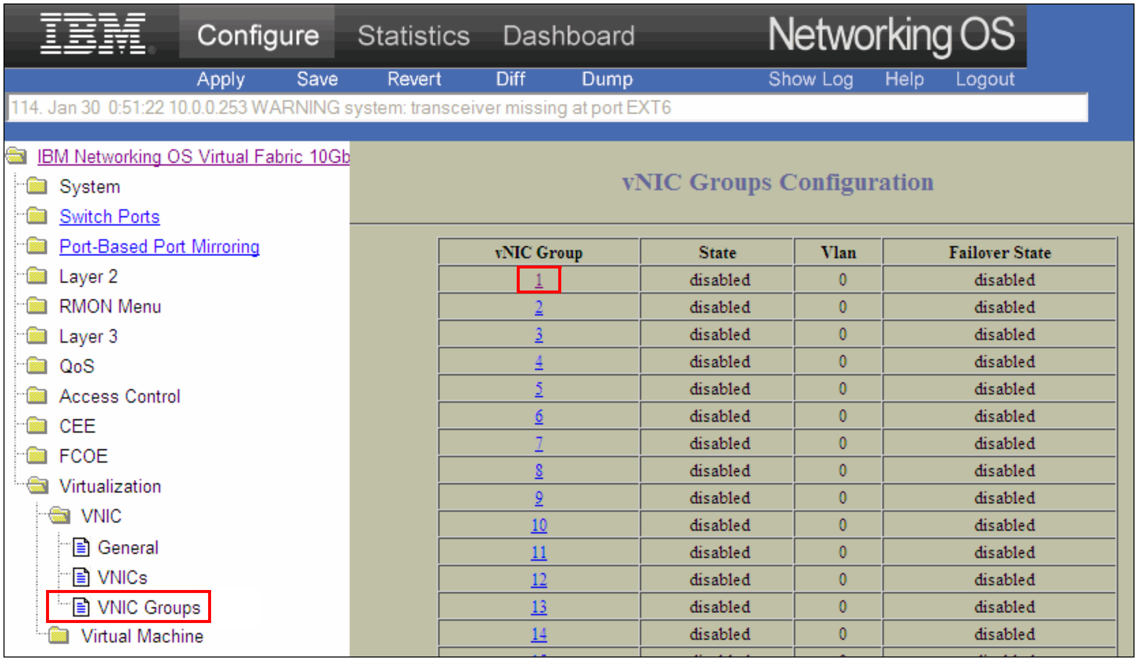

vNIC Group	State	Vlan	Failover State
1	disabled	0	disabled
2	disabled	0	disabled
3	disabled	0	disabled
4	disabled	0	disabled
5	disabled	0	disabled
6	disabled	0	disabled
7	disabled	0	disabled
8	disabled	0	disabled
9	disabled	0	disabled
10	disabled	0	disabled
11	disabled	0	disabled
12	disabled	0	disabled
13	disabled	0	disabled
14	disabled	0	disabled

Figure 6-5 vNIC Groups Configuration panel

9. In the vNIC Group 1 Configuration pane (Figure 6-6):
   a. For vNIC Group State, select **enabled**.
   b. For vNIC Group Van, set the VLAN to `127`.
   c. In the vNICS Available area, select the two vNICs. Then click **Add** to move them to the vNICs in Group area.

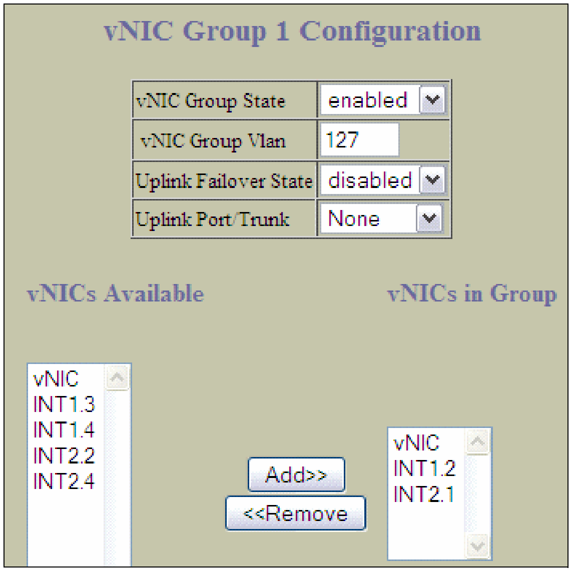

Figure 6-6 vNIC Group 1 Configuration panel

6.2 Scenario 2: vNICs only with one uplink

Scenario 2 creates group VGRP 3. This group consists of vNIC ports and one uplink port and has the following configuration:

- vNIC INT1.3 (blade in slot 1, vNIC 3) of the switch in I/O bay 7, set to a maximum bandwidth of 2.5 Gbps (25% of 10 Gbps)
- vNIC INT2.2 (blade 2, vNIC 2) of the switch in I/O bay 7, set to a maximum bandwidth of 7 Gbps (70% of 10 Gbps)
- One vNIC group that contains the two vNICs (vNIC Group VLAN 500)
- One uplink port, EXT2

You can use the IBM Networking OS command script in Example 6-3 to implement this configuration.

Example 6-3 IBM Networking OS script for Scenario 2

```
/c/virt/vnic
        on
/c/virt/vnic/port INT1/vnic 3
        ena
        bw 25
/c/virt/vnic/port INT2/vnic 2
        ena
        bw 70
/c/virt/vnic/vnicgrp 3
        ena
        vnicvlan 500
        addvnic INT1.3
        addvnic INT2.2? addport EXT2
apply
```

You can also use the isCLI command script in Example 6-4 to implement this configuration.

Example 6-4 isCLI script for Scenario 2

```
vnic enable
vnic port INT1 index 3
   enable
   bandwidth 25 (default)
vnic port INT2 index 2
   enable
   bandwidth 70
vnic vnicgroup 3
   enable
   vlan 500
   member 1.3
   member 2.2
   port EXT2
   exit
```

In addition, you can use the Virtual Fabric Switch BBI to implement this configuration:

1. On the **Configure** tab of the BBI, to enable vNIC operations, expand **Virtualization** → **VNIC**, and select **General** (Figure 6-7).
2. In the VNIC Configuration pane on the right side, for Global VNIC On/Off, select **On**, and then click **Submit**.
3. Select **Apply** → **Save**.

> **Tip:** After *every* change, click **Submit**, and then select **Apply** → **Save** to ensure that the setting persists after any switch reboot.

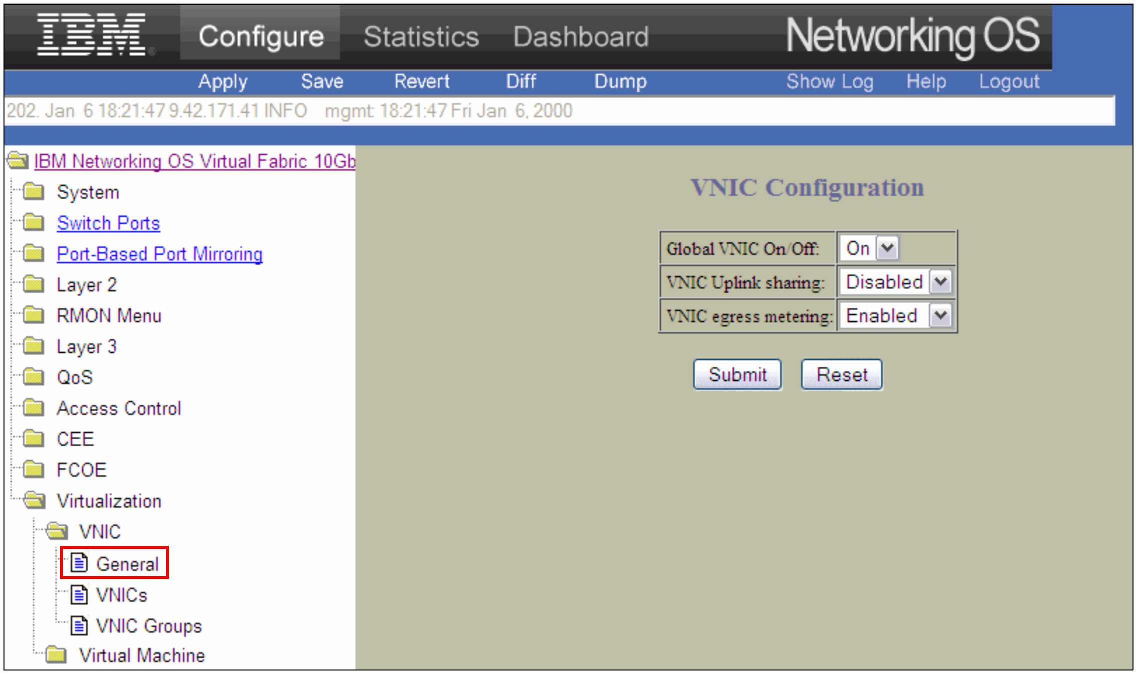

Figure 6-7 Enabling vNIC operations on the vNIC Configuration panel

4. Enable and configure each vNIC. Expand **Virtualization** → **VNIC**, and then select **VNICs** (Figure 6-8).

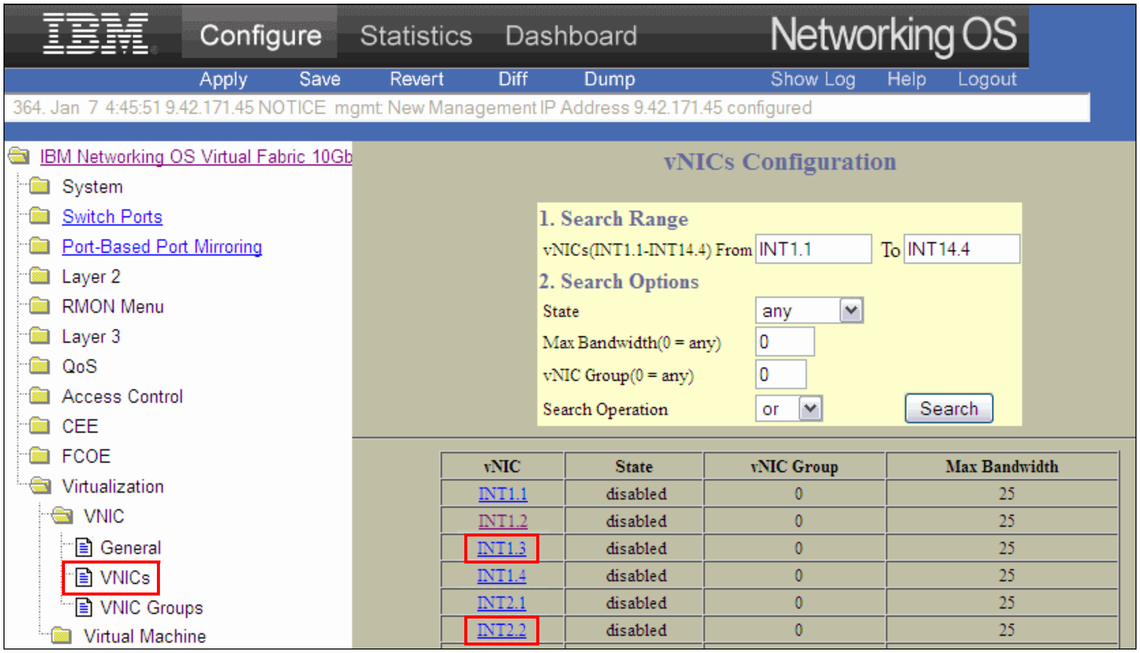

Figure 6-8 vNICs Configuration panel

5. For the first vNIC (Figure 6-9), for State, select **enabled**. Then enter the Max Bandwidth allocation. In this example, we enter `25`. Then click **Submit**.

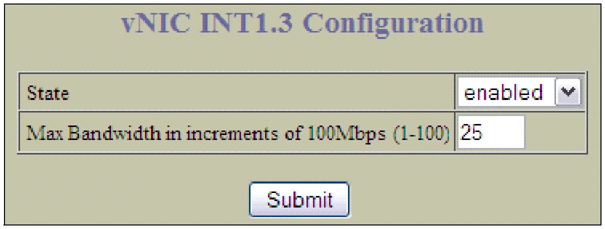

Figure 6-9 vNIC INT1.3 Configuration panel

6. For the second vNIC (Figure 6-10), for State, select **enabled**, and then enter the Max Bandwidth. In this example, we enter `70`. Then click **Submit**.

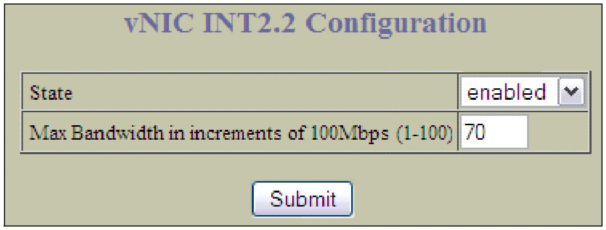

Figure 6-10 vNIC INT2.2 Configuration panel

7. Enable and configure the vNIC group. Expand **Virtualization** → **VNIC**, and then select **VNIC Groups** (Figure 6-11).
8. In the right pane, click **vNIC Group 3**.

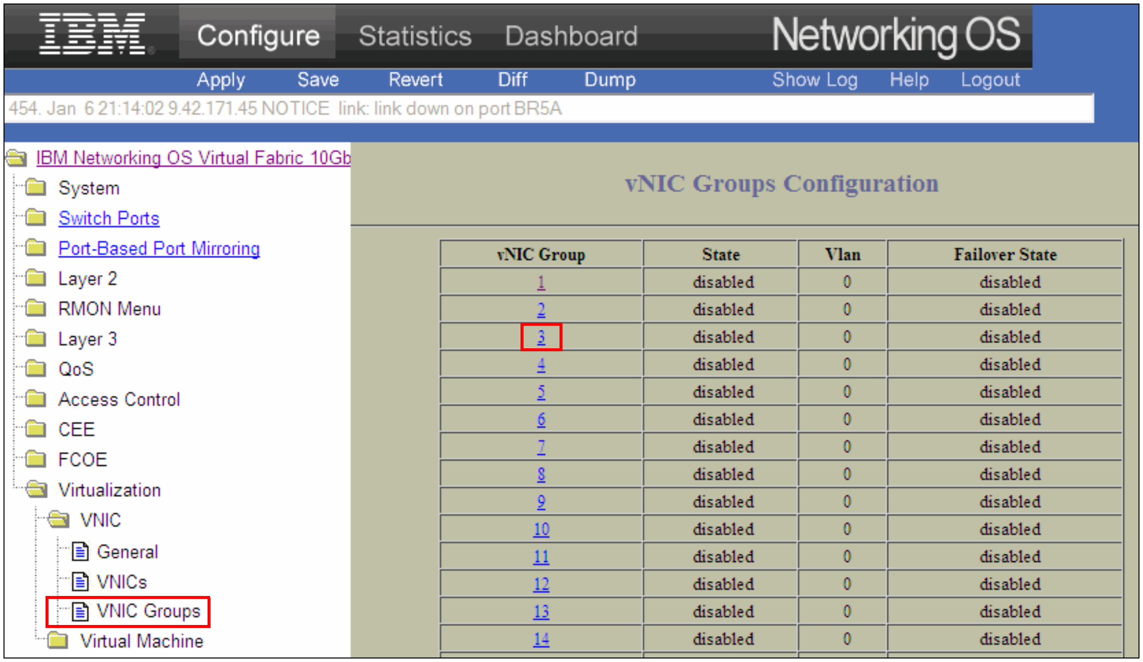

Figure 6-11 vNIC Groups Configuration panel

9. In the vNIC Group 3 Configuration pane (Figure 6-12):
   a. For vNIC Group State, select **enabled**.
   b. For vNIC Group VLAN, set the VLAN to `500`.
   c. For Uplink Port/Trunk, select **EXT2**.
   d. In the vNICS Available area, select the two vNICs. Then click **Add** to move them to the vNICs in Group area.

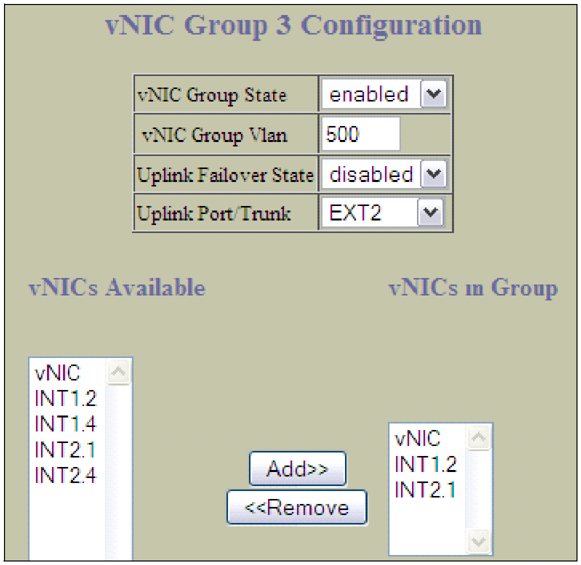

Figure 6-12 vNIC Group 3 Configuration panel

6.3 Scenario 3: vNICs only with one uplink trunk

Scenario 3 creates group VGRP 5. This group consists of vNIC ports and one uplink trunk and has the following configuration:

- vNIC INT1.1 (blade in slot 1, vNIC 1) of the switch in I/O bay 9, set to a maximum bandwidth of 3 Gbps (30% of 10 Gbps)
- vNIC INT2.2 (blade 2, vNIC 2) of the switch in I/O bay 9, set to a maximum bandwidth of 3.3 Gbps (33% of 10 Gbps)
- One vNIC group that contains the two vNICs (vNIC Group VLAN 925)
- One uplink trunk that consists of external ports EXT1 and EXT3

You can use the IBM Networking OS command script in Example 6-5 to implement this configuration.

Example 6-5 IBM Networking OS script for Scenario 3

```
/c/l2/trunk 1
        ena
        add EXT1
        add EXT3
/c/virt/vnic
        on
/c/virt/vnic/port INT1/vnic 1
        ena
        bw 30
/c/virt/vnic/port INT2/vnic 2
        ena
        bw 33
/c/virt/vnic/vnicgrp 5
        ena
        vnicvlan 925
        addvnic INT1.1
        addvnic INT2.2
        addtrnk 1
apply
```

You can also use the isCLI command script in Example 6-6 to implement this configuration.

Example 6-6 isCLI script for Scenario 3

```
vnic enable
portchannel 1 port EXT1
portchannel 1 port EXT3
portchannel 1 enable
vnic port INT1 index 1
   enable
   bandwidth 30
vnic port INT2 index 2
   enable
   bandwidth 33
vnic vnicgroup 5
   enable
   vlan 925
   member 1.1
   member 2.2
   trunk 1
   exit
```

In addition, you can use the Virtual Fabric Switch BBI used to implement this configuration:

1. On the **Configure** tab of the BBI, to enable vNIC operations, expand **Virtualization** → **VNIC**, and then select **General** (Figure 6-13).
2. In the VNIC Configuration pane on the right side, for Global VNIC On/Off, select **On**, and then click **Submit**.
3. Select **Apply** → **Save**.

> **Tip:** After *every* change, click **Submit**, and then select **Apply** → **Save** to ensure that the setting persists after any switch reboot.

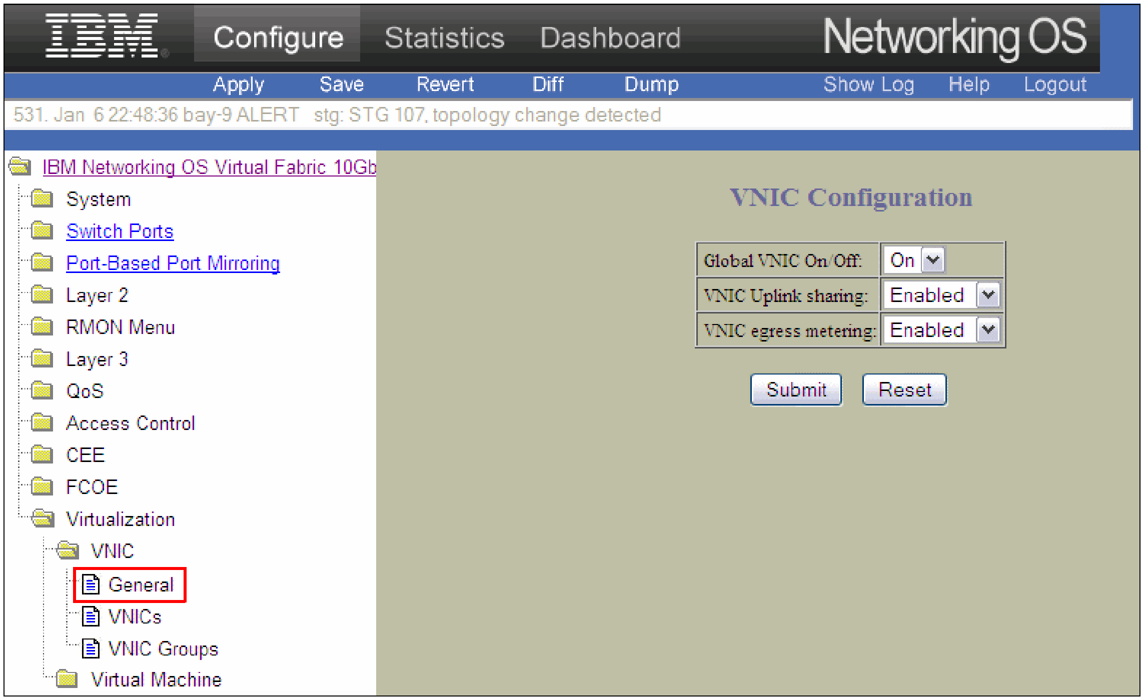

Figure 6-13 Enabling vNIC operations on the vNIC Configuration panel

4. Enable and configure each vNIC. Expand **Virtualization** → **VNIC**, and then select **VNICs** (Figure 6-14).

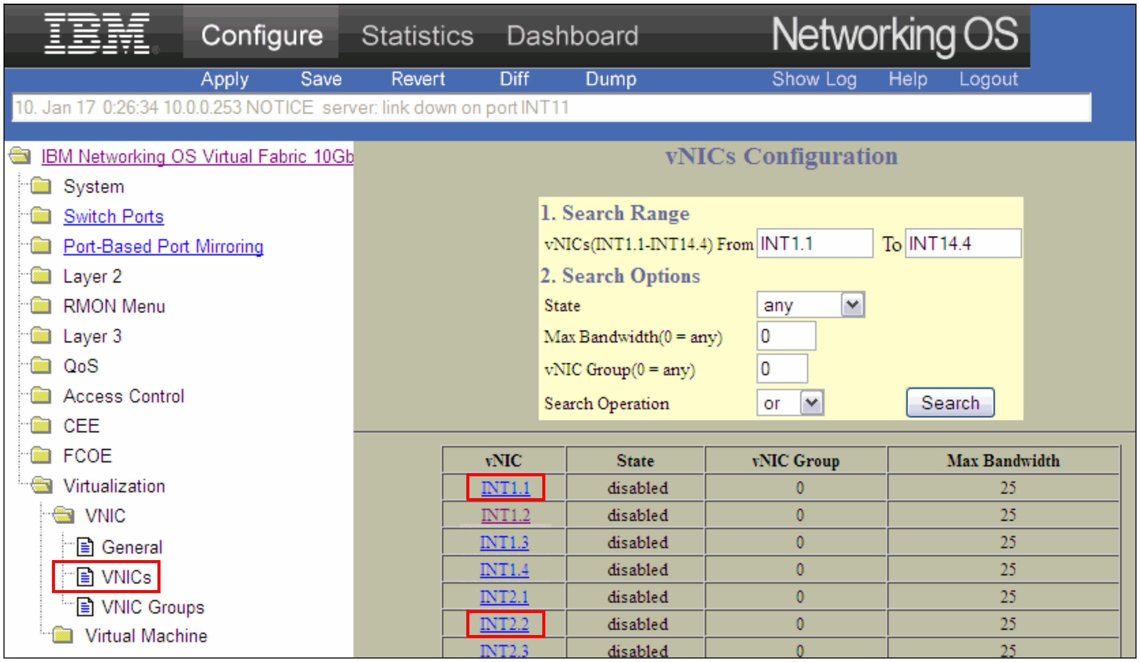

Figure 6-14 vNICs Configuration panel

5. For the first vNIC (Figure 6-15), for State, select **enabled**. Then enter the Max Bandwidth allocation. In this example, we enter 30. Then click **Submit**.

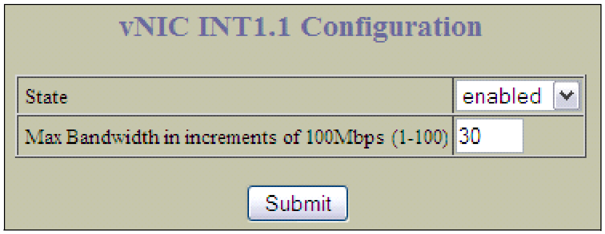

Figure 6-15 Scenario 5: vNIC INT1.1 Configuration panel

6. For the second vNIC (Figure 6-16), for State, select **enabled**, and then enter the Max Bandwidth. In this example, we enter 33. Then click **Submit**.

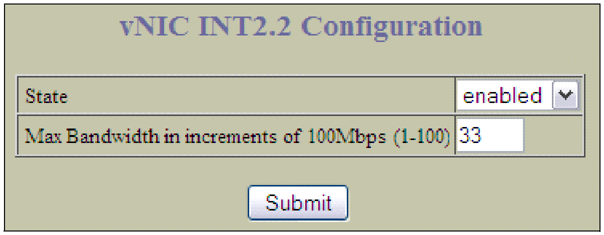

Figure 6-16 vNIC INT2.2 Configuration panel

7. Enable the trunk group. Expand **Layer 2,** and select **Trunk Groups**.
8. In the Trunk Groups Configuration pane (Figure 6-17), click **Trunk Group 1**.

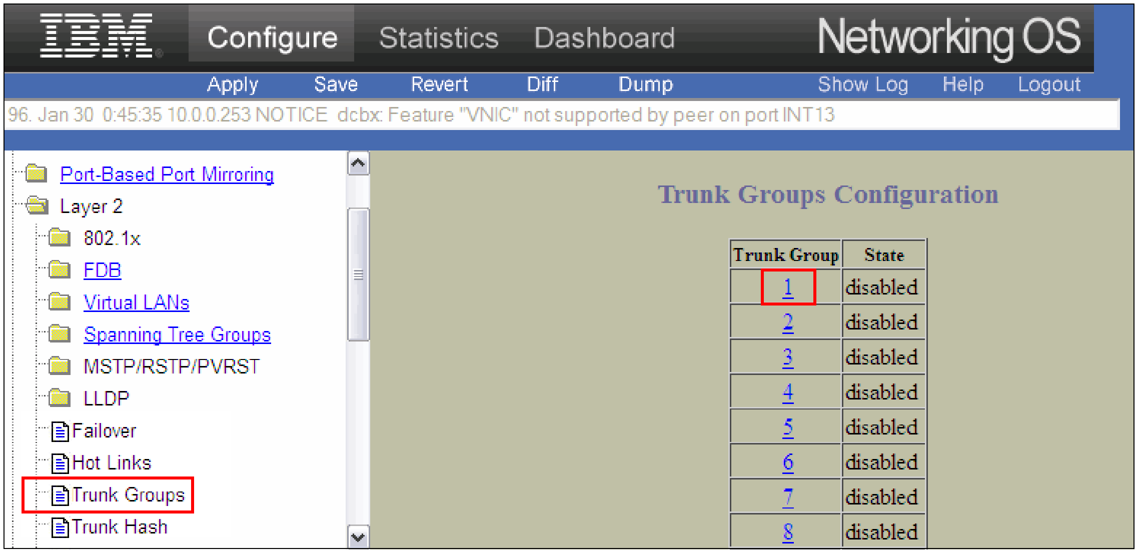

Figure 6-17 Scenario 5: Trunk Groups Configuration panel

9. In the Switch Trunk Group 1 Configuration pane (Figure 6-18):
   a. Set Trunk State to **Enabled**.
   b. In the Ports Available area, select **EXT1** and **EXT3**. Then click **Add** to move them to the Ports added to Trunk area.
   c. Click **Submit**.

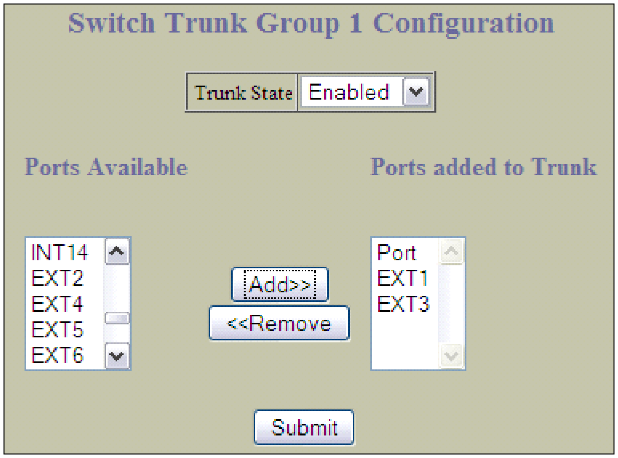

Figure 6-18 Switch Trunk Group 1 Configuration pane

10. Enable and configure the vNIC group. Expand **Virtualization** → **VNIC**, and then select **VNIC Groups** (Figure 6-19).

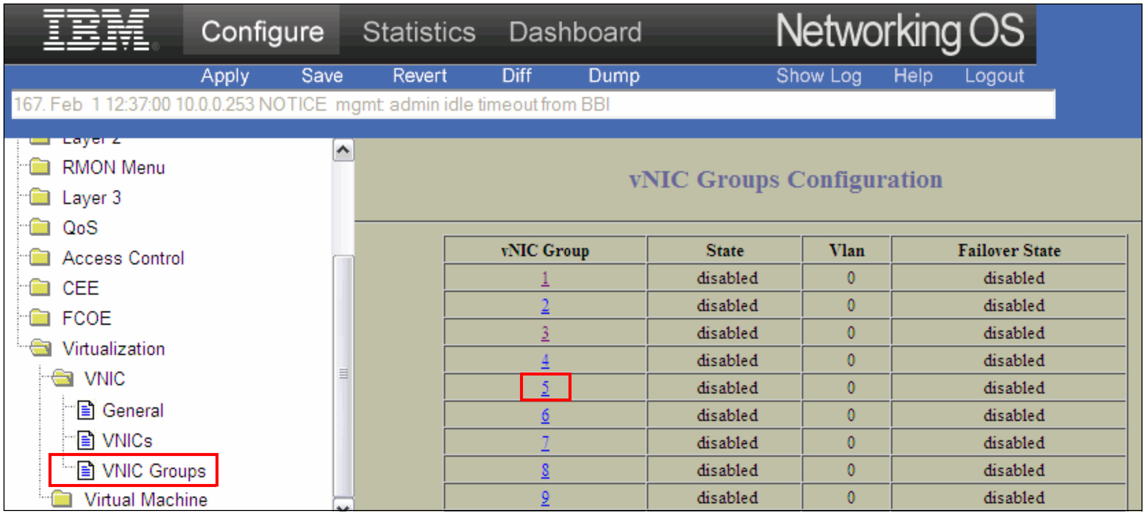

Figure 6-19 vNIC Groups Configuration pane

11. In the vNIC Groups Configuration pane, click **vNIC Group 5**.
12. In the vNIC Group 5 Configuration pane (Figure 6-20):
    a. For vNIC Group State, select **enabled**.
    b. For vNIC Group VLAN, set the VLAN to `925`.
    c. For Uplink Port/Trunk, select **Trunk 1**.
    d. In the vNICS Available area, select the two vNICs. Then click **Add** to move them to the vNICs in Group area.

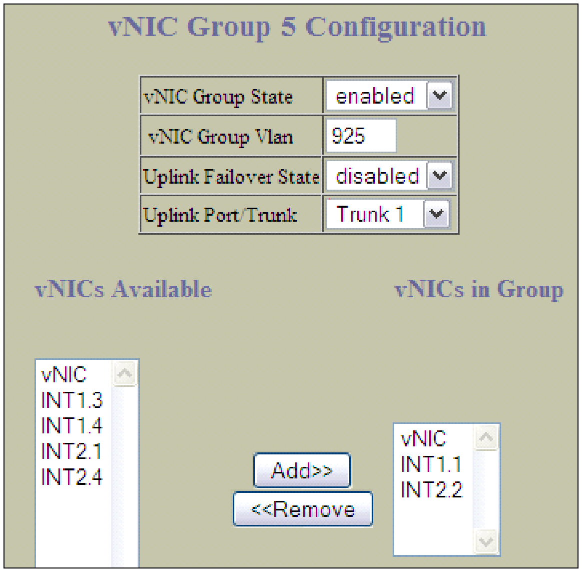

Figure 6-20 vNIC Group 5 Configuration pane

6.4 Scenario 4: vNIC FCoE on Emulex II, BNT, and QLogic with IBM DS5300

This section highlights the steps for implementing the Virtual Fabric Mode for Fibre Channel over Ethernet (FCoE) on the following adapter and modules:

- Emulex Virtual Fabric Adapter II
- BNT Virtual Fabric 10Gb Switch Module
- QLogic Virtual Fabric Extension Module connecting to the IBM DS5300

6.4.1 Configuring the QLogic Virtual Fabric Extension Module

The QLogic Virtual Fabric Extension Module is a critical part of the IBM first in-chassis FCoE solution. When paired with the BNT 10GE Switch Module, it provides the Fibre Channel Forwarder (FCF) piece of the overall solution. It strips the IP part from the Fibre Channel (FC) part in the Converged Enhanced Ethernet (CEE) packet. It then sends the FC packet to either the SAN FC fabric or any FC Storage to which it might be attached.

When matched with the BNT 10GE Switch Module, the QLogic Virtual Fabric Extension Module provides IBM with its first in-chassis FCoE solution. It eliminates the need for, and cost of, an external top-of-rack FCoE switch. The easy-to-install module provides the key component to the full FCoE solution, which is the FCF. Its four internal ports attach to four ports of the BNT Switch Module to provide the BladeCenter six ports of SAN connectivity. In addition, the BNT switch provides six ports of IP connectivity.

The QLogic Virtual Fabric Extension Module (and BNT 10GE Switch Module) provides six external FC connections between BladeCenter servers with QLogic 2-port converged network adapter (CNA) CFFh Expansion Cards and external SAN fabrics and FC devices.

6.4.2 Implementing the FCoE-enabled BladeCenter with Virtual Fabric Switch and Extension Module

This section explains how to enable FCoE host access to the FC SAN-attached DS5300 storage. This procedure uses the QLogic Virtual Fabric Extension Module and BNT Virtual Fabric 10Gb Switch Modules installed in the BladeCenter H chassis.

Figure 6-21 on page 122 shows the I/O topology that is internal to the BladeCenter H chassis. Bridge bays 3 and 5 have internal connections to the high-speed I/O bays 7 and 9.

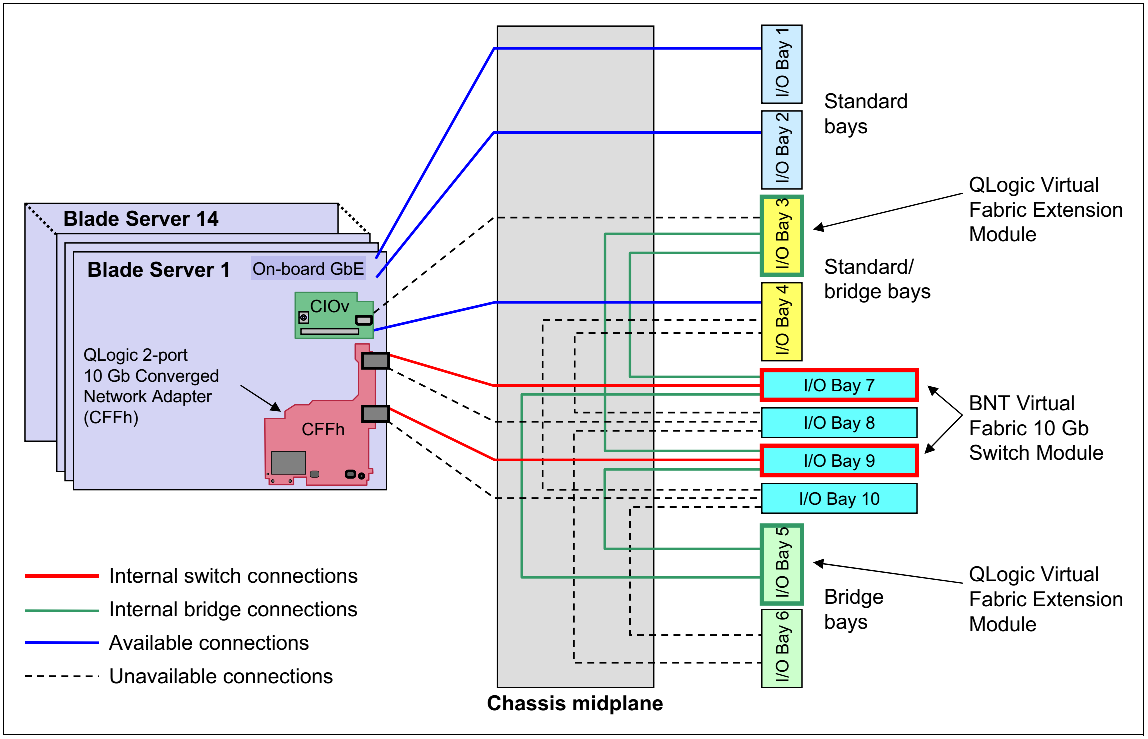

Figure 6-21 BladeCenter H internal connections with the QLogic Virtual Fabric Extension Modules

The QLogic CNA installed on the blade server connects to BNT 10G Virtual Fabric Switch Module in bays 7 and 9. The QLogic Virtual Fabric Extension Module that is installed in bays 3 and 5 has internal connections to the BNT 10G switch modules.

The FC storage is directly connected to the external ports on the QLogic Virtual Fabric Extension Module. Alternatively, it is connected through a basic FC switch that is connected to the QLogic Virtual Fabric Extension Module running in transparent mode (N_Port Virtualization (NPV)).

The QLogic Virtual Fabric Extension Module acts as an FCF. It decapsulates and encapsulates the FCoE frames and forwards the frames to FC devices that are attached to the external ports. Eight internal ports are on the QLogic Virtual Fabric Extension Module, but a maximum of four ports can be active at any time. The QLogic Virtual Fabric Extension Module can function as a full fabric switch. This module provides basic FC functions, such as FLOGI, Name Server, and Address assignment. It can also be configured in transparent mode to connect to external top-of-the-rack full fabric switch. The default configuration on the QLogic Virtual Fabric Extension Module is full fabric to allow support for direct-attach FC storage on the external ports.

Table 6-1 lists the supported configurations with the QLogic Virtual Fabric Extension Module and BNT 10G switch installed in the BladeCenter H chassis.

Table 6-1 Ethernet to bridge port mapping

Configuration	Switch bay-7	Switch bay-9	Connects to:	Ports on the BNT Virtual Fabric 10Gb Switch Module
A	20/40 Gbps		Bridge bay-5 only	BR5A-BR5D replaces EXT1-EXT4.
B		20/40 Gbps	Bridge bay-3 (or bay-5)	BR3D-BR3A replaces EXT7-EXT10.
C	40 Gbps	40 Gbps	Bridge bay-5 and bay-3	BR5A-D replaces EXT1-4, and BR3D-A replaces EXT7-10.
D	20 Gbps	20 Gbps	Bridge bay-5 only	BR5A replaces EXT1, and BR5B replaces EXT2.

The maximum supported internal bandwidth between each QLogic Virtual Fabric Extension Module and the BNT 10G switch module is 40 Gbps. It uses four 10-Gbps external ports on the high-speed switch module installed in bay 7 or 9. The minimum supported bandwidth between each QLogic Virtual Fabric Extension Module and the BNT 10G switch module is 20 Gbps. It uses two 10-Gbps external ports on the high-speed switch module installed in bay 7 or 9.

Tip: For configuration of option D in Table 6-1 with two BNT Virtual Fabric 10Gb Switch Modules and one QLogic Virtual Fabric Extension Module, you can configure two bridge ports for each switch module. In making this configuration, you achieve a total bandwidth of 40 Gbps by distributing it across the two switch modules.

The bridge ports on the QLogic Virtual Fabric Extension Module use the same internal path as the external ports on the 10 Gbps BNT switch module. Therefore, enabling the bridge ports results in automatically disabling some of the external ports on the BNT 10 Gbps switch module. If the bandwidth assigned for the bridge ports is 40 Gbps, four external ports are disabled. If the bandwidth assigned on the bridge ports is 20 Gbps, two external ports on the BNT 10 Gbps switch module are disabled.

At a high level, complete the following checklist to establish end-to-end FCoE connectivity:

- On the host, install and connect an FCoE adapter (CNA) to the BNT Virtual Fabric 10G Switch Module switch module that is installed in bay 7 and 9.
- On the host, install the Ethernet and FC device drivers.

- Ensure that the FCoE host has a link on the 10 Gbps CEE ports.
- Ensure that the FC device driver on the host logs in to the QLogic Virtual Fabric Extension Module as a VN_port.
- Connect the FC storage device to the fabric.
- Ensure that the FC target devices are online as N port devices.
- Zone the FCoE host and FC target worldwide port names (WWPNs).
- Map the FC LUN to the host.
- Ensure that the host detects the LUNs from the disk manager.

The deployment of the Virtual Fabric FCoE solution needs minimum manual configuration.

On the BNT Virtual Fabric 10G Switch Module, to create the configuration shown in Figure 6-21 on page 122:

1. Create the bridge interfaces:
   a. Assign bandwidth.
   b. Enable the bridge interfaces.
   c. Reboot the switch.
2. Enable CEE mode.
3. Enable Federal Information Processing Standard (FIPS).
4. Enable Link Layer Discovery Protocol (LLDP).
5. Create an FCoE VLAN (1002).
   a. Assign the FCoE VLAN to the appropriate server ports (internal interfaces).
   b. Set the FCoE VLAN as the Port VLAN ID (PVID) on the bridge interfaces.
6. Enable tagging on the bridge interfaces, and tag the PVID.
7. Disable spanning tree on the bridge interfaces.

On the QLogic Virtual Fabric Extension Modules, to configure the network connectivity across the devices:

1. Configure the FCoE VLAN if one any other than VLAN 1002 is used.
2. Configure and enable zoning, or set it to **Transparent Mode**.

To configure the FC storage subsystem, define host and LUN masking on the storage system.

On the blade server, to connect to the storage disk:

1. Install the MPIO driver.
2. Discover and configure the disks.

6.4.3 Defining the FCoE and FC fabric topology

This IBM Redbooks publication uses two QLogic Virtual Fabric Extension Modules and two IBM Virtual Fabric Switch Modules to set up a fully redundant topology with two separate FCoE and FC SAN fabrics. In regard to Table 6-1 on page 123, this book uses configuration scenario D.

In this case, you can configure the Virtual Fabric Extension Module by using either of the following options:

- A full fabric switch with its own domain ID providing Fibre Channel basic functions, which include Name Server and Zoning
- In Transparent mode to connect to an external FC fabric (Figure 6-22)

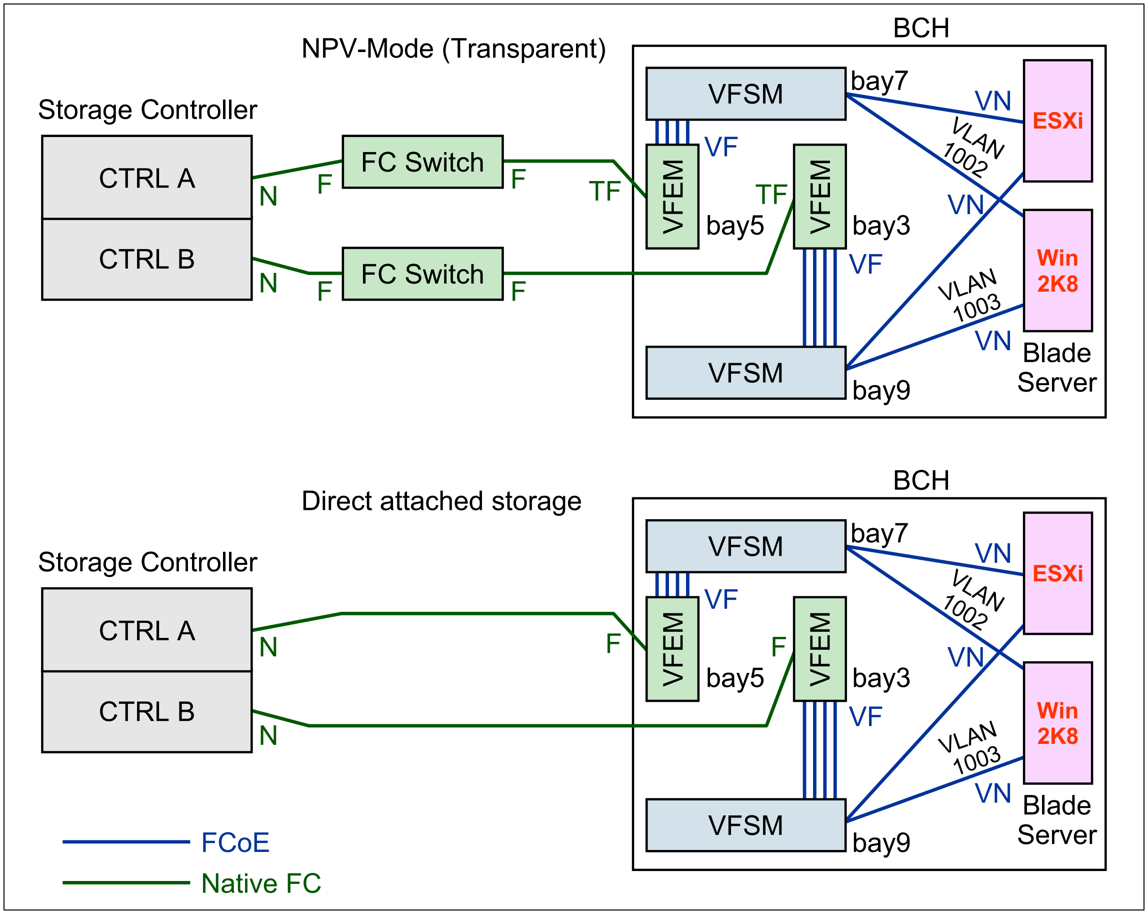

Figure 6-22 FCoE and FC fabric topology

The full fabric provides support for direct attach FC storage on the external ports. Transparent mode might be used to avoid concerns about interoperability or to scale the FC fabric without adding domain IDs.

The Virtual Fabric Extension Module Transparent mode is based on NPV so that the upstream Fibre Channel switch must support NPIV. The TF_Port (Transparent Fabric Port, which corresponds to the NP_Port in NPV) acts similar to an NPIV-enabled host to the upstream F_Port.

In either case, FCoE is kept internal to the BladeCenter, with Virtual N_Ports (VN) for the end-node port of the FC and FCoE fabric and Virtual F_Ports (VF) as the fabric ports on the switch side. The external Fibre Channel connections are all running at 8 Gbps.

As shown in Figure 6-22 on page 125, the switch in bay 7 will connect to the extension module in bay 5, and the switch in bay 9 will connect to the extension module in bay 3. Each connection results in 40 Gbps of aggregated bandwidth, which in sum, is 80 Gbps full duplex for FCoE traffic to the FC Gateway. Although the BNT Virtual Fabric 10 Gb Switch Module has 10 external 10 Gb ports, this configuration allows use of up to only six external ports on each Virtual Fabric Switch Module. Four external ports of each switch module are disabled (ports EXT1-4 in bay 7 and ports EXT7-10 in bay 9). Those circuits are rerouted to each QLogic Virtual Fabric Extension Module as bridge ports BR5A-BR5D bay 7 and BR3D-BR3A in bay 9.

6.4.4 Configuring the BNT Virtual Fabric 10Gb Switch Modules

To configure the BNT Virtual Fabric 10Gb Switch Modules:

1. Connect the BNT Virtual Fabric 10Gb Switch Module to the Virtual Fabric Extension Module (also called *bridge module*).

   Initiate this connection from the BNT Virtual Fabric 10Gb Switch Module with the options to choose the bandwidth (20 or 40 Gbps) and the location of the Virtual Fabric Extension Module (bay 3 or bay 5). Example 6-7 shows the configuration steps for a connection from the BNT Virtual Fabric 10Gb Switch Module in bay 7 to Virtual Fabric Extension Module in bay 5 with the maximum bandwidth of 40 Gbps.

   Example 6-7 Connection configuration between Virtual Fabric Switch Module and Virtual Fabric Extension Module

   ```
   bay-7>enable
   bay-7#conf t
   bay-7(config)#boot bridge-module 5 bandwidth 40
   Please disable spanning tree on the bridge module 5 ports BR5A,
   BR5B, BR5C and BR5D after the switch reboots.
   ```

```
bay-7(config)#boot bridge-module 5 enable
Bridge module 5 is enabled. Reboot the switch in order for the
settings to take effect.
bay-7(config)#reload
```

2. Reload the BNT Virtual Fabric 10Gb Switch Module as shown in Example 6-7 on page 126.
3. Check the connectivity settings by entering the **show bridge-module** command.
4. View the active bridge ports, which are listed in place of the corresponding external ports, by using the **show interface link** command (Example 6-8).

Example 6-8 Verifying Virtual Fabric Extension Module connectivity

```
bay-7#show bridge-module
Bridge module 5 is set to 40Gbs.
Bridge module 3 is disabled.
bay-7#
bay-7#show interface link
------------------------------------------------------------------
Alias   Port   Speed    Duplex     Flow Ctrl      Link
------- ----   -----   --------  --TX-----RX--   ------
[...]
BR5A     17     10000     full      no     no       up
BR5B     18     10000     full      no     no       up
BR5C     19     10000     full      no     no       up
BR5D     20     10000     full      no     no       up
EXT5     21     10000     full      no     no       up
EXT6     22     10000     full      no     no       up
[...]
------------------------------------------------------------------
Alias   Speed
-----   -----
BR3     40Gbs
bay-7#
```

5. Enable the Ethernet enhancements that are summarized as Data Center Bridging (DCB) or CEE by using the **cee enable** command.

   This **cee enable** command automatically enables LLDP, which is required for the DCB Exchange protocol (DCBX). The **cee enable** command also disables the previous Ethernet flow control, because priority-based flow control (PFC) is used instead.

   The settings for the bandwidth management (Enhanced Transmission Selection (ETS)) and PFC might be left on the default values. That is, for

FCoE, traffic has a 50% guaranteed bandwidth and priority flow control for Ethernet class of service (COS) 3.

All these configuration parameters are transferred to the CNA by the DCBX Protocol, so that nothing has to be configured on the CNA itself.

6. Enable the FCoE part, which relies on two protocols:

   – FCoE, as the data plane protocol, carries the FC command frames and the SCSI traffic.

   – FCoE Initialization Protocol as the control plane protocol, which is used for the following reasons:

     - VLAN Discovery
     - FCF Discovery
     - FLOGI/FDISC
     - KeepAlives

   The BNT Virtual Fabric 10Gb Switch Module is passing only the FCoE frames. However, to control FCoE traffic by dynamically creating access control lists (ACLs), you must have knowledge about the active FCoE session. You must also understand what is obtained by snooping with the FCoE Initialization Protocol, which you can configure by using the **fc0e fips enable** command.

   The configuration of the FC Gateway is explained in 6.4.6, "Switching the Virtual Fabric Extension Module to N-Port Virtualization mode if connected to an existing FC fabric" on page 137.

7. Configure a dedicated VLAN for FCoE, which by default is 1002, with all the internal ports and all bridge ports as members.

   The FCoE VLAN is sent tagged by the CNA, which corresponds to the default setup of the BNT Virtual Fabric 10Gb Switch Module.

8. Configure the bridge ports as **tagged** ports to ensure communication to the Virtual Fabric Extension Module.

   The **tag-pvid** command ensures that only tagged frames are transmitted.

   Example 6-9 summarizes the commands for the BNT Virtual Fabric 10Gb Switch Module in bay 7.

Example 6-9 Enabling the Ethernet enhancements

```
bay-7bay-7(config)#cee enable
### bay-7(config)#lldp enable - done automatically with "cee enable"
###
bay-7(config)#fcoe fips enable
bay-7(config)#vlan 1002
bay-7(config-vlan)#member INT1-14,BR5A-BR5D
```

```
bay-7(config-vlan)#enable
bay-7(config)#interface port BR5A-BR5D
bay-7(config-if)#pvid 1002
bay-7(config-if)#tagging
bay-7(config-if)#tag-pvid
bay-7(config)#exit
```

9. To ensure cleaner network traffic:

   a. Remove all bridge ports from any other VLAN than the FCoE VLAN. Example 6-10 shows the commands for VLAN 1.

   Example 6-10 Removing bridge ports from VLAN 1

   ```
   bay-7(config)#vlan 1
   bay-7(config-vlan)#no member BR5A-BR5D
   bay-7(config)#exit
   ```

   b. Disable Spanning Tree Protocol on all bridge ports as shown in Example 6-11.

   Example 6-11 Disabling STP for all bridge ports

   ```
   bay-7(config)#int port BR5A-BR5D
   bay-7(config-if)#no spanning-tree stp 106 enable
   bay-7(config)#exit
   ```

   > **Tip:** Check the Spanning Tree instance that is used for the FCoE VLAN by using the **`show running | i stp`** command.

The necessary configuration tasks on the BNT Virtual Fabric 10Gb Switch Module in I/O bay 7 are now completed.

Repeat the same configuration process on the BNT Virtual Fabric 10Gb Switch Module in I/O bay 9. This time, connect to Virtual Fabric Extension Module I/O bay 3, instead of I/O bay 5, and VLAN 1003, instead of VLAN 1002 as FCoE VLAN.

Example 6-12 shows the final running configuration from the BNT Virtual Fabric 10Gb Switch Module in bay 7.

Example 6-12 Configuration of BNT Virtual Fabric 10Gb Switch Module

```
bay-9#sh run
Current configuration:
version "6.8.0.66"
```

```
switch-type "IBM Networking OS Virtual Fabric 10Gb Switch Module for
IBM BladeCenter"
!
snmp-server name "bay-9"
hostname "bay-9"
!
interface port INT1
pvid 99
no flowcontrol
exit
!
!
! ### repeats for interface port INT2-INT13 ###
!
!
interface port INT14
pvid 99
no flowcontrol
exit
!
interface port BR3D
tagging
tag-pvid
pvid 1003
exit
!
interface port BR3C
tagging
tag-pvid
pvid 1003
exit
!
interface port BR3B
tagging
tag-pvid
pvid 1003
exit
!
interface port BR3A
tagging
tag-pvid
pvid 1003
exit
!
vlan 1
```

```
member INT1-INT14,EXT1-EXT6,EXT11
no member BR3D-BR3A
!
vlan 99
enable
name "VLAN 99"
member INT1-INT14
!
vlan 1003
enable
name "VLAN 1003"
member INT1-INT14,BR3D-BR3A
!
spanning-tree stp 1 vlan 1
spanning-tree stp 1 vlan 99

spanning-tree stp 107 vlan 1003

interface port BR3D
no spanning-tree stp 107 enable
exit
interface port BR3C
no spanning-tree stp 107 enable
exit
interface port BR3B
no spanning-tree stp 107 enable
exit
interface port BR3A
no spanning-tree stp 107 enable
exit
!
fcoe fips enable
cee enable
lldp enable
!
end
bay-9#
```

Use the **show interface links** command to verify that the link status is UP for all active host ports and bridge interfaces.

6.4.5 Configuring the QLogic Virtual Extension Modules

By default, the Virtual Fabric Extension Module has the necessary Ethernet configuration built-in. This configuration enables FCoE hosts to log in to the name server if the FCoE configuration on the Virtual Fabric Switch Modules was completed by using the default VLAN ID 1002 for FCoE.

All basic configurations tasks are completed. Only the nondefault FCoE VLAN ID 1003 must be configured on the Virtual Fabric Extension Module in bay 3. A different VLAN ID to 1002 might be necessary if redundant FC fabrics with separate VLANs are used for the FC fabrics. A violation of the customer’s compliance rules is another reason for choosing a nondefault VLAN ID.

To configure the Virtual Fabric Extension Module FCoE VLAN:

1. Point your browser to the IP address of the Virtual Fabric Extension Module.

> **Tip:** Use the default login of USERID and PASSW0RD, with a zero in the word “Password” for the letter *O*.

2. To access the VLAN Manager, select **Switch** → **FCoE** → **VLAN Manager** (Figure 6-23).

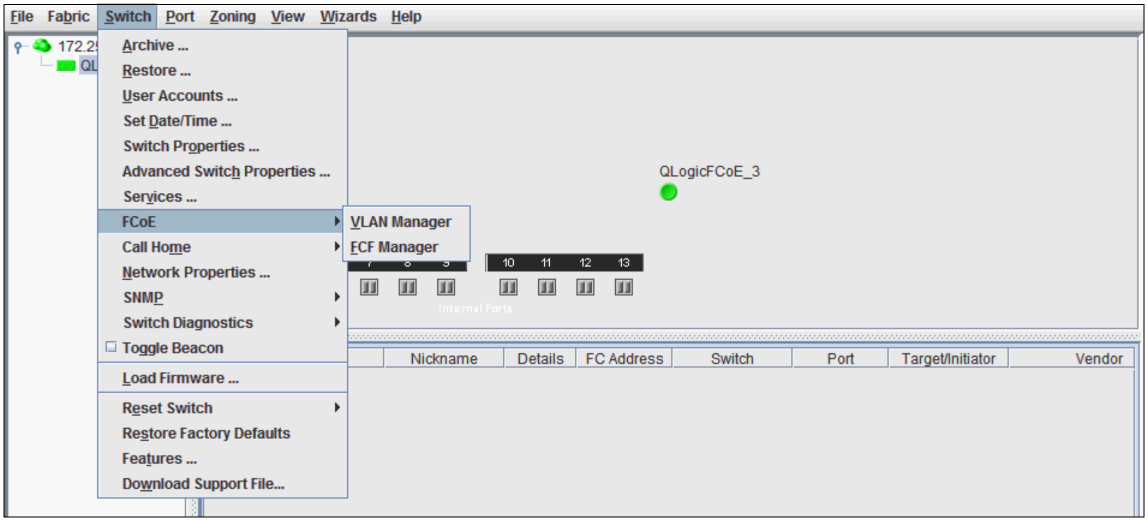

Figure 6-23 Virtual Fabric Extension Module GUI

3. In the VLAN Manager window (Figure 6-24):
   a. Click **Add VLAN**.
   b. Enter 1003 as the additional VLAN ID for FCoE.
   c. Back in the VLAN Manager Window, click the arrows buttons next to ports SM 9A:10, SM 9B:11, SM 9C:12, SM 9D:13, and select **VLAN 1003** for each one.
   d. Click **Apply**.

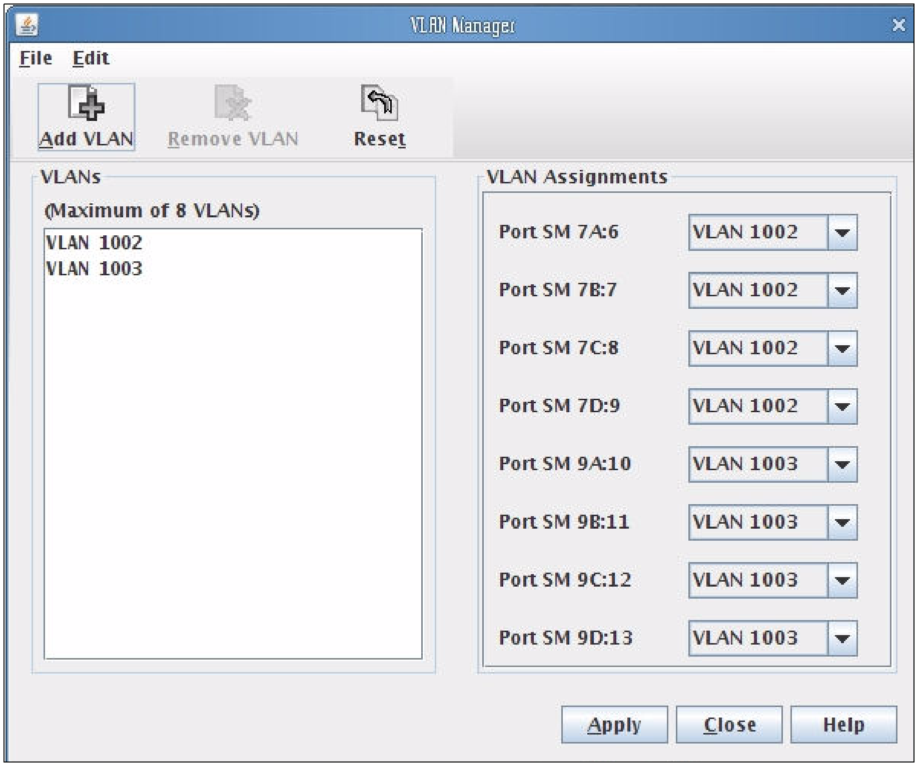

Figure 6-24 Virtual Fabric Extension Module VLAN Manager

4. Select **Switch** → **FCoE** → **FCF Manager**.

5. In the FCF Configuration Manager window (Figure 6-25), in the VLAN List box, change the FCoE mapping to VLAN 1003.

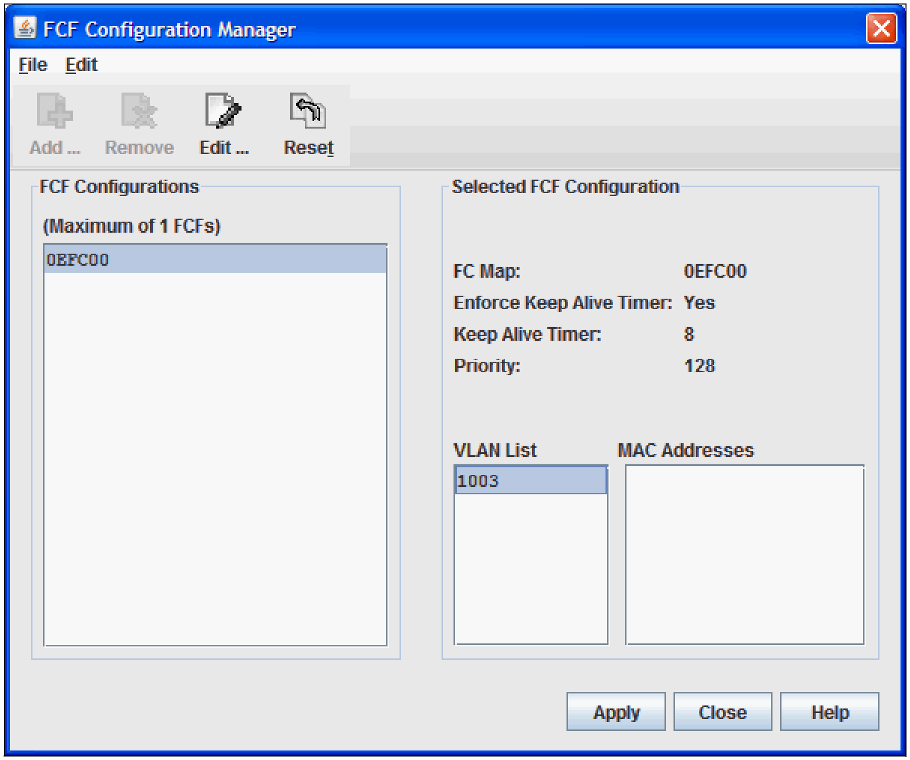

Figure 6-25 Virtual Fabric Extension Module FCF Configuration Manager

In the FCF Configurations list box, if you select **0EFC00**, and then click **Edit**, you see the results in the FCF Editor as shown in Figure 6-26 on page 135.

6. In the FCF Editor window (Figure 6-26):
   a. Select **VLAN ID 1002**, and click **Remove**.
   b. From the VLAN list, select **VLAN 1003l**.
   c. Click **Add**, and then click **OK**.

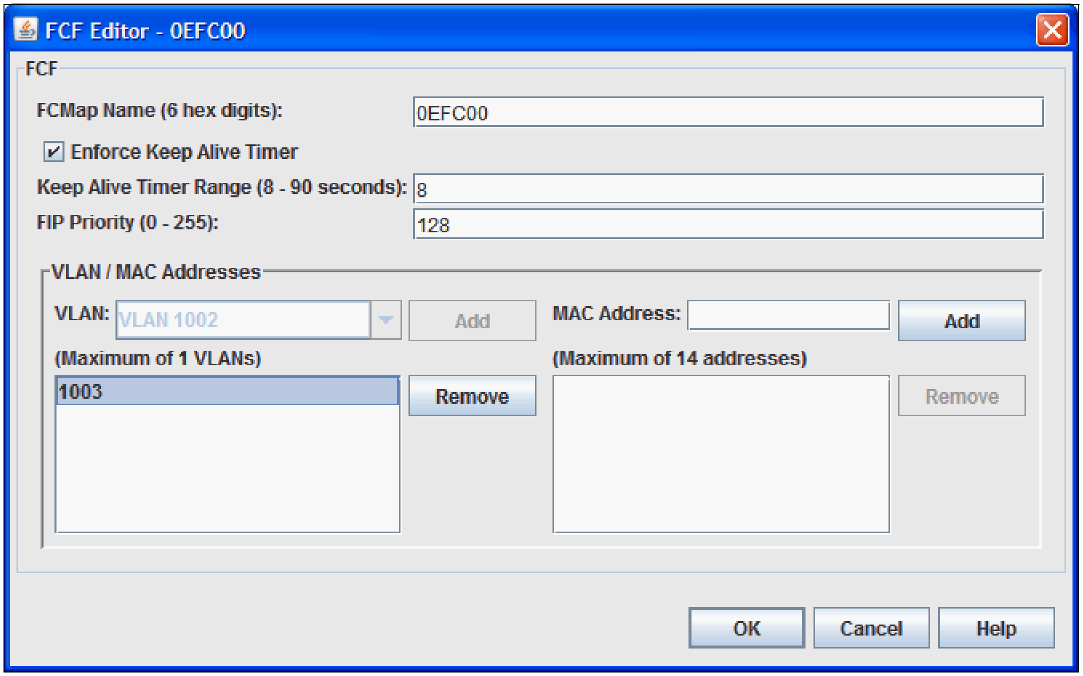

Figure 6-26 Virtual Fabric Extension Module FCF Editor

7. Back in the FCF Configuration Manager window, click **Apply**.

The change of the VLAN ID to 1003 is now completed.

Changing some FC default configurations, such as Domain ID or Zoning, are the only tasks that you might need to complete on the Virtual Fabric Extension Module. You can complete this configuration by using the GUI. However, this procedure goes beyond the scope of this publication. Another option without a Zoning or Domain ID configuration on the Virtual Fabric Extension Module is to run it in Transparent mode or to use NPV. For more information about this option, see 6.4.6, “Switching the Virtual Fabric Extension Module to N-Port Virtualization mode if connected to an existing FC fabric” on page 137.

Figure 6-27 shows an example of an active FCoE setup with Virtual Fabric Extension Module in FC Switch mode.

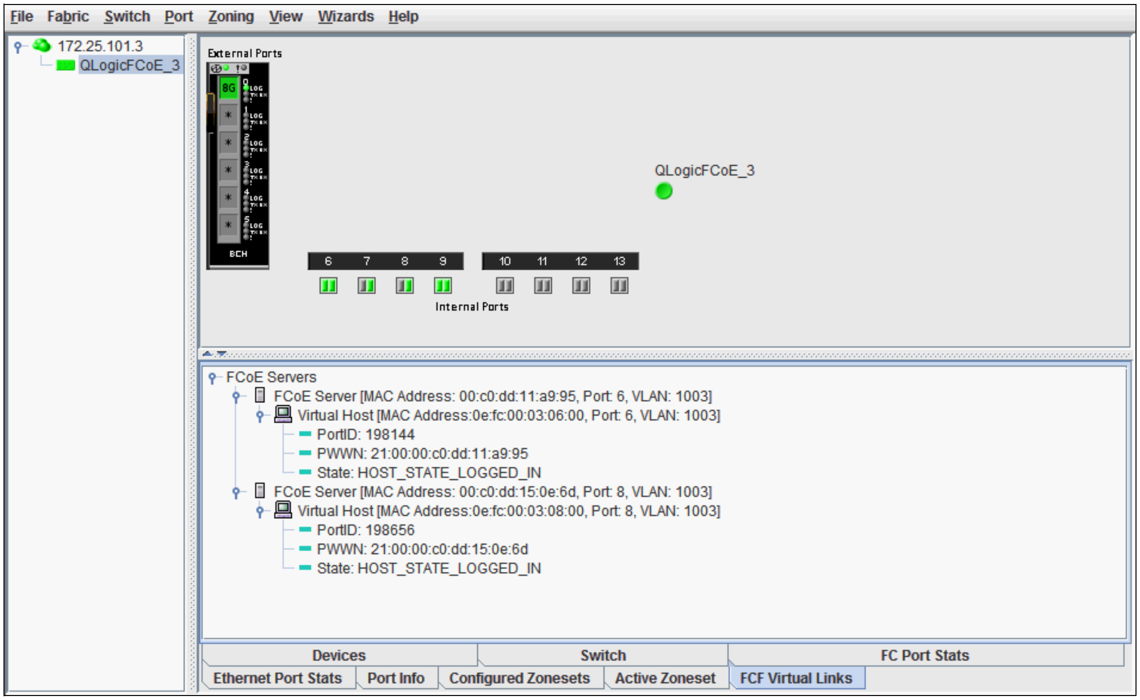

Figure 6-27 Virtual Fabric Extension Module in FC switch mode

Verify active FCoE session by using the **show fcoe fip fcf** and **show fcoe fip fcoe** commands as shown in Example 6-13.

Example 6-13 Verifying the active FCoE session

```
bay-9#sh fcoe fip fcf
     FCF MAC          Port     Vlan
-----------------------------------
00:c0:dd:13:9b:fc     BR3D     1003
00:c0:dd:13:9b:fb     BR3C     1003
00:c0:dd:13:9b:fa     BR3B     1003
00:c0:dd:13:9b:f9     BR3A     1003
bay-9#
bay-9#sh fcoe fips fcoe
```

```
   VN_PORT MAC           FCF MAC          Port    Vlan
------------------------------------------------------
0e:fc:00:05:0a:00   00:c0:dd:18:d7:15     INT11   1003
0e:fc:00:05:0b:00   00:c0:dd:18:d7:16     INT13   1003
bay-9#
```

6.4.6 Switching the Virtual Fabric Extension Module to N-Port Virtualization mode if connected to an existing FC fabric

If the storage is not directly connected to the Virtual Fabric Extension Module FC ports, use N_Port Virtualization to avoid interoperability issues and Domain ID consumption. Additionally, NPV facilitates the setup of the Virtual Fabric Extension Module, because no specific FC configuration needs to be done.

To switch the Virtual Fabric Extension Module into NPV mode (*Transparent mode* in QLogic):

1. In the main window of the Virtual Fabric Extension Module GUI (Figure 6-28), select **Switch** → **Advanced Switch Properties**.

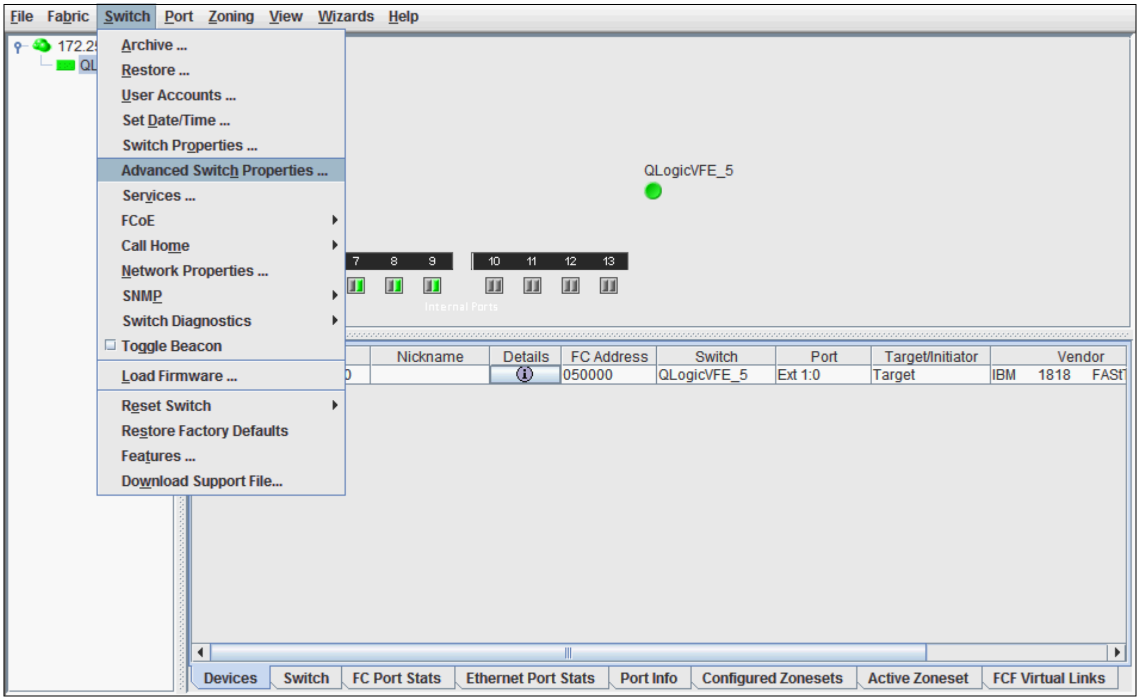

Figure 6-28 Virtual Fabric Extension Module GUI

2. In the Advanced Switch Properties window (Figure 6-29), select **Transparent Mode**, and click **OK**.

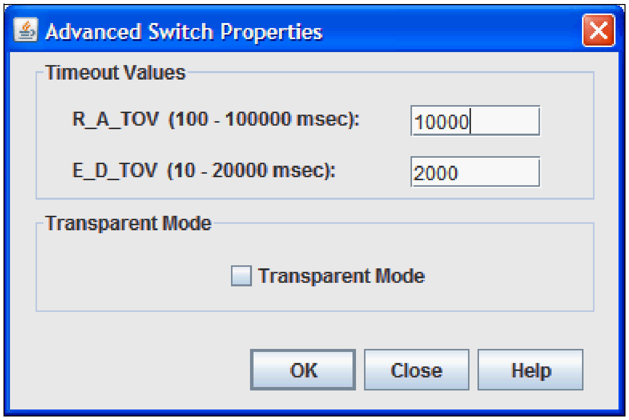

Figure 6-29 Virtual Fabric Extension Module Advanced Switch Properties window

3. Reboot the Virtual Fabric Extension Module to complete the switch to N-Port Virtualization.

Setting up Transparent mode by using the CLI

By default the QLogic Virtual Fabric Extension Module is configured for direct attach storage. To change the Virtual Fabric Extension Module to Transparent mode:

1. Use Telnet to connect to the Virtual Fabric Extension Modules IP address.
2. Enter the following commands in the order shown:

```
admin start
config edit
set config switch
```

3. Change Transparent Mode from False to True:

```
config save
config activate
```

Figure 6-30 shows an example of an active FCoE setup with a Virtual Fabric Extension Module in Transparent mode.

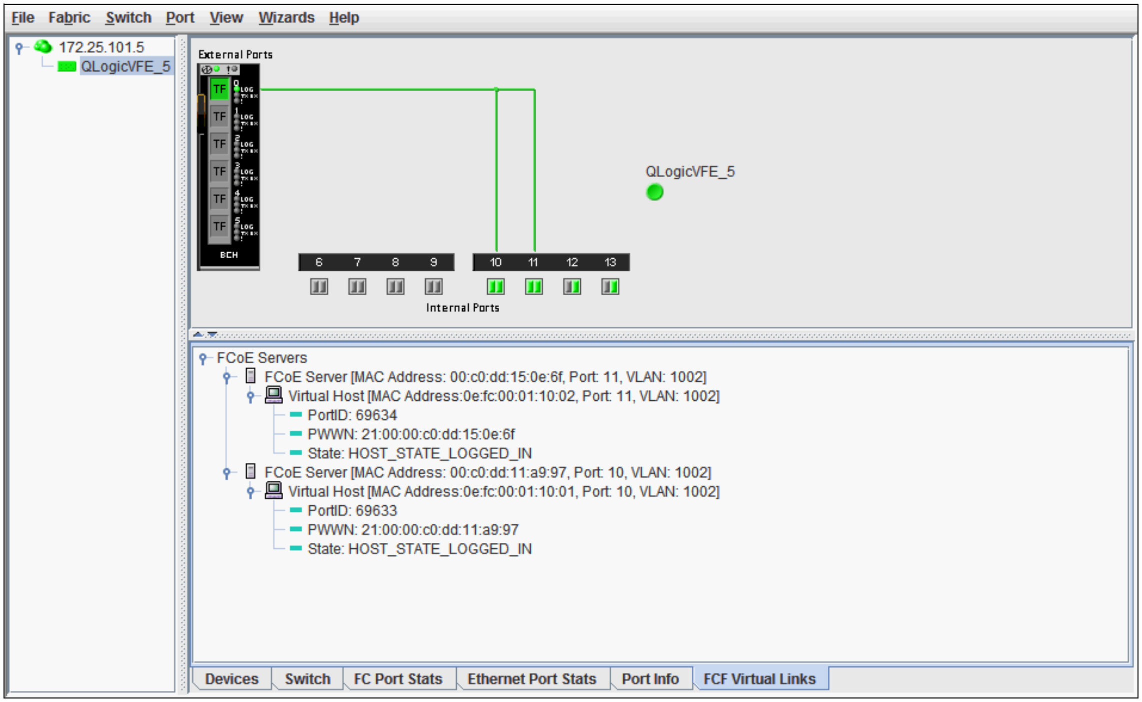

Figure 6-30 Virtual Fabric Extension Module in Transparent mode

6.4.7 Configuring the FCoE VLAN ID on the CNA

In general, you do not need to configure the VLAN ID that is used for FCoE traffic on the CNA. It is part of the FCoE Initiation Protocol to automatically discover the FCoE VLAN ID that is advertised by the FCF.

If the VLAN ID is not advertised by the FCF or is incompatible, you can manually configure the FCoE VLAN ID by using one of the following options. The example shown here is for the QLogic CNA (8100 series).

- Enter the VLAN ID in the HBA Parameters file by using the QLogic SANsurfer GUI:

   a. Start the SANsurfer GUI, and select the port of the FCoE VLAN ID to change.

   b. On the **Utilities** tab, under HBA Parameters, click **Save**. Enter a file name and location to dump the HBA Parameters into a text file on your system.

c. After you see the “`File saved successfully`” message, open the saved file in a text editor, and then search for the string “`Primary FCF VLAN ID`”.

d. Set `Primary FCF VLAN ID Match = 1`. In addition, set the FCoE VLAN ID, for example: `FCoE Primary FCF VLAN ID = 1003`. Keep the rest of the settings at their default values.

   Example 6-14 the modified parameters, which are highlighted in bold.

   Example 6-14 FCoE VLAN settings for QLogic CNA 81xx HBA parameters

```
[...]
;*******************************************
;Offset 0x080h

version                          [0-65535] = 1
Primary FCF VLAN ID Match            [0-1] = 1
Primary FCF Fabric Name Match        [0-1] = 0
reserved                            [0-63] = 0
reserved                           [0-255] = 0
reserved                         [0-65535] = 0
Primary FCF VLAN ID              [0-65535] = 1003
Primary FCF Fabric Name 0          [0-255] = 0
[...]
```

e. Return to SANsurfer, and ensure that the port is selected for the FCoE VLAN ID to change.

f. On the **Utilities** tab, under HBA Parameters, click **Update**.

g. When a message window opens, click **Yes**.

h. Open the edited file, and when prompted, enter the configuration password. (The default is `conifg`.)

The process is complete when you see the message “`HBA Parameters Update Complete`.” No reboot is needed.

- Press F1 upon starting the system to access the UEFI configuration menu (Figure 6-31). Set the Enable FCF VLAN ID to **<Enabled>**, and enter the VLAN ID.

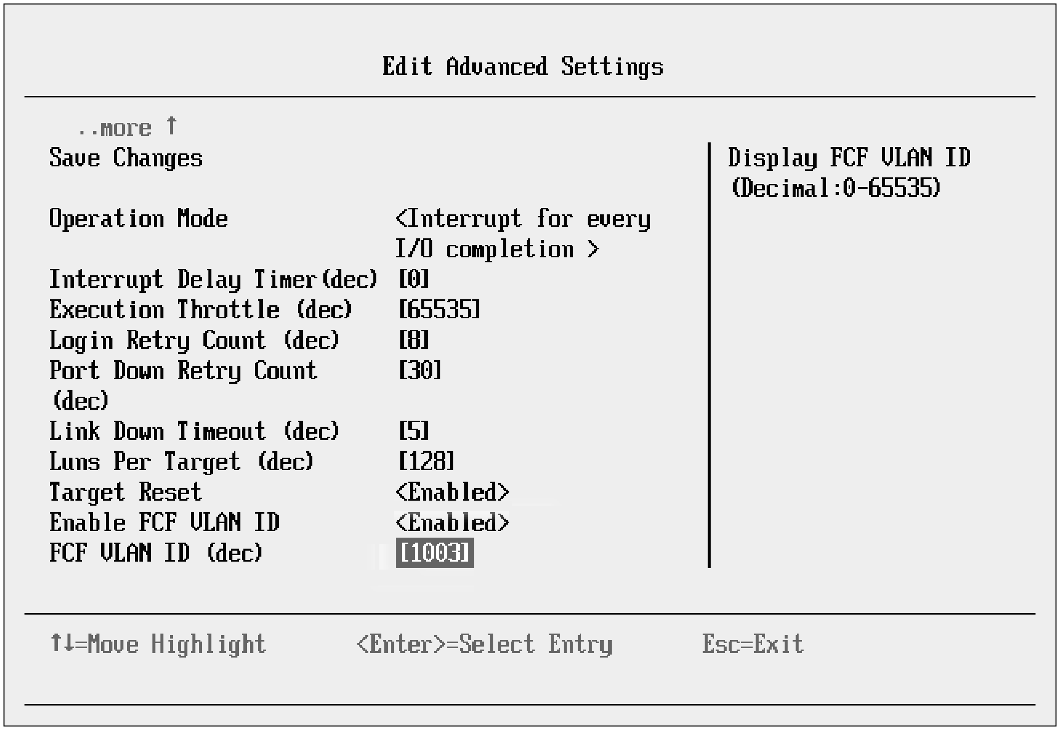

Figure 6-31 Manual FCoE VLAN ID setting for QLogic CNA in UEFI

> **Important:** Automatic FCoE VLAN discovery, based on the FCoE Initialization Protocol does not work after the FCoE VLAN ID is set manually.

6.4.8 Configuring FCoE for IBM Virtual Fabric VNIC

When the Emulex CNAs Virtual Fabric Adapter I or II are used in vNIC mode with FCoE personality, no additional configuration for FCoE is required on the Virtual Fabric Switch Module.

vNIC instances 1, 3, and 4 can be enabled for non-FCoE traffic, and at least one is necessary if the operating system needs to detect an active NIC. Example 6-15 on page 142 shows the configuration for vNIC instances 1 and 3 for the blade servers in slots 9 and 11 of the BladeCenter blade server.

Example 6-15 Configuration of BNT Virtual Fabric 10Gb Switch Module with vNIC

```
bay-7#sh run
Current configuration:
version "6.8.0.1"
switch-type "BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter"
!
! ### snip standard configuration and port INT1-INT8 ###
!
!
interface port INT9
pvid 99
no flowcontrol
exit
!
interface port INT10
pvid 99
no flowcontrol
exit
!
interface port INT11
pvid 99
no flowcontrol
exit
!
! ### snip port INT12-EXT5 ###
!
interface port EXT6
tagging
exit
!
interface port BR5A
tagging
tag-pvid
pvid 1002
exit
!
interface port BR5B
tagging
tag-pvid
pvid 1002
exit
!
interface port BR5C
tagging
```

```
tag-pvid
pvid 1002
exit
!
interface port BR5D
tagging
tag-pvid
pvid 1002
exit
!
vlan 1
member INT1-INT14,EXT1-EXT6,EXT11
no member BR5A-BR5D
vlan 99
enable
name "VLAN 99"
member INT1-INT14
!
vlan 1002
enable
name "VLAN 1002"
member INT1-INT14,BR5A-BR5D
!
vnic enable
vnic port INT9 index 1
bandwidth 25
enable
exit
!
vnic port INT9 index 3
bandwidth 25
enable
exit
!
vnic port INT11 index 1
bandwidth 25
enable
exit
!
vnic port INT11 index 3
bandwidth 25
enable
exit
!
vnic vnicgroup 1
```

```
vlan 101
enable
member INT9.1
member INT11.1
exit
!
vnic vnicgroup 3
vlan 103
enable
member INT9.3
member INT11.3
exit
!
spanning-tree stp 101 vlan 101
spanning-tree stp 103 vlan 103
spanning-tree stp 106 vlan 1002

interface port BR5A
no spanning-tree stp 106 enable
exit
!
interface port BR5B
no spanning-tree stp 106 enable
exit
!
interface port BR5C
no spanning-tree stp 106 enable
exit
!
interface port BR5D
no spanning-tree stp 106 enable
exit
!
fcoe fips enable
!
cee enable
!
lldp enable
!
end
bay-7#
```

vNIC instance 2 is fixed defined for FCoE traffic, and no vNIC instance or vNIC Group must be configured for FCoE. The configuration of FCoE between pNIC and vNIC mode is the same. For more information about vNIC and FCoE, see 10.1.2, “Switch and adapter vNIC correlation” on page 241.

6.4.9 Summary

Implementing FCoE for the BladeCenter, with IBM Virtual Fabric components Virtual Fabric 10G Switch and Virtual Fabric Extension Module, requires minimal configuration effort. The solution is flexible and scalable in bandwidth. It showed good performance, only a minor issue in regard to automatic VLAN discovery. We did not experience any compatibility issues with the external FC fabric or FC-attached storage.

No significant differences were detected by using Win2008R2, ESXi5.0, or Red Hat Enterprise Linux (RHEL) 6.1 with QLogic or Emulex CNAs.

In summary, it is easy to set up an FCoE solution within a BladeCenter chassis.

6.5 Scenario 5: vNIC iSCSI on Emulex II, BNT, and QLogic with IBM DS5300

This section includes the steps for implementing Virtual Fabric Mode for FCoE on the following adapters and modules:

- The Emulex Virtual Fabric Adapter II
- The BNT Virtual Fabric 10Gb Switch Module
- The QLogic Virtual Fabric Extension Module connecting to the IBM DS5300

Internet Small Computer System Interface (iSCSI), similar FCoE, is a block-oriented storage networking technology. The file system logic is in the client (initiator) operating system, not in the storage array. The location of this logic is a key difference between iSCSI and network-attached storage (NAS) protocols, such as Network File System (NFS) and Server Message Block (SMB) or Common Internet File System (CIFS).

Our testing included the hardware iSCSI initiator function provided by the Emulex Virtual Fabric Adapter I and II. The Emulex card offers the following support for iSCSI:

- In pNIC as one storage HBA and one NIC instance for each of the two ports
- In vNIC as one storage HBA and up to three NIC instances per port

We also tested iSCSI by using software initiators from Microsoft and VMware with the QLogic CNA, which does not provide a hardware initiator.

6.5.1 Key results

Our testing had the following key results:

- The Emulex card works successfully to constrain bandwidth on data and iSCSI vNIC instances. This result is different from FCoE vNICs, which are not configured with a bandwidth cap and do not appear to enforce one.
- You can use CEE configuration commands to assign a lossless priority group to iSCSI. However, you can use these commands only when using a hardware initiator that sets the appropriate priority at the origin. Our tests showed a behavior that is often shown about the ways in which CEE enhances iSCSI performance:
  - Without CEE, graphs of throughput showed the typical "sawtooth pattern," where bandwidth grows to a limit and then drops sharply.
  - With CEE, throughput reached at least the guaranteed bandwidth and was mostly steady at that value.
- In our iSCSI test in which we used the software initiators from Microsoft and VMware with the QLogic CNA, the VMware software initiator delivered excellent performance. For more information about storage and network convergence, and examples, see "Technology and brand comparison" in *Storage and Network Convergence Using FCoE and iSCSI*, SG24-7986. You can also use a software initiator within a guest VM under the hypervisor.

6.5.2 Configuration details for vNIC mode

Figure 6-32 shows the network topology and switch configurations used in testing the Emulex adapter in vNIC (virtual NIC) mode.

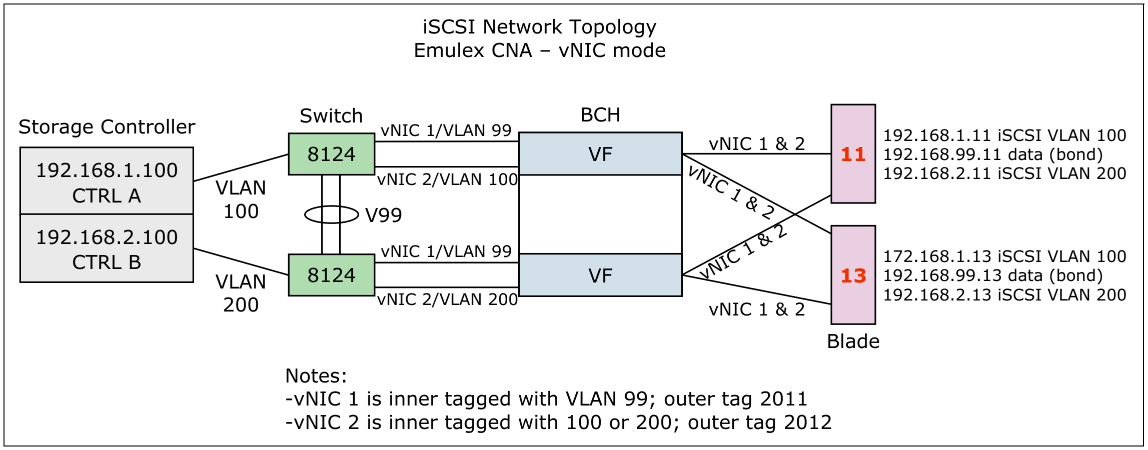

Figure 6-32 iSCSI network topology in vNIC mode

Example 6-16 shows the switch configuration for the Virtual Fabric switch in bay 7. An identical configuration was used for the switch in bay 9, except that iSCSI traffic in bay 7 used VLAN 100 and bay 9 used VLAN 200. Key parts of the configuration are highlighted in bold.

Example 6-16 Configuration of BNT switch module (bay 7 only)

```
version "6.7.4"
switch-type "BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter"



hostname "bay-7"


access user administrator-password
"f2c4b070e004a020bfadf3b323b403d2f0fc097036e20934f12feb2686ae0b65"


interface port INT1-INT14 (edited to remove repeated typing)
pvid 99
no flowcontrol
exit

interface port EXT1
shutdown
exit

interface port EXT2
shutdown
exit

interface port EXT3
shutdown
exit

interface port EXT5
tagging
exit

interface port EXT6
tagging
exit

!
```

```
interface port EXT7-EXT10
shutdown
exit

vlan 1
member INT1-INT4,INT7-INT14,EXT1-EXT4,EXT7-EXT11
no member INT5-INT6,EXT5-EXT6

vlan 99
enable
name "VLAN 99"
member INT1-INT5,INT7-INT14

vlan 100
enable
name "VLAN 100"
member INT1-INT4,INT6-INT14

vlan 2011
enable
name "VLAN 2011"

vlan 2012
enable
name "VLAN 2012"

vlan 4095
member INT1-INT4,INT7-MGT2
no member INT5-INT6


vnic enable
vnic port INT11 index 1
bandwidth 50
enable
exit

vnic port INT11 index 2
bandwidth 50
enable
exit

vnic port INT13 index 1
bandwidth 50
enable
```

```
exit

vnic port INT13 index 2
bandwidth 50
enable
exit

vnic vnicgroup 1
vlan 2011
enable
member INT11.1
member INT13.1
port EXT5
exit

vnic vnicgroup 2
vlan 2012
enable
member INT11.2
member INT13.2
port EXT6
exit


spanning-tree stp 91 vlan 2011

interface port EXT5
no spanning-tree stp 91 enable
exit

spanning-tree stp 92 vlan 2012

interface port EXT6
no spanning-tree stp 92 enable
exit

spanning-tree stp 99 vlan 99

spanning-tree stp 100 vlan 100

spanning-tree stp 107 vlan 1003


snmp-server name "bay-7"
```

```
cee enable



access-control list 1 tcp-udp destination-port 3260 0xffff
access-control list 1 action set-priority 3

access-control list 2 tcp-udp source-port 3260 0xffff
access-control list 2 action set-priority 3

interface port INT1-INT14 (edited to remove redundant display)
access-control list 1
access-control list 2

interface port EXT1-EXT6 (edited to remove redundant display)
access-control list 1
access-control list 2


lldp enable

interface ip 1
ip address 192.168.1.254 255.255.255.0
vlan 100
enable
exit

interface ip 99
ip address 192.168.99.253
vlan 99
enable
exit



end
```

The Emulex Virtual Fabric Adapter I and II support vNICs, which are not the same as vNICs that are used in hypervisor environments. Hypervisor environments use virtualized hardware NICs or paravirtualized software NICs, which are detected by the software-switching logic that is included in most hypervisors. Emulex vNICs are discovered by the operating system on the server (hypervisors and conventional OSs) by using Peripheral Component Interconnect (PCI). Each vNIC corresponds one-to-one with a PCI function code associated with the Emulex card.

vNICs are configured on an IBM Virtual Fabric switch and communicated to the Emulex card by using the DCBX Protocol (Virtual Fabric Mode). Alternatively, vNICs are configured in the UEFI of a system in Switch Independent Mode, which works with any upstream switch. We used vNIC mode in our testing. In either case, the configuration allows the specification of the vNIC instances that are active and the amount of bandwidth that is allocated to each one.

iSCSI vNIC instances are configured in the same way as data instances. Storage vNIC instances (iSCSI and FCoE) are always assigned vNIC index 2 by convention when configuring them by using the switch. When the Emulex adapter is set to a personality that includes storage functions, as part of the DCBX communication with the switch, the adapter indicates the presence of the storage function (either iSCSI or FCoE). The adapter then assigns it to vNIC instance 2 on each port.

Treatment of VLANs in vNIC mode is unusual in that the VLANs that are carried on the vNICs are not configured on the switch. For example, in our testing, VLAN 99 was used to carry non-iSCSI traffic, and VLAN 100 was used to carry iSCSI traffic. In vNIC mode, both classes of traffic were carried through vNIC groups 1 and 2. A vNIC group uses *802.1 q-in-q double tagging* with an outer tag VLAN (2011 and 2012 in this test) on all traffic that it carries through the switch. The VLANs configured on the operating system of the server and on upstream switches are inner-tag VLANs, which are invisible to the switch. Double tagging is applied on ingress to the switch and removed on egress.

The use of vNIC mode, therefore, had the following consequences in our testing:

- Configuration of VLANs 99 and 100 on the bay-7 switch had no effect. The same result was true for VLANs 99 and 200 on the bay-9 switch (configuration not shown). It does not matter that the external ports used in this text (EXT5 and EXT6) are not members of those VLANs because they are members of vNIC groups 1 and 2.
- The ACL shown in bold in the configuration was ineffective partly because the priority field is part of the same field in the header where tagging is applied. Bandwidth management in vNIC mode is achieved by assigning bandwidth to the individual vNIC instances.

Part 3

Switch Independent Mode

This part highlights the product details and implementation for the Switch Independent Mode offering.

This part includes the following chapters:

- Chapter 7, “Solution architecture” on page 155
- Chapter 8, “Supported products” on page 165
- Chapter 9, “Configuring the components” on page 195
- Chapter 10, “Usage scenarios” on page 237

7

Solution architecture

In Switch Independent Mode, the virtual channel separation on the shared 10 Gbps link is maintained in both the Virtual Fabric Adapter and the switch. The adapter network interface controllers (NICs) enforce bandwidth limits for each virtual channel. However, in Switch Independent Mode, the switches do not enforce bandwidth limits to the servers. Instead, they rely upon link-level flow control to throttle traffic when necessary.

This chapter includes the following sections:

- Mode comparison
- Solution overview
- Features

7.1 Mode comparison

This section highlights the differences between IBM Virtual Fabric Mode and Switch Independent Mode.

7.1.1 IBM Virtual Fabric Mode

IBM Virtual Fabric Mode is the same as virtual NIC (vNIC) mode on the original Emulex 10GbE Virtual Fabric Adapter. For IBM BladeCenter, this mode works only with the BNT Virtual Fabric 10Gb Switch Module. For IBM System x, this mode is supported only with IBM RackSwitch G8124 and G8264.

In IBM Virtual Fabric Mode, the Emulex Virtual Fabric Adapter communicates with the switch module to obtain vNIC parameters by using the Data Center Bridging Exchange (DCBX). Also, a special tag within each data packet is added and later removed by the NIC and switch for each vNIC group to maintain separation of the virtual channels.

In vNIC mode, each physical port is divided into four virtual ports, for a total of eight vNICs per adapter. The default bandwidth for each vNIC is 2.5 Gbps. Bandwidth for each vNIC can be configured at the BNT switch operating system (OS) from 1 Gbps to 10 Gbps, and up to a total of 10 Gbps per physical port. The vNICs can also be configured to have 0 bandwidth if you must allocate the available bandwidth to fewer than eight vNICs.

In IBM Virtual Fabric Mode, you can change the bandwidth allocations through the BNT switch OS user interfaces without rebooting the server.

When storage protocols are enabled on the Emulex 10GbE Virtual Fabric Adapter II Advanced for IBM BladeCenter, six ports are Ethernet, and two ports can be Internet Small Computer System Interface (iSCSI) or Fibre Channel over Ethernet (FCoE).

7.1.2 Switch Independent Mode

In Switch Independent Mode, the Emulex Virtual Fabric Adapter works with the following switches:

- BNT Virtual Fabric 10Gb Switch Module (for BladeCenter)
- IBM RackSwitch G8124 (for System x)
- IBM RackSwitch G8264 (for System x)
- Cisco Nexus 4001I Switch Module
- 10Gb Ethernet Pass-Thru Module and a top-of-rack switch

In IBM BladeCenter servers, these switches must be installed in I/O module bays 7 and 9 in the chassis. In addition to how the adapter communicates with the switch, Switch Independent Mode offers similar capabilities as IBM Virtual Fabric Mode. For example, they are similar in terms of the number of vNICs and the bandwidth that each can have.

Switch Independent Mode extends existing customer VLANs to the vNIC interfaces. The IEEE 802.1Q VLAN tag is essential to the separation of the vNIC groups by the NIC adapter or driver and the switch. The VLAN tags are added to the packet by the applications or drivers at each end station rather than by the switch.

The full features of the BNT and Cisco switches are still available when operating in Switch Independent Mode. In this mode, an outer VLAN tag is not required to maintain channel separation. Instead, the mode uses the media access control (MAC) address alone or with the IEEE 802.1Q VLAN tag to distinguish the individual vNIC channels on each link. Therefore, you and your network administrator must coordinate the configuration of the Virtual Fabric Adapters and the chassis switches to match the existing VLAN structure in your network.

Figure 7-1 shows an example implementation of Switch Independent Mode and the VLAN tagging that is employed.

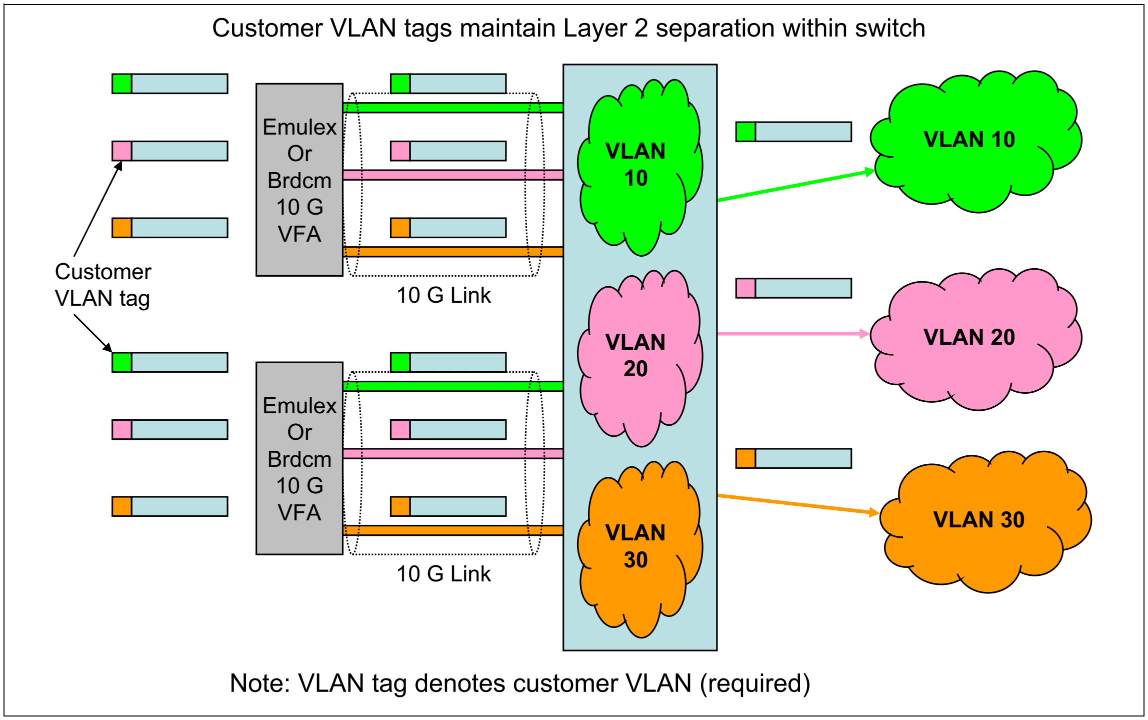

Figure 7-1 Overview of Switch Independent Mode

By using Emulex 10GbE Virtual Fabric Adapter II, your network administrator can configure one VLAN ID per virtual channel interface. Up to 31 additional VLANs can be configured within the applications stack (such as VMware) for each virtual channel. These VLANs are then incorporated with the MAC address or addresses on inbound frames to forward traffic to the correct interface.

The Broadcom 2-port 10Gb Virtual Fabric Adapter includes one MAC address per virtual channel. You can assign up to 16 MAC addresses to each vNIC partition, for a total of up to 64 MAC addresses per each physical port. You can also allocate separate bandwidth allocations to each vNIC partition. The vNIC partition bandwidth allocation allows for fine-tuning and division of available physical port bandwidth between the vNICs.

7.2 Solution overview

Table 7-1 shows the possible configurations that support Switch Independent Mode.

Table 7-1 Supported combinations by using Switch Independent Mode

	Supported switches		
Supported adapters	**BNT Virtual Fabric 10Gb Switch Module**	**Cisco Nexus 4001I Switch Module**	**Brocade Converged 10GbE Switch Module**
Emulex 10GbE Virtual Fabric Adapter, 49Y4235	Not supported	Not supported	Not supported
Emulex 10GbE Virtual Fabric Adapter Advanced, 49Y4275	Not supported	Not supported	Not supported
Emulex 10GbE Virtual Fabric Adapter II, 90Y3550	Supported	Supported	Supported
Emulex 10GbE Virtual Fabric Adapter II Advanced, 90Y3566	Supported	Supported	Supported
Broadcom 2-port 10Gb Virtual Fabric Adapter, 81Y3133	Supported	Supported	Not supported

Figure 7-2 on page 159 provides an overview of the vNIC solution architecture. The colors represent the virtual port groups to isolate traffic flows for services and applications. The thickness of each link represents the amount of assigned bandwidth for a particular traffic flow. Switch Independent Mode provides the capability to subdivide a single 10 Gbps link into up to four virtual channels.

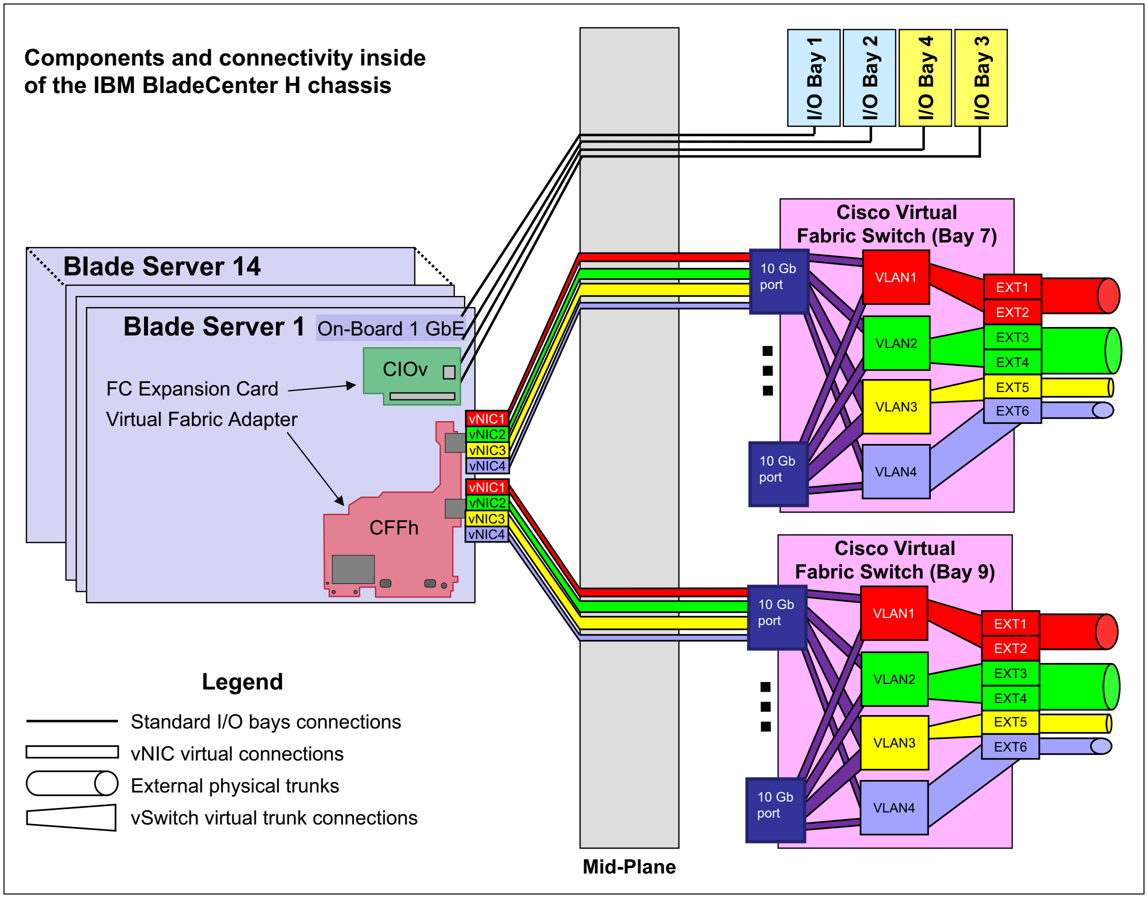

Figure 7-2 BladeCenter Virtual Fabric vNIC solution diagram with Cisco switch

> **Tip:** An uplink can support multiple VLANs (not shown in Figure 7-2).

vNICs are configured in the operating system by using the Emulex or Broadcom software in the preboot environment. The Virtual Fabric Adapter configuration and setup are performed inside the network controller rather than in the switch.

Each physical interface of a Virtual Fabric Adapter can be divided into up to four vNICs. Create a corresponding VLAN within the chassis switch for each vNIC associated with a switch port. For example, a Virtual Fabric Adapter port with four vNICs defined must correspond to four VLANs on the switch port. The switch VLANs must correspond to the VLANs that are defined throughout your network infrastructure. The vNIC interfaces become extensions of these network-wide VLANs.

You must take the following considerations into account when planning to use the components in Switch Independent Mode:

- The Virtual Fabric Adapter must operate in Switch Independent Mode (Emulex) or in NIC Partition mode (Broadcom).
- Up to four vNICs can be configured on each physical NIC (pNIC) on the Virtual Fabric Adapter (a total of eight vNICs on a two-port Virtual Fabric Adapter).
- The bandwidth for each vNIC can be assigned in a range of 100 Mb to 10 Gb in 100 Mbps increments. The default is 2.5 Gb.
- The total bandwidth of the four vNICs per port cannot exceed 10 Gbps. However, the Broadcom 2-port 10Gb Virtual Fabric Adapter supports *oversubscription*. That is, the aggregate bandwidth can be as high as 4 times 10 Gb. This amount of bandwidth allows any of the vNICs to use as much as 10 Gb where no contention exists or a lesser amount when contention exists (based on the weight assigned to that vNIC).
- The sum of the bandwidth allocated for all four vNICs on the same pNIC cannot exceed 10 Gb.
- Support by BladeCenter Open Fabric Manager is planned but was not available at the time this book was written.

> **Tip:** All parameters for Switch Independent Mode and NIC Partition are configured on the Virtual Fabric adapter.

For more information about supported configurations, see the BladeCenter Interoperability Guide (BIG) at:

http://www.ibm.com/support/entry/portal/docdisplay?brand=5000020&lndocid=MIGR-5073016

7.3 Features

The Emulex 10GbE Virtual Fabric Adapter II and Broadcom 2-port 10Gb Virtual Fabric Adapter are both high-performance, dual-port network adapters for 10 Gbps Ethernet (GbE) networks. They have the following general characteristics:

- Each server pNIC port is divided into up to four vNICs.
- The operating system configurations see eight unique NICs (2 ports x 4 vNICs). Figure 7-3 on page 161 shows how a configured Broadcom 2-port 10Gb Virtual Fabric Adapter is displayed to Windows Server 2008.

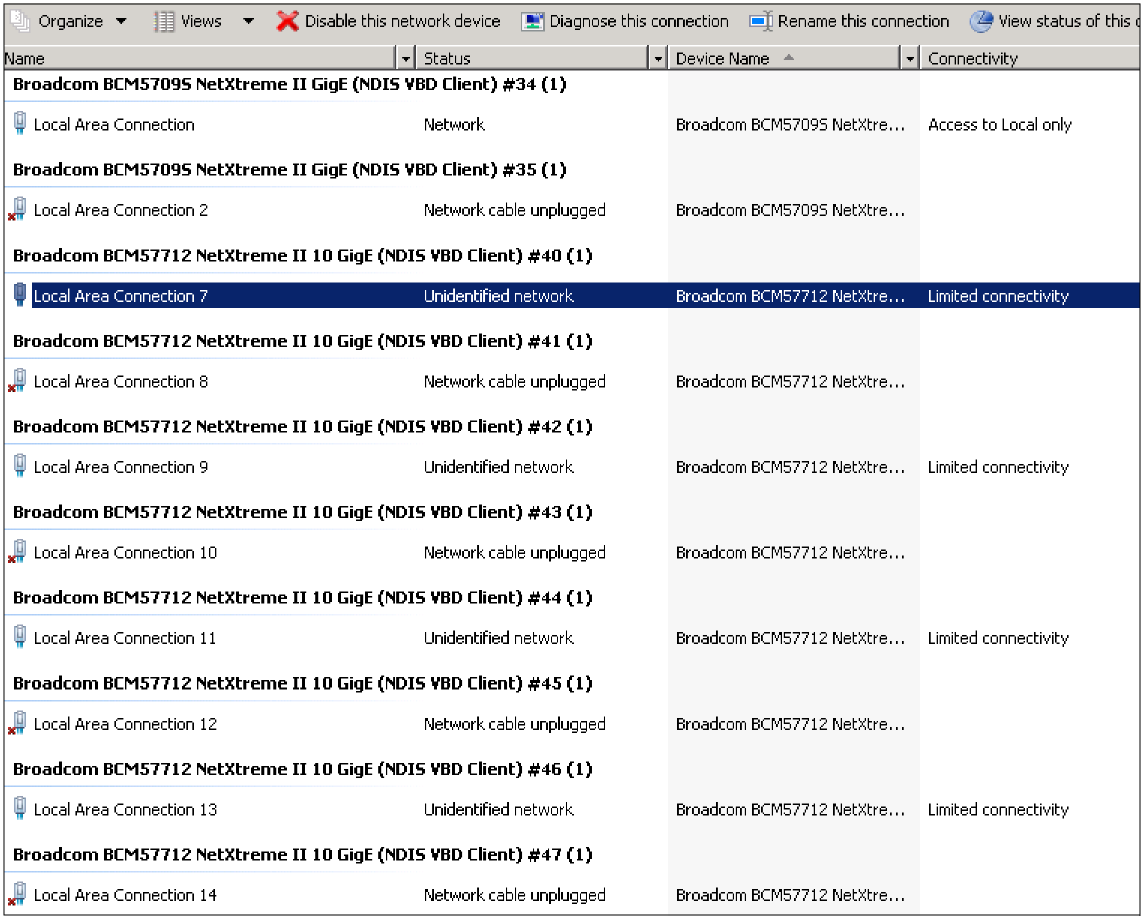

Figure 7-3 NICs in network connections for Windows Server 2008

- All vNIC parameters are configured from the Virtual Fabric Adapter user interface.
- The user enables vNICs and allocates bandwidth:
  - The allowable vNIC bandwidth range is from 100 Mbps to 10 Gbps:
    - Bandwidth increments are in 100 Mbps.
    - The default bandwidth is 2.5 Gbps.
    - For the Emulex adapter, the sum of all four vNIC bandwidths cannot exceed 10 Gbps.
    - For the Broadcom adapter, the sum of all four allocated vNIC bandwidths can exceed 10 Gbps because the adapter supports

oversubscription. The total actual used vNIC bandwidths cannot exceed 10 Gbps.

– Server-to-Virtual Fabric Switch bandwidth metering is on a per-vNIC basis.

Figure 7-4 shows an example of a possible configuration.

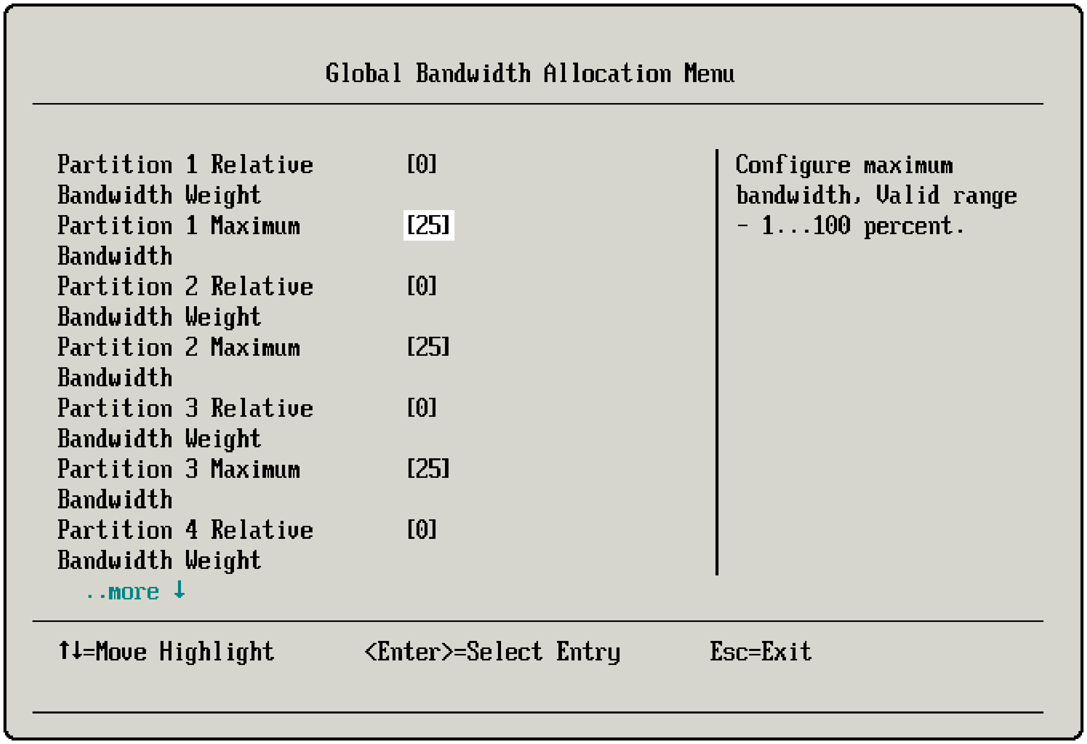

Figure 7-4 Setting bandwidth

- The user associates vNICs to VLAN groups within the switch by configuring the switch independent from the Virtual Fabric Adapter. Optionally, switch uplinks can be assigned to VLAN groups:
  – VLAN groups within the switch serve to isolate virtual NIC traffic that is flowing on the same physical port:
    - Existing VLANs within the client network are not impacted.
    - No forwarding occurs between uplinks assigned to VLAN groups.
    - An uplink (port or trunk) can belong to multiple vNIC groups.
    - A server port (pNIC or vNIC) can be associated with one or more VLAN groups.

- The failover mechanism for Switch Independent Mode is not virtual port-aware.

Concurrency must comply with the following key rules:

- vNIC and VMready

  VMready is a unique offering that enables the switches in the network to be virtual machine-aware. The network switches can be configured and managed for virtual ports (vports), rather than just for physical ports.

  By using Switch Independent Mode, you can enable vNIC mode on the Emulex 10GbE Virtual Fabric Adapter II and implement VMready in the switch for that server at the same time. For more information, see Chapter 8, “Supported products” on page 165.

- vNIC and switch stacking

  A *switch stack* is a group of BNT Virtual Fabric 10Gb Switch Modules that work as a unified system and can be managed as a single entity.

  The network views the stack as a single switch entity. The stack can be accessed and managed as a whole by using standard switch IP interfaces. After the stacking links are established, the number of ports available in a stack equals the total number of remaining ports of all the switches that are part of the stack. The number of available IP interfaces, VLANs, trunks, trunk links, and other switch attributes are not aggregated among the switches in a stack. The totals for the stack as a whole are the same as for any single switch configured in stand-alone mode.

  The BNT Virtual Fabric Switch supports stacking, with either of the vNIC modes at the same time.

- vNIC and iSCSI or vNIC and FCoE

  When storage protocols are enabled on the Emulex 10GbE Virtual Fabric Adapter Advanced II, six ports are Ethernet, and two ports are iSCSI or FCoE.

8

Supported products

This chapter describes the products that are used in Switch Independent Mode. This chapter includes the following sections:

- Overview of Switch Independent Mode
- Virtual Fabric adapters
- Virtual Fabric Switch Modules
- 10Gb Ethernet Pass-Thru Module for IBM BladeCenter

8.1 Overview of Switch Independent Mode

Switch Independent Mode consists of the supported combinations of the following adapters, switch modules, and chassis types:

- Virtual Fabric adapters
  - Emulex 10GbE Virtual Fabric Adapter II
  - Emulex 10GbE Virtual Fabric Adapter II Advanced
  - Broadcom 2-port 10Gb Virtual Fabric Adapter

> **Important:**
>
> - Support for Switch Independent Mode by the Emulex 10GbE Virtual Fabric Adapter II requires a firmware upgrade. You must have the latest firmware installed.
> - Emulex 10GbE Virtual Fabric Adapter, 49Y4235, and Emulex 10GbE Virtual Fabric Adapter Advanced, 49Y4275, are not supported in the Switch Independent Mode offering. Use the Emulex 10GbE Virtual Fabric Adapter II or II Advanced, 90Y3550 or 90Y3566.

- 10 Gbps Ethernet switch modules
  - BNT Virtual Fabric 10Gb Switch Module (for BladeCenter)
  - 10Gb Ethernet Pass-Thru Module and a top-of-rack switch
  - Cisco Nexus 4001I Switch Module
- BladeCenter H or HT chassis

Table 8-1 lists the combinations that are supported.

Table 8-1 Supported combinations by using Switch Independent Mode

	Supported switches			
Supported adapters	**BNT Virtual Fabric 10Gb Switch Module**	**Cisco Nexus 4001I Switch Module**	**10Gb Ethernet Pass-Thru Module**	**Brocade Converged 10GbE Switch Module**
Emulex 10GbE Virtual Fabric Adapter, 49Y4235	Not supported	Not supported	Not supported	Not supported
Emulex 10GbE Virtual Fabric Adapter Advanced, 49Y4275	Not supported	Not supported	Not supported	Not supported
Emulex 10GbE Virtual Fabric Adapter II, 90Y3550	Supported	Supported	Supported	Supported

Supported adapters	Supported switches			
	BNT Virtual Fabric 10Gb Switch Module	Cisco Nexus 4001I Switch Module	10Gb Ethernet Pass-Thru Module	Brocade Converged 10GbE Switch Module
Emulex 10GbE Virtual Fabric Adapter II Advanced, 90Y3566	Supported	Supported	Supported	Supported
Broadcom 2-port 10Gb Virtual Fabric Adapter, 81Y3133	Supported	Supported	Not supported	Supported

8.2 Virtual Fabric adapters

The Switch Independent Mode offering supports the use of the Emulex 10GbE Virtual Fabric Adapter II and the Broadcom 2-port 10Gb Virtual Fabric Adapter. This section describes these two adapters.

8.2.1 Emulex 10GbE Virtual Fabric Adapter II

The Emulex 10GbE Virtual Fabric Adapter II, part number 90Y3550, is a dual-port, 10 Gbps Ethernet card that supports 1 Gbps or 10 Gbps traffic, or up to eight virtual network interface controller (vNIC) devices. The vNICs are configured to meet your mix of network connectivity and throughput demands of today's complex server application environments. Each physical 10 Gbps port can be divided into four virtual ports, with bandwidth allocation in 100 Mbps increments, to a maximum of 10 Gbps per physical port. The adapter is a CFFh expansion card. Therefore, it requires installation of switch modules in I/O bays 7 and 9.

Emulex Virtual Fabric Adapter II Advanced, part number 90Y3566, is a separate adapter based on the same hardware. In addition to the features on a standard card, it has support for Fibre Channel over Ethernet (FCoE) and Internet Small Computer System Interface (iSCSI) hardware initiator functions.

The Emulex Virtual Fabric Adapter Advanced Upgrade, part number 49Y4265, is also available. By using this offering, you can upgrade the Emulex 10GbE Virtual Fabric Adapter II to support the FCoE and iSCSI capabilities of the Advanced adapter.

Figure 8-1 shows the Emulex 10GbE Virtual Fabric Adapter II.

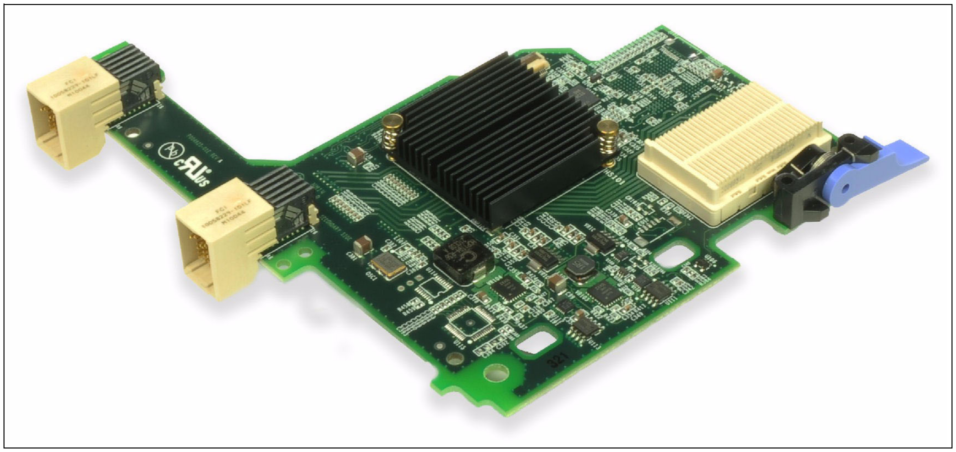

Figure 8-1 Emulex 10GbE Virtual Fabric Adapter II (CFFh)

Table 8-2 lists the part numbers and feature codes of the Emulex Virtual Fabric Adapter (CFFh).

Table 8-2 Part numbers and feature codes

Description	Part number	Feature code
Emulex Virtual Fabric Adapter II	90Y3550	A1XG
Emulex Virtual Fabric Adapter II Advanced	90Y3566	AIXH
Emulex Virtual Fabric Adapter Advanced Upgrade	49Y4265	2436

The Emulex 10GbE Virtual Fabric Adapter II has the following features:

- Connection to 1 Gbps or 10 Gbps data center infrastructure (1 Gbps and 10 Gbps auto-negotiation)
- Peripheral Component Interconnect Express (PCIe) 2.0 x8 host interface
- IBM BladeCenter CFFh form factor, which can be combined with a CIOv expansion card on the same blade server
- Operation as an 8-port virtual NIC or a 2-port, 1/10 Gbps Ethernet adapter:
  - Virtual port bandwidth allocation in 100 Mbps increments
  - Support for up to eight virtual ports
- Wake on LAN support
- Full-duplex (FDX) capability

- Bus-mastering support
- Direct memory access (DMA) support
- Preboot Execution Environment (PXE) support
- IPv4 or IPv6 TCP, User Datagram Protocol (UDP) checksum offload
  - Large send offload (LSO)
  - Large receive offload (LRO)
  - Receive Side Scaling (RSS)
  - IPv4 TCP Chimney Offload
- VLAN insertion and extraction
- Jumbo frames up to 9000 bytes
- Load balancing and failover support
- This support includes adapter fault tolerance (AFT), switch fault tolerance (SFT), adaptive load balancing (ALB), teaming support, and IEEE 802.3ad.
- Enhanced Ethernet (draft)
  - Enhanced Transmission Selection (ETS; P802.1Qaz)
  - Priority-based flow control (PFC; P802.1Qbb)
  - Data Center Bridging Capabilities Exchange (DCBX) Protocol, Cisco Intel Nuova (CIN)-DCBX, and Converged Enhanced Ethernet (CEE) DCBX (P802.1Qaz)
- FCoE and iSCSI hardware initiator support

  You can configure the adapter to operate as two physical iSCSI ports.
- Serial over LAN function support
- Support for both virtual fabric modes: IBM Virtual Fabric Mode and Switch Independent Mode

The original Emulex 10GbE Virtual Fabric Adapter can be configured in physical NIC (pNIC) or vNIC mode. The new Emulex 10GbE Virtual Fabric Adapter II now offers two vNIC modes of operation:

- IBM Virtual Fabric Mode (also called *vNIC1*), which is the same as vNIC mode on original Emulex Virtual Fabric Adapters

  This mode is used when configuring the IBM Virtual Fabric Mode offering. This mode is also called *Switch Dependent Mode*, because the Emulex adapters rely on the switch module to obtain vNIC parameters (by using DCBX). This mode works only with BNT Virtual Fabric 10Gb Switch Module.

- Switch Independent Mode (also called *vNIC2*), which does not rely on any vNIC settings that are passed to the adapter from the switch module

  In this mode, the Emulex 10GbE Virtual Fabric Adapter II can work with the following switches:

  – BNT Virtual Fabric 10Gb Switch Module
  – Cisco Nexus 4001I Switch Module

Figure 8-2 shows placement of Emulex 10GbE Virtual Fabric Adapter II in an HS22 blade server. Because it is a high-speed expansion card, you must install it in the CFFh expansion slot.

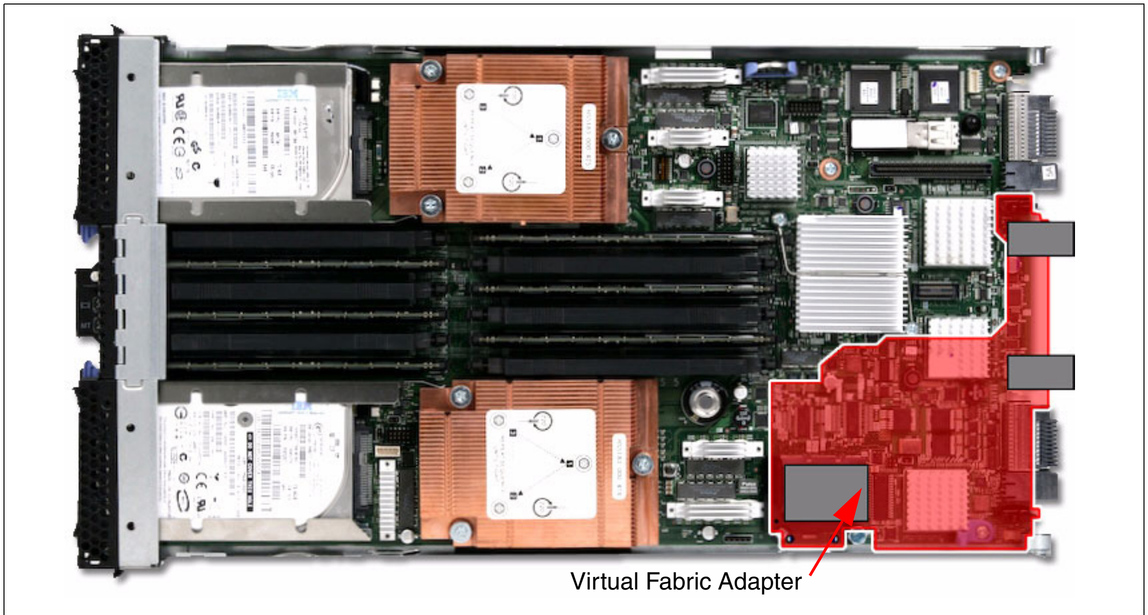

Figure 8-2 Placement of a CFFh card in the BladeCenter HS22 server

8.2.2 Broadcom 2-port 10Gb Virtual Fabric Adapter

The Broadcom 2-port 10Gb Virtual Fabric Adapter (Figure 8-3), part number 81Y3133, is a dual-port 10Gbps high-speed expansion card that operates in Switch Independent Mode.

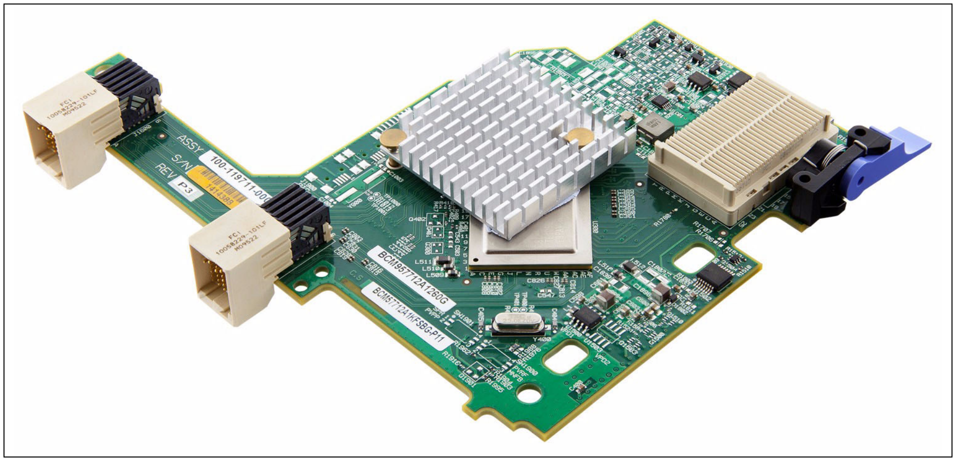

Figure 8-3 Broadcom 2-port 10Gb Virtual Fabric Adapter

You can configure up to four partitions (vNICs) on each of the two 10-Gbps physical ports. vNIC partitioning is based on media access control (MAC) addresses. In contrast with the Emulex 10GbE Virtual Fabric Adapter II, the Broadcom 2-port 10Gb Virtual Fabric Adapter does not use VLAN IDs for vNIC partitioning.

You can assign up to 16 MAC addresses to each vNIC partition, for a total of up to 64 MAC addresses per physical port. You can dynamically allocate bandwidth to each vNIC partition. vNIC partition bandwidth allocation allows for fine-tuning and division of available physical port bandwidth between the vNICs.

Each enabled partition is displayed as a dedicated network card to the operating system.

The Broadcom 2-port 10Gb Virtual Fabric Adapter has the following features:

- Integrated dual 10 Gbps MAC and dual 10GBASE-CX4
- PCIe x8 2.0, 5 GT/s compliant
- IEEE 802.3ap Clause 73 compliant backplane operation
- IEEE 802.3xx Clause 37 compliant auto-negotiation for 1Gbps
- TCP/IP Offload Engine (TOE)
- Microsoft TCP chimney compliant
- Full FASTPATH TCP offload for IPv4 and IPv6
- Network Controller Sideband Interface (NCSI)

- PXE v2.1 remote boot
- Wake on LAN
- Virtual LANs, IEEE 802.1q VLAN tagging
- IEEE 802.3x flow control

The Broadcom 2-port 10Gb Virtual Fabric Adapter form factor is CFFh. Therefore, it communicates with high-speed I/O modules installed in I/O bays 7 and 9. See Figure 8-2 on page 170 for placement of the adapter in an HS22 blade.

8.3 Virtual Fabric Switch Modules

Because the Switch Independent Mode is based on the vNIC configuration capabilities of the Virtual Fabric adapters, it does not depend on the switch module model that is used. Currently, the following switches are supported:

- BNT Virtual Fabric 10Gb Switch Module
- Cisco Nexus 4001I Switch Module

8.3.1 BNT Virtual Fabric 10Gb Switch Module

The BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter offers the most bandwidth of any blade switch. It represents the perfect migration platform for clients who are still at 1 Gb outside the chassis by seamlessly integrating into the existing 1 Gbps infrastructure. This switch is the first 10 Gbps switch for IBM BladeCenter that supports converged networking. That is, it can transmit CEE to an FCoE-capable, top-of-rack switch. This feature is available with firmware release 6.1 and later.

By using the CEE and FCoE functions, clients can transfer storage, network, Voice over Internet Protocol (VoIP), video, and other data over the common Ethernet infrastructure. Also, by using the QLogic Virtual Fabric Extension Module, clients can achieve FCoE gateway functions inside the BladeCenter chassis.

This book focuses on virtual fabric. The BNT switch module can be used in IBM Virtual Fabric Mode and Switch Independent Mode offerings. In an IBM Virtual Fabric Mode solution, all vNIC parameter configuration is performed on the BNT Virtual Fabric 10Gb Switch Module. Therefore, the switch module is a key part of the offering. By using this offering, clients can form eight vNICs from one pNIC and manage them in virtual groups. The switch module can be managed by using a command-line interface (CLI) or a web browser interface. In the future, with BladeCenter Open Fabric Manager, the switch module will provide all the benefits of I/O virtualization at 10 Gbps speeds.

For a chassis with multiple servers, with several that operate at 1 Gbps or at 10 Gbps, and several that transmit converged packets, this switch can handle all of these workloads. It can also connect to a 1 Gbps infrastructure, a 10 Gbps infrastructure, or both.

Because of the extreme flexibility of the BNT switch, you can take advantage of the technologies that they require for multiple environments. For 1 Gbps uplinks, they can take advantage of small form-factor pluggable (SFP) transceivers. For 10 Gbps uplinks, they can use either of the following options:

- Enhanced small form-factor pluggable (SFP+) transceivers (short-range or long-range) for longer distances
- Direct-attached copper (DAC) cables (also known as *twinax active cables*) for shorter distances

  DAC cables are more cost-effective, use less power, and can be up to 7 meters in length. They are ideal for connecting chassis together, connecting to a top-of-rack switch or to an adjacent rack.

Figure 8-4 shows the switch module.

Figure 8-4 BNT Virtual Fabric 10Gb Switch Module

Table 8-3 lists the part number to use to order the module.

Table 8-3 Part number and feature code

Description	Part number	Feature code
BNT Virtual Fabric 10Gb Switch Module	46C7191	1639

The part number includes the following items:

- *BNT Virtual Fabric 10Gb Switch Module Installation Guide*
- BNT user license agreement
- Documentation CD
- *Important Notices* document
- One BNT Virtual Fabric 10Gb Switch Module
- One filler module
- Three-meter, mini-Universal Serial Bus (USB)-to-DB9 serial console cable

Tip: SFP+ transceivers are not included. You must purchase them separately.

To communicate outside of the chassis, you must have SFP+ transceivers or SFP+ DAC cables connected. DAC cables have SFP+ transceivers on both ends. You have the flexibility to expand your bandwidth as desired, by using from one to ten connections per switch.

Table 8-4 lists the part number to use to order the SFP+ transceivers, FC cables, and DAC cables.

Table 8-4 IBM part numbers for ordering SFP+ transceivers, FC cables, and DAC cables

Description	Part number	Feature code
10Gb SFP+		
IBM 10GBase-SR 10GbE 850 nm Fiber SFP+ transceiver	44W4408	4942
BNT SFP+ transceiver	46C3447	5053
1Gb SFP+		
BLADE 1000BASE-T (RJ45) SFP transceiver	81Y1618	3268
BLADE 1000BASE-SX SFP transceiver	81Y1622	3269
DAC cables		
0.5 m Molex DAC SFP+ cable	59Y1932	3735
1 m Molex DAC SFP+ cable	59Y1936	3736
3 m Molex DAC SFP+ cable	59Y1940	3737
7 m Molex DAC SFP+ cable	59Y1944	3738
FC cables		
3 m Intel Connects optical cable	46D0153	3852
10 m Intel Connects optical cable	46D0156	3853

Description	Part number	Feature code
30 m Intel Connects optical cable	46D0159	3854

Table 8-5 lists additional transceivers and DAC cable options that are available directly from BNT.

Table 8-5 BNT part numbers for ordering SFP-based transceivers and cables

Description	Part number
BLADE 1000Base-T SFP	BN-CKM-S-T
BLADE 1000Base-SX SFP	BN-CKM-S-SX
BLADE 10GBase-LR SFP+	BN-CKM-SP-LR
SFP+ DAC - 1 M	BN-SP-CBL-1M
SFP+ DAC - 3 M	BN-SP-CBL-3M
SFP+ DAC - 7 M	BN-SP-CBL-7M

The BNT Virtual Fabric 10Gb Switch Module includes the following features and functions:

- Form-factor
  - Single-height, high-speed switch module
- Internal ports
  - Fourteen internal, auto-negotiating ports of 1 Gbps or 10 Gbps to the server blades
  - Two internal full-duplex 100 Mbps ports connected to the management module
- External ports
  - Up to ten 10 Gbps SFP+ ports (also designed to support 1 Gbps SFP, if required; flexibility of mixing 1 Gb and 10 Gb)

    The oversubscription ratio (14 internal ports to 10 external ports) is low. Therefore, the switch module is suitable for the most performance-intensive environments.
  - One 10/100/1000-Mb copper RJ45 for management or data
  - An RS-232 mini-USB connector for serial port that provides an additional means to install software and configure the switch module

- Scalability and performance
  - Autosensing 1 Gbps and 10 Gbps internal and external Ethernet ports for bandwidth optimization
  - Nonblocking architecture with wire-speed forwarding of traffic and full-line rate performance of 480-Gbps full duplex
  - MAC address learning, which includes automatic update and support for up to 32-Kb MAC addresses
  - Up to 128 IP interfaces per switch
  - Static, EtherChannel, and Link Aggregation Control Protocol (LACP) (IEEE 802.3ad) link aggregation, up to 100 Gb of total bandwidth per switch, up to 18 trunk groups, and up to eight ports per group
  - Support for jumbo frames (up to 12288 bytes)
  - Broadcast/multicast storm control
  - Internet Group Management Protocol (IGMP) snooping for limit flooding of IP multicast traffic (IGMP V1, V2, and V3)
  - Configurable traffic distribution schemes over trunk links, based on source and destination IP addresses, MAC addresses, or both
  - Fast port forwarding and fast uplink convergence for rapid Spanning Tree Protocol (STP) convergence
  - Stacking support for up to eight BNT Virtual Fabric 10Gb Switch Modules
- Availability and redundancy
  - Virtual Router Redundancy Protocol (VRRP) for Layer 3 router redundancy
  - IEEE 802.1D STP for providing Layer 2 redundancy with PVRST+
  - IEEE 802.1s Multiple STP (MSTP) for topology optimization, up to 128 STP instances supported by single switch
  - IEEE 802.1w Rapid STP (RSTP), providing rapid STP convergence for critical, delay-sensitive, traffic-like voice or video
  - Layer 2 trunk failover to support active and standby configurations of network adapter teaming on blades
  - Interchassis redundancy (Layer 2 and Layer 3)
- VLAN support
  - Up to 1024 VLANs supported per switch, VLAN numbers in the range 1 - 4095 (a dedicated VLAN for the management module connection only)
  - 802.1Q VLAN tagging support on all ports
  - Protocol-based VLANs

- Security
  - VLAN-based, MAC-based, and IP-based access control lists (ACLs)
  - 802.1X port-based authentication
  - Multiple user IDs and passwords
  - User access control
  - Remote Authentication Dial In User Service (RADIUS), Terminal Access Controller Access-Control System Plus (TACACS+), and Lightweight Directory Access Protocol (LDAP)
- Quality of service (QoS)
  - Up to eight queues per port
  - Support for IEEE 802.1p, IP type of service (ToS)/Differentiated Services Code Point (DSCP), and ACL-based (MAC and IP source and destination addresses, VLANs) traffic classification and processing
  - Traffic shaping and remarking based on defined policies
  - Eight Weighted Round-Robin (WRR) priority queues per port for processing qualified traffic
- Layer 3 functions
  - IP forwarding
  - IP filtering with ACLs (up to 4096 ACLs supported)
  - VRRP for router redundancy
  - Support for up to 128 static routes
  - Routing protocol support (Router Information Protocol (RIP) V1, RIP V2, Open Shortest Path First protocol (OSPF) V1, V2, and V3, BGP-4), up to 1024 entries in routing table
  - IPv6 routing, including static routes and OSPFv3 (requires firmware V6.3 or higher)
  - Support for Dynamic Host Configuration Protocol (DHCP) Relay
  - IPv6 host management
  - IPv6 forwarding based on static routes
- Manageability
  - Simple Network Management Protocol (SNMP) V1, V2, and V3
  - HTTP/HTTPS browser-based interface (BBI)
  - Industry-standard CLI (isCLI) and IBM Networking OS or AlteonOS CLI
  - Telnet interface for CLI

  - Secure Shell (SSH) v1/v2
  - Serial interface for CLI
  - Scriptable CLI
  - Firmware image update (Trivial File Transfer Protocol (TFTP) and File Transfer Protocol (FTP))
  - Network Time Protocol (NTP) for switch clock synchronization
  - BNT BLADEHarmony Manager support
- Monitoring
  - Switch LEDs for external port status and switch module status indication
  - Port mirroring to analyze network traffic that is passing through the switch
  - Change tracking and remote logging with syslog feature
  - Power-on self test (POST) diagnostic test
- Special functions
  - Serial over LAN
- Virtualization features
  - VMready
  - Virtual Fabric Adapter vNIC support
- Converged Enhanced Ethernet and FCoE features
  - Allows FC traffic to be transported over Ethernet links
  - FCoE Initialization Protocol snooping to enforce point-to-point links for FCoE traffic outside the regular FC topology
  - PFC (IEEE 802.1Qbb)

    This feature extends 802.3x standard flow control to allow the switch to pause traffic based on the 802.1p priority value in the VLAN tag of each packet
  - Enhanced Transmission Selection (ETS) (IEEE 802.1Qaz)

    ETS provides a method for allocating link bandwidth based on the 802.1p priority value in the VLAN tag of each packet.
  - DCBX (IEEE 802.1AB), which allows neighboring network devices to exchange information about their capabilities
  - Support for the QLogic Virtual Fabric Extension Module for IBM BladeCenter

    Support provides FCoE gateway functions inside the BladeCenter Chassis.

VMready is a unique solution that enables the network to be virtual machine (VM)-aware. The network can be configured and managed for virtual ports (vports), rather than just for physical ports. With VMready, as VMs migrate across physical hosts, their network attributes also migrate. Virtual machines can be added, moved, and removed, but retain the same ACLs, QoS, and VLAN attributes. VMready allows for a define-once-use-many configuration that evolves as the server and network topologies evolve. VMready works with all virtualization products, including VMware, Hyper-V, Xen, and KVM, without modification of virtual machine hypervisors or guest operating systems. VMready available as part of the 6.1 (and higher) software code.

VMready has the following compatibility with Virtual Fabric solutions:

- VMready *is not* supported with IBM Virtual Fabric Mode
- VMready *is* supported with Switch Independent Mode

The switch module supports the following Institute of Electrical and Electronics Engineers (IEEE) standards:

- IEEE 802.1D STP with PVRST+
- IEEE 802.1s MSTP
- IEEE 802.1w RSTP
- IEEE 802.1p Tagged Packets
- IEEE 802.1Q Tagged VLAN (frame tagging on all ports when VLANs are enabled)
- IEEE 802.1x port-based authentication
- IEEE 802.2 Logical Link Control
- IEEE 802.3ad Link Aggregation Control Protocol
- IEEE 802.3x Full-duplex Flow Control
- IEEE 802.3ab 1000BASE-T Gigabit Ethernet
- IEEE 802.3ae 10GBASE-SR 10Gb Ethernet fiber optics short range
- IEEE 802.3ae 10GBASE-LR 10Gb Ethernet fiber optics long range
- IEEE 802.3z 1000BASE-SX Gigabit Ethernet

The following network cables are supported for the BNT Virtual Fabric 10Gb Switch Module:

- 10GBASE-SR for 10 Gbps ports: 850Nm wavelength, multimode fiber, 50 µ or 62.5 µ (300 meters maximum), with LC duplex connector
- 1000BASE-T for RJ45 port:
  - UTP Category 6 (100 meters maximum)
  - UTP Category 5e (100 meters maximum)
  - UTP Category 5 (100 meters maximum)
  - EIA/TIA-568B 100-ohm STP (100 meters maximum)

For more information, see the following publications:

- *BNT Virtual Fabric 10Gb Switch ModuleBNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter*, TIPS0708:
- Documents in PDF format available at:

  http://www.ibm.com/support/docview.wss?uid=psg1MIGR-5080917

  - Application Guide
  - BBI (Browser-based Interface) Quick Guide
  - *BNT Virtual Fabric 10Gb Switch Module Installation Guide*
  - Command Reference
  - User License Agreement
  - Industry standard CLI Reference (isCLI Reference)
  - Release Notes

8.3.2 Cisco Nexus 4001I Switch Module

The Cisco Nexus 4001I Switch Module is a blade-switch solution for the BladeCenter H and HT chassis. It provides the server I/O solution that is required for high-performance, scale-out, virtualized, and non-virtualized x86 computing architectures. It is a line rate, low-latency, non-blocking, Layer 2, 10 Gbps Ethernet switch module that is fully compliant with FCoE and IEEE Data Center Bridging standards.

The Cisco Nexus 4001I Switch Module enables a standards-based, high-performance unified fabric that is running over 10 Gbps Ethernet in the blade server environment. This unified fabric enables consolidation of the following types of traffic over a single 10 Gbps Ethernet server network:

- LAN traffic
- Storage traffic (IP-based, such as iSCSI, network-attached storage (NAS), FC storage area network (SAN))
- High-performance computing (HPC) traffic

This offering works with BladeCenter Open Fabric Manager to provide all of the benefits of I/O virtualization at 10 Gbps speed.

Figure 8-5 shows the switch module.

Figure 8-5 Cisco Nexus 4001I Switch Module

Table 8-6 shows the part numbers to order these modules and additional options for them.

Table 8-6 Cisco Nexus 4001I Switch Module and related options part numbers and feature codes

Description	Part number	Feature code	Cisco part number
Cisco Nexus 4001I Switch Module for IBM BladeCenter	46M6071	0072	N4K-4001I-XPX
Software upgrade license for the Cisco Nexus 4001I Switch Module	49Y9983	1744	N4K-4001I-SSK9

The module part numbers include the following items:

- Cisco Console Cable RJ45-to-DB9
- Documentation CD
- Important Notices document
- One Cisco Nexus 4001I Switch Module
- One filler panel

Software upgrade license for the Cisco Nexus 4001I Switch Module

The Cisco Nexus 4001I Switch Module supports 10 Gbps Ethernet and FCoE. By having the software upgrade license for the Cisco Nexus 4001I Switch Module, part number 49Y9983, the switch can work in FCoE mode. When connected to a converged adapter in the server, this switch can route CEE packets to an upstream FCoE switch, which can then route the packets to the LAN or SAN.

SFP+ transceivers and copper cables

The Cisco Nexus 4001I Switch Module does not ship with SFP+ transceivers nor SFP+ copper cables. These transceivers and cables can be ordered from IBM, from Cisco Systems, or from authorized Cisco Systems resellers, as listed in Table 8-7.

Table 8-7 SFP+ transceivers and copper cables for the Cisco Nexus 4001I Switch Module

Description	IBM part number	IBM feature code	Cisco Systems part number
10Gb Ethernet SFP+ transceiver			
Cisco 10GBASE-SR SFP+ (MMF, 850-nm, LC)	88Y6054	A1A6	SFP-10G-SR(=)
Cisco 10GBASE-LR SFP+ (SMF, 1310-nm, LC)	None	None	SFP-10G-LR(=)
DAC Cables			
Cisco 1-m 10G SFP+ Twinax cable assembly, passive	None	None	SFP-H10GB-CU1M(=)
Cisco 3-m 10G SFP+ Twinax cable assembly, passive	None	None	SFP-H10GB-CU3M(=)
Cisco 5-m 10G SFP+ Twinax cable assembly, passive	None	None	SFP-H10GB-CU5M(=)
Gigabit Ethernet SFP+			
Cisco 1000BASE-T SFP transceiver (Cat. 5 copper wire, RJ-45)	88Y6058	A1A7	GLC-T(=)
Cisco 1000BASE-SX SFP transceiver (MMF, 850-nm, LC)	88Y6062	A1A8	GLC-SX-MM(=)
Cisco 1000BASE-LX/LH SFP transceiver (MMF/SMF, 1300-nm, LC)	None	None	GLC-LH-SM(=)

The Cisco Nexus 4001I Switch Module includes the following features and functions:

- Form-factor

  Single-height high-speed switch module

- External ports
  - Six 10 Gbps SFP+ ports that are operating at wire speed

    They support 1Gb SFP if required, with the flexibility of mixing 1 Gb and 10 Gb. Table 8-7 on page 182 lists supported transceivers and cables.

  - One 10/100/1000-Mb Ethernet copper RJ45 that is used for management
  - An RS-232 to RJ45 connector for a serial port that provides an additional means to configure the switch module

    The console cable is supplied with the switch module.
- Internal ports
  - Fourteen internal auto-negotiating ports of 1 Gbps or 10 Gbps to the server blades
  - Two internal full-duplex 100 Mbps ports connected to the management modules
- Scalability and performance
  - Autosensing 1 Gbps/10 Gbps internal and external Ethernet ports for bandwidth optimization
  - Non-blocking architecture with wire-speed forwarding of traffic and full-line rate performance of 400-Gbps full duplex
  - Forwarding rate of 300 million packets per second (mpps)
  - Low, predictable, and consistent latency of 1.5 microseconds regardless of packet size, traffic pattern, or enabled features on 10GbE interface
  - MAC address learning, including automatic update and support for up to 8 Kb MAC addresses
  - EtherChannels and LACP (IEEE 802.3ad) link aggregation, up to 60 Gbps of total uplink bandwidth per switch, up to seven trunk groups, and up to six ports per group
  - Support for jumbo frames (up to 9216 bytes)
  - Traffic suppression (unicast, multicast, and broadcast)
  - IGMP snooping to limit flooding of IP multicast traffic (IGMP V2 and V3)
  - Configurable traffic distribution schemes over EtherChannel links based on source/destination IP addresses, MAC addresses, or ports
  - Spanning Tree edge ports (formerly PortFast) for rapid STP convergence
- Availability and redundancy
  - Link state tracking
    - Mirror the state of the external ports on the internal Ethernet links
    - Allow the failover of the processor blade traffic to an operational external link on a separate Cisco Ethernet switch
  - IEEE 802.1D STP for redundant backbone connections and loop-free networks

  - IEEE 802.1s MSTP for grouping VLANs into a spanning-tree instance and for multiple forwarding paths for data traffic and load balancing.
  - IEEE 802.1w RSTP for rapid convergence of the spanning tree by immediately transitioning root and designated ports to the converting state
  - Configurable Unidirectional link detection (UDLD) for detecting and disabling unidirectional links

    This feature prevents a larger network failure if a unidirectional link is detected, reducing downtime in these situations.
- VLAN support
  - Up to 512 VLANs supported per switch with VLAN numbers 1 - 4000
  - 802.1Q VLAN tagging support on all ports
  - Private VLANs
- Security
  - VLAN-based, MAC-based, and IP-based ACLs
  - Role-based access control (RBAC) to restrict the authorization of the user to perform switch management functions
  - TACACS+, which is a proprietary feature for managing network security through a TACACS server
  - RADIUS for verifying the identity of, granting access to, and tracking activities of remote users
- Quality of service
  - Support for IEEE 802.1p class of service (COS), IP ToS/DSCP, IP Real Time Protocol, and ACL-based (MAC/IP source and destination addresses, VLANs) traffic classification and processing
  - Trust boundaries on incoming interfaces to automatically classify incoming packets into system classes based on packet COS value
  - Traffic shaping, guaranteeing bandwidth, or prioritizing based on defined policies
  - Up to eight egress queues per interface (one queue per COS) to process qualified traffic (support for strict priority and WRR COS policies)
- FCoE
  - Support for T11-compliant FCoE on all 10-Gigabit Ethernet interfaces
  - FCoE Initialization Protocol snooping to enable the switch to operate as a loss-less Ethernet bridge transparently forwarding FCoE packets
  - 802.1Q VLAN tagging for FCoE frames
  - PFC (IEEE 802.1Qbb)

PFC simplifies management of multiple traffic flows over a single network link and creates lossless behavior for Ethernet by allowing COS-based flow control.

  - ETS (IEEE 802.1Qaz)

    ETS enables consistent management of QoS at the network level by providing consistent scheduling of traffic types (such as IP and storage).

  - DCBX Protocol (IEEE 802.1AB)

    This protocol simplifies network deployment and reduces configuration errors by providing auto-negotiation of IEEE 802.1 DCB features between the NIC and the switch and between switches.

- IPv6 host support for IPv6 unicast addressing, IPv6 traffic processing, IPv6 applications support (syslog server, RADIUS, TACACS+, NTP, Telnet, and SSH)

  IPv6 traffic forwarding at Layer 3 is not supported.

- Manageability
  - CLI

    You can configure switches by using the CLI from an SSH V2 session, a Telnet session, or the console port. SSH provides a secure connection to the device.

  - XML Management Interface over SSH

    You can configure switches by using the XML management interface. This interface is a programming interface based on the NETCONF protocol that complements the CLI functions. For more information, see the *Cisco NX-OS XML Interfaces User Guide* at:

    http://www.cisco.com/en/US/docs/switches/datacenter/sw/nx-os/xml/user/guide/nxos_xml_interface.html

  - Cisco Data Center Manager support
  - SNMP V1, 2, and 3 support
  - Protected Mode feature to isolate switch management from the advanced management module (AMM) for increased security of the switch
  - Cisco Discovery Protocol (CDP) Versions 1 and 2 to aid in troubleshooting and reporting the misconfiguration of ports that connect to other devices that support CDP

- Monitoring
  - Switch LEDs for external port status and switch module status indication
  - Remote Network Monitoring (RMON)
  - Change tracking and remote logging with the syslog feature
  - Online diagnostic tests

  - Cisco Fabric Services
  - Session Manager
  - SPAN for local traffic monitoring

- Special functions

  - Serial over LAN

The switch module supports the following IEEE standards:

- IEEE 802.1D: Spanning Tree Protocol
- IEEE 802.1p: COS Prioritization
- IEEE 802.1Q: VLAN Tagging
- IEEE 802.1s: Multiple VLAN Instances of Spanning Tree Protocol
- IEEE 802.1w: Rapid Reconfiguration of Spanning Tree Protocol
- IEEE 802.3ab: 1000Base-T (10/100/1000 Ethernet over copper)
- IEEE 802.3ad: Link Aggregation Control Protocol
- IEEE 802.3ae 10 Gbps Ethernet over fiber optics

  - 10GBASE-SR
  - 10GBASE-LR

- IEEE 802.3z Gigabit Ethernet over fiber optics

  - 1000BASE-SX
  - 1000BASE-LX

- IEEE 802.3ab 1000BASE-T over copper twisted pair

The network cable has the following requirements:

- 10GBASE-SR

  Table 8-8 lists the 10GBASE-SR cabling specifications.

  Table 8-8 10GBASE-SR cabling specifications

Wavelength	Cable type	Core size (microns)	Modal bandwidth (MHz/km)	Maximum cable length
850 nm	MMF	62.5	160	85 feet (26 m)
850 nm	MMF	62.5	200	108 feet (33 m)
850 nm	MMF	50	400	217 feet (66 m)
850 nm	MMF	50	500	269 feet (82 m)
850 nm	MMF	50	2000	984 feet (300 m)

- 10GBASE-LR

  Table 8-9 lists the 10GBASE-LR cabling specifications.

Table 8-9 10GBASE-LR cabling specifications

Wavelength	Cable type	Core size (microns)	Modal bandwidth (MHz/km)	Maximum cable length
1310 nm	SMF	G.652	Not applicable	6.2 miles (10 km)

- 1000BASE-SX

  Table 8-10 lists the 1000BASE-SX cabling specifications.

Table 8-10 1000BASE-SX cabling specifications

Wavelength	Cable type	Core size (microns)	Modal bandwidth (MHz/km)	Maximum cable length
850 nm	MMF	62.5	160	722 feet (220 m)
850 nm	MMF	62.5	200	902 feet (275 m)
850 nm	MMF	50	400	1640 feet (500 m)
850 nm	MMF	50	500	1804 feet (550 m)

- 1000BASE-LX

  Table 8-11 lists the 1000BASE-LX cabling specifications.

Table 8-11 1000BASE-LX cabling specifications

Wavelength	Cable type	Core size (microns)	Modal bandwidth (MHz/km)	Maximum cable length
1310 nm	MMF	62.5	500	1804 feet (550 m)
1310 nm	MMF	50.0	400	1804 feet (550 m)
1310 nm	MMF	50.0	500	1804 feet (550 m)
1310 nm	SMF	G.652	Not applicable	6.2 miles (10 km)

- 1000BASE-T
  - UTP Category 6 (100 meters maximum)
  - UTP Category 5e (100 meters maximum)
  - UTP Category 5 (100 meters maximum)
  - EIA/TIA-568B 100-ohm STP (100 meters maximum)

For more information, see the following documentation:

- *Configuration Guide for the Cisco Nexus 4001I Switch Module*
- *Hardware Installation Guide for the Cisco Nexus 4001I Switch Module*
- *Getting Started Guide for the Cisco Nexus 4001I Switch Module*
- *Command Reference Guide for the Cisco Nexus 4001I Switch Module*

To download these publications, go to the Cisco Nexus 4001I Switch Module - IBM BladeCenter page at:

http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5082494

In addition, for information about the Cisco Nexus 4001I Switch Module, see *Cisco Nexus 4001I Switch Module for IBM BladeCenter*, TIPS0754.

8.3.3 Brocade Converged 10GbE Switch Module for IBM BladeCenter

The Brocade Converged 10GbE Switch Module and Brocade 2-Port 10Gb CNA are part of a leading converged Ethernet solution for IBM BladeCenter. It offers FC investment protection, maximum bandwidth and performance, and simplicity in a converged environment. The Brocade Converged 10GbE Switch Module also features Dynamic Ports on Demand capability through the Port Upgrade Key. With this capability, you can enable any combination of Fibre Channel and Ethernet ports based on your infrastructure requirements. You can also experience converged Ethernet benefits without significant investment.

The Brocade CNAs and Brocade Converged 10GbE Switch Modules for IBM BladeCenter are the first converged FCoE BladeCenter solution in the industry with basic Fibre Channel and Ethernet ports built into a single switch module. The converged switch module provides up-front reduction in SAN or LAN cost and complexity. It requires fewer adapters, switch modules, and cabling, which means an up-front reduction in your network investment.

Figure 8-6 shows the Brocade Converged 10GbE Switch Module for IBM BladeCenter.

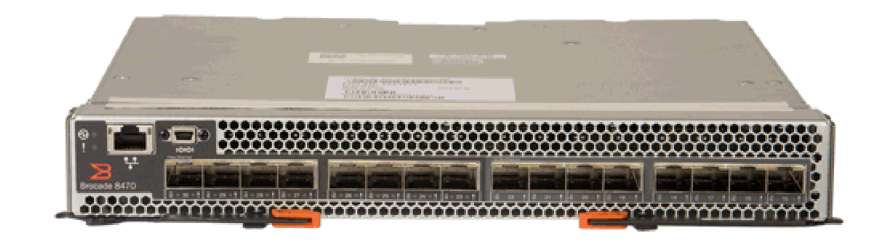

Figure 8-6 Brocade Converged 10 GbE Switch Module for IBM BladeCenter

Benefits

The Brocade Converged 10GbE Switch Module with a supported CNA provides the following benefits:

- Highly integrated BladeCenter switch module

  The Brocade Converged 10GbE Switch Module offers one of the best integrated I/O solutions of the industry. The compact design incorporates Ethernet and Fibre Channel switching. It provides a total of 30 ports, which include eight external 10 Gbps Ethernet CEE ports for LAN connectivity, eight external 8Gb FC ports for storage and SAN connectivity, and virtual ports. With the high integration of this module, you can achieve all of your networking and storage I/O needs with a single module.

- Unique flexible ports and server deployment

  By using the Dynamic Port on Demand (DPOD) feature of the Brocade Converged 10GbE Switch Module, you ca activate any combination of internal or external ports. You can balance server and I/O port assignments to accommodate the workloads of applications and satisfy business needs. With DPOD, you can enable only 10GbE ports for networking needs, only FC ports for storage support, or any combination of Ethernet and FC ports.

- Support for Data Center Fabric Manager

  IBM System Storage® Data Center Fabric Manager (DCFM) is available for use with the switch. DCFM is fully integrated with IBM Systems Director for end-to-end data center management. You can download a 75-day trial version from:

  `http://www.ibm.com/systems/storage/san/b-type/dcfm/`

- Highest bandwidth and low latency

  The total available bandwidth of 144 Gbps (80 Gbps Ethernet and 65 Gbps FC) accessible through external ports means that the switch module can support high demanding applications and storage solutions. In addition to the high bandwidth available on external ports, the switch module supports an additional 140 Gbps on 14 internal 10GbE ports. The switch module uses cut-through and nonblocking to deliver high performance and low latency for demanding virtual applications and high-speed environments.

- Low total cost of ownership

  The highly integrated design delivers the lowest total cost of ownership (TCO). The Brocade Converged 10GbE Switch Module is two switches in one, Ethernet and Fibre Channel, which reduces the cost of acquisition. The switch module is built around a pay-as-you-go model, so that you have a lower price entry point and can add ports as your business needs dictate.

  The switch module ships with two 10GbE SFP+ transceivers to further lower initial investments. Because it is designed for seamless integration into existing networking and storage environments, this module protects existing investments. Overall, with this integrated switch module, your server infrastructure will use fewer hardware components with lower costs and higher reliability, which further drives a lower TCO.

Features

The converged switch modules have the following features:

- Eight external 10 Gb Converged Enhanced Ethernet ports
- Eight external auto-negotiated FC ports (2 Gbps, 4 Gbps, or 8 Gbps)
- Fourteen internal auto-negotiated 10 GbE ports (1 Gbps or 10 Gbps)
- Two internal full-duplex 100 Mbps Ethernet interfaces for management purposes
- One external RS232 console port with a mini-USB interface for serial console management
- One external 10/100/1000 Mb RJ45 Ethernet copper port for management
- CEE features
- PFC: IEEE 802.1Qbb
- ETS: IEEE 802.1Qaz
- DCBX
- Layer 2 features
- Layer 2 VLANs: 4096
- VLAN Encapsulation 802.1Q
- RSTP
- MSTP (802.1s): 16 instances

- LACP IEEE- 802.3ad
- Brocade enhanced frame-based trunking
- Advanced PortChannel hashing based on Layer 2, 3, and 4 information
- Pause Frames (802.3x)
- Storm Control (unicast, multicast, and broadcast)
- Address Resolution Protocol (ARP) RFC 826
- Layer 2 security
- Ingress ACLs
- Standard and extended Layer 2 ACLs
- VLAN-based ACLs (VACLs)
- Port-based ACLs (PACLs)
- ACL statistics
- Port-based Network Access Control: IEEE 802.1X
- Layer 2 QoS
- Eight priority levels for QoS
- IEEE 802.1p COS
- Eight hardware queues per port
- Per-port QoS configuration
- COS trust: IEEE 802.1p
- Per-port virtual output queuing
- COS-based egress queuing
- Egress strict priority queuing
- Egress port-based scheduling: WRR
- Layer 3 features
- Static IP routes
- External FC ports that can operate as fabric ports (F_ports), fabric loop ports (FL_ports), or expansion ports (E_ports)
- FC fabric services
- Simple Name Server (SNS)
- Registered State Change Notification (RSCN)
- Dynamic Path Selection (DPS)
- Enhanced Group Management (EGM)

- POST diagnostics and status reporting
- Optional: Advanced Performance Monitor
- Optional: ISL Trunking
- Optional: Fabric Watch

The following software features come with the switch modules:

- Brocade Web Tools
- 16-port licensing

The switch supports the following fabric management. All management connections go through the management module, except the direct serial connection, which goes through the mini-USB port.

- Web interface through Web Tools
- CLI through the Telnet program
- A terminal emulation program connection to the mini-USB port interface
- IBM System Storage DCFM application
- Switch SNMP agent

For more information, see the following references:

- *Brocade Converged 10GbE Switch Module for IBM BladeCenter*, TIPS0789
- IBM Systems Networking

  http://www.ibm.com/systems/networking/index.html

8.4 10Gb Ethernet Pass-Thru Module for IBM BladeCenter

The 10Gb Ethernet Pass-Thru Module for IBM BladeCenter is ideal for clients who want to enable end-to-end non-blocking 10 Gbps setup within the chassis. This device supports both Ethernet and CEE packets, so that clients can connect a BladeCenter chassis to an FCoE-capable top-of-rack switch.

The 14 10-Gbps Uplink ports are based on optical SFP+ technology to offer the highest performance and maintain industry standard connectivity. This offering also works with BladeCenter Open Fabric Manager, by providing all the benefits of I/O Virtualization at 10 Gbps speeds.

This module is ideal if you want to keep the networking outside of the chassis and use only external switches. Although this module is not as expensive as a switch module, when you compare total price per ports of a solution, this model might be more expensive. However, this module can give you extra flexibility by providing one port per blade and low maintenance.

Figure 8-7 Shows the Ethernet Pass-Thru Module for IBM BladeCenter

Figure 8-7 10 Gb Ethernet Pass-Thru Module for IBM BladeCenter

Features and specifications

The 10Gb Ethernet Pass-Thru Module includes the following features and functions:

- Single-wide high-speed switch module
- 14 internal 10 Gbps ports to the server blades (no auto-negotiation)
- Up to 14 10-Gbps SFP+ uplink ports (SFP+ modules are not included)
- Support for SR, LR, and DAC copper cables
- Direct one-to-one mappings of external and internal ports (no configuration required)
- Simple setup and monitoring through AMM
- Management-module inter-integrated Circuit (I2C) interface that provides vital product data (VPD) and register access
- Provides self-test and diagnostics capability
- No Serial over LAN support

For more information, see the following references:

- *10Gb Ethernet Pass-Thru Module for IBM BladeCenter*, TIPS0715
- IBM Systems Networking

  `http://www.ibm.com/systems/networking/index.html`

9

Configuring the components

When configuring the IBM Virtual Fabric Mode solution (see Part 2, “IBM Virtual Fabric Mode” on page 17), most of the steps are performed on the BNT Virtual Fabric 10Gb Switch Module. The switch module transfers configuration settings to the Emulex 10GbE Virtual Fabric Adapter by using Data Center Bridging Exchange (DCBX). The Emulex 10GbE Virtual Fabric Adapter requires little configuration. You need only to enable virtual network interface card (vNIC) mode by using the Unified Extensible Firmware Interface (UEFI) utility of the blade server.

In contrast, for Switch Independent Mode, you must set vNIC parameters for the adapter by using the blade server UEFI utility. You also must configure the switch module to match the virtual local area network (VLAN) ID configuration on the vNICs and the extended network infrastructure.

This chapter explains how to configure components in Switch Independent Mode. It includes the following sections:

- Configuring Virtual Fabric adapters
- Configuring the BNT Virtual Fabric 10Gb Switch Module
- Upgrading the firmware for the Cisco Nexus 4001I Switch Module
- Creating vNICs with the Broadcom 2-port 10Gb Virtual Fabric Adapter
- Configuring vNIC VLANs
- vNIC and VMready
- Understanding vNIC representation in an operating system

9.1 Configuring Virtual Fabric adapters

This section explains the firmware upgrade and adapter configuration in UEFI for Emulex 10GbE Virtual Fabric Adapter II and the Broadcom 2-port 10Gb Virtual Fabric Adapter.

9.1.1 BladeCenter Open Fabric Manager Advanced

BladeCenter Open Fabric Manager Advanced supports configuration of vNIC parameters on Virtual Fabric Adapters. You can configure the channel bandwidth from the Templates panel for the virtual NIC ports. The following vNIC parameters are supported by BladeCenter Open Fabric Manager:

- MAC address of virtual NIC
- VLAN ID
- Minimum bandwidth
- Maximum bandwidth
- Priority

The next revision of this book will include detailed descriptions of vNIC parameter configuration by using BladeCenter Open Fabric Manager. For more information about BladeCenter Open Fabric Manager Advanced, see the “Using OFM-Advanced upgrade” topic in the IBM Director Information Center at:

http://pic.dhe.ibm.com/infocenter/director/v5r2/index.jsp?topic=/bofm_1.00/btp0_bofm_tp_using_bofm_advanced.html

9.1.2 Upgrading the firmware for Emulex 10GbE Virtual Fabric Adapter II

For information about upgrading Emulex 10GbE Virtual Fabric Adapter II, see 5.1.1, “Upgrading firmware” on page 52.

Important: Support for Switch Independent Mode by the Emulex 10GbE Virtual Fabric Adapter II requires a firmware upgrade. Make sure that you have the latest firmware installed.

9.1.3 Configuring the Emulex 10GbE Virtual Fabric Adapter II in UEFI

For information about configuring Emulex 10GbE Virtual Fabric Adapter II in UEFI, see 5.1.2, “Configuring the adapter in the UEFI” on page 54.

9.1.4 Configuring the Broadcom 2-port 10Gb Virtual Fabric Adapter in UEFI

To configure vNIC partitions on the physical ports of the Broadcom 2-port 10Gb Virtual Fabric Adapter, we use the blade server UEFI to configure the vNIC parameters. Alternatively, you can use the ROM configuration utility option of the Broadcom Virtual Fabric Adapter. To access this option, press Ctrl+S when prompted. For more information, see 9.4, “Creating vNICs with the Broadcom 2-port 10Gb Virtual Fabric Adapter” on page 209.

1. Power on or restart the blade.
2. When prompted during POST, press F1.
3. In the System Configuration and Boot Management panel (Figure 9-1) of the UEFI utility, select **System Settings**.

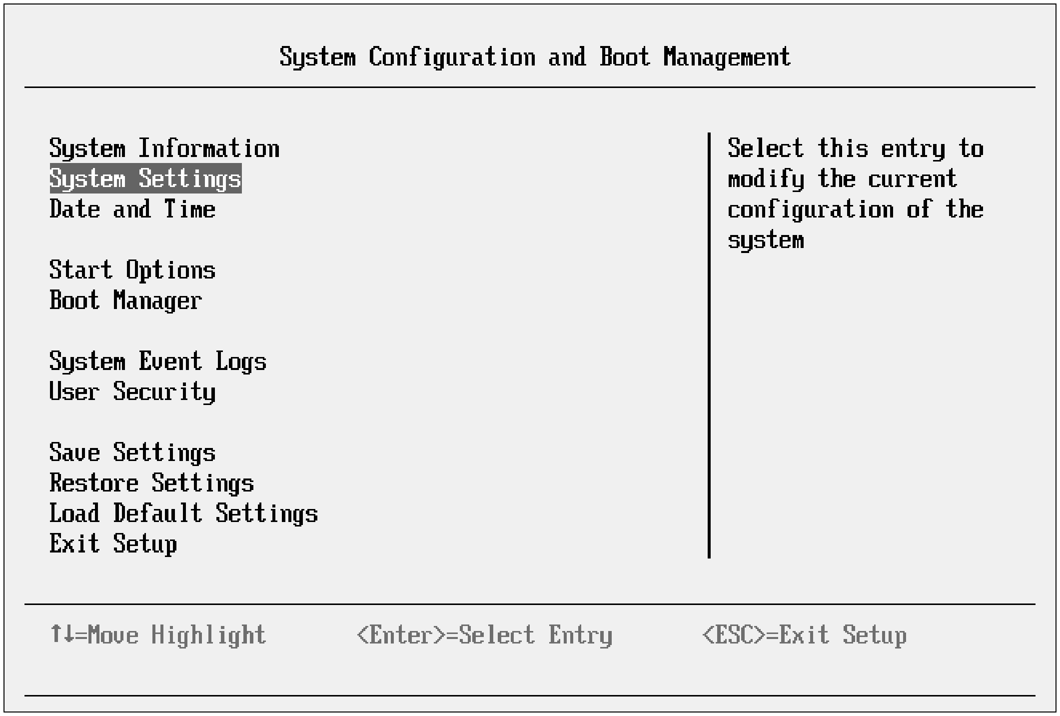

Figure 9-1 System Configuration and Boot Management panel of the Blade server UEFI utility

4. In the System Settings panel (Figure 9-2), select **Network**.

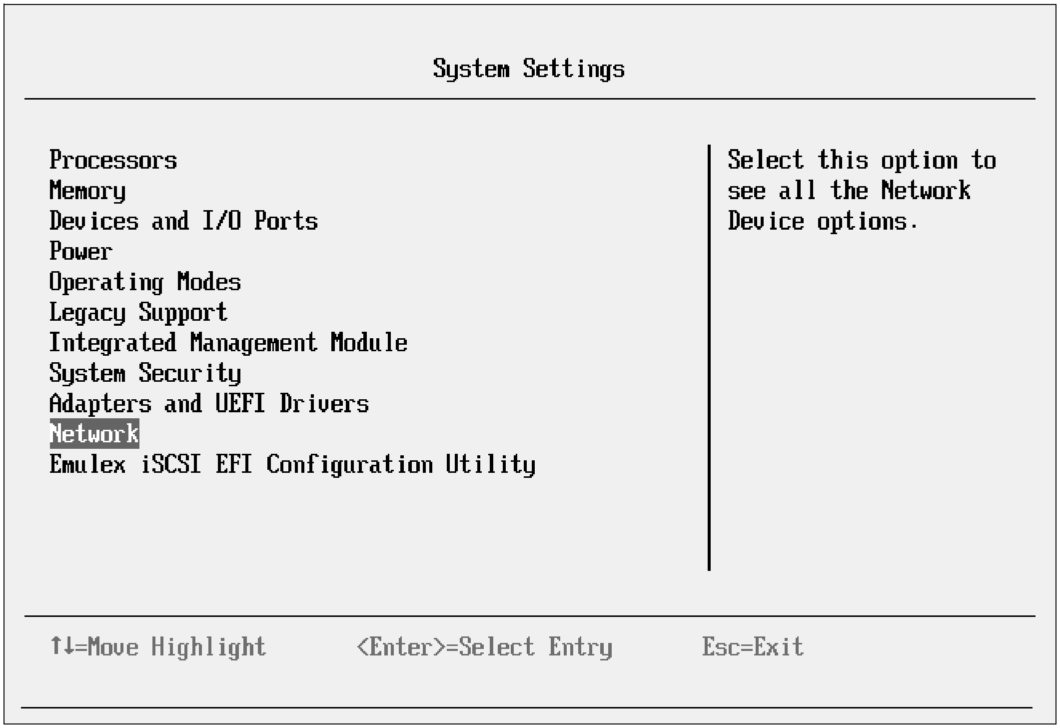

Figure 9-2 System Settings panel

5. In the Network panel (Figure 9-3), under Network Device List, select the first 10 Gbps network port.

   Notice that four Broadcom NetXtreme II network ports are listed. Two 1-Gb Ethernet ports belong to the onboard Broadcom network controller. The two 10-Gbps network ports belong to the Broadcom 2-port 10Gb Virtual Fabric Adapter.

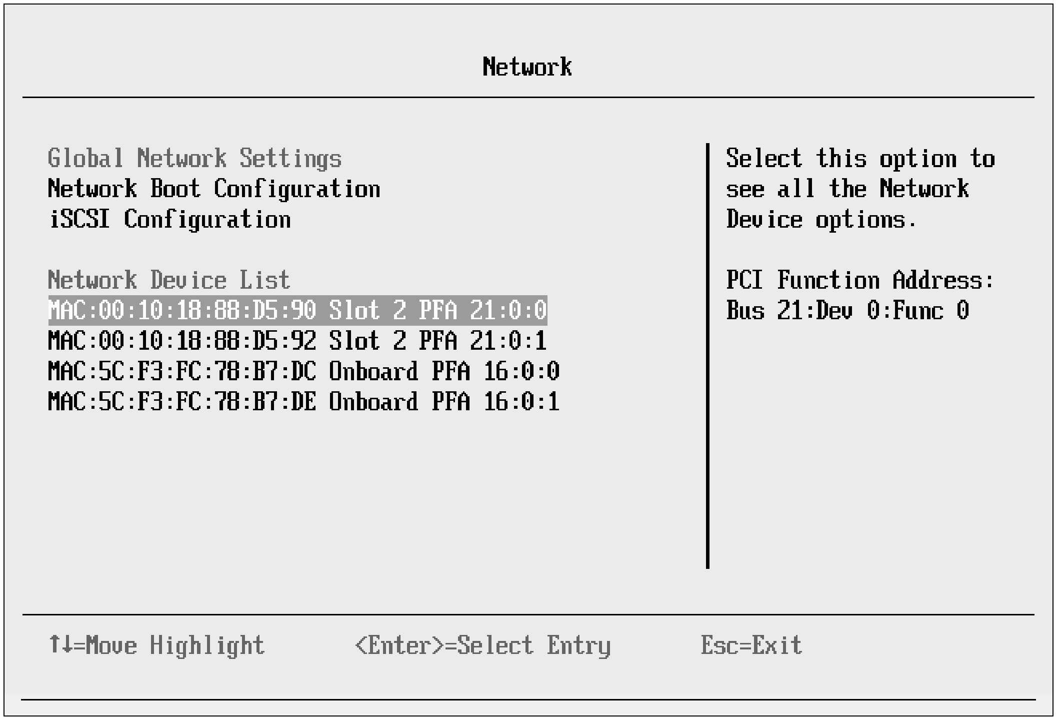

Figure 9-3 Network panel

6. In the Broadcom Main Configuration panel (Figure 9-4), click **Broadcom NeXtreme II 10 Gigabit Ethernet - 00:10:18:88:D5:90**.

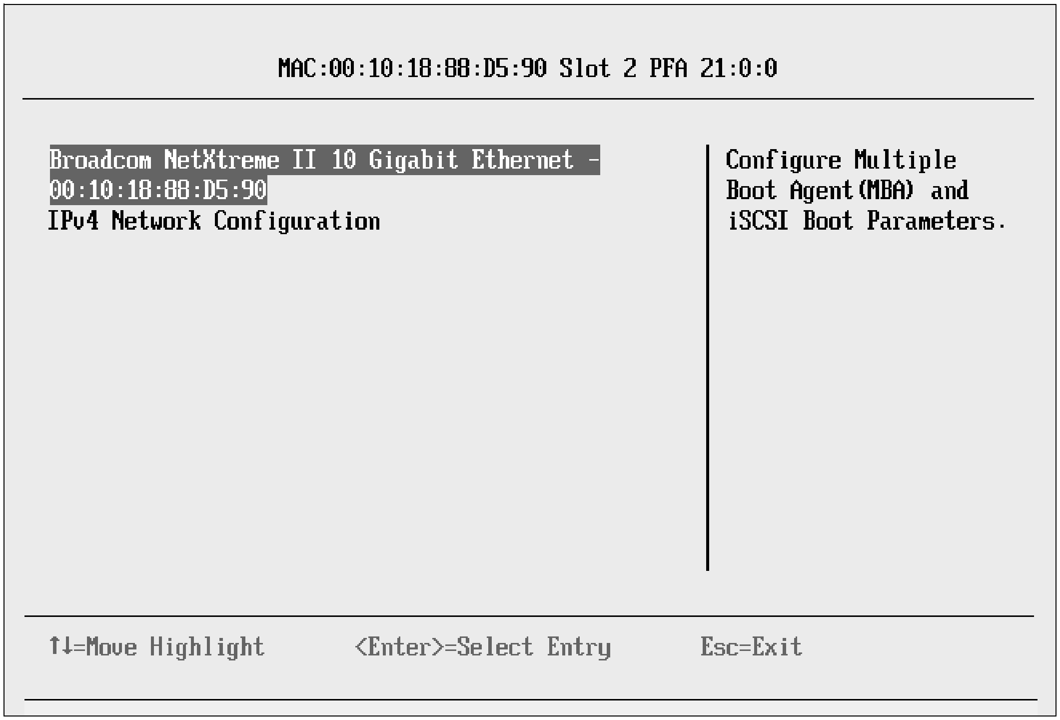

Figure 9-4 Broadcom selection panel

7. In the Broadcom Main Configuration Page panel (Figure 9-5), select **Device Configuration Menu** to enable the NIC partition.

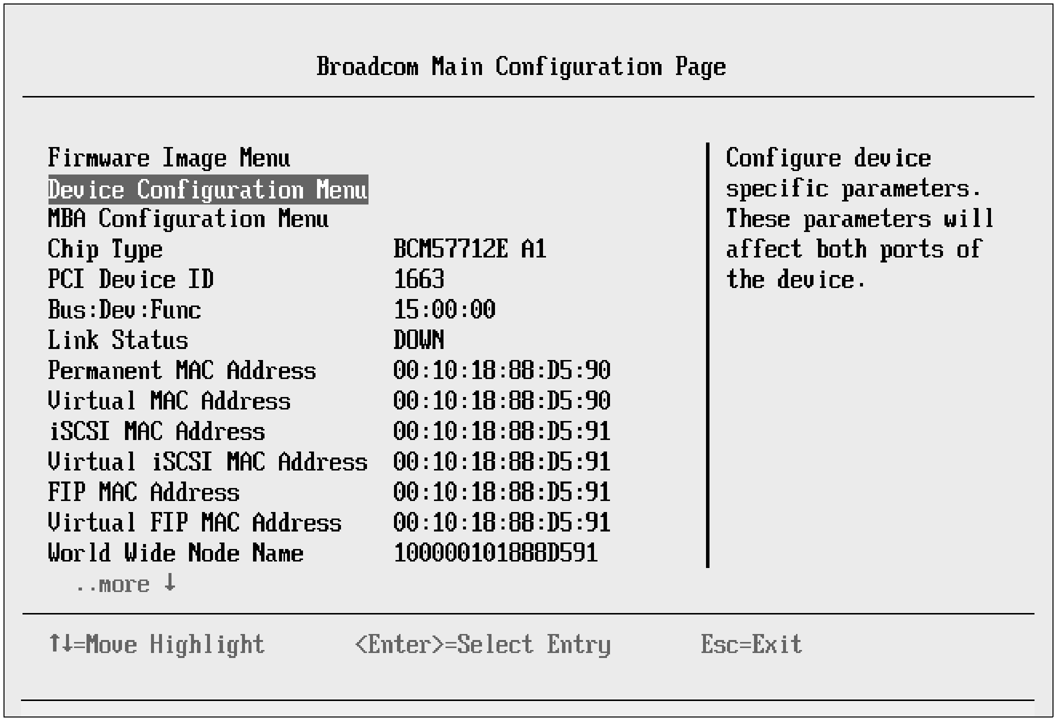

Figure 9-5 Broadcom Main Configuration Page panel

8. In the Device Configuration panel (Figure 9-6), for NIC Partition, select **Enabled**. With the vNIC partitions enabled, press Esc.

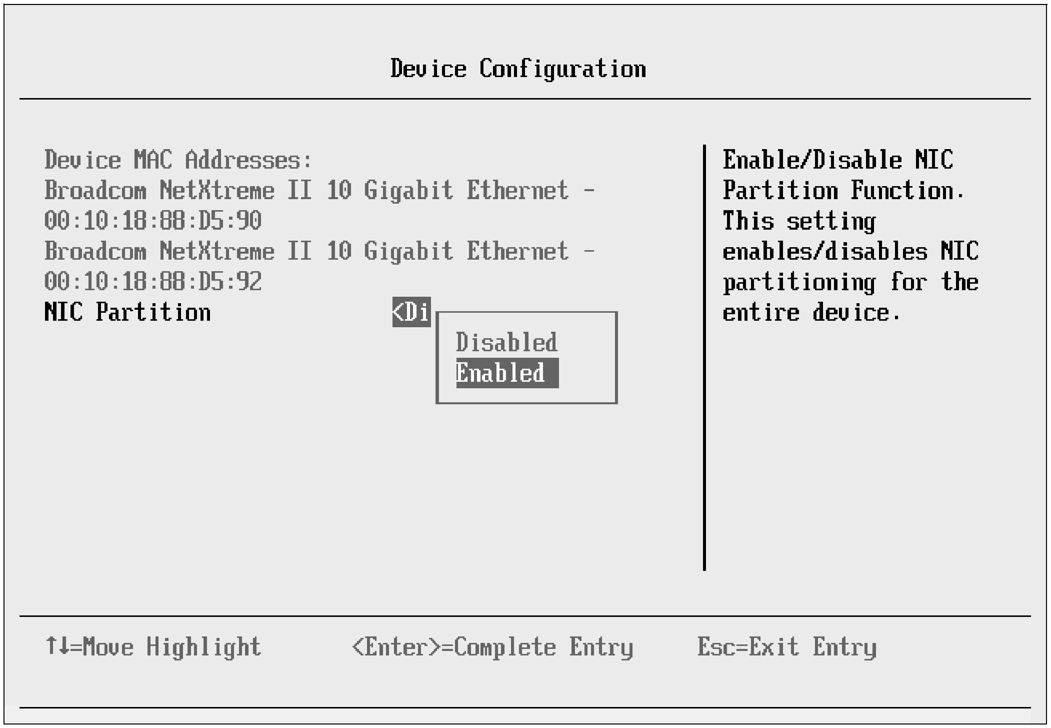

Figure 9-6 Device Configuration panel to enable the NIC partition

9. Back in the Broadcom Main Configuration Page panel (Figure 9-7), click **NIC Partitioning Configuration Menu**.

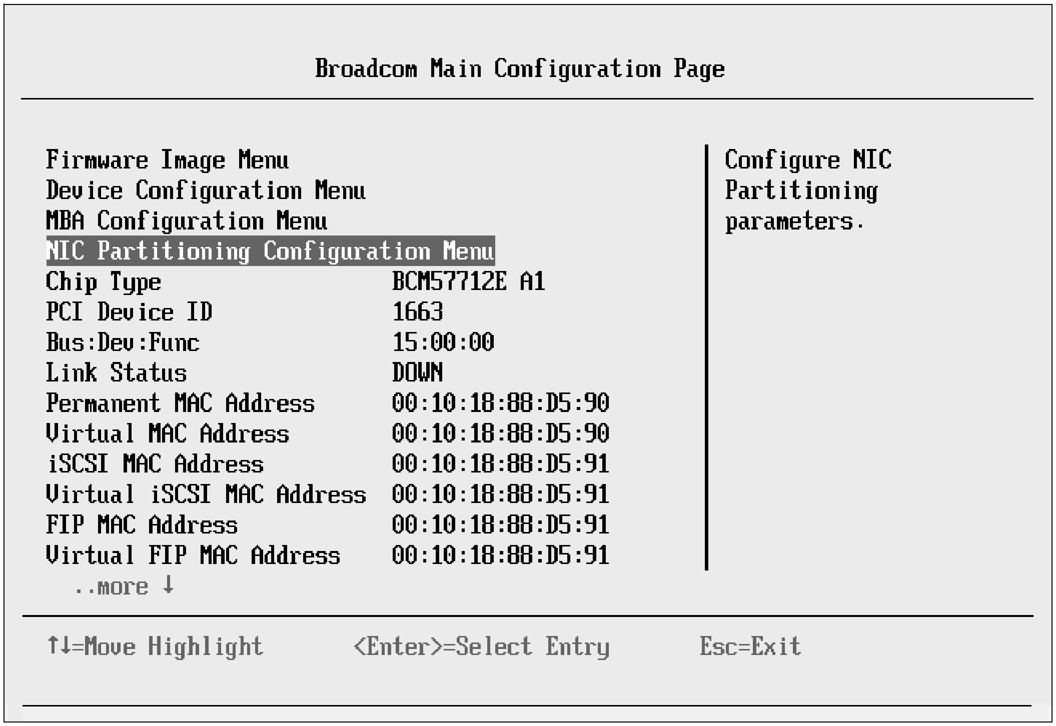

Figure 9-7 NIC Partitioning Configuration Menu selection window

10. In the NIC Partition Configuration panel (Figure 9-8), click **Global Bandwidth Allocation Menu**.

Notice that in addition to Global Bandwidth Allocation Menu option to configure bandwidth allocations, this panel shows four vNICs.

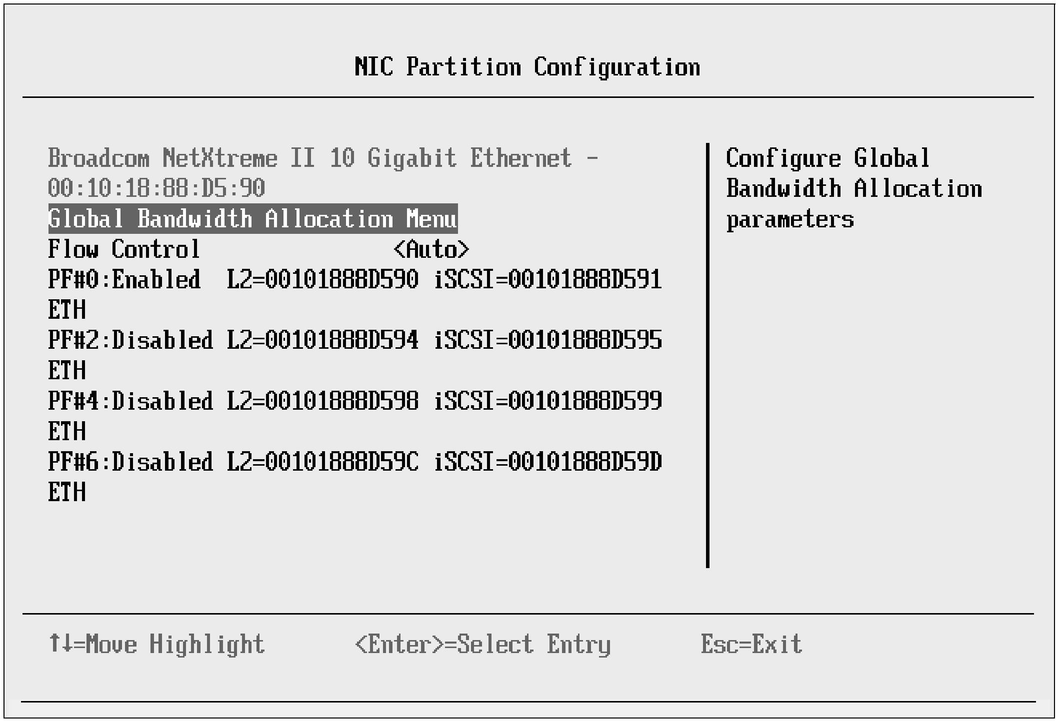

Figure 9-8 NIC Partition Configuration panel

11. In the Global Bandwidth Allocation Menu panel (Figure 9-9), set the global bandwidth for each port. The total of all four partitions must equal 0 or 100. For our example, the relative bandwidth weight is `0`, and our partition maximum is `25`.

    a. Set the value for Relative Bandwidth Weight.

       You can use the Relative Bandwidth Weight parameter to specify the guaranteed bandwidth in an overload situation. For example, if the parameter is set to 20, the NIC partition is guaranteed to transmit 2 Gbps of bandwidth (20% of total bandwidth). The sum of all four NIC partitions must be either 0 or 100. If it is zero, the relative bandwidth weight mechanism is disabled, and a round-robin bandwidth allocation mechanism is used instead.

    b. Set the value for Maximum Bandwidth.

       You can use this parameter to set the maximum bandwidth as a percentage of total physical port bandwidth that the NIC partition can transmit. You can set Maximum Bandwidth to 100 on all four partitions. In this case, a partition can use all 10 Gbps of bandwidth when no other partition is communicating. If this parameter is set to a value lower than 100, this maximum bandwidth is available to the NIC partition.

```
                    Global Bandwidth Allocation Menu
-------------------------------------------------------------------
Partition 1 Relative      [0]            | Configure relative
Bandwidth Weight                         | bandwidth weight.
Partition 1 Maximum       [25]           | Valid range - 0...100
Bandwidth                                | percent. The sum total
Partition 2 Relative      [0]            | of all four partitions
Bandwidth Weight                         | must be equal 0 or
Partition 2 Maximum       [25]           | 100. A total of 0
Bandwidth                                | indicates Relative
Partition 3 Relative      [0]            | BandWidth Weight is
Bandwidth Weight                         | disabled.
Partition 3 Maximum       [25]           |
Bandwidth                                |
Partition 4 Relative      [0]            |
Bandwidth Weight                         |
  ..more ↓
-------------------------------------------------------------------
↑↓=Move Highlight      <Enter>=Select Entry      Esc=Exit
```

Figure 9-9 Global Bandwidth Allocation Menu panel

12. After you enable the NIC partitions and set the bandwidth allocation parameters, exit to the main menu.
13. Save your settings.

9.2 Configuring the BNT Virtual Fabric 10Gb Switch Module

For more information about configuring the BNT Virtual Fabric 10Gb Switch Module, see 5.2, “Configuring the BNT Virtual Fabric 10Gb Switch Module” on page 59. The BNT Virtual Fabric 10Gb Switch Module supports Switch Independent Mode.

9.3 Upgrading the firmware for the Cisco Nexus 4001I Switch Module

At the time this book was written, the latest Cisco NX-OS version was 4.1.2.E1.1F for the Cisco Nexus 4001I Switch Module. The firmware update package is available on the IBM support website at:

http://www.ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5086542

When upgrading firmware on the BNT Virtual Fabric 10Gb Switch Module, you transfer the new code directly from the TFTP server to an image bank on the switch module. In contrast, with Cisco, you can transfer the image from an FTP or TFTP server to a directory in the NX-OS file system.

To upgrade the firmware by installing the new kickstart and system image files:

1. Copy the following image files from the FTP or TFTP server to the Cisco Nexus 4001I Switch Module file system:
   - **`n4000-bk9-kickstart.4.1.2.E1.1f.bin`**, which is the NX-OS kickstart image file.
   - **`n4000-bk9.4.1.2.E1.1f.bin`**, which is the NX-OS system image file.

   The IP address of the FTP server in our example is 9.42.171.50. The destination directory in the NX-OS file system is the `bootflash:` directory. We specify this information as parameters in the following **`copy`** commands:

   - `copy ftp://9.42.171.50/n4000-bk9-kickstart.4.1.2.E1.1f.bin bootflash: vrf chassis-management`

– `copy ftp://9.42.171.50/n4000-bk9.4.1.2.E1.1f.bin bootflash: vrf chassis-management`

Figure 9-10 shows the output of these commands. At the prompt, enter the user name and password to access the FTP server.

```
ibm-switch-2# copy ftp://9.42.171.50/n4000-bk9-kickstart.4.1.2.E1.1f.bin bootflash: vrf
chassis-management
Enter username: cisco
Password:
***** Transfer of file Completed Successfully *****
ibm-switch-2# copy ftp://9.42.171.50/n4000-bk9.4.1.2.E1.1f.bin bootflash: vrf
chassis-management
Enter username: cisco
Password:
***** Transfer of file Completed Successfully *****
```

Figure 9-10 Copying the kickstart and system image files to a directory

2. Before you install the files, verify that they are displayed in the `bootflash:` directory, as shown in Figure 9-11.

```
ibm-switch-2# dir
         49     Jul 03 03:58:50 2010  ..tmp-kickstart
         39     Jul 03 03:58:55 2010  ..tmp-system
   20624384     Jul 03 03:58:50 2010  .tmp-kickstart
   74274426     Jul 03 03:58:55 2010  .tmp-system
      49152     Jul 03 03:58:49 2010  lost+found/
   20624384     Jul 03 03:54:12 2010  n4000-bk9-kickstart.4.1.2.E1.1d.bin
   20699136     May 18 13:49:36 2011  n4000-bk9-kickstart.4.1.2.E1.1f.bin
   74274426     Jul 03 03:54:49 2010  n4000-bk9.4.1.2.E1.1d.bin
   74281634     May 18 13:50:21 2011  n4000-bk9.4.1.2.E1.1f.bin
       4096     Jul 03 03:56:30 2010  vdc_2/
       4096     Jul 03 03:56:31 2010  vdc_3/
       4096     Jul 03 03:56:31 2010  vdc_4/

Usage for bootflash://
  333664256 bytes used
  290123776 bytes free
  623788032 bytes total
```

Figure 9-11 The bootflash directory

3. To upgrade the NX-OS with the new files, use the **install all** command. Specify the **kickstart** and **system** keywords and image files as parameters, as in the following example:

```
install all kickstart bootflash:n4000-bk9-kickstart.4.1.2.E1.1f.bin
system bootflash:n4000-bk9.4.1.2.E1.1f.bin
```

Figure 9-12 shows an example of using the **install all** command with the **kickstart** and **system** keywords and image files as parameters.

```
ibm-switch-2# install all kickstart bootflash:n4000-bk9-kickstart.4.1.2.E1.1f.bin system
bootflash:n4000-bk9.4.1.2.E1.1f.bin

Verifying image bootflash:/n4000-bk9-kickstart.4.1.2.E1.1f.bin for boot variable
"kickstart".
[####################] 100% -- SUCCESS

Verifying image bootflash:/n4000-bk9.4.1.2.E1.1f.bin for boot variable "system".
[####################] 100% -- SUCCESS

Lines deleted for clarity

Switch will be reloaded for disruptive upgrade.
Do you want to continue with the installation (y/n)?  [n]  y

Install is in progress, please wait.

Setting boot variables.
[####################] 100% -- SUCCESS

Performing configuration copy.
[####################] 100% -- SUCCESS

Module 1: Refreshing compact flash and upgrading bios/loader/bootrom.
Warning: please do not remove or power off the module at this time.
[####################] 100% -- SUCCESS

Finishing the upgrade, switch will reboot in 10 seconds.
```

Figure 9-12 Upgrading firmware for the Cisco Nexus 4001I Switch Module

The switch module automatically reboots when the upgrade is finished.

For more information about Cisco Nexus 4001I Switch Module CLI commands, see the *Command Reference Guide for Cisco Nexus 4001I Switch Module* at:

http://www.ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5082494

You can configure the Cisco Nexus 4001I Switch Module by using Cisco Nexus OS (NX-OS) command-line interface (CLI) commands. For information about configuring this switch module, see 10.2.4, “Configuring the Cisco Nexus 4001I Switch Module” on page 272.

9.4 Creating vNICs with the Broadcom 2-port 10Gb Virtual Fabric Adapter

vNICs that use Switch Independent Mode are created and configured on the Virtual Fabric adapter directly, by using the UEFI utility (Emulex and Broadcom adapters) or a preboot environment (available with the Broadcom adapter only).

Configuration and support of the virtual NICs are carried out using a Broadcom application, available in the pre-boot environment (Ctrl+S during boot when prompted), or by using the UEFI utility. This section describes the preboot method.

To configure the Broadcom 2-port 10Gb Virtual Fabric Adapter by using the preboot utility:

1. Power on or restart the blade.
2. When prompted, press Ctrl+S.

3. From the Device List panel (Figure 9-13), highlight the device of interest and press Enter.

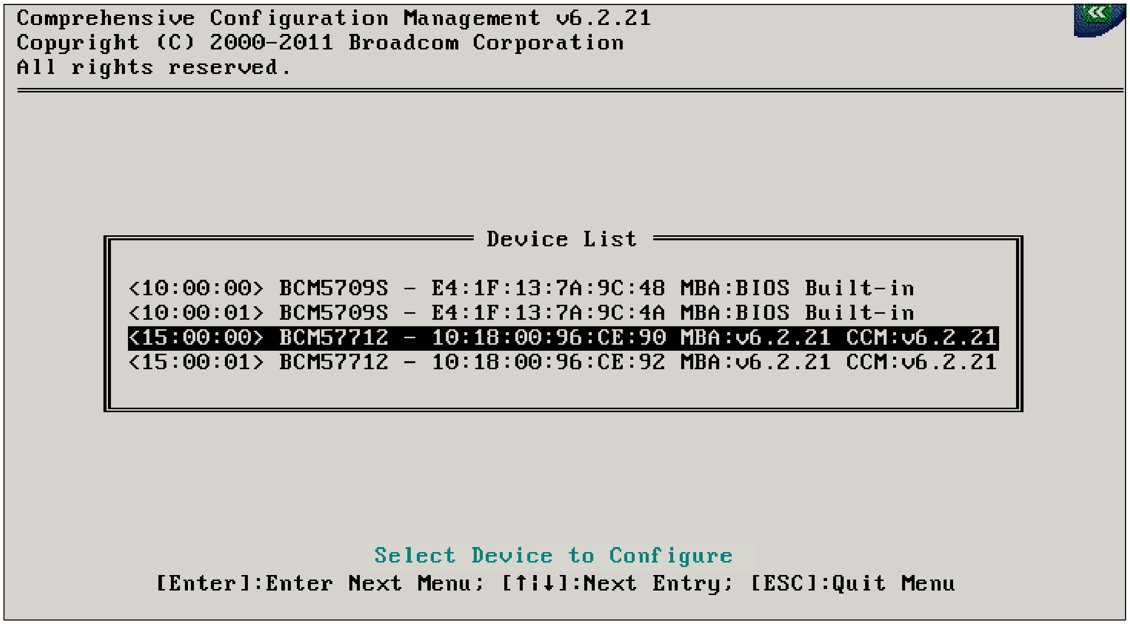

Figure 9-13 Device List panel

4. From the Main Menu panel (Figure 9-14), select **NIC Partition Configuration**, and then press Enter.

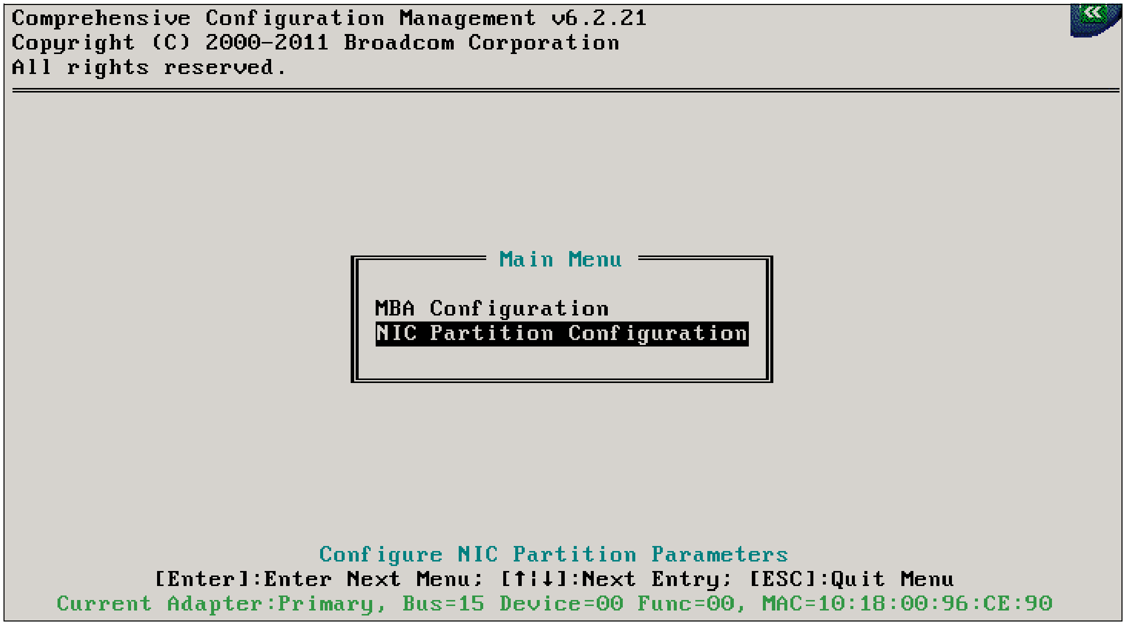

Figure 9-14 Device Main Menu panel

5. In the NIC Partition Configuration panel (Figure 9-15), highlight the NIC Partition field, and press the right arrow to change the value to **Enabled**.

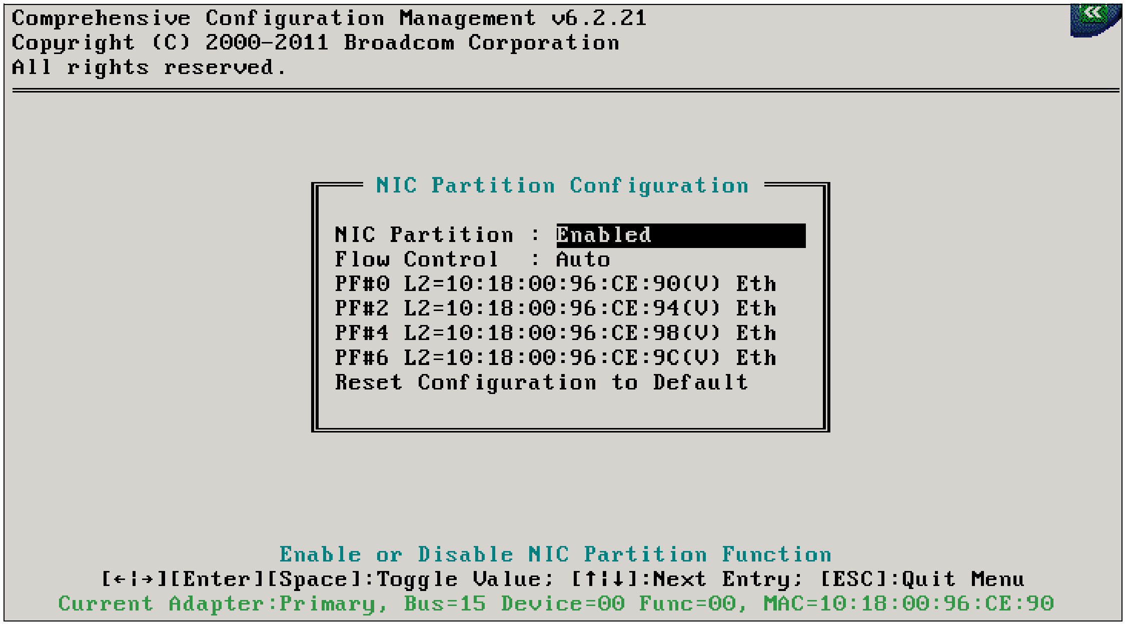

Figure 9-15 NIC Partition Configuration panel

The four NIC partitions are displayed, with their MAC addresses:

- PF#0, MAC address 10:18:00:96:CE:90
- PF#2, MAC address 10:18:00:96:CE:94
- PF#4, MAC address 10:18:00:96:CE:98
- PF#6, MAC address 10:18:00:96:CE:9C

6. Highlight one of the partitions that you want to configure. In this example, we selected **PF#0**. Then press Enter.

7. In the PF# 0 panel (Figure 9-16), configure the bandwidth parameters:
   a. Set the Bandwidth Weight to specify the guaranteed bandwidth in an overload situation.

      For example, if the parameter is set to 20, the NIC partition is guaranteed to transmit 2 Gbps of bandwidth (20% of total bandwidth). The sum of all four NIC partitions must be 0 or 100. If the sum is 0, the relative bandwidth weight mechanism is disabled, and a round-robin bandwidth allocation mechanism is used instead.
   b. Set the Maximum Bandwidth as percentage of total physical port bandwidth that the NIC partition can transmit.

      You can set this parameter to 100 on all four partitions. In this case, a partition can use all 10 Gbps of bandwidth when no other partition is communicating. If this parameter is set to a value lower than 100, this value indicates the maximum bandwidth available to the NIC partition.
   c. After you set the bandwidth values for NIC partition PF#0, press Esc to exit the panel.

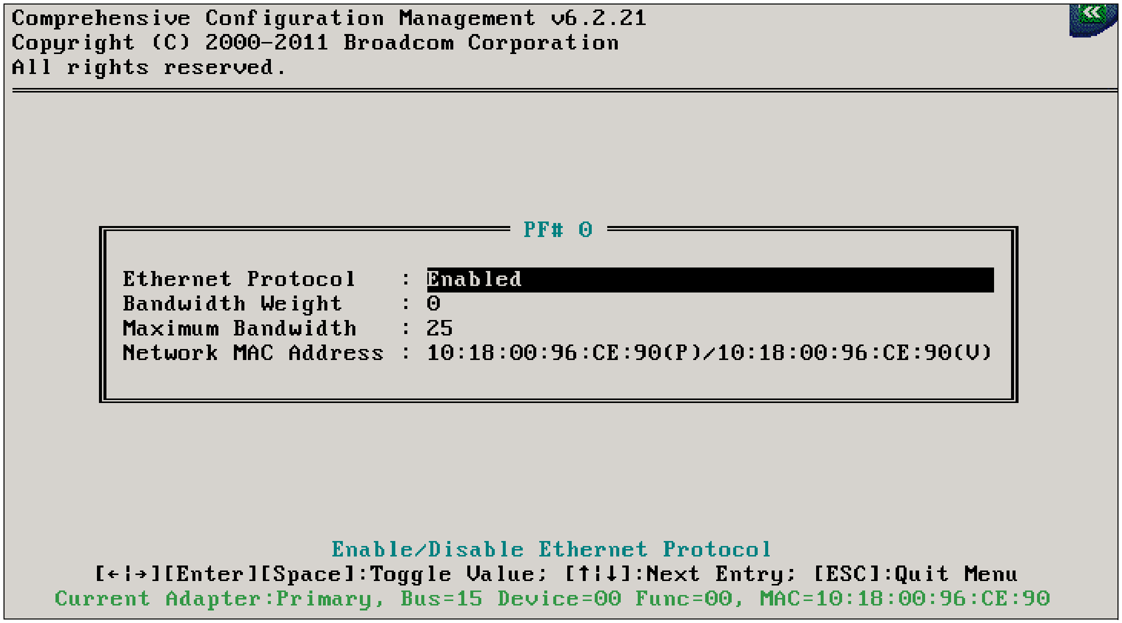

Figure 9-16 Partition function menu

8. Set the bandwidth parameters for the other three NIC partitions (PF#2, PF#4 and PF#6) by repeating steps 6 on page 212 and 7.

9.5 Configuring vNIC VLANs

Switch Independent Mode extends the existing customer VLANs to the NIC interfaces. Switch Independent Mode vNICs do not require an outer tag to maintain channel separation. Instead, they use the MAC address and the *IEEE 802.1Q VLAN tag* to distinguish the individual vNIC channels on each link. With this approach, you must coordinate the configuration of the Virtual Fabric Adapters and the chassis switches to match the VLAN structure in the network.

Conventional VLANs are configured by an operating system on a server blade (such as with the **vconfig** command in Linux) and are implemented as IEEE 802.1q VLAN tags. The VLAN configuration on upstream switches must match the VLAN configuration on the blades (Figure 9-17).

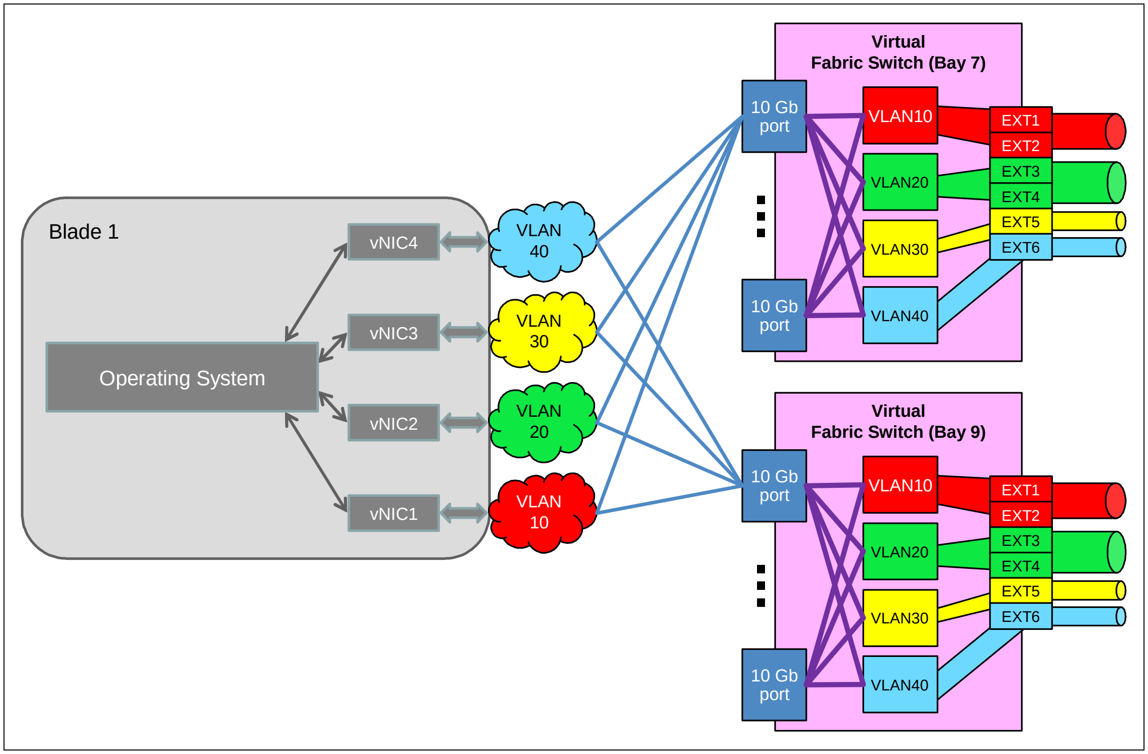

Figure 9-17 Diagram of Switch Independent Mode

You also configure the switch ports by using switch configuration methods, such as GUI, CLI, or Simple Network Management Protocol (SNMP). Presuming that multiple internal switch ports share common VLANs by default, you create distinct VLAN port groups by using conventional port VLAN schemes. Therefore, most ports are mapped to multiple VLANs, with the ports configured as VLAN trunks. The VLANs can then be mapped to physically separate external ports, or two or more VLANs can share an uplink port (for example, a VLAN trunk).

To configure the switch ports:

1. Assign internal ports to each VLAN based on vNIC-to-VLAN mapping. A single port is mapped to multiple VLANs.
2. Assign corresponding external ports to their corresponding VLAN. A VLAN can be mapped to a single port or a port group. Optionally, multiple VLANs can be mapped to a single port or a port group.

The switch normally provides basic Layer 2 forwarding among ports that share a VLAN. Only the BNT Virtual Fabric 10Gb Switch Module can handle L3 forwarding between separate VLANs within the chassis.

9.6 vNIC and VMready

You can use VMready only with the BNT Virtual Fabric 10Gb Switch Module. For more information, see 5.8, “vNIC and VMready” on page 101.

9.7 Understanding vNIC representation in an operating system

This section explains how vNICs are represented within an operating system.

9.7.1 vNICs on the Emulex 10GbE Virtual Fabric Adapter II

As expected, eight vNICs are listed in the Network Connections window (Figure 9-18), with four on each 10 Gbps physical port.

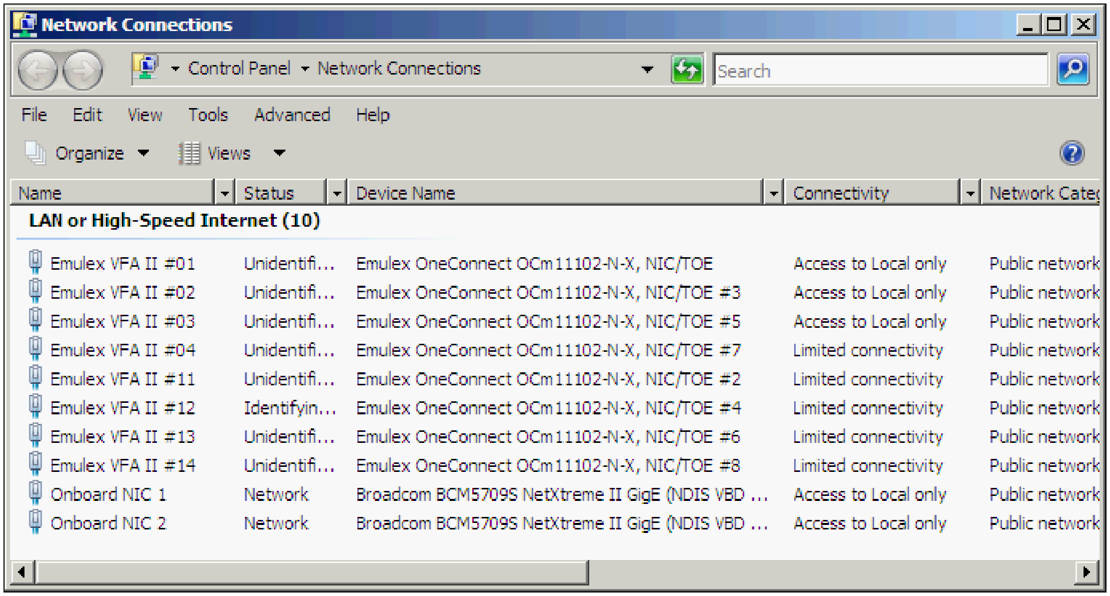

Figure 9-18 Network Connections window

For clarity, the names of the NICs were changed from the default values to the names listed in Table 9-1. You can use the MAC addresses to match each vNIC with the Virtual Fabric Adapter that is shown in the Network Connections window.

Table 9-1 NIC names and matching vNICs

NIC name	10Gb physical port	vNIC	MAC address
Emulex VFA II #01	0	1	00:00:C9:B2:57:02
Emulex VFA II #02	0	2	00:00:C9:B2:57:03
Emulex VFA II #03	0	3	00:00:C9:B2:57:04

NIC name	10Gb physical port	vNIC	MAC address
Emulex VFA II #04	0	4	00:00:C9:B2:57:05
Emulex VFA II #11	1	1	00:00:C9:B2:57:06
Emulex VFA II #12	1	2	00:00:C9:B2:57:07
Emulex VFA II #13	1	3	00:00:C9:B2:57:08
Emulex VFA II #14	1	4	00:00:C9:B2:57:09

To verify parameters, such as the MAC address and link speed, right-click a NIC, and select **Status**. Figure 9-19 shows the status panel for the Emulex VFA II #01.

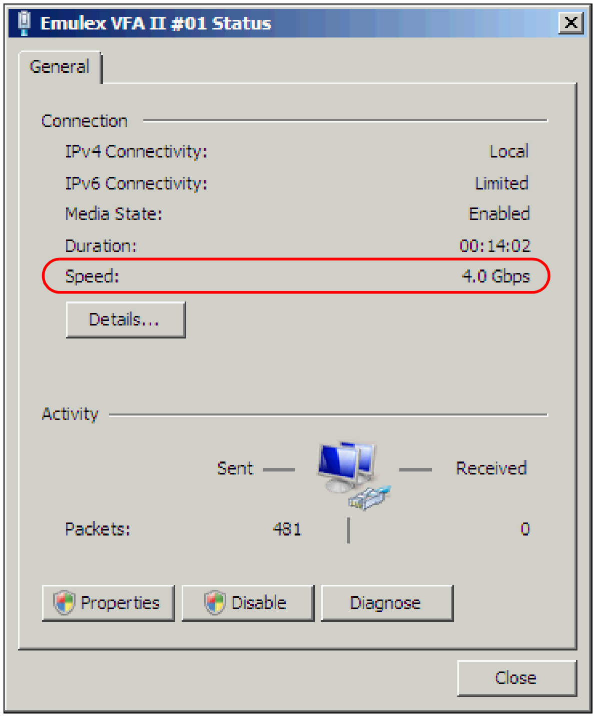

Figure 9-19 Emulex VFA II #01 Status window

Tip: Speed is reported as 4.0 Gbps, meaning that the bandwidth was set to 40 when configuring this vNIC in the UEFI utility.

To see the MAC address of this vNIC, click **Details**. Figure 9-20 shows the Network Connection Details window, which displays the MAC address in the Physical Address field. This MAC address represents vNIC 1 on 10 Gbps physical port 0.

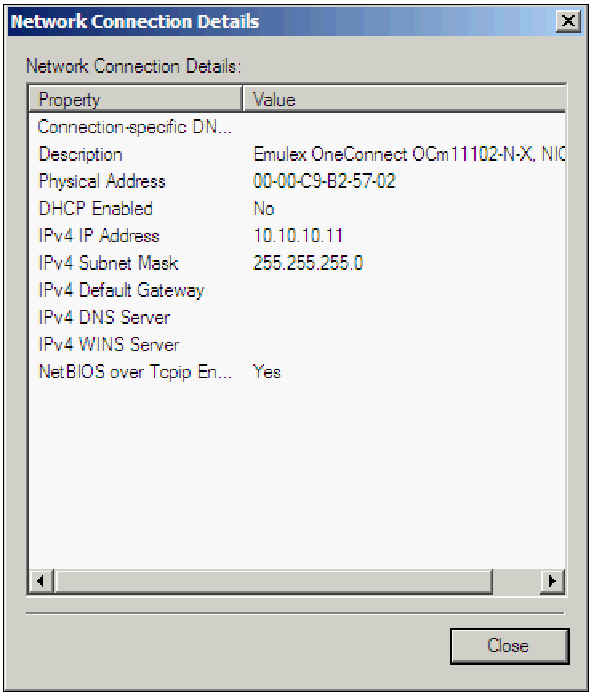

Figure 9-20 Network Connection Details window

You can identify and verify MAC addresses and link speeds of the other vNICs in the same manner.

9.7.2 vNICs on the Broadcom 2-port 10Gb Virtual Fabric Adapter

The NIC partitions on the Broadcom 2-port 10Gb Virtual Fabric Adapter are displayed to the operating system (Windows Server 2008). This section begins by showing the built-in networking tools for Windows and then by showing the Broadcom Advanced Control Suite.

Windows networking tools

Figure 9-21 shows how the NIC partitions are displayed in the Network and Sharing Center.

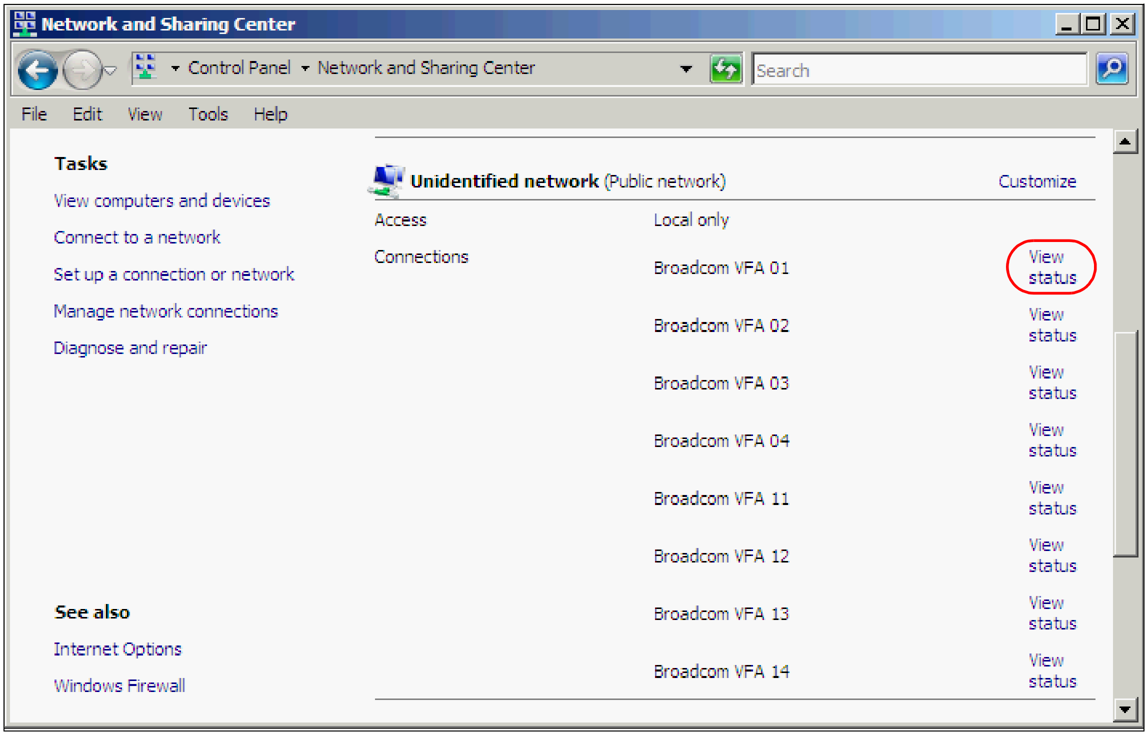

Figure 9-21 Broadcom 2-port 10Gb Virtual Fabric Adapter NIC partitions in Windows Server 2008

For clarity, we changed the NIC names from the defaults as shown in Table 9-2. We used the MAC addresses to match each NIC with the actual NIC partition.

Table 9-2 NIC names and matching NIC partitions

NIC name	10 Gbps physical port	NIC partition	MAC address
Broadcom VFA 01	0	1	10:18:00:96:E3:80
Broadcom VFA 02	0	2	10:18:00:96:E3:84
Broadcom VFA 03	0	3	10:18:00:96:E3:88
Broadcom VFA 04	0	4	10:18:00:96:E3:8C
Broadcom VFA 11	1	1	10:18:00:96:E3:82
Broadcom VFA 12	1	2	10:18:00:96:E3:86

NIC name	10 Gbps physical port	NIC partition	MAC address
Broadcom VFA 13	1	3	10:18:00:96:E3:8A
Broadcom VFA 14	1	4	10:18:00:96:E3:8E

By clicking **View Status** (highlighted in Table 9-21 on page 219), you can verify parameters for each of the NIC partitions, such as link speed and MAC address. Figure 9-22 shows the parameters for NIC Broadcom VFA 01.

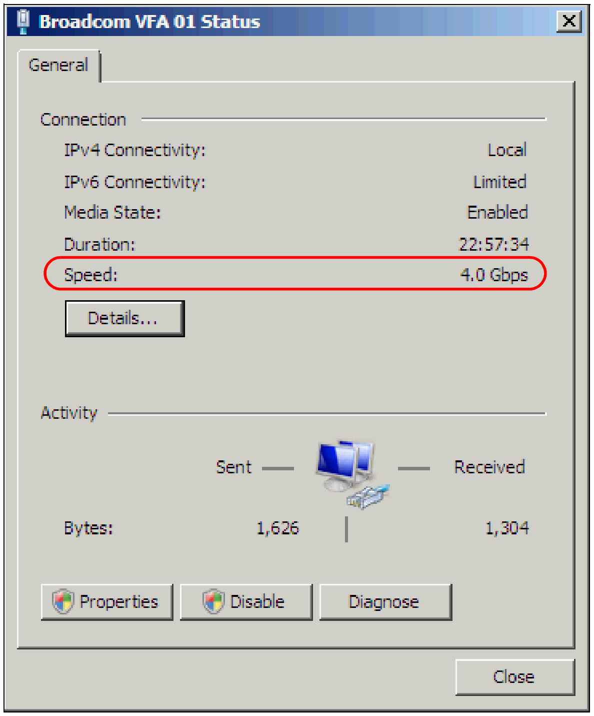

Figure 9-22 Broadcom VFA 01 Status window

In the example shown in Figure 9-22, speed is reported as 4.0 Gbps. This speed indicates that the bandwidth was set to 40 on this NIC partition when configuring parameters in the UEFI utility.

You can click **Details** to see the MAC address of this NIC partition.

The Network Connection Details window (Figure 9-23) shows the MAC address. This MAC address is for NIC partition 1 on 10 Gbps physical port 0.

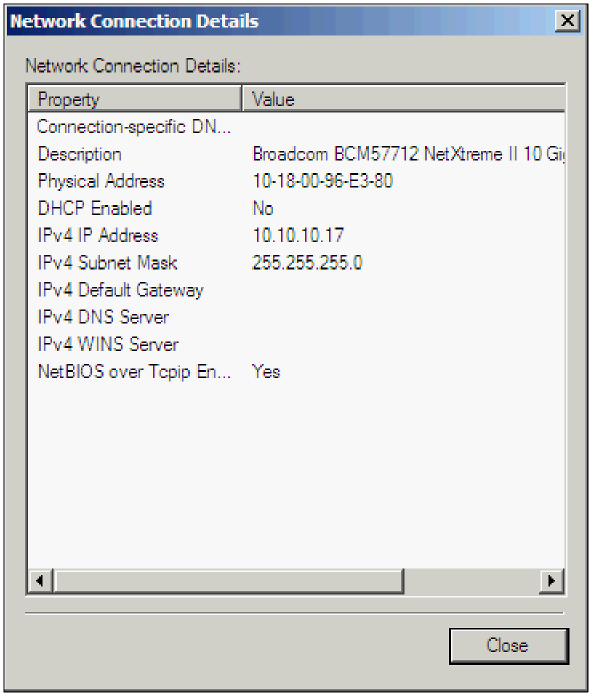

Figure 9-23 Network Connection Details window with the MAC address

You use the same method to verify the MAC addresses and speed for the other NIC partitions on the Broadcom 2-port 10Gb Virtual Fabric Adapter.

Broadcom Advanced Control Suite

Broadcom Advanced Control Suite is a management utility that is based on Windows. You use it to administer, configure, and diagnose Broadcom Virtual Fabric Adapters and other Broadcom NICs.

Figure 9-24 shows an example from Broadcom Advanced Control Suite. The left pane shows eight NIC partitions. NIC partition 1 on physical port 0 highlighted. The speed is 4 Gbps, which indicates a bandwidth parameter value of 40 in the UEFI utility.

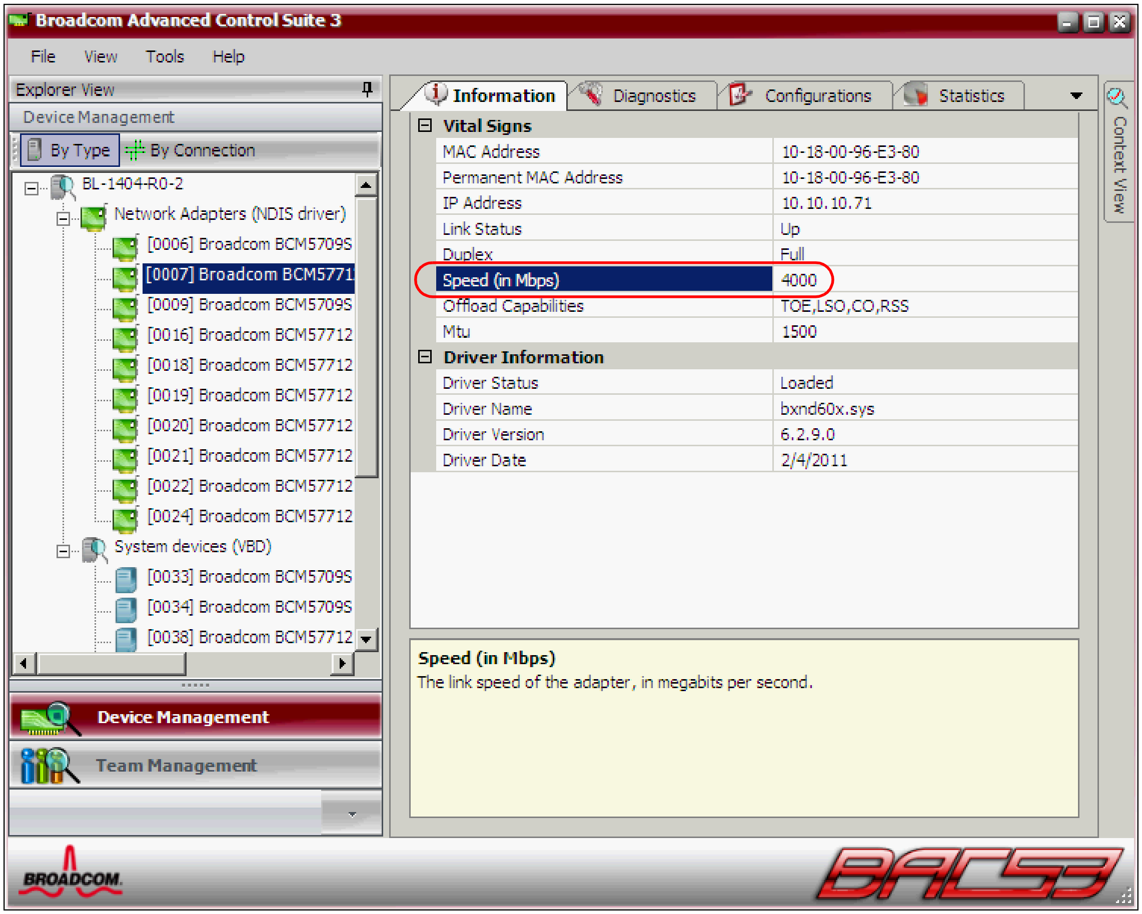

Figure 9-24 NIC partition 1 in Broadcom Advanced Control Suite

By using the Broadcom Advanced Control Suite utility, you can test the configuration by running a basic connectivity test. You click the **Diagnostics** tab (Figure 9-25), and then ping to an IP gateway or a user-specified IP address.

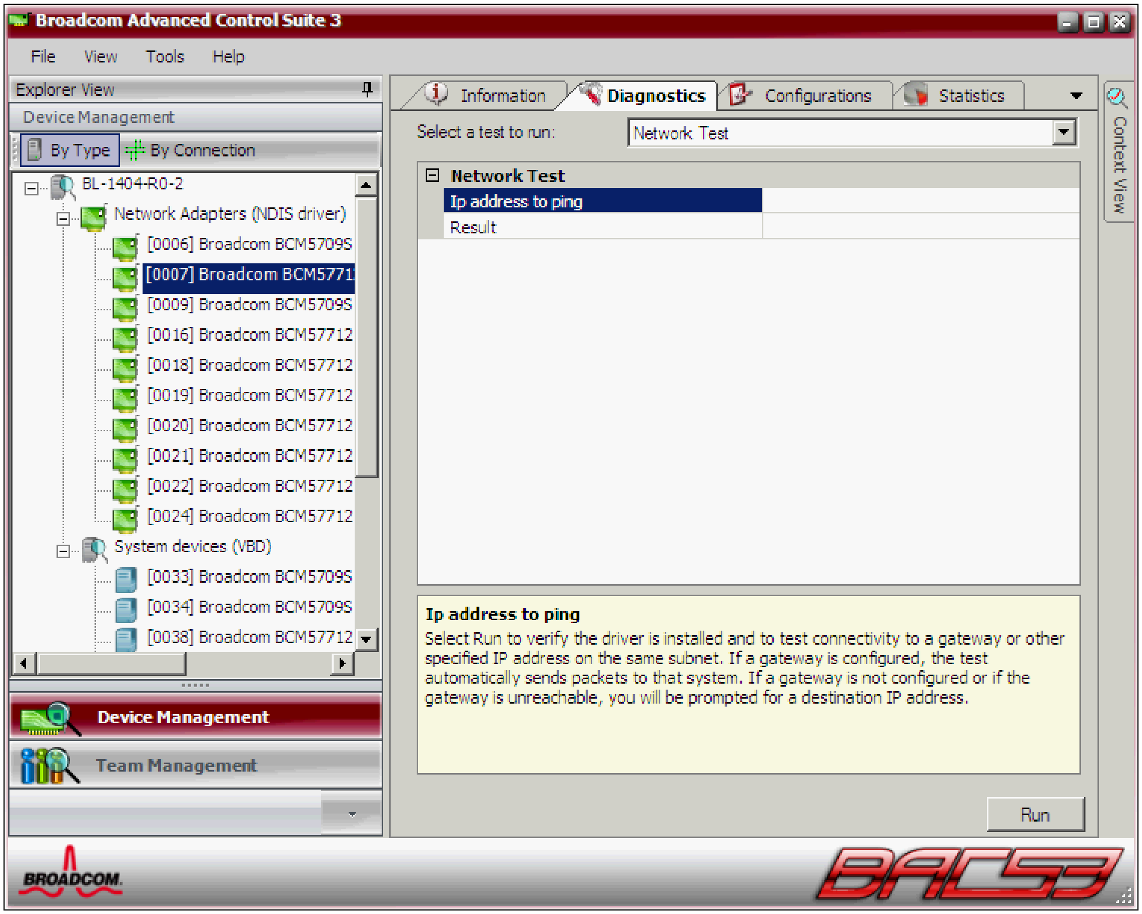

Figure 9-25 Diagnostics tab in the Broadcom Advanced Control Suite

The **Statistics** tab (Figure 9-26) shows the network traffic counters. These counters are useful when monitoring or troubleshooting NIC behavior.

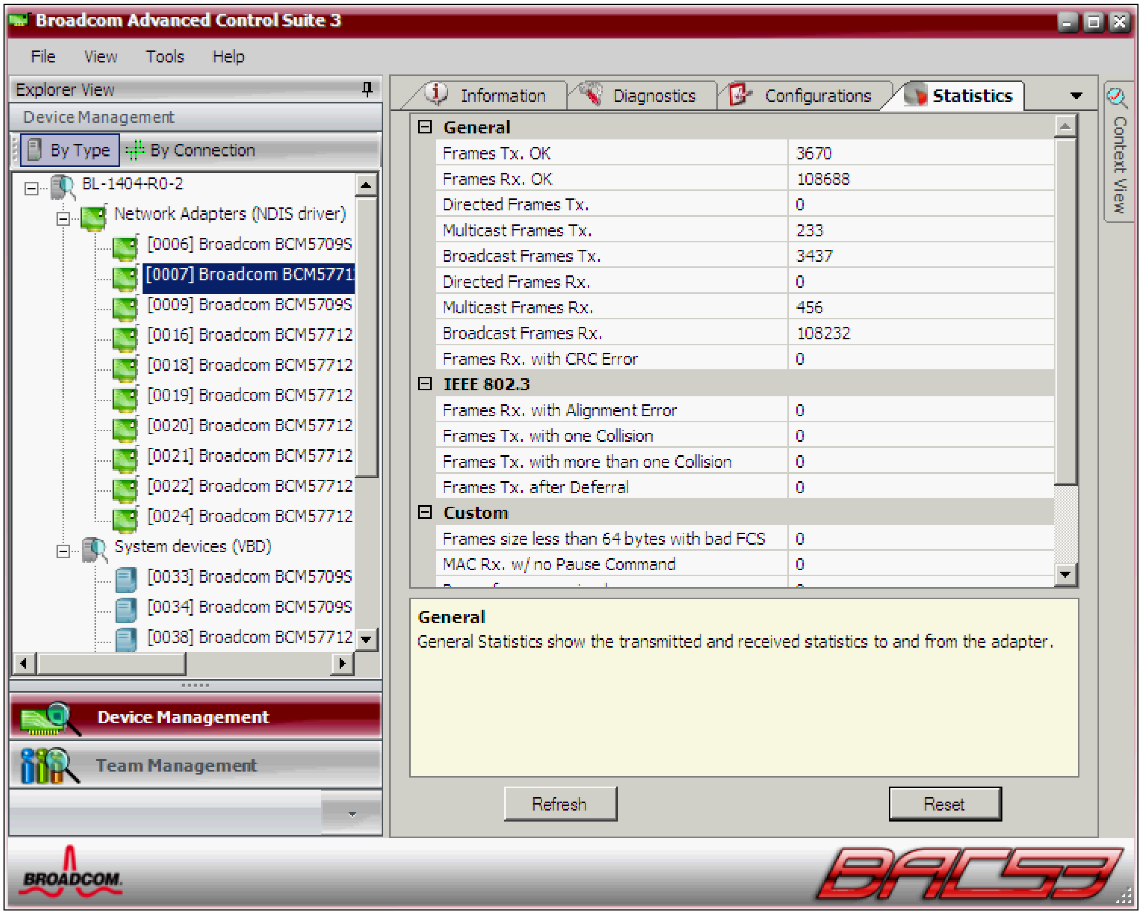

Figure 9-26 Statistics tab in the Broadcom Advanced Control Suite

One of the functions of the Broadcom Advanced Control Suite utility configures NIC teaming. You can use NIC partitions in teaming configurations the same way you might use physical NICs. To set up NIC teaming, click **Team Management** in the lower-left corner (Figure 9-26 on page 224).

For more information about NIC teaming by using the Broadcom Advanced Control Suite, see 9.7.4, “Configuring teams with Broadcom Advanced Control Suite” on page 229.

9.7.3 Configuring teams with Emulex OneCommand Manager

To configure NIC teaming with the Emulex 10GbE Virtual Fabric Adapter II, you can use the Emulex OneCommand Manager utility, which runs in Windows and Linux. It also supports VMware ESX/ESXi and Microsoft Hyper-V hypervisors. For more information, see the OneCommand Manager home page at:

http://www.emulex.com/products/management-software/device-management/onecommand-manager/

To create a team with two members, Emulex VFA II #01 and Emulex VFA II #02:

1. In the OC NIC Teaming and VLAN Manager window (Figure 9-27), click **Create Team**.

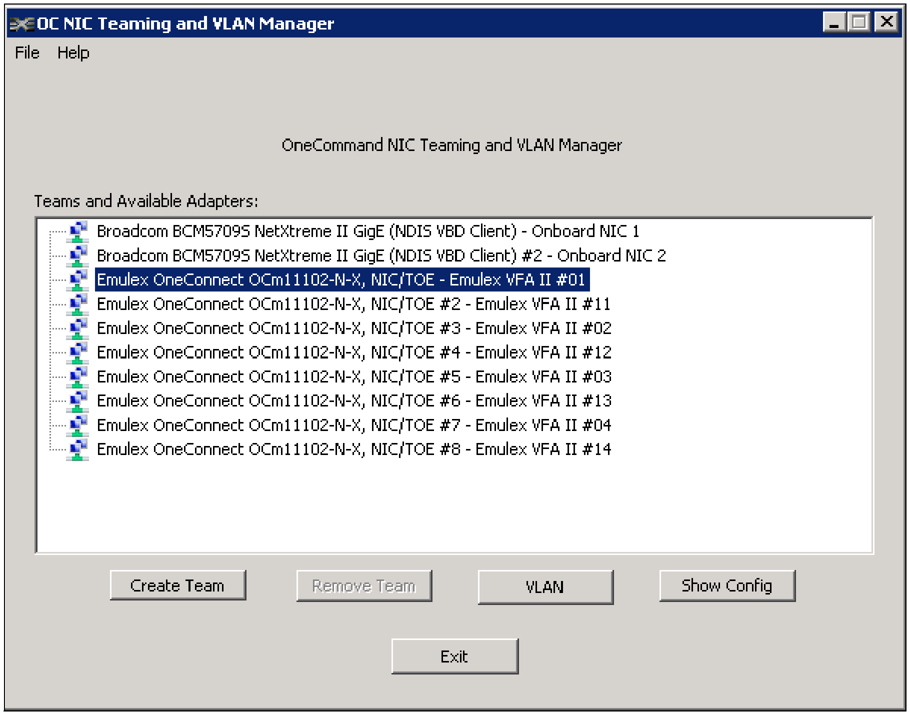

Figure 9-27 Emulex OneCommand Manager utility

2. In the Create team window (Figure 9-28), specify the team parameters:

   a. Enter a team name.

   b. Select a team type. In this example, we use the simplest form of teaming, which is *FailOver*. Therefore, select **FailOver**. This setting means that the primary member will carry all the traffic, and the secondary member will have a standby role. Other team type options are available:

      - LoadBalancing
      - 802.3ad static
      - LACP

   c. From the Available Network Adapters list, select each adapter you want to form into the team. Click **Add** to move it to the Team Member Adapters list.

   d. When done you are done setting parameters and selecting the team members, click **OK**.

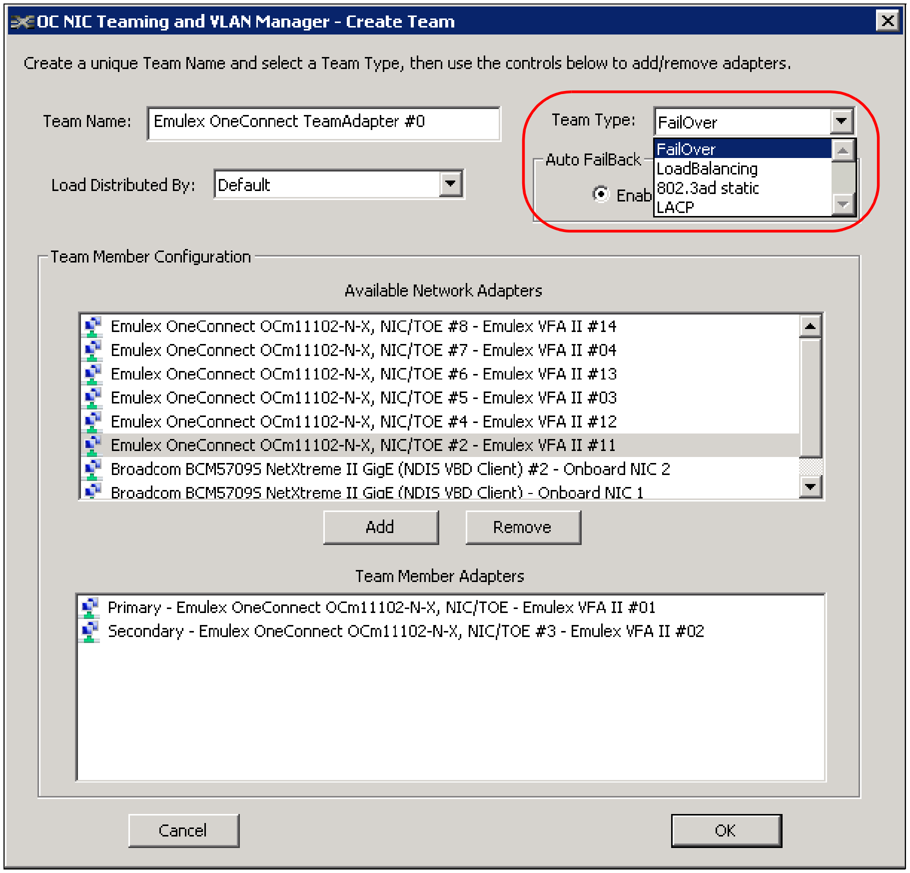

Figure 9-28 Create team window in Emulex OneCommand Manager

3. Highlight the new team and click **VLAN** to configure the VLAN ID for the team. Figure 9-29 shows the newly created team.

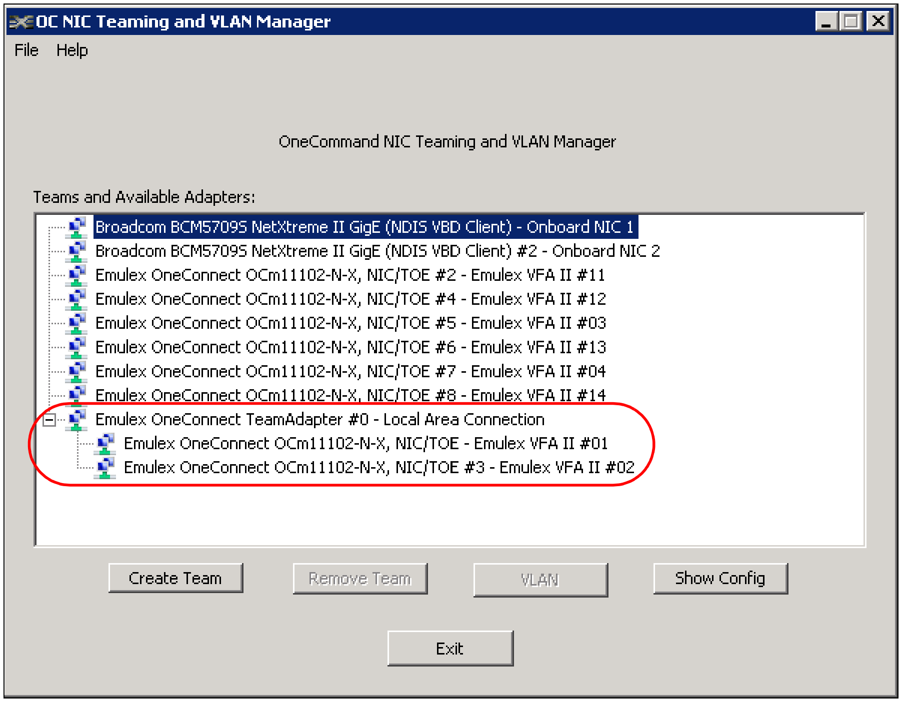

Figure 9-29 New team created in Emulex OneCommand Manager

4. In the Add/Remove VLAN window (Figure 9-30), enter the VLAN ID and click **Add**.

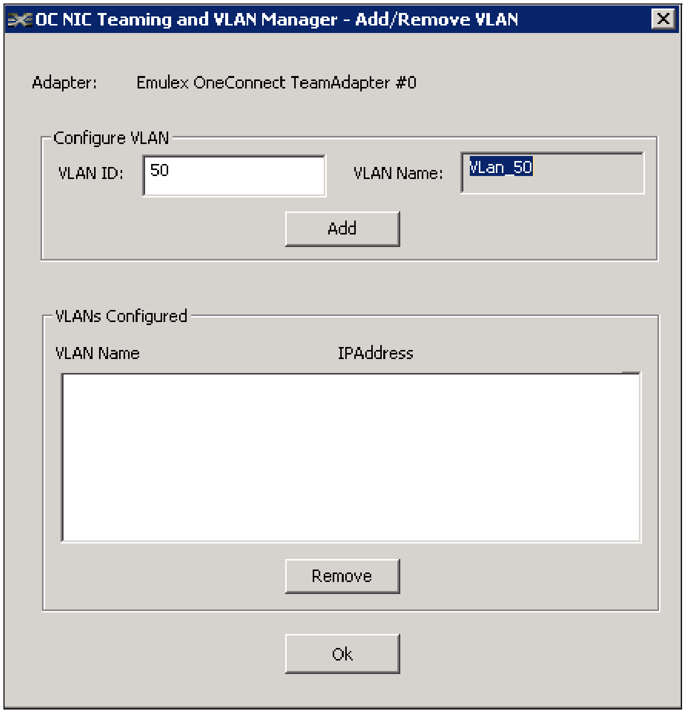

Figure 9-30 Adding the VLAN in Emulex OneCommand Manager

The VLAN is now configured as shown in Figure 9-31.

5. Click **OK** to return to the main window.

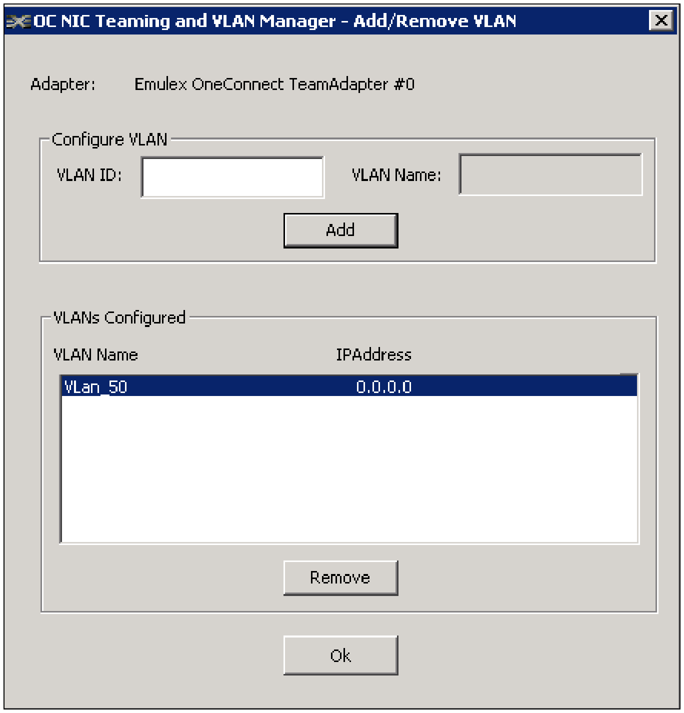

Figure 9-31 Emulex OneCommand Manager showing the VLAN configured

9.7.4 Configuring teams with Broadcom Advanced Control Suite

A *team* (or *bond* in Linux terminology) of adapters functions as a single virtual network interface and looks the same as a nonteamed adapter to other network devices. By using Broadcom Advanced Control Suite, you can team up to eight network adapters in a single group. This group becomes a virtual network adapter with its own IP address.

Teaming has the following benefits:

- Virtual adapters can balance a large network load across several adapters.
- The use of teams improves system availability by providing a failover path for critical network connections.

The following types of teams are supported by Broadcom Advanced Control Suite:

- Switch-dependent
  - Link Aggregation Control Protocol (LACP) or IEEE 802.3ad
  - Gigabit EtherChannel (GEC) or Fast EtherChannel (FEC)

  Switch-dependent NIC teaming is not related to IBM Virtual Fabric Mode. This teaming type requires that LACP or EtherChannel is configured on an Ethernet switch module. Modules must also be stacked to act as a single switch.

- Switch-independent
  - Smart Load Balancing (SLB)
  - SLB with Auto Failover Disable

  Switch-independent NIC teaming is not related to Switch Independent Mode. The name indicates that switch-independent NIC teaming does not require any LACP or EtherChannel configuration on an Ethernet switch module.

To configure a new team in Windows by using Broadcom Advanced Control Suite application:

1. Start the Broadcom Advanced Control Suite application.
2. In the main Broadcom Advanced Control Suite window (Figure 9-24 on page 222), in the lower-left corner, click **Team Management** to start the Broadcom Teaming Wizard.

3. In the Welcome window (Figure 9-32), click **Next**.

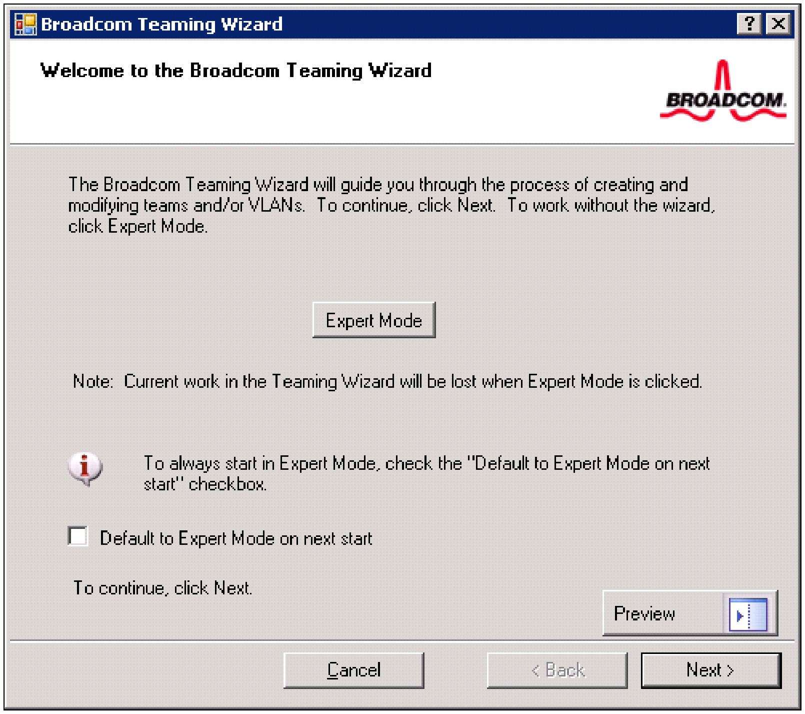

Figure 9-32 Welcome window in the Broadcom Teaming Wizard

4. In the Creating/Modifying a Team: Team Name window (Figure 9-33), enter the name for the new team. Then click **Next**.

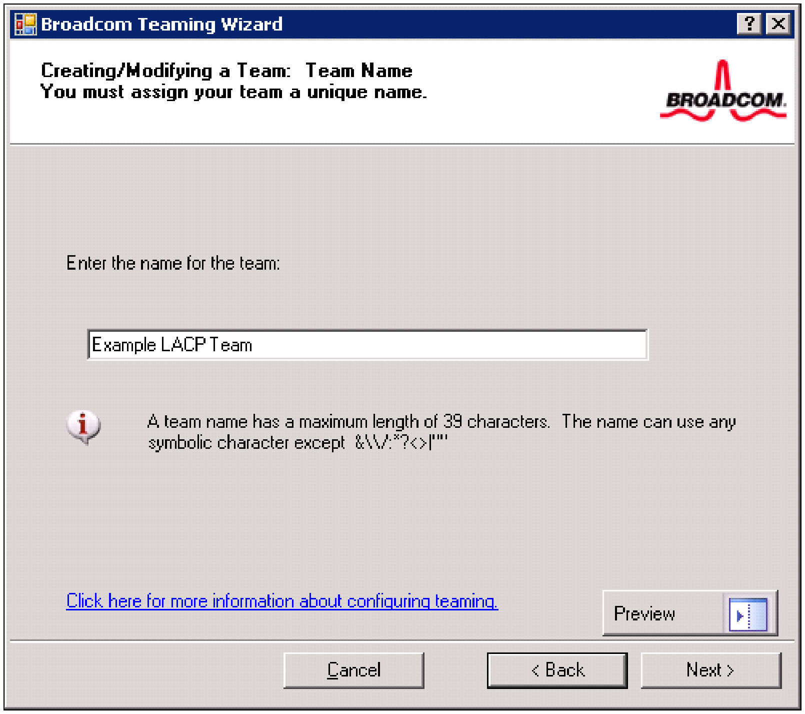

Figure 9-33 Creating/Modifying a Team: Team Name window

5. In the Team Type window (Figure 9-34), select the type of team that you want to create, and then click **Next**.

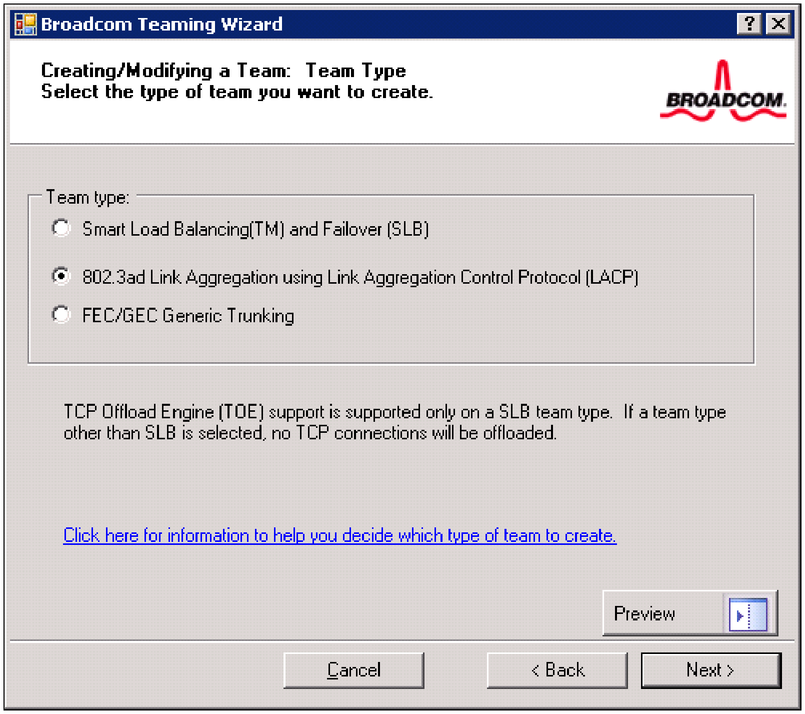

Figure 9-34 Creating/Modifying a Team: Team Type window

6. When you see the verification message window (Figure 9-35), verify and confirm that the switch is configured correctly. This message applies only to teaming modes that require LACP or EtherChannel configuration on the switch modules. Then click **OK**.

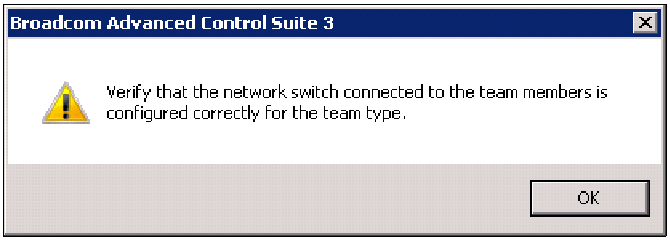

Figure 9-35 Verification message window

7. In the Creating/Modifying a Team window (Figure 9-36), assign team members:

   a. Highlight the adapter you want to add to the team.
   b. Click **Add**.
   c. Repeat steps a and b to add the remaining team members.
   d. Click **Next**.

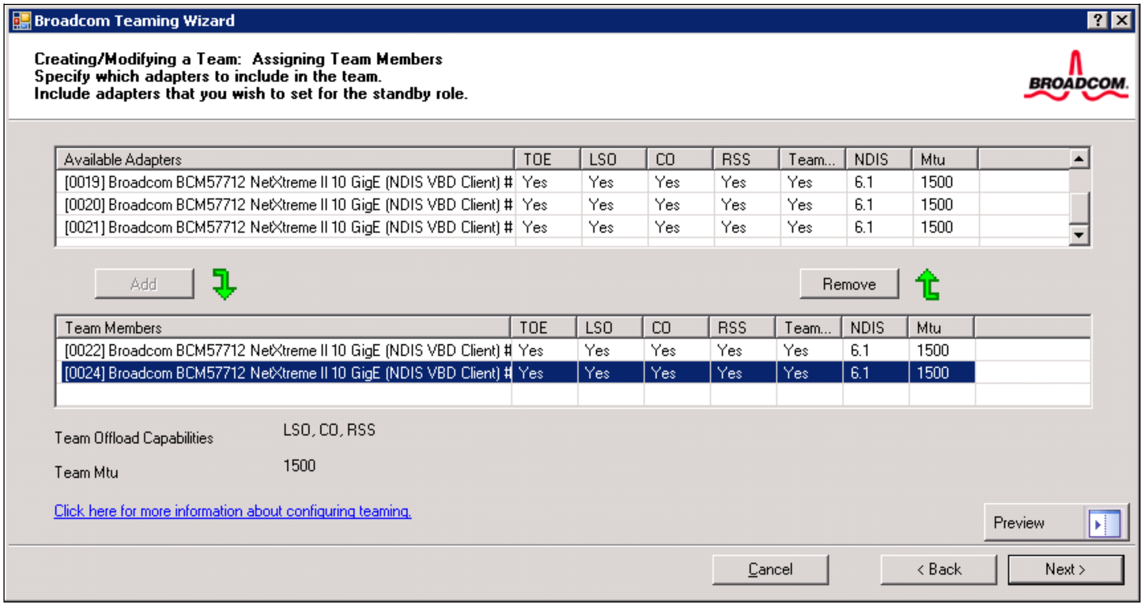

Figure 9-36 Creating/Modifying a Team: Assigning Team Members window

8. In the Commit Changes window (Figure 9-37), if needed, select **Save changes and continue to manage more teams**. In this example, we select **Commit changes to system and Exit the wizard**. Then click **Finish**.

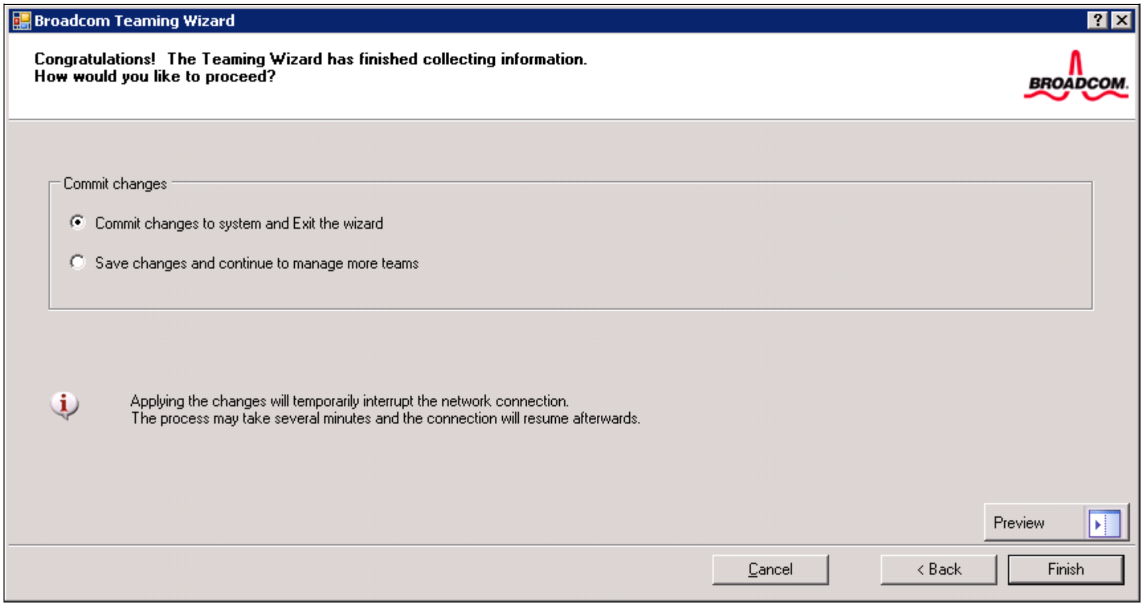

Figure 9-37 Commit changes window

The Team Management window confirms the newly added team (Figure 9-38).

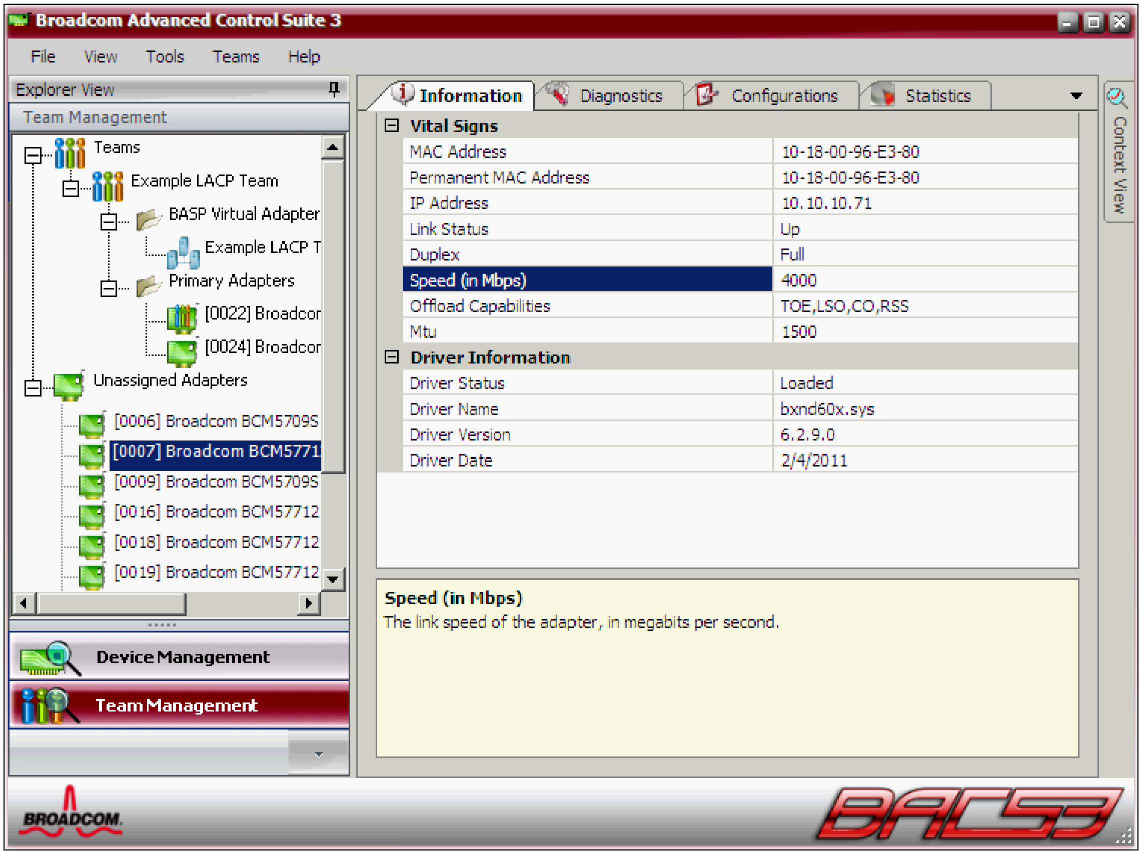

Figure 9-38 Teaming established

You now completed the configuration of NIC teaming with the Broadcom Advanced Control Suite application.

10

Usage scenarios

This chapter shows examples of Switch Independent Mode configurations by using the following combinations:

- Cisco Nexus 4001I Switch Module with the following adapters:
  - Emulex 10GbE Virtual Fabric Adapter II
  - Broadcom 2-port 10Gb Virtual Fabric Adapter
- BNT Virtual Fabric 10Gb Switch Module with the following adapters
  - Emulex 10GbE Virtual Fabric Adapter II
  - Broadcom 2-port 10Gb Virtual Fabric Adapter
- Brocade Converged Switch with
  - Emulex 10GbE Virtual Fabric Adapter II

This chapter includes the following sections:

- Overview of the configurations
- Scenario 1: Cisco Nexus 4001I Switch Module configurations
- Scenario 2: BNT Virtual Fabric 10Gb Switch Module configurations
- Scenario 3: Brocade Converged Switch configuration

10.1 Overview of the configurations

We use the following equipment to demonstrate the four configurations:

- Two BladeCenter H chassis (BC4 and BC5)
- Two Cisco Nexus 4001I Switch Modules in chassis BC4
- Two BNT Virtual Fabric 10Gb Switch Modules in chassis BC5
- Two HS22 blades in each chassis
- Two Emulex 10GbE Virtual Fabric Adapter II cards
- Two Broadcom 2-port 10Gb Virtual Fabric Adapter cards

Chassis BC4 is used to demonstrate the following configurations, shown in Figure 10-1:

- Emulex 10GbE Virtual Fabric Adapter II and Cisco Nexus 4001I Switch Module
- Broadcom 2-port 10Gb Virtual Fabric Adapter and Cisco Nexus 4001I Switch Module

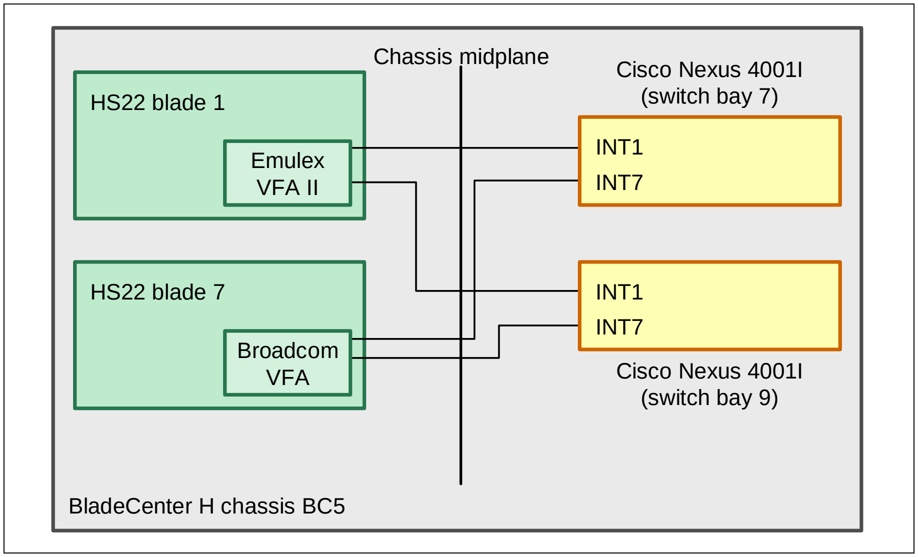

Figure 10-1 BladeCenter H chassis BC5

Chassis BC5 is used to demonstrate the following configurations, shown in Figure 10-2:

- Emulex 10GbE Virtual Fabric Adapter II and BNT Virtual Fabric 10Gb Switch Module
- Broadcom 2-port 10Gb Virtual Fabric Adapter and BNT Virtual Fabric 10Gb Switch Module

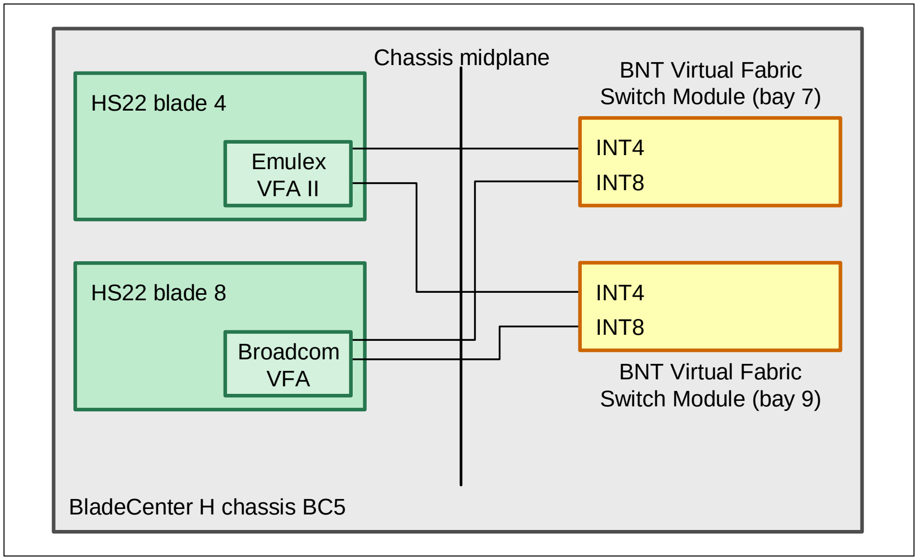

Figure 10-2 BladeCenter H chassis BC5

10.1.1 VLAN definitions

We enable the following virtual local area networks (VLANs) on Virtual Fabric Switch Modules:

- VLAN 10

  Both blades use this VLAN for communication through external port 2 on the Virtual Fabric Switch. Virtual network interface controller (vNIC) 1 on each blade server communicates on this VLAN.

- VLAN 20

  The blades use this VLAN to communicate through external port 4. vNIC 2 on each blade server communicates on this VLAN.

- VLAN 30

  This VLAN does not use any external port on the Virtual Fabric Switch Module. vNIC 3 on each blade server use this VLAN to communicate between themselves.

- VLAN 40

  vNIC 4 on blade server 1 in chassis BC4 communicates through external port 2 on the Virtual Fabric Switch Module. We use VLAN 40 for this purpose.

  Similarly, vNIC 4 on blade server 4 in chassis BC5 communicates through external port 2 on the BNT 10 Gbps Virtual Fabric Switch Module.

- VLAN 50

  vNIC 4 on blade server 7 in chassis BC4 communicates through external port 4 on the Virtual Fabric Switch Module. We use VLAN 50 for this purpose.

  Similarly, vNIC 4 on blade server 8 in chassis BC5 communicates through external port 4 on the BNT 10 Gbps Virtual Fabric Switch Module.

Table 10-1 shows the VLANs on Cisco Nexus 4001I Switch Module in chassis BC4.

Table 10-1 VLANs on Cisco Nexus 4001I Switch Module

VLAN	Internal ports	External ports
10	1, 7	2
20	1, 7	4
30	1, 7	
40	1	2
50	7	4

Table 10-2 shows the VLANs on BNT 10 Gbps Virtual Fabric Switch Module in chassis BC5.

Table 10-2 VLANs on BNT Virtual Fabric 10Gb Switch Module

VLAN	Internal ports	External ports
10	4, 8	2
20	4, 8	4
30	4, 8	
40	4	2
50	8	4

Because each internal (blade-facing) port on the switch modules needs to pass traffic in multiple VLANs, configure these ports in trunk mode. You can configure external ports in trunk mode when necessary (when traffic from multiple VLANs passes through a particular external port).

10.1.2 Switch and adapter vNIC correlation

In IBM Virtual Fabric Mode (switch-dependent mode), vNICs on the Emulex Virtual Fabric Adapter have a direct correlation with vNIC definitions on the BNT Virtual Fabric 10Gb Switch Module. This correlation results because vNICs are configured on the switch module, and the settings are transferred to the adapter by using DCBX.

In Switch Independent Mode, this correlation is different. The vNICs are configured on the adapter itself, for example, by using the Unified Extensible Firmware Interface (UEFI) utility. No vNIC settings need to be configured on the switch module. However, internal ports must be configured as trunk ports. Also one or more VLANs must be associated with each vNIC channel, allowing for multiple VLANs on each port to ensure that the traffic flows to the correct interfaces. The intent is to extend specific VLANs to specific vNIC interfaces within each server blade.

10.2 Scenario 1: Cisco Nexus 4001I Switch Module configurations

This section shows examples that use the Cisco Nexus 4001I Switch Module. In these examples, the modules are installed in chassis BC4. The blade in slot 1 has Emulex 10GbE Virtual Fabric Adapter II, and the blade in slot 7 has Broadcom 2-port 10Gb Virtual Fabric Adapter. The scenario begins with the Emulex configuration on blade 1, then shows the Broadcom configuration, and concludes with setting up the switch module.

This section includes the following topics:

- vNICs summary
- Configuring the Emulex adapter
- Configuring the Broadcom adapter
- Configuring the Cisco Nexus 4001I Switch Module

10.2.1 vNICs summary

The HS22 blade server in slot 1 contains an Emulex 10GbE Virtual Fabric Adapter II. On this adapter, we enable vNICs and set the following VLAN IDs for them:

- VLAN 10 for vNIC 1
- VLAN 20 for vNIC 2
- VLAN 30 for vNIC 3
- VLAN 40 for vNIC 4

The HS22 blade server in slot 7 contains a Broadcom 2-port 10Gb Virtual Fabric Adapter. We enable vNICs on the adapter, but do not specify any VLAN IDs on the Broadcom 2-port 10Gb Virtual Fabric Adapter. However, we set the following VLANs for vNICs in the operating system:

- VLAN 10 for vNIC 1
- VLAN 20 for vNIC 2
- VLAN 30 for vNIC 3
- VLAN 50 for vNIC 4

10.2.2 Configuring the Emulex adapter

Table 10-3 shows the configuration that we implement on the Emulex 10GbE Virtual Fabric Adapter II in HS22 blade 1.

Table 10-3 vNIC configuration: Emulex 10GbE Virtual Fabric Adapter II

vNIC function	Bandwidth	VLAN ID	MAC address
1	4 Gbps	10	00:00:C9:B2:57:02
2	3 Gbps	20	00:00:C9:B2:57:03
3	2 Gbps	30	00:00:C9:B2:57:04
4	1 Gbps	40	00:00:C9:B2:57:05

We can identify the MAC addresses for each vNIC function by using the UEFI utility (see “vNICs in the operating system” on page 252). You must know the MAC addresses so that you can identify each vNIC function in the operating system.

To configure parameters on the Emulex adapter:

1. Power on or restart HS22 blade 1.
2. When prompted during POST, press F1.

3. In the System Configuration and Boot Management panel (Figure 10-3) of the UEFI utility, select **System Settings**.

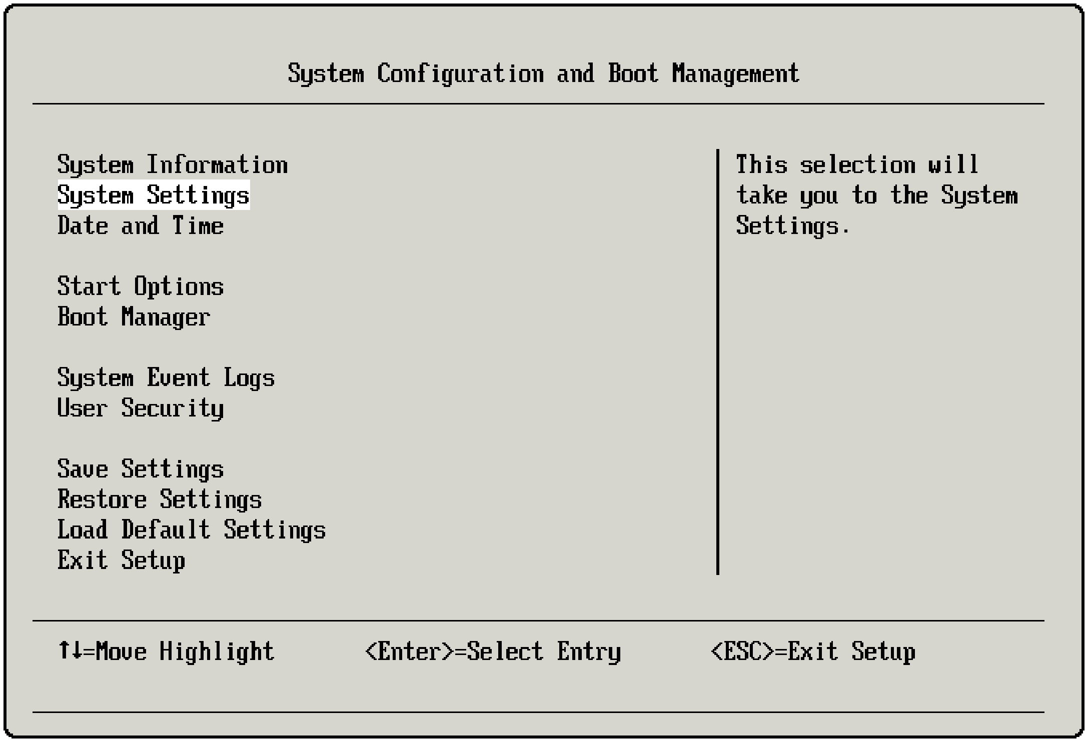

Figure 10-3 UEFI utility main menu

4. In the System Settings panel, select **Network**.

5. In the Network panel (Figure 10-4), which shows two Emulex 10 Gbps ports, select the first port on which to configure vNICs.

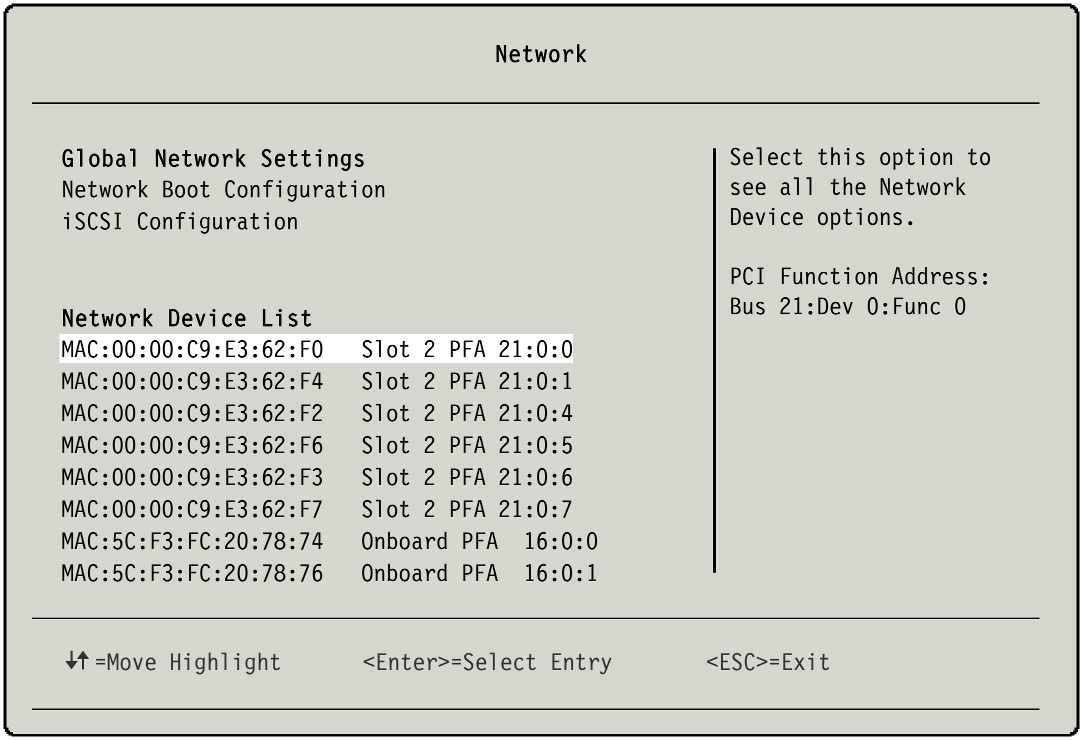

Figure 10-4 Network configuration panel in the blade UEFI utility

6. In the Emulex NIC Selection panel (Figure 10-5), verify the following settings:
   a. Set Multichannel to **Enabled**.
   b. In the Switch Configuration field, select **Switch Independent Mode**.
   c. Select **Multichannel Configuration** to configure VLAN and bandwidth settings for each of the four vNICs.

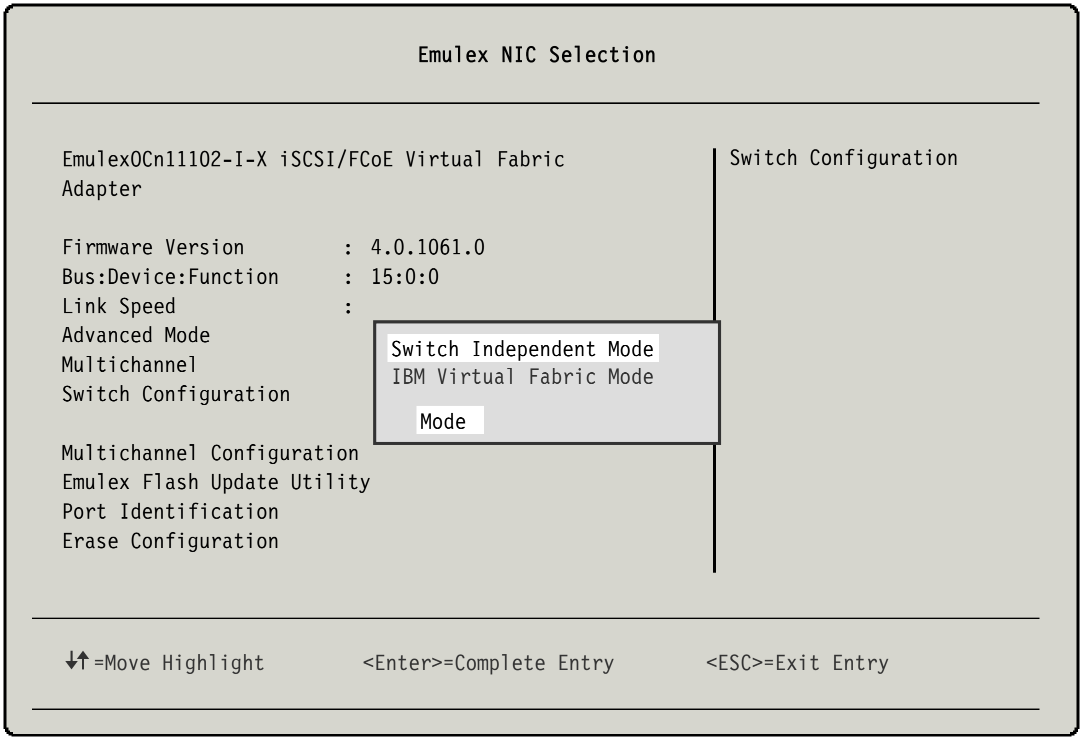

Figure 10-5 Emulex vNIC Switch Independent Mode / IBM Virtual Fabric Mode selection

7. In the Function Configuration panel (Figure 10-6), configure each of the four vNIC functions. To begin, select **Function 0** and press Enter.

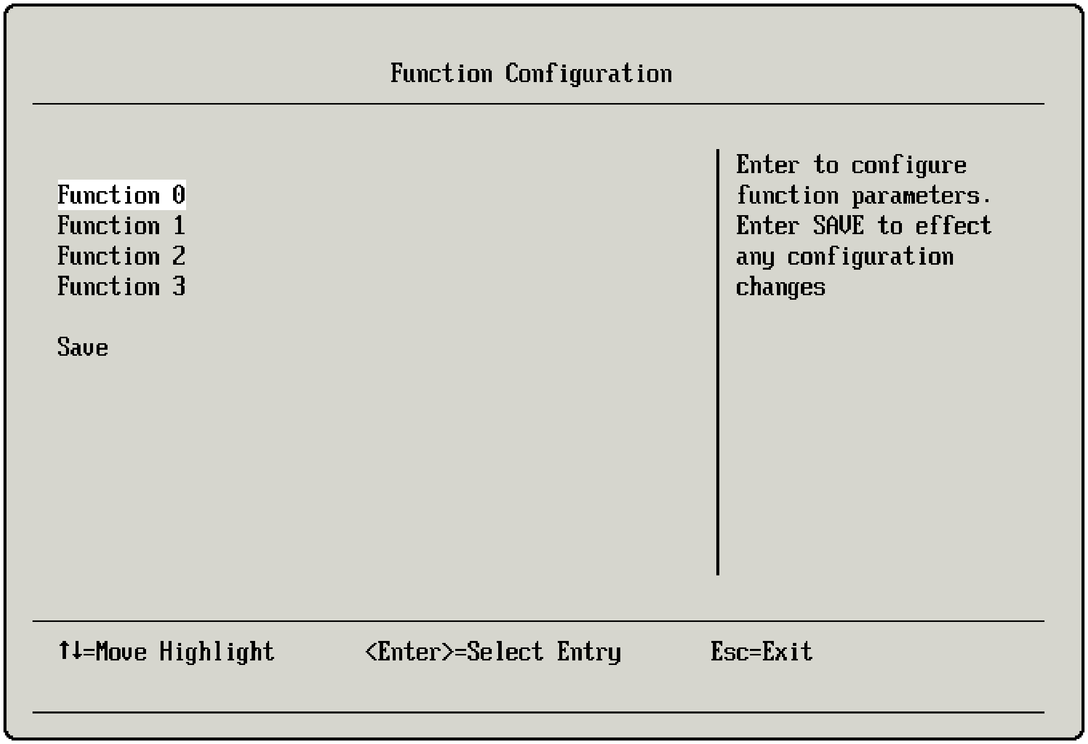

Figure 10-6 Function Configuration panel

a. In the Multichannel Configuration panel (Figure 10-7), for the first vNIC function, select MAC address **00:00:C9:B2:57:02**, and then set Bandwidth to 4 Gbps, for 40% of the total bandwidth. This vNIC uses VLAN 10.

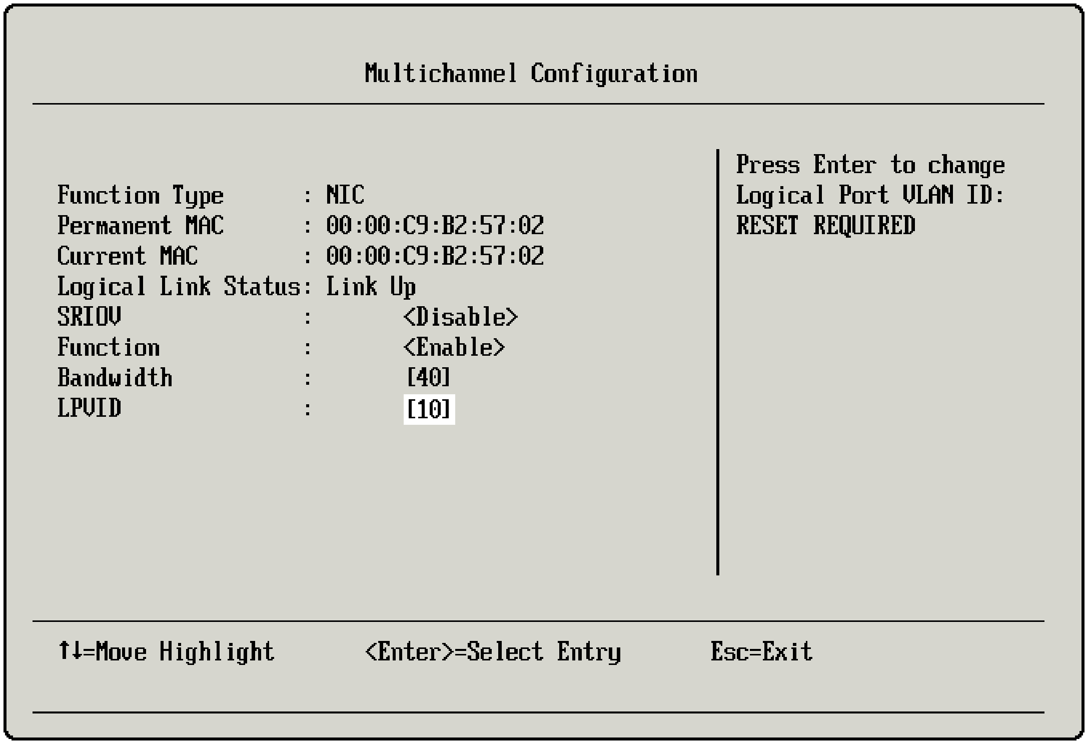

Figure 10-7 Bandwidth and VLAN settings for the first vNIC

b. For the second vNIC function (Figure 10-8), select MAC address **00:00:C9:B2:57:03**, and then set Bandwidth to 3 Gbps, for 30% of the total bandwidth. This vNIC uses VLAN 20.

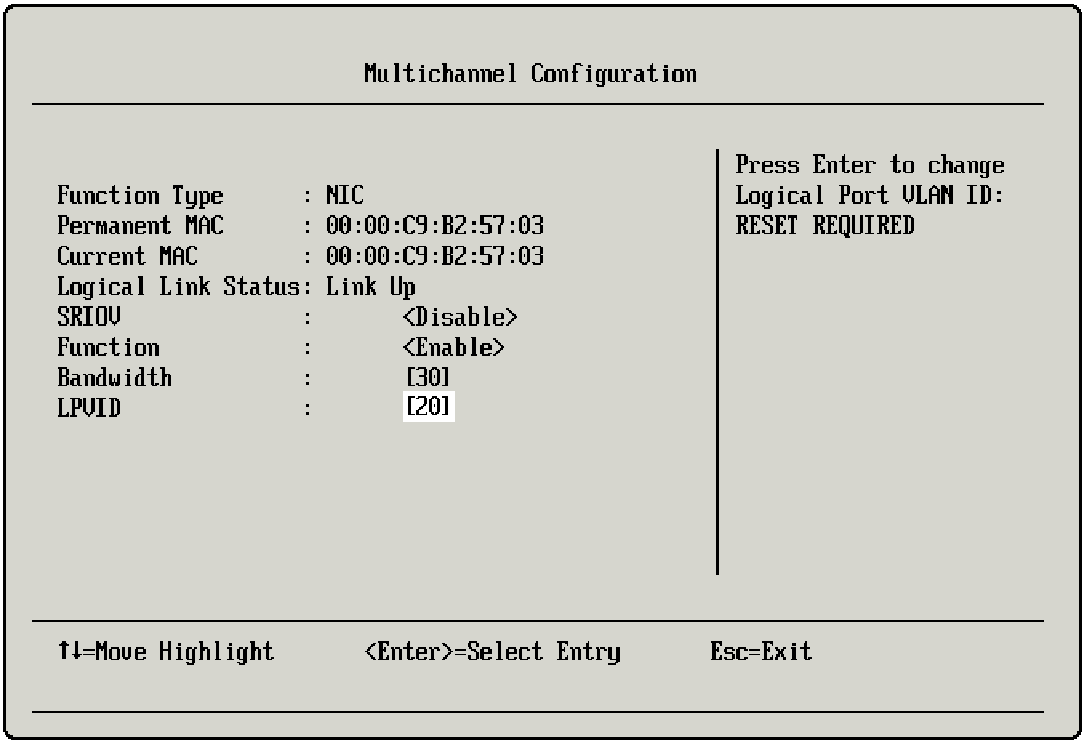

Figure 10-8 Bandwidth and VLAN settings for the second vNIC

c. For the third vNIC function (Figure 10-9), select MAC address **00:00:C9:B2:57:04**, and then set Bandwidth to 2 Gbps, for 20% of the total bandwidth. This vNIC uses VLAN 30.

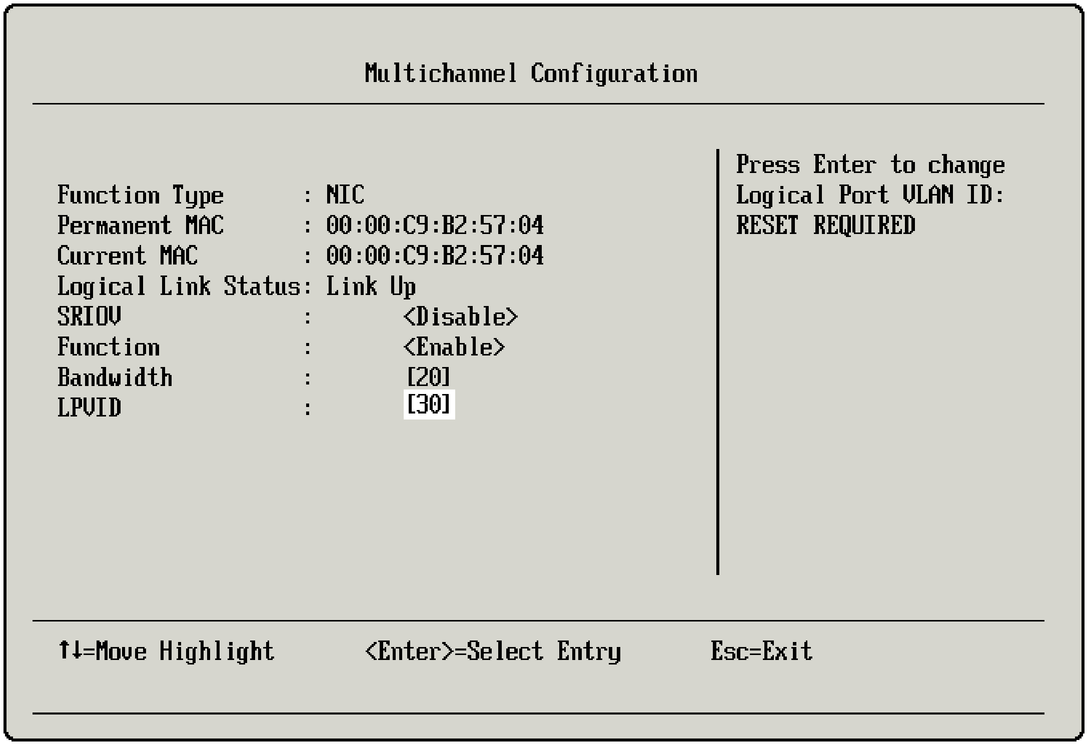

Figure 10-9 Bandwidth and VLAN settings for the third vNIC

d. For the fourth vNIC function (Figure 10-10), select MAC address **00:00:C9:B2:57:05**, and then set Bandwidth to 1 Gbps, for 10% of the total bandwidth. This vNIC uses VLAN 40.

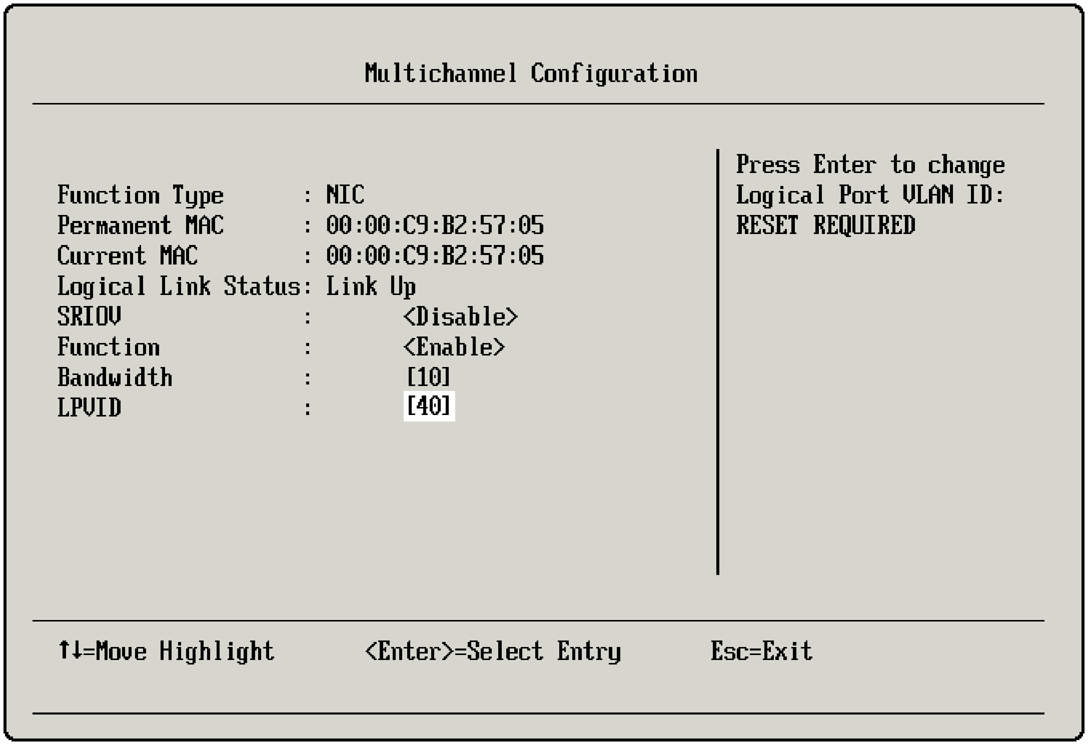

Figure 10-10 Bandwidth and VLAN settings for the fourth vNIC

8. With all four vNIC functions configured, back in the Function Configuration panel (Figure 10-11), click **Save** to save your changes.

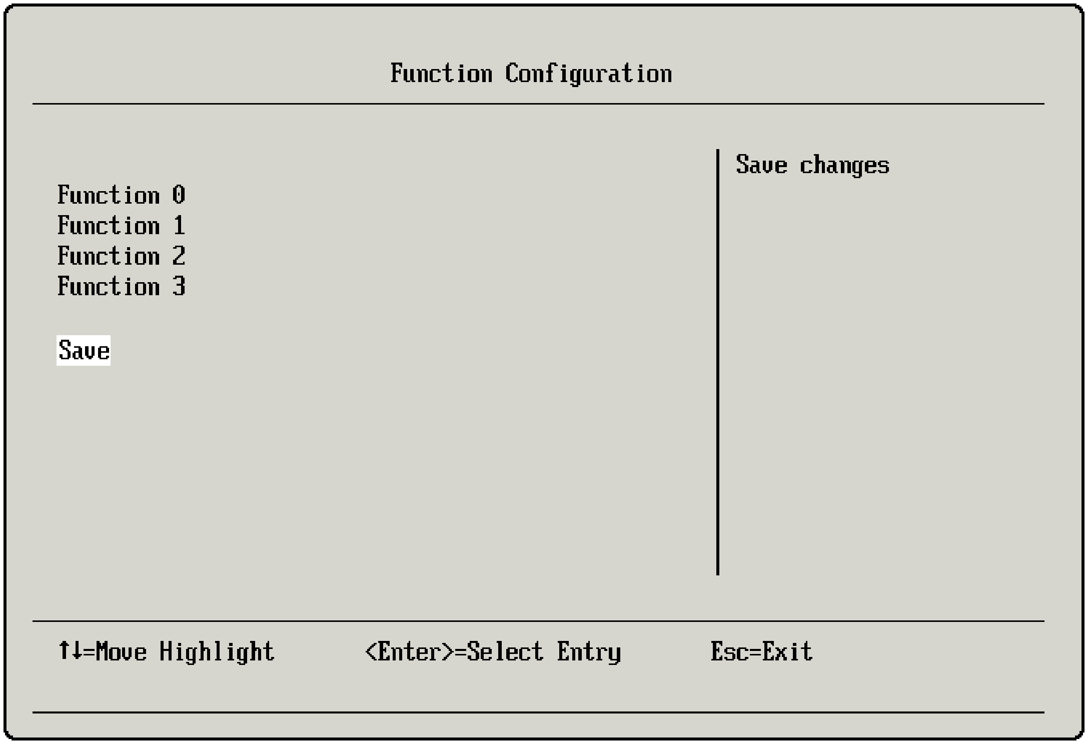

Figure 10-11 Saving the settings before exiting

9. Repeat the steps in this task to configure vNIC functions on the second 10 Gbps physical port.
10. Back in the System Configuration and Boot Management panel (Figure 10-3 on page 244), save your settings before exiting the UEFI utility.

vNICs in the operating system

HS22 blade server 1 runs Windows Server 2008. Now we check how the vNICs are displayed to the operating system.

The Network Connections window (Figure 10-12) lists eight vNICs, four on each 10 Gbps physical port.

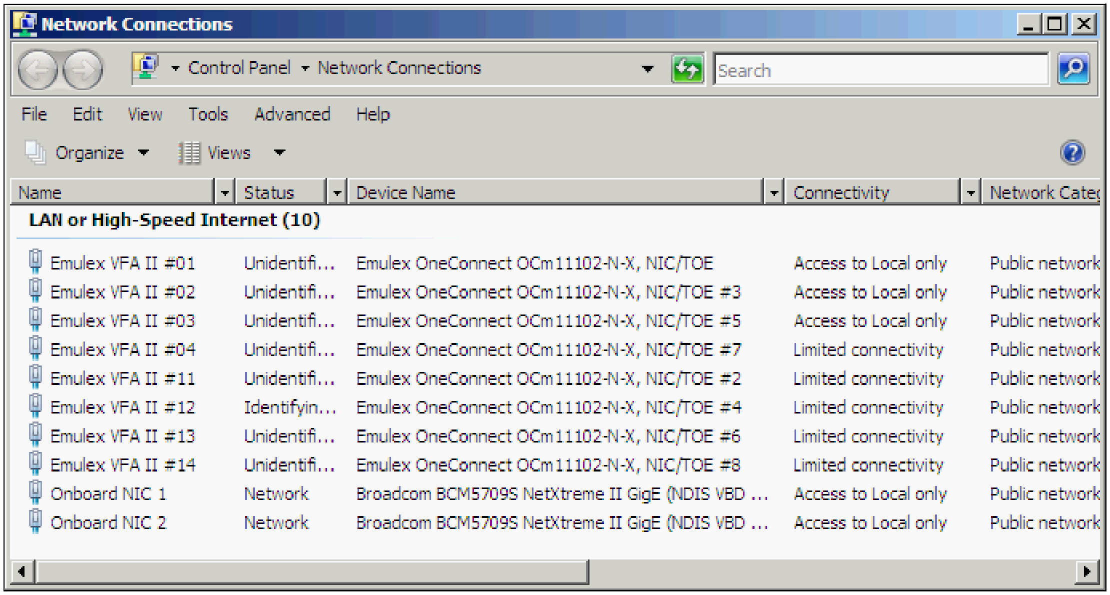

Figure 10-12 Network Connections window

For clarity, we changed the names of NICs from the default values to the values listed in Table 10-4. We used the MAC addresses to match each vNIC with the adapter shown in the Network Connections window.

Table 10-4 NIC names and matching vNICs

NIC name	10 Gbps physical port	vNIC	MAC address
Emulex VFA II #01	0	1	00:00:C9:B2:57:02
Emulex VFA II #02	0	2	00:00:C9:B2:57:03
Emulex VFA II #03	0	3	00:00:C9:B2:57:04
Emulex VFA II #04	0	4	00:00:C9:B2:57:05
Emulex VFA II #11	1	1	00:00:C9:B2:57:06
Emulex VFA II #12	1	2	00:00:C9:B2:57:07
Emulex VFA II #13	1	3	00:00:C9:B2:57:08
Emulex VFA II #14	1	4	00:00:C9:B2:57:09

To verify parameters, including MAC address and link speed, right-click a NIC (for example, the Emulex VFA II #01) and select **Status** (Figure 10-13).

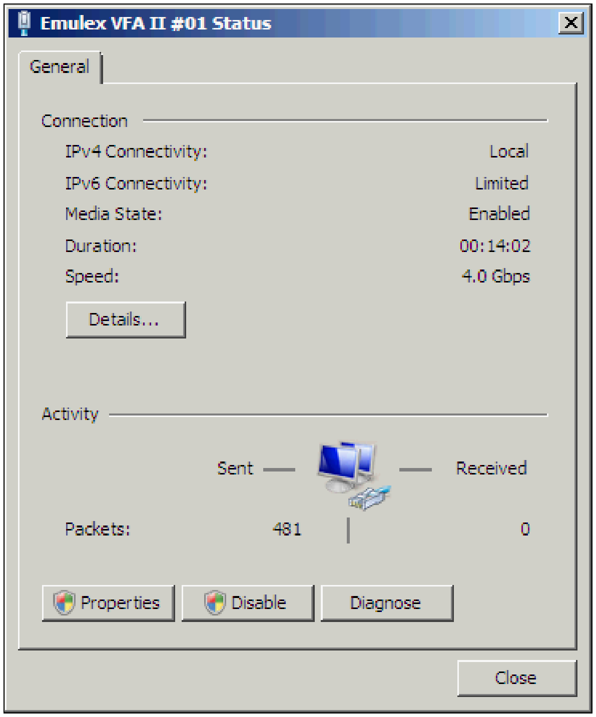

Figure 10-13 Emulex VFA II #01 status

Speed is reported as 4.0 Gbps, because we set the bandwidth to 40 for this vNIC (see Figure 10-7 on page 248).

To view the MAC address, click **Details**. The Network Connection Details window (Figure 10-14) contains the Physical Address field, which shows the MAC address. This MAC address belongs to vNIC 1 on 10 Gbps physical port 0.

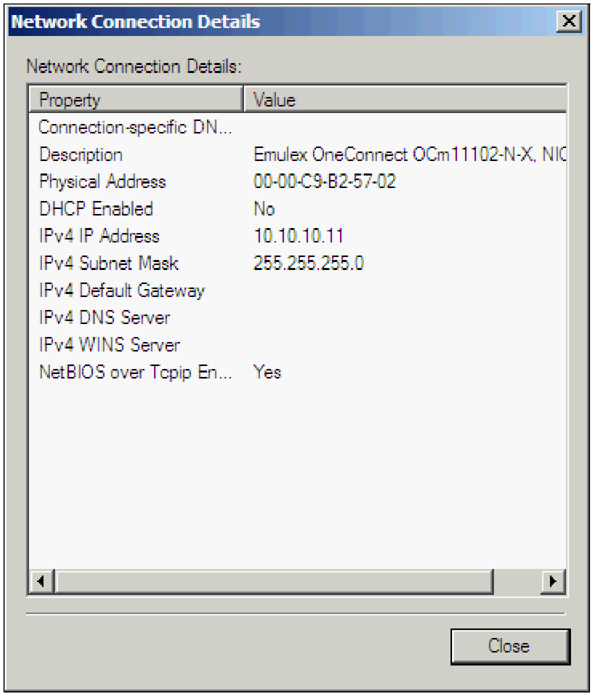

Figure 10-14 Network Connection details

Identify and verify MAC addresses and link speeds of all other vNICs in the same manner.

10.2.3 Configuring the Broadcom adapter

We set up the Broadcom 2-port 10Gb Virtual Fabric Adapter in a manner similar as for the Emulex adapter, with one difference. That is, we do not specify any VLAN ID on the Broadcom 2-port 10Gb Virtual Fabric Adapter, which uses only MAC addresses for virtual channel traffic separation.

Table 10-5 lists the settings for the Broadcom 2-port 10Gb Virtual Fabric Adapter.

Table 10-5 Broadcom 2-port 10Gb Virtual Fabric Adapter vNIC configuration

NIC partition	Bandwidth	MAC address
1	4 Gbps	10:18:00:96:E3:80
2	3 Gbps	10:18:00:96:E3:84
3	2 Gbps	10:18:00:96:E3:88
4	1 Gbps	10:18:00:96:E3:8C

We can identify the MAC addresses for each NIC partition by using the UEFI utility (see Figure 10-7 on page 248). It is important to know the MAC addresses so that you can identify each NIC partition in the operating system.

To configure the Broadcom 2-port 10Gb Virtual Fabric Adapter parameters in the UEFI utility:

1. Power on or restart the blade server.
2. When prompted during POST, press F1.
3. In the System Configuration and Boot Management panel of the UEFI utility (Figure 10-3 on page 244), select **System Settings**.
4. In the System Settings panel, select **Network**.

5. In the Network panel (Figure 10-15), select the first of the two 10 Gbps physical ports to configure.

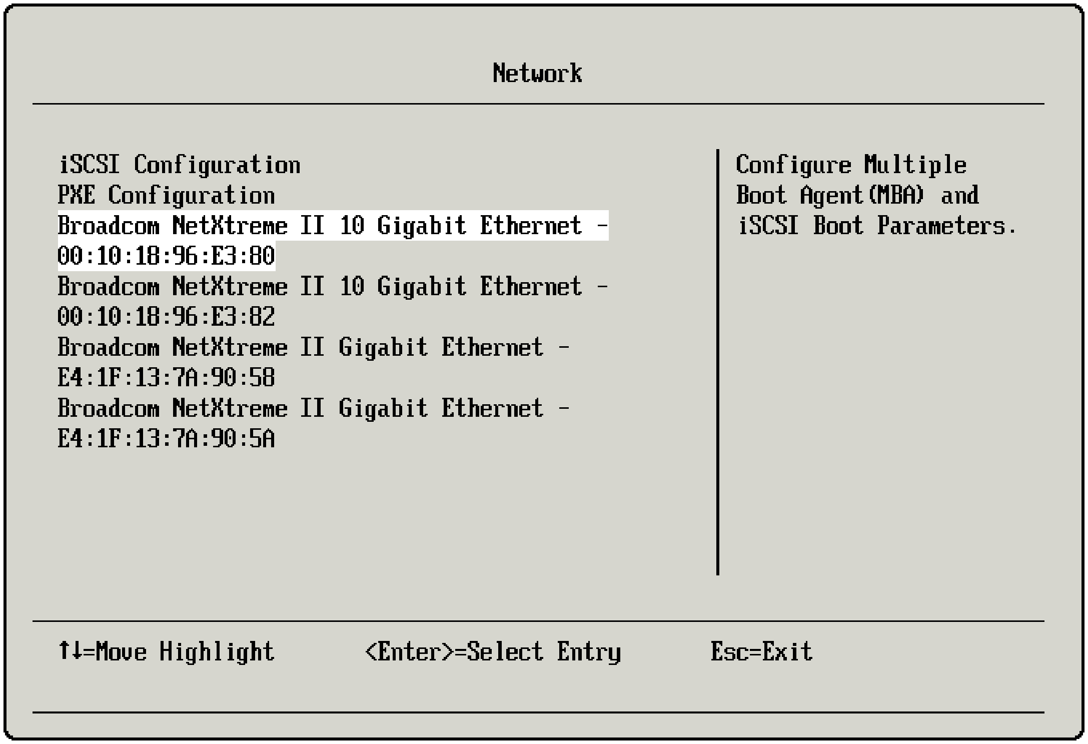

Figure 10-15 Network configuration: Selecting the first 10 Gbps port

6. In the Broadcom Main Configuration panel (Figure 10-16), select **Device Configuration Menu**.

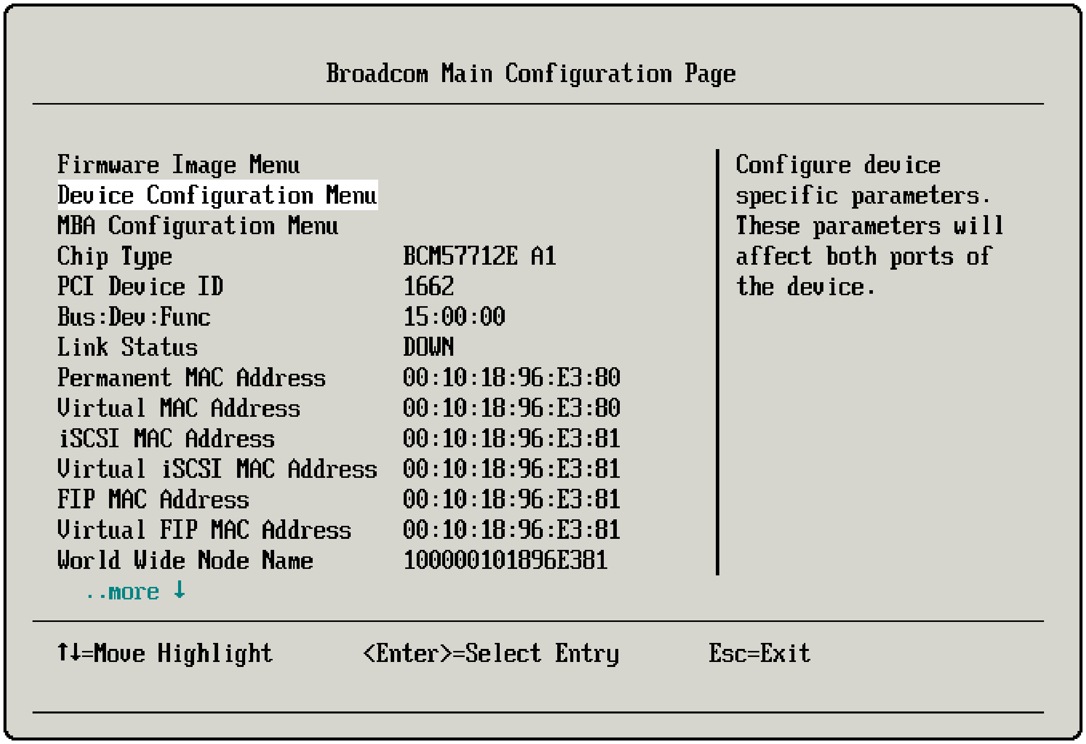

Figure 10-16 Broadcom Main Configuration panel

7. In the Device Configuration panel (Figure 10-17), highlight NIC Partition, and set it to **Enabled**. Press Esc.

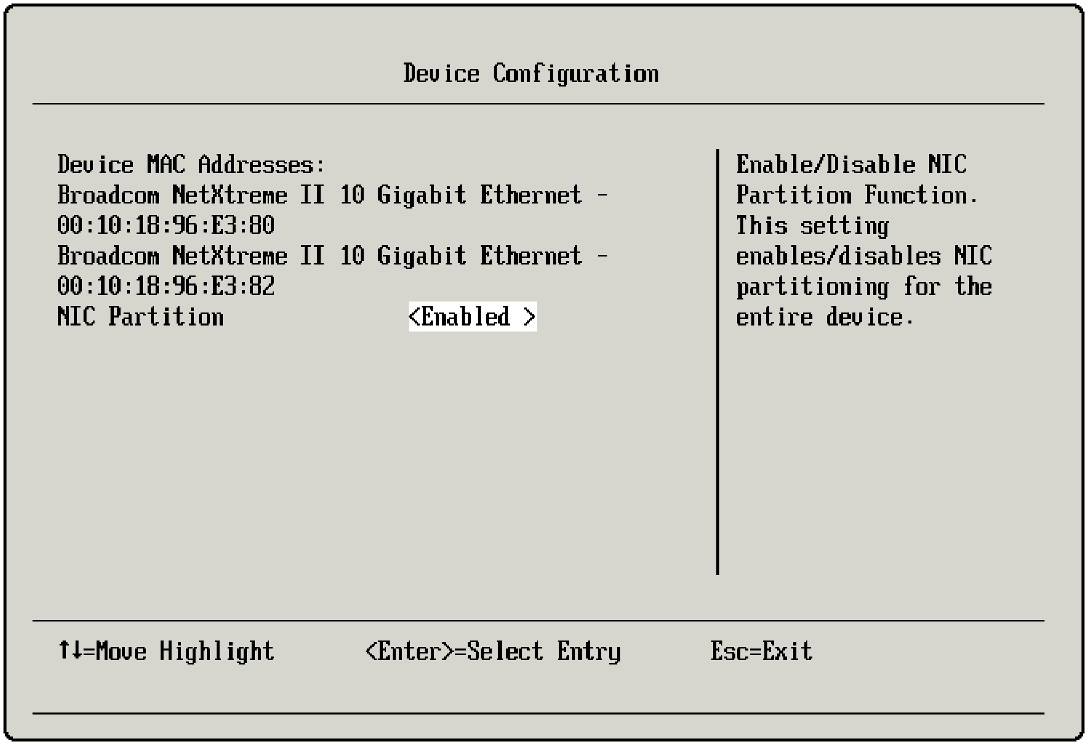

Figure 10-17 Device Configuration panel with NIC Partition set to Enabled

8. Back in the Broadcom Main Configuration panel (Figure 10-18), select **NIC Partitioning Configuration Menu** option. This option was added to this panel based on the previous settings.

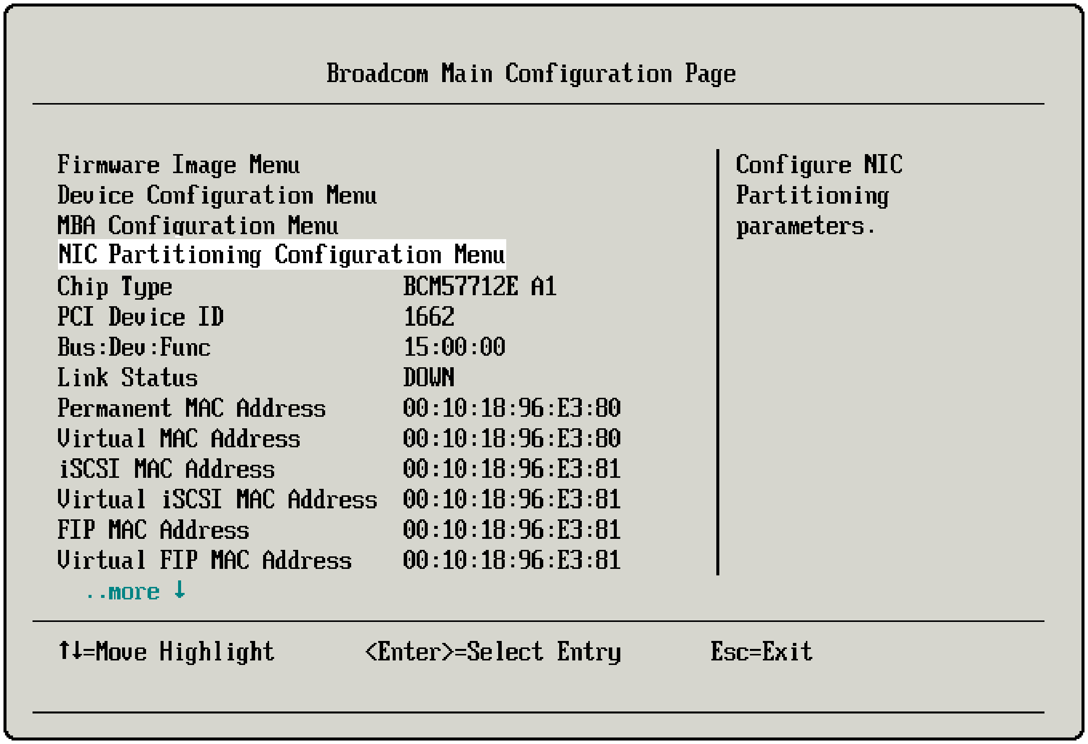

Figure 10-18 New configuration option for NIC partitioning

9. In the NIC Partition Configuration panel (Figure 10-19), identify the MAC addresses of all four NIC partitions:
   - PF#0 represents the first NIC partition, `L2=10180096E380`, which means the MAC address is 10:18:00:96:E3:80.
   - PF#2 is the second NIC partition, `L2=10180096E384`, which means the MAC address is 10:18:00:96:E3:84.
   - PF#4 is the third NIC partition, `L2=10180096E388`, which means the MAC address is 10:18:00:96:E3:88.
   - PF#6 is the fourth NIC partition, `L2=10180096E38C`, which means the MAC address is 10:18:00:96:E3:8C.

   You must know these MAC addresses so that you can match the NICs that are presented in the operating system to the NIC partitions.

   Select **Global Bandwidth Allocation Menu**, and then press Enter.

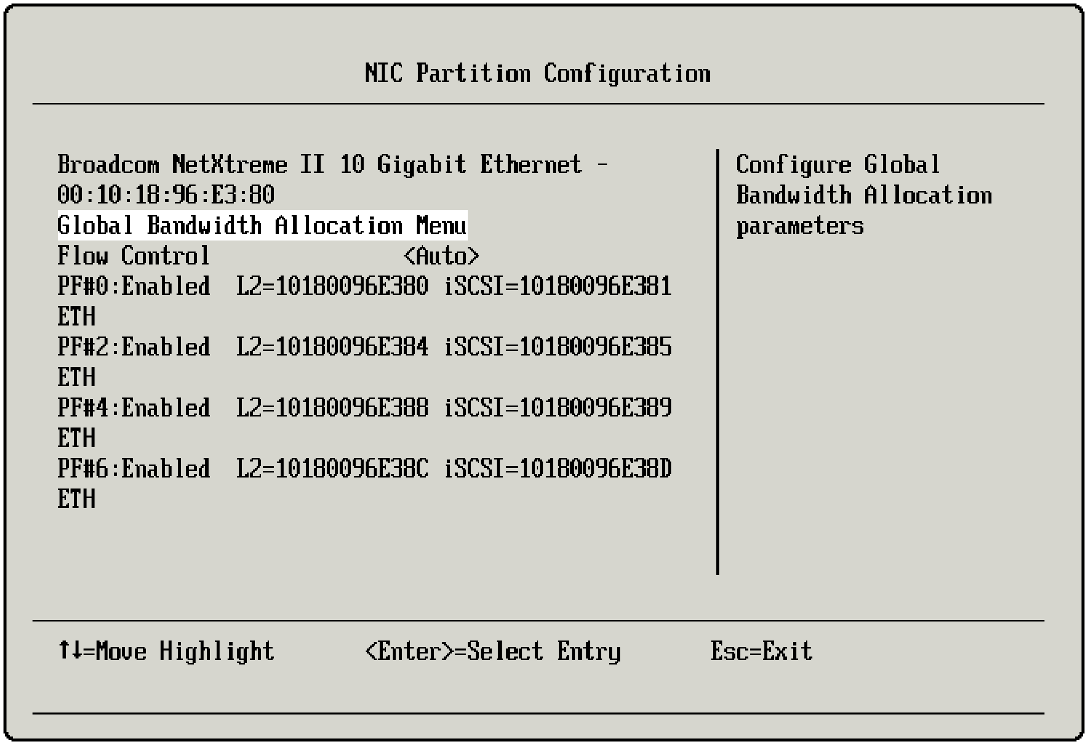

Figure 10-19 NIC Partition Configuration panel

10. In the Global Bandwidth Allocation Menu panel (Figure 10-20), set the bandwidth allocation parameters. In this example, we do not use relative bandwidth weight, so we leave these values at zero. However, we configure maximum bandwidth, as follows:

    a. Set NIC partition 1 to 4 Gbps (40% of total bandwidth).
    b. Set NIC partition 2 to 3 Gbps (30% of total bandwidth).
    c. Set NIC partition 3 to 2 Gbps (20% of total bandwidth).
    d. Set NIC partition 4 to 1 Gbps (10% of total bandwidth).

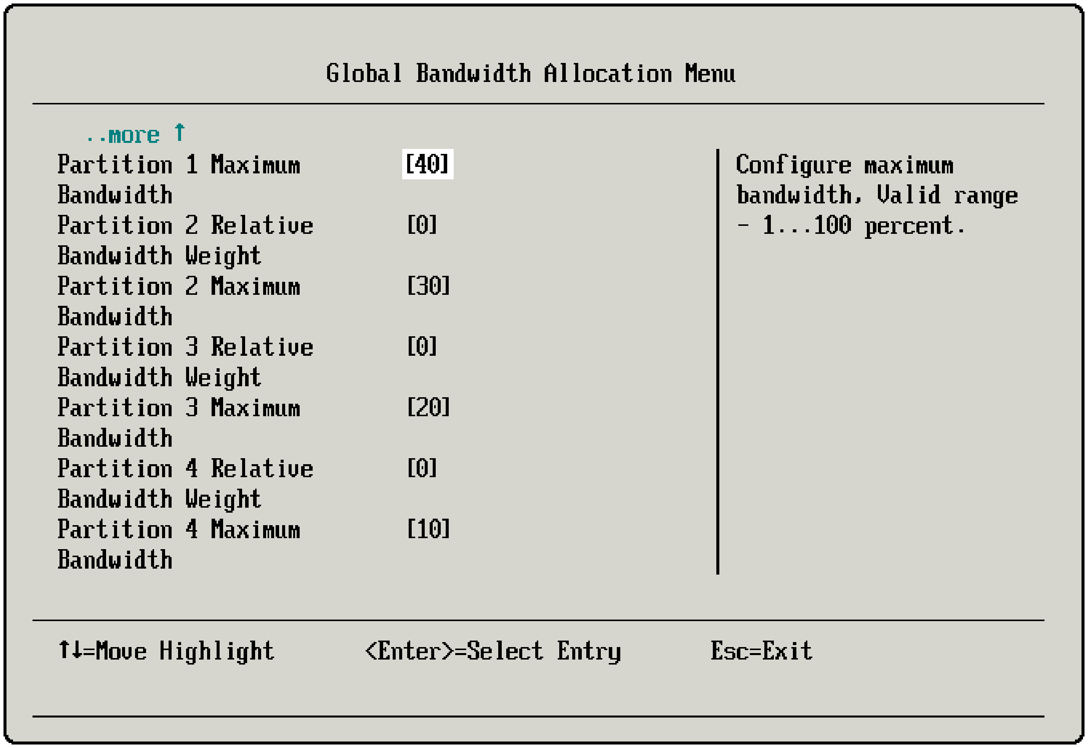

Figure 10-20 Global Bandwidth Allocation Menu panel

11. Configure NIC partitions and bandwidth parameters for the second 10 Gbps physical port, starting with Figure 10-15 on page 257.

12. Back in the System Configuration and Boot Management panel (Figure 10-3 on page 244), save your settings before exiting the UEFI utility.

vNICs in the operating system

We now verify how the vNIC partitions are displayed to the operating system. Blade 7 runs Windows Server 2008. This section highlights the Windows built-in networking tools and the Broadcom Advanced Control Suite.

Windows networking tools

Figure 10-21 shows how the NIC partitions are displayed in the Network and Sharing Center window.

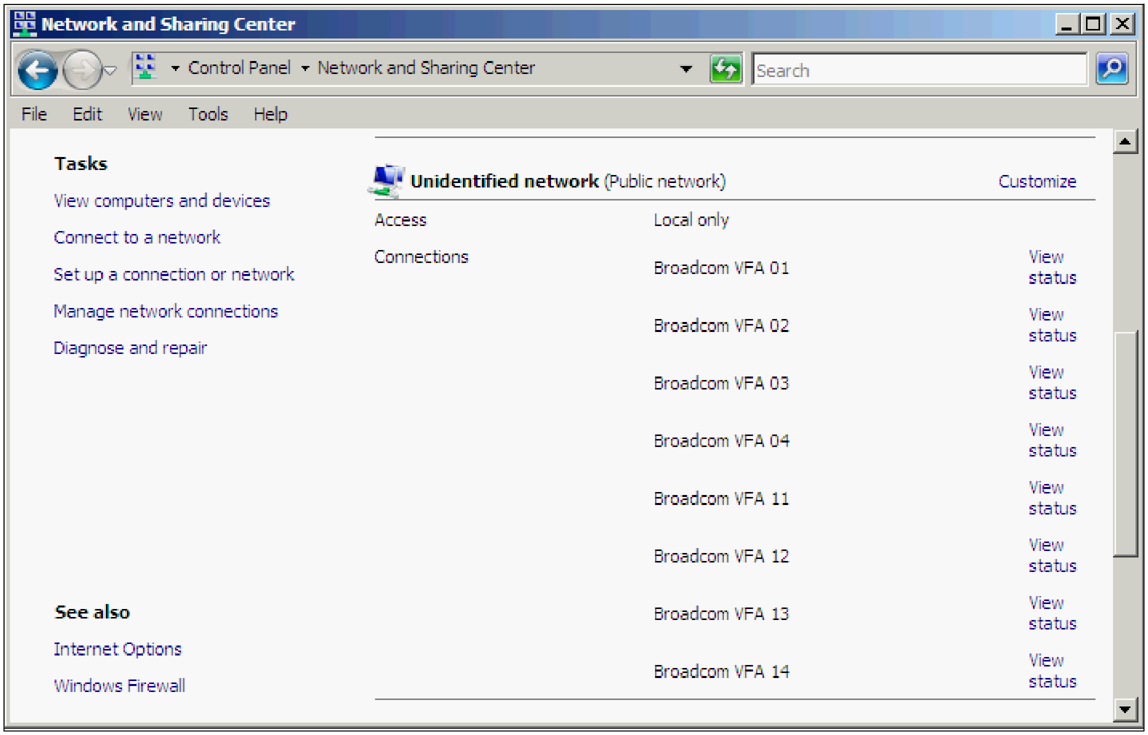

Figure 10-21 Broadcom 2-port 10Gb Virtual Fabric Adapter NIC partitions in Windows Server 2008

For clarity, we changed the NIC names from the default values as shown in Table 10-6. We used MAC addresses to match each NIC with the actual NIC partition.

Table 10-6 NIC names and matching NIC partitions

NIC name	10 Gbps physical port	NIC partition	MAC address
Broadcom VFA 01	0	1	10:18:00:96:E3:80
Broadcom VFA 02	0	2	10:18:00:96:E3:84
Broadcom VFA 03	0	3	10:18:00:96:E3:88
Broadcom VFA 04	0	4	10:18:00:96:E3:8C
Broadcom VFA 11	1	1	10:18:00:96:E3:82

NIC name	10 Gbps physical port	NIC partition	MAC address
Broadcom VFA 12	1	2	10:18:00:96:E3:86
Broadcom VFA 13	1	3	10:18:00:96:E3:8A
Broadcom VFA 14	1	4	10:18:00:96:E3:8E

To verify the parameters for each NIC partition, such as link speed and MAC address, in the Network and Sharing Center window, click **View Status**. Figure 10-22 shows the status for NIC Broadcom VFA 01.

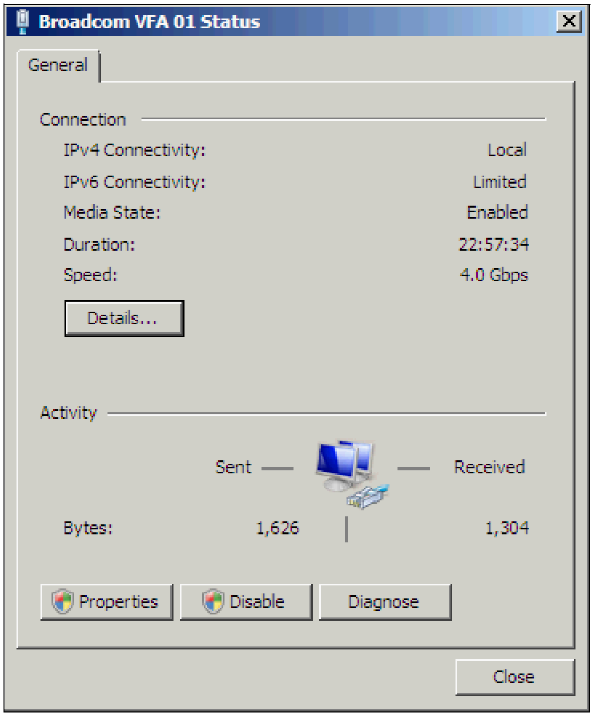

Figure 10-22 Broadcom VFA 01 status

As you can see, speed is reported as 4.0 Gbps, because we set bandwidth to 40 on this NIC partition (see Figure 10-20 on page 262).

To view the MAC address of this NIC partition, in the Status window, click **Details**.

The Network Connection Details window (Figure 10-23) shows that this address is the MAC address of NIC partition 1 on 10 Gbps physical port 0.

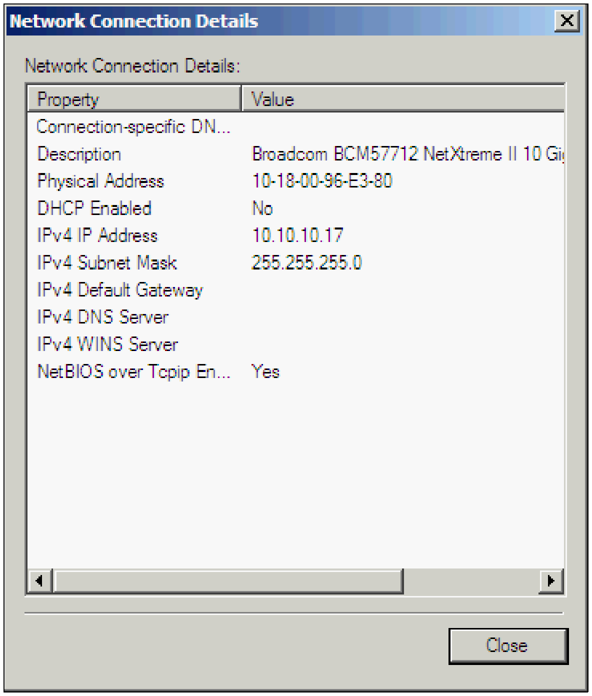

Figure 10-23 Network Connection Details showing the MAC address

We can use the same method to verify MAC addresses and speed for all other NIC partitions on the Broadcom 2-port 10Gb Virtual Fabric Adapter.

VLAN IDs in the operating system

You can define VLAN IDs for each NIC partition in the operating system. For example, to define VLAN ID 10 on NIC named Broadcom VFA 01:

1. In the Windows Device Manager window, open the Broadcom VFA 01 adapter.
2. In the Broadcom VFA 01 Properties window (Figure 10-24), click **Configure**.

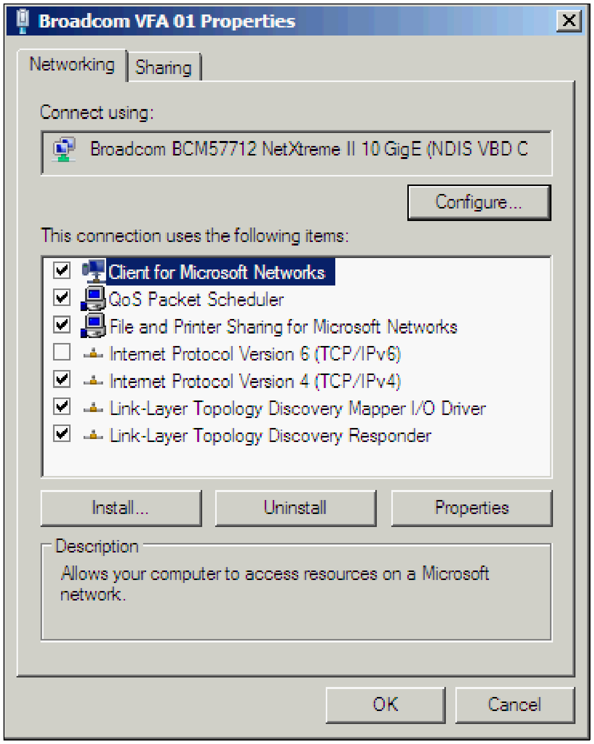

Figure 10-24 Broadcom VFA 01 properties window

3. On the **Advanced** tab (Figure 10-25), configure NIC settings, such as the VLAN ID. In this example, we set the VLAN ID to 10. Click **OK** to save your changes.

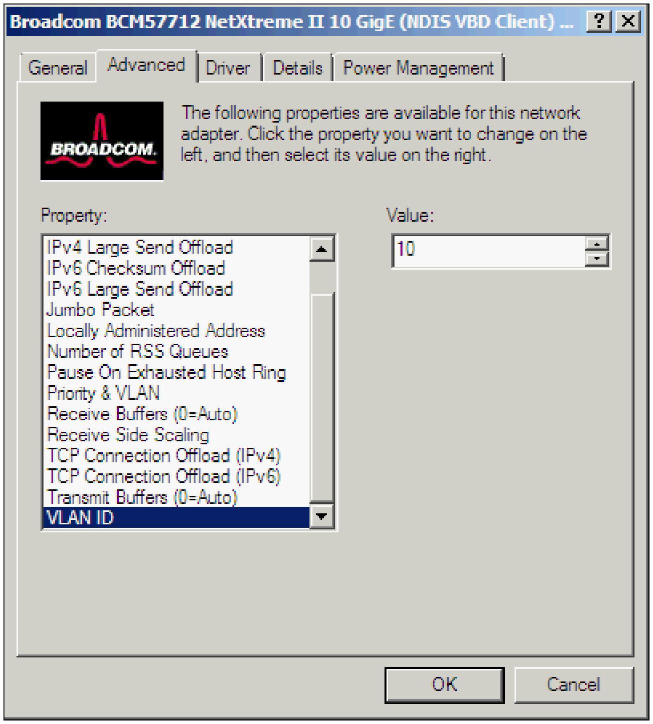

Figure 10-25 VLAN ID configuration

4. Repeat steps 1 on page 266 through 3 to configure the VLAN ID on other NIC partitions on 10 Gbps physical port 0, as follows:
   - VLAN ID 20 for Broadcom VFA 02
   - VLAN ID 30 for Broadcom VFA 03
   - VLAN ID 40 for Broadcom VFA 04
5. Configure the VLAN IDs on NIC partitions on 10 Gbps physical port 1 in the same way, starting with step 1 on page 266.

Broadcom Advanced Control Suite

Broadcom Advanced Control Suite is a management utility that you can use to administer, configure, and diagnose Broadcom Virtual Fabric Adapters and other Broadcom NICs. The example in Figure 10-26 shows eight NIC partitions in the left pane, with NIC partition 1 on physical port 0 highlighted. The speed is 4 Gbps, which was set previously (see Figure 10-20 on page 262).

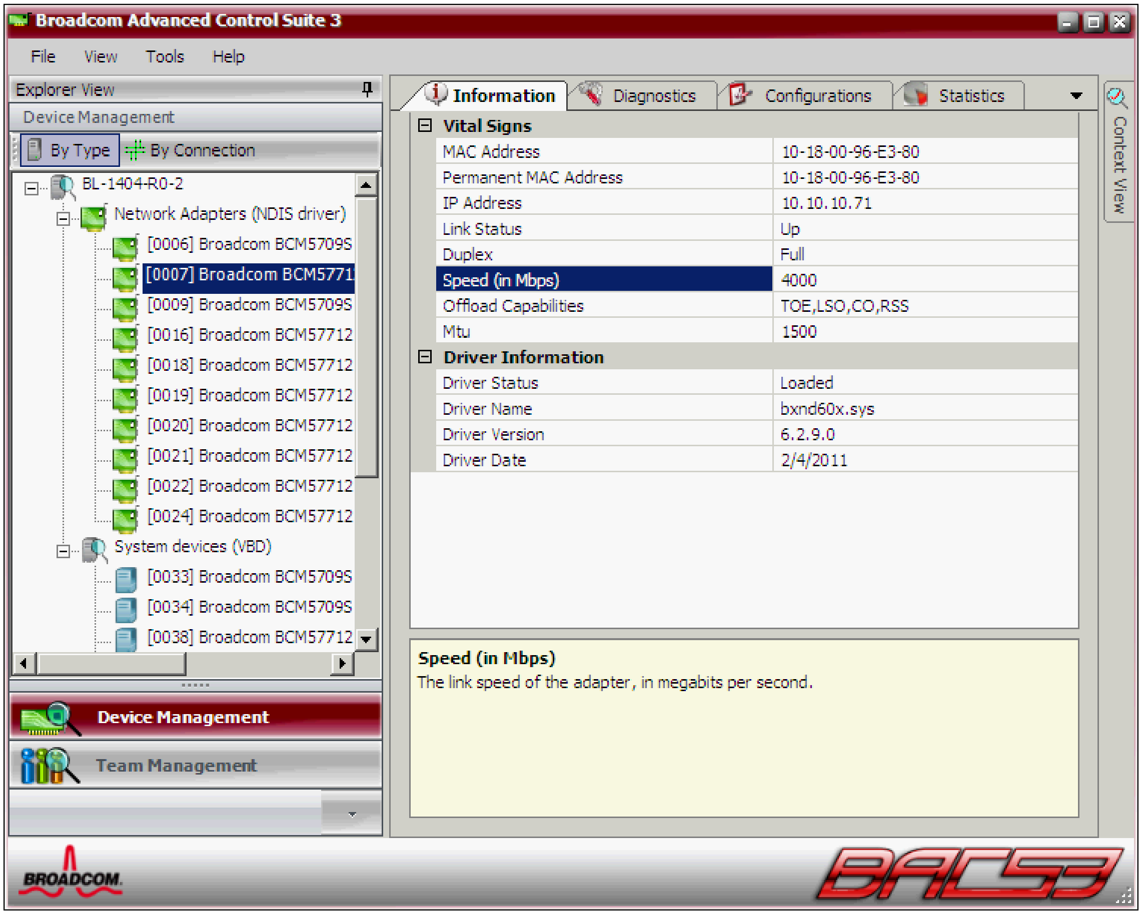

Figure 10-26 NIC partition 1 in Broadcom Advanced Control Suite

On the **Configurations** tab, you can check and set NIC parameters (Figure 10-27). Notice that VLAN ID is highlighted and is set to 10 (as we set it in Figure 10-25 on page 267). You can use Broadcom Advanced Control Suite to configure the VLAN ID to a different value or to change any other parameter.

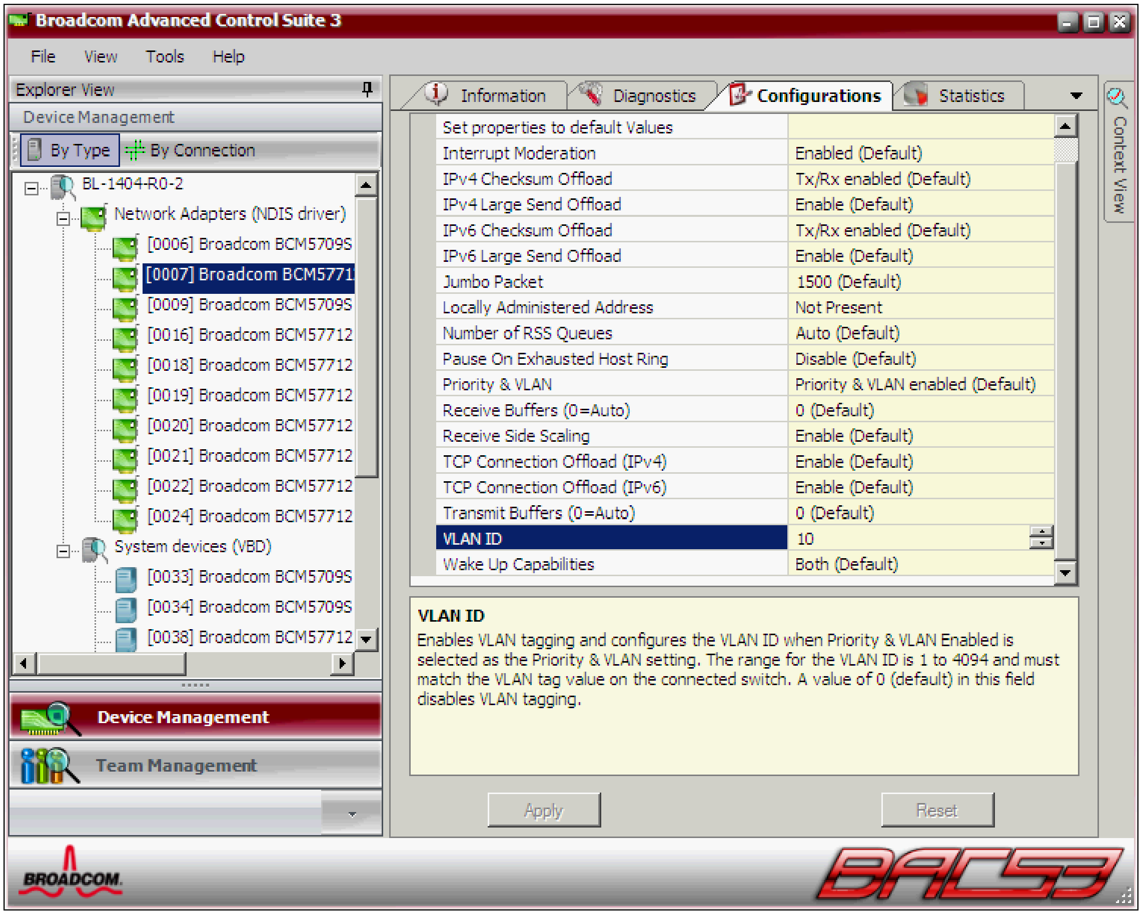

Figure 10-27 Configuration tab in Broadcom Advanced Control Suite

On the **Diagnostics** tab (Figure 10-28), you can test configurations by running a basic connectivity test, which is to ping to an IP gateway or to a user-specified IP address.

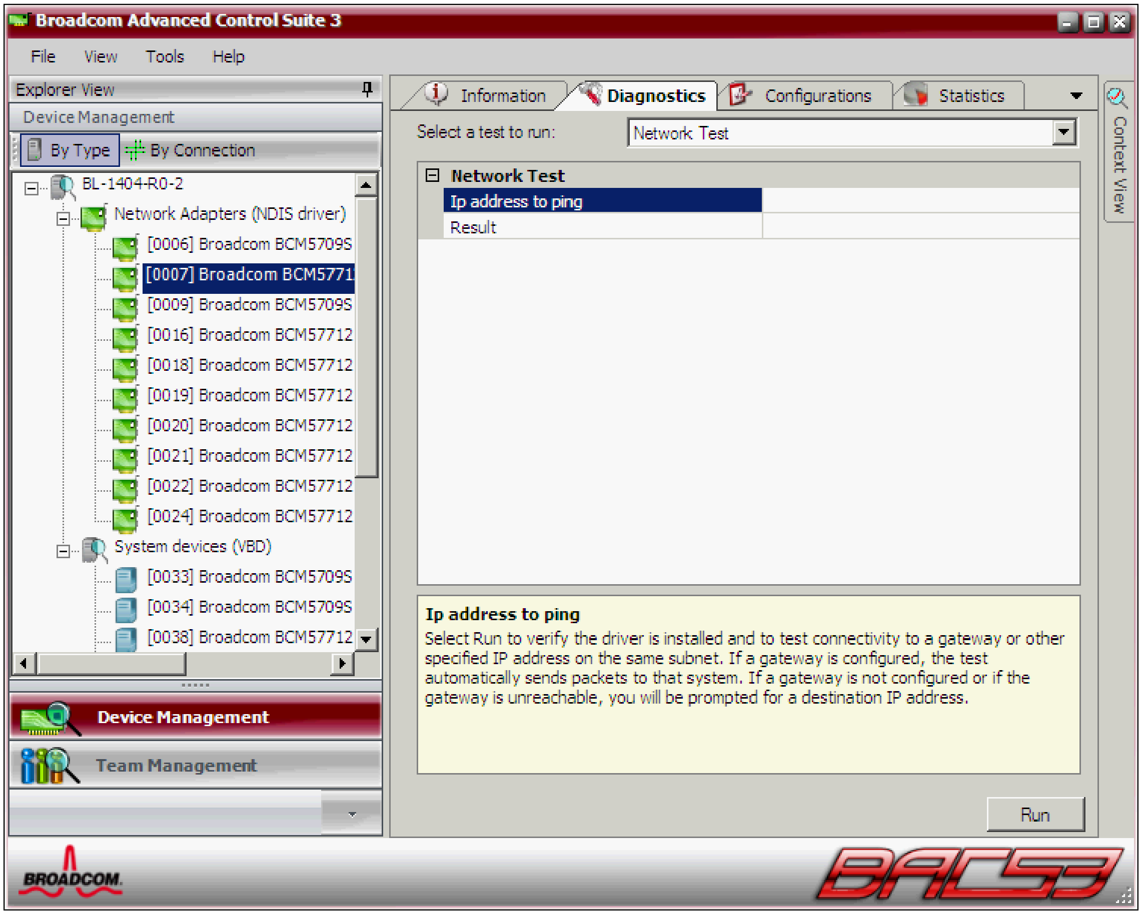

Figure 10-28 Diagnostics tab in Broadcom Advanced Control Suite

The **Statistics** tab (Figure 10-29) shows the network traffic counters. These counters are useful when monitoring or troubleshooting NIC behavior.

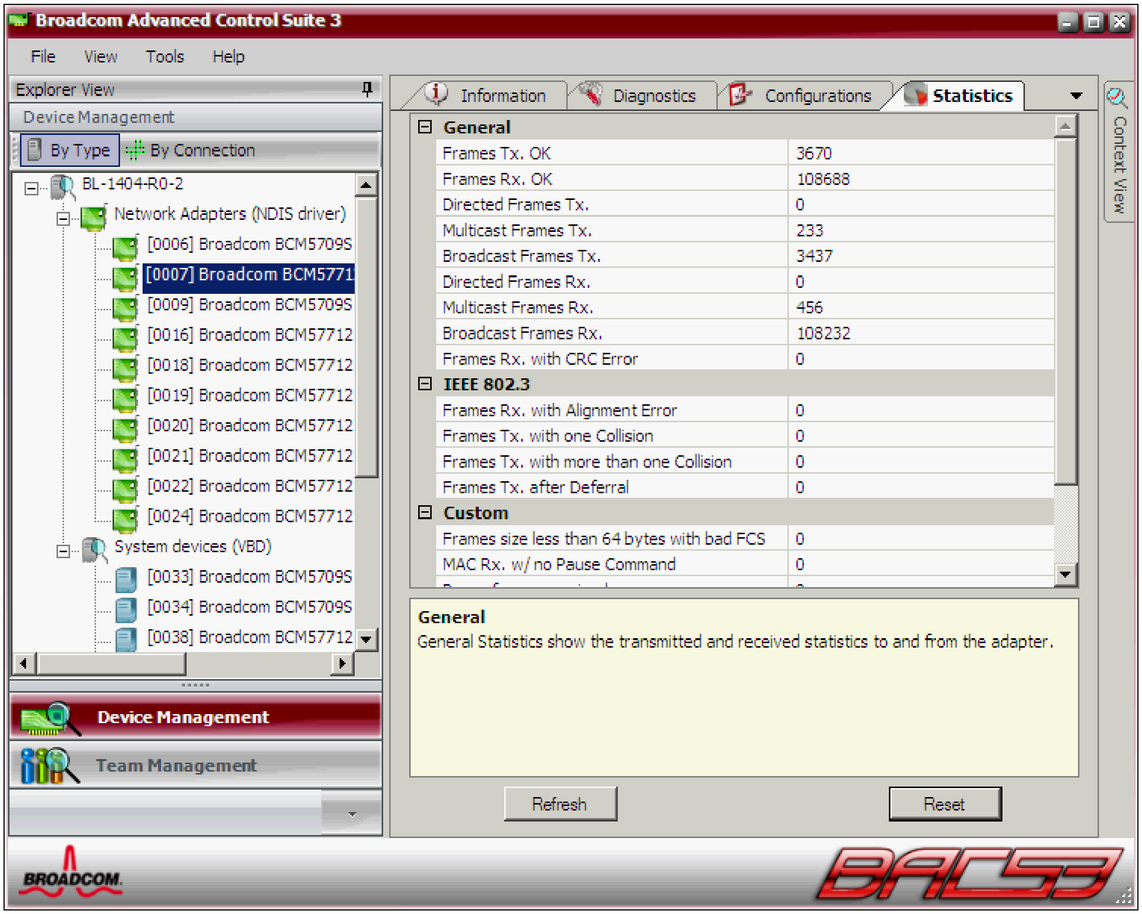

Figure 10-29 Statistics tab in the Broadcom Advanced Control Suite

One of the main functions of the Broadcom Advanced Control Suite utility is to configure NIC teaming. You can use NIC partitions in teaming configurations in the same way you use physical NICs.

To set up NIC teaming, click **Team Management** (lower-left corner) to start the teaming wizard (Figure 10-29 on page 271). From here, you can select the NICs (or NIC partitions) for the team, and then configure teaming parameters.

For more information about the teaming wizard, see 9.7.4, “Configuring teams with Broadcom Advanced Control Suite” on page 229.

10.2.4 Configuring the Cisco Nexus 4001I Switch Module

To configure the VLANs on the Cisco Nexus 4001I Switch Module:

1. Log in to the switch module by using Telnet and the default login credentials:
   - `USERID`
   - `PASSW0RD` (where the number 0 replaces the letter *O*)

   These same credentials are used on the advanced management module (AMM).
2. Configure the switch module by using the Cisco Nexus OS (NX-OS) command-line interface (CLI) as shown in Figure 10-30.

```
Nexus 4000 Switch
ibm-switch-2 login: USERID
Password:
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2010, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php
ibm-switch-2#
```

Figure 10-30 Cisco NX-OS CLI

When configuring a Switch Independent Mode solution, configure the VLANs on the Cisco Nexus 4001I Switch Module to match the existing VLAN structure in the network.

The HS22 blade in slot 1 contains the Emulex adapter, which is configured to use VLANs 10, 20, 30, and 40 on vNICs. Therefore, we need to set the internal (blade-facing) port 1 on the switch module to trunk mode and to allow VLANs 10, 20, 30, and 40 on it.

Figure 10-31 shows the commands to perform these actions.

```
ibm-switch-2(config)# configure terminal
ibm-switch-2(config)# interface Ethernet 1/1
ibm-switch-2(config-if)# switchport mode trunk
ibm-switch-2(config-if)# switchport trunk allowed vlan 10, 20, 30, 40
```

Figure 10-31 Configuring internal port 1 on Cisco Nexus 4001I Switch Module

The HS22 blade in slot 7 contains the Broadcom adapter. No VLAN ID configuration is in the UEFI settings for the Broadcom card, but we set VLANs 10, 20, 30, and 50 in the operating system. We now must configure internal port 7 on the switch module to pass traffic in these VLANs.

Figure 10-32 shows the commands.

```
ibm-switch-2(config-if)# configure terminal
ibm-switch-2(config)# interface Ethernet 1/7
ibm-switch-2(config-if)# switchport mode trunk
ibm-switch-2(config-if)# switchport trunk allowed vlan 10, 20, 30, 50
```

Figure 10-32 Configuring internal port 7 on Cisco Nexus 4001I Switch Module

We need to allow VLANs also on external ports, so that the blade servers can communicate with devices outside of the chassis. If any external port needs to carry traffic from multiple VLANs, you must set trunk port mode and allow the corresponding VLANs on it.

In this example, two external ports are used:

- External port 2 must carry VLANs 10 and 40, as shown by the commands in Figure 10-33.

  ```
  ibm-switch-2(config-if)# configure terminal
  ibm-switch-2(config)# interface Ethernet 1/16
  ibm-switch-2(config-if)# switchport mode trunk
  ibm-switch-2(config-if)# switchport trunk allowed vlan 10, 40
  ```

  Figure 10-33 External port 2 VLANs

- External port 4 must carry VLANs 20 and 50, as shown by the commands in Figure 10-34.

```
ibm-switch-2(config-if)# configure terminal
ibm-switch-2(config)# interface Ethernet 1/18
ibm-switch-2(config-if)# switchport mode trunk
ibm-switch-2(config-if)# switchport trunk allowed vlan 20, 50
```

Figure 10-34 External port 4 VLANs

The configuration of VLANs on the Cisco Nexus 4001I Switch Module is now completed.

10.3 Scenario 2: BNT Virtual Fabric 10Gb Switch Module configurations

This section shows examples with the BNT Virtual Fabric 10Gb Switch Module. These modules are installed in chassis BC5 in our lab examples. The blade in slot 4 has Emulex 10GbE Virtual Fabric Adapter II, and the blade in slot 8 has the Broadcom 2-port 10Gb Virtual Fabric Adapter. This scenario begins with an Emulex configuration, shows the Broadcom configuration, and concludes with setting up the switch module.

This section includes the following topics:

- Configuring the Emulex adapter
- Configuring the Broadcom adapter
- Configuring the BNT Virtual Fabric 10Gb Switch Module
- Configuring Switch Independent Mode for Emulex II, BNT, QLogic, and IBM Storwize V7000
- Configuring Switch Independent Mode: Emulex Virtual Fabric Adapter II, BNT, and IBM Storwize V7000 or iSCSI

10.3.1 Configuring the Emulex adapter

Table 10-7 shows the configuration we implemented on Emulex 10GbE Virtual Fabric Adapter II on the HS22, blade 4.

Table 10-7 Emulex Virtual Fabric Adapter II vNIC configuration

vNIC function	Bandwidth	VLAN ID	MAC address
1	4 Gbps	10	00:00:C9:B2:5B:3A
2	3 Gbps	20	00:00:C9:B2:5B:3B
3	2 Gbps	30	00:00:C9:B2:5B:3C
4	1 Gbps	40	00:00:C9:B2:5B:3D

You can identify the MAC addresses for each vNIC function by using the UEFI utility (see Figure 10-19 on page 261). With the MAC addresses, you can then identify each vNIC function in the operating system.

To configure the parameters on the Emulex adapter:

1. Power on or restart HS22 blade 4.
2. When prompted during POST, press F1.
3. In the System Configuration and Boot Management panel (Figure 10-3 on page 244), select **System Settings**.
4. In the System Settings panel, select **Network**.

5. In the Network panel, which shows two Emulex 10 Gbps ports (Figure 10-35), select the first port to configure the vNICs, and then press Enter.

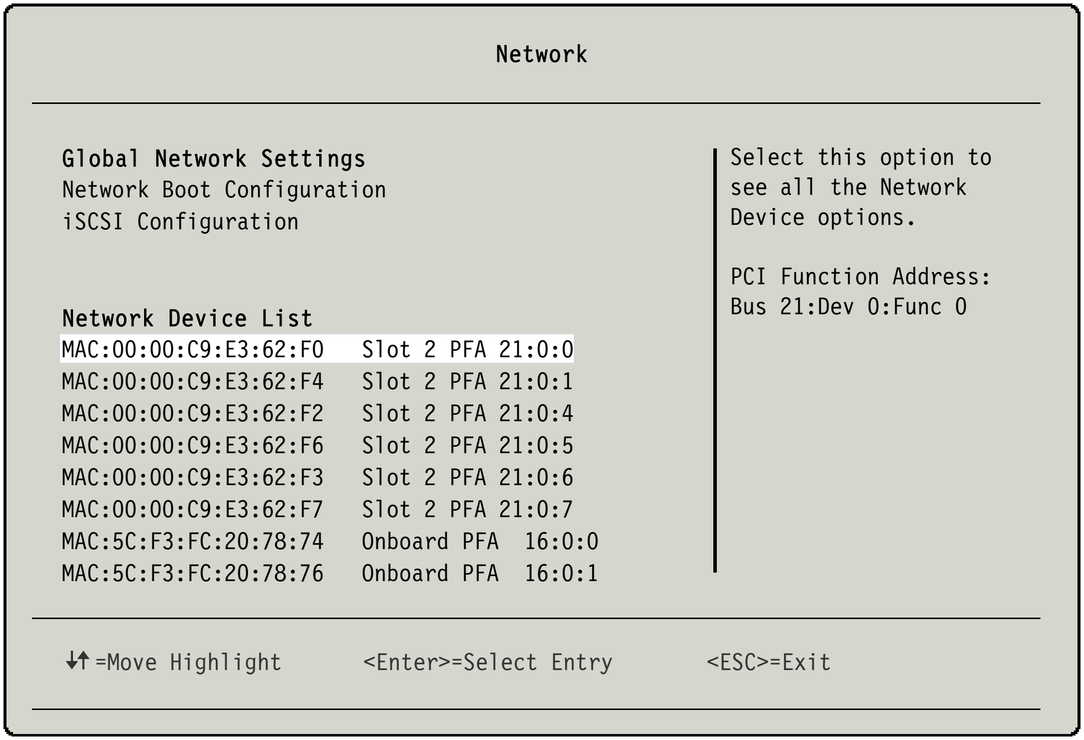

Figure 10-35 Network configuration panel

6. In the Emulex NIC Selection panel (Figure 10-36):
   a. Set Multichannel to **Enabled**.
   b. Select **Switch Independent Mode**.
   c. Highlight **Multichannel Configuration** and press Enter to configure the VLAN and bandwidth settings for each of the four vNICs.

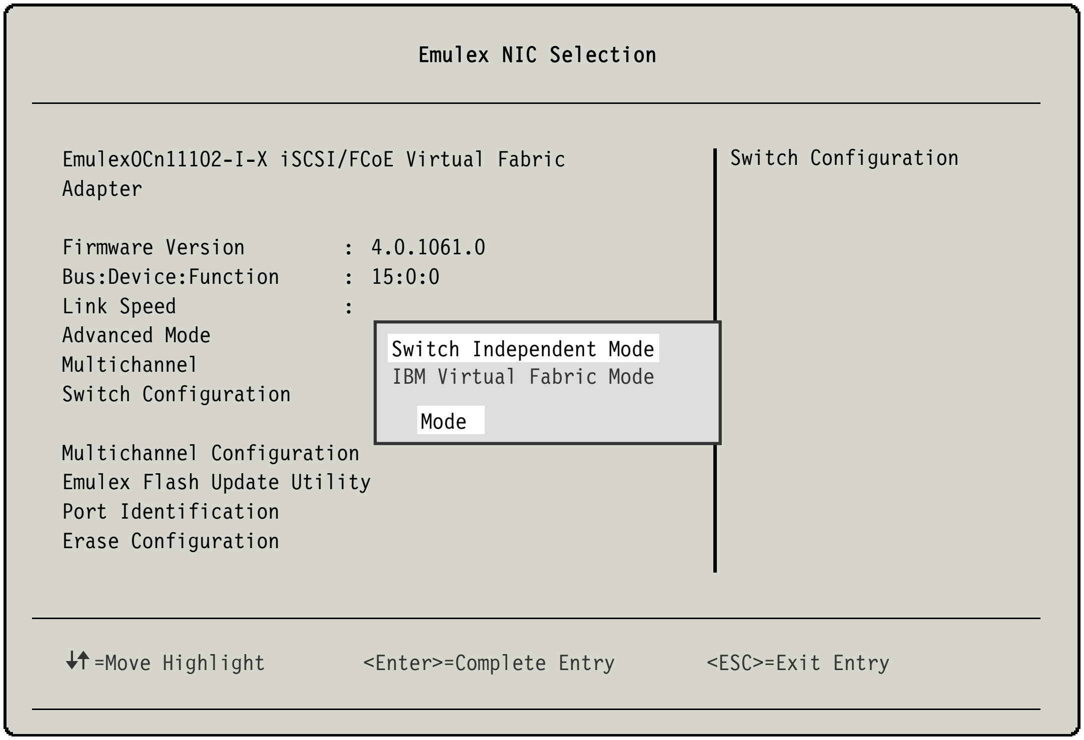

Figure 10-36 Emulex NIC Selection panel

7. In the Function Configuration panel (Figure 10-37), configure each of the four vNIC functions. To begin, select **Function 0**, and then press Enter.

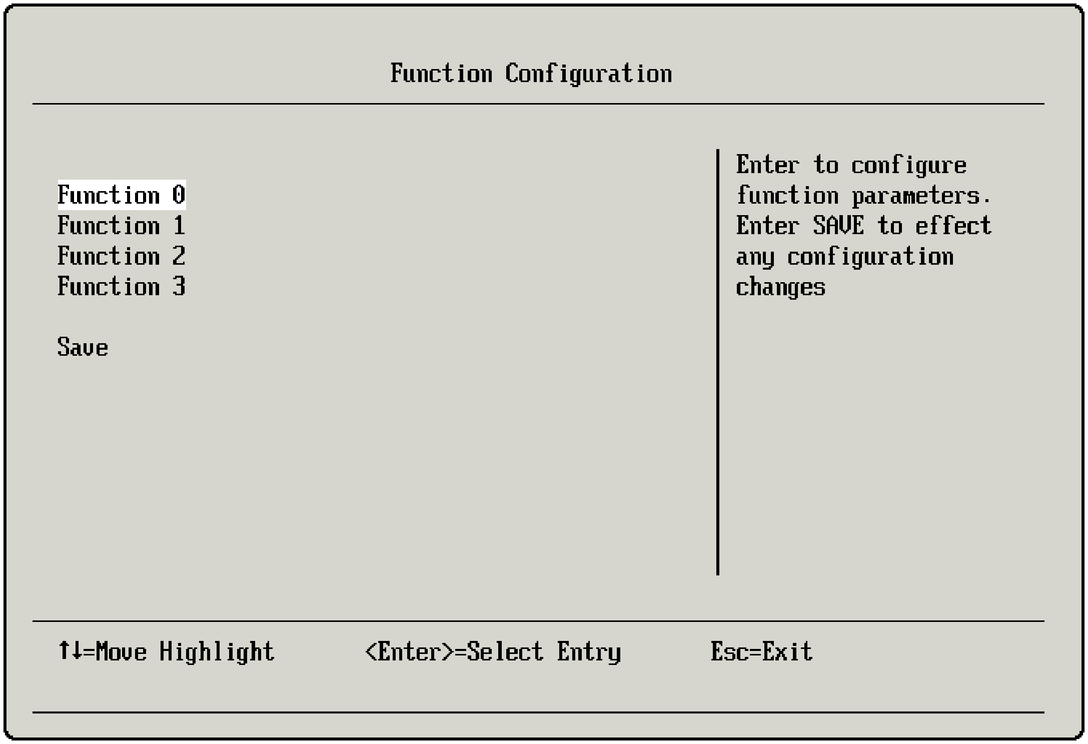

Figure 10-37 Function Configuration panel

a. In the Multichannel Configuration panel (Figure 10-38), for the first vNIC, select MAC address **00:00:C9:B2:5B:3A**, and then set Bandwidth to 4 Gbps, for 40% of the total bandwidth. This vNIC uses VLAN 10.

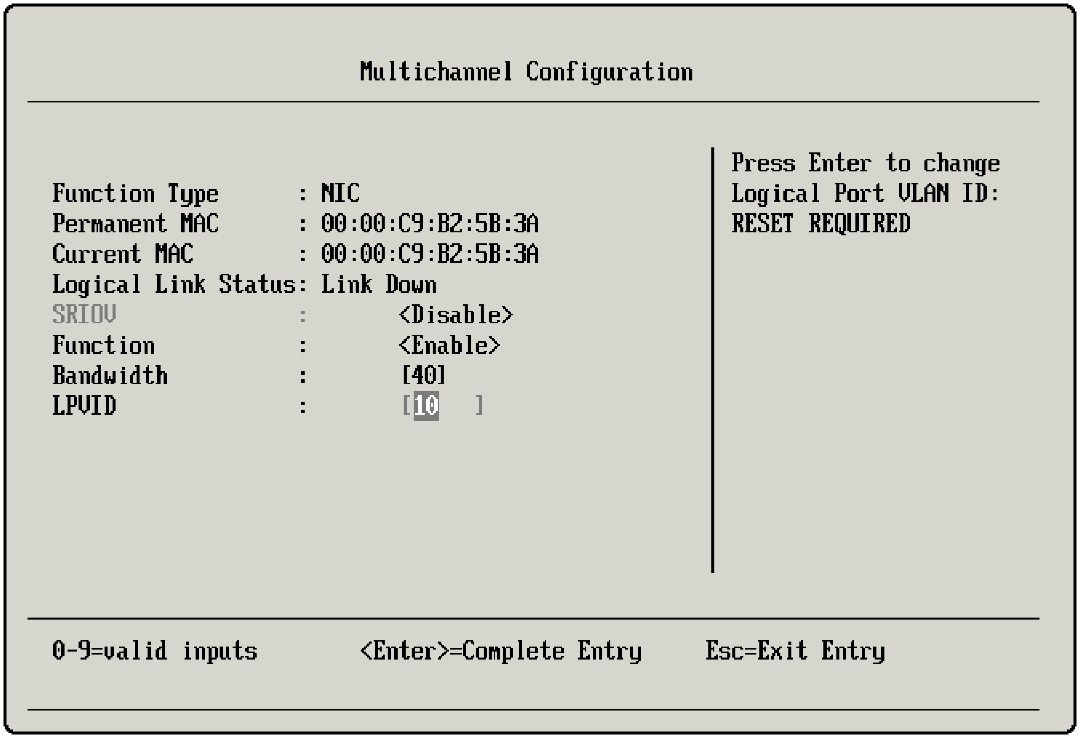

Figure 10-38 Bandwidth and VLAN settings for the first vNIC

b. For the second vNIC (Figure 10-39), select MAC address **00:00:C9:B2:5B:3B**, and then set Bandwidth to 3 Gbps, for 30% of the total bandwidth. This vNIC uses VLAN 20.

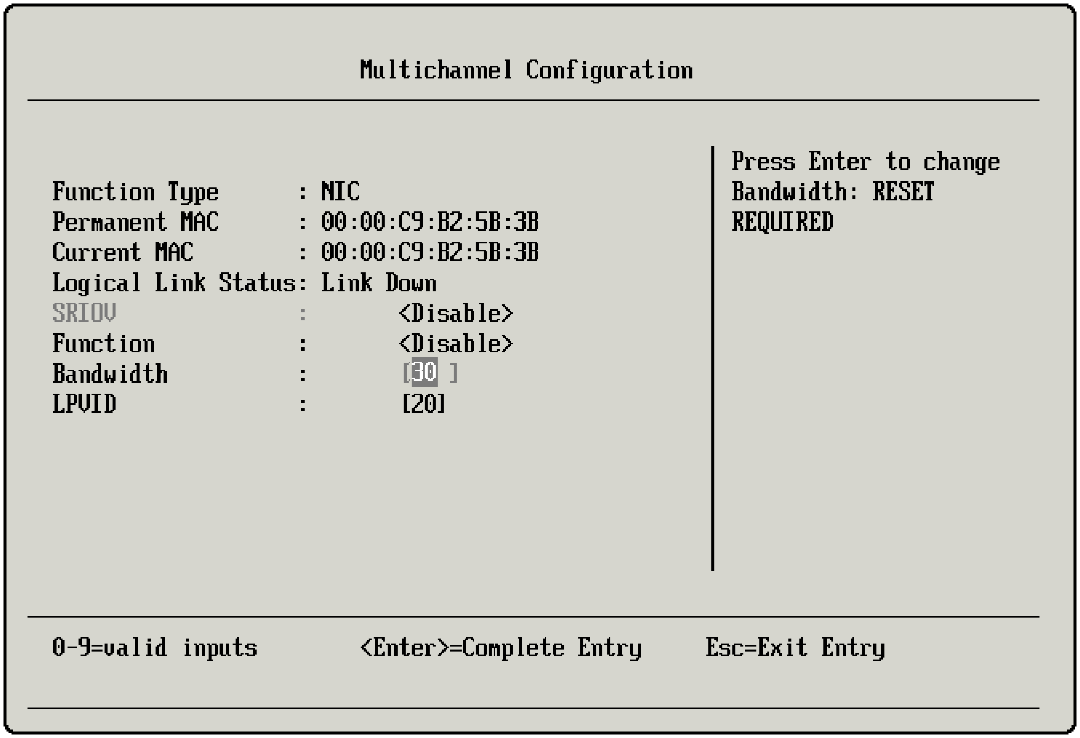

Figure 10-39 Bandwidth and VLAN settings for the second vNIC

c. For the third vNIC (Figure 10-40), select MAC address **00:00:C9:B2:5B:3C**, and then set Bandwidth to 2 Gbps, for 20% of the total bandwidth. This vNIC uses VLAN 30.

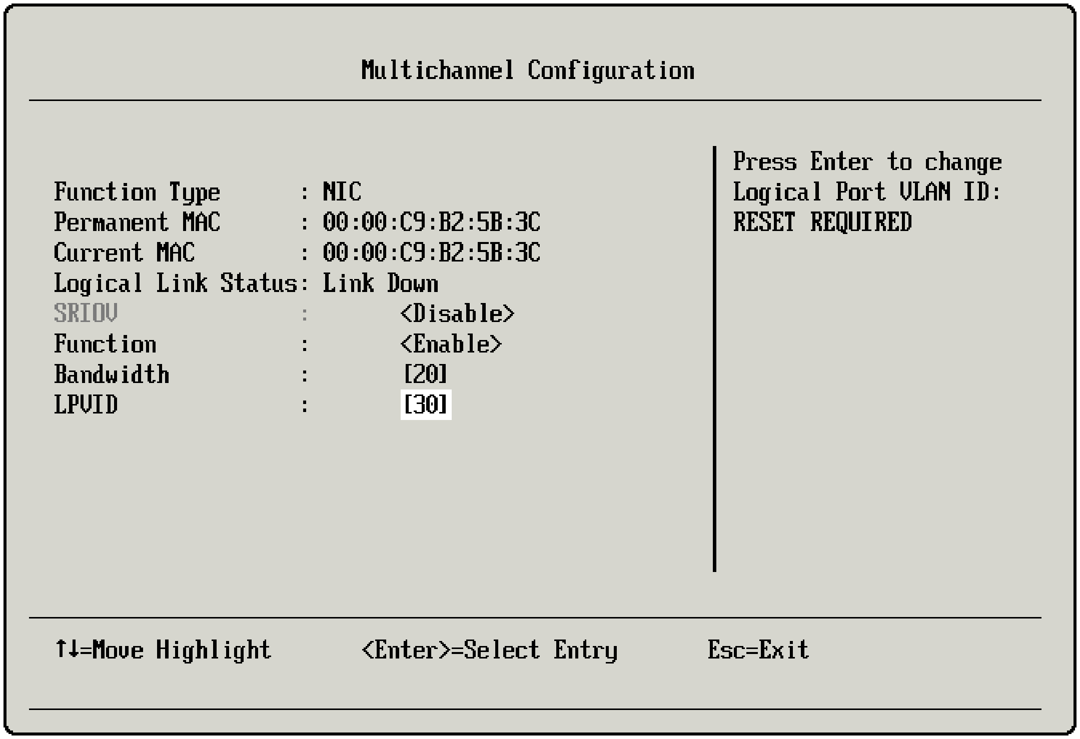

Figure 10-40 Bandwidth and VLAN settings for the third vNIC

d. For the fourth vNIC (Figure 10-41), select MAC address **00:00:C9:B2:5B:3D**, and then set Bandwidth to 1 Gbps, for 10% of the total bandwidth. This vNIC uses VLAN 40.

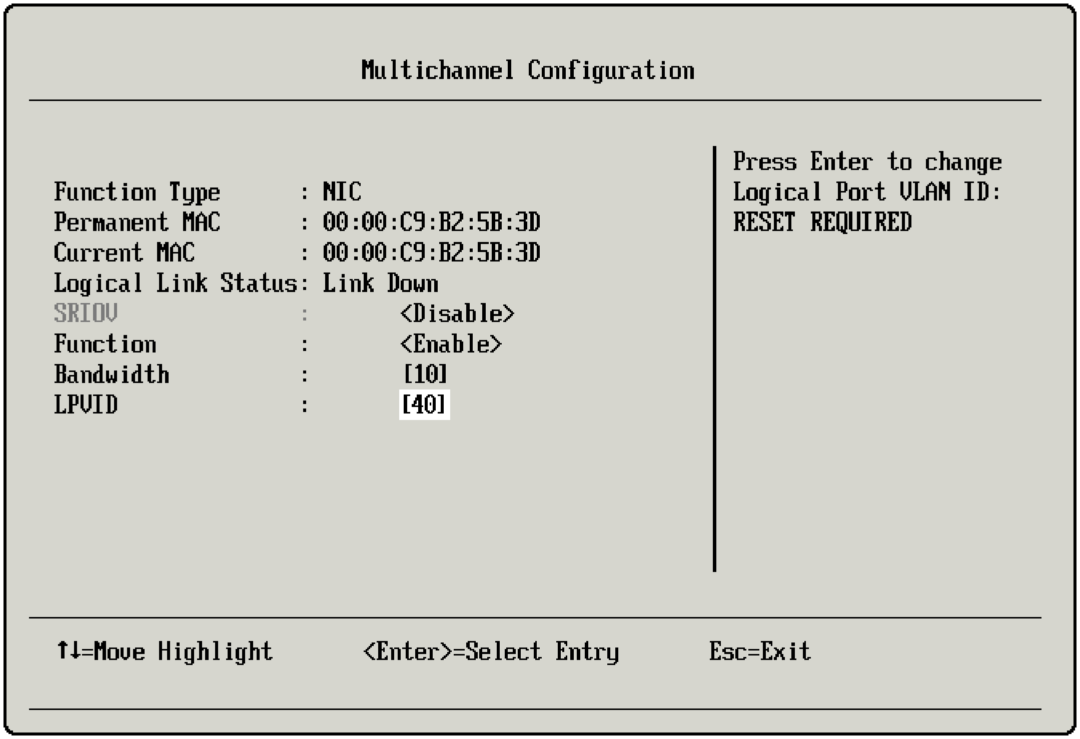

Figure 10-41 Bandwidth and VLAN settings for the fourth vNIC

8. With all four vNIC functions configured, back in the Function Configuration panel (Figure 10-42), click **Save** to save your changes.

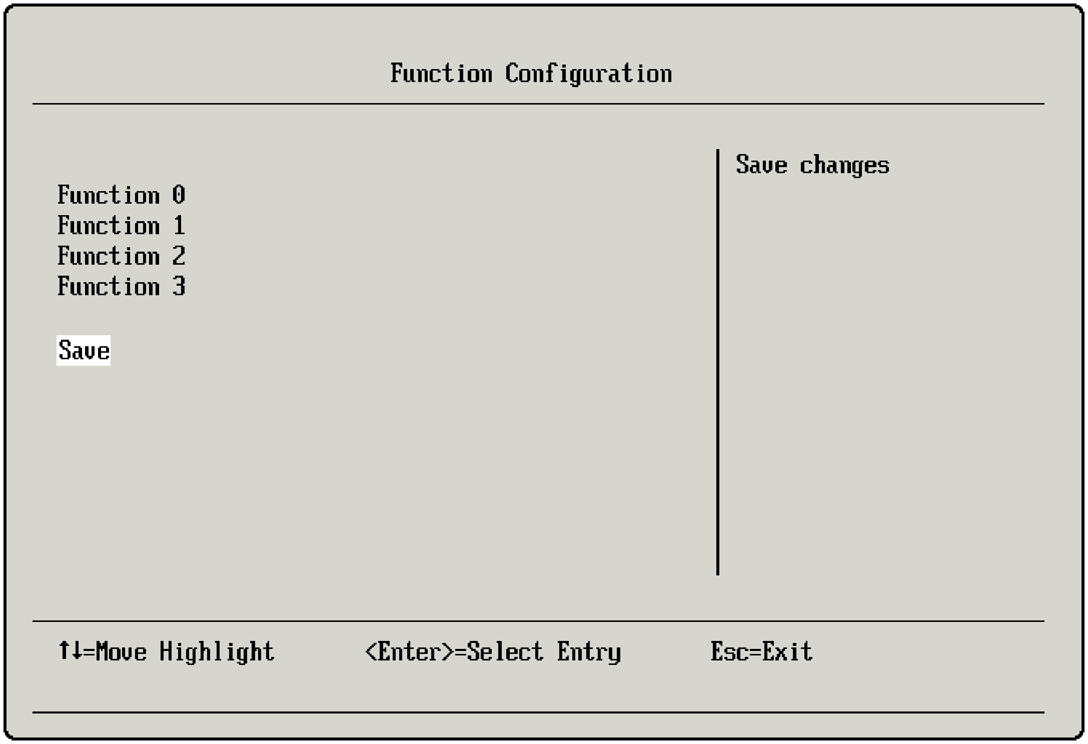

Figure 10-42 Save settings before exiting

9. Repeat the steps in this task to configure vNIC functions on the second 10 Gbps physical port shown in Figure 10-35 on page 276.
10. Back in the System Configuration and Boot Management panel (Figure 10-3 on page 244), save your settings before exiting the UEFI utility.

vNICs in the operating system

The HS22 blade in bay 4 runs SUSE Linux Enterprise Server 11 (SLES 11). Now you see how the vNICs are displayed to the operating system. The Network

Settings panel (Figure 10-43) lists eight vNICs, four on each 10 Gbps physical port.

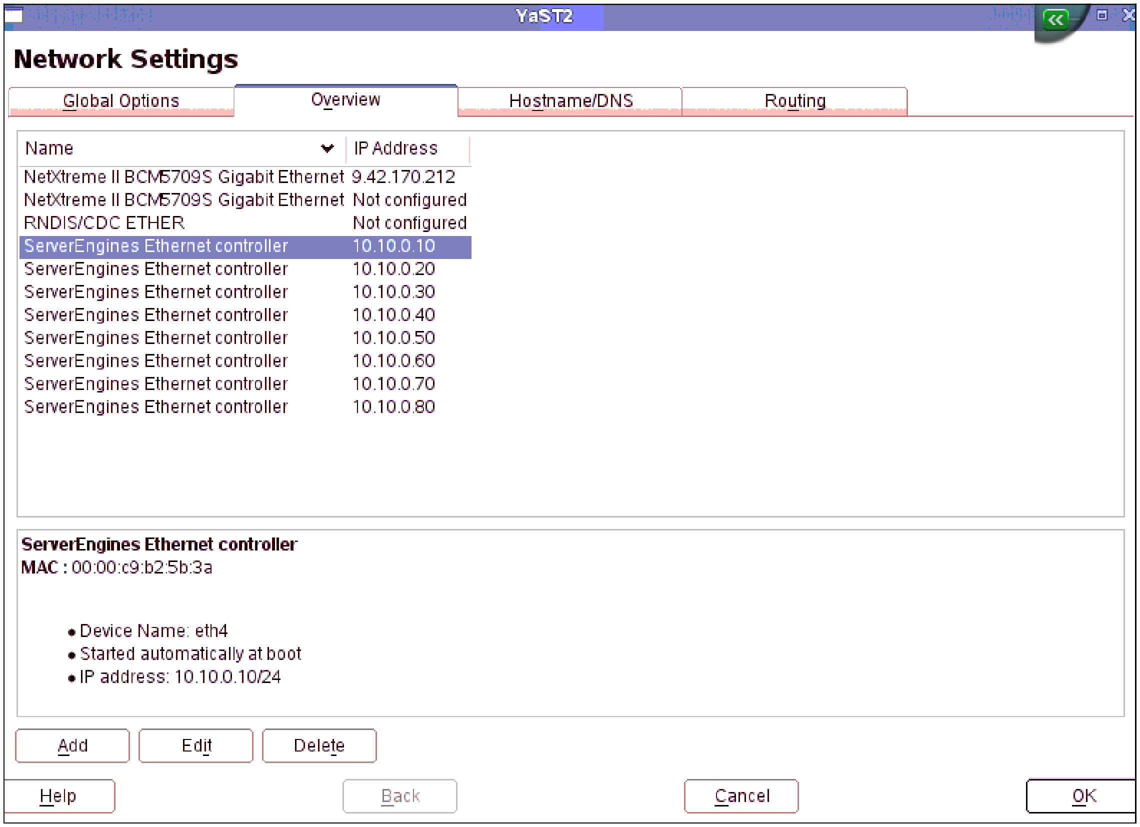

Figure 10-43 Network Settings

Tip: At the time of writing this book, Linux and VMware both reported the Emulex 10GbE Virtual Fabric Adapter II adapter as having a ServerEngines chip.

For clarity, we changed the names of the NICs from the default values to those listed in Table 10-8. We used the MAC addresses to match each vNIC with the adapter as displayed on the Network Connections panel.

Table 10-8 NIC names and matching vNICs

NIC name	10 Gbps physical port	vNIC	MAC address
eth4	0	1	00:00:C9:B2:5B:3A
eth6	0	2	00:00:C9:B2:5B:3B
eth8	0	3	00:00:C9:B2:5B:3C
eth10	0	4	00:00:C9:B2:5B:3D
eth5	1	1	00:00:C9:B2:5B:3E
eth7	1	2	00:00:C9:B2:5B:3F
eth9	1	3	00:00:C9:B2:5B:40
eth11	1	4	00:00:C9:B2:5B:41

To determine the MAC addresses, you can use OneCommand Manager, as explained in the following section, or the **`ifconfig`** command (see “The ifconfig command” on page 287).

OneCommand Manager

You can use OneCommand Manager to verify parameters such as MAC addresses and link speed. Figure 10-44 shows the status panel for NIC eth4.

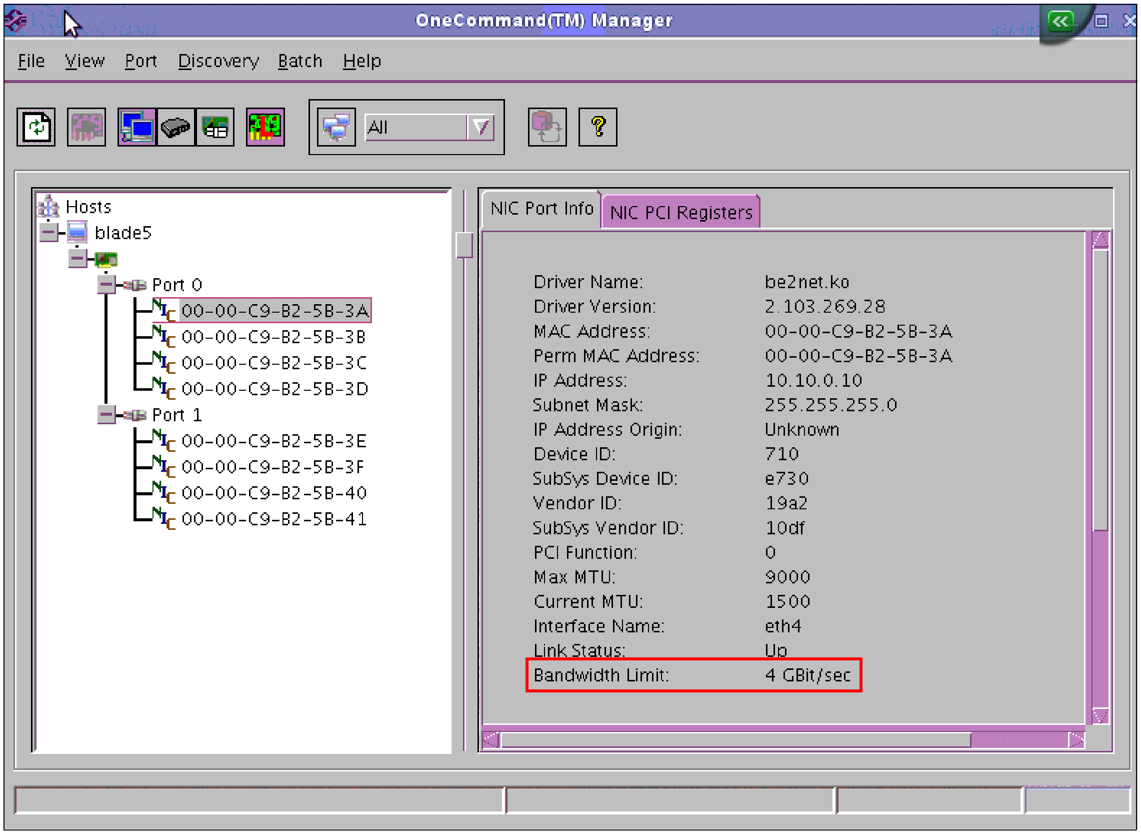

Figure 10-44 Eth4 status

The speed is reported as 4.0 Gbps (see Figure 10-13 on page 254) because we set the bandwidth to 40 on this vNIC in step 7 on page 278.

The ifconfig command

Another method to display the MAC addresses is to use the **ifconfig** command. Figure 10-45 shows the results of the **ifconfig** command. The MAC addresses are displayed in the Hwaddr field.

```
eth10     Link encap:Ethernet  HWaddr 00:00:C9:B2:5B:3D
          inet addr:10.10.0.70  Bcast:10.10.0.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
          Memory:95a80000-95aa0000

eth11     Link encap:Ethernet  HWaddr 00:00:C9:B2:5B:41
          inet addr:10.10.0.80  Bcast:10.10.0.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:0 errors:0 dropped:0 overruns:0 frame:0
          TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
          Memory:95ac0000-95ae0000
```

Figure 10-45 Results of the ifconfig command for viewing MAC addresses

We confirmed that this MAC address belongs to vNIC 4 on 10 Gbps physical port 0 (eth10) and vNIC 4 on 10 Gbps physical port 1 (eth11). You can identify and verify the MAC addresses and link speeds of all other vNICs in the same manner.

10.3.2 Configuring the Broadcom adapter

The Broadcom 2-port 10Gb Virtual Fabric Adapter is set up in a similar manner as the Emulex 10GbE Virtual Fabric Adapter II. The difference is that you do not specify a VLAN ID for the Broadcom adapter. The Broadcom adapter uses only MAC addresses for virtual channel traffic separation. Table 10-9 lists the settings for the Broadcom 2-port 10Gb Virtual Fabric Adapter.

Table 10-9 Broadcom Virtual Fabric Adapter vNIC configuration

NIC partition	Bandwidth	MAC address
1	4 Gbps	10:18:00:96:CE:90
2	3 Gbps	10:18:00:96:CE:94
3	2 Gbps	10:18:00:96:CE:98

NIC partition	Bandwidth	MAC address
4	1 Gbps	10:18:00:96:CE:9C

You can identify the MAC addresses for each NIC partition by using the UEFI utility (see Figure 10-19 on page 261). It is important to determine the MAC addresses so that you can identify each NIC partition in the operating system.

To configure the Broadcom adapter parameters by using the UEFI utility:

1. Power on or restart the blade server.
2. When prompted during POST, press F1.
3. In the System Configuration and Boot Management panel (Figure 10-3 on page 244) of the UEFI utility, select **System Settings**.
4. In the System Settings panel, select **Network**.
5. In the Network configuration panel (Figure 10-46), select the first of the two **Broadcom NetXtreme II 10 Gigabit Ethernet** entries, and press Enter.

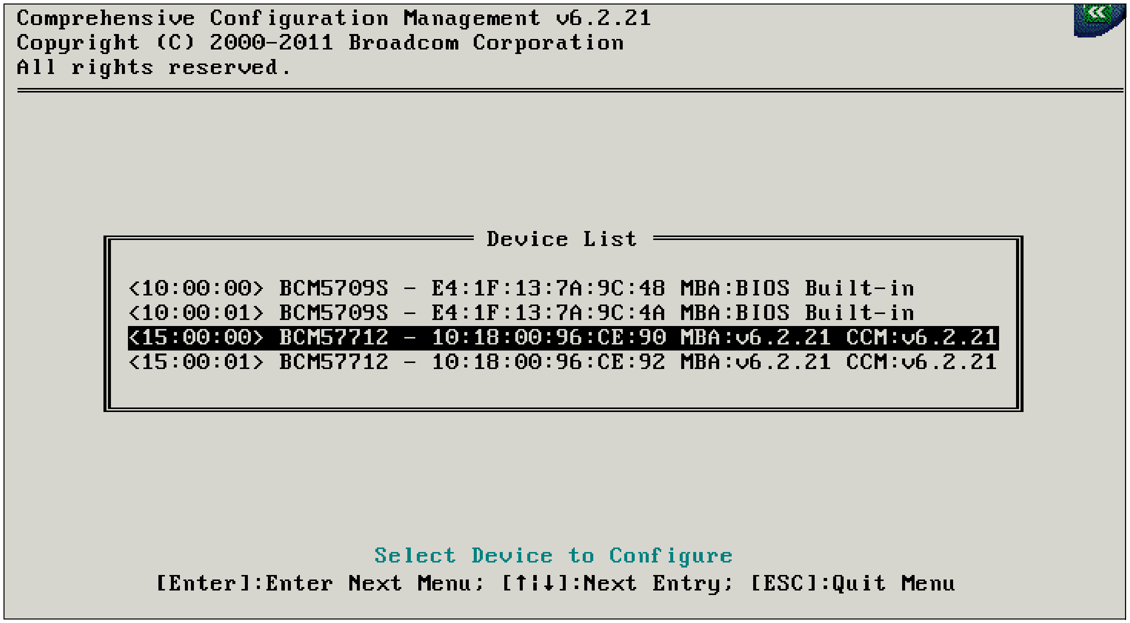

Figure 10-46 Network configuration panel

6. In the Broadcom Main Menu panel (Figure 10-47), select **NIC Partition Configuration Menu**, and press Enter.

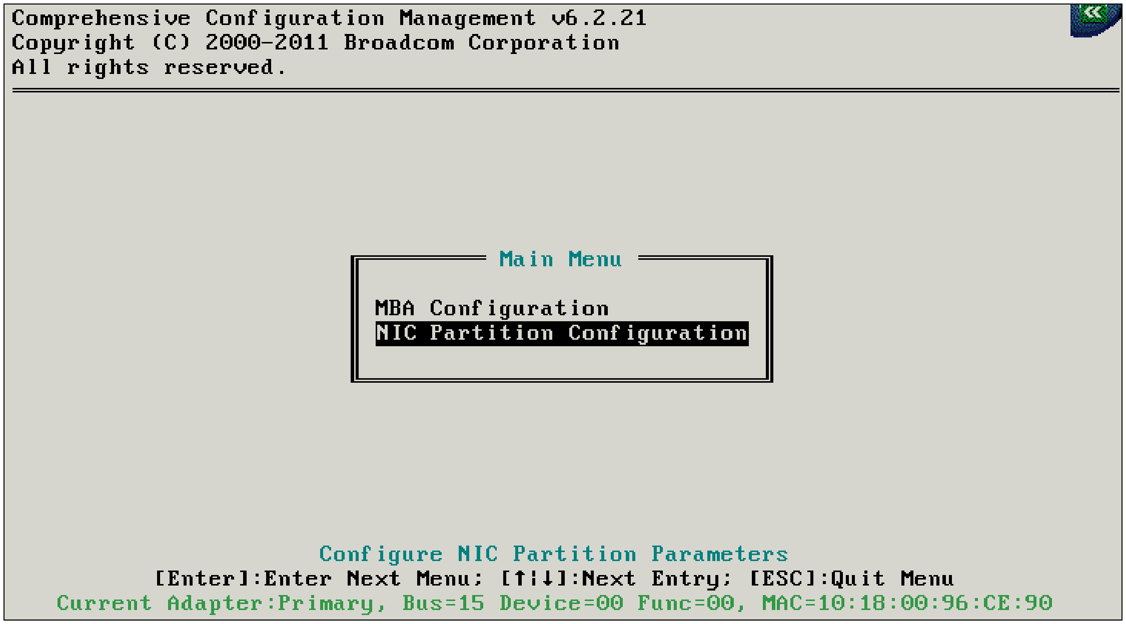

Figure 10-47   Broadcom Main Menu panel

7. In the NIC Partition Configuration panel (Figure 10-48), ensure that NIC Partition is set to **Enabled**. Then press Esc.

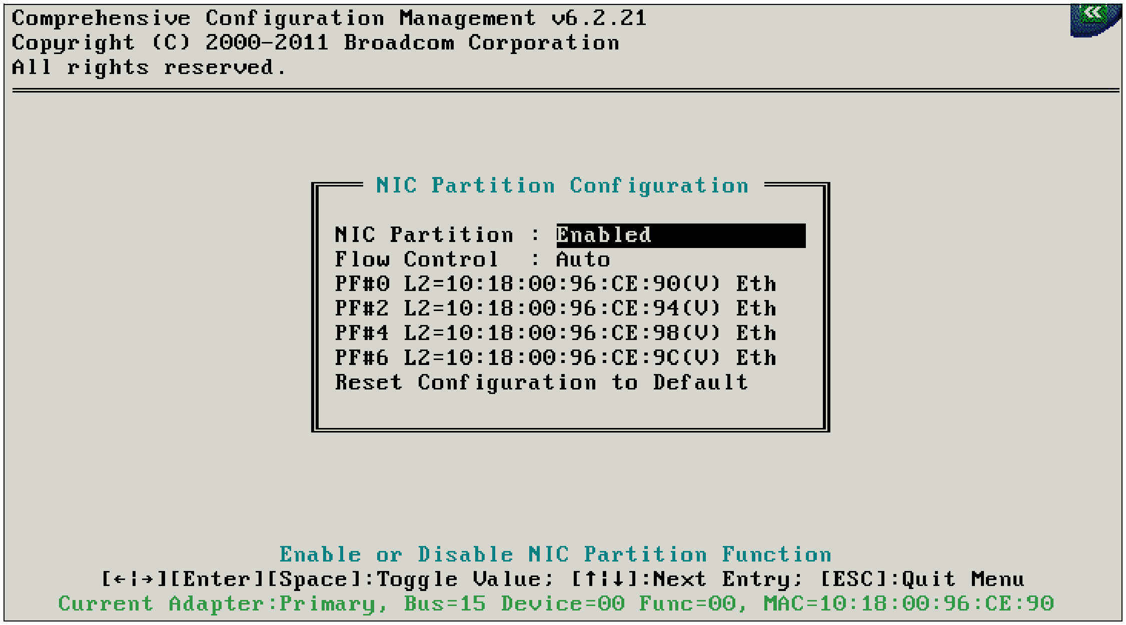

Figure 10-48 Device Configuration panel with NIC Partition set to Enabled

8. In the configuration panel for PF# 0 (Figure 10-49), a new option is available now, NIC Partitioning Configuration Menu. Highlight it and press Enter.

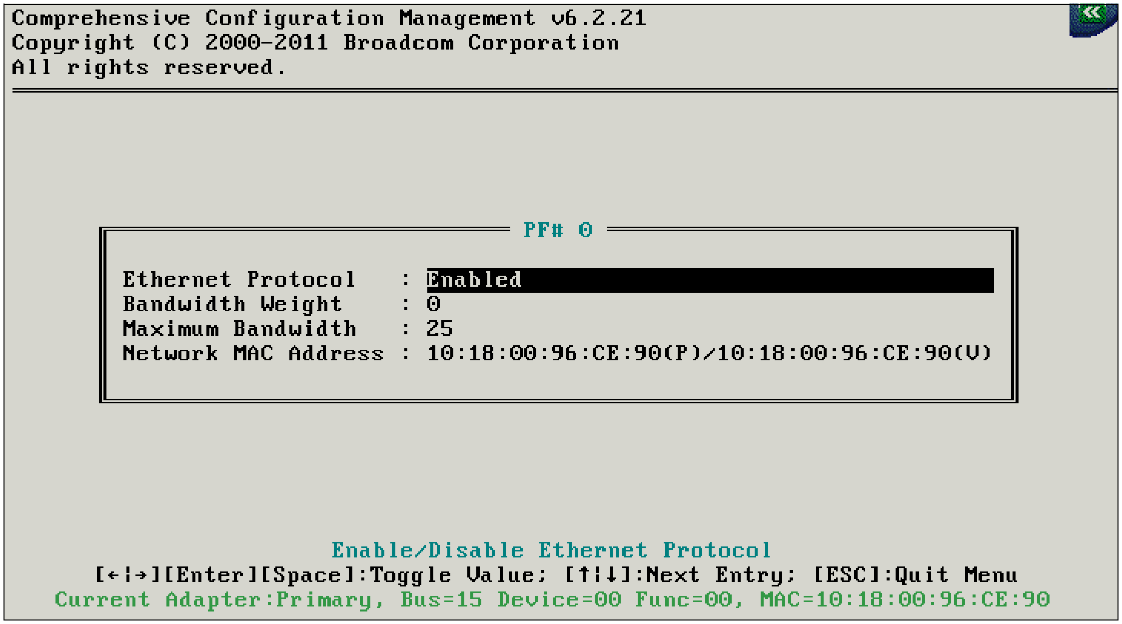

Figure 10-49 New configuration option: the NIC Partitioning Configuration Menu

9. In the NIC Partition Configuration panel (Figure 10-50), identify the MAC addresses of all four NIC partitions on this panel:
   - PF#0 represents the first NIC partition, `L2=10180096CE90`, which means the MAC address is 10:18:00:96:CE:90.
   - PF#2 is the second NIC partition, `L2=10180096CE94`, which means the MAC address is 10:18:00:96:CE:94.
   - PF#4 is the third NIC partition, `L2=10180096CE98`, which means the MAC address is 10:18:00:96:CE:98.
   - PF#6 is the fourth NIC partition, `L2=10180096CE9C`, which means the MAC address is 10:18:00:96:CE:9C.

   It is important to know these MAC addresses so that you can match the NICs presented in the operating system to NIC partitions.

   Select **Global Bandwidth Allocation Menu**, and press Enter.

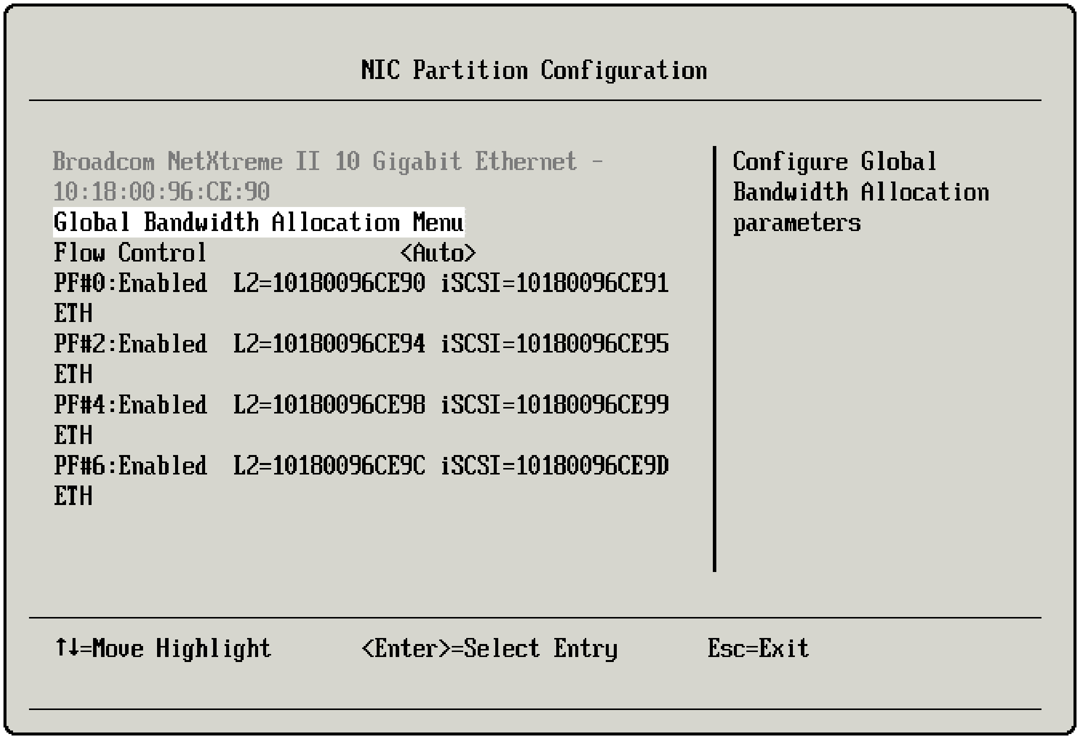

Figure 10-50 NIC Partition Configuration panel

10. In the Global Bandwidth Allocation Menu panel (Figure 10-51), set the bandwidth allocation parameters in this panel. In this example, we use the following configuration:

   a. Because we do not use a relative bandwidth weight, leave the Relative Bandwidth Weight fields set to 0.

   b. Configure maximum bandwidth as follows:

      i. Set NIC partition 1 to 4 Gbps (40% of total bandwidth).
      ii. Set NIC partition 2 to 3 Gbps (30% of total bandwidth).
      iii. Set NIC partition 3 to 2 Gbps (20% of total bandwidth).
      iv. Set NIC partition 4 to 1 Gbps (10% of total bandwidth).

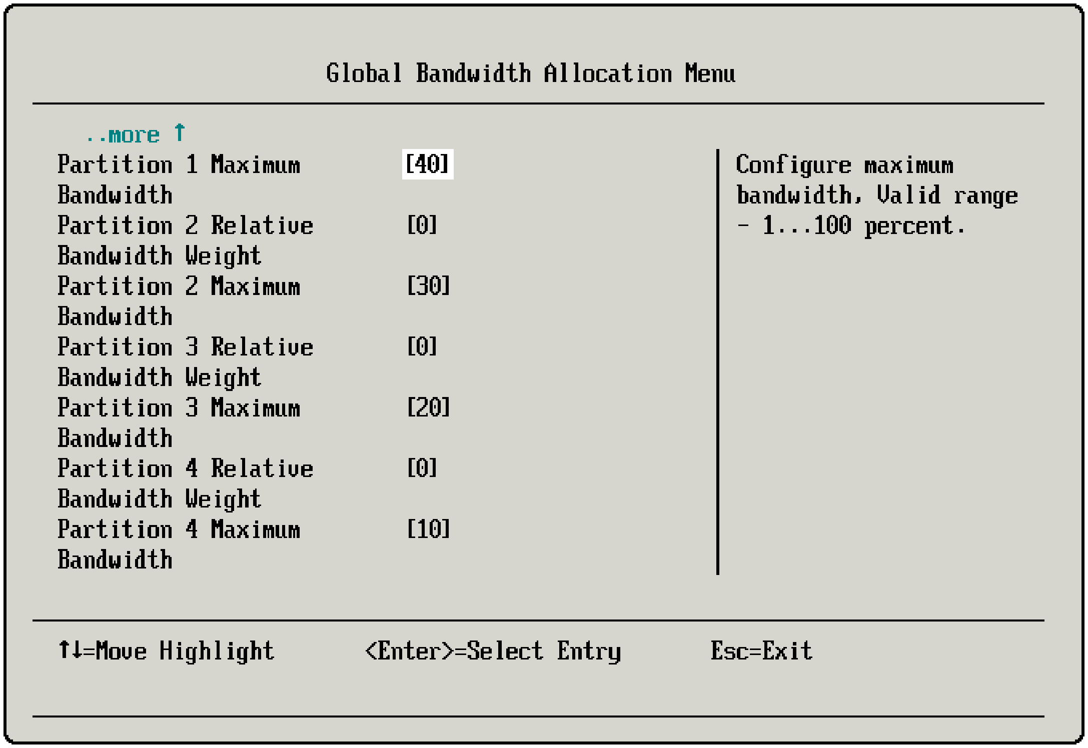

Figure 10-51 Global Bandwidth Allocation Menu panel

11. Repeat the steps, starting at step 6 on page 289, to configure NIC partitions and bandwidth parameters for the second 10 Gbps physical port.

12. Back in the System Configuration and Boot Management panel (Figure 10-3 on page 244), save your settings before exiting the UEFI utility.

NIC partitions in the operating system

You can define how the NIC partitions are displayed to the operating system. Blade 7 runs on Windows Server 2008. This section highlights the Windows built-in networking tools and the Broadcom Advanced Control Suite.

Windows networking tools

Figure 10-52 shows the NIC partitions in the Network and Sharing Center.

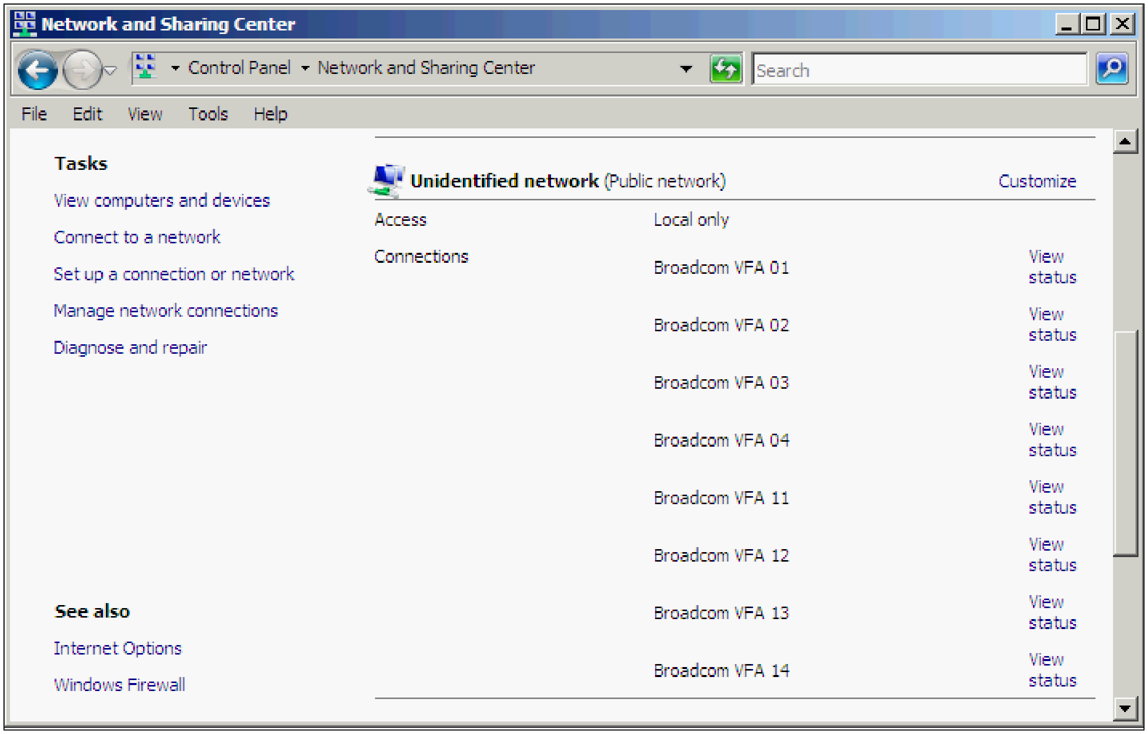

Figure 10-52 Broadcom Virtual Fabric Adapter NIC partitions in Windows Server 2008

For clarity, we changed the NIC names from the defaults as shown in Table 10-10. We used MAC addresses to match each NIC with the actual NIC partition.

Table 10-10 NIC names and matching NIC partitions

NIC name	10 Gbps physical port	NIC partition	MAC address
Broadcom VFA 01	0	1	10:18:00:96:CE:90
Broadcom VFA 02	0	2	10:18:00:96:CE:94
Broadcom VFA 03	0	3	10:18:00:96:CE:98
Broadcom VFA 04	0	4	10:18:00:96:CE:9C
Broadcom VFA 11	1	1	10:18:00:96:CE:92
Broadcom VFA 12	1	2	10:18:00:96:CE:96
Broadcom VFA 13	1	3	10:18:00:96:CE:9A
Broadcom VFA 14	1	4	10:18:00:96:CE:9E

In Windows, complete the following steps for each partition in Table 10-10 on page 295 to view for correctness:

1. In the Network and Sharing Center window (Figure 10-52 on page 294), click **View Status** for each partition.
2. In the Status window (Figure 10-53 on page 296), verify the parameters, such as link speed and MAC address. Figure 10-53 shows the parameters of the NIC Broadcom VFA 01. Speed is reported as 4.0 Gbps, because we set the bandwidth to 40 on this NIC partition. Click **Details**.

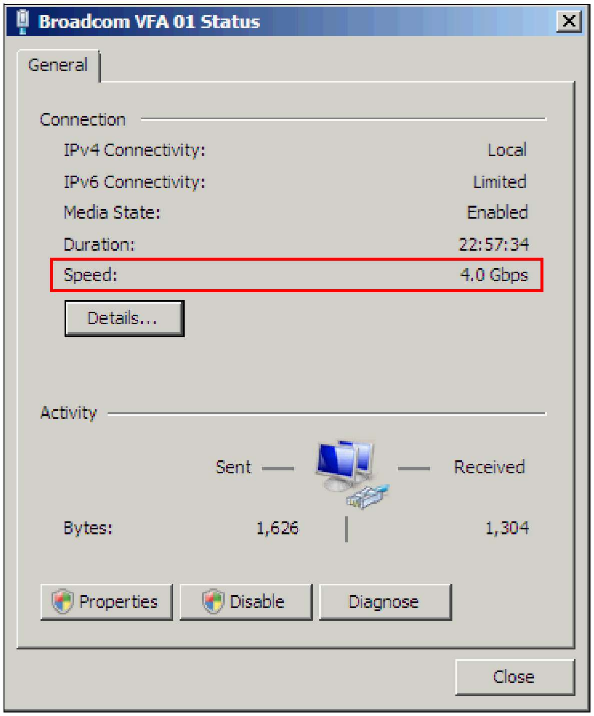

Figure 10-53 Broadcom VFA 01 status

3. In the Network Connection Details window (Figure 10-54), review the MAC address of the selected NIC partition. Then click **Close**.

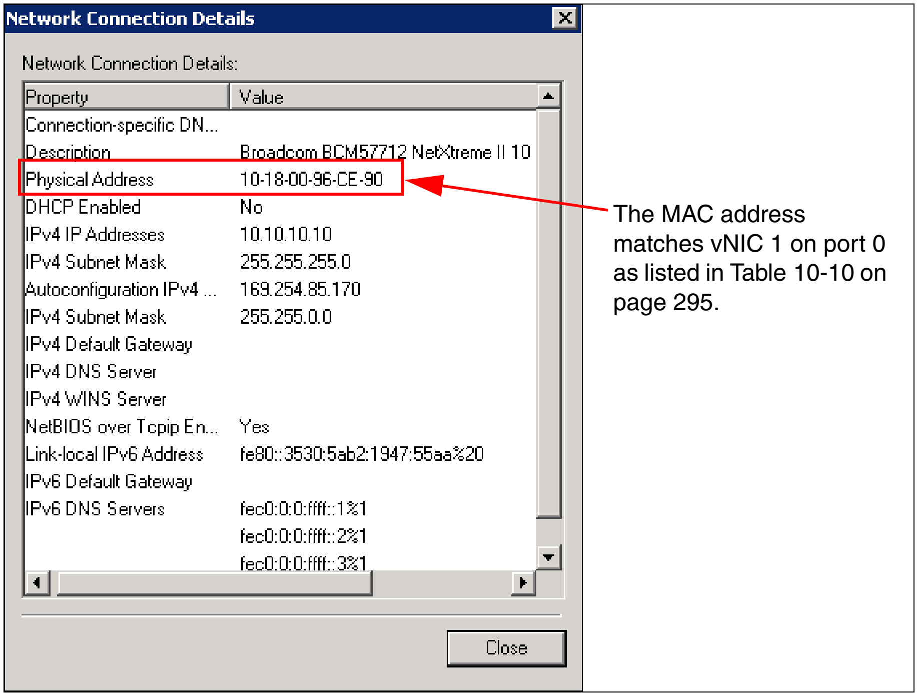

Figure 10-54 Network Connection Details panel for a selected NIC partition

We can confirm from Figure 10-54 that this address is the MAC address of NIC partition 1 (the Broadcom DCM57712 NeXtreme II) on 10 Gbps physical port 0.

4. Repeat steps 1 on page 296 through 3 on page 297 to verify the MAC address and speed for the other NIC partitions on the Broadcom Virtual Fabric Adapter.

Setting VLAN IDs in the operating system in Windows

You can define VLAN IDs for each NIC partition in the operating system. For example, to define VLAN ID 10 for NIC Broadcom VFA 01:

1. In the Windows Device Manager window, open the Broadcom VFA 01 adapter.
2. In the Broadcom FVFA 01 Properties window (Figure 10-55), click **Configure** to configure the vNIC.

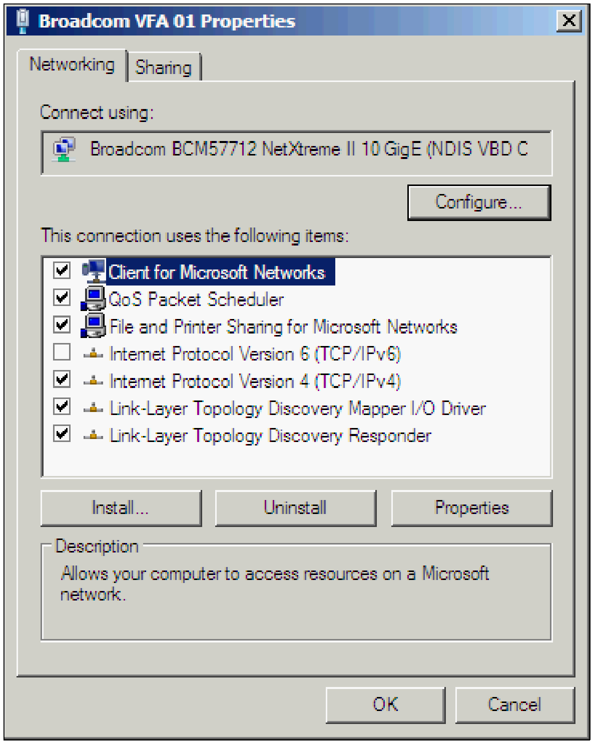

Figure 10-55 Broadcom VFA 01 properties window

3. On the **Advanced** tab (Figure 10-56), in the Property box, scroll down to and select **VLAN ID**. Set the VLAN ID value to **10**. Click **OK**.

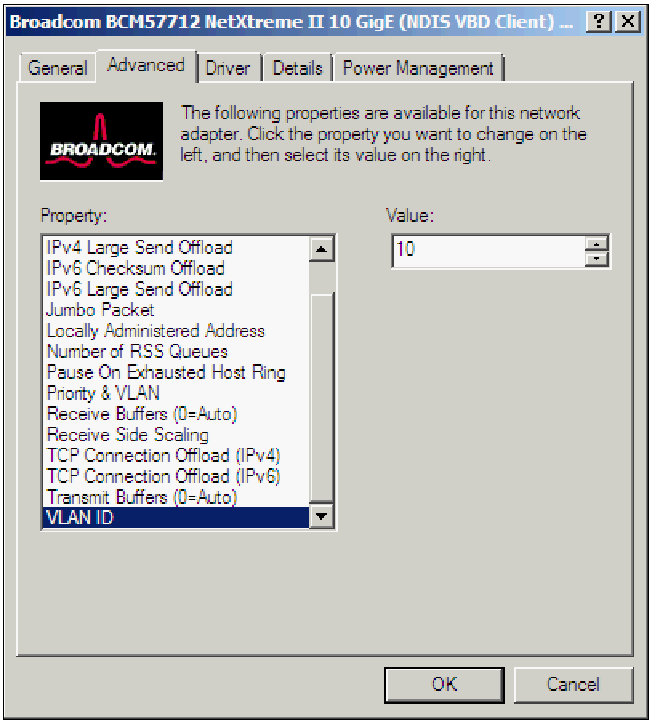

Figure 10-56 VLAN ID configuration

4. Repeat step 1 on page 298 to step 3 to configure the remaining VLAN IDs for the other NIC partitions on 10 Gbps physical port 0, as follows:
   – Set Broadcom VFA 02 to LAN ID 20.
   – Set Broadcom VFA 03 to VLAN ID 30.
   – Set Broadcom VFA 04 to VLAN ID 40.

Broadcom Advanced Control Suite

By using the Broadcom Advanced Control Suite management utility, you can administer, configure, and diagnose the Broadcom 2-port 10Gb Virtual Fabric Adapter and other Broadcom NICs.

Figure 10-57 shows an example of Broadcom Advanced Control Suite showing eight NIC partitions in the left pane. NIC partition 1 on physical port 0 is highlighted. The speed is 4 Gbps, which was set in the bandwidth parameter in the UEFI utility.

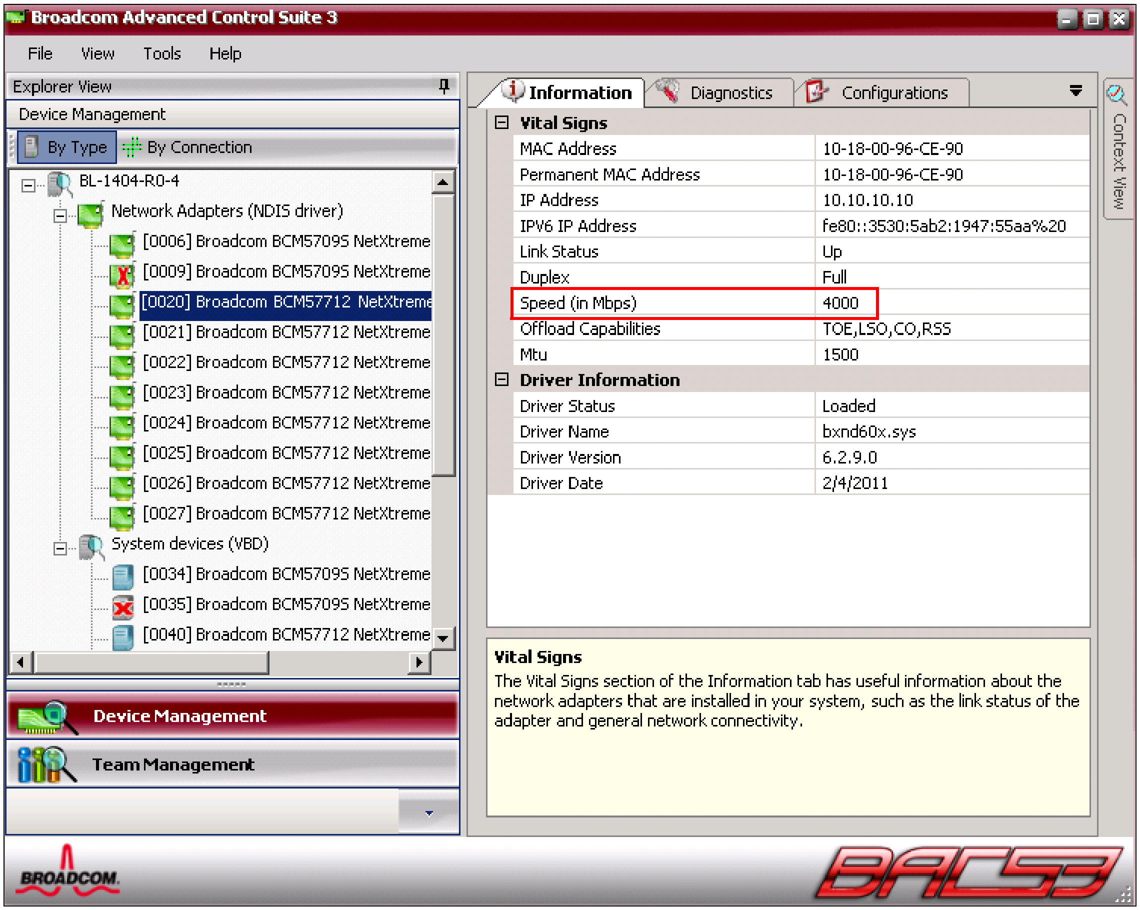

Figure 10-57 NIC partition 1 in Broadcom Advanced Control Suite

On the **Configurations** tab (Figure 10-58), you can check and set NIC parameters. The VLAN ID is set to 10 as shown in Figure 10-25 on page 267. Use the Broadcom Advanced Control Suite utility to configure the VLAN ID to a different value or to change any other parameter.

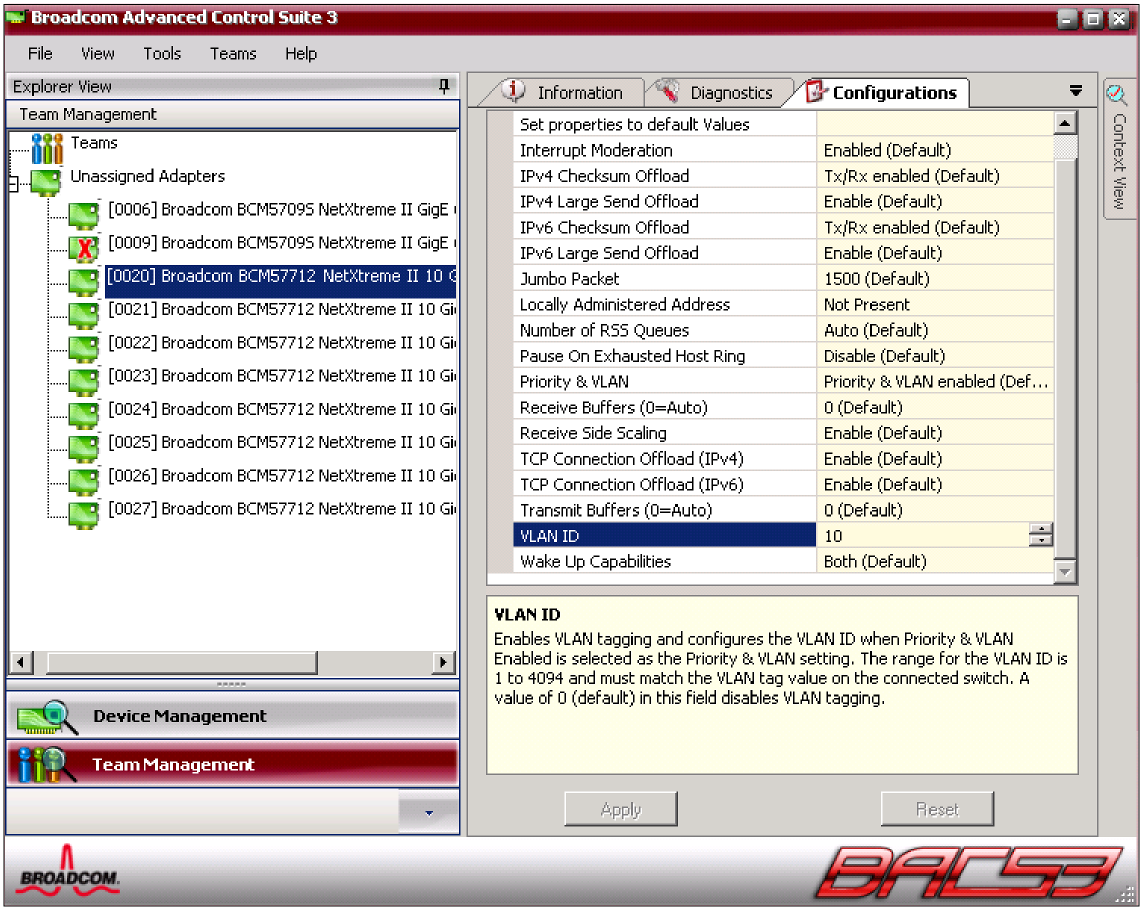

Figure 10-58 Configuration tab in Broadcom Advanced Control Suite

On the **Diagnostics** tab (Figure 10-59), you can test a configuration by running a basic connectivity test by pinging the IP gateway or a user-specified IP address.

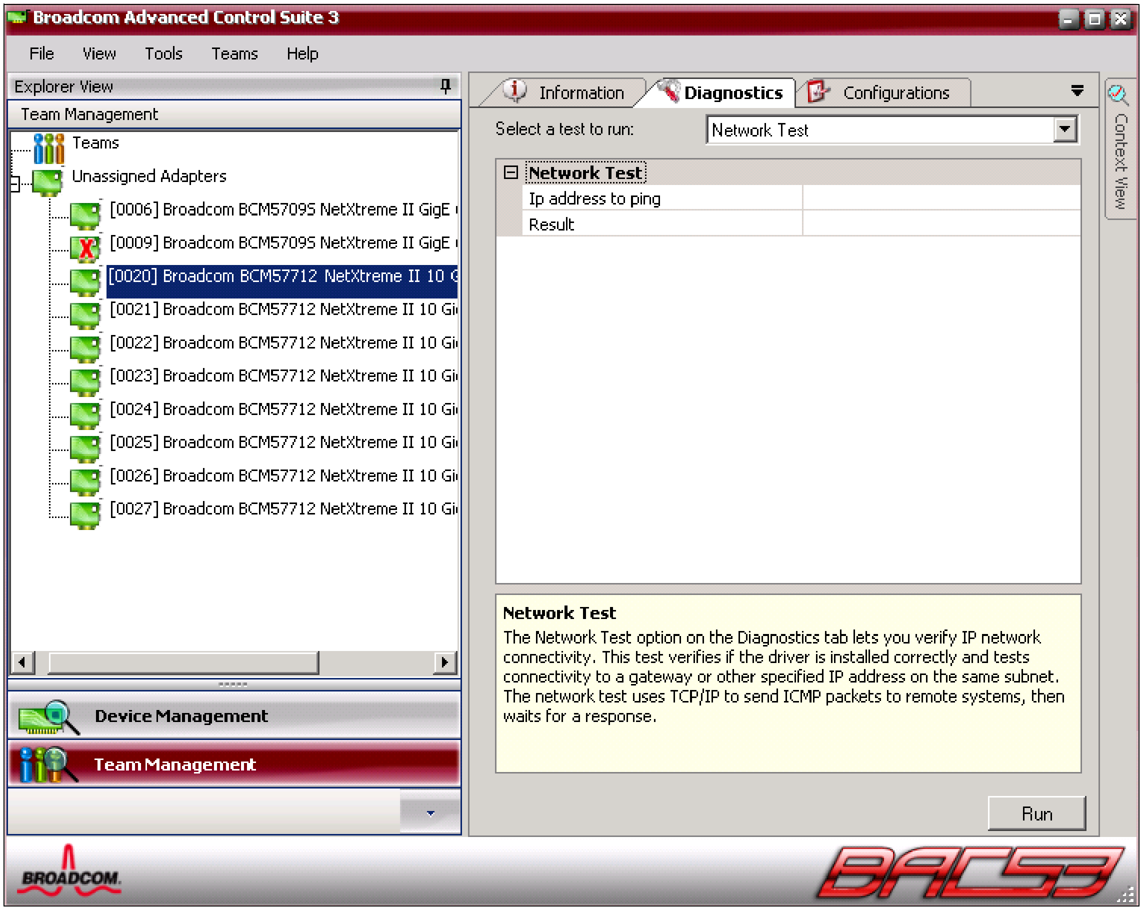

Figure 10-59 Diagnostics tab in Broadcom Advanced Control Suite

The **Statistics** tab (Figure 10-60) shows the network traffic counters, which are useful when monitoring or troubleshooting NIC behavior.

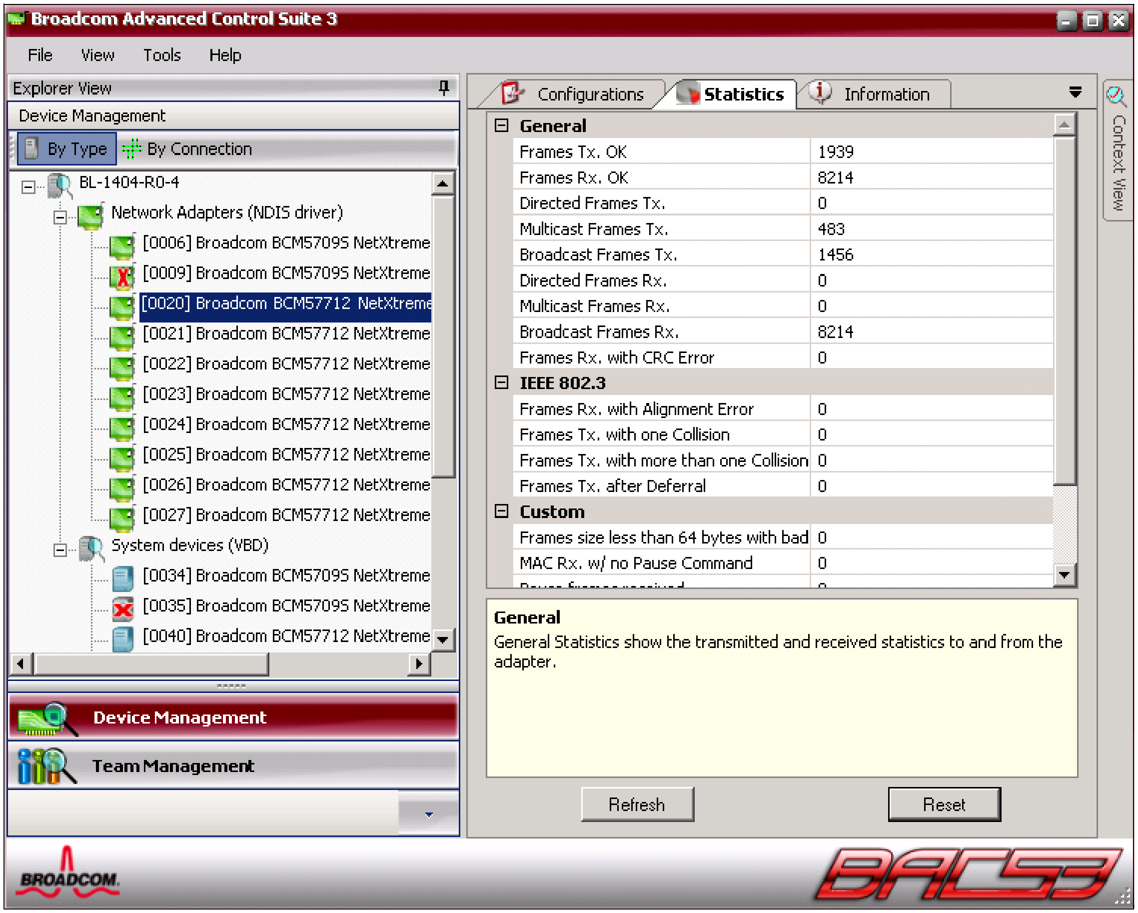

Figure 10-60 Statistics tab in the Broadcom Advanced Control Suite

NIC Teaming by using Broadcom Advanced Control Suite

One of the main functions of the Broadcom Advanced Control Suite utility is to configure NIC teaming. You use NIC partitions in teaming configurations in the same way as with physical NICs.

By clicking **Team Management** in the left pane, you can start the Teaming Wizard. For more information, see 9.7.4, “Configuring teams with Broadcom Advanced Control Suite” on page 229.

10.3.3 Configuring the BNT Virtual Fabric 10Gb Switch Module

You are now ready to configure VLANs on the BNT Virtual Fabric 10Gb Switch Module. The HS22 blade server in slot 4 contains Emulex 10GbE Virtual Fabric Adapter II, which is configured to use VLANs 10, 20, 30, and 40. Internal (blade-facing) port 4 (INT4) on the BNT switch module needs to pass traffic in each of these VLANs. Therefore, VLAN tagging must be enabled on this port.

The HS22 blade server in slot 8 contains Broadcom Virtual Fabric Adapter, configured to use VLANs 10, 20, 30, and 50. Internal (blade-facing) port 8 (INT8) on the BNT switch module needs to pass traffic in each of these VLANs. Therefore, VLAN tagging must be enabled on this port.

External port 2 (EXT2) must pass traffic in VLANs 10 and 40, and external port 4 (EXT4) must pass traffic in VLANs 20 and 50. Both EXT2 and EXT4 must have tagging enabled.

By default, VLAN tagging is enabled on internal ports, but disabled on external ports. Therefore, you must enable tagging on ports EXT2 and EXT4. Example 10-1 shows the industry-standard command-line interface (isCLI) commands to enable this tagging.

Example 10-1 Enabling tagging on EXT2 and EXT4

```
interface port EXT2
tagging
exit
interface port EXT4
tagging
exit
```

The VLAN structure is defined in Table 10-2 on page 241. Example 10-2 shows the sequence of isCLI commands to configure these VLANs.

Example 10-2 VLAN definitions on BNT Virtual Fabric Switch Module

```
vlan 10
enable
name "VLAN 10"
member INT4
member INT8
member EXT2

vlan 20
enable
name "VLAN 20"
member INT4
member INT8
```

```
member EXT4

vlan 30
enable
name "VLAN 30"
member INT4
member INT8

vlan 40
enable
name "VLAN 40"
member INT4
member EXT2

vlan 50
enable
name "VLAN 50"
member INT8
member EXT4
```

You have now configured the BNT Virtual Fabric Switch Module.

10.3.4 Configuring Switch Independent Mode for Emulex II, BNT, QLogic, and IBM Storwize V7000

The Emulex Virtual Fabric Adapters II improve the virtual fabric solutions by providing more features and extended vNICs support for Fibre Channel over Ethernet (FCoE) and Internet Small Computer System Interface (iSCSI). Virtual Fabric Adapter II has the following features:

- Support for the Emulex OneCommand Manager application
- Support for Serial over LAN (OCm11102-N-X and OCm11102-N-XI Virtual Fabric Adapters only)
- Dual port 10 Gbps adapter supporting multiple virtual I/O functions
- vNIC

  Physical ports can be divided into a maximum of eight virtual NICs per adapter with bandwidth allocation in 100 Mbps increments. Each virtual NIC is displayed as an individual adapter to the operating system.
- FCoE functionality
- Support for Boot from Network by using iSCSI software or hardware initiator
- Support for FCoE boot from SAN
- Support for hardware licensing

Enabling the Virtual Fabric Adapter Switch Independent Mode function on the UEFI utility

To enable the Virtual Fabric Adapter Switch Independent Mode function:

1. Turn on the system.
2. When prompted, press F1 to enter UEFI setup.
3. In the System Configuration and Boot Management panel (Figure 10-3 on page 244), click **System Settings**.
4. In the System Settings panel (Figure 10-61), choose **Emulex Configuration Utility Ver:5.01a8**, then press Enter.

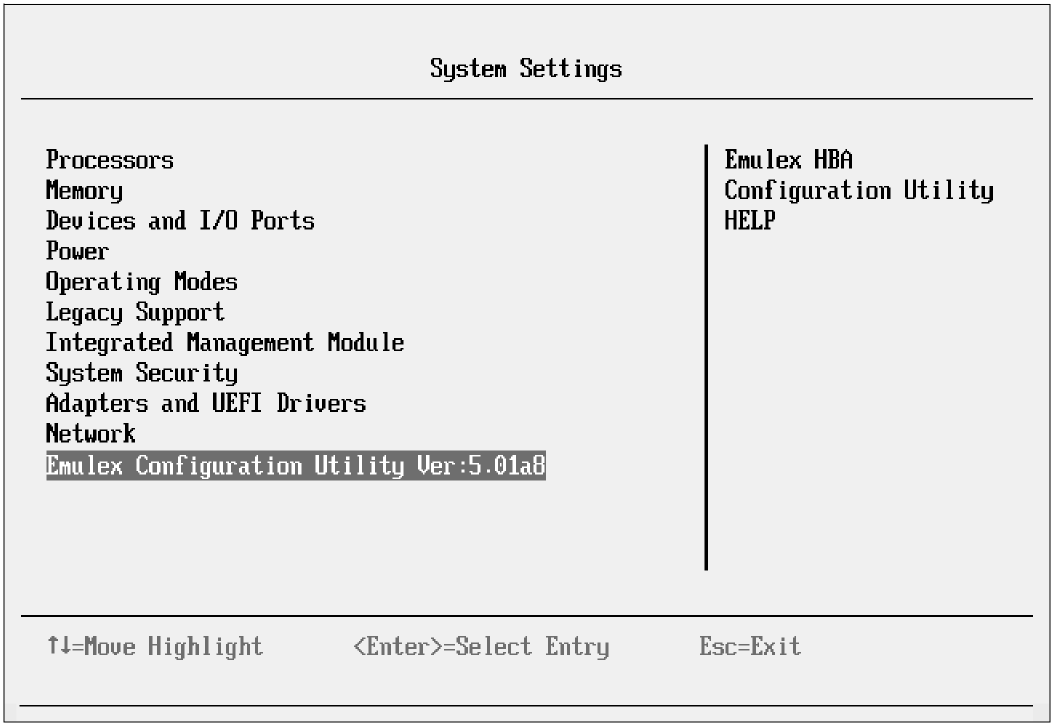

Figure 10-61 System Settings panel

5. In the Adapter Selection panel (Figure 10-62), select **001** for port one.

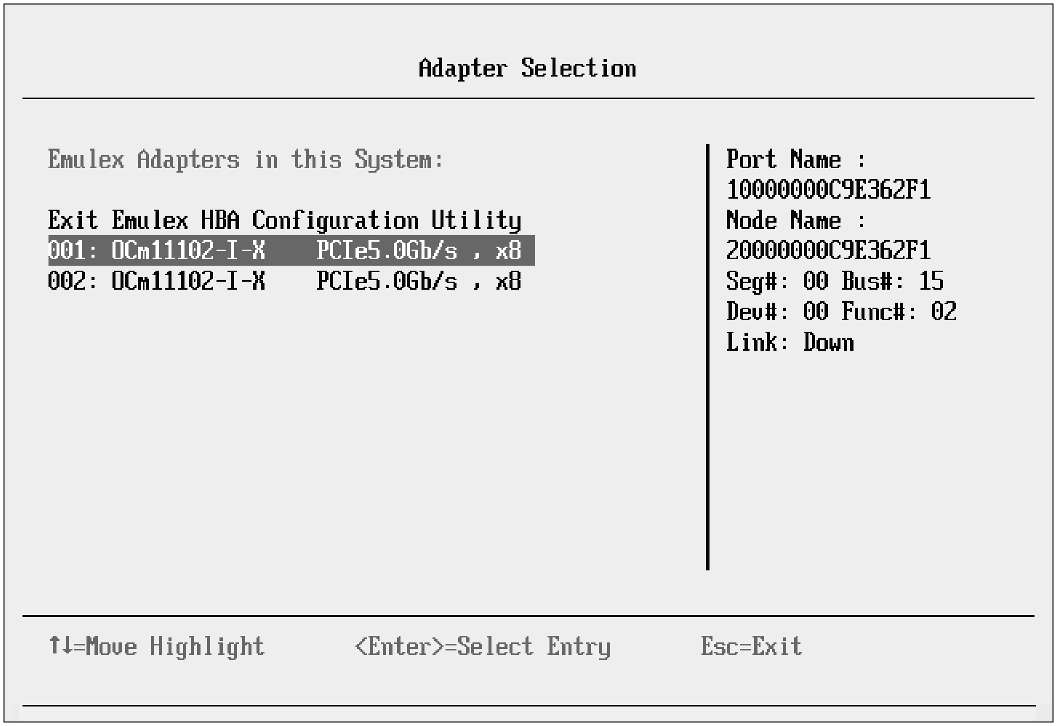

Figure 10-62 Adapter Selection panel

6. In the Emulex Adapter Configuration Main Menu panel (Figure 10-63), for Configure DCBX Mode, select **CEE**. Press Esc to exit.

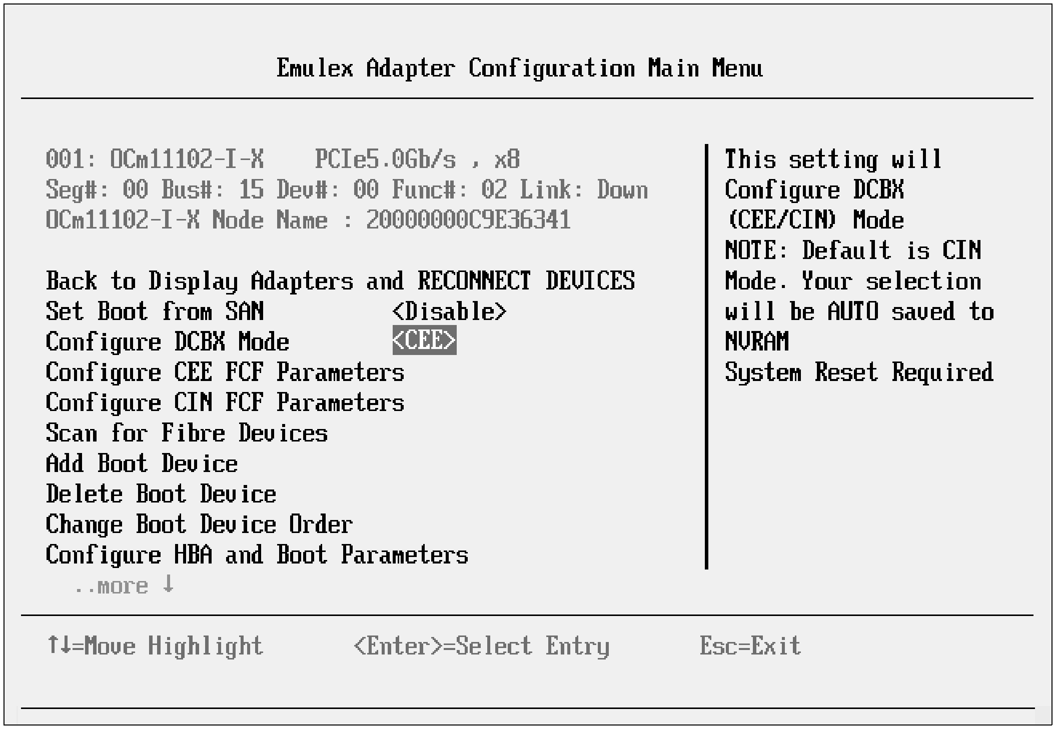

Figure 10-63 Emulex Adapter Configuration Main Menu panel

7. Back in the System Configuration and Boot Management panel (Figure 10-3 on page 244), save your settings before exiting the UEFI utility.

Installing the Windows 2008 drivers and Management Emulex OCM

The following sections explain how to install the device drivers and Management Emulex OCM.

Installing the device drivers

To install the device driver:

1. Download the device driver from Emulex website at:

   `http://www.emulex.com/downloads/ibm/vfa-software-kits.html`

   For this example, we downloaded the `lxdrvr-fc-fcoe-V3.703.397.3806` FCoE device driver file and the `elxocm-windows-x86-5.1.42.4-1(32 bit)` Emulex OneCommand Manager driver file.

2. Run the `elxdrvr-fc-fcoe-V3.703.397.3806` FCoE device driver file.

> **Version tip:** At the time of the writing of this book, we used the `elxdrvr-fc-fcoe-V3.703.397.3806` version for Emulex OneCommand Manager. Always obtain the latest version for your configuration.

Click **Next**.

3. If you plan to use boot from SAN or to create a Windows Preinstallation Environment (PE) disk image, in the Installation options display (Figure 10-64), select the **Unpack All Drivers**. For this example, we did not select **Unpack All Drivers**. Click **Install**.

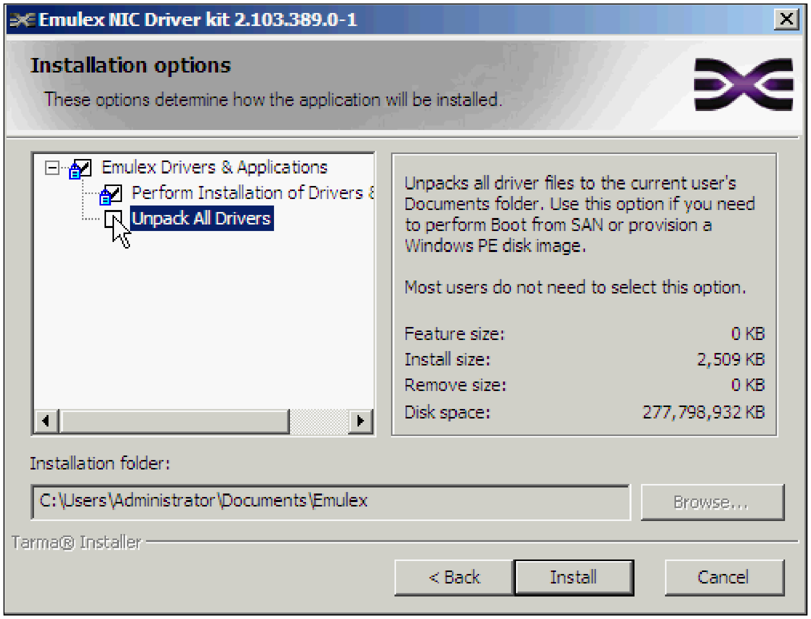

Figure 10-64 Installation options window

4. Click **Finish**.

5. After the AutoPilot Installer starts, when you see the list of installed host bus adapters (Figure 10-65), click **Next**.

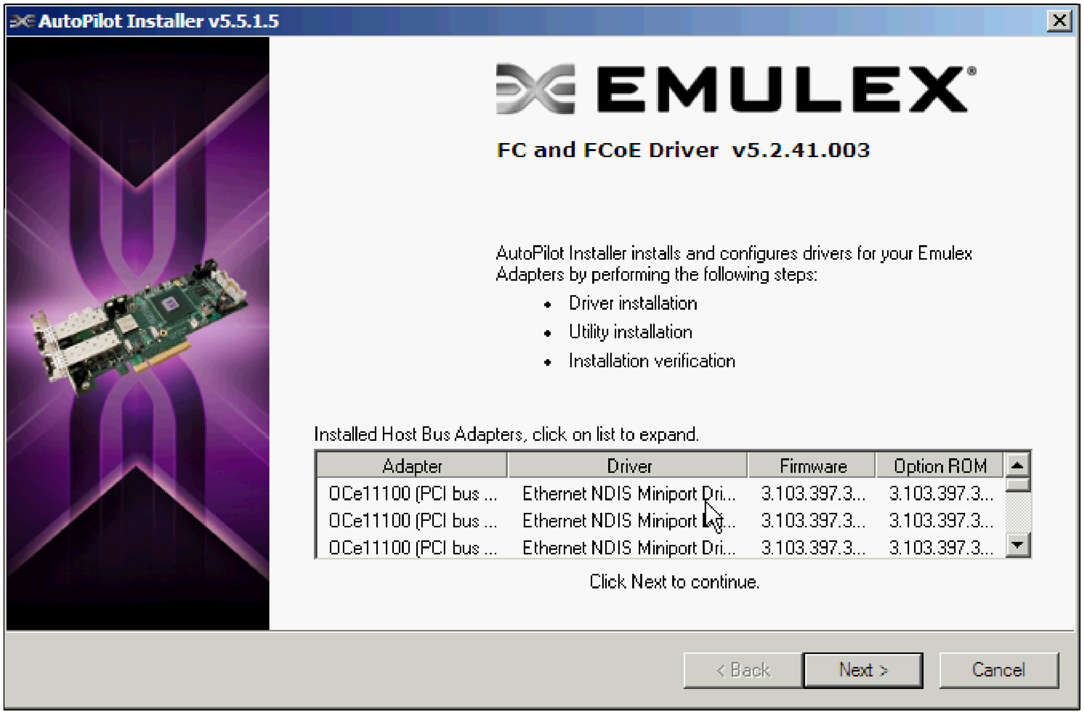

Figure 10-65 FC and FCoE Driver window

6. Click **Finish** to complete the installation.

Installing the Management Emulex OCM

To install the Emulex OCM:

1. Run the `elxocm-windows-x86-5.1.42.4-1(32 bit)` file.
2. In the first Emulex OCManager Enterprise window (Figure 10-66), click **Next**.

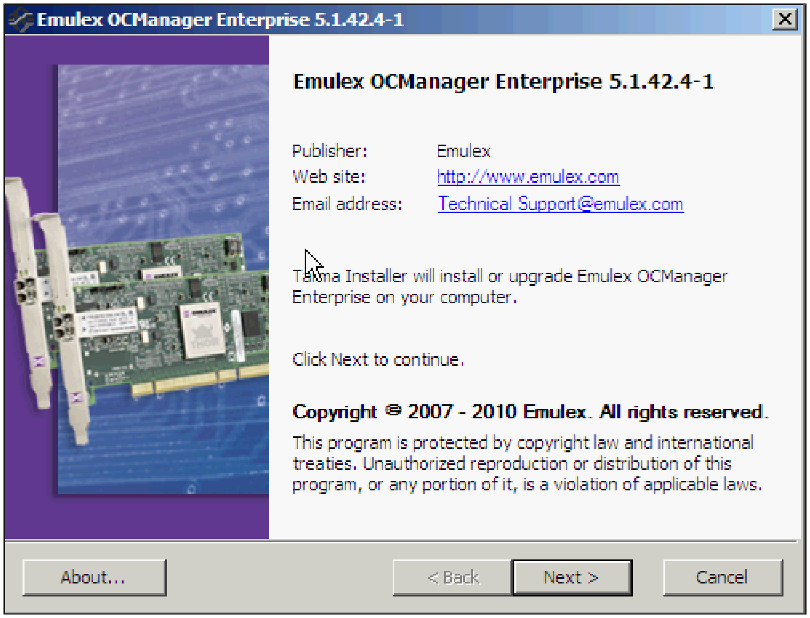

Figure 10-66 Emulex OCManager Enterprise window

3. In the Installation options display (Figure 10-67), accept the default options by leaving them selected and clicking **Install**.

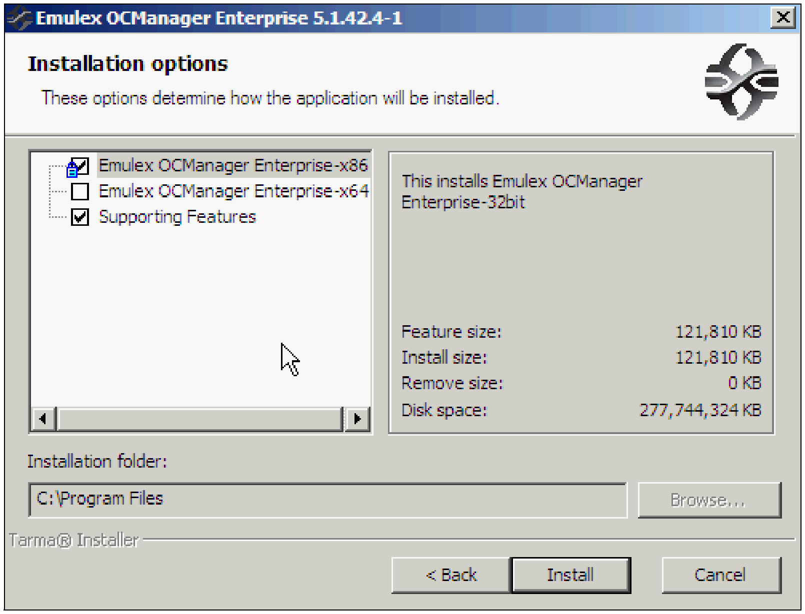

Figure 10-67 Installation options window

4. In the OneCommand Manager Management Mode window (Figure 10-68), select one of three modes operations:
   – Strictly Local Management (only management of local adapters)
   – Local Management Plus (only management of local adapters but allows management from other hosts)
   – Full Management (management of local and remote hosts and allows management from others hosts)

   In this example, we selected **Full Management**.

   Click **OK**.

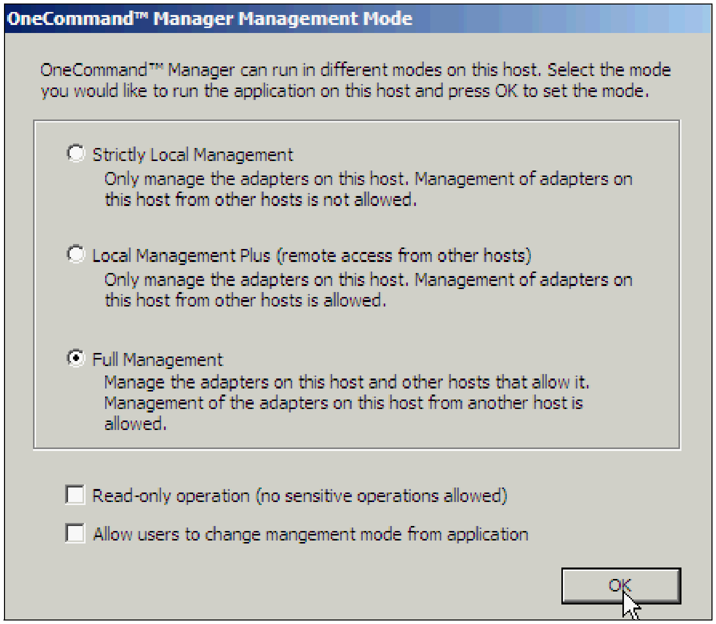

Figure 10-68 OneCommand Manager Management Mode window

Updating the Virtual Fabric Adapter firmware

Important: You must update the firmware to the latest levels.

To update the Virtual Fabric Adapter firmware:

1. Download the firmware file from Emulex support website at:

   `http://www.emulex.com/downloads/ibm/vfa-software-kits.html`

2. Double-click the file to open the OneCommand Manager application.
3. In the Emulex OCManager Enterprise window (Figure 10-69), click **Next**.

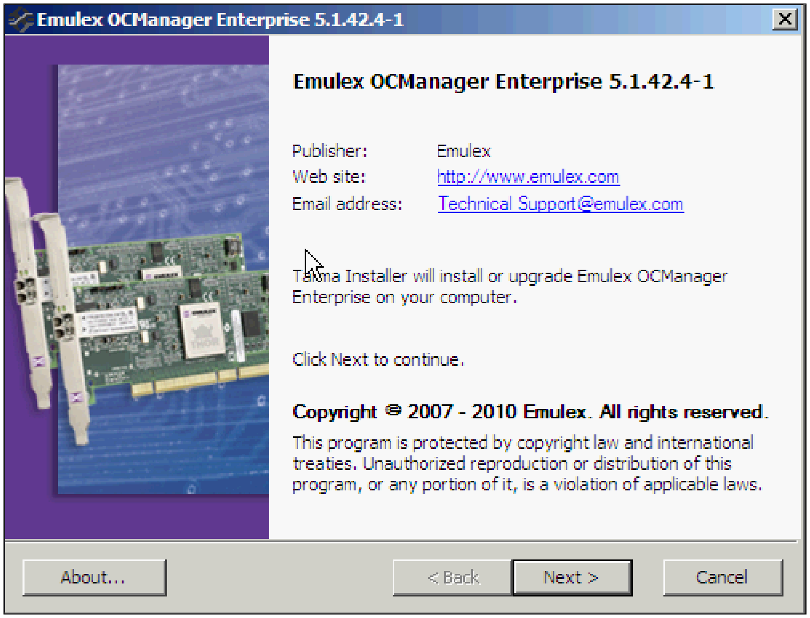

Figure 10-69 Emulex OCManager Enterprise 5.1.42.4-1 window

4. After OneCommand Manager discovers the hosts and adapters, if you see the ports on the **Discover Information** tab in the OneCommand Manager window (Figure 10-70), click the adapter that you want update. For this example, we click **OneConnectOCe11100**.

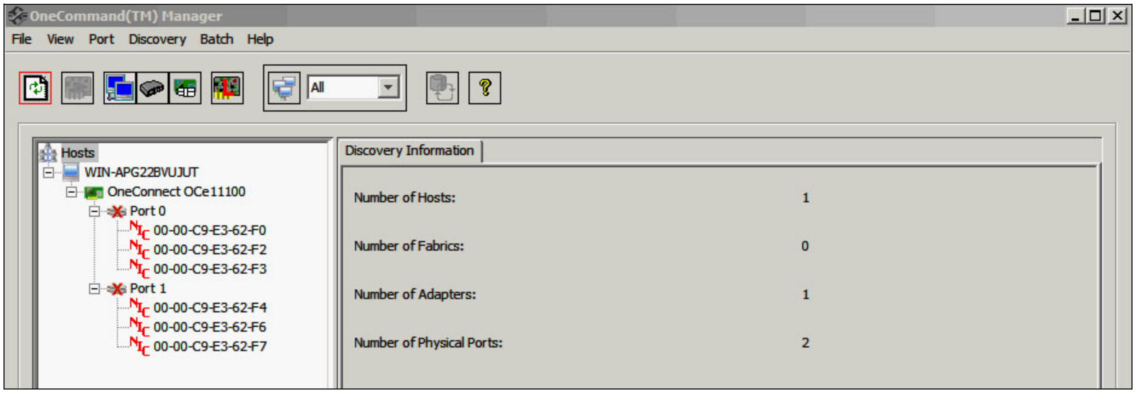

Figure 10-70 OneCommand Manager window

5. From the menu bar, select **Batch** → **Download Firmware**.
6. Select the firmware file:
   a. In the OneCommand Manager Batch Firmware Download window (Figure 10-71), click **Browse**.
   b. In the Firmware File Selection window (Figure 10-71), select the `xxx.UFI` update file. For this example, we selected the `T33973806.UFI` file. Click **OK**.

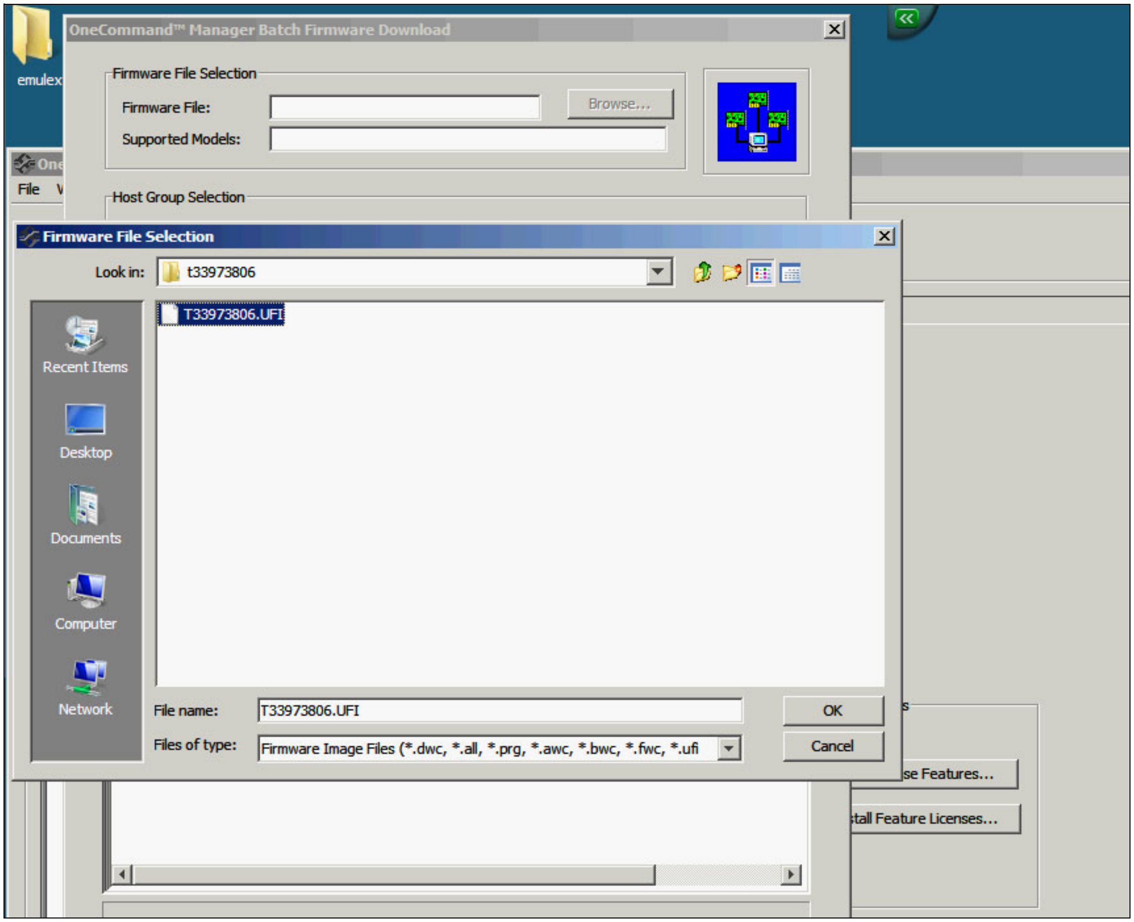

Figure 10-71 Firmware File Selection window

c. Back in the OneCommand Manager Batch Firmware Download window (Figure 10-72), next to the box where the adapters are selected, click **Start Download**.

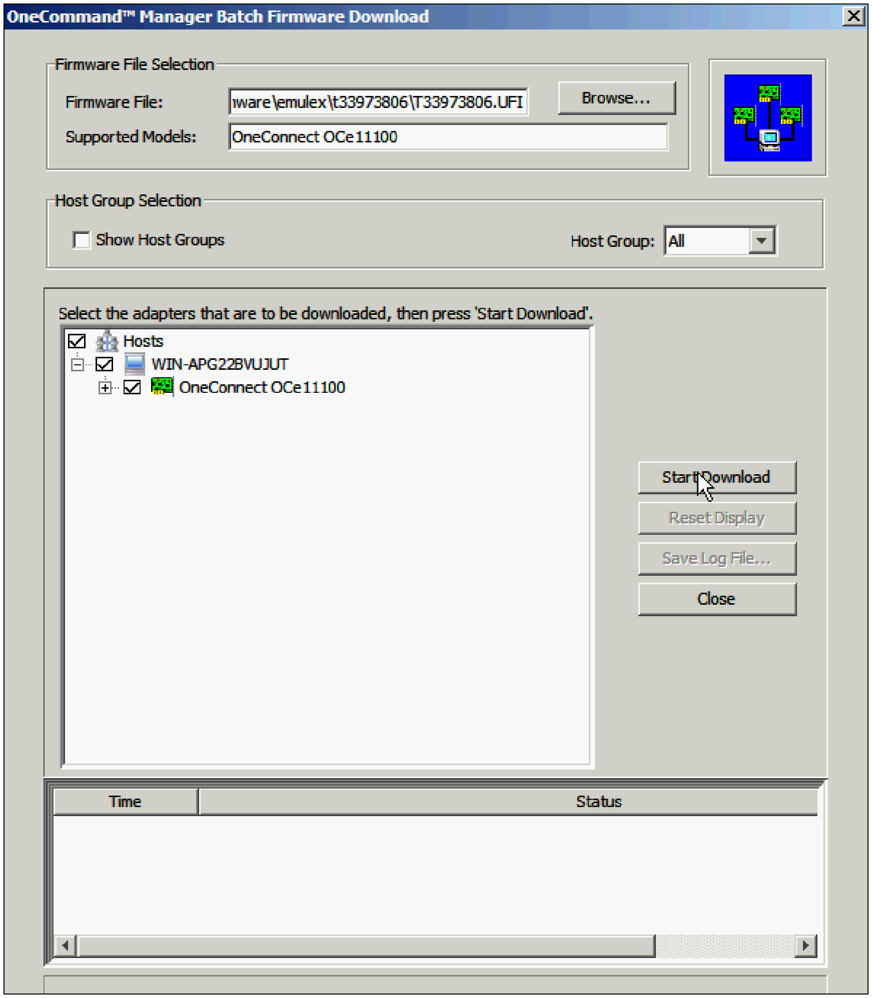

Figure 10-72 OneCommand Manager Batch Firmware Download window

Do not restart or turn off the system until the process is finished.

7. When prompted to restart the system, click **OK**.
8. Check whether the update finished successfully.

9. Again in the OneCommand Manager Batch Firmware Download window (Figure 10-73), click **Close**.

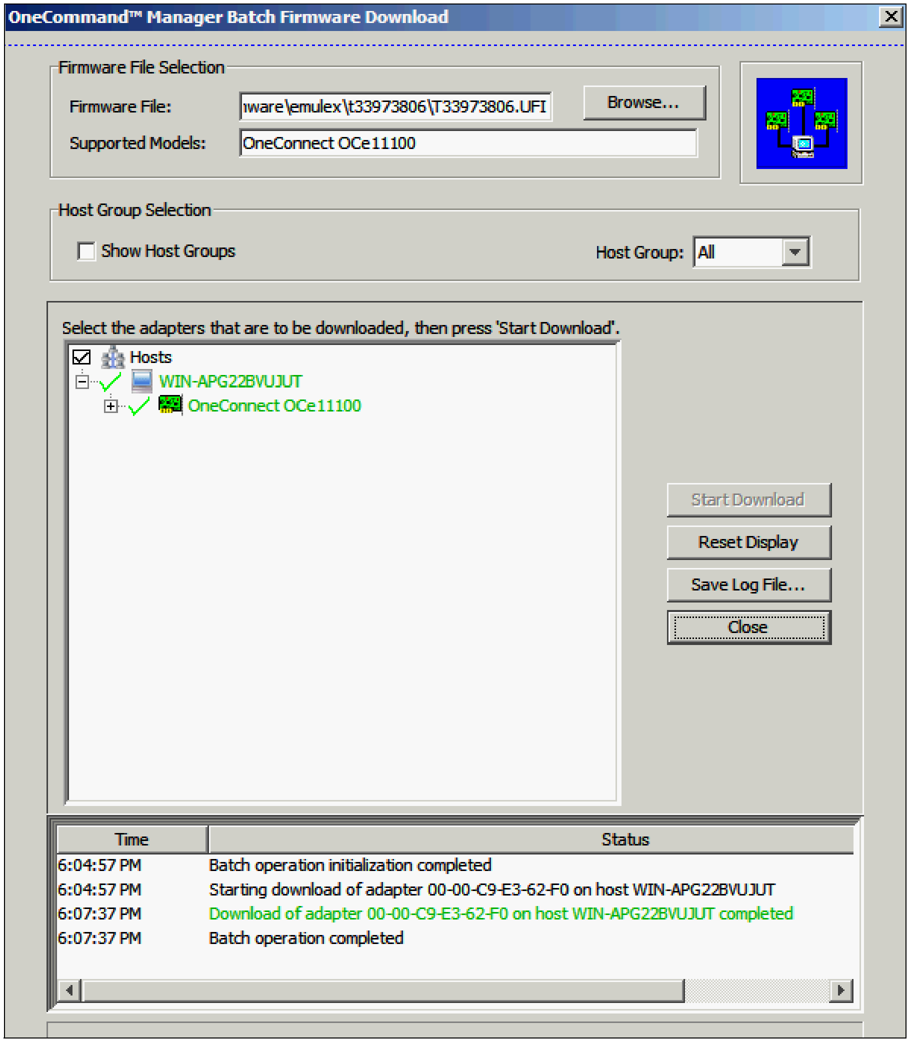

Figure 10-73 OneCommand Manager Batch Firmware Download window

10. Restart your system.

Configuring ports operation mode

To configure the ports in operation mode:

1. Open the OneCommand Manager application.

   After OneCommand Manager discovers the hosts and adapters, the ports might be displayed as unavailable as shown in Figure 10-74.

2. In the left pane, click the adapter.

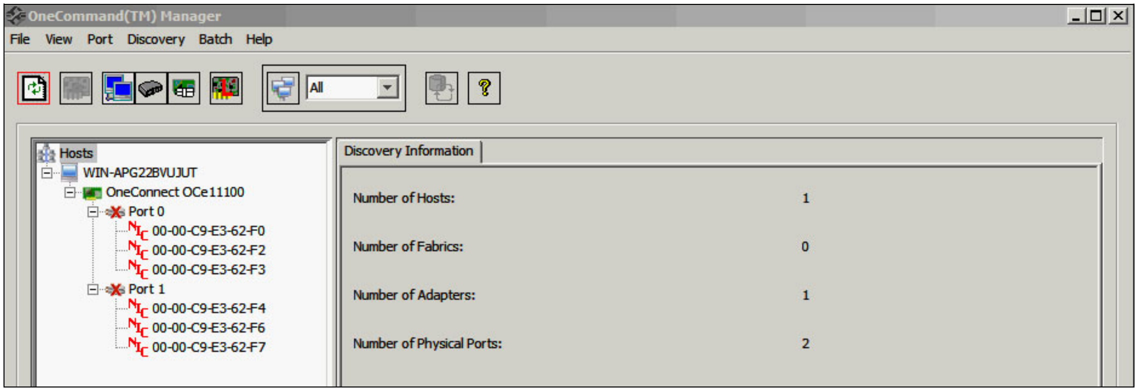

Figure 10-74 OneCommand Manager windows

3. On the **Adapter Information** tab (Figure 10-75), in the Personality area, under After Reboot, choose the disk connection type **FCoE**, and then click **Apply**.

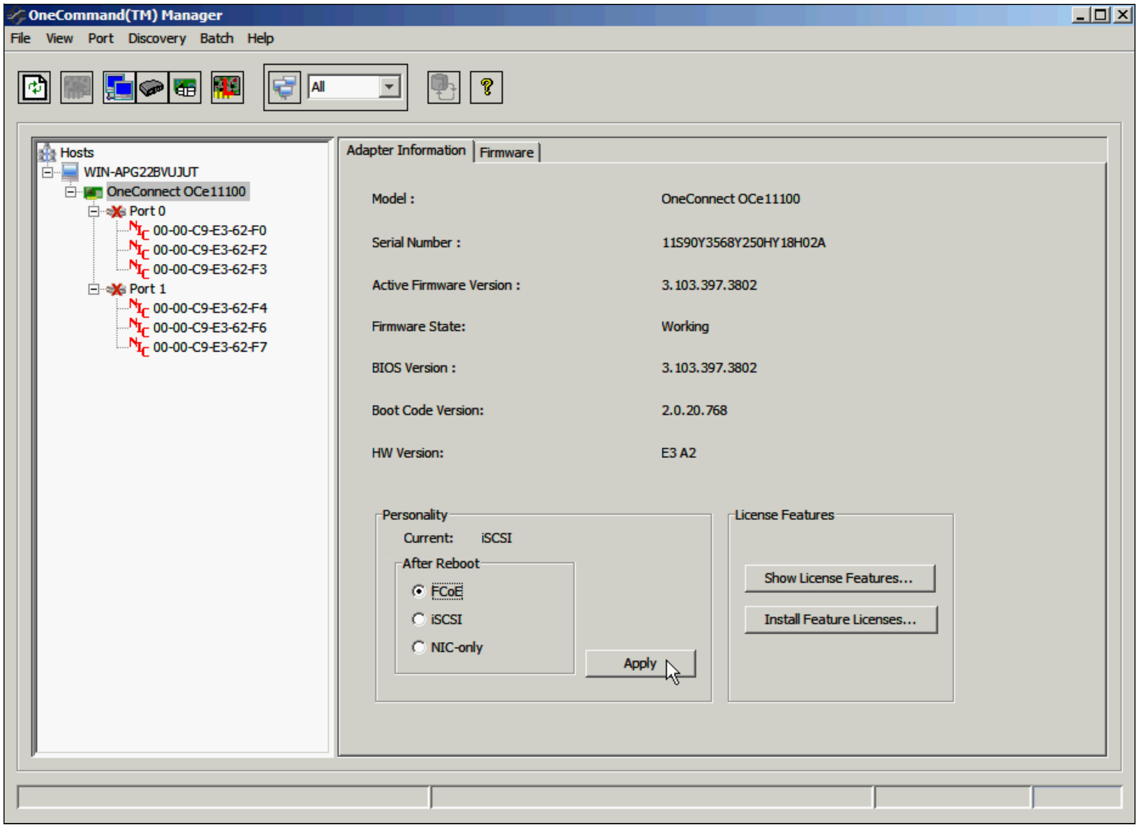

Figure 10-75 OneCommand Manager window

4. When prompted to restart your system (Figure 10-76), click **OK**.

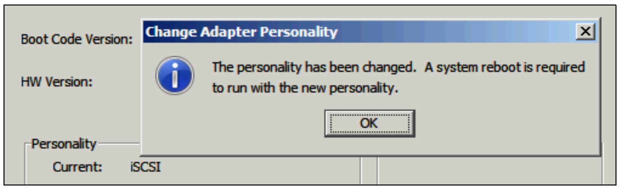

Figure 10-76 Change Adapter Personality dialog box

Configuring the BNT switch

To configure the BNT switch:

1. Log on the BNT switch by using the advanced management module (AMM) or by using a web browser. Enter the default user ID (`admin`) and password (`admin`).
2. Set bridge module bandwidth:
   a. As shown in Figure 10-77, select **Configure**.
   b. In the left navigation pane, expand **System**, and select **Bridge Module**.
   c. In the Bridge Module Configuration pane:
      i. For Bridge Module, select **3**.
      ii. For Bandwidth, enter `40` Gbps.
      iii. For Status, select **Enabled**.
      iv. Click **Submit**.

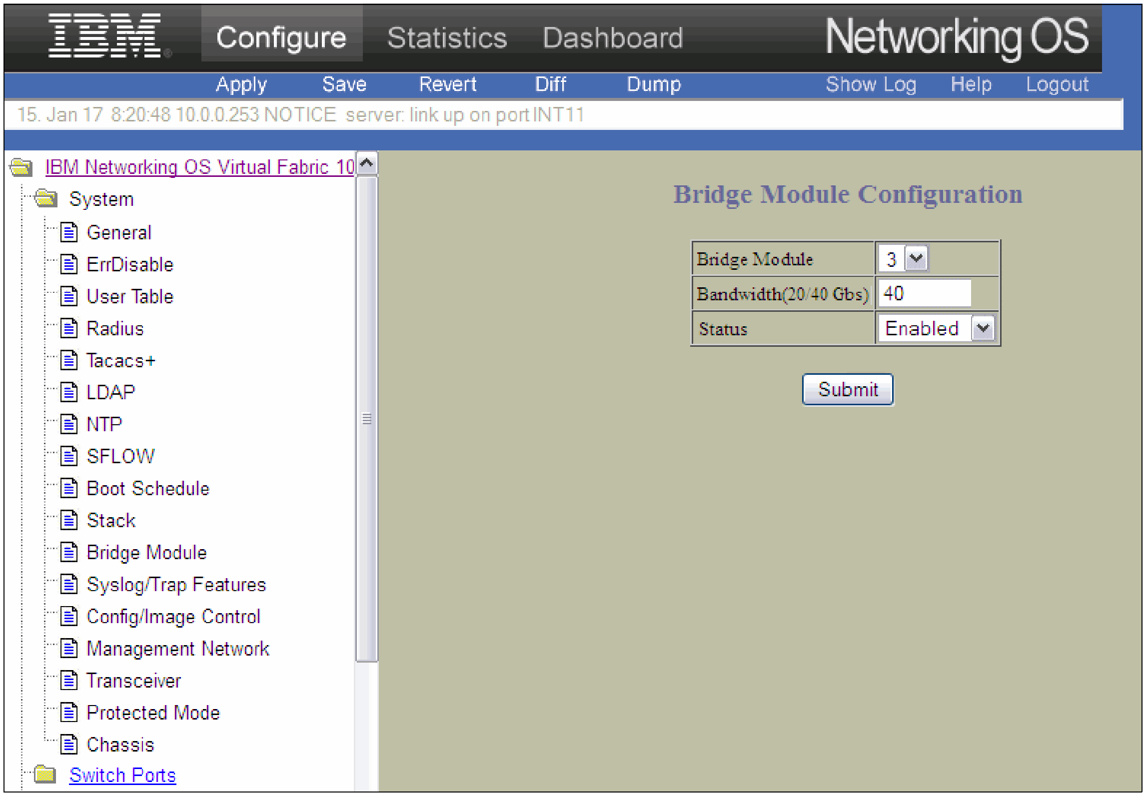

Figure 10-77 Bridge Module Configuration window

3. Disable spanning tree on the Bridge ports:

   a. Again on the **Configure** tab (Figure 10-78), in the left navigation pane, expand **Layer 2** → **Spanning Tree Group**.

   b. In the Switch Spanning Tree Groups Configuration pane on the right side, in the Spanning Tree Group column, select group **1**.

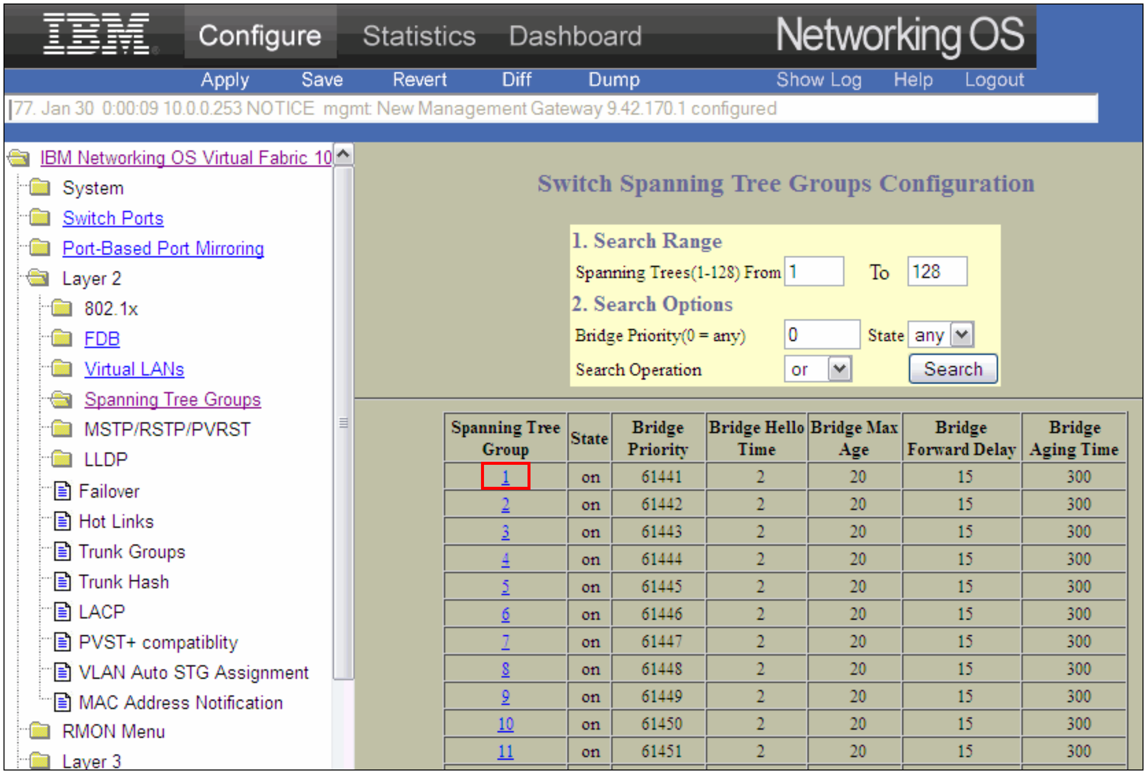

Figure 10-78 Switch Spanning Tree Groups Configuration window

4. When you see the external ports, disable external ports 1 - 4, and then click **Submit** (Figure 10-79).

INT5	128	0	off
INT6	128	0	off
INT7	128	0	off
INT8	128	0	off
INT9	128	0	off
INT10	128	0	off
INT11	128	0	off
INT12	128	0	off
INT13	128	0	off
INT14	128	0	off
EXT1	128	0	on
EXT2	128	0	on
EXT3	128	0	on
EXT4	128	0	on
EXT5	128	0	on
EXT6	128	0	on
EXT7	128	0	on
EXT8	128	0	on
EXT9	128	0	on
EXT10	128	0	on
EXT11	128	0	on

Figure 10-79 External ports

5. Click **Submit**.
6. At the top of the page, click **Apply**.
7. Reboot the switch.

8. On the **Dashboard** page (Figure 10-80), in the left navigation pane, select **Switch ports** to check the ports status and speed.

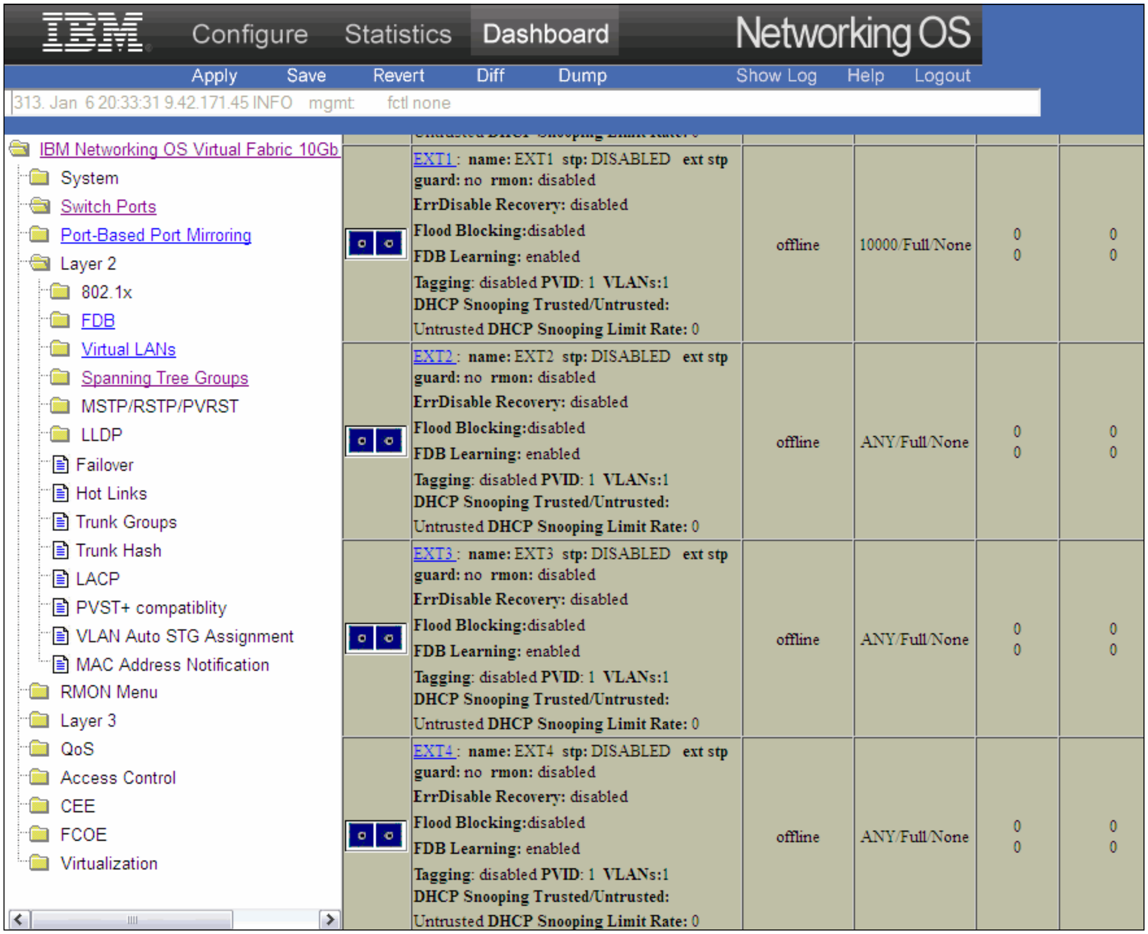

Figure 10-80 IBM Networking OS Dashboard page

9. Enable CEE global ports (Figure 10-81):
   a. Select **Configure**.
   b. In the left navigation pane, expand **CEE**, and select **General**.
   c. In the CEE Configuration pane on the right side, for Global CEE On/Off, select **On**. Click **Submit**.

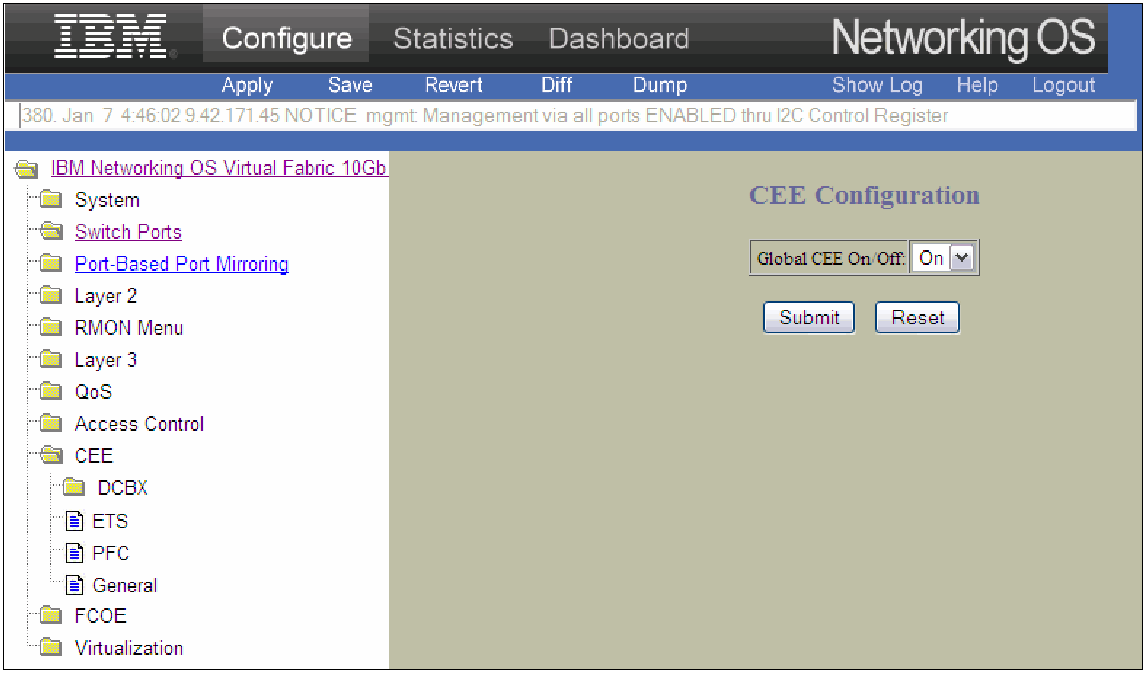

Figure 10-81 CEE Configuration pane

10. Turn on the FCoE Federal Information Processing Standard (FIPS) protocol (Figure 10-82):

   a. On the **Configure** page, in the left navigation pane, expand **FCoE**, and then select **FIP Snooping**.

   b. In the FIP Snooping Configuration pane on the right side, for Global FIPS State, select **Enabled**.

   c. Click **Submit**.

   d. Click **Apply**.

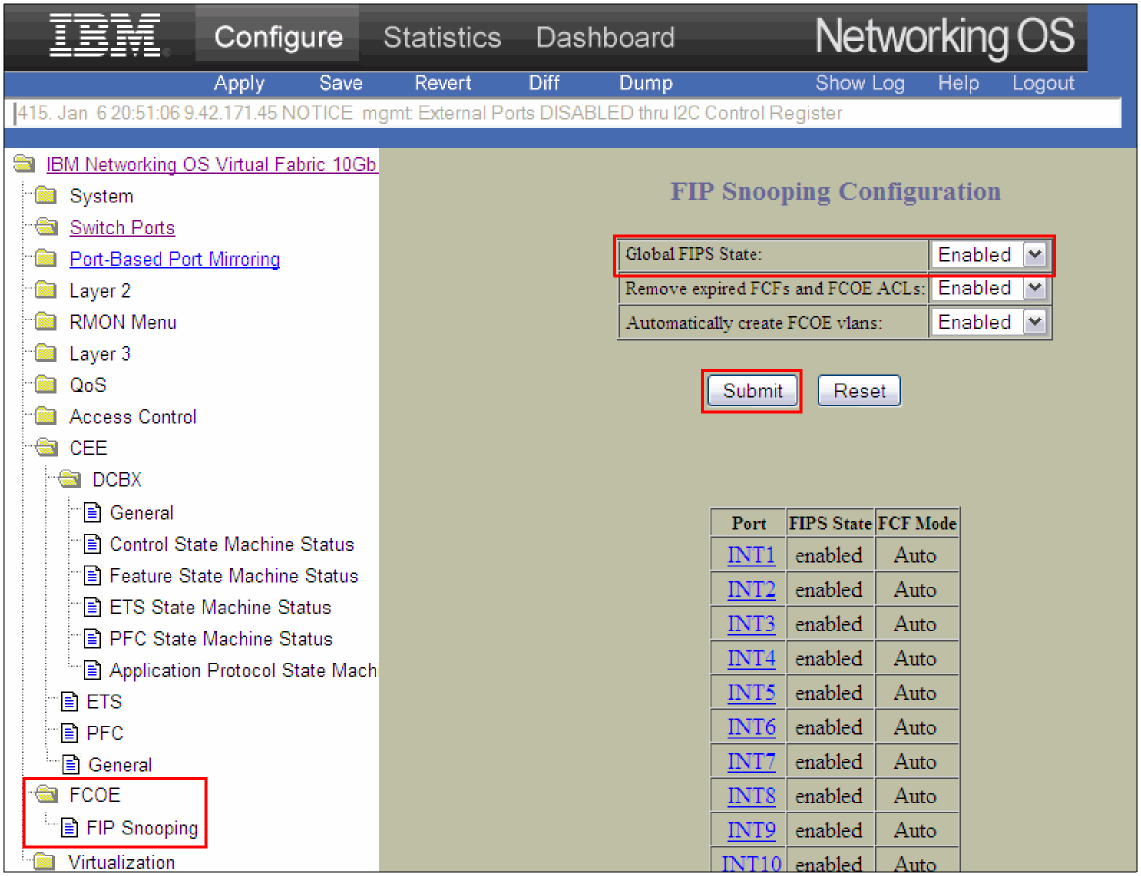

Figure 10-82 FTP Snooping Configuration pane

11. Save the configuration.

12. Create a VLAN 1002 and add INT1-14 ports and external port 4 on this VLAN (Figure 10-83):

   a. In the left navigation pane, expand **Layer 2** → **Virtual LANs**, and select **Add VLAN**.

   b. In the VLAN “New” Configuration pane on the right side:

      i. For VLAN ID, enter `1002`.
      ii. For VLAN Name, enter `VLAN 1002`.
      iii. For VLAN State, select **Enabled**.
      iv. For Management VLAN State, select **Disabled**.
      v. Under Ports Available, select ports **INT1 - INT14** and **EXT4**, and click **Add** to move them to the Ports in VLAN list.

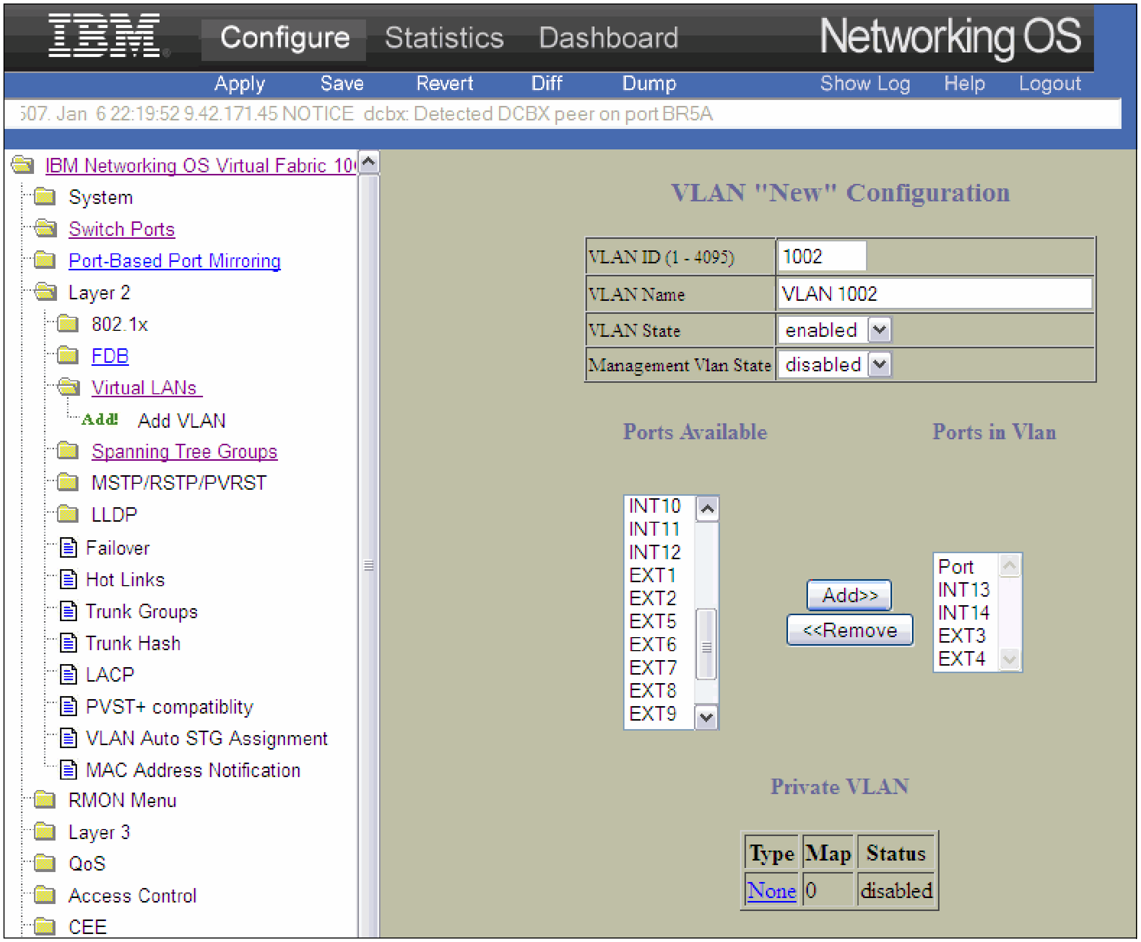

Figure 10-83 VLAN “New” Configuration pane

13. Enable tags and PVID on bridge ports (Figure 10-84):

   a. On the **Configure** page, in the left navigation pane, select **Switch Ports**.

   b. Select **EXT3**.

   c. In the Switch Port EXT3 Configuration pane:

      i. For VLAN Tagging, select **Enabled**.
      ii. Scroll down, and click **Submit**.
      iii. Click **Apply**.
      iv. For PVID Tagging, select **Enabled**.
      v. Scroll down, and click **Submit**.
      vi. Click **Apply**.

   d. Repeat steps b and c for **EXT4**.

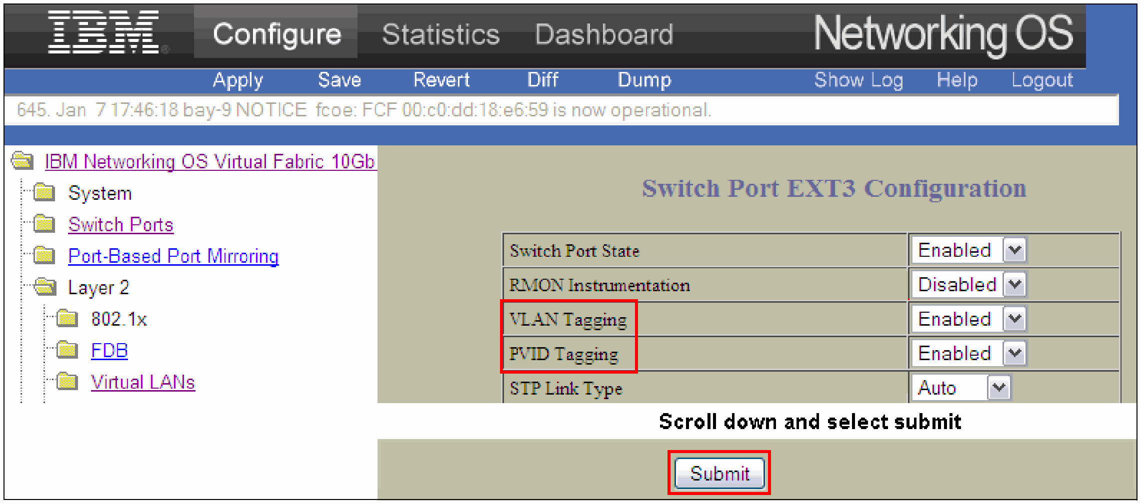

Figure 10-84 Switch Port Configuration pane

14. Remove the external ports from the other VLANs. In the VLAN "New" Configuration pane (Figure 10-83 on page 326):

   a. For VLAN ID, enter `1002`.

   b. For VLAN Name, enter `VLAN 1002`.

   c. For VLAN State, select **Enabled**.

   d. For Management VLAN State, select **Disabled**.

   e. In the Ports in VLAN list, select ports **EXT3** and **EXT4**, and then click **Remove** to move them to the Ports Available list.

   f. Click **Submit**.

15. At the top of page, click **Apply**, and then click **Save**.

BNT switch configuration: The BNT switch on bay 7 accesses the bridge module on bay 5 (internal ports BR5A, BR5B, BR5C, and BR5D). The BNT switch on bay 9 accesses the bridge on bay 3 (internal ports BR3A, BR3B, BR3C, and BR3D). Ensure that you have the correct association of BNT switches to internal ports.

Configuring the QLogic Virtual Fabric Extension Module

To configure the QLogic Virtual Fabric Extension Module, see 6.4.5, “Configuring the QLogic Virtual Extension Modules” on page 132.

Configuring IBM Storwize V7000

Now you create a logical driver and map a LUN on the IBM Storwize V7000 midrange storage. To perform the physical installation, configure the IP of the management console, create a cluster, and define the IP of cluster management, see the IBM Storwize® V7000 Product Manuals at:

https://www.ibm.com/support/docview.wss?uid=ssg1S7003318

After you configure the adapters and create a zone group, add the worldwide name (WWN) storage port on the same zone to guarantee LUN access.

To configure IBM Storwize V7000:

1. Log on to the cluster management IP.
2. On the left side of the IBM Storwize V7000 window (Figure 10-85), click the **Volumes** icon (circled in the figure).

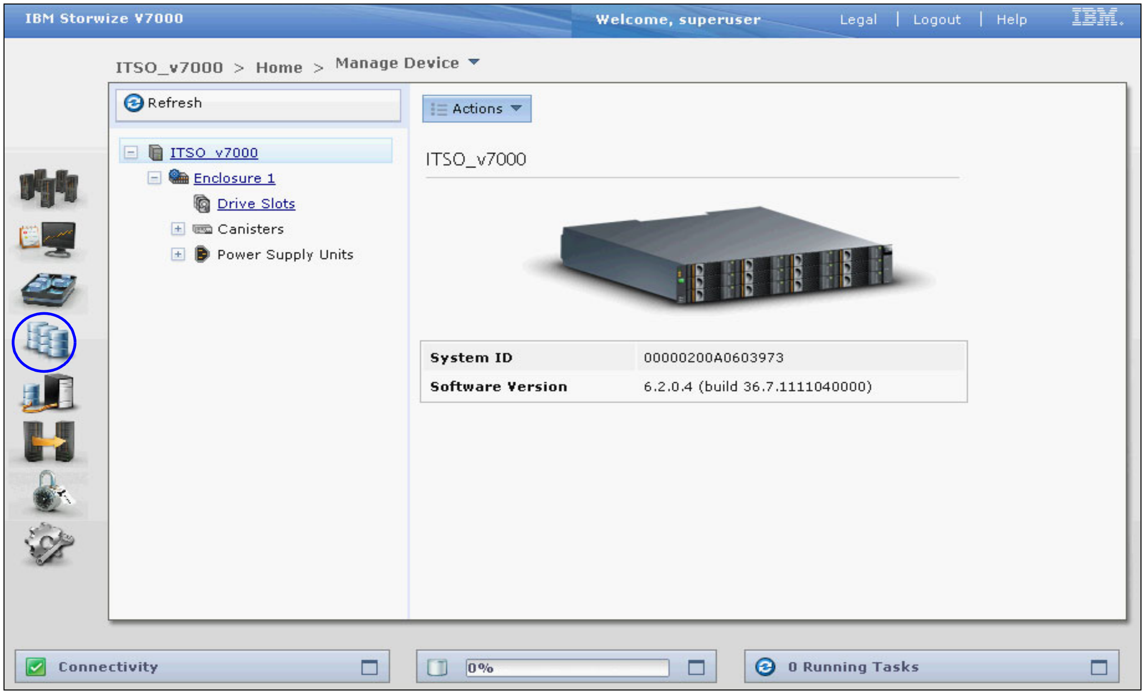

Figure 10-85 IBM Storwize V7000 window

3. Select **All Volumes** → **New Volume** (Figure 10-86).

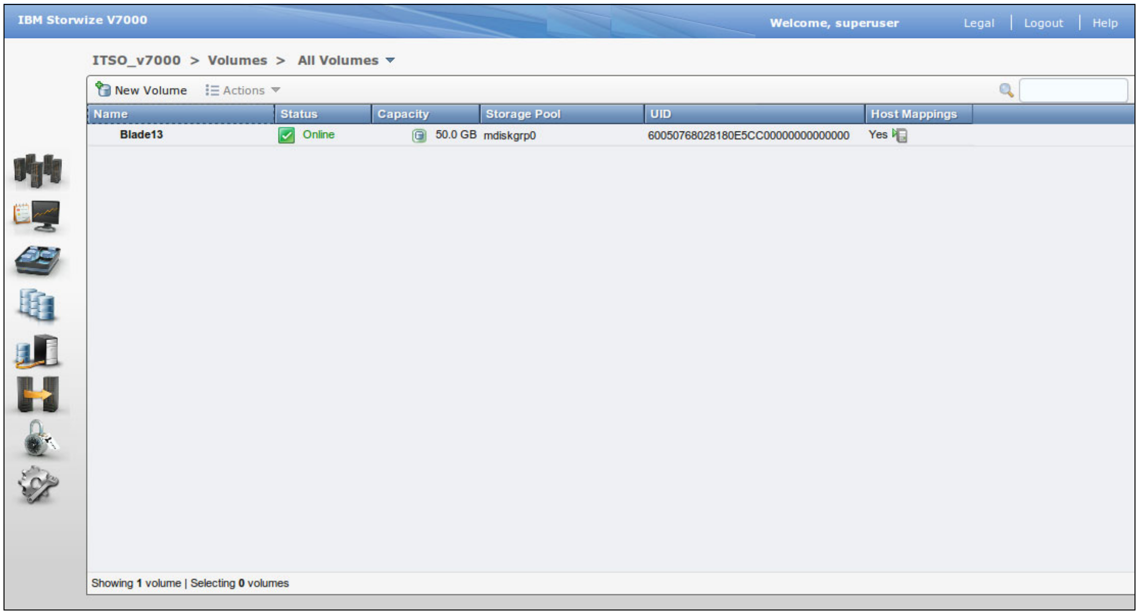

Figure 10-86 IBM Storwize V7000 window

4. In the New Volume pane (Figure 10-87), select the preset type. In this example, we selected **Thin Provision** (the best choice). Then, select a pool. In this example, we selected **mdiskgrp0**. The pane refreshes.

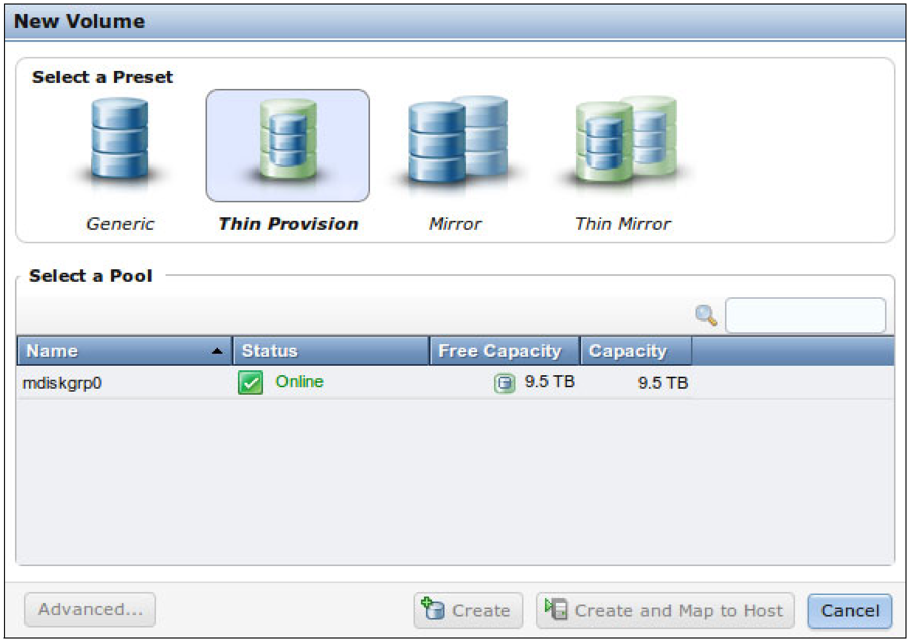

Figure 10-87 Selecting the preset type in the New Volume pane

5. In the Select Names and Sizes area (Figure 10-88), enter the name and size of the volume. In this example, for Volume Name, we enter `blade14`, and for Size, we enter `100` **GB**. Then click **Create**.

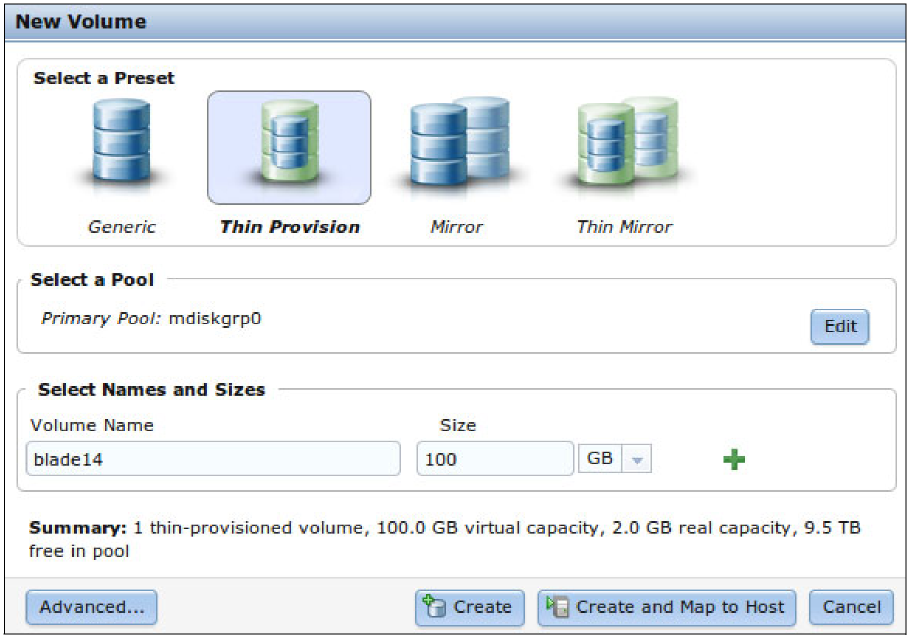

Figure 10-88 Entering the name and size in the New Volume pane

6. Click the **Hosts** icon ().
7. From the Hosts menu, select all hosts, and then click **New host**.

8. In the Create Host pane (Figure 10-89), select **Fibre Channel Host** for FC/FCoE. The pane refreshes.

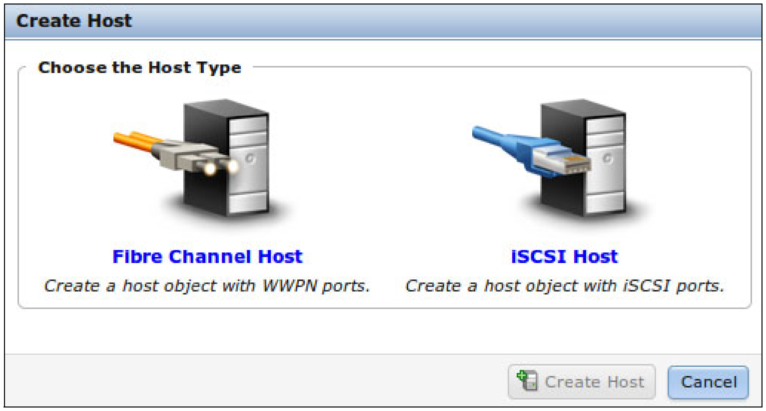

Figure 10-89 Create Host pane

9. In the expanded Create Host pane (Figure 10-90):

   a. Enter the name of host. In this example, we entered `blade 14`.
   b. Click **Rescan**.
   c. Choose the WWN.
   d. Click **Create Host**.

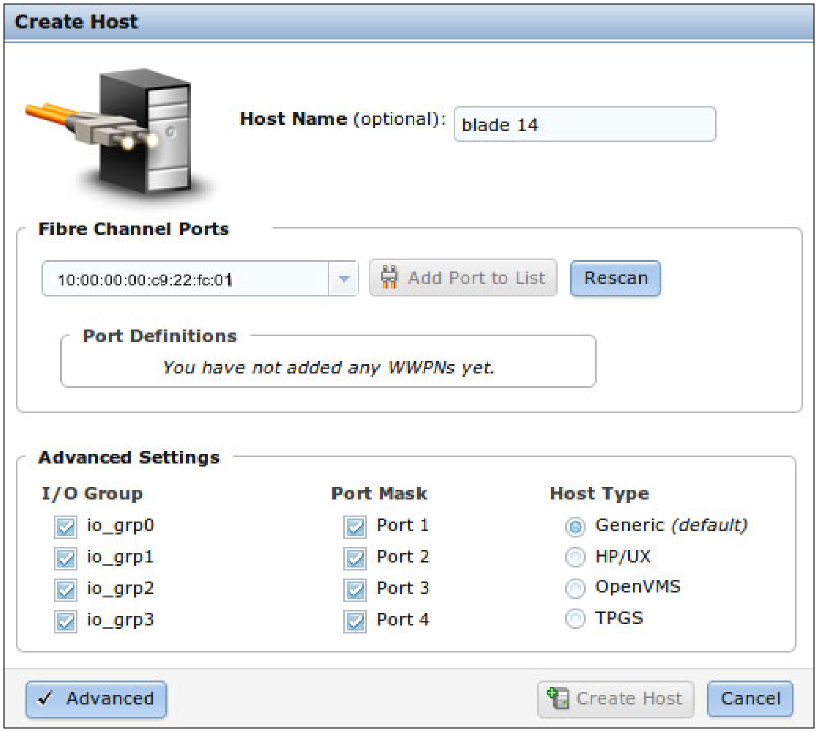

Figure 10-90 Create Host pane

10. In the Volumes panel, select the volume that you created, which in this case is **Blade13**. Then select **Actions** → **Map to Host** (Figure 10-91).

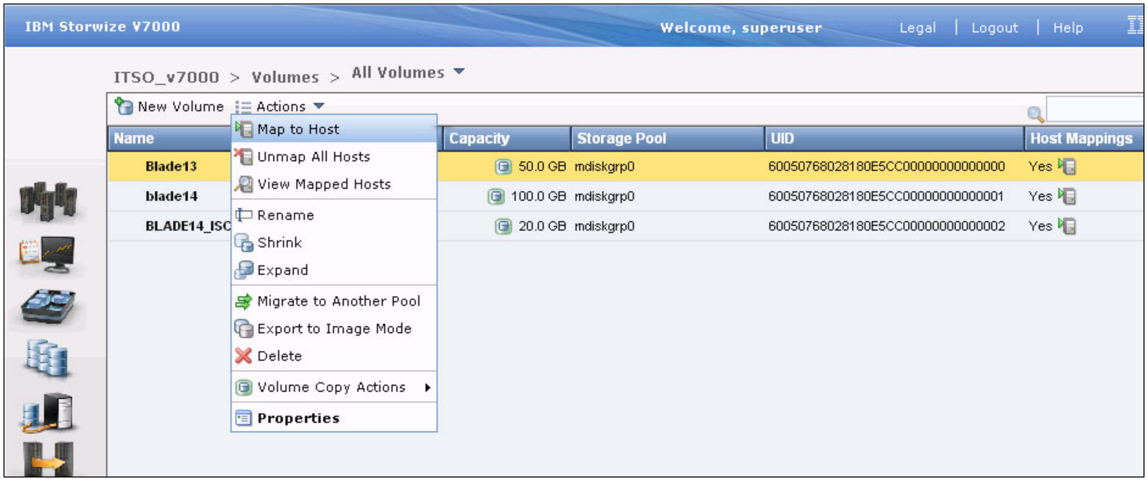

Figure 10-91 Selecting the Map to Host action

11. Select a host, and then click **map volume**.

Now that the operating system has access to the disk, you can create the partitions and format the new driver.

Firmware for this configuration

For this configuration, use the following firmware versions:

- For AMM, use BPET62C.
- For BNT Virtual Fabric 10Gb Switch Module, use version 6.7.5.0 (BOOT and OS).
- For HS22, choose from the following versions:
  - BIOS 1.15
  - IMM 1.30
  - Virtual Fabric Adapter V3.703.397.3806
- For IBM Storwize V7000, use firmware version 6.2.0.4 of the IBM Storwize V7000 Code.

Network device driver settings

Depending on your application, the device driver and system UEFI settings in the following sections can enhance your application performance.

Windows environment

> **Important:** Complete the following steps on each Universal Converged Network Adapters (UCNA) interface in the system. You must modify the advanced driver properties for each Emulex UCNA interface listed in the Device Manager.

In a Windows environment, configure the following settings:

1. Disable the TCP/IP offload engine (TOE) on all pNIC or vNIC interface instances.
2. Disable receive-side scaling (RSS) on all pNIC or vNIC interface instances.
3. In the device manager, for each UCNA interface, right-click and select **Properties**. On the **Advanced** tab:
   a. Select the **Interrupt Moderation** option.
   b. Enter the value `6` (for the lowest number of interrupts per second).
   c. Select the **RSS** option, and then choose **Disable**.
   d. In the Protocol Offloads section, expand the **IPv4** section, select **TCP Connection Offload**, and then select **Disable**.

Linux environment

In Red Hat Enterprise Linux 6 (RHEL6) and SLES 11, disable RSS on all UCNA interfaces by using the multi_rxq option during driver loading:

```
# modprobe -r be2net && modprobe be2net multi_rxq=0
```

> **Tip:** You can automate this step each time you boot by adding the following line to the `/etc/modprobe.conf` file:
>
> `options be2net multi_rxq=0`

In all supported Linux releases, adjust the interrupt moderation value on all UCNA interfaces by using the **`ethtool`** command:

```
# ethtool -C interface_name adaptive-rx off rx-usecs 96
```

Where *`interface_name`* matches with the ethX name that is assigned by the OS (such as eth2, eth3, and eth4).

> **Tip:** Repeat this **`ethtool`** command for each UCNA interface in the system. You can automate running this command during each boot by adding the **`ethtool`** command to the `/etc/rc.d/rc.local` file. This script runs after all other init scripts are run as part of the system boot process.

UEFI settings

From the main UEFI menu, set Operating Mode to **Max Performance**. By selecting this mode, you ensure that memory and QPI are set to use the full performance capability of the installed processors and memory.

References

For more information, see the following references:

- *Brocade FCoE Switch Module for IBM BladeCenter Installation and User's Guide*

  http://download2.boulder.ibm.com/sar/CMA/XSA/015bu/0/Brocade_CEE_Switch_Module_for_IBM_BladeCenter_Installation_and_Users_Guide_60Y1582.pdf

- *Emulex OneConnect Virtual Fabric Adapters (CFFh) for IBM BladeCenter Release Notes for OneConnect SW Version 5.1.42*

  http://www.emulex.com/fileadmin/files/downloads/hardware/oce10102/ibm/oneconnect_release_notes.pdf

10.3.5 Configuring Switch Independent Mode: Emulex Virtual Fabric Adapter II, BNT, and IBM Storwize V7000 or iSCSI

This section explains how to configure Switch Independent Mode for iSCSI storage by using Emulex Virtual Fabric Adapter II, the BNT Virtual Fabric 10Gb Switch Module, and IBM Storwize V7000.

Disabling Virtual Fabric Adapter vNIC II on the UEFI utility

To disable Virtual Fabric Adapter vNIC II:

1. Turn on the system.
2. When prompted, press F1 to enter the UEFI setup.
3. In the System Configuration and Boot Management panel (Figure 10-3 on page 244), select **System Settings**.

4. In the System Settings panel (Figure 10-92), select **Network**.

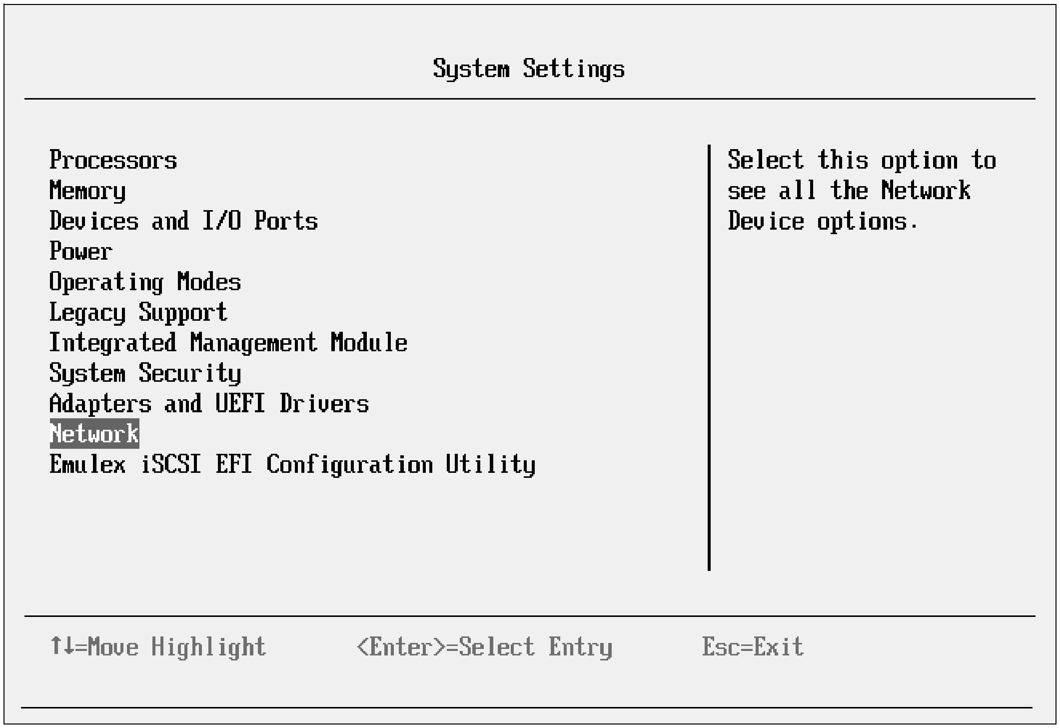

Figure 10-92 System Settings panel

5. In the Network panel (Figure 10-93), select a port on the Emulex adapter.

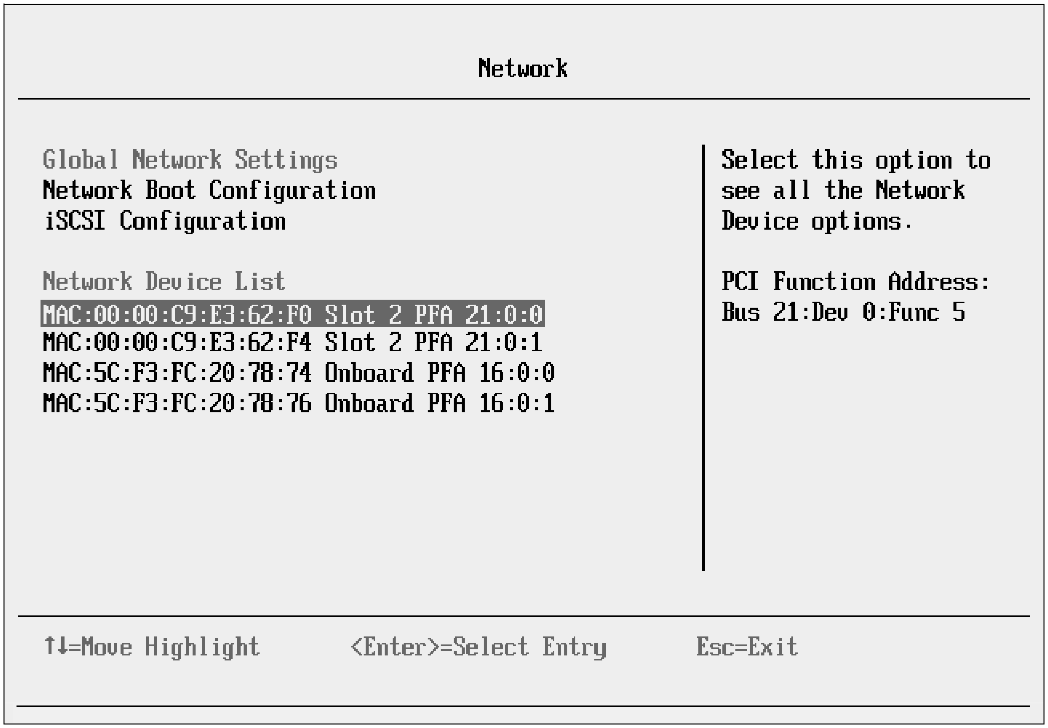

Figure 10-93 Network panel

6. Enable the port (Figure 10-94).

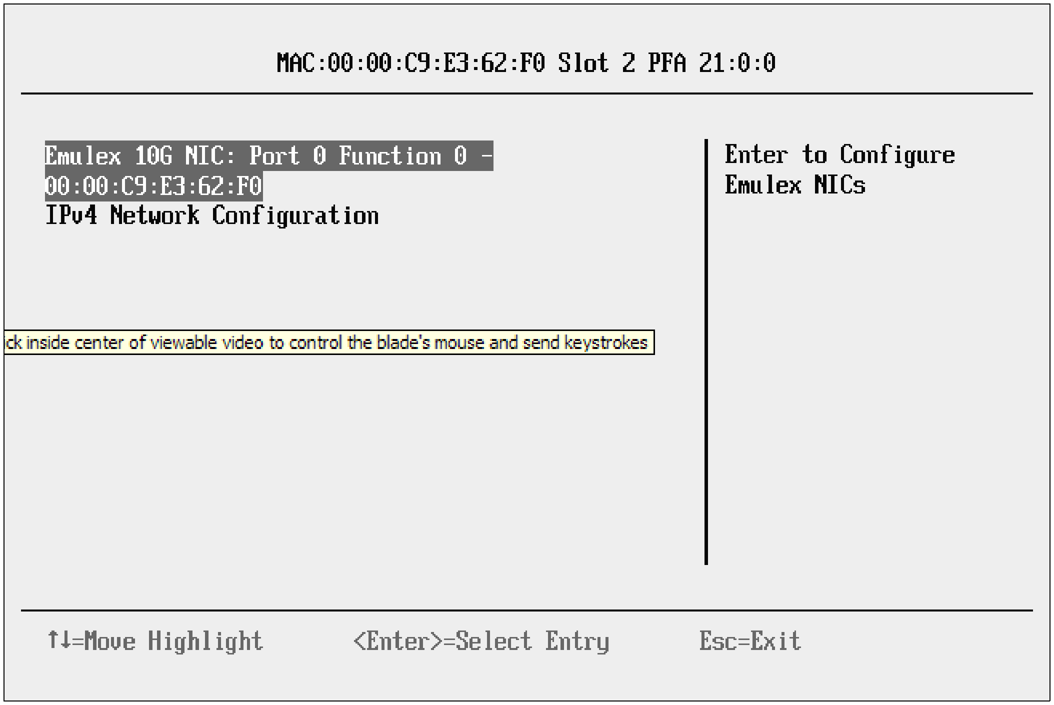

Figure 10-94 Emulex 10G NIC port

7. In the Emulex NIC Selection panel (Figure 10-95), for Mode, select **Switch Independent Mode**.

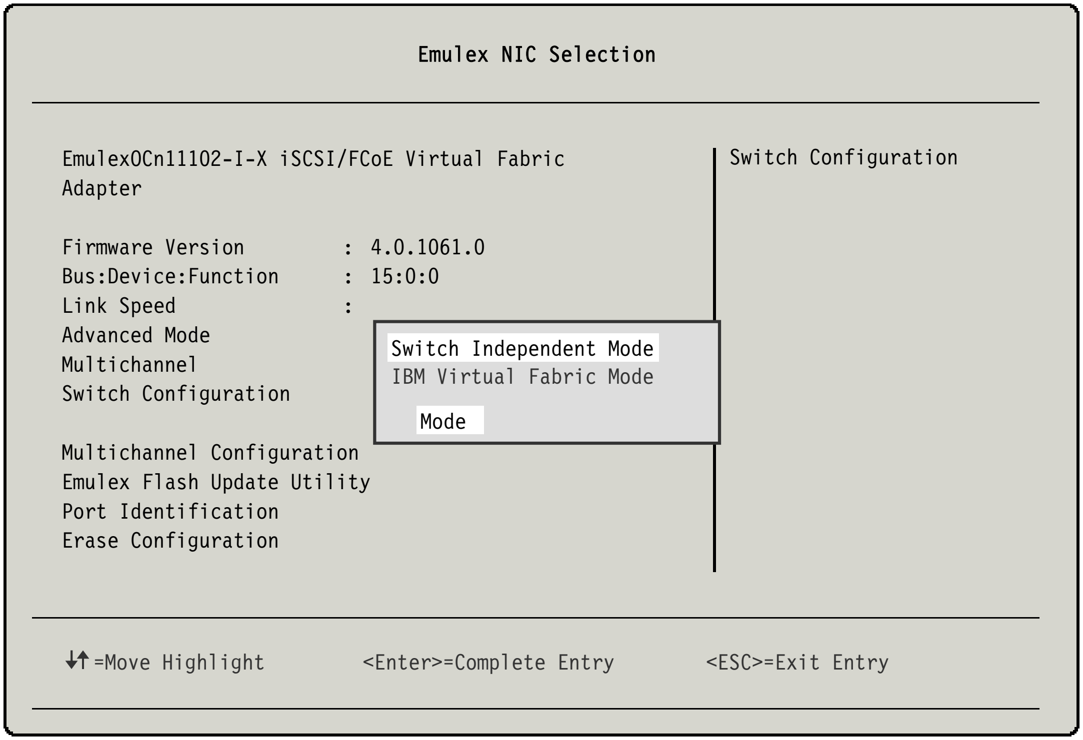

Figure 10-95 Emulex NIC Selection panel

8. Back in the System Configuration and Boot Management panel (Figure 10-3 on page 244), save your settings, and then exit the UEFI utility.

After the installation is complete, add the device drivers as explained in the following section.

Installing the Windows 2008 drivers and management software

Next, you install the device drivers and management software.

Installing the device drivers

To install the device drivers:

1. Download the device driver from the Emulex website at:

   http://www.emulex.com/downloads/ibm/vfa-software-kits.html

   Download the NIC drivers for network and ISCSI device drivers and for Emulex OneCommand Manager.

In this example, we downloaded the following NIC and iSCSI device drivers:

- `elxdrvr-nic-2.103.389.0-1`
- `elxdrvr-nic-teaming-2.0.3-4`
- `elxdrvr-iscsi-2.103.386.0-1`

In addition, we downloaded Emulex OneCommand Manager driver `elxocm-windows-x86-5.1.42.4-1` (32-bit).

2. Run the `elxdrvr-nic-2.103.389.0-1` NIC device driver. Then click **Next**.

> **Latest version:** At the time this book was written, we used Emulex OneCommand Manager versions `elxdrvr-iscsi-2.103.386.0-1` and `elxdrvr-fc-fcoe-V3.703.397.3806`. Always obtain the latest version for your configuration.

3. If you are planning to use boot from SAN or to create a Windows PE disk image, select **Unpack All Drivers** (Figure 10-96). For our example, we did not select **Unpack All Drivers**. Then click **Install**.

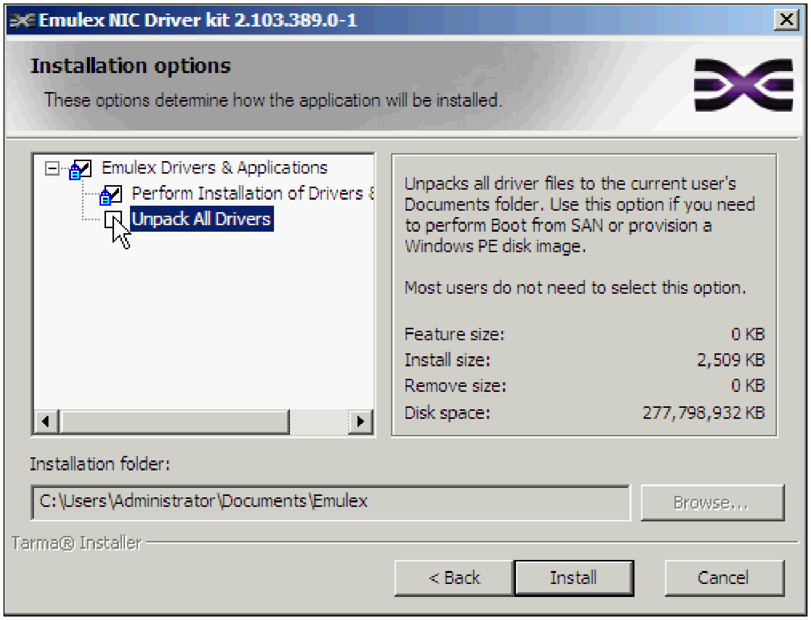

Figure 10-96 Installation options window

4. Click **Finish**. The AutoPilot Installer starts.

5. In the Emulex iSCSI Driver kit window (Figure 10-97), click **Next**.

Figure 10-97 FC and FCoE Driver window

6. Click **Finish** to complete the installation.

If you are using the iSCSI driver `elxdrvr-iscsi-2.103.386.0-1` file, repeat step 2 on page 342 through step 6. For redundant network purposes, also install the teaming driver `elxdrvr-nic-teaming-2.0.3-4` file.

Installing the Emulex OneCommand Manager

To install the Emulex OneCommand Manager:

1. Run the `elxocm-windows-x86-5.1.42.4-1` (32 bit) file.
2. In the Emulex OCManager Enterprise window (Figure 10-98), click **Next**.

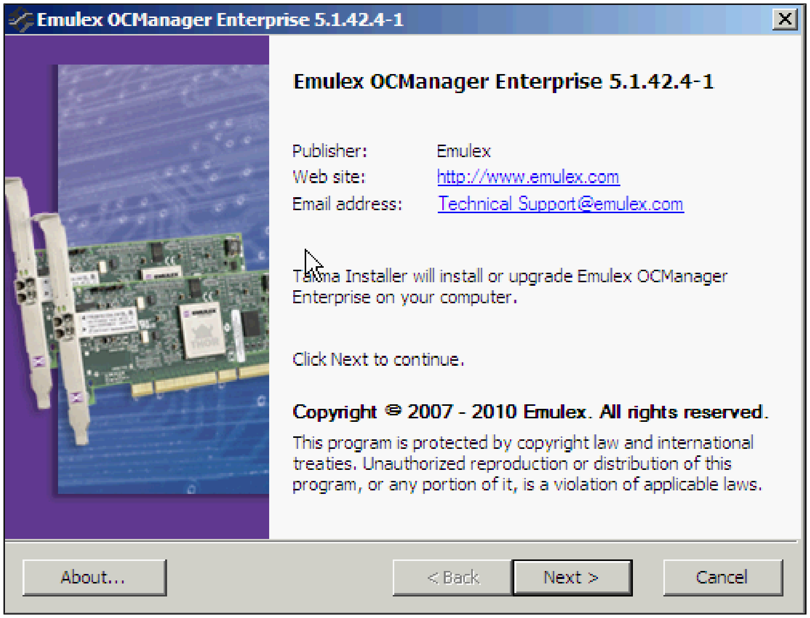

Figure 10-98 Emulex OCManager Enterprise window

3. In the Installation options window (Figure 10-99), click **Install** to accept the default options.

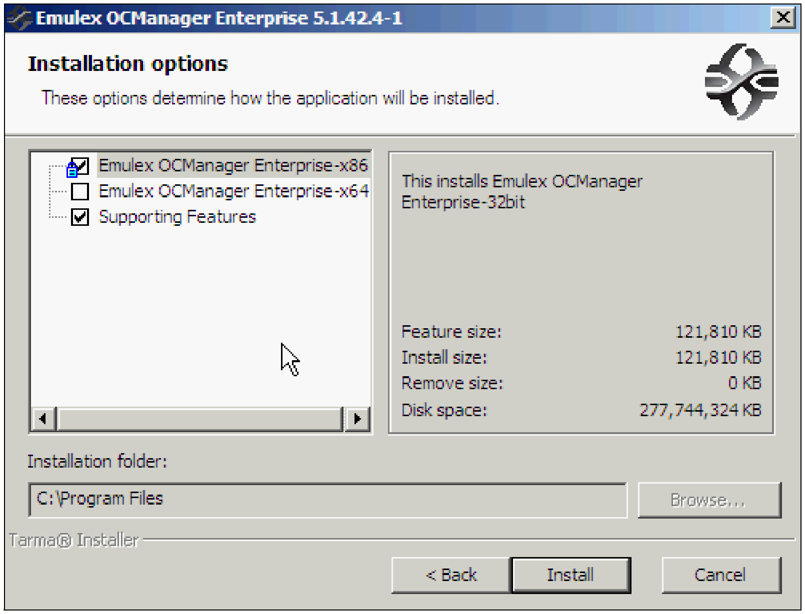

Figure 10-99 Installation options window

4. In the OneCommand Manager Management Mode window (Figure 10-100), select one of the following modes of operation:
   - Strictly Local Management (only management of the local adapters)
   - Local Management Plus (only management of the local adapters but allows management from other hosts)
   - Full Management (management of local and remote hosts and allows management from other hosts)

   In this example, we selected **Full Management**.

   Click **OK**.

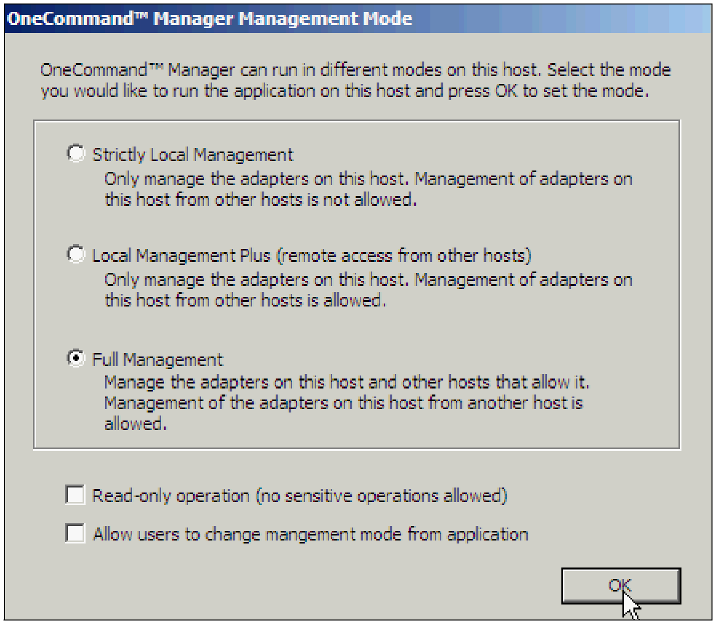

Figure 10-100 OneCommand Manager window

Updating the Virtual Fabric Adapter firmware

Important: You must update the firmware to its latest release level.

To update the Virtual Fabric Adapter firmware:

1. Download Virtual Fabric Adapter file from the Emulex website at:

   `http://www.emulex.com/downloads/ibm/vfa-software-kits.html`

2. Open the OneCommand Manager application.

3. In the Emulex OCManager Enterprise window (Figure 10-101), click **Next**.

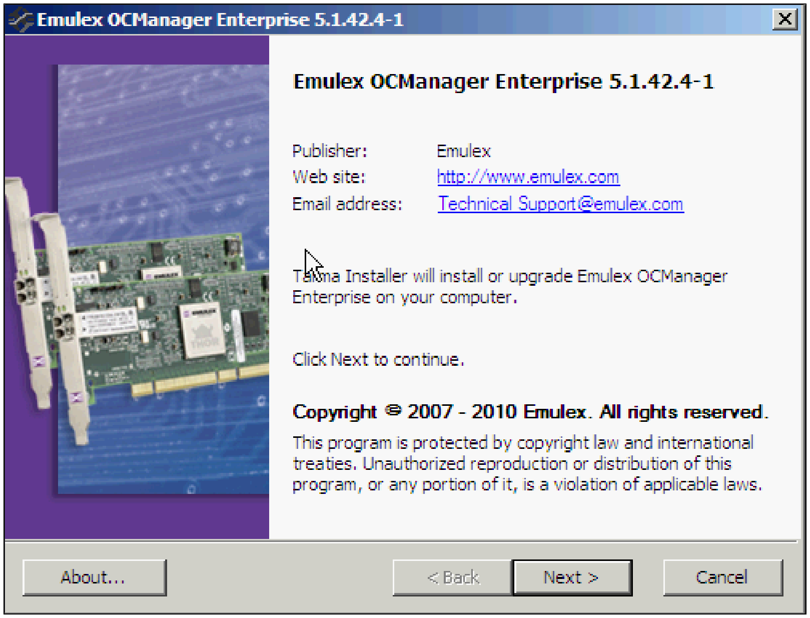

Figure 10-101 Emulex OCManager Enterprise 5.1.42.4-1 window

4. After OneCommand Manager discovers the hosts and adapters, if you see the ports on the **Discover Information** tab in the OneCommand Manager window (Figure 10-102), click the adapter that you want update. For this example, we click **OneConnectOCe11100**.

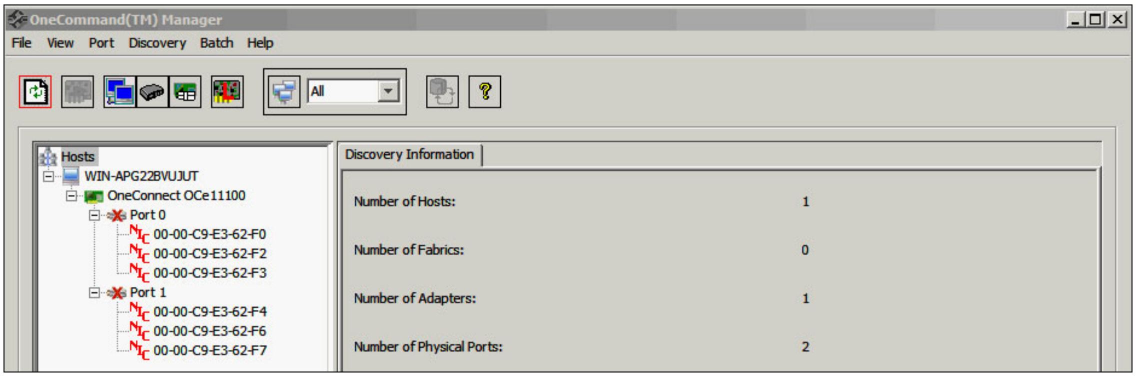

Figure 10-102 OneCommand Manager window

5. From the menu bar, select **Batch** → **Download Firmware**.

6. Select the firmware file:

   a. In the OneCommand Manager Batch Firmware Download window (Figure 10-71), click **Browse**.

   b. In the Firmware File Selection window (Figure 10-103), select the ***xxx*.UFI** update file. For this example, we selected the **T33973806.UFI** file. Click **OK**.

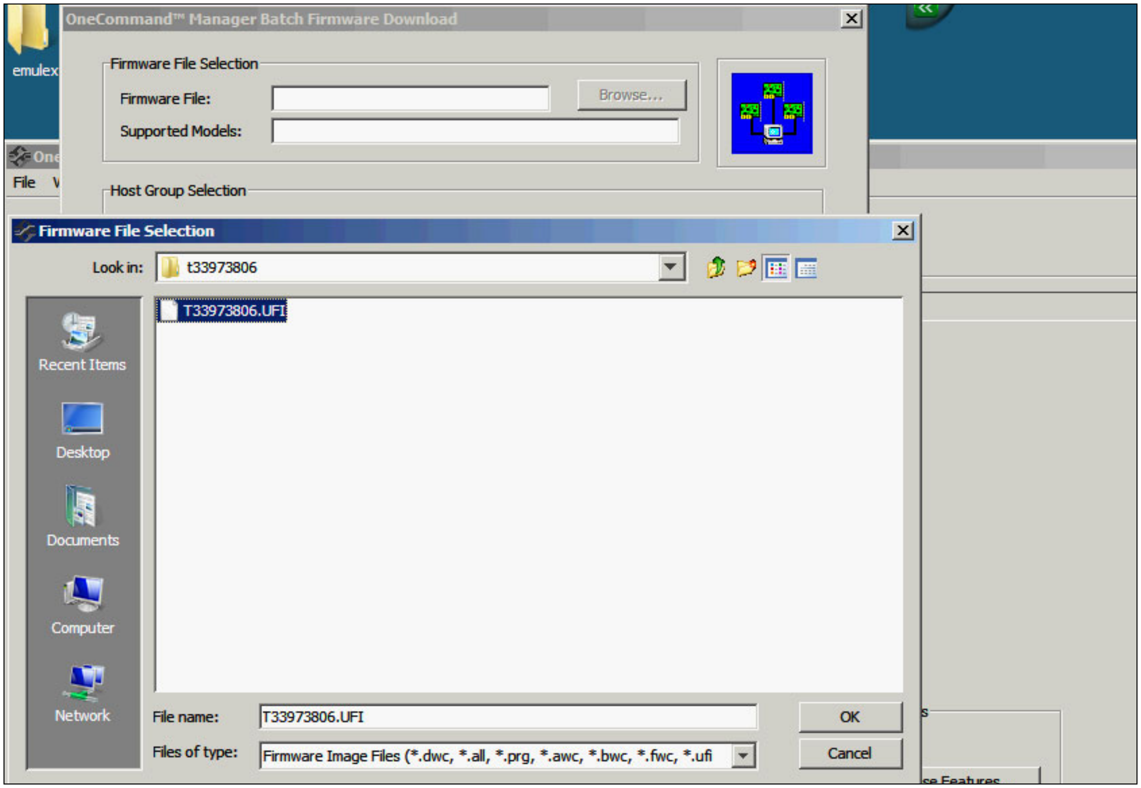

Figure 10-103 Firmware File Selection window

c. Back in the OneCommand Manager Batch Firmware Download window (Figure 10-104), next to the box where the adapters are selected, click **Start Download**.

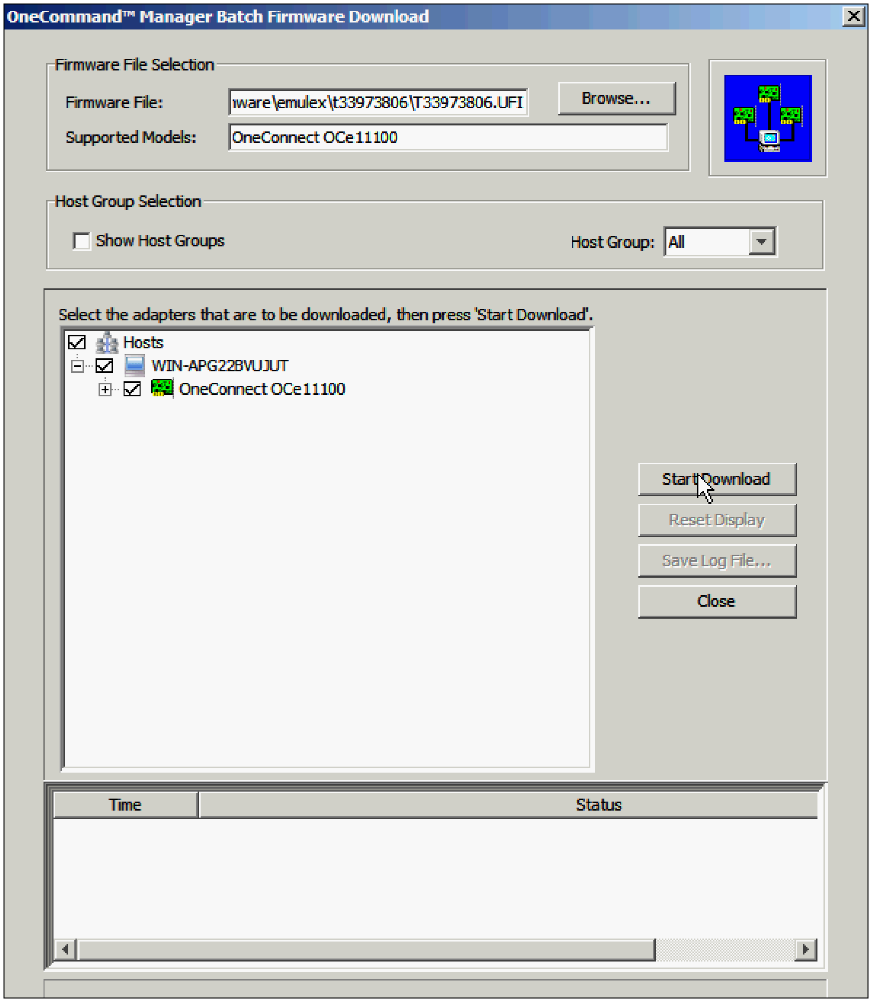

Figure 10-104 OneCommand Manager Batch Firmware Download window

Do not restart or turn off the system until the process is finished.

7. When prompted to restart the system, click **OK**.
8. Check whether the update finished successfully.

9. Again in the OneCommand Manager Batch Firmware Download window (Figure 10-105), click **Close**.

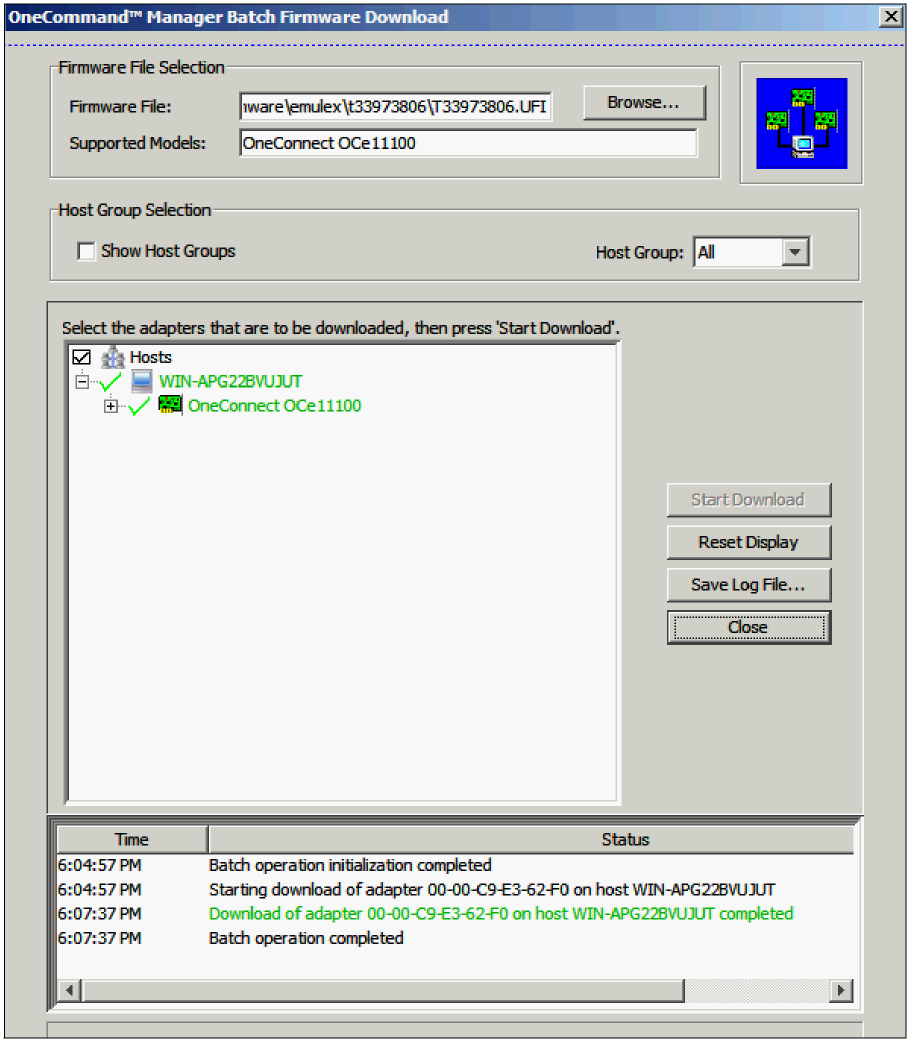

Figure 10-105 OneCommand Manager Batch Firmware Download window

10. Restart your system.

Configuring ports operation mode

To install the device driver:

1. Open the OneCommand Manager application.

   After OneCommand Manager discovers the hosts and adapters, the ports might be displayed as unavailable as shown in Figure 10-106.

2. In the left pane, click the adapter.

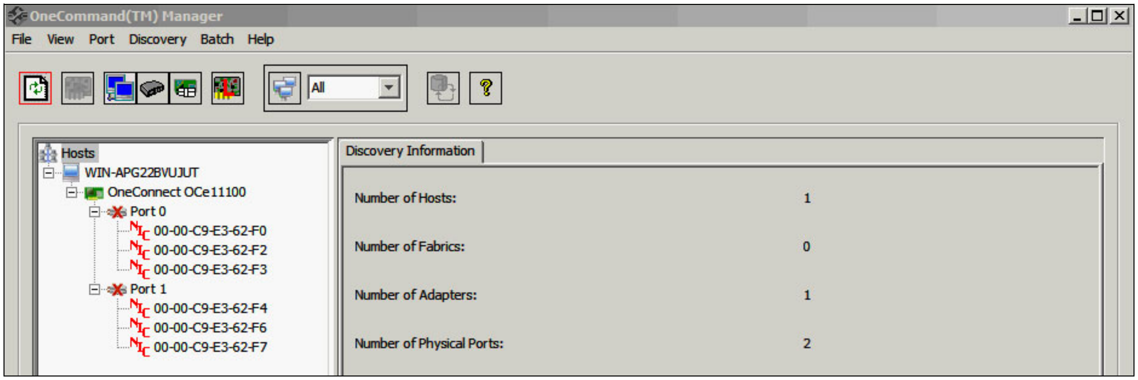

Figure 10-106 OneCommand Manager windows

3. On the **Adapter Information** tab (Figure 10-107), in the Personality area, under After Reboot, choose the disk connection type **iSCSI**, and then click **Apply**.

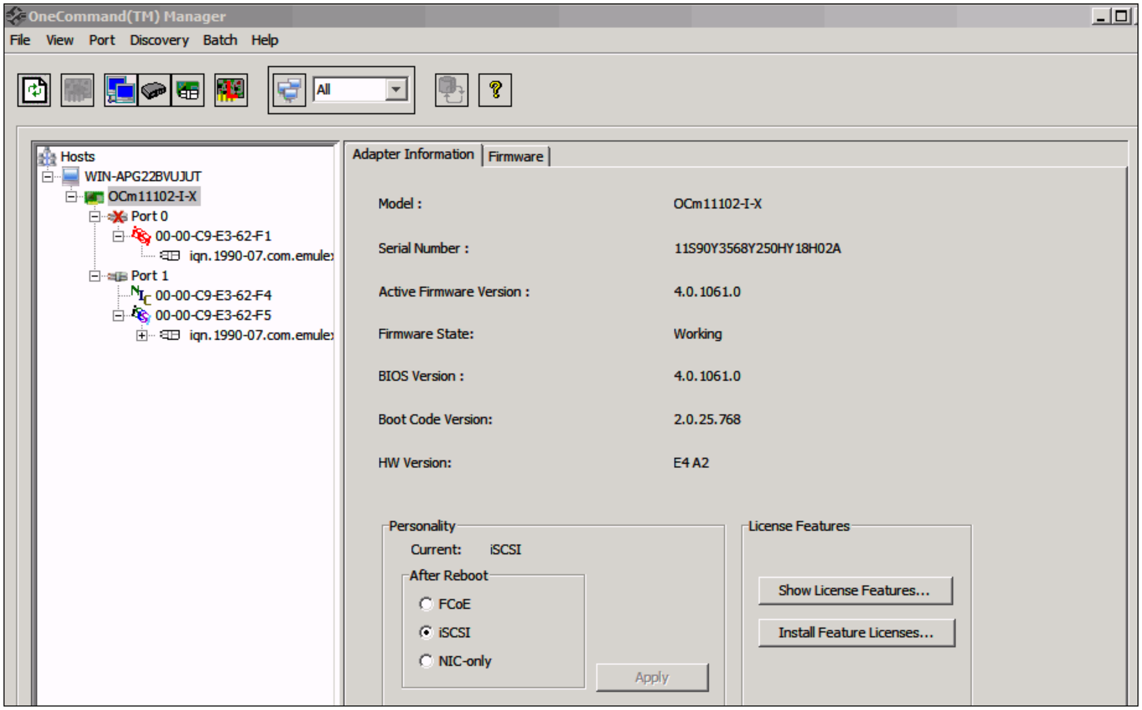

Figure 10-107 OneCommand Manager

4. When prompted to restart your system (Figure 10-108), click **OK**.

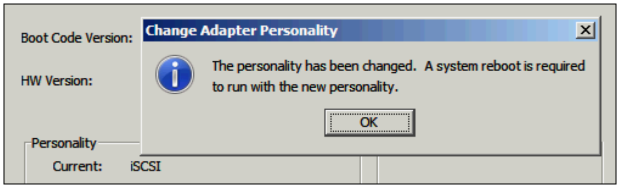

Figure 10-108 Change Adapter Personality dialog box

5. Back in the OneCommand Manager window (Figure 10-109):
   a. Select the port to configure.
   b. Select **iSCSI Initiator**.
   c. On the **iSCSI Port Info** tab, click **Modify**.
   d. In the Modify TCP/IP Configuration window (inset in Figure 10-109), enter the VLAN and IP address configuration. Click **OK**.

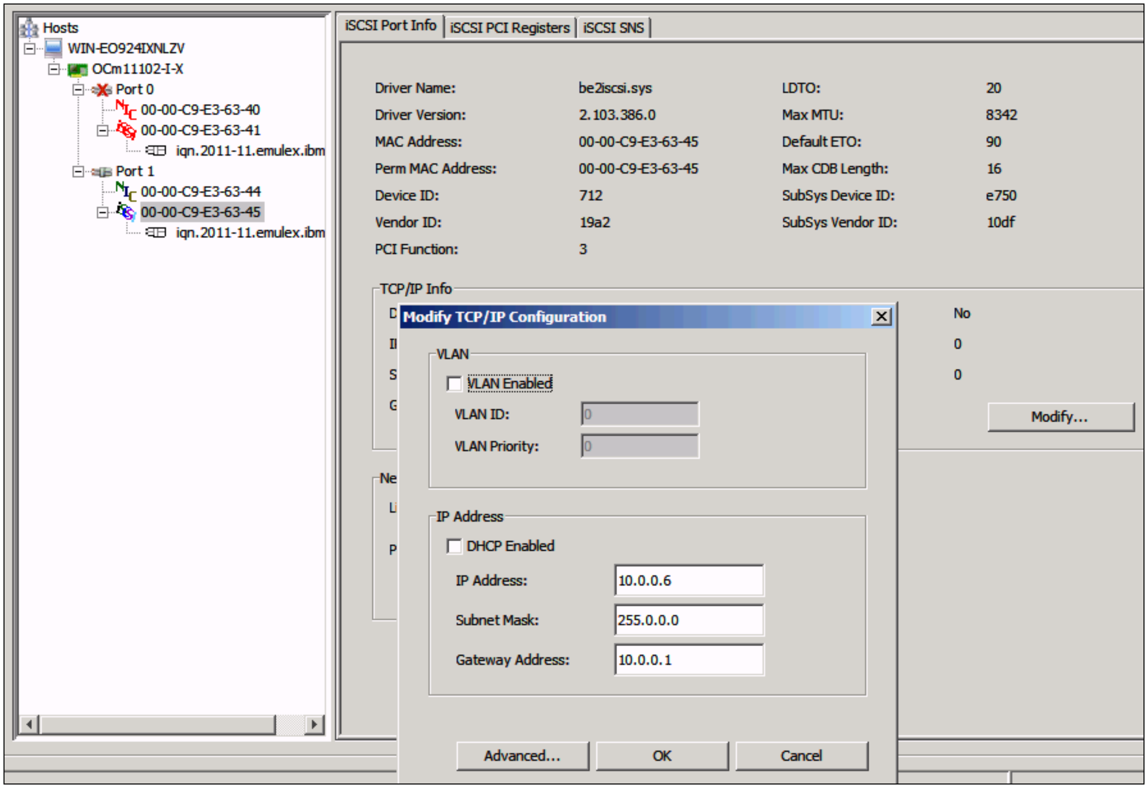

Figure 10-109   Modify TCP/IP Configuration window

e. In left pane, select the iSCSI port (Figure 10-110).

f. On the **iSCSI Initiator Login Options** tab, enter the initiator name and alias. Then click **Apply Changes**.

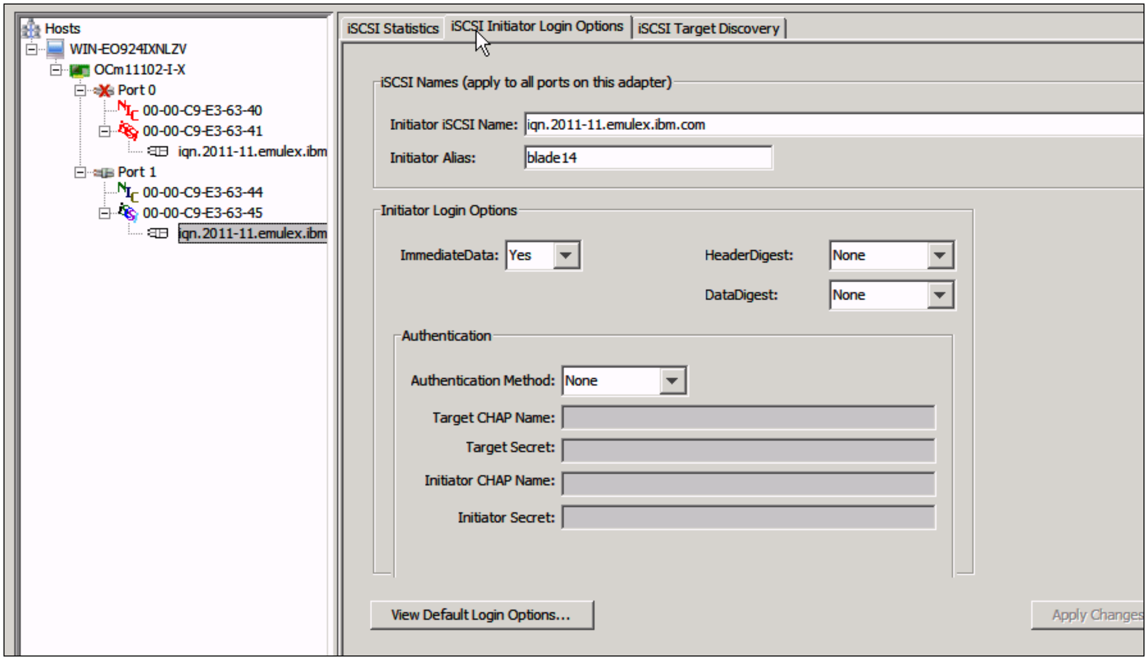

Figure 10-110 iSCSI Initiator Login Options window

g. On **iSCSI Target Discovery** tab (Figure 10-111), click **Add Portal**.

h. In the Add Target Portal window (inset in Figure 10-111), enter the IP address of the storage iSCSI. Leave the default port as `3260`. Click **OK**.

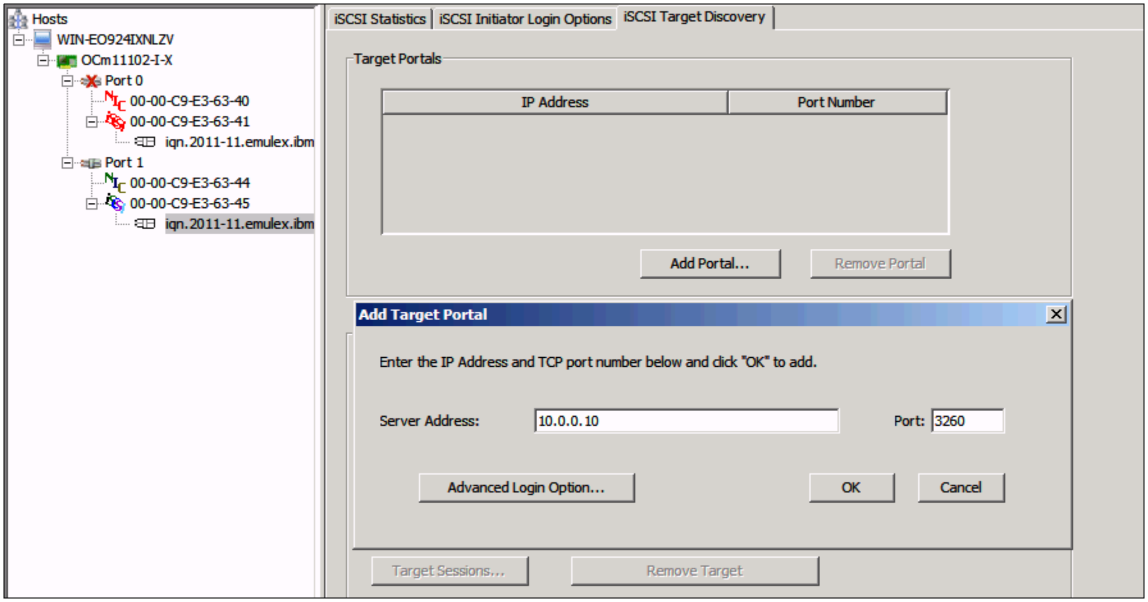

Figure 10-111 Add Target Portal window

6. In the window that shows the message, "`Success adding send target portal appear`," click **OK**.

The Targets section of the iSCSI Target Discovery tab now shows the iSCSI target (Figure 10-112).

7. Click **Target Login**.

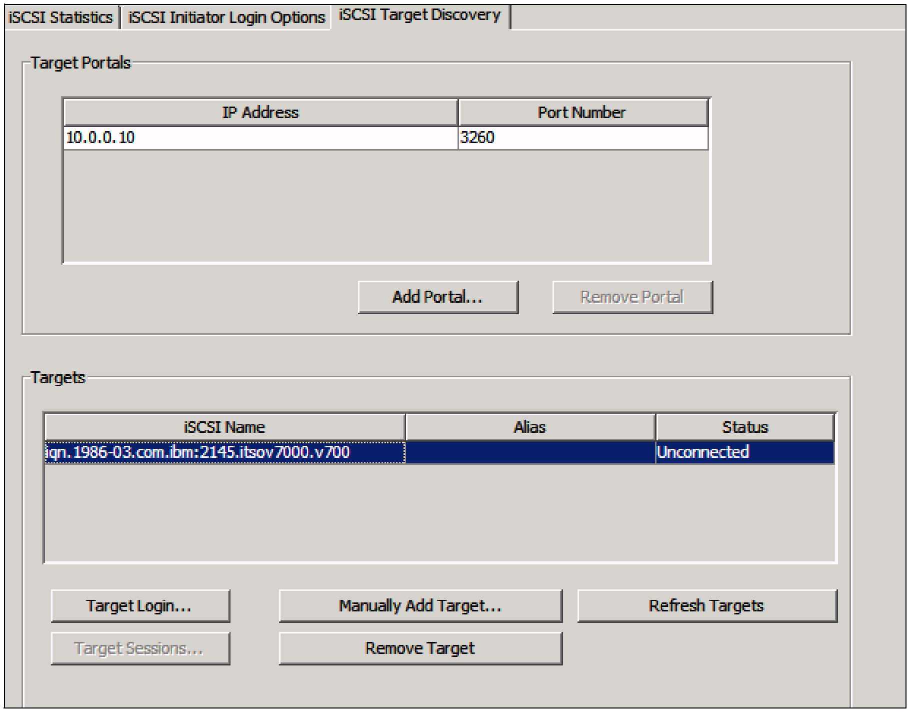

Figure 10-112 iSCSI Target Discovery window

8. In the Target Login window (Figure 10-113), in the Target Name field, confirm that the name is the same name for iSCSI storage IBM Storwize V7000. Click **OK**.

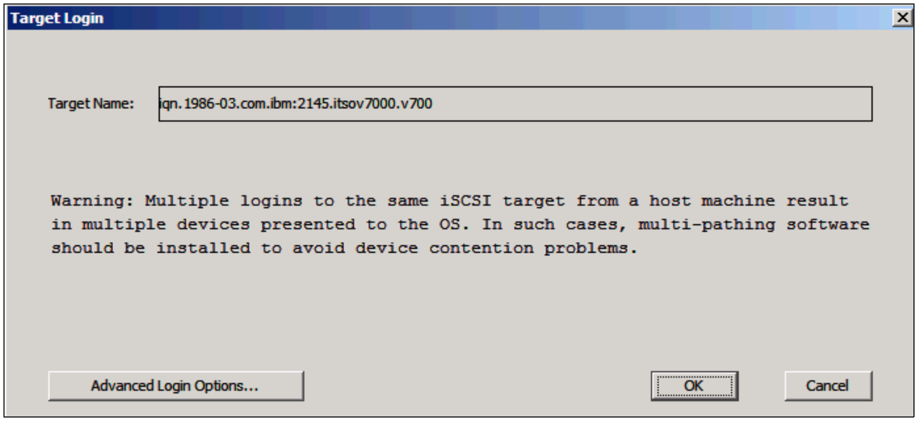

Figure 10-113 Target Login window

The target is displayed with status *Connected* (Figure 10-114).

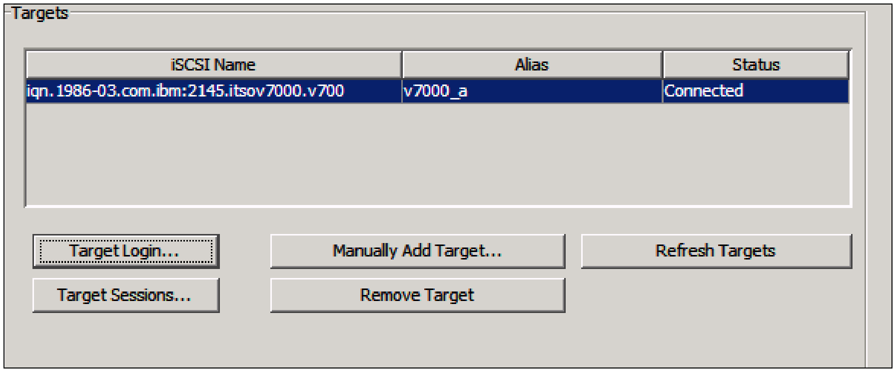

Figure 10-114 Targets window

Configuring the BNT switch

To configure the BNT switch module:

1. Log on the BNT switch by using AMM or use a web browser. Enter the default user ID (`admin`) and password (`admin`).
2. Enable CEE global ports (Figure 10-115):
   a. Select **Configure**.
   b. In the left navigation pane, expand **CEE**, and select **General**.
   c. In the CEE Configuration pane, set Global CEE On/Off to **On**.
   d. Click **Submit**.

Figure 10-115 CEE Configuration pane

3. Click **Apply**, and then click **Save**.

4. Create a VLAN, and add INT1 - 14 ports and any external ports to which you plan to connect storage (Figure 10-116):

   a. In the left navigation pane, expand **Layer 2** → **Virtual LANs**, and select **Add VLAN**.

   b. In the VLAN “New” Configuration pane on the right side:

      i. For VLAN ID, enter `1002`.
      ii. For VLAN Name, enter `VLAN 1002`.
      iii. For VLAN State, select **enabled**.
      iv. For Management VLAN State, select **disabled**.
      v. Under Ports Available, select ports **INT1 - INT14** and **EXT4**, and click **Add** to move them to the Ports in VLAN list.

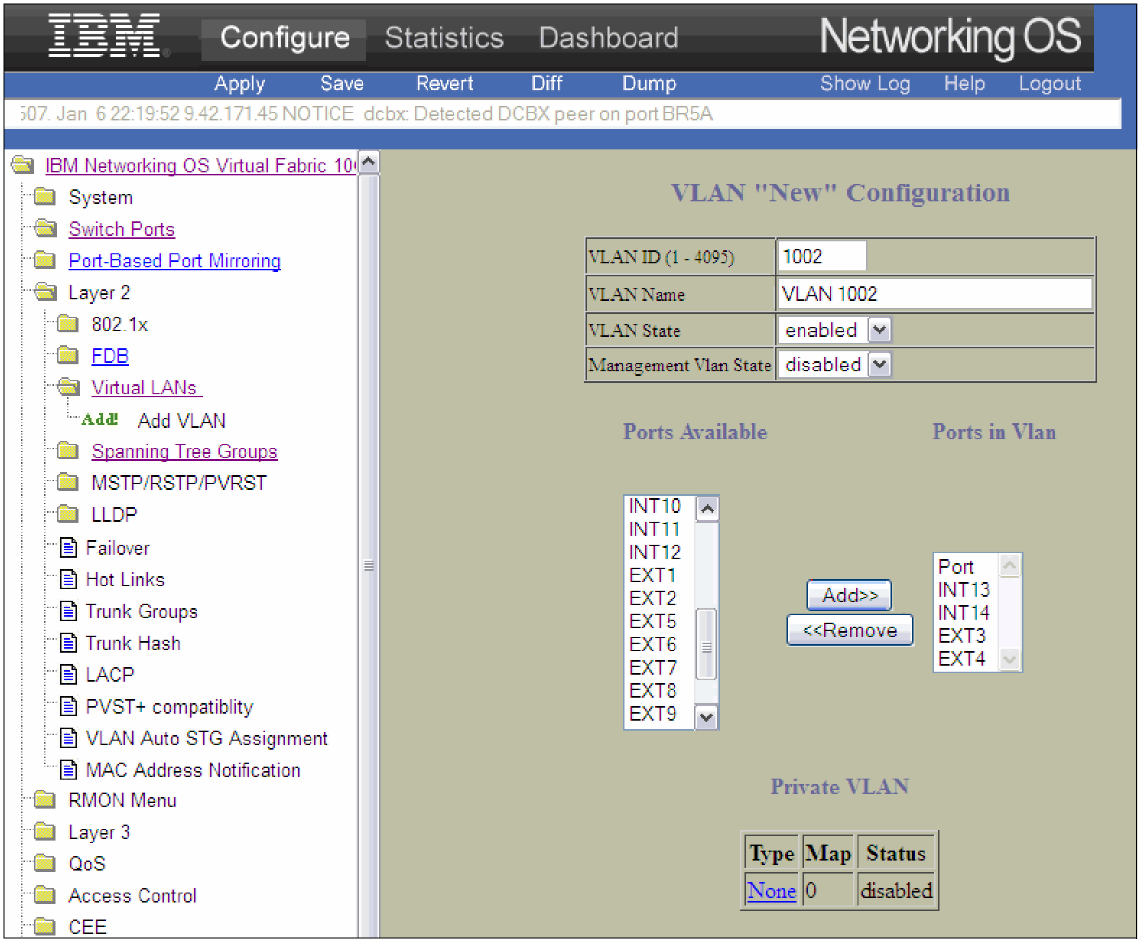

Figure 10-116 VLAN “New” Configuration pane

5. Disable the tags and PVID on the external ports:
   a. On the **Configure** page, in the left navigation pane, select **Switch Ports**.
   b. Select the external ports (EXT1 - 11).
   c. In the Switch Port EXT3 Configuration pane:
      i. For VLAN Tagging, select **Disabled**.
      ii. For PVID Tagging, select **Disabled**.
      iii. Scroll down, and click **Submit**.
      iv. Click **Apply**.
6. Click **Save** at the top of window.

> **Port TAGs:** In this configuration, Port TAGs are supported according to the network topology only to direct connections between the network switch and the IBM Storwize V7000. However, storage does not work with TAGS.

Configuring IBM Storwize V7000 for iSCSI

Now you create a logical driver and map a LUN on IBM Storwize V7000 midrange storage.

Now you create a logical driver and map a LUN on the IBM Storwize V7000 midrange storage. To perform the physical installation, configure the IP of the management console, create a cluster, and define the IP of cluster management, see the IBM Storwize V7000 Product Manuals at:

`https://www.ibm.com/support/docview.wss?uid=ssg1S7003318`

After you configure the adapters and create a zone group, add the WWN storage port on the same zone to guarantee LUN access.

To configure IBM Storwize V7000 for iSCSI:

1. Log on to the cluster management IP.
2. From the Configuration menu, select **Network** → **iSCSI** (Figure 10-117).

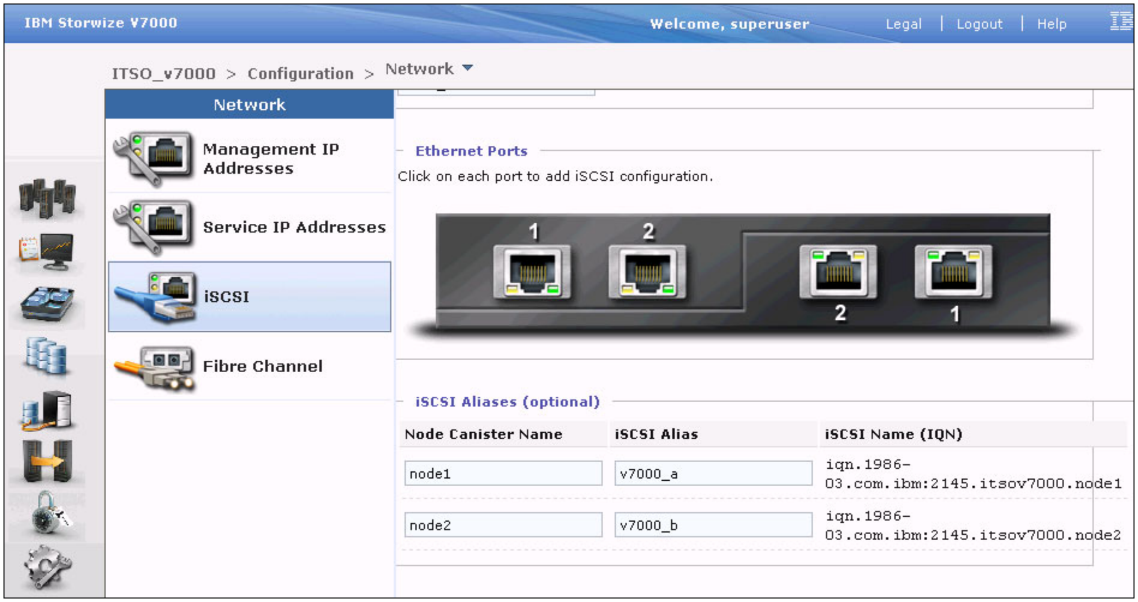

Figure 10-117 Ethernet Ports window

3. Under Ethernet Ports, choose the ports that you want to use for iSCSI connections:

   a. For each port, enter the IP address, subnet mask, and gateway (Figure 10-118). Then click **OK**.

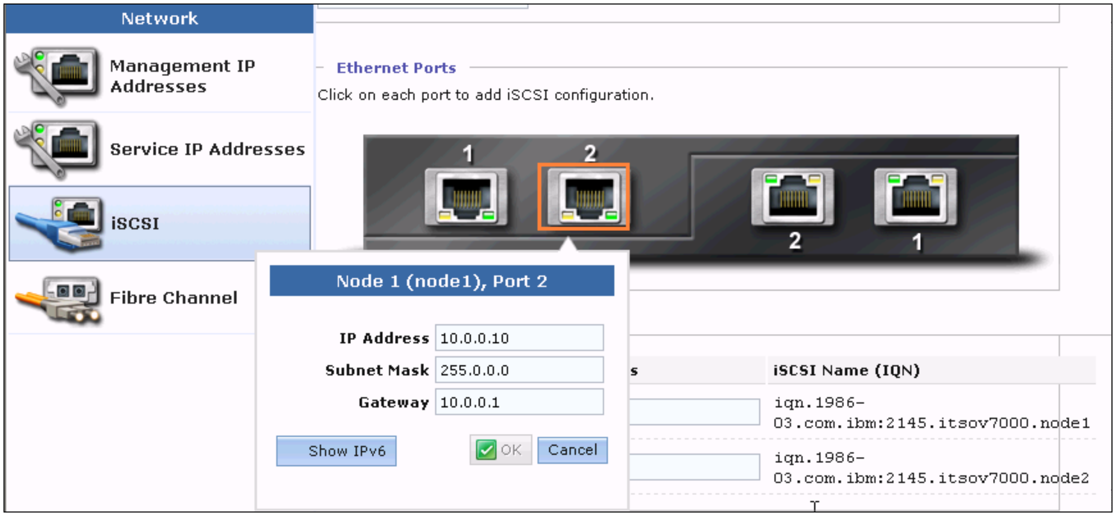

Figure 10-118 Node 1

   b. Define the Canister Names and iSCSI Alias (see Figure 10-117 on page 361).

   c. Copy the iSCSI qualified name (IQN) for future configurations.

   d. Apply the changes.

4. From the **Volumes** menu, select **All Volumes** → **New Volume** (Figure 10-119).

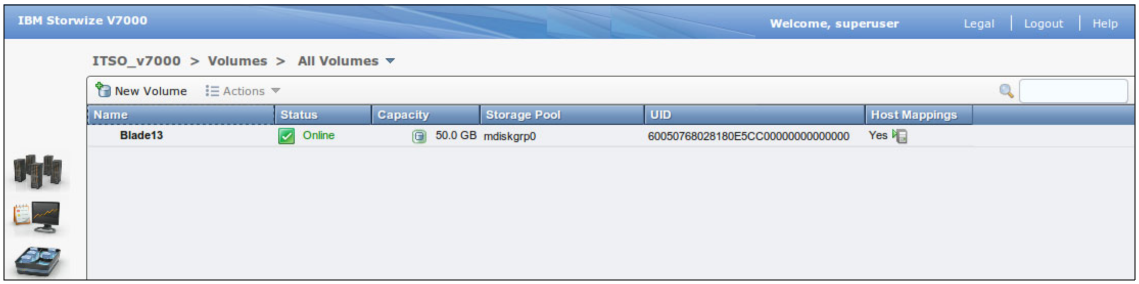

Figure 10-119 IBM Storwize V7000 window

5. In the New Volume panel (Figure 10-120), select **Thin Provision** as the preset type, and then select a pool. For this example, the pool is **mdiskgrp0**. The panel refreshes to show more detail.

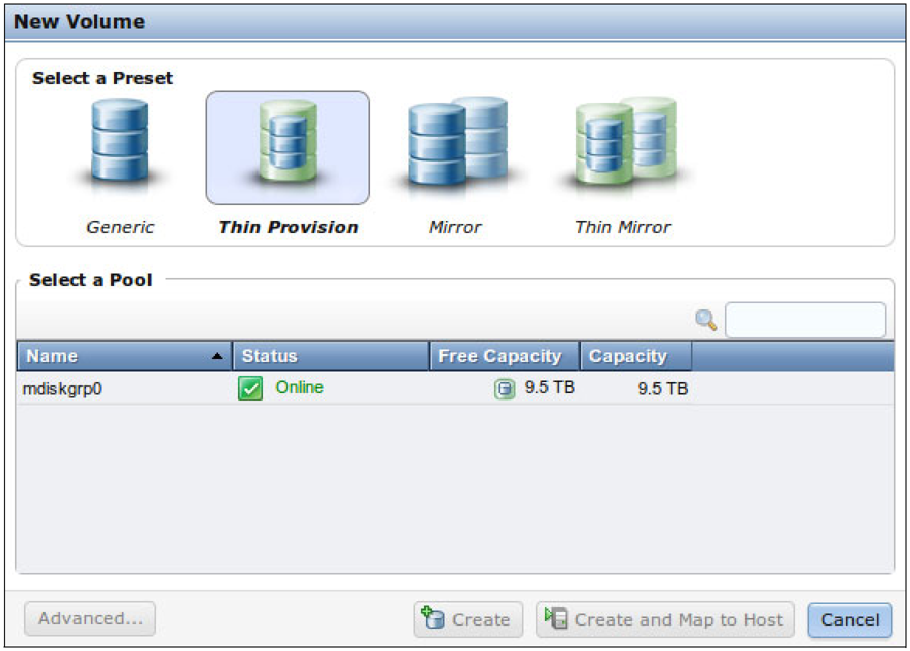

Figure 10-120 New Volume panel

6. Enter the name and size of the volume. For this example, Volume Name is `blade14`, and the size is `100` **GB**. Then click **Create** (Figure 10-121).

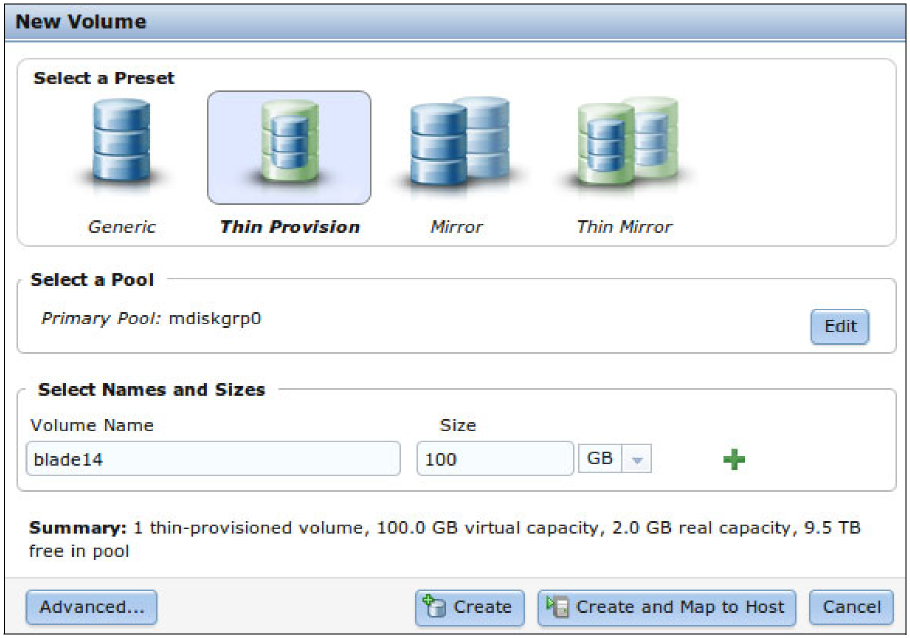

Figure 10-121 Selecting the names and sizes in the New Volume panel

7. From the **Hosts** menu, select **All Hosts**, and then click **New Host**.
8. In the Create Host panel (Figure 10-122), select **ISCSI Host** as the connection. The panel refreshes to show more information.

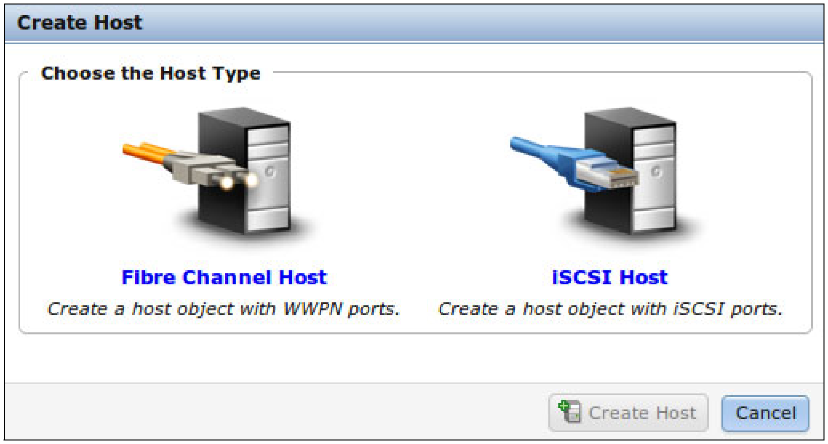

Figure 10-122 Choosing the host type in the Create Host panel

9. In the Host name field, enter the IQN that you defined on the host configuration. For this example, we used `BLADE14_ISCSI` (Figure 10-123). Click **Create Host**.

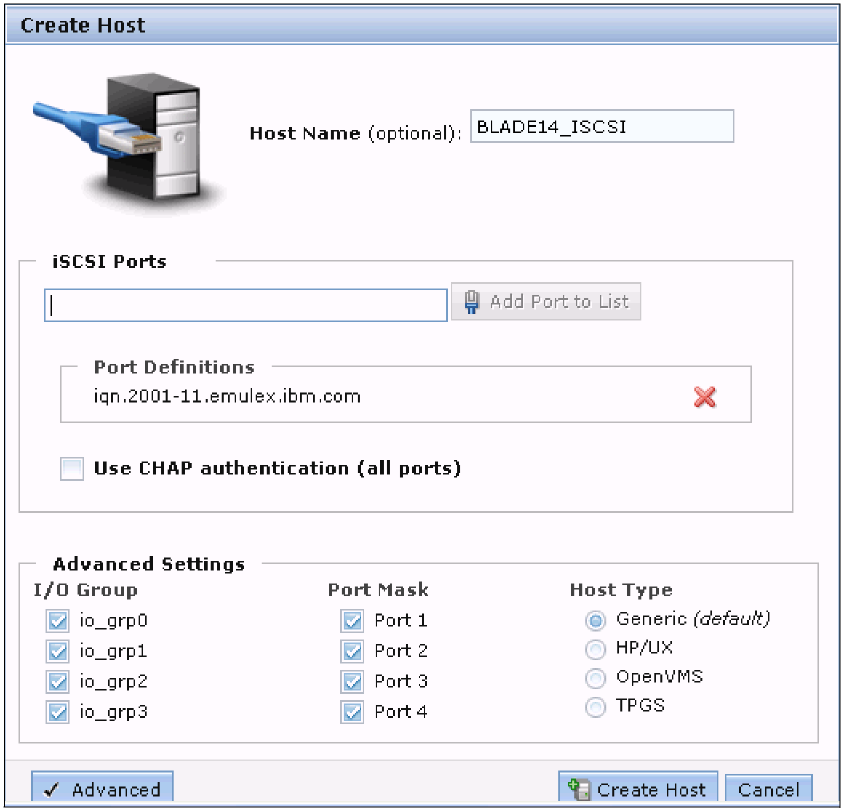

Figure 10-123 Create Host window

10.From the **Volumes** menu (Figure 10-124), select the volume that you created, and click **Actions** → **Map to Host**.

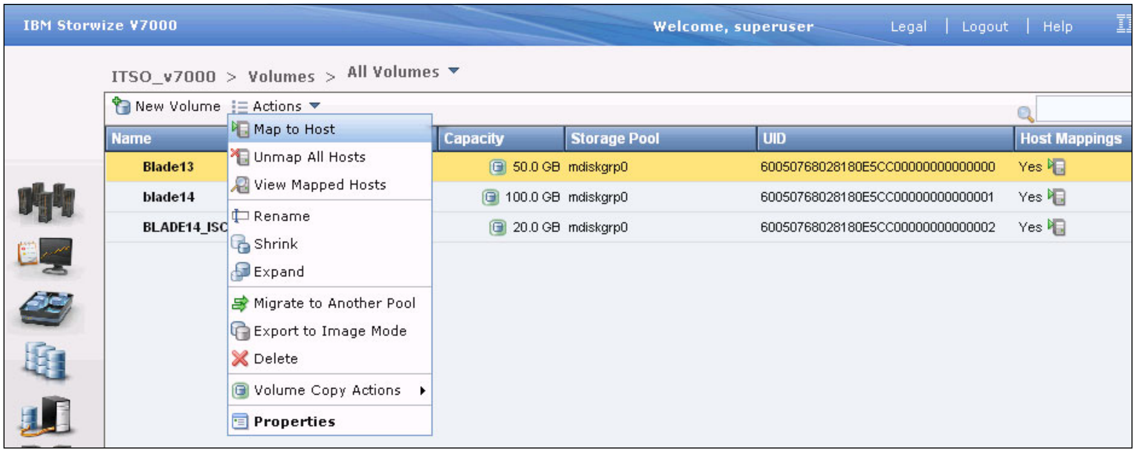

Figure 10-124 Map to Host window

11.Select a host on drop box. For this example, we select **BLADE14_ISCSI**. Then click **Map Volume**.

Your operating system now has access to the disk.

Firmware for this configuration

For this configuration, use the following firmware versions:

- For AMM, use BPET62C.
- For BNT Virtual Fabric 10Gb Switch Module, use version 6.7.5.0 (BOOT and OS).
- For HS22, choose from the following versions:
  - BIOS 1.15
  - IMM 1.30
  - Virtual Fabric Adapter V3.703.397.3806
- For IBM Storwize V7000, use firmware version 6.2.0.4 of the IBM Storwize V7000 Code.
- For Emulex, see *OneConnect Virtual Fabric Adapters (CFFh) for IBM BladeCenter Release Notes for OneConnect SW Version 5.1.42* at:

  http://www.emulex.com/fileadmin/files/downloads/hardware/oce10102/ibm/oneconnect_release_notes.pdf

10.4 Scenario 3: Brocade Converged Switch configuration

This section explains how to configure Switch Independent Mode by using Emulex Virtual Fabric Adapter II, Brocade Converged Switch Module, and the IBM Storwize V7000.

10.4.1 Enabling the Virtual Fabric Adapter Switch Independent Mode function on UEFI

To enable the Virtual Fabric Adapter for Switch Independent Mode:

1. Turn on the system.
2. When prompted, press F1 to enter the UEFI setup.
3. In the System Configuration and Boot Management panel (Figure 10-3 on page 244) of the UEFI utility, select **System Settings**.
4. In the System Settings panel, select **Network**.
5. In the Network panel, select a port on the Emulex adapter on which the vNIC will be enabled.
6. Change to **Switch Independent Mode**.
7. Select **Save** and press Enter.
8. Back in the System Configuration and Boot Management panel (Figure 10-3 on page 244), save your settings before exiting the UEFI utility.
9. Restart your system.

10.4.2 Installing the Windows 2008 drivers

To install the device drivers:

1. Download device driver from Emulex website at:

   http://www.emulex.com/downloads/ibm/vfa-software-kits.html

2. Download NIC drivers for network and ISCSI/FCoE device drivers, in addition to Emulex OneCommand Manager.
3. Run the `elxdrvr-nic-X.XXX.XXX` file.
4. In the window that opens, click **Next**.
5. If you plan to use boot from SAN or to create a Windows PE disk image, select **Unpack All Drivers**. Click **Install**.

6. Click **Finish**. The AutoPilot Installer starts.
7. Click **Next**.
8. Click **Finish** to complete the installation.

Now repeat step 3 on page 367 through step 8 for the next driver. For iSCSI, use the `elxdrvr-iscsi-X.XXX.XXX` file. For FcoE, use the `elxdrvr-fc-fcoe-X.XXX.XXX` file.

For network redundancy, also install the teaming driver `elxdrvr-nic-teaming-X.XXX.XXX` file.

10.4.3 Installing Emulex OneCommand Manager

To install the Emulex OneCommand Manager:

1. Download and run the `elxocm-windows-x86-X.X.XX.X-X` file.
2. In the Emulex OCManager Enterprise window (Figure 10-125), click **Next**.

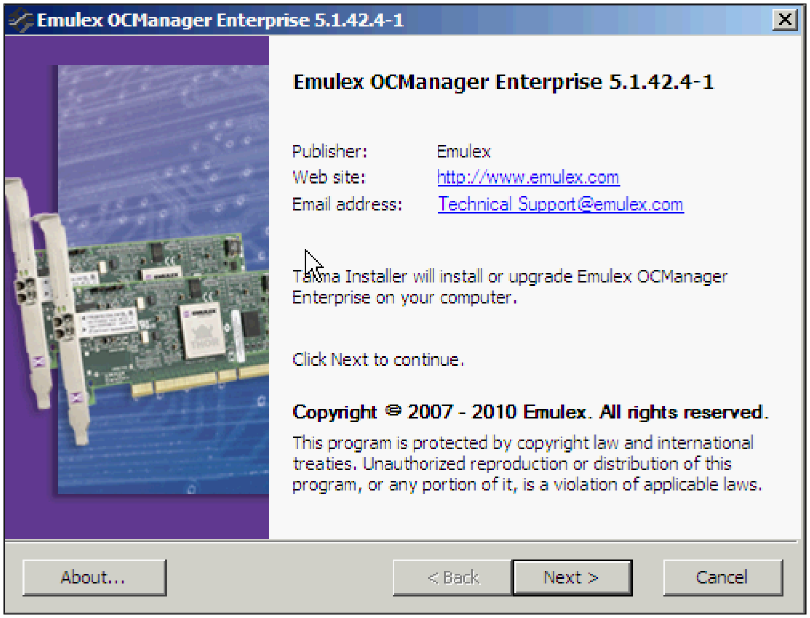

Figure 10-125 Emulex OCManager Enterprise 5.1.42.4-1 window

3. In the Installation options window (Figure 10-126), to **Install** to accept the default options.

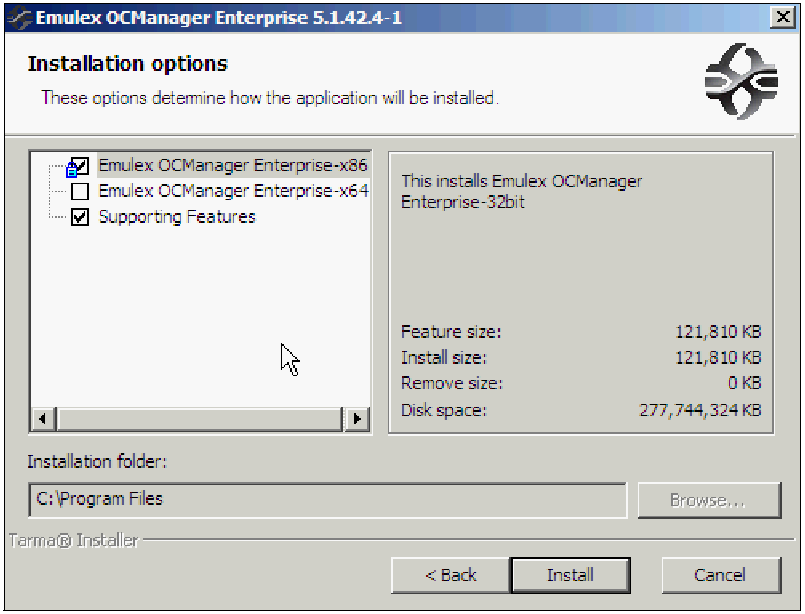

Figure 10-126 Installation options window

4. When prompted, select one of the following modes operations (see Figure 10-127):
   – Strictly Local Management (only management of local adapters)
   – Local Management Plus (only management of local adapters but allows management from other hosts)
   – Full management (management of local and remote hosts; allows management from others hosts)

   In this example, we select **Full management**. Then click **OK**.

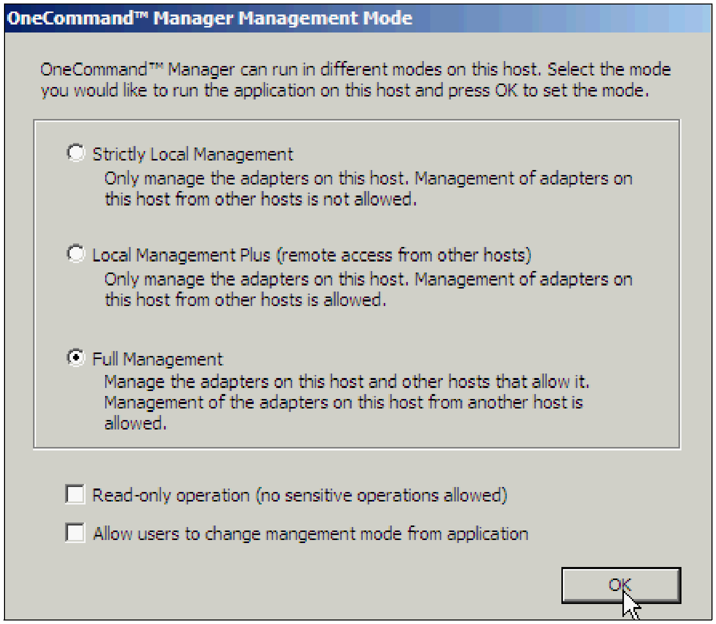

Figure 10-127 OneCommand Manager

10.4.4 Updating the Virtual Fabric Adapter firmware

To download and configure the Virtual Fabric Adapter firmware:

1. Download the file from the Emulex support site at:

   http://www.emulex.com/downloads/ibm/vfa-software-kits.html

2. Open the OneCommand Manager application.

   After OneCommand Manager discovers the hosts and adapters, the ports might be displayed.

3. Click the adapter that you want update (Figure 10-128).

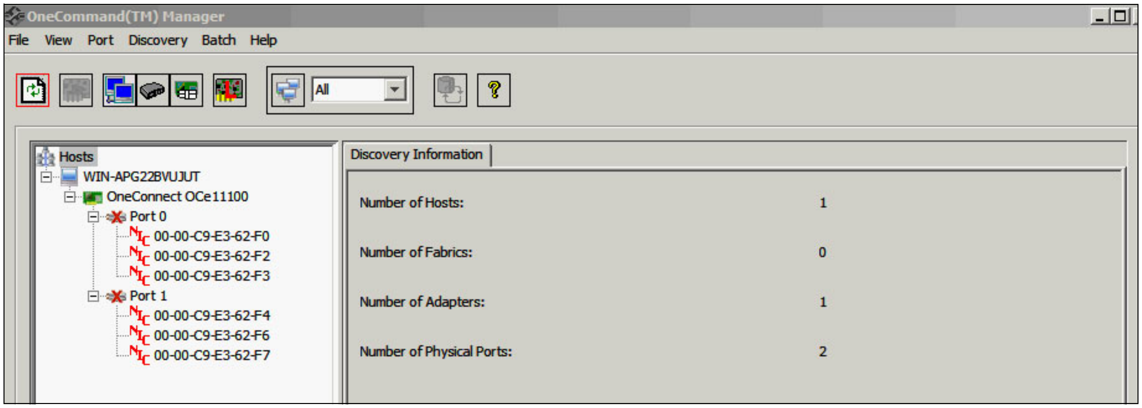

Figure 10-128 OneCommand Manager - Discovery Information

4. Select **Batch** → **Download Firmware**.

5. Select the firmware file:

   a. In the OneCommand Manager Batch Firmware Download window (Figure 10-129), click **Browse**.

   b. In the Firmware File Selection window (Figure 10-129), select the ***xxx*.UFI** update file. For this example, we selected the **T33973806.UFI** file. Click **OK**.

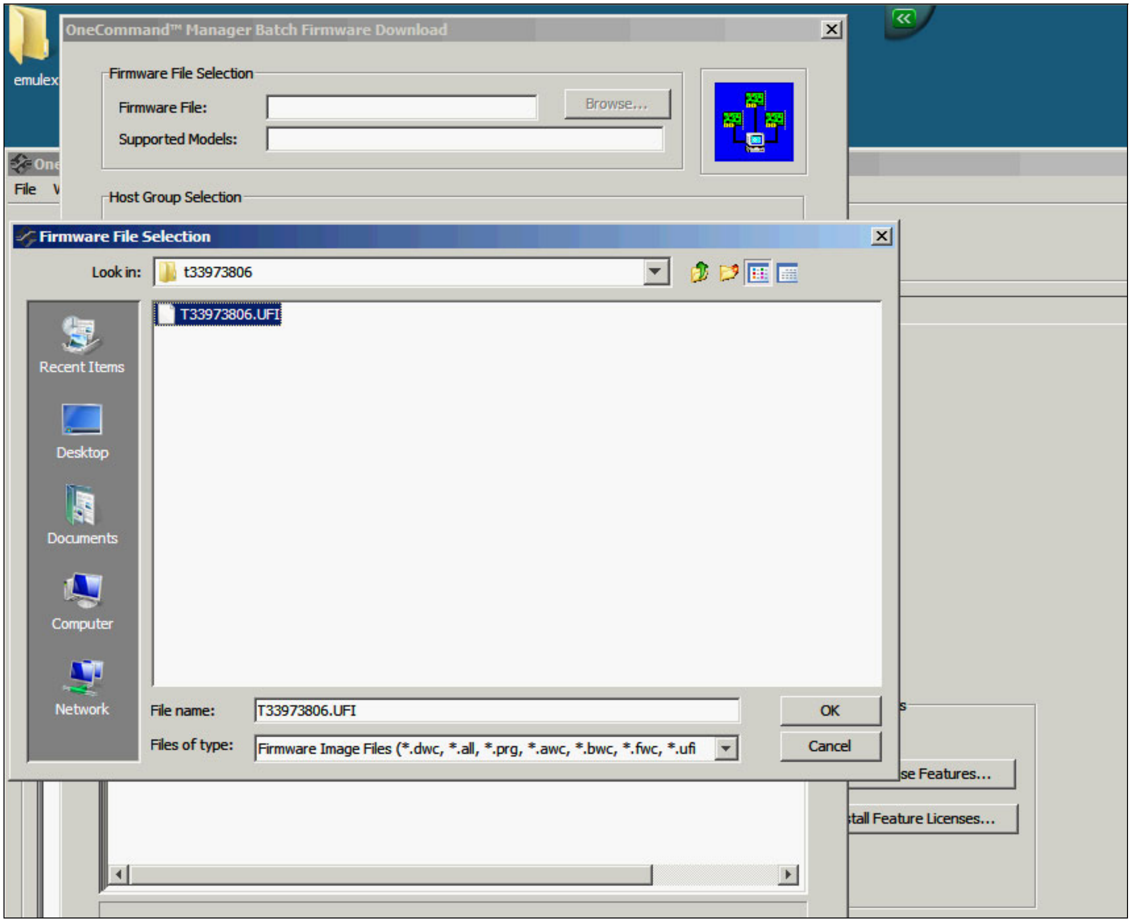

Figure 10-129 Firmware File Section window

c. Back in the OneCommand Manager Batch Firmware Download window (Figure 10-130), next to the box where the adapters are selected, click **Start Download**.

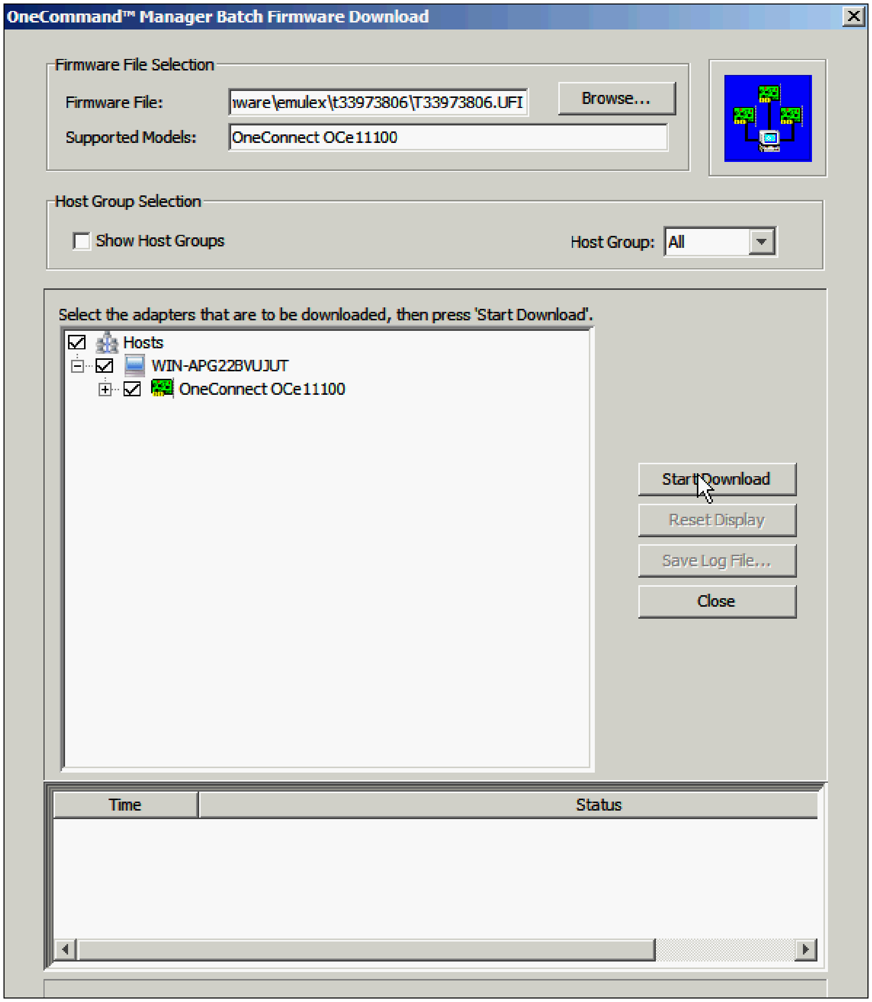

Figure 10-130 OneCommand Manager Batch Firmware Download window

Do not restart or turn off the system until the process is finished.

6. When prompted to restart the system, click **OK**.
7. Check whether the update finished successfully.

8. Again in the OneCommand Manager Batch Firmware Download window (Figure 10-131), click **Close**.

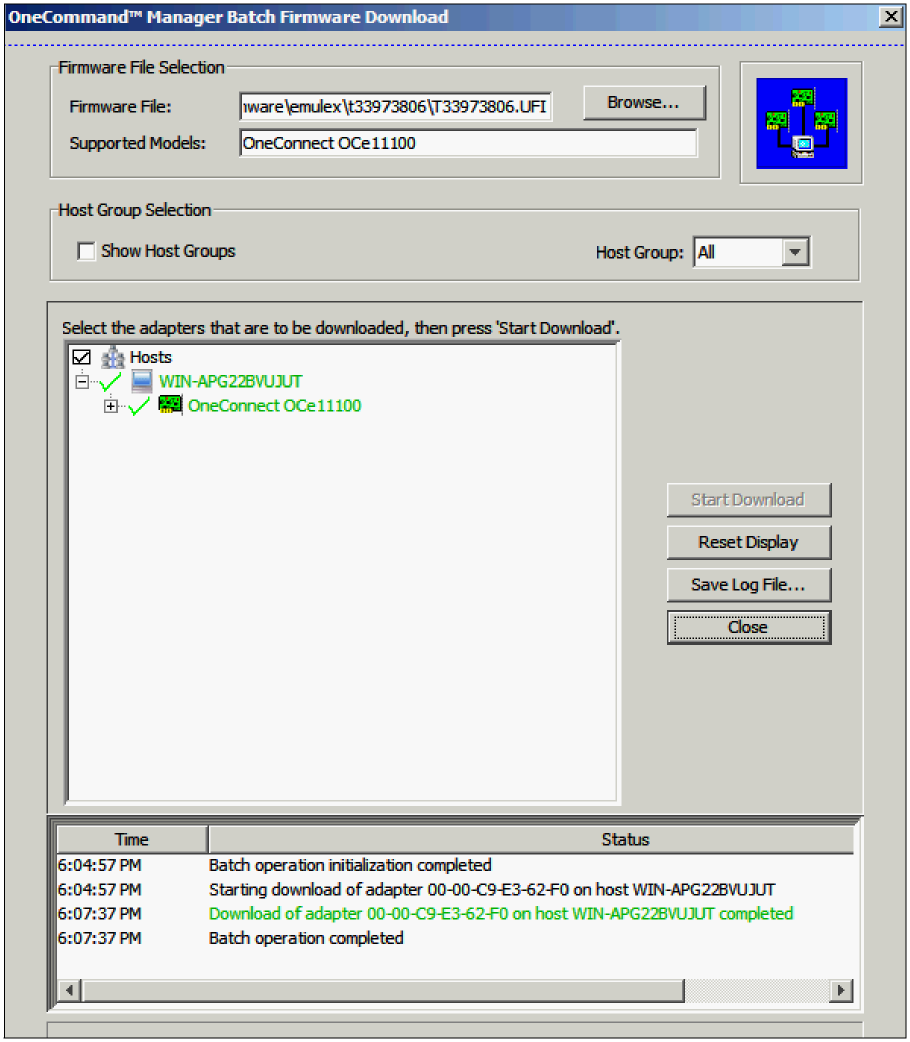

Figure 10-131 Firmware download file complete

9. Restart your system.

10.4.5 Configuring the ports in operation mode

> **Important:** Before you configure your production environment, update firmware to the latest levels.

To configure the ports in operation mode:

1. Open the OneCommand Manager application.

   After OneCommand Manager discovers the hosts and adapters, the ports might be displayed as unavailable as shown in Figure 10-132.

2. In the left pane, click the adapter. In this example, we select **OneConnect OCe11100**.

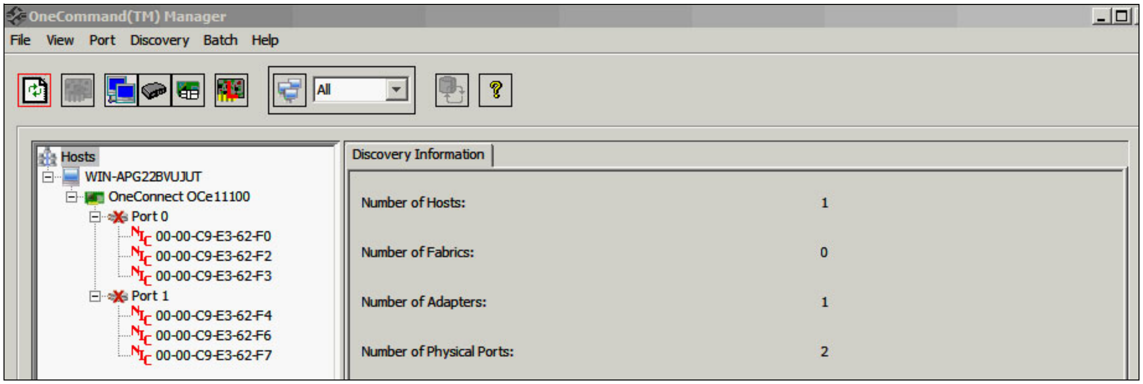

Figure 10-132 OneCommand Manager Discovery Information window

3. In the Personality area (Figure 10-133), under After Reboot, select **FCoE** for the disk connection type. Then click **Apply**.

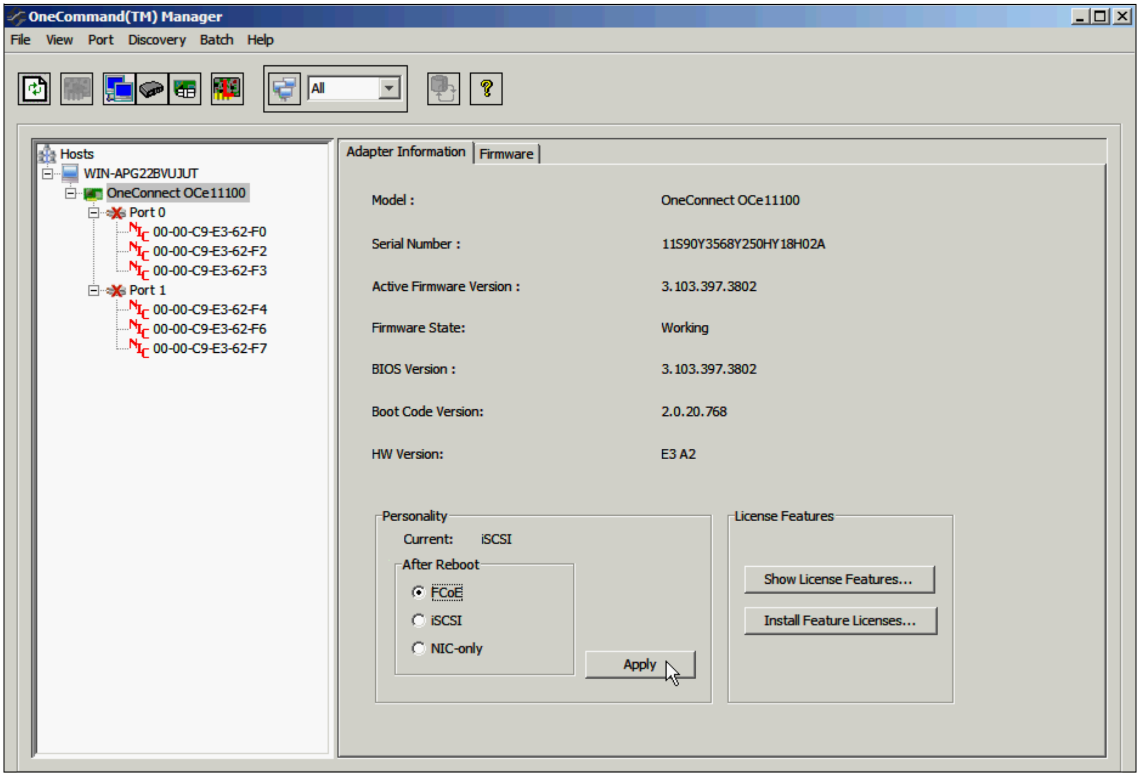

Figure 10-133 Personality area of the OneCommand Manager window

4. When you see the Change Adapter Personality message (Figure 10-134), click **OK** to reboot your system.

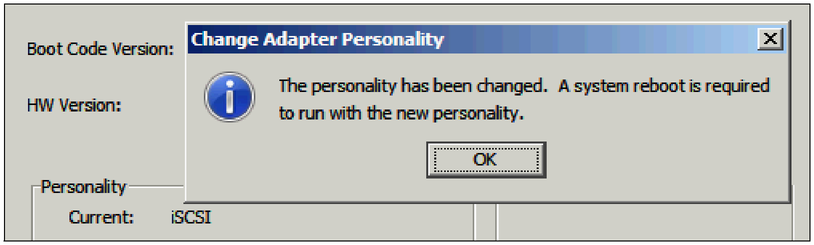

Figure 10-134 Change Adapter Personality window

10.4.6 Virtual Fabric Adapter II: Configuring vNICs with FCoE and Brocade converged switch to IBM Storwize V7000

This section explains how to configure a vNIC with FCoE and Brocade converged switch to connect to the IBM Storwize V7000.

Configuring the Brocade Converged switch

The initial configuration of the switch has the CEE ports shut down. To configure the ports for FCoE operation, you must access the CEE command shell and configure the internal and external CEE ports. After you log in to the switch, use the **`cmsh`** command to access the CEE command shell.

To configure the Brocade converge switch to support FCoE:

1. Log in to the switch by using Telnet.
2. From the command prompt, type **`cmsh`**.
3. Type **`enable`**.
4. Type **`conf t`**.
5. Type **`int int 0/`*`x`***, where ***x*** is the internal port you want to change. After the first **`int`**, press the Tab key, and then type the second **`int`**.
6. When you are in the interface, type **`fcoeport`**.
7. Type **`no shut`**.
8. Type **`exit`**.
9. Repeat steps 5 - 8 for any other blade port that you want to configure.
10. After you are finished with all of the blade ports, type **`exit`**. You should still be in the CMSH.
11. Type **`write mem`**.
12. Answer `yes` to overwrite the startup file.
13. Type **`copy run start`**
14. Answer `yes` to overwrite.

Creating SAN zoning

To create SAN zoning:

1. Log in to the switch by web tool. Log in to the AMM. Select **I/O Module Tasks** → **Configuration**. Select the slot where switch exists, and click **Advanced Options**. Scroll to Start Web Session, and then click **Start Session**.
2. Enter the default user name (`USERID`) and password (`PASSWORD`).

3. In the left pane (Figure 10-135), click **Zone Admin**.

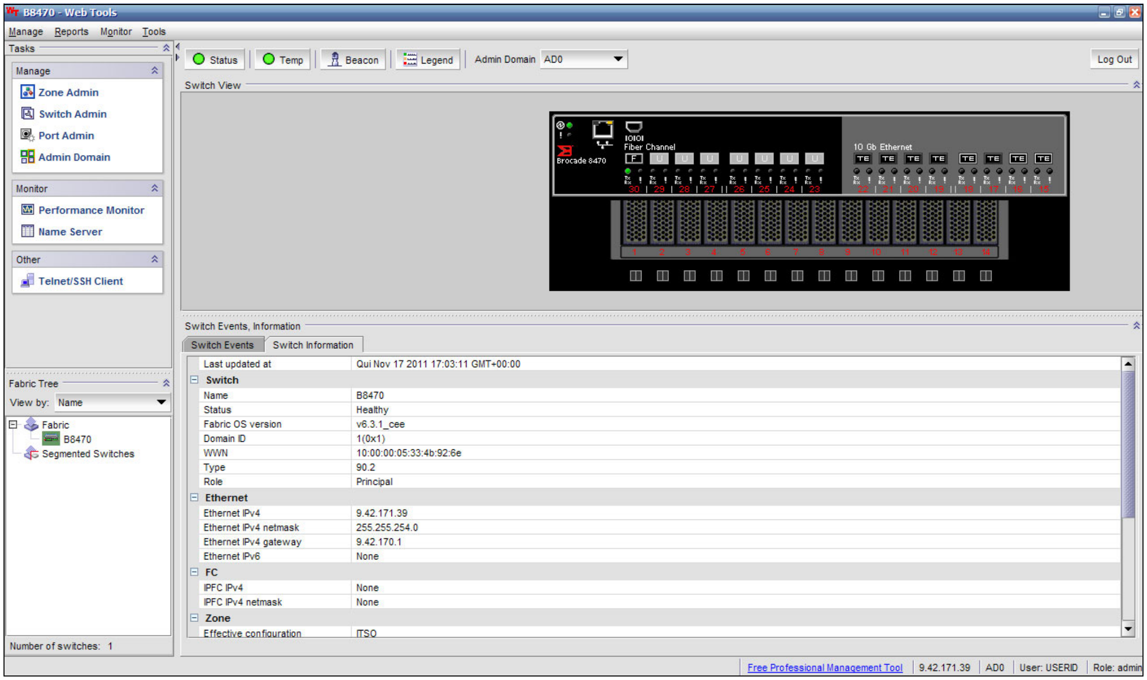

Figure 10-135 WebTools window

4. Create an alias for each port in the Zone Administration window (Figure 10-136):

   a. On the **Alias** tab, click **New Alias**.
   b. Enter a name for the alias.
   c. Enter the WWN member list. For our alias, we used BLADE13.

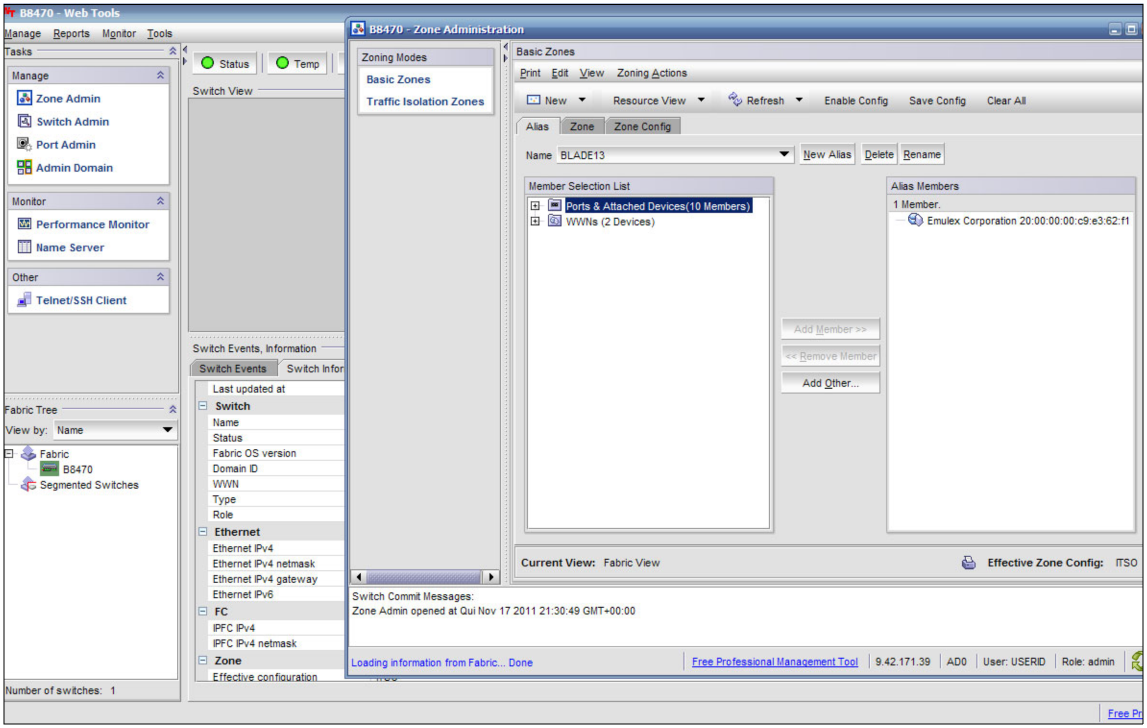

Figure 10-136 Alias page in the Zone Administration window

5. Create a zone:

   a. On the **Zone** tab (Figure 10-137), click **New Zone**.
   b. Enter the name of the zone. For this example, our name is `ITSO_TEST`.
   c. Select alias and external ports that are part of zone.

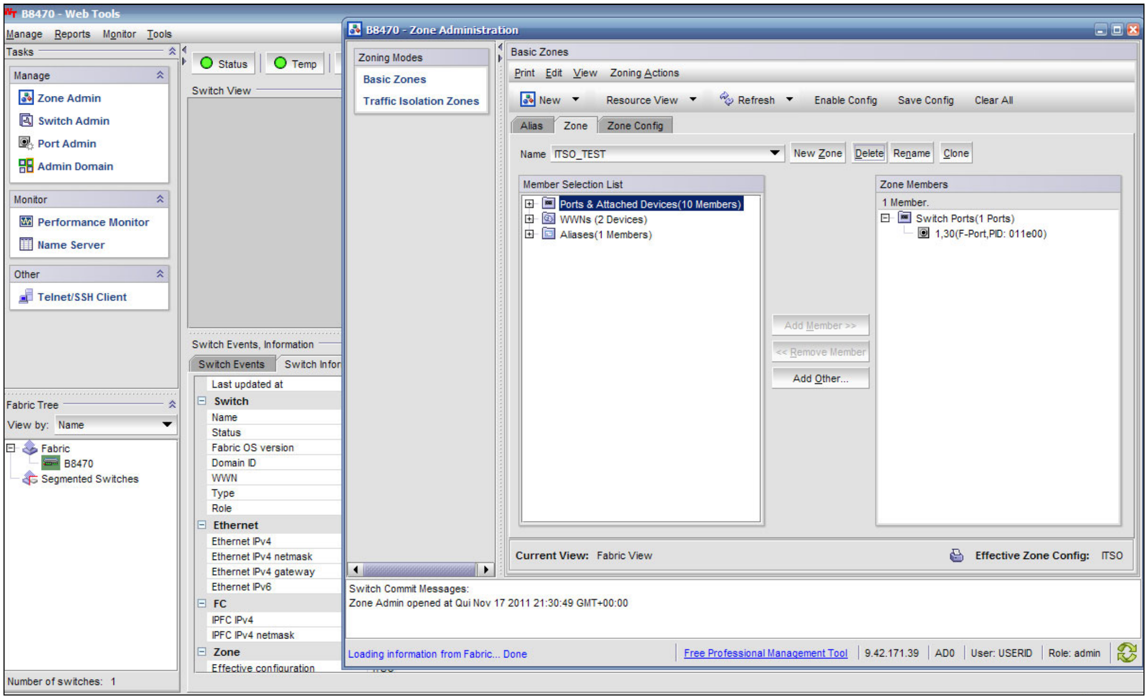

Figure 10-137 Zone page in the Zone Administration window

6. Create a zone configuration:

   a. On the **Zone Config** tab (Figure 10-138), click **New Zone config**.
   b. Click **Add Zone**, and then click **Finish**.

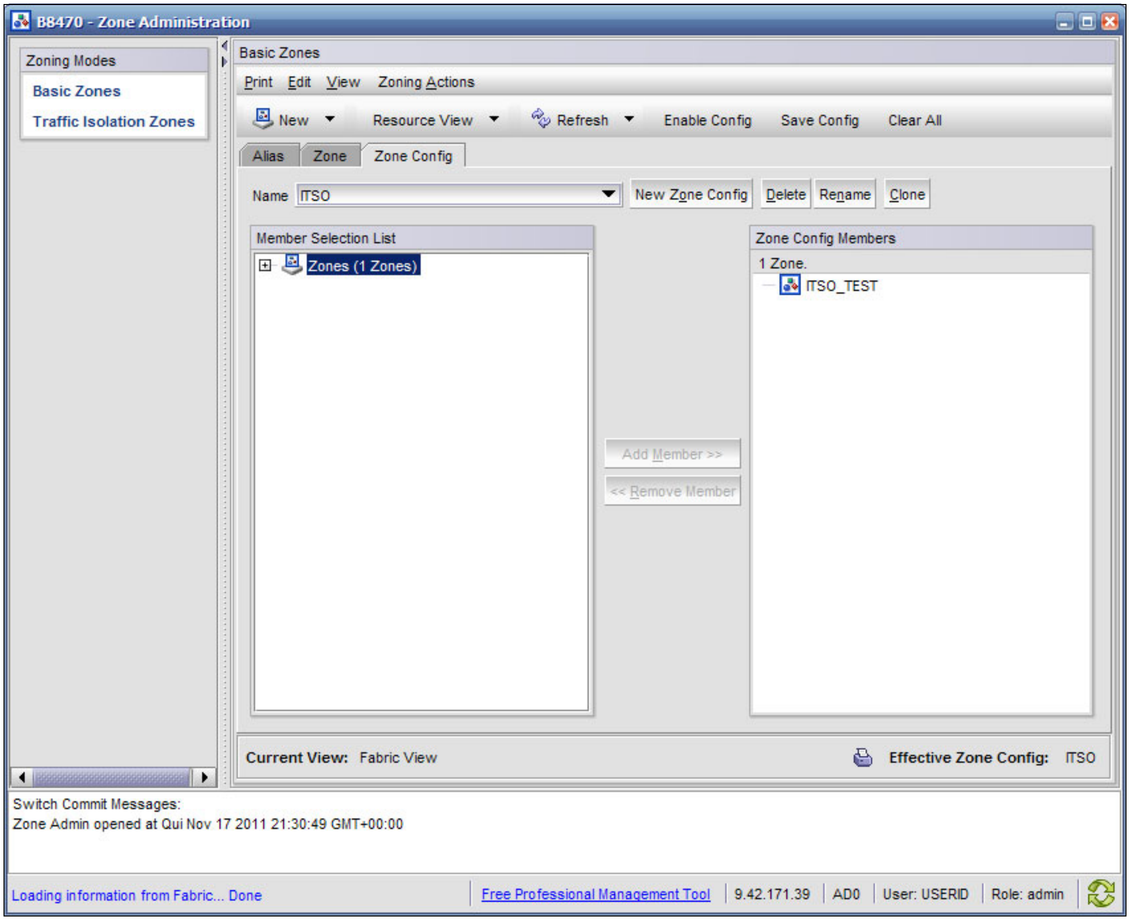

Figure 10-138 Zone Config page in the Zone Administration window

7. Click **Enable Config** to activate the new zone.

Next you proceed to define the host and LUN masking on the storage. However, this section serves only as an example to this point and does not cover the additional configuration of the storage devices. For more information, contact your storage specialist.

Part 4

Abbreviations, related publications, and index

This part includes additional information that support the topic of this book. It includes the following information:

- “Abbreviations and acronyms” on page 385
- “Related publications” on page 389
- “Index” on page 391

Abbreviations and acronyms

AC	alternating current
ACL	access control list
AFT	adapter fault tolerance
ALB	adaptive load balancing
AMM	advanced management module
ARP	Address Resolution Protocol
ASIC	application-specific integrated circuit
BBI	browser-based interface
CD-ROM	compact disc read only memory
CDB	command descriptor block
CDP	Cisco Discovery Protocol
CEE	Converged Enhanced Ethernet
CIFS	Common Internet File System
CIN	Cisco Intel Nuova
CLI	command-line interface
CNA	converged network adapter
COS	class of service
CRC	cyclic redundancy check
DAC	direct-attached copper
DC	domain controller
DCB	Data Center Bridging
DCBX	Data Center Bridging Exchange
DCFM	Data Center Fabric Manager
DHCP	Dynamic Host Configuration Protocol
DMA	direct memory access
DNS	Domain Name Service
DPOD	Dynamic Port on Demand
DPS	Dynamic Path Selection
DSA	Dynamic System Analysis
EGM	Enhanced Group Management
ETS	Enhanced Transmission Selection
ETSI	European Telecommunications Standard Industry
FAN	Fabric Address Notification
FC	Fibre Channel
FCF	Fibre Channel Forwarder
FCoE	Fibre Channel over Ethernet
FDX	full duplex
FEC	Fast EtherChannel
FIPS	Federal Information Processing Standard
FTP	File Transfer Protocol
GEC	Gigabit EtherChannel
GPL	GNU General Public License
GT	Gigatransfers
GUI	graphical user interface
HBA	host bus adapter
HPC	high performance computing
HSSM	high speed switch module
I/O	input/output
I2C	inter-integrated circuit
IBM	International Business Machines
ID	identifier
IEEE	Institute of Electrical and Electronics Engineers
IGMP	Internet Group Management Protocol

IP	Internet Protocol
IPTV	Internet Protocol Television
IPsec	Internet Protocol security
IQN	iSCSI qualified name
IT	information technology
ITSO	International Technical Support Organization
KVM	keyboard video mouse
LACP	Link Aggregation Control Protocol
LAN	local area network
LDAP	Lightweight Directory Access Protocol
LED	light emitting diode
LGPL	Lesser General Public License
LLDP	Link Layer Discovery Protocol
LPVID	Logical Port VLAN ID
LR	long range
LRO	large receive offload
LSO	large send offload
MAC	media access control
MB	megabyte
MBA	Master of Business Administration
MSTP	Multiple Spanning Tree Protocol
NAS	network-attached storage
NC-SI	Network Controller Sideband Interface
NCSI	Network Controller Sideband Interface
NDCLA	Non-Disruptive Code Load Activation
NEBS	Network Equipment Building System
NFS	Network File System
NGN	next-generation network
NIC	network interface card
NTP	Network Time Protocol
OS	operating system
OSI	Open Systems Interconnect
OSPF	Open Shortest Path First
PACLs	port-based ACLs
PCI	Peripheral Component Interconnect
PCIe	Peripheral Component Interconnect Express
PDF	Portable Document Format
PDU	protocol data unit
PE	Preinstallation Environment
PFC	priority-based flow control
POST	power-on self test
PVID	Port VLAN ID
PXE	Preboot Execution Environment
QoS	quality of service
RADIUS	Remote Authentication Dial In User Service
RBAC	Role-based access control
RHEL	Red Hat Enterprise Linux
RHEL6	Red Hat Enterprise Linux 6
RHEV	Red Hat Virtualization
RIP	Routing Information Protocol
RMON	Remote Monitoring
ROM	read-only memory
RSA	Remote Supervisor Adapter
RSCN	Registered State Change Notification
RSS	receive-side scaling
RSTP	Rapid Spanning Tree Protocol
SAN	storage area network
SAS	serial-attached SCSI
SATA	Serial ATA
SFP	small form-factor pluggable

SFT	switch fault tolerance
SLB	Smart Load Balancing
SLES	SUSE Linux Enterprise Server
SMB	Server Message Block
SNMP	Simple Network Management Protocol
SNS	Simple Name Server
SR	short range
SSH	Secure Shell
STP	Spanning Tree Protocol
TACACS	Terminal Access Controller Access Control System
TADDR	TFTP address
TCO	total cost of ownership
TCP	Transmission Control Protocol
TCP/IP	Transmission Control Protocol/Internet Protocol
TF	Transparent Fabric
TFTP	Trivial File Transfer Protocol
TOE	TCP/IP Offload Engine
ToS	type of service
UCNA	Universal Converged Network Adapters
UDLD	unidirectional link detection
UDP	User Datagram Protocol
UEFI	Unified Extensible Firmware Interface
USB	universal serial bus
USERID	user name
UTP	unshielded twisted pair
UXSPI	UpdateXpress Systems Pack Installer
VACL	VLAN-based ACL
VE	Virtualization Engine
VF	Virtual F_Ports
VFA	Virtual Fabric Adapter
VFS	Virtual Fabric Switch
VGRP	vNIC Group
VLAN	virtual local area network
VM	virtual machine
VN	Virtual N_Port
VOIP	Voice over Internet Protocol
VPD	vital product data
VRRP	Virtual Router Redundancy Protocol
VoIP	Voice over IP
WRR	weighted round-robin
WWN	worldwide name
WWPN	worldwide port name
XML	Extensible Markup Language
XOR	exclusive OR
iSCSI	Internet Small Computer System Interface
isCLI	industry-standard command-line interface
pNIC	physical NIC
vNIC	virtual network interface card

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *Emulex 10GbE Virtual Fabric Adapter and Virtual Fabric Adapter Advanced for IBM BladeCenter*, TIPS0748
- *BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter*, TIPS0708
- *Cisco Nexus 4001I Switch Module for IBM BladeCenter*, TIPS0754
- *Broadcom 2-port 10Gb Virtual Fabric Adapter for IBM BladeCenter*, TIPS0827
- *Emulex 10GbE Virtual Fabric Adapter II for IBM BladeCenter*, TIPS0828

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following web site:

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources:

- Publications for the BNT Virtual Fabric 10Gb Switch Module

  http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5080917
- Product publications for the Cisco Nexus 4001I Switch Module

  http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5082494
- Installation Guide for IBM 10 Gb Ethernet Passthru Module

  http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5080966

Online resources

These websites are also relevant as further information sources:

- BNT Virtual Fabric 10 Gigabit switch module firmware update v6.5.3.0

  http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5086594

- Cisco support and downloads

  http://www.cisco.com/cisco/web/support/index.html

- Emulex OneCommand Manager

  http://www.emulex.com/products/management-software/device-management/onecommand-manager/

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

Numerics

46C7191 40, 173
46M6071 181
49Y4235 33–34
49Y4265 34, 38, 168
49Y4275 33–34
49Y9983 181
81Y3133 171
90Y3550 33, 38, 167–168
90Y3566 33, 38, 167–168

A

architecture
- IBM Virtual Fabric Mode 19, 22
- Switch Independent Mode 158

B

bandwidth 4
- IBM Virtual Fabric Mode 21
- Switch Independent Mode 160–161

BBI 66
- examples 103
- logon 68
- Quick Guide 67
- Switch Dashboard 69
- VLANs 87
- vNIC
  - creation 69, 76
  - mode 71

BladeCenter Open Fabric Manager 10
- IBM Virtual Fabric Mode 30
- Switch Independent Mode 160, 196

BladeCenter Virtual Fabric solution concept 4

BladeOS
- BNT Virtual Fabric 10Gb Switch Module 59
- command reference 59
- examples 103
- logon 59
- VLANs 86
- vNIC
  - creation 69, 77
  - groups 80
  - mode 73

BNT Virtual Fabric 10Gb Switch Module 39, 172
- BladeOS 59
- cables 46, 180
- configuration 59, 196, 206
- configuration tools 51
- DAC cables 41, 174
- direct-attached cables 40, 173
- examples 274
- features 42, 175
- firmware update 60
- IEEE standards 46, 179
- isCLI 65
  - privilege levels 65
- Layer 3 functions 44, 177
- management 45, 177
- performance 43, 176
- ports 42, 175
- publications 59
- QoS 44, 177
- redundancy 43, 176
- reset 64
- scalability 43, 176
- scenarios 274
- security 44, 177
- SFP+ transceivers 40, 173
- shipgroup 41, 174
- stacking 25
- transceivers 40–41, 173–174
- VLANs 44, 176
- vNIC
  - configuration 23
  - creation 69
  - group creation 79

Broadcom 2-port 10Gb Virtual Fabric Adapter 171
- bandwidth 205
  - weight 213
- Broadcom Advanced Control Suite 218, 221
- configuration 197
- Device Configuration Menu 201
- features 160, 171
- MAC address maximums 158
- MAC addresses 171
- Maximum Bandwidth parameter 213

oversubscription 160
relative bandwidth 205
Switch Independent Mode 158
teaming 229
vNIC creation 209
Broadcom Advanced Control Suite 218, 221

C
channels 8
characteristics
IBM Virtual Fabric Mode 24
Switch Independent Mode 157
Cisco Nexus 4001I Switch Module 180
BladeCenter Open Fabric Manager 181
cables 186
command reference 208
configuration 209
DAC cables 182
examples 242
FCoE 181, 184
firmware 206
IEEE standards 186
license upgrade 181
performance 183
ports 182
publications 188
scenarios 242
security 184
Software Upgrade License 181
Telnet 272
transceivers 182
VLAN 272
support 184
compatibility 6
components
IBM Virtual Fabric Mode 32
Switch Independent Mode 165
configuration
BNT Virtual Fabric 10Gb Switch Module 59, 196, 206
Broadcom 2-port 10Gb Virtual Fabric Adapter 197
Cisco Nexus 4001I Switch Module 209
Emulex 10GbE Virtual Fabric Adapter II 52
IBM Virtual Fabric Mode 51

D
Data Center Bridging Exchange (DCBX) protocol 51
double tag scheme, IBM Virtual Fabric Mode 21

E
elxflash command 53
Emulex 10GbE Virtual Fabric Adapter 33, 37, 51
comparison 39
configuration 52
features 34
firmware 52
minimum bandwidth 29
modes supported 35
UEFI 54
vNIC mode in UEFI 55
Emulex 10GbE Virtual Fabric Adapter II 33, 167
configuration 52
driver 217
features 160, 168
firmware 52
minimum bandwidth 29
modes of operation 38
Multichannel Configuration 57
Switch Independent Mode 158
UEFI 54
vNIC mode in UEFI 55
Emulex OneCommand Manager 94, 225
Emulex Virtual Fabric Adapters II 305
examples
BNT Virtual Fabric 10-Gb Switch Module 274
Cisco Nexus 4001I Switch Module 242
IBM Virtual Fabric Mode 103
Switch Independent Mode 237
existing tagging 9

F
failover 9
IBM Virtual Fabric Mode 23
Switch Independent Mode 163
FCoE 9
firmware
BNT Virtual Fabric 10Gb Switch Module 60
Cisco Nexus 4001I Switch Module 206
Emulex 10GbE Virtual Fabric Adapters 52

G
Global Configuration Mode 65
groups 9

gtimg command 63

H

hardware
IBM Virtual Fabric Mode 31
Switch Independent Mode 165
high availability, IBM Virtual Fabric Mode 81
hypervisor 5

I

IBM Storwize V7000 305, 337
IBM Virtual Fabric Mode 6, 51
architecture 19, 22
bandwidth 21
BladeCenter Open Fabric Manager 30
BNT switch configuration 59
BNT Virtual Fabric 10Gb Switch Module 39
characteristics 24
comparison 8
components 32
double tag scheme 21
Emulex 10GbE Virtual Fabric Adapter 33, 37
Emulex 10GbE Virtual Fabric Adapter II 33
examples 103
failover 23
firmware updates 53, 60
hardware 6, 31
high availability 81
isCLI 65
Layer 2/3 functions disabled 21
LPVID 58
mapping of vNICs 23
minimum bandwidth 29
multichannel configuration 57
optics 32
outer tag 21
overview 20
products 31
rules 24
scenarios 103
selection 9
SFP and transceivers 32
solution configuration 51
stacking 22, 25
teams 93
transceivers 32
trunk failover 23
trunks 30, 81
virtual channels 20
VLAN
configuration 85
tagging 21
VMready 24, 101
vNIC
creation 69
enablement 71
mode 21
vNIC groups 30
creation 79
vSwitches 21
IEEE 802.1q 7, 9
VLAN tag 157
ifconfig command 287
isCLI
BNT Virtual Fabric 10Gb Switch Module 65
command reference 65
examples 103
privilege levels 65
VLAN 86
vNIC
creation 69, 76
groups 80
mode enablement 72
iSCSI 9

J

jumbo frames 34

L

Layer 2/3 functions disabled, IBM Virtual Fabric Mode 21
limits 8
Logical Port VLAN ID (LPVID), IBM Virtual Fabric Mode 58

M

mapping of vNICs
IBM Virtual Fabric Mode 23
Switch Independent Mode 159
maximum VLANs 8
minimum bandwidth
IBM Virtual Fabric Mode 29
Switch Independent Mode 161
mode selection 9
modes 6

Multichannel Configuration 57

O
OneCommand 94
optics
- IBM Virtual Fabric Mode 32
- Switch Independent Mode 175, 182

outer tag
- IBM Virtual Fabric Mode 21
- Switch Independent Mode 157

oversubscription, Broadcom 2-port 10-Gb Virtual Fabric Adapter 160
overview 5
- IBM Virtual Fabric Mode 20
- Switch Independent Mode 158

P
planning for Switch Independent Mode 160
Privileged EXEC mode 65
products
- IBM Virtual Fabric Mode 31
- Switch Independent Mode 165

R
Redbooks website 389
- Contact us xii

rules
- IBM Virtual Fabric Mode 24
- Switch Independent Mode 163

S
scenarios
- BNT Virtual Fabric 10-Gb Switch Module 274
- Cisco Nexus 4001I Switch Module 242
- IBM Virtual Fabric Mode 103
- Switch Independent Mode 237

SFP and transceivers
- IBM Virtual Fabric Mode 32
- Switch Independent Mode 175, 182

stacking 9, 22
- IBM Virtual Fabric Mode 25
- Switch Independent Mode 163

Switch Independent Mode 155, 337
- about 155
- architecture 158
- bandwidth 160–161
  - limits 155
- BladeCenter Open Fabric Manager 160, 196
- BladeOS 59
- BNT Virtual Fabric 10Gb Switch Module 166
- Broadcom 2-port 10Gb Virtual Fabric Adapter 158, 166, 197, 229
- Broadcom Advanced Control Suite 218, 221
- Cisco Nexus 4001I Switch Module 166
- comparison 8
- configuration 195
- Emulex 10GbE Virtual Fabric Adapter II 158, 166
- Emulex OneCommand Manager 225
- examples 237
- failover 163
- features 160
- firmware 196
- hardware 7, 165
- IEEE 802.1Q VLAN tag 157
- introduction 6
- MAC addresses 171
- outer tag 157
- overview 158
- planning 160
- products 165
- scenarios 237
- selection 10
- stacking 163
- VLAN 159, 214
  - groups 162
  - tagging 157
- VMready 163, 215
- vNIC creation 209

T
teaming
- Broadcom 2-port 10Gb Virtual Fabric Adapter 229
- creation 94
- Emulex 10GbE Virtual Fabric Adapter II 225
- IBM Virtual Fabric Mode 93, 225

Telnet 59
- Cisco Nexus 4001I Switch Module 272

total ports 7
transceivers, IBM Virtual Fabric Mode 32
trunk failover, IBM Virtual Fabric Mode 23
trunking 9
trunks, IBM Virtual Fabric Mode 30, 81

U
UEFI (Unified Extensible Firmware Interface) 51
Unified Extensible Firmware Interface (UEFI) 51
User EXEC mode 65

V
virtual channels 8
virtual groups 9
virtual NIC
    See vNIC
virtualization 4
VLAN groups
    IBM Virtual Fabric Mode 30
    Switch Independent Mode 162
VLAN tagging 4
    IBM Virtual Fabric Mode 21
    Switch Independent Mode 157
VLANs
    maximum 8
    Switch Independent Mode 214
VMready 101
    IBM Virtual Fabric Mode 24
    Switch Independent Mode 163, 215
vNIC 5
    creation
        IBM Virtual Fabric Mode 69
        isCLI 69, 76
    groups, IBM Virtual Fabric Mode 30
vNIC2
    See Switch Independent Mode
vSwitches 21

IBM

Redbooks

IBM BladeCenter Virtual Fabric Solutions

(0.5” spine)
0.475”<->0.875”
250 <-> 459 pages

IBM

IBM BladeCenter Virtual Fabric Solutions

Redbooks

Highlights IBM Virtual Fabric Mode and Switch Independent Mode

Defines the concepts and benefits of virtual NICs (vNICs)

Outlines steps to create various vNIC solutions

The deployment of server virtualization technologies in data centers requires significant efforts in providing sufficient network I/O bandwidth to satisfy the demand of virtualized applications and services. For example, every virtualized system can host several dozen network applications and services. Each service requires certain bandwidth (or speed) to function properly. Furthermore, the traffic flows of different network traffic patterns that are relevant to various service types can interfere with each other. They can lead to serious network problems.

The IBM Virtual Fabric solution for IBM BladeCenter addresses these issues. The solution is based on the IBM BladeCenter H chassis with a 10 Gbps Converged Enhanced Ethernet infrastructure. This infrastructure is built on 10 Gbps Ethernet switch modules in the chassis and the Emulex or Broadcom Virtual Fabric Adapters in each blade server.

This IBM Redbooks publication provides configuration scenarios that use technology from five global leaders in network switch and adapter technology: Broadcom, Brocade, Cisco, Emulex, and IBM. These scenarios demonstrate the usage of IBM System Networking technology and products from Broadcom, Brocade, Cisco, and Emulex. This book is for clients who want to learn how to implement an IBM Virtual Fabric solution by using the IBM Virtual Fabric Mode offering or the Switch Independent Mode offering. This book explains step-by-step how to configure the adapters and switches.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7966-01 0738436712